

1 2 9 0



UNIVERSIDADE D  
COIMBRA

Paulo Alexandre Brinca da Costa Braz

SENSITIVITY TO THE  $0\nu\beta\beta$  DECAY OF  $^{136}\text{Xe}$   
AND DEVELOPMENT OF MACHINE  
LEARNING TOOLS FOR PULSE  
CLASSIFICATION FOR THE LUX-ZEPLIN  
EXPERIMENT

Tese no âmbito do Doutoramento em Física, Astrofísica orientada pelo Professor Doutor Francisco Filipe Bento Neves, pelo Professor Doutor Alexandre Miguel Ferreira Lindote e pelo Professor Doutor José Carvalho Maneira e apresentada ao Departamento de Física da Faculdade de Ciências e Tecnologia da Universidade de Coimbra.

Dezembro de 2020





FACULDADE DE  
CIÊNCIAS E TECNOLOGIA  
UNIVERSIDADE DE  
**COIMBRA**

Paulo Alexandre Brinca da Costa Braz

**Sensitivity to the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  and  
development of Machine Learning tools for pulse  
classification for the LUX-ZEPLIN experiment**

Tese no âmbito do Doutoramento em Física, Astrofísica orientada pelo Professor Doutor Francisco Filipe Bento Neves, pelo Professor Doutor Alexandre Miguel Ferreira Lindote e pelo Professor Doutor José Carvalho Maneira e apresentada ao departamento de Física da Faculdade de Ciências e Tecnologia da Universidade de Coimbra.

Dezembro de 2020

*“This thesis includes the reviews and suggestions done by the Juri”*



*Sensitivity to the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  and development of  
Machine Learning tools for pulse classification for the  
LUX-ZEPLIN experiment*

## Abstract

An elusive form of matter that does not interact via electromagnetic or strong forces permeates the known Universe, and is therefore designated as “dark”. This dark matter (DM) is responsible for the evolution of cosmic structures, the cohesion of galaxies and galaxy clusters, and represents around a quarter of the total content of the Universe. Several state-of-the-art experiments are currently searching for dark matter in the form of weakly interacting massive particles (WIMPs), using ultra-low background “observatories” where one of these particles could interact with a material target and produce a readable signature.

The LZ experiment is a 10 tonne dark matter detector expected to begin operations in early 2021, that aims to surpass the current world-leading limit on the WIMP-nucleon interaction cross section by more than one order of magnitude. The design of LZ features a dual-phase xenon time projection chamber (TPC) and two additional instrumented veto detectors encompassing the TPC for improved background reduction and active shielding. The projected sensitivity of LZ to the spin-independent WIMP-nucleon scattering cross-section is  $1.4 \times 10^{-48} \text{ cm}^2$  for a  $40 \text{ GeV c}^{-2}$  mass WIMP. LZ has the potential to study and discover a wide range of new physics. The inner portions of the TPC of LZ will be one of the most “quiet” environments where rare event searches can be performed.

The ultra-low background required for dark matter searches allows LZ to be potentially sensitive to other rare events such as neutrinoless double beta decay of some xenon isotopes, axion interactions or coherent neutrino-nucleus scattering from solar neutrinos, all of which were not yet observed. The projected sensitivity of LZ to the half-life of the neutrinoless double beta decay of  $^{136}\text{Xe}$  is presented in this document. For an exposure of 1360 kg·year, a sensitivity to the half-life of  $1.06 \times 10^{26}$  years with a 90% confidence level is obtained. The projected sensitivity to this same decay from a dedicated run with a 90%  $^{136}\text{Xe}$  enriched target and an exposure of

13.8 tonne·year is  $1.06 \times 10^{27}$  years.

The development of pulse classification tools for the data processing framework of LZ (LZap) is also presented in this document. These tools represent the groundwork for pulse classification in LZ, both in the form of dedicated heuristics algorithms and machine learning implementations. The *Heuristics Algorithm for Discrimination of Event Substructures* (HADES) developed in the context of this work is currently the default pulse classification tool in LZap, and provides a measured overall classification accuracy of 98.58% across all pulse topologies in LZ simulated data. The *RFClassifier* and the *TriNet* pulse classification tools are two machine learning implementations that use a random forest model and an ensemble of artificial neural networks, respectively, that are aimed at assisting HADES and potentially replacing it in LZap. The *RFClassifier* algorithm achieved a classification accuracy of 99.37% over LZ simulated data when combined with a powerful clustering analysis using Gaussian mixture models (GMMs). The *TriNet* algorithm was trained using the results from HADES and achieved a classification accuracy of 95.56% against the GMM clustering results, but demonstrated that it could generalize its results beyond HADES.

*Sensitivity to the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  and development of  
Machine Learning tools for pulse classification for the  
LUX-ZEPLIN experiment*

## Resumo

Uma forma de matéria que não interage através das forças eletromagnética e forte, por isso designada por matéria “escura”, permeia o universo visível. Esta matéria escura (ME) representa cerca de um quarto do conteúdo total do universo e é responsável pela evolução das estruturas cósmicas e pela coesão das galáxias e dos aglomerados de galáxias. Várias experiências de ponta procuram pela matéria escura na forma de WIMPs (weakly interacting massive particles na sigla inglesa), usando “observatórios” com fundos radiogénicos e cosmogénicos extremamente reduzidos onde uma destas partículas pode interagir com um material alvo e produzir um sinal mensurável.

A experiência LZ é um detetor de matéria escura com 10 toneladas que deverá iniciar operações no início de 2021 e cujo principal objetivo é melhorar o atual limite de exclusão da secção eficaz de interação WIMP-nucleão por mais de uma ordem de grandeza. LZ é composto por uma câmara de projeção temporal (TPC) de duas fases de xénon e por dois detetores adicionais que envolvem a TPC e são usados como vetos, a fim de reduzirem ativamente sinais indesejados (fundos) da experiência. A sensibilidade estimada de LZ à secção eficaz da interação WIMP-nucleão independente de spin é de  $1.4 \times 10^{-48} \text{ cm}^2$  para uma WIMP de  $40 \text{ GeV } c^{-2}$  de massa. Para além da matéria escura, LZ tem o potencial de estudar, e talvez descobrir, uma grande variedade de novos processos físicos raros. A região mais interna da TPC de LZ será um dos ambientes mais “calmos” onde o estudo destes processos raros é possível.

Os fundos extremamente baixos de LZ permitem-lhe ter uma boa sensibilidade a processos raros nunca observados como o decaimento beta duplo sem emissão de neutrinos de alguns isótopos de xénon, interações de axiões ou dispersão elástica coerente neutrino-núcleo de neutrinos solares. A sensibilidade de LZ à meia-vida do decaimento beta duplo sem emissão de neutrinos do  $^{136}\text{Xe}$  é apresentada neste documento. Para uma exposição de 1360 kg·ano, a sensibilidade estimada

é de  $1.06 \times 10^{26}$  anos com um intervalo de confiança de 90%. A sensibilidade estimada para um *run* dedicado subsequente, com enriquecimento isotópico de 90% de  $^{136}\text{Xe}$  e exposição de 13.8 toneladas·ano é de  $1.06 \times 10^{27}$  anos.

O desenvolvimento de ferramentas de classificação de sinais para a cadeia de processamento de dados de LZ (LZap) é também apresentado neste documento. Estas ferramentas representam a base para classificação de sinais em LZ, tanto na forma de algoritmos heurísticos dedicados como implementações de *Machine Learning*. O HADES (Heuristics Algorithm for Discrimination of Event Substructures na sigla inglesa), desenvolvido no contexto deste trabalho, é atualmente a principal ferramenta de classificação de sinais em LZap e consegue uma exatidão global de 98.58% para todas as topologias de sinais presentes nos dados de simulação de LZ. As ferramentas de classificação *RFClassifier* e *TriNet* são duas implementações de *Machine Learning* que usam, respectivamente, um modelo de *random forests* e um *ensemble* de redes neuronais para auxiliar o desenvolvimento do HADES e potencialmente substituí-lo na cadeia de LZap. O algoritmo *RFClassifier* consegue uma exatidão de classificação de 99.37% sobre os dados simulados de LZ quando combinado com *Gaussian mixture models* (GMMs), uma técnica de *clustering* poderosa. O algoritmo *TriNet* foi treinado usando os resultados obtidos pelo HADES e consegue uma exatidão de classificação de 95.56% comparando com resultados do *clustering* com GMM, mas demonstrou que consegue generalizar os seus resultados para além do HADES.



# Acknowledgements

This work is dedicated to my mother, my father and my sister, whose endless support will always be cherished; to my family, for the constant bliss they pour into my life; and especially to Cris, whose love and support are the very foundation on which this document is built.

Special thanks to my supervisors: Francisco Neves, Alexandre Lindote and José Maneira, for the guidance along this journey.

I would like to thank the University of Coimbra and the Laboratório de Instrumentação e Física Experimental de Partículas (LIP) for hosting this work. A very special thanks to the LIP Dark Matter Group for their support and tireless help, and especially to Professor Doctor Maria Isabel Silva Ferreira Lopes for inviting me into this journey.

I would also like to thank the LZ family for their inclusion and acceptance, and for the opportunity to work on the frontiers of physics and technology.

One last appreciation to the Jury for the effort of reading this thesis.

This work is supported by a Ph.D. scholarship within the DAEPHYS doctoral programme and funded by the Fundação para a Ciência e Tecnologia (FCT), with the reference PD/BD/114114/2015. My gratitude goes towards FCT and the DAEPHYS doctoral programme for the opportunity provided.

**Cofinanciado por:**



In addition, I would like to thank the BigDataHEP project (POCI/01-0145-FEDER-029147 PTDC/FIS-PAR/29147/2017, financiado por fundos OE/FCT, Lisboa2020, Compete2020, Portugal 2020, FEDER).



# Contents

<b>Abstract</b>	<b>ii</b>
<b>Resumo</b>	<b>iv</b>
<b>List of Contents</b>	<b>x</b>
<b>1 Overview</b>	<b>1</b>
<b>2 Dark Matter</b>	<b>5</b>
2.1 The $\Lambda$ CDM Model of the Universe . . . . .	6
2.2 The Evidences of Dark Matter . . . . .	9
2.2.1 Cosmic Microwave Background . . . . .	10
2.2.2 Baryon Acoustic Oscillations . . . . .	11
2.2.3 Big Bang Nucleosynthesis . . . . .	13
2.2.4 Large-scale Structure of the Universe . . . . .	15
2.2.5 Gravitational Lensing and Collision of Clusters . . . . .	17
2.2.6 The Agreement of the Different Evidences . . . . .	19
2.3 Dark Matter Candidates . . . . .	19
2.3.1 MACHOs . . . . .	19
2.3.2 Neutrinos . . . . .	20
2.3.3 Axions and Axion-like Particles . . . . .	21
2.3.4 WIMPs . . . . .	22
2.3.5 Modified Gravity . . . . .	23
2.4 Search and Detection of Dark Matter . . . . .	24
2.4.1 Production at Particle Colliders . . . . .	24
2.4.2 Indirect Detection of Annihilation or Decay Products . . . . .	25

2.4.3	Direct Detection of Dark Matter Scattering . . . . .	26
2.5	The Canonical Halo model . . . . .	27
2.6	WIMP Interaction Cross Section . . . . .	28
2.6.1	Nuclear Form Factor Correction . . . . .	29
2.6.2	Event Rate . . . . .	30
<b>3</b>	<b>Neutrinos</b>	<b>33</b>
3.1	Neutrino Hypothesis and Discovery . . . . .	34
3.2	Flavour Oscillations and Neutrino Mass . . . . .	36
3.2.1	Discovery of Neutrino Oscillations . . . . .	36
3.2.1.1	The Solar Neutrino Problem . . . . .	36
3.2.1.2	The Atmospheric Neutrino Anomaly . . . . .	38
3.2.2	The Mechanism of Neutrino-Mass-Induced Flavour Oscillations . . . . .	40
3.2.3	Current results from oscillations . . . . .	42
3.2.4	Current constraints on the neutrino masses . . . . .	43
3.3	Theoretical Aspects of Neutrino Masses . . . . .	45
3.3.1	Neutrino Mass Models . . . . .	46
3.3.2	Leptogenesis . . . . .	50
3.4	Double Beta Decay . . . . .	51
3.4.1	Neutrinoless Double Beta Decay . . . . .	52
3.4.1.1	Experimental signature of $0\nu\beta\beta$ decay . . . . .	55
3.4.1.2	Twin isotopes and probing the $0\nu\beta\beta$ decay mechanism . . . . .	56
3.4.1.3	Past, current and future searches for the $0\nu\beta\beta$ decay . . . . .	57
3.4.2	The importance of $0\nu\beta\beta$ decay searches . . . . .	57
<b>4</b>	<b>The LUX-ZEPLIN experiment</b>	<b>59</b>
4.1	Dual-phase Xenon TPC . . . . .	62
4.1.1	TPC Operation Principle . . . . .	64
4.1.2	Energy Reconstruction and Light and Charge Yields . . . . .	65
4.1.3	Discrimination of Electron and Nuclear Recoils . . . . .	69
4.2	The Liquid Xenon Skin System . . . . .	70
4.3	The Outer Detector System . . . . .	70
4.4	Background Modelling and Simulations . . . . .	72
<b>5</b>	<b>Sensitivity of a <math>0\nu\beta\beta</math> decay search on <math>^{136}\text{Xe}</math> with LZ</b>	<b>75</b>
5.1	Assumptions About Detector Performance . . . . .	75

5.1.1	Energy Resolution at $Q_{\beta\beta}$ . . . . .	76
5.1.2	Position Resolution, Multi-site Interaction Identification and Minimal Vertex Separation . . . . .	77
5.1.3	Impact of S2 Signal Saturation . . . . .	78
5.1.4	Veto Systems . . . . .	79
5.2	Detector Calibration . . . . .	80
5.3	Background Model . . . . .	81
5.3.1	Detector Components . . . . .	84
5.3.2	Davis Cavern Walls . . . . .	86
5.3.2.1	Simulation strategy for external $\gamma$ -ray sources . . . . .	88
5.3.2.2	High-energy $\gamma$ -rays from the cavern walls . . . . .	89
5.3.3	Internal Radon . . . . .	90
5.3.4	Internal $^{137}\text{Xe}$ . . . . .	92
5.3.4.1	Muon-induced neutron activation of $^{136}\text{Xe}$ . . . . .	92
5.3.4.2	Radiogenic neutron activation of $^{136}\text{Xe}$ . . . . .	93
5.3.5	Physics Backgrounds . . . . .	94
5.3.5.1	Neutrino-electron elastic scattering . . . . .	95
5.3.5.2	Charged-current neutrino-nucleus interactions . . . . .	95
5.3.5.3	$2\nu\beta\beta$ decay of $^{136}\text{Xe}$ . . . . .	96
5.4	Sensitivity Projection . . . . .	97
5.4.1	Sensitivity Projection with Natural Abundance of $^{136}\text{Xe}$ . . . . .	100
5.4.2	Projection with 90% $^{136}\text{Xe}$ Enrichment . . . . .	101
5.5	Conclusions . . . . .	103
<b>6</b>	<b>Pulse Classification in LZ</b> . . . . .	<b>105</b>
6.1	LZap - the LZ Analysis Programme . . . . .	106
6.1.1	The Architecture of LZap . . . . .	107
6.1.1.1	PhotonDetection processing chain . . . . .	109
6.1.1.2	InteractionDetection processing chain . . . . .	112
6.2	Pulse Classification in LZap . . . . .	114
6.2.1	Main Pulse Topologies in LZ . . . . .	116
6.2.2	Classification Efficiency and Other Metrics of Performance . . . . .	122
6.2.3	Pulse Classifier HADES . . . . .	124
6.2.3.1	Construction and tuning of HADES . . . . .	127

6.2.3.2	Overall efficiency of HADES . . . . .	128
6.2.3.3	Limitations of HADES and future improvements . . . . .	132
6.2.4	Pulse Classifier COMPACT . . . . .	133
6.2.4.1	Results from COMPACT . . . . .	135
6.2.4.2	Final remarks on COMPACT . . . . .	136
<b>7</b>	<b>Machine Learning methods for Classification in LZap</b>	<b>139</b>
7.1	Introduction to Data Science . . . . .	140
7.1.1	Artificial Intelligence and Machine Learning . . . . .	143
7.1.2	The Learning Process . . . . .	144
7.1.3	A Crisis of Reproducibility and the Role of ML in Science . . . . .	146
7.2	Overview of the LZap Dataset . . . . .	147
7.2.1	Data-driven Limitations to ML Algorithms . . . . .	150
7.3	Clustering Analysis of LZap Data . . . . .	155
7.3.1	Gaussian Mixture Models . . . . .	157
7.3.2	Clustering of LZap Data using GMMs . . . . .	162
7.4	Random Forest Classifier . . . . .	168
7.4.1	Decision Tree Models . . . . .	168
7.4.1.1	Bootstrap aggregation and random forests . . . . .	172
7.4.1.2	Boosting and boosted decision trees . . . . .	174
7.4.2	The RFClassifier Pulse Classification Tool . . . . .	177
7.4.2.1	Feature importance ranking . . . . .	184
7.5	Neural Network Ensemble Classifier . . . . .	187
7.5.1	Artificial Neural Networks . . . . .	187
7.5.2	TriNet Classifier . . . . .	203
7.5.2.1	Alternative methods of training and model implementation . . . . .	224
7.6	Comparison Between Techniques . . . . .	225
7.7	Final Remarks . . . . .	226
<b>8</b>	<b>Conclusions</b>	<b>229</b>
	<b>References</b>	<b>261</b>
<b>A</b>	<b>Decay Schemes and Decay chains</b>	<b>281</b>
<b>B</b>	<b>Reduced Quantities (RQ) of LZap data</b>	<b>293</b>
<b>C</b>	<b>Useful Nomenclature</b>	<b>301</b>

# Symbols and Abbreviations

$\Lambda$ CDM	$\Lambda$ Cold Dark Matter model
$0\nu\beta\beta$	Neutrinoless double beta decay
$0\nu ECEC$	Neutrinoless double electron-capture
$0\nu EC\beta^+$	Neutrinoless electron capture with positron emission
$0\nu\beta^+\beta^+$	Neutrinoless double positron emission
$2\nu\beta\beta$	Two-neutrino double beta decay
$2\nu ECEC$	Two-neutrino double electron capture
$2\nu EC\beta^+$	Two-neutrino electron capture with positron emission
$2\nu\beta^+\beta^+$	Two-neutrino double positron emission
<b>2fDC</b>	Two-fold coincidence dark count
<b>3fDC</b>	Three-fold coincidence dark count
<b>AdaBoost</b>	Adaptive Boosting
<b>AdaGrad</b>	Adaptive Gradient Algorithm
<b>Adam</b>	Adaptive Moment estimation
<b>Adamax</b>	Adam variant based on the infinity norm
<b>ADMX</b>	Axion Dark Matter eXperiment
<b>AF</b>	Analysis Framework
<b>AGI</b>	Artificial General Intelligence
<b>AI</b>	Artificial Intelligence

<b>ANN</b>	Artificial Neural Network
<b>AP</b>	Afterpulsing
<b>ALP</b>	Axion-like particle
<b>BAO</b>	Baryon Acoustic Oscillations
<b>BBN</b>	Big Bang Nucleosynthesis
<b>BDT</b>	Boosted Decision Tree
<b>BG</b>	Background
<b>CAST</b>	CERN Axion Solar Telescope
<b>CDM</b>	Cold Dark Matter
<b>CL</b>	Confidence Level
<b>CMB</b>	Cosmic Microwave Background
<b>CNN</b>	Convolutional Neural Network
<b>COMPACT</b>	Comprehensive Pulse Analysis and Classification Tool
<b>CWoLA</b>	Classification Without Labels
<b>DBSCAN</b>	Density-Based Spatial Clustering of Applications with Noise
<b>DER</b>	Detector Electronics Response
<b>DM</b>	Dark Matter
<b>DOG</b>	Difference of Gaussians
<b>DPF</b>	Data Processing Framework
<b>DRU</b>	Differential Rate Unit
<b>DSG</b>	Dwarf Spheroidal Galaxy
<b>ELU</b>	Exponential Linear Unit
<b>EM</b>	Expectation-Maximization
<b>ER</b>	Electron Recoil
<b>FGST</b>	Fermi Gamma-Ray Space Telescope
<b>GdLS</b>	Gadolinium-loaded Liquid Scintillator
<b>GMM</b>	Gaussian Mixture Model



<b>GR</b>	General Relativity
<b>HADES</b>	Heuristic Algorithm for Discrimination of Event Substructures
<b>HPC</b>	High Performance Computing
<b>IF</b>	Isolation Forest
<b>IGM</b>	Intergalactic Medium
<b>LHC</b>	Large Hadron Collider
<b>LS</b>	Liquid Scintillator
<b>LSP</b>	Lightest supersymmetric particle
<b>LUX</b>	Large Underground Xenon
<b>LXe</b>	Liquid Xenon
<b>LZ</b>	LUX-ZEPLIN
<b>LZap</b>	LZ Analysis Programme
<b>MENACE</b>	Machine Educable Noughts And Crosses Engine
<b>MDC</b>	Mock Data Challenge
<b>MOND</b>	Multi-Layer Perceptron
<b>MLP</b>	Modified Newtonian Dynamics
<b>MSSM</b>	Minimal Supersymmetric Standard Model
<b>Nadam</b>	Nesterov-accelerated Adaptive Moment Estimation
<b>PPV</b>	Positive Predictive Value
<b>NR</b>	Nuclear Recoil
<b>OOB</b>	Out-of-Bag
<b>PDF</b>	Probability Density Function
<b>PHE</b>	Photoelectron
<b>PHD</b>	Photons detected
<b>PLR</b>	Profile Likelihood Ratio
<b>PMT</b>	Photomultiplier Tube
<b>NPV</b>	Negative Predictive Value

<b>PTFE</b>	Polytetrafluoroethylene
<b>ReLU</b>	Rectified Linear Unit
<b>RF</b>	Random Forest
<b>RMSprop</b>	Root Mean Squared Propagation
<b>RQ</b>	Reduced Quantity
<b>SAMME</b>	Stagewise Additive Modeling using a Multi-class Exponential loss function
<b>SD</b>	Spin-dependent
<b>SDSS</b>	Sloan Digital Sky Survey
<b>SGD</b>	Stochastic Gradient Descent
<b>SI</b>	Spin-independent
<b>SM</b>	Standard Model
<b>SLAC</b>	Stanford Linear Accelerator Center
<b>SLP</b>	Single Layer Perceptron
<b>SURF</b>	Sanford Underground Research Facility
<b>SUSY</b>	Supersymmetry
<b>TPC</b>	Time Projection Chamber
<b>WIMP</b>	Weakly Interactive Massive Particle

# Chapter 1

## Overview

According to the most widely accepted models of cosmology, Dark Matter and Dark Energy are responsible for roughly 95% of the total mass-energy content of the Universe [CK09]. Furthermore, cosmological observations indicate that dark matter dominates over regular (baryonic) matter by a factor of more than 5 in the total matter content of the Universe, and is a fundamental ingredient in the formation of the cosmic structures observed today, both at large and small distance scales.

The overwhelming number of astronomical evidences extends through several different observables, from galactic rotation curves and galactic motion in clusters to formation of large scale structures, cosmic microwave background (CMB) fluctuations, relative abundance of light elements and gravitational lensing [BCnG10]. These evidences, obtained over the past 80 years, led to the establishment of dark matter detection as one of the primary goals in physics research of the XXI century.

A large number of dark matter experiments have been built, especially in the past three decades, either aiming to directly detect interactions within a target material, observing the annihilation or decay products of dark matter particles or by trying to produce it at a particle collider [Sch11, JKG96, A<sup>+</sup>13a, CHL15]. Some tonne-scale dark matter direct detection experiments using dual-phase noble element technology have already completed their science runs [A<sup>+</sup>17k, W<sup>+</sup>20] and other multi-tonne scale experiments will start collecting data in the near future [A<sup>+</sup>16a, A<sup>+</sup>20k]. The LUX-ZEPLIN (LZ) detector [A<sup>+</sup>20g] is one of these upcoming detectors already in the final stages of installation. It will feature a 10 tonne dual-phase xenon time projection chamber (TPC) with ultra-low background requirements aimed to detect, or at least further exclude, dark matter in the form of weakly interacting massive particles (WIMPs) by progressively constraining the WIMP-nucleon interaction cross section. LZ will also look for dark matter in the form of axions, axion-like particles and other possible candidates, as well as search for other rare events such as

neutrino interactions and rare decays of some xenon isotopes.

The LZ collaboration resulted of the merging of two remarkable dark matter detection experiments, LUX and ZEPLIN. The ZEPLIN programme dates back over twenty years, and operated in the Boulby mine, North Yorkshire, United Kingdom. The ZEPLIN-II and ZEPLIN-III detectors were pioneers in the use of dual-phase noble gas TPC technology for rare event searches, a technology widely used today by several DM experiments, including LUX and LZ. ZEPLIN-III established the potential of dual-phase xenon TPCs in discriminating between electron recoils (ER) and nuclear recoils (NR) based on the combined analysis of the scintillation and ionization channels, with ZEPLIN-III achieving a remarkable  $>99.99\%$  ER rejection due to its strong electric field [A<sup>+</sup>07a]. LUX was the most sensitive dark matter direct detection experiment from 2013, the year when it completed the first science run, to 2017, reaching a lower limit for the WIMP-nucleon spin-independent scattering cross section of  $1.1 \times 10^{-46} \text{ cm}^2$  (0.11 zb) for a WIMP mass of  $50 \text{ GeV c}^{-2}$  [A<sup>+</sup>17h]. The LZ detector is expected to reach a sensitivity to the spin-independent WIMP-nucleon cross section of  $1.4 \times 10^{-48} \text{ cm}^2$  at 90% confidence level for a  $40 \text{ GeV c}^{-2}$  WIMP [A<sup>+</sup>20d], an improvement over the limit set by LUX by two orders of magnitude and over the current best limit by a factor of 30 [A<sup>+</sup>18i].

The inner region of the LZ detector will be an incredibly quiet laboratory where it will be possible to observe some rare physical phenomena aside from dark matter interactions. Therefore LZ presents an opportunity to explore different rare processes and physics beyond the Standard Model, like the yet unobserved neutrinoless double beta decay of  $^{136}\text{Xe}$ . The observation of this decay would provide the first evidence of fundamental Majorana particles and could help to solve the mystery of matter-antimatter asymmetry observed in the universe. LZ will be in a prime position to search for this rare decay due to its low-background environment, the natural high abundance of the decaying isotope and the excellent energy resolution provided by dual-phase xenon technology. The search for this decay will be the most important physics goal of LZ besides searching for WIMP dark matter.

Coupled to the physics analyses of the data obtained with LZ is an entire data processing framework that can convert the raw output of the detector into information ready for physics analysis. The LZ analysis program (LZap) is a modular chain with a wide variety of algorithms that perform sequential low level data processing and return reduced quantities with meaningful physics information ready to be analysed. The development, tuning and testing of such algorithms follows the live-time of the experiment but it is crucial on the early commissioning stages to ensure that the processing framework is ready for science data. Among those low level processing tasks, the identification and classification of detector signals (pulses) based on their physical origin is a critical step on which the remaining analysis heavily depends. Several methods are explored to perform the low level computational tasks such as pulse classification, from heuristics algorithms to advanced data-mining techniques like Machine Learning and Deep Learning.

The search for dark matter is the main scientific goal of LZ, and thus Chapter 2 contains a brief introduction to the dark matter problem and cosmological evidences, an overview of the several

detection techniques from which it is possible to study dark matter, and a quick description of the mathematical formalism for direct detection techniques. A brief history of neutrino physics and introduction to the neutrinoless double beta decay process is given in Chapter 3. Chapter 4 formally presents the LZ dark matter experiment and provides an overview of the detector systems in which the work presented in this document is centred. The study of the sensitivity of LZ to the neutrinoless double beta decay of  $^{136}\text{Xe}$  is presented in Chapter 5. Chapter 6 presents a description of the processing framework to be used within the LZ data analysis programme and the detailed description of the pulse classification algorithms developed and currently implemented in this framework. Some advanced pulse classification algorithms and data processing tools based on Machine Learning techniques are discussed in Chapter 7. This chapter will feature extensive in-depth discussions of the techniques explored during the research and development of the classifications tools. Finally, Chapter 8 is left for some final remarks about the work done.



# Chapter 2

## Dark Matter


Several scientific evidences agree with the existence of dark matter (DM), a kind of matter that permeates the universe and does not interact via the electromagnetic or strong forces [CK09, S<sup>+</sup>10, CBG<sup>+</sup>06]. Dark matter and dark energy account for roughly 95% of the total mass-energy content of the universe, but the nature of both these components is still unknown. Nevertheless, the effects of these two dark components on cosmological structures and across different distance scales are evident.

The first clue for the existence of dark matter was found by Fritz Zwicky in 1933, where he noticed that there was a large discrepancy in the mass of the Coma cluster (Abell 1656) when calculated using galactic motion or inferred from the luminous matter. He proposed that some form of non-luminous matter was present in the Coma cluster and that its effects on the motion of the galaxies accounted for the observed discrepancy [Zwi33, Zwi37]. The mass of the Coma cluster calculated by Zwicky was  $M_c = 4.5 \times 10^{13} M_\odot$ , but the mass that he could infer from luminosity measurements was almost 170 times smaller [Zwi37]. The real Coma cluster mass is now known to be closer to  $M_c \approx 2 \times 10^{15} M_\odot$  and less than 12% of this value is relative to stars and interstellar gas [Ryd17]. Zwicky called this missing mass “*dunkle materie*” – German for “dark matter”.

The dutch astronomer Jan Oort had postulated, one year before, that additional matter was necessary to explain the motion of stars in our system neighbourhood [Oor32], but his results were later understood to be due to interstellar dust structures and contributions from white dwarf stars. Nevertheless, the involvement of Oort in this subject gathered the attention of the scientific community.

Another very important contribution came from the work of Vera Rubin on the precise measurement of galaxy rotation curves. The flat galactic velocity profiles at large radii, observed

for nearly all galaxies, is a strong evidence for the existence of a massive dark matter halo that encompasses the galactic disks. Rubin observed that there is a strong disagreement between the measured velocity profiles of several galaxies and those expected from Newtonian dynamics – the velocity of stars and gas clouds remained approximately constant at large radii, as seen in Figure 2.1 for the galaxy NGC3198. A possible explanation for these flat rotation curves is the existence of a dark matter halo that reaches far beyond the galactic luminous disk that contributes gravitationally to the observed dynamics of stars. In our Milky Way galaxy, this halo is expected to compose  $\sim 88.5\%$  of its total mass [KSLB14].



**Figure 2.1:** NGC3198 rotation curve, obtained by measuring the redshift of the 21 cm emission line of neutral hydrogen gas. The upper curve represents the best fit of the velocities to the data. The curves marked with “halo” and “disk” represent the mass density contributions of the dark matter halo and galactic disk (stars, gas and dust) obtained from this fit, respectively. Picture from [vBBS85].

A recent study of the Coma cluster resulted in the discovery of more than 800 dark galaxies [KYYK15]. These galaxies have an abnormal amount of dark matter, around 98% of their total mass content. It becomes apparent from these findings that dark matter structures the size of the Milky Way are more common than previously thought, and that its presence is fundamental to understand how the large-scale structures of the cosmos were formed.

## 2.1 The $\Lambda$ CDM Model of the Universe

The observation that most extragalactic objects seem to be receding from each other, made by Georges Lemaître and Edwin Hubble in the late 1920s, led to the conclusion that the universe is expanding, an idea previously proposed by Alexander Friedmann in 1922. Under the assumption that at large scales the universe is isotropic and homogeneous (*Cosmological Principle*) and



its content behaves as a perfect fluid with energy density  $\rho(t)$  and pressure  $p(t)$  [T<sup>+</sup>18], the Friedmann equations can be derived from Einstein's field equations of General Relativity:

$$\left(\frac{\dot{a}}{a}\right)^2 = H^2 = \frac{8\pi G}{3}\rho - \frac{\kappa}{a^2} + \frac{\Lambda}{3} \quad (2.1)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}, \quad (2.2)$$

where  $a(t)$  is the cosmic scale factor that represents the size of the Universe at the time  $t$ ,  $H(t)$  is the Hubble parameter,  $\Lambda$  is the cosmological constant associated with the vacuum energy,  $\kappa$  is the curvature parameter, and  $G$  is the Newtonian gravitational constant. These equations make use of the Friedmann-Lemaître-Robertson-Walker solution to the field equations of General Relativity [Lem31, Rob35, Fri99a, Fri99b].

The critical density  $\rho_c$  is derived from Equation 2.1 considering  $\kappa = 0$  and  $\Lambda = 0$ :

$$\rho_c = \frac{3H^2}{8\pi G}. \quad (2.3)$$

The cosmological density parameter  $\Omega$  is then defined as the total energy density relative to the critical density,  $\Omega = \rho/\rho_c$ , and Equation 2.1 can be rewritten as

$$\frac{\kappa}{a^2} = H^2(\Omega - 1) + \frac{\Lambda}{3}, \quad (2.4)$$

where  $\Omega$  contains both the contributions from matter ( $\Omega_m$ ) and radiation ( $\Omega_r$ ). The contribution from the vacuum energy density can be defined as  $\Omega_\Lambda = \Lambda/3H^2$ , and Equation 2.4 becomes

$$\frac{\kappa}{a^2} = H^2(\Omega_r + \Omega_m + \Omega_\Lambda - 1). \quad (2.5)$$

At the current time  $t_0$ , the cosmic scale factor is considered to be  $a(t = t_0) = 1$  by definition, and the current rate of expansion is then given by  $H_0 = H(t = t_0)$ . The dimensionless density parameters  $\Omega_i$ , with  $i$  being each of the components considered in the model, are known as present-day density parameters [Sch15] and can be written explicitly as


$$\Omega_i = \frac{8\pi G}{3H_0^2}\rho_i \quad ; \quad \Omega_\Lambda = \frac{\Lambda}{3H_0^2}. \quad (2.6)$$

The curvature of space-time is often presented as a density parameter of the form  $\Omega_{\kappa,0} = -\kappa/a_0^2 H_0^2$  and added to the total energy density parameter, ensuring that  $\Omega_r + \Omega_m + \Omega_\Lambda + \Omega_\kappa = 1$  at the current time. The density parameters are often presented in the form  $\Omega_i h^2$ , where  $h = H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$  is the scaled Hubble parameter.

All evidences associated with Big-Bang cosmology seem to indicate that the universe today is flat ( $\Omega_\kappa \sim 0$ ) and dominated by the contributions of the cosmological constant  $\Lambda$ , associated with dark energy, and a non-baryonic, non-relativistic (cold) dark matter component (CDM) to

its total energy density, with only a small contribution from baryonic matter<sup>1</sup> [T<sup>+</sup>18, A<sup>+</sup>18a]. The contribution of radiation to the energy density today is negligible ( $\Omega_r \sim 0$ ). This is the foundation of the  $\Lambda$ CDM model, considered the Standard Model of Cosmology.

Supernova surveys can be used to determine the acceleration history of the universe and infer the relevant cosmological parameters that influence those changes in acceleration [S<sup>+</sup>12, P<sup>+</sup>97, F<sup>+</sup>01]. Figure 2.2 displays the results from several Supernovae (SNe) Ia surveys and their agreement with the  $\Lambda$ CDM model. These observations indicate that the universe is expanding faster today than in the past, meaning that the energy density associated with the cosmological constant is  $\Omega_\Lambda > 0$  and providing the first direct evidence of a non-vanishing dark energy component<sup>2</sup> [Sch15].



**Figure 2.2:** Combined data of type Ia supernovae from several surveys. The black line indicates the best fit to the cosmological parameters of the  $\Lambda$ CDM model. Figure from [S<sup>+</sup>12].


The Planck Collaboration provides the current best estimates to the values of the  $\Lambda$ CDM model [A<sup>+</sup>18a]. The latest results are summarized in Table 2.1. The current best estimate of the mass-energy content of the universe is schematically represented in Figure 2.3.

<sup>1</sup>The total matter content ( $\Omega_m$ ) contains the contributions from dark matter ( $\Omega_c$ ) and baryonic matter ( $\Omega_b$ ).

<sup>2</sup>2011 Nobel Prize in Physics for Saul Perlmutter, Brian Schmidt, and Adam Riess.

**Table 2.1:** Latest results from the Planck Collaboration for the cosmological parameters of the  $\Lambda$ CDM model. The values presented have 68% error bands and are obtained using a combined likelihood on the CMB temperature (TT) power spectrum, E-mode polarization (EE) power spectrum, high-multipole temperature-polarization cross-correlation (TE), low-multipole polarization (lowE), lensing and baryon acoustic oscillations (BAO) data [A<sup>+</sup>18a].

Parameter	Symbol	Value
Hubble constant [ $\text{km s}^{-1} \text{Mpc}^{-1}$ ]	$H_0$	$67.66 \pm 0.42$
Baryon energy density	$\Omega_b h^2$	$0.02242 \pm 0.00014$
Cold Dark Matter energy density	$\Omega_c h^2$	$0.11933 \pm 0.00091$
Total matter energy density	$\Omega_m$	$0.3111 \pm 0.0056$
Dark energy density	$\Omega_\Lambda$	$0.6889 \pm 0.0056$
Curvature	$\Omega_{\kappa,0}$	$0.0007 \pm 0.0019$
Sum of neutrino masses [eV]	$\sum m_\nu$	$< 0.12$
Age of the universe [Gy]		$13.787 \pm 0.020$



**Figure 2.3:** Cosmic recipe of our universe according to the latest results from Planck [A<sup>+</sup>18a].

Recent results from surveys of Cepheid variables<sup>3</sup> in the Large Magellanic Cloud provide a local measurement of the Hubble constant of  $H_0 = 74.03 \pm 1.42 \text{ km s}^{-1} \text{ Mpc}^{-1}$  [RCY<sup>+</sup>19], in strong disagreement ( $4.4\sigma$  significance) with the values predicted from Planck cosmic microwave background data from the early universe, considering the standard  $\Lambda$ CDM model [A<sup>+</sup>18a]. This tension on the  $H_0$  value between the early and present time could be an indication of a cosmological feature beyond the standard  $\Lambda$ CDM model.

## 2.2 The Evidences of Dark Matter

Aside from the evidences for the existence of dark matter found in galactic and cluster dynamics discussed at the beginning of this section, several important evidences are also found in the measurement of the anisotropies in the cosmic microwave background (CMB), formation of large-scale structures, measurement of baryon acoustic oscillations, Big Bang nucleosynthesis (BBN) and observation of non-luminous matter structures through gravitational lensing. All

<sup>3</sup>A Cepheid variable is a type of pulsating star whose brightness oscillation period is tightly correlated to their absolute luminosity and so it can be used as a “standard candle” [Rod57].

these evidences seem to agree with each other over the  $\Omega_m$ - $\Omega_\Lambda$  phase-space with astounding precision.

### 2.2.1 Cosmic Microwave Background

The thermal radiation left over from the early stages of the universe following the Big-Bang permeates the universe as a cosmic microwave background (CMB). This radiation is today composed of microwave wavelength photons with a mean temperature of  $2.725 \pm 0.001$  K [A<sup>+</sup>18a] due to the expansion of the universe. Discovered by chance in 1965 by the radio-astronomers Arno Penzias and Robert Wilson<sup>4</sup> of the Bell Telephone Laboratories [PW65], the CMB is one of the strongest evidences in favour of the Big Bang [T<sup>+</sup>18, Gam48, AH48].


The decoupling of photons trapped in the photon-baryon-electron plasma of the infant universe resulted in these particles permeating the universe and travelling almost unimpaired. The large temperatures of the primordial plasma resulted in photodissociation and recombination of protons and electrons to form neutral hydrogen atoms to occur at the same rate, keeping all species in thermal equilibrium. As the universe expanded (adiabatically) and the temperatures began to decrease, photons no longer had enough energies to ionize hydrogen and propagated without scattering significantly. This moment is called the “time of last scattering”, and occurred at a redshift of  $z \approx 1100$  (around 380 thousand years after the Big Bang) [PA<sup>+</sup>15]. These photons carry the information about the state of the universe at the time of decoupling, including a map of the denser and therefore hotter regions of the primordial plasma. Figure 2.4 shows the sky map of the CMB as measured by the Planck Observatory (2015 data) [PA<sup>+</sup>15].

The thermal anisotropies of the CMB are formed due to density fluctuations in the photon-baryon plasma at the time of decoupling, with hotter regions correlating to larger local density fluctuations and vice-versa. The values of the cosmological parameters at that time dictated the scale of those density fluctuations in the primordial plasma and are therefore encoded in the anisotropies of the CMB. The CMB angular power spectrum (or temperature spectrum) measures the angular scale of the fluctuations using an expansion of spherical harmonics ( $l$ ) [Sch15] and its shape can be used to constrain the cosmological parameters of the  $\Lambda$ CDM model, namely the total matter density ( $\Omega_m$ ), vacuum energy density ( $\Omega_\Lambda$ ), overall curvature ( $\Omega_k$ ) and baryon energy density ( $\Omega_b$ ). Figure 2.5 represents how the global shape of the power spectrum varies as a function of some of these parameters. Information about the non-baryonic matter density can be extracted from the CMB spectrum by measuring the total matter density and the baryonic matter independently. The relative sizes of the first, second and third peaks of the CMB power spectrum are very sensitive to the baryonic density parameter  $\Omega_b$  [Hu08]. The total matter contribution  $\Omega_m$  to the acoustic oscillations in the CMB can be obtained by fitting the overall amplitude of the power spectrum. The amount of non-baryonic matter (cold dark matter) in the universe can then be inferred considering  $\Omega_c = \Omega_m - \Omega_b$ .

The temperature power spectrum obtained in the analysis of the CMB anisotropy data from

---

<sup>4</sup>Arno Allan Penzias and Robert Woodrow Wilson received the Nobel Prize in Physics 1978 “for their discovery of cosmic microwave background radiation”.




**Figure 2.4:** Mollweide projection of the CMB map measured by the Planck satellite (2015 data), with monopole (mean temperature) and dipole (Doppler shift) contributions removed. Figure from [Col15].

Planck 2018 results is represented in Figure 2.6. A summary of the 2018 results from the Planck collaboration [A<sup>+</sup>18a] for the  $\Lambda\text{CDM}$  cosmological parameters is presented in Table 2.1. These results, together with the results obtained from baryon acoustic oscillations and CMB polarization and gravitational lensing measurements, indicate that a non-baryonic matter component is responsible for 84% of the matter content and 26% of the total mass-energy content of the universe.

### 2.2.2 Baryon Acoustic Oscillations

Baryon acoustic oscillations (BAO) are pressure waves formed in denser regions within the primordial plasma that seeded the formation of dense cosmological structures as the universe expanded, resulting in the large scale structures observed today [SZ70]. In the primordial plasma, denser regions gravitationally attract more matter, and the resulting gravitational collapse produces an increase in temperature at the center of those regions. This produces an increase in the outward pressure from radiation due to it being strongly coupled to baryonic matter. The repulsive effect of the radiation pressure counteracting the gravitational collapse of the system results in oscillations in the fluid that produce a baryon-photon acoustic wave moving away from the density center. The weakly-interacting dark matter is not affected by this outward radiative pressure and remains in the center of the density perturbation. The outward-moving baryon-photon wave front produces a density shell, resembling the effect of a shock wave in the fluid. When the fluid temperature decreases, due to expansion, and the photons decouple from the baryons, the pressure on the acoustic wave is relieved, leaving the shell of baryonic material




**Figure 2.5:** Illustrative sensitivity of the acoustic peaks of the CMB theoretical power spectrum to the baryon energy density  $\Omega_b h^2$  (left), total matter energy density  $\Omega_m h^2$  (center) and dark energy density  $\Omega_\Lambda h^2$  (right), as predicted by the  $\Lambda$ CDM model. The default cosmological parameters considered here are  $\Omega_k = 0$ ,  $\Omega_\Lambda = 0.65$ ,  $\Omega_b h^2 = 0.02$  and  $\Omega_m h^2 = 0.147$ , not the current best estimates. Figure adapted from Reference [HD02].

at a characteristic defined radius called the sound horizon [ESW07].

Figure 2.7 represents the evolution of the acoustic oscillation. The wave front of the baryon-photon fluid can be seen in the upper-left panel, streaming away from the density center where most of the dark matter is concentrated. After the photon decouple from the baryons (upper-right panel), the pressure of the acoustic wave ceases and the baryonic acoustic wave stalls behind, leaving an excess of baryonic matter at the sound horizon distance from the density fluctuation. Due to mutual gravitational influences, over time the dark matter distribution equalizes with the distribution of baryonic matter and the mass distribution profiles become similar [ESW07].

The larger matter density distribution characterized by the acoustic peak effectively seeded the formation of large scale structures, meaning that a larger number of structures are expected to be found separated by a distance similar to the sound horizon than by smaller length scales. Statistically analysing the separation of a great number of galaxies and measuring the abundance of intergalactic  $\text{H}_2$  gas by redshifted Lyman- $\alpha$  emissions, *i.e.*, the absorption of distant quasar light on neutral hydrogen along the light path that produces a redshifted set of spectral lines, it is possible to measure the sound horizon length scale [SIK<sup>+</sup>13], as it is demonstrated in the last panel of Figure 2.7. The peak is obtained using a two-point distance correlation function that represents the average amount of matter separated by a distance  $r$ , convoluted with all points in space [SIK<sup>+</sup>13]. The baryonic acoustic peak at  $100 - 130 h^{-1}$  Mpc scale is clearly visible.


The length-scale of the baryon-acoustic peak is sensitive to the acceleration rate of the universe and the baryon fraction of the total matter content of the universe. The fraction of baryonic matter in the total matter content obtained using BAO is  $\Omega_b/\Omega_m \approx 0.2115$  [T<sup>+</sup>06], consistent with data from CMB and BBN. Recent results from galaxy surveys [A<sup>+</sup>17i] and Lyman- $\alpha$  forest measurements [PD<sup>+</sup>15] also support these results.



**Figure 2.6:** Planck temperature power spectrum of the CMB, obtained using multipolar analysis over the CMB map. The blue line represents the  $\Lambda$ CDM model theoretical spectrum prediction that best fits the data. The first acoustic peak is clearly visible at multipole  $l \sim 200$ , followed by the secondary peaks and damping tail, at higher  $l$ . Figure from Reference [A<sup>+</sup>18a].

### 2.2.3 Big Bang Nucleosynthesis

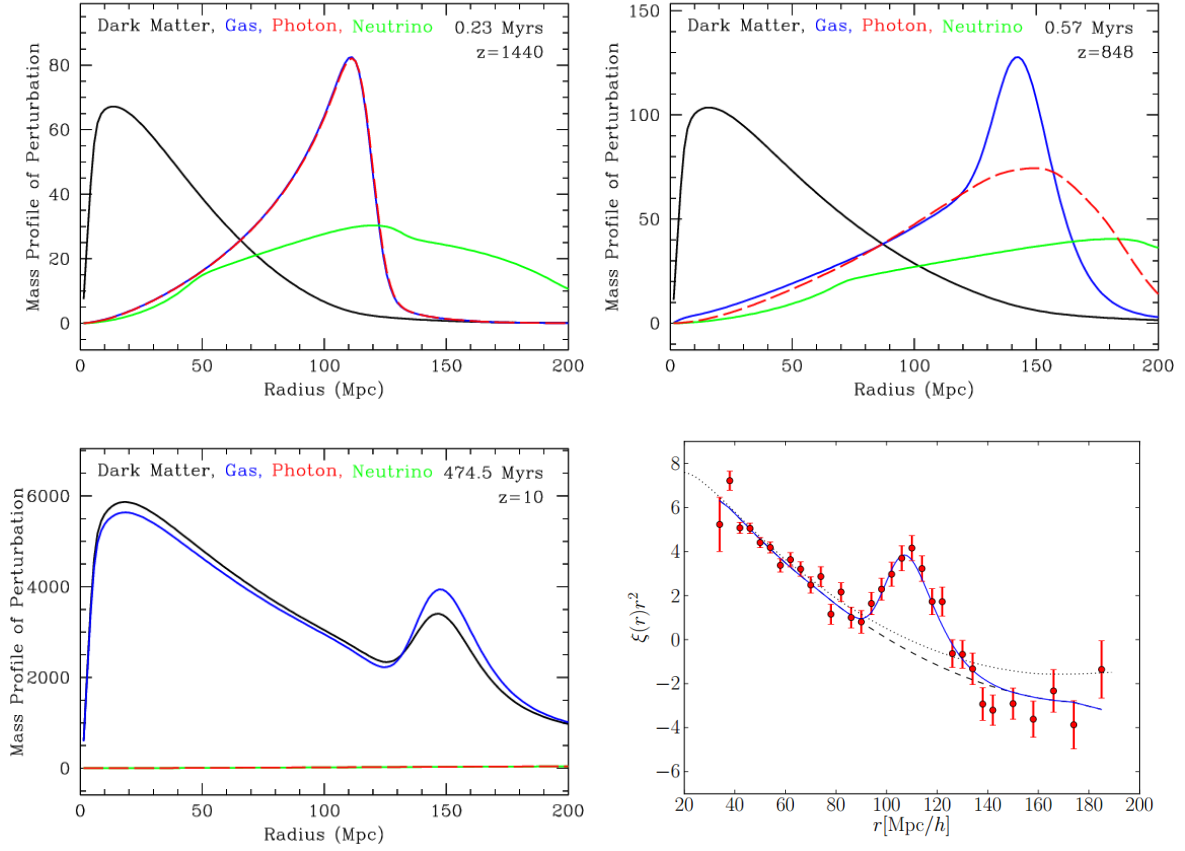
The production of light elements seconds after the Big Bang, called Big Bang Nucleosynthesis (BBN), began when the temperature of the primordial baryon plasma dropped below  $T \sim 1$  MeV and nuclear photodissociation ceased, being essentially over in just a few minutes [T<sup>+</sup>18]. Before this period, the free protons and neutrons were in thermal equilibrium through the reactions presented in Equations 2.7 to 2.9 [Sch15].



At a time  $t \sim 1$  s the primordial plasma temperature decreased sufficiently and the neutron-proton conversion rate became lower than the Hubble expansion rate, “freezing out” the neutron-proton fractions at around  $n/p \simeq 1/6$  [T<sup>+</sup>18]. As free neutrons naturally decay with a 885 s half-life [End00], after neutrinos effectively decouple from the baryons the fraction of neutrons to protons decreases to  $n/p \simeq 1/7$ .

After the “freeze out”, sequential nuclear fusion reactions involving the leftover protons and neutrons form light nuclei. The first and most important step in primordial nucleosynthesis is the formation of deuterium (D) which is described by Equation 2.10 [Sch15].





**Figure 2.7:** Evolution of an acoustic wave in the primordial plasma (upper-left), showing the different species present and their behaviour (mass distribution) over different moments: baryons (named “gas”, blue), dark matter (black), photons (red), and neutrinos (green). The upper-right panel represents the decoupling of photons and baryons, where photons disperse and the baryonic peak stalls. Dark matter and baryonic matter then rearrange due to gravity, resulting in the merging of the respective curves as portrayed in the lower-left figure [ESW07]. The lower-right figure shows the baryon acoustic peak measured from the Lyman- $\alpha$  forest [SIK<sup>+</sup>13]. The peak in the correlation function is clearly visible at length scale  $r = 100 - 130 h^{-1}$  Mpc.


After some minutes, deuterium and free neutrons and protons begin to form other nuclei, starting with  ${}^3\text{He}$  and  ${}^4\text{He}$ . The main reactions for generating helium nuclei are presented in Equations 2.11 to 2.15 [Ryd16].



where  $T$  represents tritium ( ${}^3\text{H}$ ) that is produced in small quantities by deuterium reactions. Tritium and  ${}^3\text{He}$  are quickly converted to  ${}^4\text{He}$  via the reactions presented in Equations 2.16 to



2.19.



The creation of nuclei heavier than helium is highly suppressed due to the large binding energy of the  ${}^4\text{He}$  nucleus, meaning that the primordial fraction of  ${}^4\text{He}$ ,  $Y_p$ , remains almost unchanged. After the production of  ${}^3\text{He}$ ,  ${}^4\text{He}$ , deuterium and tritium, the heavier nuclei  ${}^6\text{Li}$ ,  ${}^7\text{Li}$ ,  ${}^7\text{Be}$  start being produced as well.


The current estimates for the nuclei abundance yields are  $D/H = (2.569 \pm 0.027) \times 10^{-5}$ ,  $Y_p = 0.245 \pm 0.003$  and  ${}^7\text{Li}/H = (1.6 \pm 0.3) \times 10^{-10}$  [T<sup>+</sup>18]. The ratio  $D/H$  is sensitive to the baryon-to-photon ratio  $\eta$ , an important cosmological parameter that influences the early or late beginning of BBN. Figure 2.8 compares the predicted primordial abundances of light elements, according to the standard model of Big-Bang nucleosynthesis, with the observed abundances. Comparing the measured abundances of the light elements with the expected from BBN, a range of values of  $(5.6 < \eta < 6.6) \times 10^{-10}$  is obtained. This baryon-to-photon ratio provides a measure of the baryon energy density of  $0.021 \leq \Omega_b h^2 \leq 0.024$  (95% CL) [T<sup>+</sup>18], in agreement with the results obtained from the CMB [A<sup>+</sup>18a]. Considering that the most recent CMB estimates of the total matter energy density point to  $\Omega_m \approx 0.3$ , this result implies that most of the matter in the universe is non-baryonic.

### 2.2.4 Large-scale Structure of the Universe

Surveys of distances of galaxies and clusters can be used to construct detailed maps of the matter distribution in the local universe. These observations can then be compared to simulation models of structure formation to determine the conditions and cosmological parameters needed to reproduce the visible features of the large scale structure of the universe. Figure 2.9 shows the distribution of galaxies measured by the Sloan Digital Sky Survey [R<sup>+</sup>10] up to a comoving distance of  $315 \text{ Mpc h}^{-1}$  (left) and two time slices of the Millennium-II simulation [BKS<sup>+</sup>09] assuming  $\Omega_m = 0.25$  and  $\Omega_\Lambda = 0.75$ . The bottom-right image of the Millennium simulation represents the same distance scale as the red slice over the left image of the Sloan Digital Sky Survey. A comparison of both images shows that the simulated model is able to reproduce the observed features of matter distribution over large scales.

In order to match simulated models of structure formation with the observed content of galactic haloes, galaxy cluster dynamics and large-scale structures, a weakly-interacting cold dark matter component in excess of baryonic matter must be present since the early universe [S<sup>+</sup>10, Sch15]. This conclusion is in line with those drawn from BAO and CMB data.


Despite the success in reproducing the filament-like matter structures at the correct scales,



**Figure 2.8:** Predicted 95% CL bands of the primordial abundances of  ${}^4\text{He}$  (top pink), D (center blue),  ${}^3\text{He}$  (center red), and  ${}^7\text{Li}$  (bottom green) as a function of the baryon-to-photon ratio  $\eta$ . The boxes over each band indicate the observed range of values for the light element abundances. The vertical blue dashed band represents the 95% CL measurement of the baryon density from the CMB, and the vertical red dashed band indicates the 95% CL constraints on the baryon content from BBN considering the observed abundances of D and  ${}^4\text{He}$  only. Figure from [T<sup>+</sup>18].

there are some observations that are not in agreement with most simulations, namely the lack of abundance of dwarf galaxies predicted by most models, the observed low mass of dwarf galaxies and the flat dark matter profile at the center of galaxies [S<sup>+</sup>10]. These problems can be partially solved by assuming a non-zero density of warm or hot dark matter, *e.g.*, neutrinos, mixed with the dominant cold dark matter component.

Since the abundance and distribution of neutral hydrogen in the intergalactic medium (IGM) is expected to follow the underlying mass distribution of large-scale structures, the Lyman- $\alpha$  forest can provide estimates to the underlying mass distributions and constraints on the amount of warm dark matter. They can also be used to measure the sound horizon as mentioned in Section 2.2.2. However, the detailed results from galaxy surveys and Lyman- $\alpha$  absorption spectroscopy seem to miss a large fraction of the expected baryonic density from CMB data [S<sup>+</sup>10]. Measurements of the dispersion level of localised fast radio bursts can be used to estimate the electron column density in the warm-hot intergalactic medium, thus providing evidence that the missing baryons are in the form of fully ionized hydrogen. These findings result in a baryonic matter density of  $\Omega_b = 0.051^{+0.021}_{-0.025} h^{-1}$ , consistent with the expected baryon density from the CMB and BAO data [M<sup>+</sup>20].




**Figure 2.9:** Comparison of the measured distribution of large-scale structures of galaxies from the Sloan Digital Sky Survey (left) and the Millennium-II simulations for a redshift of  $z = 6.2$  (top) and today (bottom), assuming a cosmological model with  $\Omega_m = 0.25$  and  $\Omega_\Lambda = 0.75$ . The side length of the Millennium-II boxes is  $100 \text{ Mpc h}^{-1}$  in comoving units. The same distance scale is represented in the left image by the red slice centred on the observer.

### 2.2.5 Gravitational Lensing and Collision of Clusters

According to General Relativity, the presence of a massive object generates local distortions of space-time. Dense matter distributions that happen to be in the optical path of light from a distant source will distort space, resulting in the formation of a distorted image of the light source from the perspective of the observer [S<sup>+</sup>10, Sch15]. These effects, referred to as “gravitational lensing”, are observed in two distinct regimes: strong and weak lensing. Strong gravitational lensing occurs when the apparent luminosity and shape of astronomical sources is modified by the gravitational effects of a foreground mass, resulting in visible distortions of the image of the light source. Weak gravitational lensing results in only slight distortions of the images of the background light sources, and the effect can only be perceived using statistical analysis [Mel99]. In both cases the overall distribution of matter distorting the background light, as well as an estimate of its mass, can be obtained. Microlensing is a distinct gravitational lensing effect caused by the transit of a massive and compact object in front of a distant light source that results in a sudden increase of luminosity of the source. Several surveys of strong and weak gravitational lensing of galaxies and galaxy clusters indicate that these structures are dominated by non-luminous matter, with this matter exceeding the baryonic matter by around an order of magnitude for galaxies and by two orders of magnitude for clusters [FSW05, GTR<sup>+</sup>07, FSB08].

The observation of collisions of clusters of galaxies provides one of the most compelling evidence for the existence of dark matter. Combining X-ray spectrometry and gravitational lensing analysis allows the precise measurements of the nature and abundance of the constituents of the

clusters, as well as infer their gravitational properties from a transversal observation viewpoint. The analysis performed to the Bullet Cluster (1E 0657-558) was of great importance for dark matter research [CBG<sup>+</sup>06]. The collision of two clusters into what is now known as Bullet Cluster is displayed in Figure 2.10. The contributions of the X-rays emitted by the hot intergalactic gas (pink) and that of the mass that would generate the weak gravitational lensing observed (blue) are superimposed to the image in the visible region of the spectrum.



**Figure 2.10:** Image of the Bullet Cluster where the contributions of the intergalactic gas obtained from its X-ray emission (pink) and the majority of the matter in the cluster inferred from weak gravitational lensing (blue) are represented. Credit to: X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al [NAS06].

A simple inspection of the composite image reveals that the hot intergalactic gas has a clear position offset relatively to the galaxies in the cluster plus the majority of its mass content (generating the gravitational lensing). As the two clusters collided, the hot interstellar gas decelerated due to friction and heated up to higher temperatures, while the dispersed galaxies passed by each other without colliding because of the large average distances between them. However, the majority of the mass in each cluster is observed to accompany the luminous galaxies without any hindering. The overall lensing effect observed cannot be completely justified by the mass of the luminous (baryonic) matter in the cluster, meaning that a large fraction of the mass of the individual clusters is attributed to non-luminous matter. Whatever matter comprises the majority of the mass content of the clusters does not interact significantly with itself neither with regular matter, justifying its unimpaired motion. These conclusions imply that most of the mass of the cluster is not from its baryonic content. The Bullet Cluster is not an isolated case where these phenomena are observed. Several cluster collisions analysed (MACS J0416.1-2403, MACS J0152.5-2852, MACS J0717.5+3745, Abell 370, Abell 2744 and ZwCl 1358+62) seem to be in agreement with the findings in the Bullet Cluster.

Alternative explanations of the observed cluster collision events involving some modification of

General Relativity and no dark matter have been proposed [BM07, IM18] but no model has stood capable to explain all observations in a unified and generalizable manner. Despite the strong evidences for a non-interacting, non-luminous matter provided by the Bullet Cluster and other cluster collisions, this alone is not enough to provide definitive proof of the existence of a dominant non-baryonic dark matter in the universe. The strength of the dark matter claim comes from the strong agreement of all the evidences in favour of it.

### 2.2.6 The Agreement of the Different Evidences

Figure 2.11 shows the confidence level contours in the  $\Omega_m - \Omega_\Lambda$  space from SNe Ia, CMB and BAO data prior to 2012. The various contours converge into a single confidence region with an excellent agreement between the data. This image is representative of the agreement between a wide number of observations and, despite not providing the most recent results, it displays how well the  $\Lambda$ CDM model fits observations and stands as the standard model of cosmology<sup>5</sup>. The evidences presented in this Section all point to the existence of a weakly-interacting, non-baryonic matter density that dominates the matter content of the universe. The last piece of the dark matter puzzle is to understand its nature. To do so, several experiments using different techniques have been searching for dark matter, either directly or indirectly, as explained in more detail in Section 2.4.

## 2.3 Dark Matter Candidates

### 2.3.1 MACHOs

MACHO stands for Massive Astrophysical Compact Halo Object. These are compact, non-luminous objects of regular baryonic matter that could explain the unaccounted mass in astrophysical systems [T<sup>+</sup>07]. Objects like neutron stars, white dwarfs, brown dwarfs, Jupiter-like planets or black holes<sup>6</sup> contribute to a galactic MACHO density.


Limits on the number density of these compact objects in our galaxy can be obtained by gravitational microlensing (see Section 2.2.5). The transit of such a compact object in front of a star would result in a measurable increase of its luminosity [S<sup>+</sup>10]. In order to determine the MACHO contribution to dark matter, the EROS-2 collaboration surveyed 7 million stars, only one of which presented a microlensing candidate. This constrains the contribution of MACHOs to dark matter at less than 4% to low mass stars and planets, and less than 1% to objects with approximately the mass of the Sun [T<sup>+</sup>07].

MACHOs fail to explain simple astronomical observations, such as the Bullet Cluster or BAO. Furthermore, given enough time, this non-luminous matter would clump together into enormous

---

<sup>5</sup>There are still some disagreements in recent measurements of the Hubble constant between the CMB and distant Supernovae, and recent surveys of Cepheid variables, as explained at the beginning of this Chapter.

<sup>6</sup>Not quite of baryonic nature, but still a valid contribution to the mass of a system from baryonic sources.



**Figure 2.11:** Combined constraints of the cosmological parameters  $\Omega_m$  and  $\Omega_\Lambda$  from type Ia Supernovae, CMB and BAO from cosmological observations prior to 2012. Despite the outdated results displayed in the image, the same agreement on  $\Omega_m$ ,  $\Omega_\Lambda$  and curvature can be found in the most recent data. The constraints from CMB and BAO results have improved significantly due to the Planck measurements [A<sup>+</sup>18a]. The region marked with “No Big Bang” means that for that set of cosmological parameters the universe starts in a contracting state, resulting in a bouncing universe [CPT92]. Figure from [S<sup>+</sup>12].

structures that would be easily observed. Nevertheless it is important to stress that, despite MACHOs not being the most favourable candidate, there is a small contribution to the  $\Omega_m$  from these compact objects<sup>7</sup>.

### 2.3.2 Neutrinos

Aside from photons, neutrinos are the most abundant particles in the universe. They are weakly-interacting neutral leptons with non-zero mass, therefore it is reasonable to consider them as dark matter candidates that, at least in part, may explain the dark matter density observed. The mass scale of neutrinos is much smaller than that of the other fundamental particles, with constraints on the sum of the neutrino masses placed by cosmological observables such as the CMB at  $\sum m_\nu < 0.12$  eV [A<sup>+</sup>18a] or by nuclear decay experiments such as KATRIN at  $m_\nu < 1.1$  eV [A<sup>+</sup>19d]. Therefore, neutrinos can only contribute to the dark matter density as a form of hot dark matter, *i.e.*, as relativistic weakly-interacting particles.

<sup>7</sup>Their contribution should be accounted for in the baryonic matter density  $\Omega_b$  from CMB data.

Arguments for large scale structure formation disfavour neutrinos, or other forms of hot dark matter, as a complete explanation for dark matter [Sch15]. Simulation models indicate that a dense and abundant fluid of non-relativistic (cold) dark matter is required in order to form structures with distance scales compatible to those observed in the universe. Relativistic particles are not as tightly bound to gravitational potentials and would produce large scale structures that are more diffuse than those observed. Furthermore, galactic halos cannot be fully composed of neutrinos as well, as these particles would disperse due to their high velocities, easily surpassing the escape velocity of the galaxy. The constraints to the total mass of all three neutrino flavours from CMB data result in a neutrino contribution to the hot dark matter density of  $\Omega_\nu h^2 \leq 0.0062$  [T<sup>+</sup>18].

### 2.3.3 Axions and Axion-like Particles

The axion,  $A^0$ , is an hypothetical pseudoscalar particle that was first introduced by Peccei and Quinn in 1977 to justify why CP-violating processes are not observed in strong interactions when, according to the SM, there is no apparent reason for them not to [PQ77a, PQ77b]. Any CP-violating phase in quantum chromodynamics (QCD) would induce a magnetic dipole moment in baryons, but recent measurements of the magnetic dipole moment of the neutron reveal that it is at most incredibly small, of order  $|d_n| < 2.9 \times 10^{-26}$  e cm [B<sup>+</sup>06]. This implies that any CP-violating term in QCD is “fined-tuned” to vanish. The hidden global symmetry introduced by Peccei and Quinn – PQ symmetry – is considered to be spontaneously broken, which results in the creation of a pseudo-Goldstone boson, the axion. This mechanism effectively compensates the CP-violating phase of QCD, making it vanishingly small. The axion is doubly motivated as a dark matter candidate because it would not only solve the dark matter problem but would also solve the CP violation problem in strong interactions. Primordial axions would not be produced thermally and therefore are not expected to be relativistic [B<sup>+</sup>12b]. Therefore, if axions exist they would mostly contribute to the matter density as cold dark matter. However, other sources of cosmic axions can produce relativistic and semi-relativistic particles that would contribute to the hot and warm dark matter energy density, respectively [Mar18].

An axion-like particle (ALP) is a generalization of the concept of the axion but with no relationship between mass and fermion coupling strength [T<sup>+</sup>18]. Both the couplings of axions and ALPs to SM particles are suppressed by the energy scale of the symmetry breaking. This energy scale can be large enough so that axions and ALPs are weakly interacting.

In the presence of a strong magnetic field, axions could be effectively converted into photons by the Primakoff effect [T<sup>+</sup>18]. The *Axion Dark Matter eXperiment* (ADMX) collaboration uses a microwave resonance cavity that explores this effect in order to detect resonant radio frequencies corresponding to the Compton wavelength of the axion. The latest results from ADMX exclude axions with masses between 2.66–3.31  $\mu\text{eV}$  [B<sup>+</sup>20]. The inverse Primakoff effect, where two photons are converted into an axion is expected to occur inside the Sun, resulting in the irradiation of axions that could be detected on Earth [FRS<sup>+</sup>14]. The *CERN Axion Solar Telescope* (CAST) helioscope places strong constraints on the axion-photon coupling at

$G_{A\gamma\gamma} < 6.6 \times 10^{-11} \text{ GeV}^{-1}$  for an axion mass lower than 0.02 eV [A<sup>+</sup>17j]. The LUX experiment also provided a limit to the axion-electron coupling of  $g_{Aee} < 3.5 \times 10^{-12}$  for axions and  $g_{Aee} < 4.2 \times 10^{-13}$  for galactic ALPs by exploring the axio-electric effect in liquid xenon [A<sup>+</sup>17e].

### 2.3.4 WIMPs

Weakly Interacting Massive Particles, or WIMPs, are a generic class of chargeless, massive and stable particles that do not interact strongly with regular matter. As previously stated, cosmological evidences suggest that dark matter particles are of non-baryonic nature, do not interact via the electromagnetic and strong forces, are abundant enough to account for the missing matter density and are stable in the time scale of the age of the universe [JKG96]. Any model that predicts a particle with the characteristics of a WIMP particle provides a valid candidate to explain dark matter [T<sup>+</sup>18].

Soon after the Big Bang, thermal annihilation of energetic particle-antiparticle pairs in the primordial plasma would result in the production of large amounts of WIMP particles. The total WIMP density would be in equilibrium as pairs of WIMP particles annihilate with the same rate as they are produced. However, the expansion of the universe resulted in the cooling of the primordial plasma and particles would no longer have enough energy to produce WIMP particles, leading to a decrease of the WIMP density via annihilation until the rate of expansion of the universe overcomes the annihilation rate, “freezing out” the population of WIMP particles at that point. This results in a thermal relic density of WIMPs that permeate the universe and would still be present today, assuming that the WIMP particle is stable. This WIMP relic density at the current time depends on the strength of the annihilation cross section, as a higher annihilation would result in fewer WIMPs today, and can be estimated by:

$$\Omega_\chi h^2 \approx \frac{3 \times 10^{-27} \text{ cm}^2 \text{ s}^{-1}}{\langle \sigma_{Av} \rangle} \approx \frac{0.1 \text{ pb c}}{\langle \sigma_{Av} \rangle}, \quad (2.20)$$

where  $\langle \sigma_{Av} \rangle$  is the thermally averaged WIMP annihilation cross section. This result is independent of  $m_\chi$  and the WIMP density is inversely proportional to the annihilation cross section. Interestingly, if the annihilation cross section is of the order of the weak interactions,  $\langle \sigma_{Av} \rangle \sim 10^{-25} \text{ cm}^2 \text{ s}^{-1}$ , the resulting relic abundance for WIMPs is approximately  $\Omega_\chi h^2 \sim 0.1$ , remarkably similar to the non-baryonic matter density contribution observed today,  $\Omega_{DM,0} h^2 \approx 0.1188$ . There is no *a priori* reason to consider a weak-scale interaction, making this result quite remarkable [DAP15].

Several models beyond the Standard Model (SM) of particle physics predict new particles with the properties of WIMPs. Supersymmetric (SUSY) extensions of the SM introduce a new global symmetry that relates fermions and bosons [JKG96, T<sup>+</sup>18], and imply that every SM particle has an associated superpartner – fermions have bosonic superpartners and bosons have fermionic superpartners. If supersymmetry was unbroken, each superpartner would have the same mass as the respective SM particle. No superpartners were yet observed at particle colliders, meaning that if they do exist, they must be very massive, and SUSY is spontaneously broken [JKG96,



T<sup>+</sup>18]. Minimal Supersymmetric extensions of the Standard Model (MSSM) introduce R-parity to explain the stability of the proton. Conservation of this R-parity means that the lightest stable particle (LSP) of the model cannot decay [JKG96]. Most MSSM models indicate the *neutralino* (a linear *eigenstate* combination of the *photino*, *zino* and *higgzino* – the supersymmetric partners of the photon,  $Z^0$  boson and the Higgs boson, respectively) as the LSP and as a possible WIMP candidate, with masses ranging from 10–1000 GeV and with weak-scale interaction cross section with matter [JKG96]. So far no evidence for any supersymmetric extension of the SM has been found at the energy scales up to the TeV range [T<sup>+</sup>18].

### 2.3.5 Modified Gravity

For completeness, a brief overview of modified theories of gravity is presented. These theories do not contemplate dark matter and in fact discard its contributions to cosmological observables completely. As most of the observational evidences for missing mass in cosmological systems involve a particular formalisation of gravity – in order to compute the motion of stellar objects – modification of the way gravity behaves at large scales can reproduce the observations without including unseen mass in the system.

One can look at this debate of adding unknowns or modifying the underlying theories in an historical way: when faced with the troubling fact that the orbit of Uranus presented some discrepancies relative to the laws of Kepler and Newton, Urbain Le Verrier decided to propose a new unseen planet beyond Uranus that could be the cause of such irregularities. His precise calculations of the orbit of the new planet lead to the discovery of Neptune by Johann Gottfried Galle in the same night that he received the coordinates by letter from Le Verrier. This was a remarkable feat for celestial mechanics. When Le Verrier was faced with the orbital irregularity of Mercury, the precession of the perihelion, and in trying to repeat his feat with Uranus, he proposed that a new planet, Vulcan, with an orbit smaller than the one of Mercury, was disturbing the motion of the celestial body. Later on, the precession problem was solved not by Le Verrier but by Albert Einstein, that successfully modified Newtonian gravity and introduced General Relativity as a new theory of gravity. This historical example shows us that both approaches should always be considered until one of them is validated or ultimately discarded.

Modified Newtonian Dynamics (MOND) [Sca06] is one of many theories of modified gravity. This theory states that the effects of additional mass that are observed are in fact the result of the misconception that Newtonian gravity works the same way locally and at large distances. MOND was developed to explain the galactic rotation curves avoiding the need for new exotic, non interacting particles. Despite being very efficient at parametrizing several galactic rotation curves, MOND fails at predicting the total mass of clusters. This is due to MOND not being general enough to cover medium-scale and large-scale effects simultaneously, and usually a substantial amount of unseen matter is still necessary to justify the observed dynamics. Gravitational lensing evidences also disfavour MOND, because these effects have large discrepancies with the theory. Finally, the Bullet Cluster collision cannot be fully justified using modified gravity alone [AFZ06] and any analysis of the CMB power spectrum in the framework of modi-

fied gravity seems to require some form of non-baryonic matter component to be present to fully justify observations [BD12].

## 2.4 Search and Detection of Dark Matter

As discussed in Section 2.2, there are several cosmological evidences that agree on the existence of a non-baryonic matter component that dominates the matter content of the universe. However, evidences from cosmological observations alone are not enough to probe the properties and nature of dark matter, that remains unknown to this present date. A detection experiment that can undoubtedly measure dark matter interactions either in annihilation or scattering processes could shed some light on the properties of dark matter particles and test cosmological and particle physics models.

Dark matter detection can be grouped into three distinct methods: indirect detection, direct detection and production at particle colliders. An indirect detection method would be to observe and measure dark matter annihilation or decay products from cosmological sources rich with dark matter. Direct detection refers to the measurement of a signal arising from a WIMP interaction with a material target on Earth.


### 2.4.1 Production at Particle Colliders

The production of dark matter particles in particle colliders would produce a specific signature in the form of a large missing energy and momentum. This is the result of the low interaction cross section of a dark matter particle that, upon creation, would leave the detector unimpaired [S<sup>+</sup>10, Sch15]. According to some models beyond the Standard Model, it should be possible to produce dark matter particle-antiparticle pairs at a particle collider, providing that the energy available in the colliding particle beam is larger than twice the mass of the hypothetical dark matter particle,  $E_{th} > 2m_\chi$ . This method provides more information about the properties of dark matter particles, if observed, than any other method. Not only would it provide a direct measurement of the self-annihilation or interaction cross-sections with the particles involved but it could also reveal the nature of the particles and in that way test different standard model extensions that predict dark matter candidates, like supersymmetry (SUSY) models [JKG96].

If a supersymmetric extension of the SM exists, events with gluon jets, leptons and a large amount of missing transverse energy should be observed at electroweak-scale energies. So far no solid evidence for SUSY particle creation was found. Some particle physics experiments like the Tetravon at Fermilab and ATLAS and CMS at CERN have set constraints on dark matter mass and production cross section, at the same time constraining some SUSY models [BFH10, S<sup>+</sup>19, S<sup>+</sup>20, A<sup>+</sup>19a, A<sup>+</sup>19b].

Figure 2.12 shows the current best limits on spin-independent WIMP-nucleon scattering cross section (left) and spin-dependent WIMP-proton scattering cross section (right) from the ATLAS experiment [A<sup>+</sup>19b]. The current best constraints from CMS and ATLAS on the spin-dependent

WIMP-proton cross section are the most stringent amongst direct detection and production experiments by several orders of magnitude over the entire 1 – 1000 GeV mass range. The same is true for the inferred WIMP-nucleon spin-independent cross section for WIMP masses below 5 GeV [A<sup>+</sup>19b].



**Figure 2.12:** Limits on the spin-independent WIMP-nucleon (left) and spin-dependent WIMP-proton (right) scattering cross sections from the ATLAS experiment [A<sup>+</sup>19b].

### 2.4.2 Indirect Detection of Annihilation or Decay Products

Probing dark matter through indirect detection methods would provide valuable information about the dark matter particles, such as their mass range, annihilation cross section or half-life. Annihilation of such particles would result in the creation of particle-antiparticle pairs with high energy that could reach the Earth and be detected. These particles could be high-energy  $\nu\bar{\nu}$  pairs,  $e^-e^+$  pairs, high energy  $\gamma$ -rays or even proton-antiproton pairs. The proton-antiproton asymmetry of cosmic rays – the lack of antiprotons when compared to the proton cosmic ray flux on earth – is a good example of an indirect dark matter probe as provides good constraints in the abundance and annihilation cross section of heavy dark matter particles [FMV14].

An example of a dark matter indirect detection effort is the probing of Dwarf Spheroidal Galaxies (DSG)<sup>8</sup> performed by the Fermi Gamma-Ray Space Telescope (FGST, former GLAST) with the objective of detecting high energy  $\gamma$ -ray signatures from annihilation of dark matter particles [CHL15].

Another example of an indirect method is the observation of the decay of axions emitted by the Sun into  $\gamma$ -rays in the magnetic field of the Earth. In September 2014, the XMM-Newton space observatory claimed the detection of an axion decay signature in the magnetosphere of the Earth [FRS<sup>+</sup>14], but the findings could be explained by other effects unrelated to solar axion production, as stated by some authors [RT15]. Despite the advantages of an indirect detection signal, the information obtained is very model-dependent and alternative explanations can often be found, as it was the case of the XMM-Newton 2014 findings.

<sup>8</sup>Galaxies with a relative large amount of dark matter and a small number of  $\gamma$ -ray sources.

### 2.4.3 Direct Detection of Dark Matter Scattering

Direct detection of dark matter is defined as the measurement of dark matter interaction events in an Earth-bound detector, and can provide constraints on the interaction cross section of these particles with normal matter and their mass. Of all the three methods mentioned here, this is the one which provides less information about the nature of the particles being detected. Nevertheless, it could present the strongest evidence for the existence of a non-baryonic component of matter in the universe as well as a value for the interaction cross section of such matter with baryonic matter. The rest of this chapter will focus on WIMP direct detection only.

WIMP direct detection experiments aim to measure the rate and energies of nuclear recoils caused by the scattering of dark matter particles with the nuclei of a target material in a detector [Sch11]. If such interactions occur, the deposited energy can be measured if the energy transfer occurs above the threshold limit of the detector. Furthermore, due to the weakly interacting nature of dark matter, a very low rate of events is expected. Therefore, a dark matter direct detection experiment needs to have both a low energy threshold and high efficiency at suppressing background events in order to allow the rare WIMP scattering signals to be observed.

The excitation of the target medium caused by a nuclear recoil from a WIMP scatter can be measured by three main energy channels: ionization charge, phonons<sup>9</sup> and scintillation. Furthermore, a detector that uses more than one energy channel can combine the different measured signals to provide extra discrimination methods, as the nuclear recoils from WIMPs and interactions from electromagnetic background sources often produce distinct signatures [Sch11]. By measuring one or more energy channels, detectors can be divided into four main types:

1. Threshold detectors, such as the COUPP experiment [B<sup>+</sup>12a], use superheated fluid, tuning some thermodynamic parameters to measure nucleation events from ionization of the target material, similar to bubble chambers.
2. Crystal scintillator detectors measure light signals arising from dark matter particles interacting within the target material. Usual crystal materials are *NaI(Tl)* and *CsI*. DAMA and DAMA/Libra [B<sup>+</sup>03, B<sup>+</sup>13] experiments use ultra-pure *NaI* crystals. These detectors cannot discriminate between nuclear recoils and electron recoils.
3. Two-phase liquid noble element detectors such as LUX, XENON1T, DarkSide, PandaX and LZ (described in more detail in Chapter 4), measure scintillation and ionization<sup>10</sup> from interactions with the liquid target volume of a dual-phase TPC [A<sup>+</sup>17h, A<sup>+</sup>18b, A<sup>+</sup>18i, W<sup>+</sup>20, A<sup>+</sup>20d].
4. Semiconductor detectors such as EDELWEISS, CRESST and SuperCDMS, measure heat and ionization or scintillation in cryogenic crystals of ultra-pure materials [A<sup>+</sup>20n, A<sup>+</sup>19c,

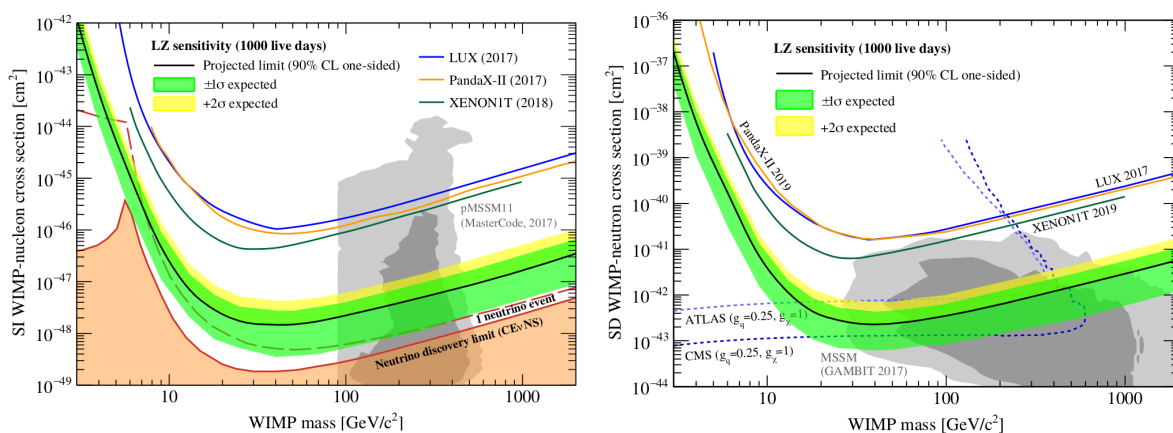
---

<sup>9</sup>Vibrations of the target lattice, resulting in increased heat that can be detected in ultra-cooled semiconductor detectors.

<sup>10</sup>Single-phase set-ups measure scintillation only, while dual-phase (liquid and gas) detectors measure scintillation and ionization by drifting electrons with an electric field into the gas phase to be measured.

$A^{+13b}$ ,  $A^{+17b}$ ]. The most common materials are germanium, silicon or  $CaWO_4$  crystals operated at  $mK$  temperatures. Two particular advantages of these detectors are their high energy resolution and low threshold, *e.g.*, the CRESST detector can reach an energy threshold as low as 30.1 eV [ $A^{+19c}$ ].

If the dark matter halo is composed of WIMPs, then the WIMP flux passing through the Earth would result in some measurable scattering events with a target of large enough mass. The expected signal and event rate for a dark matter direct detection experiment can be obtained considering models for the dark matter halo, such as the canonical model described below.



**Figure 2.13:** Projected sensitivity on the spin-independent WIMP-nucleon (left) and spin-dependent WIMP-neutron (right) scattering cross sections of LZ [ $A^{+20d}$ ]. Solid lines represent the most recent limits from direct detection experiments [ $A^{+17h}$ ,  $W^{+20}$ ,  $A^{+18i}$ ,  $A^{+17g}$ ,  $A^{+19f}$ ,  $A^{+19h}$ ,  $X^{+19}$ ].

## 2.5 The Canonical Halo model

The “Canonical” halo model for the dark matter distribution encompassing our Milky Way galaxy is approximated to an isothermal spherical distribution, considered to have no bulk rotation, of non-interacting particles that behave as an ideal gas, moving with no preferable direction within the halo [LS96]. The velocity profile of these particles follows a Maxwell-Boltzmann distribution.

Furthermore, the Earth moves across the dark matter halo due to its motion around the Sun and the motion of the Sun around the galaxy. The induced annual modulation due to the rotation of the Earth around the Sun results in a variation of the expected event rate and nuclear recoil spectrum that can be measured by a dark matter experiment [ $B^{+13}$ , LS96].

The WIMP velocity,  $v_0$ , can be approximated to the circular velocity of an object orbiting the Milky Way, as  $\langle v_0 \rangle \approx \langle v_r \rangle$  [LS96]. At the location of the Sun,  $R_\odot \approx 8.5$  kpc, this velocity is taken to be  $v_r = 220$  km s $^{-1}$  [ $P^{+14}$ ,  $B^{+12c}$ ] but some authors suggest a higher number, Lewin and Smith [LS96] use  $\sim 230$  km s $^{-1}$ . On the following calculations, the value considered is  $v_r = 220 \pm 20$  km s $^{-1}$  [Sch11].

The kinetic energy cut-off of the dark matter particles in the halo is set by the galactic escape velocity. This velocity at the radial distance of the Sun is measured to be from  $492 < v_{\text{esc}} < 587 \text{ km s}^{-1}$  in [P<sup>+</sup>14] and  $v_{\text{esc}} = 550.9^{+32.4}_{-22.1} \text{ km s}^{-1}$  in [KSLB14].

The density of the Milky Way dark matter halo at the position of the Solar system is estimated by Gates, Gyuk and Turner [GGT95] between  $0.3 < \rho_{\text{DM}} < 0.7 \text{ GeV cm}^{-3}$ . Recent surveys point to a value of  $\rho_{\text{DM}} = 0.35 \pm 0.08 \text{ GeV cm}^{-3}$  [KSLB14]. The value we will use here is  $\rho_{\text{DM}} = 0.4 \text{ GeV cm}^{-3}$ , as considered in Lewin and Smith [LS96].

As an illustration of this particle density, if WIMP particles have a mass of  $M_\chi = 100 \text{ GeV}$ , a litre bottle would enclose on average 4 WIMP particles, and the average WIMP flux would be of  $\Phi_\chi \approx 10^7/M_\chi \text{ s}^{-1} \text{ cm}^{-2}$ , or around 10 million particles passing through the area of a fingernail every second. This high flux is the reason why direct detection is thought to be feasible despite the weakly interacting nature of dark matter.

It is unlikely that most of the dark matter halos that surround galaxies are exactly spherical, they are most likely irregular ellipsoids [WHC<sup>+</sup>15, KSLB14]. Also, galactic dark matter structures are expected to have some rotation due to conservation of angular momentum as the dark matter distribution collapses to form the halo. This rotation would result in some flattening of the halo with respect to the rotational axis [KK98]. Furthermore, the halo could have an alternative radial profile, rather than the  $\rho_\chi \propto r^{-2}$  density distribution arising from an isothermal sphere model. Essentially, all the empirical information we have on the halo is provided by the galactic rotation curves [KK98]. Any halo parametrization that is consistent with the observables is a valid one. The canonical model has become the “standard” halo model used by dark matter direct detection experiments in order to fairly compare their results.

## 2.6 WIMP Interaction Cross Section

The calculations presented in this chapter follow the work of Lewin and Smith [LS96]. The WIMP-nucleus interaction cross section can be divided into a zero-momentum transfer cross section term,  $\sigma_{0WN}$ , that contains all the dependencies on the physics model of the interaction mechanism, and a form factor term  $F^2(q)$  containing the dependence on the momentum transfer  $q$  [Sch11]. This separation of the energy dependence is represented in Equation 2.21.

$$\frac{d\sigma_{WN}(q)}{dq^2} = \frac{\sigma_{0WN}}{4\mu_A^2 v^2} F^2(q). \quad (2.21)$$

Here  $v$  is the WIMP velocity,  $\mu_A$  is the WIMP-nucleus reduced mass,  $\mu_A = M_\chi M_A / (M_\chi + M_A)$ , for the mass of a WIMP particle  $M_\chi$  and the mass of the nucleus  $M_A$  with a mass number  $A$ . The zero momentum transfer WIMP-nucleus cross section can be separated into a spin-dependent (SD) and a spin-independent (SI) component [KK04, Sch11]. For SI interactions, the effective couplings to the protons and the neutrons are expected to be identical, meaning that in a SI interaction the WIMP particle will effectively interact with all nucleons.

As dark matter experiments have a broad range of target materials, it is useful to express the WIMP interaction cross section in a target-independent way, considering the WIMP-nucleon cross section, as shown in Equation 2.22 [LS96]. This allows for the comparison of results from experiments using different targets.

$$\sigma_{0WN,SI} = \frac{4\mu_n^2 f_n^2 \mu_A^2}{\pi \mu_n^2} A^2 = \sigma_{SI} \frac{\mu_A^2}{\mu_n^2} A^2. \quad (2.22)$$

The term  $\mu_n$  is the WIMP-nucleon reduced mass,  $f_n$  is the effective WIMP-nucleon coupling and  $\sigma_{SI}$  is the SI WIMP-nucleon cross section. The dependence on the target material is called the coherence factor, expressed in the term  $A^2 \mu_A^2/\mu_n^2$ . The (target independent) SI WIMP-nucleon interaction cross section is expressed by

$$\sigma_{SI} = \frac{4\mu_n^2 f_n^2}{\pi}. \quad (2.23)$$

### 2.6.1 Nuclear Form Factor Correction

The sensitivity of the WIMP particle to the structure of the target nucleus is dependent on the momentum of the WIMP particle. This dependence results in a modification of the WIMP-nucleus cross section in order to account the internal nuclear structure. This modification is parametrized by a form factor correction, which is unity for  $q \rightarrow 0$ . For a nucleus with finite size, the form factor is approximately the Fourier transform of the spacial density distribution of nucleons (scattering centres) in the nucleus,  $\rho_n(r)$  and, considering a spherically symmetric distribution of nucleons inside the nucleus, can be simplified to be:

$$F(q) = \frac{4\pi}{q} \int_0^\infty r \cdot \sin(qr) \rho_n(r) dr. \quad (2.24)$$

For the SI and SD terms of the total cross section, it is useful to consider the form factor obtained by Fourier transform of, respectively:

- (a) a solid sphere that represents the whole nucleus.
- (b) a thin shell that represents the unpaired nucleon in the nucleus.

For the thin shell approximation, the Fourier transform returns the first spherical Bessel function,  $j_0$

$$F(q) \text{ (thin shell)} = j_0(qr_N) = \frac{\sin(qr_N)}{qr_N}, \quad (2.25)$$

where  $r_N$  is the effective nuclear radius. For the solid sphere approximation, the resulting form factor is:

$$F(q) \text{ (solid sphere)} = \frac{3j_1(qr_N)}{qr_N} = \frac{3}{(qr_N)^3} [\sin(qr_N) - qr_N \cos(qr_N)], \quad (2.26)$$

where  $j_1$  is the second spherical Bessel function. These approximations return satisfactory form factor corrections, but some different nuclear density distributions return better results, such as a Fermi distribution, considered to be the most realistic, or the truncated Gaussian distribution proposed by Helm [Hel56]. A detailed discussion of these form factors can be found in [LS96], but goes beyond the scope of this work.

## 2.6.2 Event Rate

The WIMP-nucleus event rate per unit mass is given by:

$$R = \frac{N_A}{A} \sigma_0 \int v dn \equiv \frac{N_A}{A} \sigma_0 n_0 \langle v \rangle \quad (2.27)$$

Here  $\sigma_0$  is the zero-momentum transfer cross section per nucleus,  $A$  is the atomic mass of the target,  $n$  is the WIMP number density and  $N_A = 6.022 \times 10^{23} \text{ mol}^{-1}$  is the Avogadro number. The differential event rate is then given by:

$$dR = R_0 \frac{1}{2\pi v_0^4} v f(\vec{v} + \vec{v}_E) d^3 v. \quad (2.28)$$

The term  $R_0$  is the time-independent rate per unit mass for zero Galactic velocity  $v_E = 0$  and  $v_{\text{esc}} = \infty$ , which is defined as:

$$R_0 = \frac{2}{\sqrt{\pi}} \frac{N_A}{A} \frac{\rho_{\text{DM}}}{M_\chi} \sigma_0 v_0 \quad (2.29)$$

The total rate  $R_0$  is expressed in  $\text{kg}^{-1} \text{ d}^{-1}$  or “*tru*” for total rate units. Considering a WIMP mass of  $M_\chi = 100 \text{ GeV}$ , a target made of xenon ( $A = 131$ ) and a WIMP-nucleon interaction cross section  $\sigma_{0Wn,SI} = 0.9 \times 10^{-46} \text{ cm}^2$ , using Equation 2.29 the total rate obtained is  $R_0 \approx 0.214 \text{ t}^{-1} \text{ d}^{-1}$ , or around 6 events per tonne of target material per month.

The minimum particle energy that can generate a recoil of energy  $E_R$  is  $E_{\text{min}} = E_R/r$ , with  $r$  a dimensionless parameter that is 1 when  $M_A = M_\chi$ . The most likely WIMP energy  $E_0$  is the energy of a WIMP particle with velocity  $v_0$ , *i.e.*  $E_0 = 1/2 M_\chi v_0^2 = (v_0^2/v^2)E$ . The minimum velocity  $v_{\text{min}} = \sqrt{2E_{\text{min}}/M_\chi} = \sqrt{E_R/E_0} r v_0$ .

The nuclear recoil spectrum is then given by:

$$\frac{dR}{dE_R} = \frac{R_0}{E_0 r} \frac{1}{2\pi v_0^2} \int_{v_{\text{min}}}^{v_{\text{max}}} \frac{1}{v} f(\vec{v} + \vec{v}_E) d^3 v. \quad (2.30)$$

$$\frac{dR(v_E = 0, v_{\text{esc}} = \infty)}{dE_R} = \frac{R_0}{E_0 r} e^{-\frac{E_R}{E_0 r}}. \quad (2.31)$$



For a WIMP mass of 100 GeV with expected income velocity  $v_0 = 220 \text{ km s}^{-1}$  scattering with a xenon nucleus ( $A = 131$ ), an average recoil energy  $\langle E_R \rangle = E_0 r \approx 25 \text{ keV}$  is obtained.

Due to the expected low event rates and deposited energies for dark matter interactions, a dark matter detector must have low energy threshold and excellent background rejection in order to be sensitive to these events.



# Chapter 3

## Neutrinos

Neutrino physics has grown remarkably since the end of the last century, and has been one of the most active fields of study in physics for some decades. The characterization of neutrino-matter interactions is of extreme importance for the understanding of neutrino oscillations, supernova dynamics, and even dark matter searches [dG<sup>+</sup>13]. A couple of important questions regarding the neutrino are still unanswered, and some solutions hint to physics beyond the Standard Model. Perhaps the most important questions regard their fundamental nature (Dirac or Majorana fermions), the origin and absolute values of their masses and their role in the matter-antimatter asymmetry observed in the universe today. The best probes for the Majorana nature of the neutrino are searches for lepton-number violating processes such as the neutrinoless double beta decay, whose observation would greatly impact a variety of fields, from particle physics to cosmology.

Section 3.1 will briefly present the history of the neutrino, from its discovery to the establishment of the electroweak theory, the solar neutrino problem and atmospheric neutrino anomaly, and how these observations led to the first evidences of neutrino flavour oscillations. Section 3.2 will discuss the *neutrino-mass-induced flavour oscillations* in detail, as well as the neutrino hierarchy problem and the current best limits for the neutrino mass and oscillation parameters. The origin of the neutrino masses is discussed in Section 3.3. This chapter closes with Section 3.4 that will provide an overview of the double beta decay process and its neutrinoless mode, the main topic of this work, and the implication of its observation.

### 3.1 Neutrino Hypothesis and Discovery

The first nuclear models of  $\beta$ -decay predicted the emission of a monoenergetic electron, but instead a continuous energy spectrum was observed. This result seemed to imply that energy and momentum were not being conserved in this three-body process. Niels Bohr was prepared to abandon energy conservation for these subatomic processes [HH00], but in 1930 Wolfgang Pauli proposed the existence of a very light, spin  $1/2$ , neutral particle that was partaking in the decay and carried away part of the available energy without being detected [Pau]. Four years later, Enrico Fermi revised the theory of  $\beta$ -decay by incorporating the unobserved particle proposed by Pauli, that he named *neutrino* (loosely from Italian for “tiny neutral one”) in order to distinguish it from the heavier neutron discovered in the meantime by James Chadwick [Cha32]. In Fermi’s model, a neutron ( $n$ ) spontaneously converts into a proton ( $p$ ) with the emission of an electron ( $e^-$ ) and a neutrino<sup>1</sup> ( $\nu$ ) [Fer34],

$$n \longrightarrow p + \nu + e^- . \quad (3.1)$$

This 4-fermion theory of weak interaction proved successful in describing the  $\beta$ -decay and led to the wide acceptance of the existence of the neutrino. In 1936, using Fermi’s theory, Hans Bethe postulated that the inverse  $\beta$ -decay process could occur:

$$p + \nu \longrightarrow n + e^+ , \quad (3.2)$$


which provided a channel for the direct detection of the neutrino. However, the estimated cross section for this process, of the order of  $10^{-44} \text{ cm}^2$  (10 zb) for a 1 MeV neutrino [GK07], seemed too low to make any experimental observation feasible.

Almost two decades later, Clyde Cowan and Frederick Reines, urged by Bruno Pontecorvo, devised an experiment to look for the neutrino by detecting the inverse  $\beta$ -decay process [Pon91, RC97]. To overcome the low cross section of this process, Cowan and Reines resorted to the large flux of antineutrinos ( $1.2 \times 10^{13} \text{ cm}^{-2} \text{ s}^{-1}$ ) produced by a nuclear reactor [Rei96]. Their detector consisted of 1400 litres of cadmium-doped liquid scintillators layered in between water targets and observed by light sensors. The water tanks provided a proton-rich target where the inverse  $\beta$ -decay was expected to occur. Cadmium has a large (low-energy) neutron capture cross section, so doping the liquid scintillator with cadmium effectively increases the chance of capturing the neutrons produced in the reaction [RC97]. Figure 3.1 illustrates the expected signature of such an event. The two products of the reaction, a positron and a neutron, produce characteristic signals that can be used to identify the event: the positron promptly annihilates within the target and produces two back-to-back 511 keV  $\gamma$ -rays, while the neutron is later captured by a  $^{108}\text{Cd}$  atom, releasing 9 MeV of energy in  $\gamma$ -rays. And finally, in 1956 Cowan and Reines provided the first direct experimental evidence for the existence of the neutrino<sup>2</sup>

---

<sup>1</sup>The  $\beta^-$ -decay involves the emission of an antineutrino ( $\bar{\nu}$ ), but at this point the distinction between neutrino and antineutrino was not established.

<sup>2</sup>Reines received the Nobel Prize in Physics 1995 “for the detection of the neutrino”. Unfortunately, Cowan passed away in 1974.



**Figure 3.1:** Detection of inverse  $\beta$ -decay on hydrogen (proton) in the reactor-neutrino experiment by Cowan and Reines [CRH<sup>+</sup>56]. The signature of this interaction is the delayed ( $\sim$ ms) observation of  $\gamma$ -rays from neutron capture that follows the detection of characteristic  $\gamma$ -rays from  $e^+$  annihilation. Figure from [RC97].

[CRH<sup>+</sup>56].

The detection of the muon neutrino in 1962 by Schwartz, Lederman and Steinberger provided evidence that not all neutrinos are equal, and as with the charged leptons the neutrinos presented different flavour states<sup>3</sup> [DGG<sup>+</sup>62]. In 1973 the Gargamelle Collaboration at CERN found evidence for neutral currents involved in neutrino interactions, which resulted in the confirmation of the unification of electromagnetism and the weak interactions [H<sup>+</sup>74]. Ten years later the UA1 and UA2 experiments at CERN discovered the  $W^\pm$  and  $Z$  bosons that mediated weak processes and the Electroweak theory was established [Jak94]. The discovery of the tau neutrino in the year 2000 by the DONUT collaboration [K<sup>+</sup>01] has completed the lepton sector as it is perceived today, with three generations of leptons represented on Table 3.1 arranged in weak isospin doublets. Notice that the weak doublets are all represented by left-handed fields (hence the subscript “ $L$ ”).

**Table 3.1:** The three generations of the left-handed lepton sector of the Standard Model (SM) and their the electromagnetic charge ( $Q$ ). The left-handed leptons can be arranged in  $SU(2)_L$  doublets of the weak sector, while right-handed particles are  $SU(2)$  singlets and therefore do not undergo weak interactions.

1 <sup>st</sup>	2 <sup>nd</sup>	3 <sup>rd</sup>	$Q$
$\begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L$	$\begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix}_L$	$\begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix}_L$	0 -1

<sup>3</sup>Schwartz, Lederman and Steinberger received the Nobel Prize in Physics 1988 “for the neutrino beam method and the demonstration of the doublet structure of the leptons through the discovery of the muon neutrino”.

## 3.2 Flavour Oscillations and Neutrino Mass

It has been established, beyond reasonable doubt, that neutrinos can oscillate between flavour states as they propagate [Z<sup>+</sup>20] and the explanation of flavour oscillations is deeply connected to neutrino masses.

### 3.2.1 Discovery of Neutrino Oscillations

#### 3.2.1.1 The Solar Neutrino Problem


The Sun is the largest continuous source of neutrinos with energies above the 1 eV range. A series of fusion reactions at the core of the Sun fuse hydrogen into helium throughout the *pp* chain and the CNO cycle [BSB05]. These reactions are depicted in Table 3.2 together with the typical energy of the emitted neutrinos. The solar fusion reactions release copious amounts of electron neutrinos, roughly 2% of the total solar energy output [Ian14].

**Table 3.2:** Sources of neutrinos from the solar fusion cycle according to the Standard Solar Model [BSB05, GK07].  $\langle E_\nu \rangle$  is the average neutrino energy and  $E_\nu^{max}$  is the endpoint energy of the reaction. Notice that neutrinos with discrete energies are the result of two body interactions, while neutrinos with a continuous distribution of energies (for which only  $E_\nu^{max}$  is given) are the result of three body interactions.

Reaction	$\langle E_\nu \rangle$ [MeV]	$E_\nu^{max}$ [MeV]
Proton-proton (pp) Chain:		
$p + p \rightarrow d + e^+ + \nu_e$	0.2668	$0.423 \pm 0.03$
$p + e^- + p \rightarrow d + \nu_e$	-	1.445
${}^3\text{He} + p \rightarrow {}^4\text{He} + e^+ + \nu_e$	9.628	18.778
${}^7\text{Be} + e^- \rightarrow {}^7\text{Li} + \nu_e$	-	0.8631
${}^8\text{B} \rightarrow {}^8\text{Be}^* + e^+ + \nu_e$	$6.735 \pm 0.036$	$\sim 15$
Carbon-Nitrogen Cycle:		
${}^{13}\text{N} \rightarrow {}^{13}\text{C} + e^+ + \nu_e$	0.7063	$1.1982 \pm 0.0003$
${}^{15}\text{O} \rightarrow {}^{15}\text{N} + e^+ + \nu_e$	0.9964	$1.7317 \pm 0.0005$

As can be seen in Figure 3.2, the *pp* neutrino flux is the main contribution to the total flux of solar neutrinos, but these neutrinos have relatively low energies (under 0.42 MeV) and are well below the energy threshold of most neutrino detectors. The *hep* and  ${}^8\text{B}$  neutrinos have a substantially smaller flux but the higher energies allow their detection even in water Cherenkov detectors with high energy thresholds. These neutrinos are a vital tool to understand the dynamics of the Sun and the nuclear reactions that are involved, allowing some solar models to be tested. As John N. Bahcall noted, neutrinos can be used in a practical experiment “...to see into the interior of a star and thus verify directly the hypothesis of nuclear energy generation in stars” [GK07, Wil15].

Following an idea by Bruno Pontecorvo, Raymond Davis Jr. envisioned an experiment to measure the solar- $\nu$  flux and interaction cross sections, and test the predictions of the Standard



**Figure 3.2:** The Solar neutrino flux as a function of the neutrino energy, highlighting the contributions from different nuclear fusion reactions [SGN<sup>+</sup>19].

Solar Model that John N. Bahcall had helped establish [Bah87, BP04]. The Homestake Solar  $\nu$  experiment consisted of a tank with 615 tonnes of tetrachloroethylene ( $C_2Cl_4$ ) at a depth of 1.5 km in the Homestake mine in Lead, South Dakota, USA [CDDJ<sup>+</sup>98]. The rock overburden shielded the experiment from cosmic ray muons, and a water shield surrounding the detector provided shielding from environmental neutrons. Neutrinos coming from the sun interact with the chlorine in the detector and produce, via inverse beta decay, radioactive argon atoms that would be collected and counted. The reaction is given by Equation 3.3.




The detector was only sensitive to higher energy neutrinos, mostly above 5.8 MeV, due to the 0.814 MeV threshold of the reaction, meaning that only *hep* and  ${}^8\text{B}$  neutrinos could contribute to the event rate.

The argon produced in the chlorine target was collected with helium purging, and the number of argon atoms was obtained by measuring the number of  ${}^{37}\text{Ar}$  decays with a proportional counter. Despite the overwhelming difficulty of the task, Davis managed to detect solar neutrinos for the first time<sup>4</sup>. The expected rate of neutrino interactions in the chlorine tank was 7.6 Solar Neutrino Units ( $\text{SNU} \equiv 10^{-36} \nu N_{\text{Cl}}^{-1} \text{s}^{-1}$ ) but Davis only observed around 2.56 SNU, roughly one third of the expected signal from solar neutrinos. This discrepancy between the experimental data and the solar model used to be known as the *Solar Neutrino Problem*, and was later confirmed by the Kamiokande II experiment [H<sup>+</sup>91], and later by other solar neutrino experiments like Gallex [A<sup>+</sup>93], Sage [Gav01] and the upgraded Super-Kamiokande [F<sup>+</sup>02].

The Solar Neutrino Problem was later solved by the SNO collaboration [A<sup>+</sup>01, A<sup>+</sup>02] that could measure the total neutrino flux from the Sun by being sensitive to all neutrino flavours

<sup>4</sup>Raymond Davis Jr. and Masatoshi Koshiba received the Nobel Prize in Physics 2002 “for pioneering contributions to astrophysics, in particular for the detection of cosmic neutrinos”.

via neutral current interactions of neutrinos with deuteron in a heavy water target. Figure 3.3 shows the flux of muon and tau neutrinos,  $\phi(\nu_{\mu\tau})$ , versus the flux of pure electron neutrinos,  $\phi(\nu_e)$ , obtained by SNO [A<sup>+</sup>01]. Since only  $\nu_e$  are produced in the Sun, the detection of all three different neutrino flavours are a strong evidence for oscillations between flavour states as these propagate, which in turn implies that neutrinos have non zero mass [MS16, WX16].



**Figure 3.3:** Solar  ${}^8\text{B}$  neutrino fluxes, showing the measured muon and tau neutrino flux ( $\phi(\nu_{\mu\tau})$ ) obtained from neutral current interactions in SNO versus the combined results from SNO and Super-Kamiokande for the electron neutrino flux ( $\phi(\nu_e)$ ) obtained from charged current interactions ( $\nu_e e$  elastic scattering). The diagonal bands represent the total flux predicted by the BPB01 solar model [BPB01] (dashed) and the total flux measured by SNO and Super-Kamiokande (full). An updated figure showing the most recent combined results can be found in Reference [Z<sup>+</sup>20]. Figure from Reference [A<sup>+</sup>01].

### 3.2.1.2 The Atmospheric Neutrino Anomaly

The atmosphere of the Earth is constantly bombarded with energetic particles. These cosmic rays come from different sources and have energies spanning from a few GeV to hundreds of PeV [A<sup>+</sup>15a, Gai16, R<sup>+</sup>16]. Most of the cosmic rays that reach the Earth are free-streaming protons, and when they hit the Earth's atmosphere they collide with nuclei and generate hadron showers [Bla13]. These showers are cascades of particles being created by other energetic particles as they lose energy in collisions with the material in the atmosphere. In a typical cosmic ray-induced shower a large number of neutrinos are produced due to the creation and subsequent decay of unstable particles within the shower. The relevant neutrino production mechanisms are

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (\text{and h.c.}) \quad (3.4)$$


$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (\text{and h.c.}), \quad (3.5)$$



where h.c. stands for the Hermitian conjugate. The continuous bombardment of the atmosphere by cosmic rays results in a large neutrino flux with typical energies ranging from some tens of MeV up to tens of TeV [R<sup>+</sup>16, F<sup>+</sup>98a]. Equations 3.4 and 3.5 predict that, on average, two muon neutrinos ( $\nu_\mu$  and  $\bar{\nu}_\mu$ ) and one electron neutrino ( $\nu_e$ ) are produced for each pion decay. If all muons decay away, the flux of muon neutrinos should be twice that of the electron neutrino flux<sup>5</sup> [F<sup>+</sup>98a]

$$\frac{\phi(\nu_\mu + \bar{\nu}_\mu)}{\phi(\nu_e + \bar{\nu}_e)} \approx 2. \quad (3.6)$$

Several experiments reported a deficit of muon neutrinos and a lower ratio of muon to electron neutrino events than the one expected under the assumption of no oscillations [RCJ<sup>+</sup>65, A<sup>+</sup>65, H<sup>+</sup>88]. This observation was named the *Atmospheric Neutrino Anomaly*. Super-Kamiokande also observed a deficit of muon neutrinos [F<sup>+</sup>98a], but by measuring the flux of neutrinos coming from the top of the detector and from the bottom, it was able to determine that the source of the atmospheric neutrino anomaly is largely due to disappearing muons, not excess electrons, and that the most likely cause of this effect is neutrino flavour oscillations. These results are displayed in Figure 3.4, where it can be seen that the measured up-to-down asymmetry of  $\mu$ -like events strongly rejects the no-oscillation hypothesis (represented by the hatched bands). Since



**Figure 3.4:** The up-to-down (U-D)/(U+D) asymmetry as a function of momentum for single ring fully contained (FC) and partly-contained (PC) e-like (top) and  $\mu$ -like (bottom) events in Super-Kamiokande. FC events have both the neutrino interaction vertex and resulting particle tracks entirely within the fiducial volume, while PC events only have the former. The hatched bands represent the expected asymmetry under the assumption of no oscillations. The red dashed line on the bottom plot represents the expected asymmetry for  $\mu \leftrightarrow \tau$  oscillations considering maximal mixing. Figure from Reference [F<sup>+</sup>98a].

neutrinos produced in atmospheric showers on the other side of the Earth need to propagate through  $\sim 1.3 \times 10^4$  km of matter, oscillations have time to manifest. On the other hand, the

<sup>5</sup>For higher energies this ratio is larger than 2 since the muon can reach the ground before decaying in flight.

atmospheric neutrinos produced directly above the detector do not propagate a large enough distance to manifest large flavour mixing. With this observation the Super-Kamiokande experiment provided the conclusive evidence that the observed atmospheric neutrinos oscillate between flavour states [F<sup>+</sup>98a].

The observation of *neutrino-mass-induced flavour oscillations* by the Super-Kamiokande and SNO experiments [F<sup>+</sup>98a, F<sup>+</sup>98b, A<sup>+</sup>02], and the consequent demonstration that neutrinos have mass<sup>6</sup>, was the first evidence of physics beyond the Standard Model (BSM), and deeply impacted the path of neutrino physics for years to come. This mechanism is briefly explained below.

### 3.2.2 The Mechanism of Neutrino-Mass-Induced Flavour Oscillations

The neutrino flavour eigenstates  $\nu_l$ ,  $l = e, \mu, \tau$ , that couple in the weak interactions are not the same as the mass eigenstates  $\nu_i$ ,  $i = 1, 2, 3$ . The flavour eigenstates that describe the neutrino weak interactions can be related to the mass eigenstates by a unitary transformation of the type [Pon58, MNS62]

$$|\nu_l\rangle = \sum_{i=1}^3 U_{li}^* |\nu_i\rangle, \quad (3.7)$$

or explicitly:

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix}, \quad (3.8)$$

where  $U_{li}$  is a unitary matrix, called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, that encodes the mixing between the neutrino mass eigenstates and flavour eigenstates.

The unitarity of the PMNS matrix is preserved if the three mass states and the three known flavour states are the only ones contribution to the mixing. Under these conditions, for Dirac neutrinos, the PMNS matrix can be written in terms of three rotation angles  $\theta_{ij}$ , with  $i, j = 1, 2, 3$  and  $i \neq j$ , and a single complex charge-parity (CP) violating phase,  $\delta_{CP}$ , in the form [GK07]

$$\begin{aligned} U &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ -s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{bmatrix}, \quad (3.9) \end{aligned}$$

where  $c_{ij} = \cos \theta_{ij}$  and  $s_{ij} = \sin \theta_{ij}$ .

Considering the neutrinos as free particles, the mass eigenstates evolve in time as plane waves

---

<sup>6</sup>Takaaki Kajita (SK) and Arthur B. McDonald (SNO) received the Nobel Prize in Physics 2015 “for the discovery of neutrino oscillations, which shows that neutrinos have mass”.

[GK07, Z<sup>+</sup>20]:

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i\rangle, \quad (3.10)$$

where  $E_i = \sqrt{\mathbf{p}^2 + m_i^2}$  is the energy associated to the mass eigenstate of  $\nu_i$  with mass  $m_i$ . The time evolution of a flavour eigenstate  $|\nu_a\rangle$  that was created at time  $t = 0$  with a distinct flavour  $a$ , can be derived from Equations 3.7 and 3.10 as

$$|\nu_a(t)\rangle = \sum_i U_{ai}^* e^{-iE_i t} |\nu_i\rangle. \quad (3.11)$$

Using the unitary relation  $U^\dagger U = \mathbf{1}$  the mass eigenstates can be expressed in terms of flavour eigenstates such that

$$|\nu_i\rangle = \sum_a U_{ai} |\nu_a\rangle. \quad (3.12)$$

Combining Equations 3.11 and 3.12, but now considering two different flavour states  $a$  and  $b$ , returns:

$$|\nu_a(t)\rangle = \sum_{b=e,\mu,\tau} \left( U_{ai}^* e^{-iE_i t} U_{bi} \right) |\nu_b\rangle, \quad (3.13)$$

where it is clear that the initial pure flavour state  $|\nu_a(t)\rangle$  transitions to  $|\nu_a(t)\rangle$ , a combination of all three flavour states, as the neutrino propagates (at  $t > 0$ ). The probability of this transition, for a neutrino in the vacuum, is given by

$$P_{\nu_a \rightarrow \nu_b} = |\langle \nu_b | \nu_a(t) \rangle|^2 = \sum_{i,j} U_{ai}^* U_{bi} U_{aj} U_{bj}^* e^{-i(E_i - E_j)t}. \quad (3.14)$$

The energy of ultrarelativistic<sup>7</sup> neutrinos can be approximated to [GK07]

$$E_i \simeq E + \frac{m_i^2}{2E} \quad \Rightarrow \quad E_i - E_j \simeq E + \frac{\Delta m_{ij}^2}{2E}, \quad (3.15)$$

where  $\Delta m_{ij}^2 = m_i^2 - m_j^2$  is the squared-mass difference. The expression for the transition probability  $P_{\nu_a \rightarrow \nu_b}$  as a function of the distance travelled  $L$  and neutrino energy  $E_\nu$ , is then given by

$$P_{\nu_a \rightarrow \nu_b}(L, E_\nu) = \sum_{i,j} U_{ai}^* U_{bi} U_{aj} U_{bj}^* \exp\left(-i \frac{\Delta m_{ij}^2 L}{2E_\nu}\right), \quad (3.16)$$

where the approximation  $t = L$  was used since these neutrinos are in the ultrarelativistic regime [GK07]. If the mass values of the mass eigenstates are not degenerate, the wave functions of each of these states will propagate differently and a neutrino that was created in a specific flavour eigenstate (corresponding to a superposition of mass eigenstates) will oscillate between flavour states as it propagates.

The relevant parameters that describe the flavour transitions are the elements of the PMNS matrix and the squared-mass differences  $\Delta m_{ij}^2$ . Since only two of the squared-mass differences are

---

<sup>7</sup>The high energy thresholds required to detect neutrinos in oscillation experiments, and their very small mass, ensure that they are in an ultrarelativistic regime, so this approximation is valid.

independent ( $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$ ), only  $\Delta m_{21}^2$  and  $\Delta m_{32}^2$  need to be determined experimentally. Solar neutrino oscillations are particularly sensitive to  $\Delta m_{21}^2$  and  $\theta_{21}$ , while atmospheric neutrino oscillations are more sensitive to  $\Delta m_{32}^2$  and  $\theta_{32}$ . Therefore,  $\Delta m_{21}^2$  and  $\Delta m_{32}^2$  are often called the “solar mass difference” and “atmospheric mass difference”, respectively. Reactor neutrino experiments are particularly sensitive to  $\Delta m_{31}^2$  and complement the other measurements, providing overall stronger constraints on the oscillation parameters [Z<sup>+</sup>20].

### 3.2.3 Current results from oscillations

The signal of  $\Delta m_{ij}^2 = m_i^2 - m_j^2$  reveals which of the two mass eigenstates,  $\nu_i$  and  $\nu_j$ , is more massive. However, the absolute values of the neutrino masses cannot be directly probed by oscillation experiments [QV15].

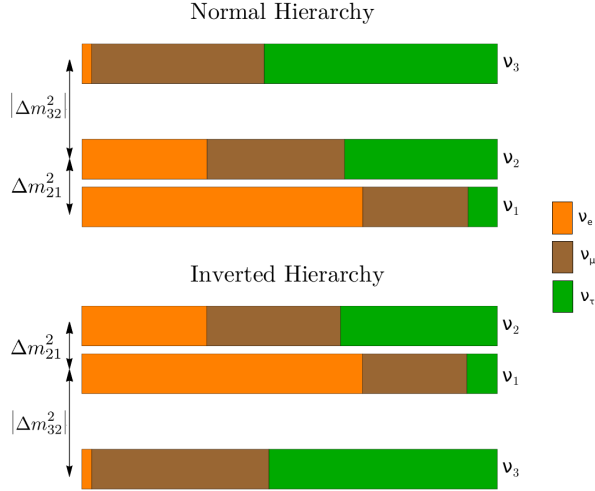
Data from Solar neutrino oscillations has determined that  $\Delta m_{21}^2$  is positive, meaning that  $m_2^2 > m_1^2$ , considering  $\nu_1$  the dominant mass eigenstate of the electron neutrino  $\nu_e$ . The sensitivity to the sign of  $\Delta m_{21}^2$  is due to the Mikheyev-Smirnov-Wolfenstein (MSW) effect [Wol78, MS85], where, in a medium with varying matter density such as the interior of the Sun, the neutrino flavour oscillations are enhanced due to elastic forward scattering interactions between the propagating neutrino and the electron cloud of the medium. Therefore, the behaviour of neutrino oscillations as these propagate through matter is different from that of neutrino propagating in vacuum [Z<sup>+</sup>20] and the difference in the oscillation pattern depends on whether  $m_1$  is larger than  $m_2$  [QV15].

The sign of  $\Delta m_{32}^2$  is still unknown. The size of the matter oscillation length for atmospheric neutrinos crossing the Earth is not enough to determine the sign of either  $\Delta m_{32}^2$  or  $\Delta m_{31}^2$ . Due to the fact that  $\Delta m_{21}^2 \ll |\Delta m_{31}^2|$  (see discussion in Section 3.2.3) the sign of  $\Delta m_{31}^2$  and  $\Delta m_{32}^2$  need to be the same [QV15]. This leads to two possible scenarios:

- If  $\Delta m_{31}^2 > 0$  then  $m_1 < m_2 < m_3$  and the mass hierarchy is referred to as “normal hierarchy” (NH) and the  $\nu_1$  mass eigenstate is the lightest of the three.
- If  $\Delta m_{31}^2 < 0$  then  $m_3 < m_1 < m_2$  and the mass hierarchy is referred to as “inverted hierarchy” (IH) with the  $\nu_3$  state being the lightest.

The “normal” and “inverted” terms are used to maintain an analogy with the mass ordering of the charged leptons, *i.e.*, the normal hierarchy would be the one where the heaviest  $\nu_i$  mass eigenstates contains the lower mixture of the  $\nu_l$  flavour corresponding to the lightest charged lepton. Figure 3.5 shows the two possible neutrino mass distributions and the mixing of mass and flavour eigenstates.

Table 3.3 shows the most recent experimental results of the  $\sin^2 \theta_{ij}$ ,  $\Delta m_{ij}^2$  and  $\delta_{CP}$  oscillation parameters, combined from experiments measuring solar, atmospheric, reactor and accelerator neutrino oscillations [Z<sup>+</sup>20]. The agreement between the different analyses is good, in particular for  $\theta_{12}$ ,  $\theta_{13}$ ,  $\Delta m_{21}^2$  and  $|\Delta m_{32}^2|$ , which are the better known parameters of the model.



**Figure 3.5:** Graphic representation of the normal (top) and inverted (bottom) neutrino mass hierarchies. Oscillation experiments are only sensitive to the differences of squared-masses  $\Delta m_{21}^2 = m_2^2 - m_1^2$  (from solar neutrino oscillations) and  $\Delta m_{32}^2 = m_3^2 - m_2^2$  (from atmospheric neutrino oscillations). Since  $\Delta m_{21}^2 \ll |\Delta m_{32}^2|$ , the squared-mass difference  $\Delta m_{32}^2$  can be redefined as  $\Delta m_{32}^2 \approx m_3^2 - (m_1^2 + m_2^2)/2$ . The sign of  $\Delta m_{32}^2$  is unknown, leading to the two possible mass hierarchies. The colors associated to each mass eigenstate ( $\nu_1, \nu_2, \nu_3$ ) represent the mixing of flavour eigenstates ( $\nu_e, \nu_\mu, \nu_\tau$ ), *i.e.*, probability of finding one of the flavour eigenstates if the neutrino is in a certain mass eigenstate. Figure adapted from Reference [DMVV16].

The mixing angle  $\theta_{23}$  and the leptonic CP phase  $\delta_{CP}$  are not as tightly constrained as the remaining parameters, with  $\theta_{23}$  having a  $3\sigma$  variance of  $9^\circ$  and  $\delta_{CP}$  having a  $3\sigma$  range of  $157^\circ - 349^\circ$ , for the normal hierarchy scenario. It is worth noting that the  $3\sigma$  range of the CP violating phase includes the scenario where  $\delta_{CP} = 180^\circ$ , which corresponds to no CP violation in the lepton sector [Z<sup>+</sup>20]. Finally, the best fit for all parameters seems to disfavour the inverted hierarchy scenario, but the significance of this results is fairly low [Z<sup>+</sup>20].

### 3.2.4 Current constraints on the neutrino masses

As mentioned previously, the absolute value of the neutrino masses cannot be directly probed by oscillation experiments [QV15]. The values in Table 3.3 allow for a variety of scenarios for the absolute mass scale of neutrinos. For example, there is no guarantee that the lightest neutrino mass is non-zero, albeit very small, or that all masses are almost degenerate at a higher mass value, eg.  $m_1 \sim m_2 \sim m_3 \sim 0.1$  eV [dG<sup>+</sup>13].

Nevertheless, oscillation data can be used to set lower bounds to the absolute mass of each neutrino mass eigenstate by considering the mass of the lightest neutrino to be zero, *i.e.*,  $m_{lightest} = 0$ , for each hierarchy scenario. More specifically,  $m_{lightest} = m_1$  for the NH scenario and  $m_{lightest} = m_3$  for the IH scenario. Considering  $\Delta m_{32}^2 \approx m_3^2 - (m_1^2 + m_2^2)/2$ , since

**Table 3.3:** Current best combined results of the 3-neutrino oscillation parameters [Z<sup>+</sup>20].

Parameter	best fit $\pm 1\sigma$	
	Normal Hierarchy	Inverted Hierarchy
$\sin^2 \theta_{12}$	$3.20_{-0.16}^{+0.20} \times 10^{-1}$	$3.20_{-0.16}^{+0.20} \times 10^{-1}$
$\sin^2 \theta_{23}$	$5.47_{-0.30}^{+0.20} \times 10^{-1}$	$5.51_{-0.30}^{+0.18} \times 10^{-1}$
$\sin^2 \theta_{13}$	$2.160_{-0.069}^{+0.083} \times 10^{-2}$	$2.220_{-0.076}^{+0.074} \times 10^{-2}$
$\Delta m_{21}^2$ [eV <sup>2</sup> ]	$7.55_{0.16}^{+0.20} \times 10^{-5}$	$7.55_{0.16}^{+0.20} \times 10^{-5}$
$\Delta m_{32}^2$ [eV <sup>2</sup> ]	$2.424_{0.03}^{+0.03} \times 10^{-3}$	$-2.50_{0.03}^{+0.04} \times 10^{-3}$
$\delta_{CP}$ [deg]	$218_{-27}^{+38}$	$281_{-27}^{+23}$

$\Delta m_{21}^2 \ll |\Delta m_{32}^2|$ , the following limits can be set using the results from Table 3.3:

$$(m_1, m_2, m_3) \geq \begin{cases} \left( 0, \sqrt{\Delta m_{21}^2}, \sqrt{|\Delta m_{32}^2| + \frac{\Delta m_{21}^2}{2}} \right) & (NH) \\ \left( \sqrt{|\Delta m_{32}^2| - \frac{\Delta m_{21}^2}{2}}, \sqrt{|\Delta m_{32}^2| + \frac{\Delta m_{21}^2}{2}}, 0 \right) & (IH) \end{cases} \quad (3.17)$$

$$\geq \begin{cases} (0, 0.87, 4.96) \times 10^{-2} \text{ eV} & (NH) \\ (4.96, 5.04, 0) \times 10^{-2} \text{ eV} & (IH) \end{cases} \quad (3.18)$$

Upper bounds on the neutrino mass can be obtained from cosmological observations, kinematic studies of beta decay, and searches for the neutrinoless double beta ( $0\nu\beta\beta$ ) decay process (discussed in detail in Section 3.4).

Measuring the endpoint of a  $\beta$ -decay spectrum with high precision places constraints on the masses of the particles involved, namely the nucleus, the  $\beta$  particle and the (anti)neutrino, due to energy-momentum conservation. The  $\beta$ -decay of  ${}^3\text{H}$  is advantageous for this measurement for two main reasons: first, this decay is a superallowed transition, and the nuclear matrix elements (NMEs) are energy independent, meaning that the  $\beta$ -spectrum is characterized only by the endpoint energy  $Q_\beta \equiv E_0 - m_e$ , where  $E_0$  is the mass difference of the initial and final nucleus and  $m_e = 0.511$  MeV is the mass of the electron [Z<sup>+</sup>20]. For a non-zero neutrino mass, the maximum kinetic energy of the electron is given by  $T_{max} = Q_\beta - m_{\nu_e}^{eff}$  [Z<sup>+</sup>20], where  $m_{\nu_e}^{eff}$  is the effective mass of the neutrino<sup>8</sup> defined as

$$m_{\nu_e}^{eff} = \sqrt{\sum_i m_i^2 |U_{ei}|^2}. \quad (3.19)$$

Second, the low endpoint energy of the process ( $Q_\beta = 18.6$  keV) makes this decay very sensitive to  $m_{\nu_e}^{eff}$ . The most recent results from KATRIN [A<sup>+</sup>19d] do not exclude  $m_{\nu_e}^{eff} = 0$  and set an

<sup>8</sup>Sometimes the effective neutrino mass obtained from  $\beta$ -decay measurements is denoted  $m_\beta$ .

upper limit of

$$m_{\nu_e}^{eff} < 1.1 \text{ eV (90\% CL)}. \quad (3.20)$$

The current best limit on  $0\nu\beta\beta$  decay half-life comes from  $^{136}\text{Xe}$  from the KamLAND-Zen experiment [G<sup>+</sup>16]. One can define an effective Majorana mass of  $\nu_e$  that governs the  $0\nu\beta\beta$  decay process as

$$\langle m_{\beta\beta} \rangle = \left| \sum_i m_i U_{ei}^2 \right|, \quad (3.21)$$

where  $m_i$  are the neutrino masses and  $U_{ei}$  is the first row of the PMNS matrix. For the largest and smallest values of the relevant NMEs in the literature, the upper limit on the effective Majorana neutrino mass from KamLAND-Zen is

$$m_{\beta\beta} < 61 - 165 \text{ meV (90\% CL)}. \quad (3.22)$$

It is important to notice that the limits on the neutrino mass obtained from  $0\nu\beta\beta$  decay searches are under the assumption that neutrinos are Majorana particles, a necessary condition for the process to occur, and that the process is not dominantly induced by other physics processes beyond the exchange of a massive Majorana neutrino (for a more complete discussion on the possible physics mechanisms for  $0\nu\beta\beta$  decay see Section 3.4).

The analysis of the CMB power spectra, combined with analysis of BAO and lensing effects, places strong constraints on the sum of all neutrino masses. The most recent results from the Planck collaboration [A<sup>+</sup>18a] provide an upper limit to the sum of all neutrino masses at

$$\sum m_\nu < 0.12 \text{ eV (95\% CL)}. \quad (3.23)$$

### 3.3 Theoretical Aspects of Neutrino Masses

In the Standard Model (SM) neutrinos are assumed to be massless Weyl particles [WX16]. They participate in the weak charged and neutral interactions, defined by the  $SU(2)_L \times U(1)_Y$  electroweak gauge group, and are represented in doublets of that electroweak group with the respective charged leptons. The Lagrangian densities for neutrino interactions with the massive weak bosons ( $W^\pm$  and  $Z^0$ ) are defined as

$$\begin{aligned} \mathcal{L}_{weak,\nu} &= \mathcal{L}_{nc} + \mathcal{L}_{cc} \\ &= \frac{g_2}{4 \cos \theta_w} \sum_l [\bar{\nu}_l \gamma^\mu (1 - \gamma_5) \nu_l] Z_\mu + \frac{g_2}{2\sqrt{2}} \sum_l [\bar{l} \gamma^\mu (1 - \gamma_5) \nu_l W_\mu^- + \text{h.c.}], \end{aligned} \quad (3.24)$$

where  $\theta_w$  is the Weinberg angle,  $g_2$  is the coupling constant associated to  $SU(2)$ , NC stands for *neutral current*, CC for *charged current* and  $l = e, \mu, \tau$  is the lepton flavour label. Notice that, at energies below 11 GeV, neutrinos of any of the three flavours can interact with electrons via NC interactions but only an electron neutrino can interact with an electron via charged current interactions. For neutrino energies above  $m_\mu^2/2m_e \approx 11 \text{ GeV}$  and  $m_\tau^2/2m_e \approx 3 \text{ TeV}$ ,

muon and tau neutrinos can interact with an electron via CC interactions through the process  $\nu_l + e \rightarrow \nu_e + l$  known as “inverse muon decay” and “inverse tau decay”, respectively. At tree level, only electron neutrinos contribute to the inverse beta decay (Equation 3.2) via CC interactions but all neutrino flavours can contribute to neutrino-electron elastic scattering ( $\nu_l + e \rightarrow \nu_l + e$ ) and neutrino-nucleus scattering ( $\nu_l + N \rightarrow \nu_l + N$ ) via NC interactions.

It is well established today that neutrinos are not massless, as discussed in Section 3.2.2, which indicates that the Standard Model is an incomplete theory and most likely an effective low-energy theory resulting from the spontaneous symmetry breaking of a much larger symmetry group [Z<sup>+</sup>20]. It follows that some mechanism beyond the Standard Model is responsible for giving neutrinos their masses. There are two possible descriptions of the massive neutrino: Dirac and Majorana. The mass terms that arise from each description are discussed below, along with other mass generating mechanisms.

### 3.3.1 Neutrino Mass Models

As spin-1/2 particles, neutrinos can be described by four-component wavefunctions  $\psi(x)$  (spinors) that obey the Dirac equation<sup>9</sup> [Zub03]

$$(i\gamma^\mu \partial_\mu - m) \psi(x) = 0, \quad (3.25)$$

where  $m$  is the neutrino mass and  $\gamma_\mu$  are the  $4 \times 4$  gamma matrices of the form

$$\gamma_0 = \begin{pmatrix} \mathbb{1} & 0 \\ 0 & -\mathbb{1} \end{pmatrix} \quad \gamma_i = \begin{pmatrix} 0 & \sigma_i \\ -\sigma_i & 0 \end{pmatrix}, \quad (3.26)$$

where  $\mathbb{1}$  is the  $2 \times 2$  identity matrix and  $\sigma_i$  are the  $2 \times 2$  Pauli matrices<sup>10</sup>. It is useful to define  $\gamma_5 = i\gamma_0\gamma_1\gamma_2\gamma_3$ , known as the chirality matrix [GK07]. The four independent components of  $\psi(x)$  represent the four combinations obtained for particles and antiparticles with the two possible spin projections  $J_Z = \pm 1/2$  (helicities  $\mathcal{H} = \pm 1$ ). However, only left-handed neutrinos ( $\mathcal{H} = -1$ ) and right-handed antineutrinos ( $\mathcal{H} = +1$ ) are observed [Zub03]. A two-component description based on Weyl spinors can then be defined using the projection operators  $P_L = 1/2(1 - \gamma_5)$  and  $P_R = 1/2(1 + \gamma_5)$  as

$$\psi_L = P_L \psi \quad \psi_R = P_R \psi, \quad (3.27)$$

where  $\psi_{L,R}$  are the chiral projections of  $\psi$  such that [Zub03]

$$\psi = (P_L + P_R) \psi = P_L \psi + P_R \psi = \psi_L + \psi_R. \quad (3.28)$$

Notice that the projection operator  $P_L$  is present in Equation 3.24, indicating that the interacting neutrino is always left-handed and the antineutrino always right-handed.

<sup>9</sup>Using Einstein conventions.

<sup>10</sup>For a complete list of properties of the Pauli matrices and gamma matrices see, for example, Reference [GK07].



By multiplying Equation 3.25 on both sides by  $P_L$  or  $P_R$  the following field equations are obtained<sup>11</sup>

$$i\gamma^\mu \partial_\mu \psi_R = m\psi_L, \quad (3.29)$$


$$i\gamma^\mu \partial_\mu \psi_L = m\psi_R. \quad (3.30)$$

These equations demonstrate that the space-time evolutions of the chiral fields  $\psi_{L,R}$  are related by the Dirac mass  $m$  [GK07, Z<sup>+</sup>20]. If the neutrino mass is zero, the two chiral fields decouple completely and these equations become the Weyl equations, as expected for the SM massless neutrino. However, it is known today that the neutrino does have mass, a conclusion drawn from the observations of neutrino flavour oscillations (see Section 3.2). The term for the neutrino mass couples the left and right-hand chiral projections on Equations 3.29 and 3.30, which indicates that the neutrino cannot be fully described by a single chiral field with two independent components. A Dirac mass term can be constructed using the left and right-handed components of the chiral fields:

$$\mathcal{L}^D = -m_D (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L). \quad (3.31)$$

However, this mass term is not invariant under SM gauge transformations (each component transforms differently under an SU(2) transformation). A possible solution to this problem is to promote the mass to a dynamical field described by an SU(2) doublet  $\phi$ , *i.e.*, the Higgs field, as explained below.

Fermions in the Standard Model gain mass via the Higgs mechanism and the Yukawa interactions with the scalar Higgs field [Z<sup>+</sup>20]. The Higgs mechanism is tied to the phenomenon of electroweak symmetry breaking [Kan17] and the mass scale of the symmetry breaking is the vacuum expectation value (VEV) of the Higgs field,  $v \sim 246$  GeV. But neutrino masses are fundamentally different from the masses of the other fermions as one can see from analysing Figure 3.6. There is a gap of six orders of magnitude between the largest allowed neutrino mass and the mass of the next lightest fermion, the electron.



**Figure 3.6:** Masses of the SM fermions. The neutrino masses presented consider the normal hierarchy scenario. The six orders of magnitude gap between the mass scale of neutrinos and the remaining fermions could indicate a new mass scale and a different mass generating mechanism for neutrinos.

<sup>11</sup>It is important to mention the following properties of the projection operators  $P_L \gamma^0 = \gamma^0 P_R$ ,  $P_R \gamma^0 = \gamma^0 P_L$ ,  $P_L^2 = P_L$ ,  $P_R^2 = P_R$  and  $P_L P_R = P_R P_L = 0$ .

Recalling again the result on Equations 3.29 and 3.30, the simplest way to introduce a neutrino mass term is to consider the neutrino a Dirac fermion like the charged leptons, which involves introducing a right-handed field  $\nu_R$  for each neutrino state in what is called a minimal extension of the SM. The masses of the fermions in the SM are generated via Yukawa couplings of the right-handed (singlet) and left-handed (doublet) components of the fermion fields with the scalar Higgs field  $\phi$  [Z<sup>+</sup>20]. Such mass terms are of the form

$$\mathcal{L}_{\phi,l} = Y_{ij}^l \bar{L}_{Li} \phi E_{Rj} + \text{h.c.}, \quad (3.32)$$

where  $L_{Li}$  are the left-handed lepton doublets,  $E_{Rj}$  are the right-handed charged lepton fields and  $Y_{ij}^l$  are the charged lepton and neutrino Yukawa couplings. A Dirac mass term  $m_{ij}^l$  arises after spontaneous symmetry breaking [Kan17]:

$$m_{ij}^l = Y_{ij}^l \frac{v}{\sqrt{2}}. \quad (3.33)$$

This mass is proportional to the VEV of the Higgs field, similarly to the mass terms of the charged fermions in the SM. However, the masses of the neutrinos are at least  $10^6$  times smaller than the lightest charged fermion, which would indicate that the neutrino-Higgs coupling is smaller than that of the other fermions by at least the same factor [Z<sup>+</sup>20].

Another solution to the origin of neutrino mass can be found if these particles are Majorana, *i.e.*, the same fermion field describes both the particle and antiparticle states. Such a field would satisfy the Majorana condition, *i.e.*, would be invariance under charge conjugation:

$$\psi = \psi^{\mathcal{C}} = \mathcal{C} \bar{\psi}^T, \quad (3.34)$$

where  $\mathcal{C}$  is the charge conjugation operator, that can be defined as  $\mathcal{C} \equiv i\gamma^2\gamma^0$ , and has the following properties:  $\mathcal{C}^\dagger = \mathcal{C}^{-1} = \mathcal{C}^T = -\mathcal{C}$ . This operator, when acting on a particle state, will return the antiparticle state without modifying the helicity (chirality) of the state [Zub03]. Equation 3.34 only holds for a neutral fermion, and the only neutral fermion in the SM is the neutrino.

Taking the Hermitian conjugate of Equation 3.29 and multiplying it on the right by  $\gamma^0$  yields

$$-i\partial_\mu \bar{\psi}_R \gamma^\mu = m \bar{\psi}_L. \quad (3.35)$$

Now taking the transpose of Equation 3.35 and multiplying on the left by the charge conjugation operator  $\mathcal{C}$  yields

$$i\gamma^\mu \partial_\mu \mathcal{C} \bar{\psi}_R^T = m \mathcal{C} \bar{\psi}_L^T. \quad (3.36)$$

Equations 3.36 and 3.30 become identical considering  $\psi_R = \xi \mathcal{C} \bar{\psi}_L^T$ . Here  $\xi$  is an arbitrary phase and can be absorbed in the redefinition of  $\psi_R$ . Notice that  $\mathcal{C} \bar{\psi}_L^T$  is a right-handed chiral field that transforms as  $\psi_L$  under Lorentz transformations. Being right-handed, the coupling  $\bar{\psi}_L^{\mathcal{C}} \psi_L$

is non zero. Equation 3.30 can now be written as the Majorana field Equation [GK07]:

$$i\gamma^\mu \partial_\mu \psi_L = m\mathcal{C}\bar{\psi}_L^T. \quad (3.37)$$

Equation 3.28 can be rewritten using the Majorana condition in Equation 3.34 as

$$\psi = \psi_L + \psi_R = \psi_L + \mathcal{C}\bar{\psi}_L^T = \psi_L + \psi_L^{\mathcal{C}}. \quad (3.38)$$

A Majorana mass term can be constructed much alike a typical Dirac mass, such as

$$\mathcal{L}_{mass}^M = -m\bar{\psi}\psi = -\frac{1}{2}m\overline{\psi_L^{\mathcal{C}}}\psi_L. \quad (3.39)$$

The overall factor of  $1/2$  in Equation 3.39 is because  $\psi_L^{\mathcal{C}}$  and  $\bar{\psi}_L$  are not independent terms. Notice that this mass term, just like the Dirac mass term in Equation 3.31, is not invariant under SM gauge transformations and breaks  $U(1)_Y$  by two units [GK07].

One can take the approach of building an effective field theory with heavier degrees of freedom, that includes the gauge symmetries of the Standard Model, and provide a mechanism to relax the neutrino mass to the small values observed [dG<sup>+</sup>13, WX16]. Considering only the first generation of leptons, the lowest dimensional operator that can generate a Majorana mass using only the available SM content and respecting the SM symmetries is the dimension-5 Weinberg operator  $\mathcal{M}^{-1}L_L L_L \Phi \Phi$ , with  $\Phi$  and  $L_L$  being the Higgs and lepton  $SU(2)_L$  doublets, respectively. The term  $\mathcal{M}$  has dimensions of mass and represents the new physics scale. This operator can be explicitly written as

$$\mathcal{L}_5 = \frac{1}{2} \frac{g}{\mathcal{M}} \left( L_L^T \mathcal{C}^\dagger \sigma_2 \vec{\sigma} L_L \right) \cdot \left( \Phi^T \sigma_2 \vec{\sigma} \Phi \right) + \text{h.c.}, \quad (3.40)$$

where  $g$  is a dimensionless coupling coefficient and  $\vec{\sigma}$  is the Pauli vector [GK07]. This term explicitly violates lepton number by two units ( $\Delta L = 2$ ). After symmetry breaking some mass term proportional to  $v^2 \mathcal{M}^{-1}$  is generated:

$$\mathcal{L}_{mass}^M = \frac{1}{2} \frac{gv^2}{\mathcal{M}} \phi_L^T \mathcal{C}^\dagger \phi_L + \text{h.c.}, \quad (3.41)$$

where  $\phi_L$  is the left-handed neutrino chiral field. This term returns a Majorana mass

$$m = \frac{gv^2}{\mathcal{M}} \propto \frac{m_D^2}{\mathcal{M}}, \quad (3.42)$$

with  $m_D$  the Dirac mass for the neutrino obtained from the Higgs mechanism in Equation 3.33. If the scale  $\mathcal{M}$  is large, the neutrino mass becomes very small, like in the *see-saw* mechanism [GK07]. For example, if  $\mathcal{M} \sim 10^{15}$  GeV, a plausible grand unification scale, then  $m \sim 10^{-2}$  eV is consistent with current limits on neutrino masses (see Section 3.2.4).

### 3.3.2 Leptogenesis

The universe today appears to contain very little antimatter. There are strong indications that matter (baryons) and antimatter (antibaryons) should have been created in equal amounts in the Big Bang, and if that was the case then one would expect that both components would have annihilated completely, leaving only  $\gamma$ -rays in a matter-free universe. However, to the best of our understanding, all observed structures in the universe seem to be formed by baryons and not antibaryons. So it follows that some asymmetry in the laws of physics must exist that generated this baryon-antibaryon asymmetry, a scenario called “baryogenesis” [DNN08, Z<sup>+</sup>20].

The baryon asymmetry is still quite small. The density of baryons compared to the density of photons today is of the order of [Z<sup>+</sup>20]

$$\eta = \frac{n_B - n_{\bar{B}}}{n_\gamma} \sim 10^{-9}, \quad (3.43)$$

where  $n_B$ ,  $n_{\bar{B}}$  and  $n_\gamma$  are the number densities of baryons, antibaryons and photons at the current time, respectively. This asymmetry was most likely generated dynamically, otherwise it would have been a very finely tuned initial condition, which seems unlikely. Furthermore, any primordial baryon asymmetry would dilute away due to cosmic inflation [DNN08].

There are three prerequisites for baryogenesis [Sak91, DNN08]:

- Baryon number ( $B$ ) violation – self-evident condition to evolve from  $n_B - n_{\bar{B}} = 0$  to  $n_B - n_{\bar{B}} \neq 0$ .
- Out of equilibrium dynamics, since chemical/thermal equilibrium averages out asymmetries in non-conserving quantum numbers – provided by the expansion of the universe and at higher orders by the electroweak phase transition.
- Charge ( $C$ ) and charge-parity ( $CP$ ) violation – if these were conserved, then processes involving baryons would proceed at precisely the same rate as processes involving antibaryons, and no baryon asymmetry would be generated.


These prerequisites are present in the SM. However, no SM mechanism can generate the required amount of baryon asymmetry: baryon number is violated in the quark sector, but the effect is not enough to produce  $\eta \sim 10^{-9}$  alone; weak interactions violate  $C$  maximally and also violate  $CP$  via the Kobayashi-Maskawa mechanism [KM73] but this amount of  $CP$  violation is not enough to justify the observed baryon asymmetry as well [DNN08]. Other sources of  $CP$  violation must exist in order to justify observations.

The non-zero mass of the neutrino implies that the lepton sector also violates  $CP$  [Z<sup>+</sup>20]. If neutrinos are Majorana fermions, heavy right-handed neutrinos arising from the *see-saw* mechanism would be allowed to decay via their Yukawa coupling to the Higgs field, with the Majorana phases  $\xi_i$  providing a  $CP$  asymmetry that would differentiate between the coupled leptons and antileptons. This difference would result in distinct decay rates that could favour the accumulation of leptons over antileptons. This mechanism is called “leptogenesis” [Z<sup>+</sup>20].

The matter-antimatter asymmetry observed in the universe seems to favour the hypothesis that lepton number-violating processes can occur [Z<sup>+</sup>20]. Furthermore, it provides a strong argument for the Majorana nature of the neutrino, which in turn motivates the neutrinoless double beta ( $0\nu\beta\beta$ ) decay. Therefore, the search for  $0\nu\beta\beta$  decays, a lepton number-violating process that results from the Majorana nature of the neutrino, is an important test for leptogenesis, along with the observation of CP violation in neutrino oscillations [Z<sup>+</sup>20].


### 3.4 Double Beta Decay

Some isotopes are known to undergo a type of decay during which the nucleus emits two electrons and two electron antineutrinos [DMVV16, DPR19]. This decay mode is denoted two-neutrino double beta decay ( $2\nu\beta\beta$ ) and can occur in even-even nuclei when the single beta decay is energetically forbidden or at least highly suppressed [Saa13, DMVV16, DPR19]. Figure 3.7 shows a scheme of the nuclear mass as a function of the atomic number  $Z$  (mass parabolas) for a  $2\nu\beta\beta$  decay candidate isotope. Due to the nuclear pairing force, even-even nuclei are lighter than the respective isobar odd-odd ( $A, Z+1$ ) nucleus, making the single  $\beta$ -decay energetically forbidden [DMVV16].



**Figure 3.7:** Nuclear mass as a function of the atomic number  $Z$  for a  $2\nu\beta\beta$  decay candidate with  $A$  even (left) and  $A$  odd (right). For the even-even isotope (lower mass parabola on the left plot) the single  $\beta$ -decay is kinematically forbidden but the  $2\nu\beta\beta$  decay is not. For odd mass number nuclei (mass parabola on the right plot) the  $2\nu\beta\beta$  decay is strongly suppressed as the single  $\beta$ -decay would be dominant. Figure from Reference [DMVV16].

The  $2\nu\beta\beta$  decay process varies the atomic number  $Z$  of the decaying isotope by two units due to the conversion of two neutrons into two protons, releasing two electrons that can be detected and two electron antineutrinos which will avoid detection. This process is represented by Equation 3.44, where  $A$  and  $Z$  are the atomic mass and atomic number of the decaying isotope, respectively:



The  $2\nu\beta\beta$  decay is not the only four lepton decay mode that is allowed by the Standard Model. Equations 3.45, 3.46 and 3.47 represent other allowed decay modes, namely the two-neutrino double electron capture ( $2\nu\text{ECEC}$ ), the two-neutrino electron capture with positron emission ( $2\nu\text{EC}\beta^+$ ) and the two-neutrino double positron emission ( $2\nu\beta^+\beta^+$ ).

$$(A, Z) + 2e^- \longrightarrow (A, Z - 2) + 2\nu_e \quad (2\nu\text{ECEC}), \quad (3.45)$$

$$(A, Z) + e^- \longrightarrow (A, Z - 2) + e^+ + 2\nu_e \quad (2\nu\text{EC}\beta^+), \quad (3.46)$$

$$(A, Z) \longrightarrow (A, Z - 2) + 2e^+ + 2\nu_e \quad (2\nu\beta^+\beta^+), \quad (3.47)$$

The electron capture process is favoured for electrons from the atomic K-shell, with capture of electrons from higher shells being successively less probable. All these decay modes are kinematically unfavoured with respect to the  $2\nu\beta\beta$  decay due to their limited phase space [Saa13]. The decays can also occur into an excited state of the final nucleus, with the energy level of that nucleus,  $\epsilon$ , further weighing on the suppression of the decay. This results in longer half-lives that make the observation of these decay modes more challenging [Saa13]. The  $2\nu\text{EC}\beta^+$  and  $2\nu\beta^+\beta^+$  decay modes have not been experimentally observed.

Table 3.4 displays the half-lives and Q-values of several radioactive isotopes observed to decay via the  $2\nu\beta\beta$  and  $2\nu\text{ECEC}$  processes. The isotope  $^{136}\text{Xe}$  is of particular interest to the LZ experiment (presented in detail in Chapter 4) since it comprises 8.9% of naturally occurring xenon. This isotope has been shown to undergo  $2\nu\beta\beta$  with a half-life of  $2.165 \pm 0.016^{(\text{stat})} \pm 0.059^{(\text{sys})} \times 10^{21}$  years with a Q-value of  $2457.83 \pm 0.37$  keV [A<sup>+</sup>14b, RWMM07].

The two-neutrino double beta decay processes are an important tool to study the structure of atomic nuclei, most notably the experimental access to the values of nuclear matrix elements (NMEs) that are used to constrain nuclear models that are also used for other decay processes that are not allowed by the Standard Model [Saa13]. In particular, the observation of  $2\nu\beta\beta$  decay on two isotopes of the same element allows the probing of the NMEs with great precision, since the theoretical uncertainties associated with different NME models are expected to partly cancel out [BTZ14]. There are 9 elements with access to  $2\nu\beta\beta$  decay in more than one isotope:  $^{46,48}\text{Ca}$ ,  $^{80,82}\text{Se}$ ,  $^{94,96}\text{Zr}$ ,  $^{98,100}\text{Mo}$ ,  $^{114,116}\text{Cd}$ ,  $^{122,124}\text{Sn}$ ,  $^{128,130}\text{Te}$ ,  $^{134,136}\text{Xe}$  and  $^{146,148,150}\text{Nd}$ . So far, only studies on  $^{128}\text{Te}$  and  $^{130}\text{Te}$  have been performed [BC10]. LZ will be sensitive to the  $2\nu\beta\beta$  decay of both  $^{134}\text{Xe}$  and  $^{136}\text{Xe}$ , allowing the direct comparison of both decays and, if the former is observed, the probing of NMEs for these isotopes as well [LZ 20].

### 3.4.1 Neutrinoless Double Beta Decay

A different and yet unobserved “neutrinoless” double beta decay ( $0\nu\beta\beta$ ) mode by which a nucleus emits two electrons and no neutrinos could potentially occur if the neutrino is its own antiparticle, *i.e.*, a Majorana neutrino. This SM-forbidden process violates lepton number conservation ( $\Delta L = 2$ ).

**Table 3.4:** Measured half-lives and Q-values of two-neutrino double beta decay processes ( $2\nu\beta\beta$ ,  $2\nu\text{ECEC}$ ) of several isotopes. The  $2\nu\text{ECEC}$  decays of  $^{78}\text{Kr}$  and  $^{124}\text{Xe}$  are the longest decays ever observed. As a reference, the age of the Universe is estimated around  $13.8 \times 10^9$  years. The  $2\nu\beta\beta$  decay of  $^{238}\text{U}$  and the  $2\nu\text{ECEC}$  decay of  $^{130}\text{Ba}$  are observed indirectly through radiochemistry and geochemistry techniques, respectively. From Reference [Bar20].

Decay	Isotope	Experiment	$T_{1/2}^{2\nu}$ [yr]	$Q_{\beta\beta}$ [MeV]
$2\nu\beta\beta$	$^{48}\text{Ca}$	NEMO-3 [ $A^+16i$ ]	$6.4_{\pm 1.2}^{\pm 0.7(\text{stat.})} \times 10^{19}$	4.268
	$^{76}\text{Ge}$	GERDA [ $A^+15c$ ]	$(1.925 \pm 0.094) \times 10^{21}$	2.039
	$^{82}\text{Se}$	CUPID-0 [ $A^+19j$ ]	$8.6_{\pm 0.19}^{\pm 0.03(\text{stat.})} \times 10^{19}$	2.992
	$^{96}\text{Zr}$	NEMO-3 [ $A^+10$ ]	$2.35_{\pm 0.16}^{\pm 0.14(\text{stat.})} \times 10^{19}$	3.350
	$^{100}\text{Mo}$	CUPID-Mo [ $A^+20i$ ]	$7.12_{\pm 0.10}^{\pm 0.18(\text{stat.})} \times 10^{18}$	3.034
	$^{116}\text{Cd}$	Aurora [ $B^+18$ ]	$2.63_{-0.12}^{+0.11} \times 10^{19}$	2.813
	$^{130}\text{Te}$	CUORE [ $N^+20$ ]	$7.9_{\pm 0.2}^{\pm 0.1(\text{stat.})} \times 10^{20}$	2.527
	$^{136}\text{Xe}$	KamLAND-Zen [ $G^+16$ ]	$2.21_{\pm 0.07}^{\pm 0.02(\text{stat.})} \times 10^{21}$	2.458
	$^{150}\text{Nd}$	NEMO-3 [ $A^+16h$ ]	$9.34_{\pm 0.62}^{\pm 0.22(\text{stat.})} \times 10^{18}$	3.371
	$^{238}\text{U}$	Radiochemistry [TEC91]	$(2.0 \pm 0.6) \times 10^{21}$	1.1
$2\nu\text{ECEC}$	$^{78}\text{Kr}$	Gas LPC [GGK <sup>+</sup> 13]	$1.9_{\pm 0.3}^{-0.7+1.3(\text{stat.})} \times 10^{22}$	$\sim 10^{-2}$
	$^{124}\text{Xe}$	XENON1T [ $A^+19i$ ]	$1.8_{\pm 0.1}^{\pm 0.5(\text{stat.})} \times 10^{22}$	2.857
	$^{130}\text{Ba}$	Geochemistry [MHPK01]	$(2.2 \pm 0.5) \times 10^{21}$	—

Each of the double beta decay processes mentioned above will have a neutrinoless counterpart. Equations 3.48, 3.49, 3.50 and 3.51 represent the respective SM-forbidden neutrinoless decay modes for the double beta decay ( $0\nu\beta\beta$ ), double electron capture ( $0\nu\text{ECEC}$ ), electron capture with positron emission ( $0\nu\text{EC}\beta^+$ ) and double positron emission ( $0\nu\beta^+\beta^+$ ).


$$(A, Z) \longrightarrow (A, Z + 2) + 2e^- \quad (0\nu\beta\beta) \quad (3.48)$$

$$(A, Z) + 2e^- \longrightarrow (A, Z - 2) \quad (0\nu\text{ECEC}), \quad (3.49)$$

$$(A, Z) + e^- \longrightarrow (A, Z - 2) + e^+ \quad (0\nu\text{EC}\beta^+), \quad (3.50)$$

$$(A, Z) \longrightarrow (A, Z - 2) + 2e^+ \quad (0\nu\beta^+\beta^+), \quad (3.51)$$

There are several models that predict the occurrence of  $0\nu\beta\beta$  decay, some involving light-neutrino mediators with sub-eV masses while others invoke GeV to TeV heavy particle mediators like heavy right-handed neutrino exchange or R-parity violating supersymmetry ( $\mathcal{R}_p\text{SUSY}$ ) mechanisms [B<sup>+</sup>05, DPR19]. All these models have the commonality of requiring physics beyond



**Figure 3.8:** First order diagram of the  $0\nu\beta\beta$  decay process mediated by the exchange of a light massive Majorana neutrino  $\nu_M$ . Figure from Reference [DMVV16].

the Standard Model. The light neutrino exchange model is often considered the most appealing mechanism that mediates the  $0\nu\beta\beta$  decay [DMVV16]. Figure 3.8 shows the diagram of the  $0\nu\beta\beta$  decay process mediated by the exchange of a light massive Majorana neutrino,  $\nu_M$ . As there are no neutrinos emitted from this process that can carry momentum, the energy deposited by the two electrons emitted is equal to the Q-value of the decay,  $Q_{\beta\beta}$ .

The “effective Majorana mass”,  $\langle m_{\beta\beta} \rangle$ , that governs the  $0\nu\beta\beta$  decay process is given by Equation 3.21. However, since the neutrino has to be Majorana for the  $0\nu\beta\beta$  decay to occur, there are 3 additional CP violating phases  $\xi_i$  in the mixing matrix  $U$  that cannot be rotated away and that contribute to the value of  $\langle m_{\beta\beta} \rangle$ . These are called Majorana phases [DMVV16, GK07]. The PMNS matrix can be rewritten to include the Majorana phases as  $U = U^D D^M$ , where  $U^D$  is the original PMNS matrix presented in Equation 3.9 that contains the Dirac phase, and

$$D^M = \text{diag} \left( e^{i\xi_1}, e^{i\xi_2}, e^{i\xi_3} \right), \quad \xi_1 = 0, \quad (3.52)$$

is the diagonal unitary matrix containing the three Majorana phases  $\xi_i$ . Only the differences between the phases can be observed, so by convention  $\xi_1 = 0$ . The contributions of the Majorana phases to  $\langle m_{\beta\beta} \rangle$  can be made explicitly as

$$\langle m_{\beta\beta} \rangle = \left| \sum_i m_i e^{i\xi_i} U_{ei}^2 \right|, \quad (3.53)$$


where  $m_i$  are the neutrino masses and  $U_{ei}$  is the first row of the  $U^D$  PMNS matrix. Since the two leptonic vertices of the  $0\nu\beta\beta$  decay process only involve the electron flavour, only the first row of  $U^D$  contributes to  $\langle m_{\beta\beta} \rangle$ . In the light neutrino exchange mechanism, the lifetime of the  $0\nu\beta\beta$  decay process is proportional to  $\langle m_{\beta\beta} \rangle$  [DPR19],

$$\left( T_{1/2}^{0\nu} \right)^{-1} = G^{0\nu} |\mathcal{M}^{0\nu}|^2 \frac{\langle m_{\beta\beta} \rangle^2}{m_e^2}, \quad (3.54)$$

where  $G^{0\nu} = G_{(0)}^{0\nu} g_A^4$  contains both the phase-space factor  $G_{(0)}^{0\nu}$  and the axial vector coupling



constant  $g_A$ , and  $\mathcal{M}^{0\nu}$  is the nuclear matrix element (NME) for the  $0\nu\beta\beta$  process. The large uncertainties on the NMEs and the neutrino mass scale make the prediction of the half-life of  $0\nu\beta\beta$  decay challenging. On the other hand, placing constraints on the half-life of  $0\nu\beta\beta$  decay would also constrain the value of  $\langle m_{\beta\beta} \rangle$ , as discussed in Section 3.2.4, which can provide some insights on the absolute values of the neutrino masses and their hierarchy (see the discussion in Section 3.2.3) and could also be used to probe the presence of additional mass eigenstates mixing into the electron neutrino [DPR19]. Figure 3.9 shows the allowed values of  $\langle m_{\beta\beta} \rangle$  in each hierarchy scenario as a function of the lightest neutrino mass  $m_l$  ( $m_l = m_1$  for NH and  $m_l = m_3$  for the IH). Placing constraints on the value of  $\langle m_{\beta\beta} \rangle$  below the allowed region of the IH would strongly disfavour this scenario (under the assumption that light Majorana neutrino exchange is the leading cause of  $0\nu\beta\beta$  decay).




**Figure 3.9:** Predicted probability density distribution of  $\langle m_{\beta\beta} \rangle$  as a function of the lightest neutrino mass,  $m_l$ , for the normal hierarchy (left) and inverted hierarchy (right) assuming a flat prior. Solid contours indicate the allowed parameter phase-space, considering the  $3\sigma$  intervals of the neutrino oscillation parameters and assuming QRPA NMEs. Figure from Reference [ABD17].

### 3.4.1.1 Experimental signature of $0\nu\beta\beta$ decay

For the  $2\nu\beta\beta$  decay mode, the energy deposition of the two electrons in a detector presents a continuous spectrum up to the Q-value of the decay due to the unaccounted momentum carried away by the escaped neutrinos. Contrastingly, a  $0\nu\beta\beta$  decay would result in a mono-energetic peak in the summed  $\beta$ -spectrum at  $Q_{\beta\beta}$  since the electrons must carry almost all the energy of the decay, with only a very small fraction going to the recoiling nucleus (see Figure 3.10).

A detector designed to observe the  $0\nu\beta\beta$  decay of a given source will search for an excess rate of events at  $Q_{\beta\beta}$  relatively to the observed  $2\nu\beta\beta$  decay spectrum. This observation is challenging and requires a complete understanding of the backgrounds in the search region, a high abundance of the decaying element to compensate for the rare nature of this process and an excellent energy resolution at the Q-value of the decay to minimize the overlap with the  $2\nu\beta\beta$  decay tail.



**Figure 3.10:** Illustration of the spectra of the summed electron kinetic energies ( $K_e$ ) scaled with the total energy of the decay ( $Q$ ) for the  $2\nu\beta\beta$  decay (dotted line) and  $0\nu\beta\beta$  decay (solid line) convolved with an energy resolution of 5% FWHM. The area of the curves is arbitrarily set for illustration purposes and does not represent the expected relative rates for these two decay modes. Figure from Reference [EV02].

### 3.4.1.2 Twin isotopes and probing the $0\nu\beta\beta$ decay mechanism

The observation of  $0\nu\beta\beta$  decay on twin isotopes, *i.e.*, two isotopes of the same element<sup>12</sup>, capable of decaying via  $2\nu\beta\beta$  decay, would lead to constraints on the contributions of the different new-physics mechanisms that govern this nuclear process, provided that the NMEs are calculated properly [SDF02]. In particular, the ratio of the  $0\nu\beta\beta$  decay half-lives of the two isotopes is sensitive to the  $0\nu\beta\beta$  decay mechanism via the dependencies of the NMEs on the physics model, since it partially washes out the theoretical uncertainties associated with nuclear models. From Equation 3.54, the ratio  $R_{\mathcal{M}^{0\nu}}$  is then given by:

$$R_{\mathcal{M}^{0\nu}}(A_1, A_2) = \frac{|\mathcal{M}^{0\nu, A_2}|}{|\mathcal{M}^{0\nu, A_1}|} = \sqrt{\frac{G^{0\nu, A_1} T_{1/2}^{0\nu, A_1}}{G^{0\nu, A_2} T_{1/2}^{0\nu, A_2}}}, \quad (3.55)$$

where  $A_1$  and  $A_2$  denote the two isotopes. The phase-space factors  $G^{0\nu, A}$  can be calculated with high precision, meaning that measuring the half-lives of the decays would return  $R_{\mathcal{M}^{0\nu}}$  directly. It is worth noticing that this ratio is not sensitive to the new physics parameters themselves, *i.e.*, the lepton number violating parameters of the model such as  $\langle m_{\beta\beta} \rangle$  in the case of light neutrino exchange, as they cancel out, but it is sensitive to the underlying physics mechanism because the NMEs are different for different mechanisms.

Xenon is unique among other elements with more than one isotope capable of decaying via

<sup>12</sup>“Isotopes of the same element” is a redundant statement, but due to abuse of the word “isotope” this clarification seems appropriate.

$0\nu\beta\beta$  since  $^{136}\text{Xe}$  is a closed shell nucleus for neutrons, which leads to the suppression of the decay on one hand but to smaller NMEs and their uncertainties on the other [SPVF99, SDF02]. It also makes xenon particularly sensitive to the decay mechanism amongst other twin isotopes available: Assuming the renormalized quasiparticle random-phase approximation (QRPA) [SSF96] for calculations of the nuclear matrix elements, a value of  $R_{\mathcal{M}^{0\nu}}(^{136}\text{Xe}, ^{134}\text{Xe}) = 2.00$  is expected if the  $0\nu\beta\beta$  decay is mediated by light neutrino exchange, which is  $\sim 35\%$  lower than the one expected for heavy neutrino exchange ( $R_{\mathcal{M}^{0\nu}} = 3.12$ ) and for R-parity breaking  $\mathcal{R}_p\text{SUSY}$  mechanisms ( $R_{\mathcal{M}^{0\nu}} = 3.03$ ) [SDF02]. Comparatively, the values of  $R_{\mathcal{M}^{0\nu}}(^{130}\text{Te}, ^{128}\text{Te})$  for the different decay mechanisms only vary by, at most, 9% under the same assumptions.

### 3.4.1.3 Past, current and future searches for the $0\nu\beta\beta$ decay

Table 3.5 summarizes some of the past, current and future experiments searching for the  $0\nu\beta\beta$  decay process using different isotopes. Chapter 5 will be dedicated to the sensitivity analysis of the LZ detector to the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$ . The high Q-value of the  $2\nu\beta\beta$  decay of the  $^{136}\text{Xe}$  isotope,  $Q_{\beta\beta} = 2457.83 \pm 0.37$  keV [A<sup>+</sup>11a], implies a higher  $0\nu\beta\beta$  decay rate in comparison to other potential decaying isotopes due to the strong dependence of the  $0\nu\beta\beta$  decay half-life on the  $Q_{\beta\beta}$ :  $T_{1/2}^{0\nu} \propto Q_{\beta\beta}^{-5}$  [Saa13]. Furthermore, it represents an experimental advantage as it places the energy search region above many common backgrounds. Currently, the best lower limit on the half-life for  $0\nu\beta\beta$  of  $^{136}\text{Xe}$  comes from the KamLAND-Zen experiment at  $1.07 \times 10^{26}$  years [G<sup>+</sup>16], which is also the current the best limit on the half-life of any  $0\nu\beta\beta$  decay candidate studied thus far and provides the most stringent constraints on  $\langle m_{\beta\beta} \rangle$  at  $< 61 - 165$  meV.

$^{134}\text{Xe}$  is another  $2\nu\beta\beta$  and  $0\nu\beta\beta$  decay candidate of particular interest to LZ, with a relatively high abundance of over 10% in natural xenon. However, the Q-value for these decays is only  $825.8 \pm 0.9$  keV [BTZ14], meaning that both the unobserved  $2\nu\beta\beta$  and  $0\nu\beta\beta$  decay signals are expected to be heavily shadowed by backgrounds (including the  $2\nu\beta\beta$  decay of  $^{136}\text{Xe}$ ) and suppressed given the Q-value dependence of both  $2\nu\beta\beta$  and  $0\nu\beta\beta$  decays. Preliminary studies of the sensitivity of LZ to both two-neutrino and neutrinoless double beta decay modes of  $^{134}\text{Xe}$  are already in progress and show that LZ will have an unprecedented sensitivity to these decays [LZ 20].

## 3.4.2 The importance of $0\nu\beta\beta$ decay searches

The observation of  $0\nu\beta\beta$  decay would have a significant impact across several fields of physics, from particle physics to cosmology. Such observation would directly imply violation of leptonic number conservation ( $\Delta L = 2$ ) as well as  $B - L$  violation ( $\Delta B = 0$ ;  $\Delta L = 2$ ), meaning that these are not fundamental symmetries of nature. Furthermore it would be the first evidence of the existence of fundamental Majorana particles [SV82], which would further support some supersymmetric dark matter candidates such as neutralinos since these are Majorana fermions in some supersymmetric extensions of the Standard Model [JKG96]. It could also help solve some cosmological mysteries such as the matter-antimatter asymmetry observed in the Universe

and the role that leptons played in it (see Section 3.3.2). Finally, the study of the  $0\nu\beta\beta$  decay is complementary to that of neutrino oscillation experiments for the study of neutrino properties, probing their mass scale and mass hierarchy [DMVV16, DPR19].

**Table 3.5:** Experiments searching for the  $0\nu\beta\beta$  decay process in different isotopes, listing the total available (active) isotopic mass and the sensitivity limits on the half-life of the  $0\nu\beta\beta$  decay. Future experiments under commissioning or R&D stages have the expected sensitivity values marked as “projected” [Det20].

Isotope	$Q_{\beta\beta}$ [MeV]	Experiment	Isotope mass	$T_{1/2}^{0\nu}$ sensitivity [yr] (90% CL)
$^{48}\text{Ca}$	4.263	CANDLES-III [I <sup>+</sup> 16]	0.3 kg	$> 0.8 \times 10^{22}$
		CANDLES-IV [I <sup>+</sup> 16]	6.4 kg	-
$^{76}\text{Ge}$	2.039	GERDA [A <sup>+</sup> 18c]	44 kg	$> 8.0 \times 10^{25}$
		Majorana D. [LC <sup>+</sup> 19]	27.3 kg	$> 2.7 \times 10^{25}$
		LEGEND 200 [LCG19]	200 kg	$\sim 10^{27}$ (projected)
		LEGEND 1000 [LCG19]	1 t	$\sim 10^{28}$ (projected)
$^{82}\text{Se}$	2.998	SuperNEMO D. [Mac17]	7 kg	$> 5.9 \times 10^{24}$ (projected)
$^{96}\text{Zr}$	3.348	ZICOS [F <sup>+</sup> 20]	45 kg	$\geq 10^{27}$ (projected)
$^{100}\text{Mo}$	3.035	AMORE-I [Lee20]	6 kg	$\sim 10^{25}$ (projected)
		AMORE-II [Lee20]	100 kg	-
		CUPID-Mo [A <sup>+</sup> 20m]	2.264 kg	$2.43 \times 10^{24}$ (projected)
$^{130}\text{Te}$	2.527	CUORE [A <sup>+</sup> 18h]	206 kg	$> 1.5 \times 10^{25}$
		SNO+ [Fis18]	1.3 t	$1.9 \times 10^{26}$ (projected)
		Theia-Te [A <sup>+</sup> 20o]	31 t	$1.1 \times 10^{28}$ (projected)
$^{136}\text{Xe}$	2.458	KamLAND-Zen 400 [G <sup>+</sup> 16]	346 kg	$> 1.07 \times 10^{26}$
		KamLAND-Zen 800 [Gan20]	683 kg	$8 \times 10^{25}$ (preliminary)
		[Gra20]	683 kg	$5 \times 10^{26}$ (projected)
		KamLAND2-Zen [Gra20]	910 kg	$2 \times 10^{27}$
		EXO-200 [A <sup>+</sup> 19g]	74.7 kg	$> 3.5 \times 10^{25}$
		nEXO [A <sup>+</sup> 18g]	4 t	$> 9.2 \times 10^{27}$ (projected)
		NEXT-WHITE [HN19]	5 kg	-
		NEXT-100 [MA <sup>+</sup> 16]	91 kg	$> 6.0 \times 10^{25}$ (projected)
		PandaX-III [C <sup>+</sup> 17a]	$\sim 900$ kg	$> 1.0 \times 10^{27}$ (projected)
		DARWIN [A <sup>+</sup> 20a]	3.6 t	$> 2.4 \times 10^{27}$ (projected)
LZ [A <sup>+</sup> 20c]	623 t	This work		
Theia-Xe [A <sup>+</sup> 20o]	$\sim 45$ t	$2.0 \times 10^{28}$ (projected)		

# Chapter 4

## The LUX-ZEPLIN experiment

The LUX-ZEPLIN (LZ) experiment is a second-generation dark matter direct detection experiment born of the merging of two previous collaborations, LUX and ZEPLIN, both with a long history of scientific accomplishments and technological research and development for dark matter searches. The LZ detector is a 10 tonne dual-phase xenon time projection chamber (TPC) that measures scintillation light and ionization charge from particle interactions in the liquid xenon target volume [A<sup>+</sup>20g]. It is currently in the final stages of installation at the Davis Campus in the Sanford Underground Research Facility (SURF) in Lead, South Dakota, 1478 meters below the surface in the Homestake gold mine. It will occupy the same space where the LUX detector operated from late 2012 to 2016, during which it maintained a world-leading sensitivity to the spin-independent (SI) WIMP-nucleon interaction cross section [A<sup>+</sup>14a, A<sup>+</sup>16b, A<sup>+</sup>17h].


LUX reached a world-leading sensitivity of  $1.1 \times 10^{-46} \text{ cm}^2$  (0.11 zb) for the SI WIMP-nucleon scattering cross section for a  $50 \text{ GeV c}^{-2}$  WIMP [A<sup>+</sup>17h]. That limit has since been surpassed by XENON1T, that has reached a limit of  $4.1 \times 10^{-47} \text{ cm}^2$  (41 yb) for a  $30 \text{ GeV c}^{-2}$  WIMP [A<sup>+</sup>18i], and by PandaX-II, that has reached a limit of  $8.6 \times 10^{-47} \text{ cm}^2$  (86 yb) for a  $40 \text{ GeV c}^{-2}$  WIMP [C<sup>+</sup>17b]. The projected sensitivity of LZ to the spin-independent WIMP-nucleon cross section is  $1.4 \times 10^{-48} \text{ cm}^2$  (1.4 yb) at 90% confidence level for a  $40 \text{ GeV c}^{-2}$  WIMP [A<sup>+</sup>20d], which represents an improvement over the limit set by LUX by two orders of magnitude and over the current best limit from XENON1T by a factor of 30. This sensitivity is mainly the combined result of the increased target mass of LZ, the deployment of a gadolinium-loaded liquid scintillator veto around the detector, an instrumented liquid xenon skin veto to improve background rejection, the procurement of ultra-pure construction materials and the optimization of the light collection [M<sup>+</sup>17, A<sup>+</sup>20g].

As a requirement for the main scientific goal of detecting dark matter, LZ will feature an ultra-low background environment in the innermost regions of its liquid xenon target. This environment

will be sensitive to several rare physics phenomena beyond dark matter interactions, such as coherent elastic neutrino-nucleus scattering (CE $\nu$ NS), several other neutrino interactions, two-neutrino and neutrinoless double beta decay of several xenon isotopes and interactions of exotic particles like axions and axion-like particles (ALPs).

The exotic decays of xenon isotopes are a particularly important scientific topic in LZ. Despite having the main goal of searching for WIMP dark matter, LZ is expected to be competitive with dedicated rare decay experiments. It will be able to measure the half-lives of some observed decay processes with great precision, such as the two-neutrino double beta decay of  $^{136}\text{Xe}$  and the two-neutrino double electron capture in  $^{124}\text{Xe}$ , the latter recently observed by XENON1T [A<sup>+</sup>19i]. The projected sensitivity of LZ to the unobserved two-neutrino double beta decay of  $^{134}\text{Xe}$ , along with the corresponding neutrinoless mode, are the best in literature [LZ 20]. Finally, the projected sensitivity of LZ to the unobserved  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  is competitive with the current best limit imposed by KamLAND-Zen [G<sup>+</sup>16]. The search for this neutrinoless mode of the double beta decay of  $^{136}\text{Xe}$  is the main secondary scientific goal of LZ, and the sensitivity projections will be presented in detail in Chapter 5.

The LZ detector is schematically represented in Figure 4.1 (a full description can be found in Reference [A<sup>+</sup>20g]). The liquid xenon TPC occupies the center of the detector, inside a vacuum-sealed, double walled titanium cryostat vessel. The cryostat is surrounded by a gadolinium-loaded liquid scintillator detector (dubbed outer detector, or OD) composed of nine acrylic tanks (green and blue volumes in the schematic) which provides a near  $4\pi$  coverage of the detector. The whole detector system is suspended in the middle of the same water tank that housed the LUX detector, immersed in ultra-pure water for additional neutron and  $\gamma$ -ray shielding. An array of photomultiplier tubes installed inside this water tank is used to detect signals from the outer detector and Cherenkov light in the water itself. Table 4.1 summarises the dimensions and masses of some key elements of the LZ detector.



**Figure 4.1:** Schematic representation of the LZ detector. The active xenon TPC in the center (1) is surrounded by the acrylic tanks of the outer detector (2) that are observed by a cylindrical array of 120 8-inch PMTs (3). The full detector is within a water tank (4). The high-voltage umbilical (5) and the neutron source tube (6) protrude the instrumented volume of the outer detector from the sides. The main liquid xenon circulation conduit that connects the xenon purification system to the TPC is adjacent to the water tank (7). A detailed view of the innermost region of the detector is presented on the inlay in the right [A<sup>+</sup>20d, A<sup>+</sup>20g].

**Table 4.1:** Dimensions and masses of some relevant elements of the LZ detector. Table adapted from Reference [A<sup>+</sup>20d].

Parameter [units]	Value
TPC active height [m]	1.46
TPC inner diameter [m]	1.46
Active liquid xenon mass [kg]	7000
Skin detector thickness (sides) [cm]	4–8
Inner cryostat diameter [m]	1.58–1.66
Inner cryostat height [m]	2.59
Outer cryostat inner diameter [m]	1.83
Outer cryostat height [m]	3.04
GdLS acrylic tanks outer radius [m]	1.64
GdLS thickness (sides) [cm]	61
GdLS thickness (top) [cm]	40–62
GdLS thickness (bottom) [cm]	34.5–57
GdLS mass [tonne]	17.3
GdLS acrylic tanks wall thickness [cm]	2.54
Water thickness (GdLS tanks to PMTs) [cm]	84
Water tank diameter [m]	7.62
Water tank height [m]	5.92
Water mass [tonne]	228

## 4.1 Dual-phase Xenon TPC

The main detector of LZ is a dual-phase (liquid/gas) xenon TPC designed to record the energy and position information of particles that deposit energy onto the 7 tonne active liquid xenon target. This technology was pioneered for dark matter searches by the ZEPLIN [A<sup>+</sup>05, A<sup>+</sup>07b, A<sup>+</sup>07a] and XENON [A<sup>+</sup>11b, A<sup>+</sup>08] collaborations in the early 2000's, and since has been proven to be the most sensitive method for dark matter direct detection for 10-1000 GeV/c<sup>2</sup> WIMP masses [A<sup>+</sup>17h, A<sup>+</sup>18i, W<sup>+</sup>20]. The success of dual-phase xenon TPCs is mostly due to the low background, low energy threshold ( $\sim$ keV) and energy and position resolution that can be achieved with this technology.

The dual-phase xenon TPC of LZ consists of a liquid xenon target volume in equilibrium with a gaseous xenon layer on top, both observed by two arrays of light collection sensors placed at the bottom of the liquid phase and on top of the gas phase. An electric field is generated across the liquid phase for drifting charges from the interaction site towards the gas phase, where a stronger electric field generated across the liquid-gas boundary is used to extract the charges into the gas phase and produce proportional electroluminescence light. These fields are generated by high voltage electrode grids woven from thin metal wire in a mesh. This grid design ensures the maximum transparency to the scintillation light of xenon while still maintaining field uniformity. The TPC nests inside a double-walled, vacuum-insulated titanium cryostat vessel that provides structural stability and thermal insulation. The total mass of xenon within the cryostat is 10 tonnes. An ultra-pure titanium batch with a remarkably low radioactive content was selected for manufacturing the cryostat and several other high-mass titanium parts [A<sup>+</sup>20h], to ensure that the background contribution from these components is minimized. This is especially important since the cryostat is adjacent to all sensitive volumes and has a large mass of 2.6 tonnes.

The top and bottom light sensor arrays use 3-inch Hamamatsu R11410-22 photomultiplier tubes (PMTs). These PMT units were specifically developed for working in cryogenic temperatures and were designed to have very low intrinsic radioactivity for use in low background environments. The average quantum efficiency (QE) at cryogenic temperatures for the 178 nm xenon scintillation light (VUV photons) is 30.9% after accounting for the dual photoelectron emission effect [A<sup>+</sup>20g]. The downward-looking top array is installed in the gas xenon volume, containing 253 PMTs disposed in a hybrid configuration with an inner hexagonal matrix that gradually transitions to circular near the edges. The outer rows of the top PMT array optimize the position reconstruction of wall events – the probability of misreconstructing the radial position of simulated wall events beyond 4 cm from the TPC walls, using the Mercury algorithm [S<sup>+</sup>11, A<sup>+</sup>18e], is estimated to be less than  $10^{-6}$ , even for the smallest S2 signals considered (corresponding to 5 extracted electrons) [A<sup>+</sup>20d]. The upward-looking bottom array sits at the bottom of the liquid xenon volume and contains 241 PMTs in a close-packed hexagonal matrix [A<sup>+</sup>20g]. Most of the primary scintillation light produced in the liquid xenon is detected in the bottom array because of the total internal reflection in the liquid-gas interface. The layout of the bottom PMT array maximizes photocathode coverage in order to increase the light collection efficiency of small light signals from low-energy interactions.



There are four grids that generate the drift and extraction fields in the TPC (from top to bottom): anode, gate, cathode and bottom grid. Their properties are summarized in Table 4.2. The drift field is produced between the cathode and gate grids and has a strength of  $310 \text{ V cm}^{-1}$ . The drift region of the TPC measures 145.6 cm from the cathode to the gate grid and has an inner diameter of 145.6 cm. This is the main target of LZ and is also called the active volume. The uniformity of the drift field along the depth of the TPC is maintained by a set of titanium field shaping rings embedded in the field cage walls that surround the TPC. The extraction field produced across the 13 mm gap between the gate and anode grid, also called the extraction region, will reach  $10.1 \text{ kV cm}^{-1}$  in the xenon gas and approximately half this value in the liquid [M<sup>+</sup>17]. Due to the electrostatic force between the grids, the extraction region is expected to contract  $\sim 1.6 \text{ mm}$  in the center region of the TPC, leading the extraction field to reach  $11.5 \text{ kV cm}^{-1}$  at the center of the anode grid [A<sup>+</sup>20g]. A bottom grid is located 13.75 cm below the cathode grid to shield the bottom PMT array from the cathode high fields. The volume of liquid xenon between the cathode and bottom grids is referred to as the “reverse field region” [M<sup>+</sup>17]. In this region the electric field drifts any ionization electrons towards the bottom PMT grid, meaning that any interaction event in the reverse field region will produce a scintillation signal (S1) but not the electroluminescence signal (S2). These events often leave a concentrated signal of S1 light in just a few PMTs in the bottom array, due to the proximity of the bottom PMTs to the reverse field region.

**Table 4.2:** Physical properties of the four grid electrodes of the TPC of LZ: working voltage, diameter of each 304 stainless steel ultra-finish wire, pitch of the grid mesh and number of wires in the mesh [A<sup>+</sup>20g].

Electrode	Voltage [kV]	Diameter [ $\mu\text{m}$ ]	Pitch [mm]	Number
Anode	+5.75	100	2.5	1169
Gate	-5.75	75	5.0	583
Cathode	-50.0	100	5.0	579
Bottom	-1.50	75	5.0	565

The TPC is delimited on the sides by the field cage, a highly reflective wall that improves light collection significantly and provides structural stability to the field-shaping rings. The field cage is constructed by 58 polytetrafluoroethylene (PTFE) reflector panels stacked vertically. Each PTFE reflector is 25 mm tall and contains a 21 mm thick titanium field-shaping ring embedded within [A<sup>+</sup>20g]. The entire inner surface of the TPC, with the exception of the PMT windows, is covered in PTFE to maximize light collection. PTFE immersed in liquid xenon has a high reflectance for VUV photons, measured at 0.973 ( $>0.971$  at 95% C.L.) for the scintillation light of xenon [NLM<sup>+</sup>17]. The empty areas between PMTs in the top and bottom arrays are also covered with PTFE “trifoils” for maximum surface coverage. The PTFE of LZ has extremely low levels of radioactive content ( $< 0.043 \text{ mBq/kg}$  for  $^{238}\text{U}$  and  $0.013 \text{ mBq/kg}$  for  $^{232}\text{Th}$ ) [A<sup>+</sup>20h], which is particularly important since the PTFE walls are adjacent to the active volume of the TPC.

The liquid xenon in the TPC circulates through the liquid xenon tower outside of the water tank, where it undergoes several purification steps designed to remove krypton, argon, radon

and electronegative impurities before being returned to the detector [A<sup>+</sup>20d, A<sup>+</sup>20g]. The xenon circulation and purifications strategies implemented in LZ are similar to those developed for the LUX detector [A<sup>+</sup>13c, A<sup>+</sup>18f].


The LZ experiment has the most strict radioactive assay policy and highest requirement standards for cleanliness in the field. An aggressive assay strategy was implemented to ensure that the detector was built using the most radio-pure materials available. All materials used in the LZ detector have been screened for radioactivity [A<sup>+</sup>20h]. Most materials were screened directly by the LZ collaboration in an extensive radioassay campaign over the last six years, with only a small percentage being based on results from previous experiments. The goal of the radioassay campaign is to inform material selection for the construction of the detector systems and to provide the necessary information to build a detailed background model. Material assays will continue with the goal of reducing uncertainties on several components.

#### 4.1.1 TPC Operation Principle

Figure 4.2 depicts the generation of the measured signals produced by an interaction in the TPC of LZ. When a particle interacts in the active xenon volume it excites and ionizes xenon atoms at the interaction site, producing both scintillation light and ionization charge [CA13]. The scintillation light is released by excited xenon dimers (excitons) with a characteristic wavelength of 178 nm and FWHM of 14 nm [M<sup>+</sup>17] and is promptly detected by the top and bottom PMT arrays as the S1 signal. As mentioned above, most of the prompt scintillation light produced in the active volume will be collected by the bottom PMT array due to total internal reflections in the liquid-gas interface.

If no electric field is applied to the active xenon volume the free electrons completely recombine with the available xenon atoms at the interaction site, leading to more scintillation light. However, in the presence of a drift field the electrons that do not immediately recombine are removed from the interaction site and drift towards the liquid-gas interface [A<sup>+</sup>20g]. The lifetime of these free electrons will be dictated by the level of electro-negative impurities in the liquid xenon. In the absence of these impurities, all ionization electrons will reach the extraction region where the high-field extracts them into the gas phase, producing a stronger electroluminescence light signal proportional to the number of extracted electrons. This secondary signal is recorded mostly in the top PMT array as the S2 signal [DLR70, B<sup>+</sup>93, CA13].

The relative intensity of the S2 light in the PMTs of the top array is used to reconstruct the position of the initial interaction in the horizontal plane using light response functions (LRFs) of the PMTs [A<sup>+</sup>18e]. The S2 light pattern in the top PMT array is strongly correlated to the horizontal ( $xy$ ) position of the electron cloud that was extracted into the gas phase, which in turn is correlated to the original  $xy$  position of the interaction. The average horizontal position of the drifted electron cloud will follow the drift field lines. Even if drift field uniformity is not reached or degrades with time, it is possible to calibrate and correct the  $xy$  position using an calibration source mixed uniformly in the liquid xenon volume [A<sup>+</sup>17d, A<sup>+</sup>16d].



**Figure 4.2:** Schematic representation of an interaction in the liquid xenon volume of a dual-phase xenon TPC. The energy deposited by a particle in the liquid xenon volume excites and ionizes the xenon medium, releasing scintillation photons and ionization electrons. The scintillation is promptly detected in the top and bottom PMT arrays as the S1 signal, while the electrons are drifted from the interaction site by an applied electric field to the liquid-gas interface, where a stronger electric field (extraction field) produced between the gate and anode grids extracts them into the gas phase, producing the strong electroluminescence light corresponding to the S2 signal. The time difference between the S1 and S2 signals is used to determine the depth of the interaction, while the S2 light pattern in the top array can be used to reconstruct the XY position of the interaction [M<sup>+</sup>17, A<sup>+</sup>20g].

The depth ( $z$ ) of the interaction is obtained using the time difference between the S2 signal and the S1 signal, that is proportional to the free electron drift distance. The  $310 \text{ V cm}^{-1}$  uniform field in the active region will drift free electrons at a constant speed of  $1.8 \text{ mm } \mu\text{s}^{-1}$ . The electron cloud will be smeared along the drift path, in both the longitudinal and transverse directions, until it reaches the liquid surface, resulting in wider S2 signals for deeper interactions. This S2 smearing will degrade the timing information for interactions at the bottom of the TPC. Fortunately, this is a well understood Gaussian smearing that can be corrected. Combining the depth ( $z$ ) information of the interaction with the horizontal ( $xy$ ) position reconstruction obtained with the S2 light pattern results in a full 3D event reconstruction.

#### 4.1.2 Energy Reconstruction and Light and Charge Yields

The energy deposited in the active volume of the TPC by a given particle,  $E$ , will produce a certain number of excitons ( $N_{ex}$ ) and some number of electron-ion pairs ( $N_{ion}$ ) plus some heat. The energy deposited in a single interaction can be written as a function of the total number of

quanta (excitons and ions) produced:

$$\begin{aligned} E &= L(N_{ex}W_{ex} + N_{ion}W_{ion}) \\ &= LN_{ion}(\alpha W_{ex} + W_{ion}), \end{aligned} \quad (4.1)$$

where  $W_{ex}$  and  $W_{ion}$  are the average energies required to produce an excitation or ionization, respectively,  $\alpha \equiv N_{ex}/N_{ion}$  is a dimensionless constant that does not depend on the energy of the interaction [SBK<sup>+</sup>11], and  $L \in [0, 1]$  is a “quenching” factor that represents the fraction of energy lost to heat. A convenient simplification can be made by combining the work functions  $W_{ex}$  and  $W_{ion}$  into a mean work function  $W = (\alpha W_{ex} + W_{ion}) / (1 + \alpha)$ , so that the deposited energy becomes:

$$E = LW(N_{ex} + N_{ion}). \quad (4.2)$$

The averaged work function  $W$  of liquid xenon is measured at  $W = 13.7$  eV [Pla61]. A fraction  $\eta$  of initial electron-ion pairs will recombine at the interaction site, resulting in more scintillation quanta. This fraction is dependent of the strength of the drift field. The final number of scintillation quanta,  $n_{scint}$ , is equal to the sum of initial scintillation produced by the excitons and the number of recombined electron-ion pairs. Conversely, the final number of charge quanta,  $n_{ion}$ , will be reduced by the same fraction of ionization electrons that recombined:

$$n_{scint} = N_{ex} + \eta N_{ion} \quad (4.3)$$

$$n_{ion} = N_{ion}(1 - \eta). \quad (4.4)$$

The number of scintillation photons detected as the S1 signal depends on the photon detection efficiency (PDE),  $g1$ , defined as the product of the detector light collection efficiency and PMT quantum efficiency. The number of photons detected as the S2 signal depends on the electron extraction efficiency,  $\epsilon_{S2}$ , the average number of electroluminescence photons produced per extracted electron,  $G_{SE}$ , and the PDE in the gas phase,  $g1_{gas}$ .

$$S1 [phd] = n_{scint} g1 \quad (4.5)$$

$$S2 [phd] = n_{ion} \epsilon_{S2} G_{SE} g1_{gas} = n_{ion} g2, \quad (4.6)$$

where  $g2 = \epsilon_{S2} G_{SE} g1_{gas}$  is the effective charge gain [A<sup>+</sup>20d].

Since the total number of initial and final quanta remains the same due to recombination, *i.e.*,  $N_{ex} + N_{ion} = n_{scint} + n_{ion}$ , Equation 4.2 can be rewritten as

$$E = LW \left( \frac{S1}{g1} + \frac{S2}{g2} \right), \quad (4.7)$$

where Equations 4.5 and 4.6 were used. This expression relates the deposited energy with the measured S1 and S2 signals and can be used to reconstruct the energy of the interaction. The gain factors  $g1$ ,  $g2$  are detector-dependent and are determined directly from data using calibration sources of known energies.

The fraction of deposited energy that is transferred into the target as heat depends on the nature of the interaction. A scattering with an atomic electron will produce a negligible amount of heat ( $L \approx 1$ ), but the same is not true for an interaction with a xenon nucleus. In this case, the energy lost to atomic motion rather than the detectable electronic channels (exciton or ion) effectively “quenches” the deposited energy [SD11, LKM<sup>+</sup>15]. The quenching factor  $L$  is given by the Lindhard model [LS61, LNST63, LSS63]

$$L = \frac{kg(\epsilon)}{1 + kg(\epsilon)}, \quad (4.8)$$

where  $k = 0.166$  is a constant that represents the proportionality between the electronic stopping power and the velocity of the recoiling nucleus, and  $g(\epsilon)$  is an empirical function that is proportional to the ratio of electronic stopping power to nuclear stopping power. The dimensionless quantity  $\epsilon$  relates to the deposited energy  $E$  by [SD11, LKM<sup>+</sup>15]

$$\epsilon = 11.5 \frac{E}{[1 \text{ keV}]} Z^{7/3}, \quad (4.9)$$

where  $Z$  is the atomic number of the nucleus.  $g(\epsilon)$  can be parametrized as [LKM<sup>+</sup>15]

$$g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon. \quad (4.10)$$

The Lindhard model seems to agree well with experimental data from nuclear recoils down to 1.1 keV, but seems to diverge from measurements at energies below that. For higher energies, biexcitonic effects may lead to loss of light yield not accounted for in the Lindhard model. In order to improve the agreement between calibration data and the predictions of the model for nuclear recoils in both energy regimes, some modifications of the basic model are considered in LZ following previous work performed for LUX [A<sup>+</sup>18d, A<sup>+</sup>16c].

Table 4.3 summarizes some key LZ detector parameters. The large light and charge yields obtained when a particle deposits energy in liquid xenon result in a good sensitivity to low energy interactions [M<sup>+</sup>17]. The excellent light and charge collection efficiencies of LZ result in a low energy threshold down to 6 keV for nuclear recoils and 1.5 keV for electron recoils, and sub-keV threshold for S2-only analyses [A<sup>+</sup>20d]. A very low energy threshold is a requirement for dark matter searches, as the recoil energy from a dark matter interaction is expected to be of the order of a few keV (see Section 2.6.2).

**Table 4.3:** Projected detector performance parameters for the TPC of LZ. Indented entries indicate a dependency on preceding parameters (see text for details). Table from [A<sup>+</sup>20d].

Detector Parameter	Value
PTFE reflectivity in liquid	0.977
PTFE reflectivity in gas	0.85
Photon absorption length in liquid [m]	100
Photon absorption length in gas [m]	500
PMT efficiency* at 175 nm	0.27
Average PDE in liquid ( $g_1$ ) [phd/scint.]	0.12
Average PDE in gas ( $g_{1_{\text{gas}}}$ ) [phd/scint.]	0.10
Average light yield per extracted electron ( $G_{SE}$ ) [scint./e]	832
S2 electron extraction efficiency ( $\epsilon_{S2}$ )	0.95
Effective charge gain ( $g_2$ ) [phd/e]	79
Single electron size [phd]	83
Electron lifetime [ $\mu\text{s}$ ]	850

\* Including first dynode collection efficiency.

A large number of calibration sources will be used in LZ in order to measure the performance parameters of the detector at different energy scales and for different interaction channels [A<sup>+</sup>20g]. There are three main deployment strategies for calibration sources:

- Internal sources – low and high-energy sources mixed within the liquid xenon.
- External sources – high-energy  $\gamma$ -ray and photoneutron ( $\gamma, n$ ) sources deployed outside the cryostat through dedicated conduits.
- DD generator neutron source – produces 2450 keV monoenergetic neutrons that travel through the water tank and outer detector via two dedicated conduits (see Figure 4.1).

The internal sources are low-energy  $\gamma$  and  $\beta$  emitters such as  $^{83m}\text{Kr}$ ,  $^{131m}\text{Xe}$ ,  $^3\text{H}$  (in the form of tritiated methane) or  $^{14}\text{C}$ , with the exception of  $^{220}\text{Rn}$  whose progeny can emit high-energy  $\alpha$ ,  $\beta$  and  $\gamma$ -rays (see Appendix A). Since these sources are not limited by shielding from passive materials, they can be deployed for low-energy calibrations. Furthermore, by being mixed within the xenon they can be used to map the active TPC volume and characterize spacial-dependent detector parameters. However, all internal sources must be either short-lived or easily removed by the purification system.

External  $\gamma$ -ray sources such as  $^{22}\text{Na}$ ,  $^{57,60}\text{Co}$ ,  $^{54}\text{Mn}$  or  $^{228}\text{Th}$ , and some neutron sources like AmLi, AmBe and  $^{252}\text{Cf}$ , are deployed outside the cryostat through dedicated small-diameter conduits. Photoneutron ( $\gamma, n$ ) sources like  $^{88}\text{YBe}$ ,  $^{124}\text{SbBe}$  and  $^{205,206}\text{BiBe}$  are deployed through a large-diameter conduit above the cryostat and require dense shielding to limit the  $\gamma$ -ray flux.

The DD generator neutron source, which has already been used by LUX [A<sup>+</sup>16c], is used to perform a calibration of the low-energy nuclear recoil (NR) response of the detector. The scatter


kinematics of the monoenergetic neutrons can be used to characterize the light and charge yields of NR interactions for a continuous range of energies with high precision.

More details about the calibration strategy and specific sources can be found in Reference [A<sup>+</sup>20g].

### 4.1.3 Discrimination of Electron and Nuclear Recoils

In liquid xenon TPC, a particle can either interact with an atomic electron or with a xenon nucleus, producing an electron recoil (ER) or nuclear recoil (NR), respectively. These two types of interaction have differing relative yields of scintillation and ionization. A NR interaction will produce, on average, less initial ionization and more direct scintillation light than an equivalent ER interaction, leading to a lower average charge/light ratio [A<sup>+</sup>20e].

Figure 4.3 shows the ER and NR bands obtained by LUX using calibration data from tritiated methane injections and monoenergetic neutrons, respectively. A clear offset between the two distributions is observed, indicating that by combining light and charge information it is possible to discriminate between ER and NR events, with rejection efficiencies typically above 99.5% for ER events under 100 keV below the NR median [M<sup>+</sup>17, A<sup>+</sup>20d, A<sup>+</sup>20g] which is one of the most important characteristics of xenon (and xenon TPCs in particular) for WIMP search. This ER discrimination efficiency improves significantly at higher energies.



**Figure 4.3:** Electron recoil (top) and nuclear recoil (bottom) bands from LUX calibration data. The ER band was obtained from tritium calibration data, while the NR band was measured using a monoenergetic neutron source. The solid lines indicate the band medians, while the dashed lines represent the band contours that contain 80% of the events of the NR (red) and ER (blue) distributions [A<sup>+</sup>16d].

The recombination of ionization electrons with free xenon ions at the interaction site strongly affects the variance of the charge/light ratio, leading to a broadening of the ER and NR bands.

Therefore, ER discrimination improves with stronger drift fields that can displace ionization electrons away from the interaction site more effectively [A<sup>+</sup>20e].

## 4.2 The Liquid Xenon Skin System

The gap between the field cage and the inner vessel of the cryostat will be filled with liquid xenon acting as a high-voltage standoff [A<sup>+</sup>20g]. This xenon volume surrounds the entire liquid phase of the TPC. A large bulk of liquid xenon will also fill the dome space below the bottom PMT array. LZ will take advantage of these otherwise unused liquid xenon volumes to create an anti-coincidence single-phase liquid xenon detector referred to as the “skin”. As no electric field is applied to the skin volumes, only scintillation will be measured from energy depositions in these regions [M<sup>+</sup>17]. The skin detector will be able to tag interactions that may result in background events in the active volume of the TPC, working as an anti-coincidence veto. Due to the high density of liquid xenon, this detector will be especially effective at tagging  $\gamma$ -rays.

The position of the skin volumes can be seen in Figure 4.1 on the right. The side volume of the skin detector has a thickness of 4 cm at the top of the TPC that gradually increases to 8 cm at the bottom of the PTFE field cage due to the pear-shaped inner cryostat vessel. The total xenon mass in the full skin system is around 2 tonnes. A total of 131 PMTs will observe the skin detector: The side skin detector will be viewed from above by 93 1-inch Hamamatsu R8520-406 PMTs located at the top of the PTFE field cage and by 20 2-inch Hamamatsu R8778 PMTs looking upward from the bottom region. The dome skin volume below the TPC will be viewed by 18 2-inch Hamamatsu R8778 PMTs mounted horizontally below the bottom TPC PMT array, with 12 looking towards the outer regions of the dome and 6 looking inward to the center of the volume. Figure 4.4 on the left shows the position of some of the skin PMTs.


To aid the light collection of the xenon skin, the inner surface of the inner cryostat vessel was covered with PTFE reflectors. Also, all PMTs and support structures in the dome region were dressed in PTFE reflectors. Figure 4.4 on the right shows the PTFE reflectors lining of the inner cryostat vessel. Due to the complex geometry of the skin detector, light collection efficiency is highly dependent on the position of the interaction. PMT coverage and wall reflectivity allows for an energy threshold for observing energy depositions of 100 keV in more than 95% of the skin volume. For WIMP searches, an event in the TPC will be vetoed if at least 3 coincident photons are detected in the skin within a 800  $\mu$ s coincidence window before and after the TPC S1 signal, which corresponds to a skin threshold of 100 keV [A<sup>+</sup>20d].

## 4.3 The Outer Detector System

The cryostat vessel of LZ is covered by 9 acrylic vessels filled with 17.3 tonnes of Gadolinium-loaded Liquid Scintillator<sup>1</sup> (GdLS) that form the outer detector. These acrylic tanks are schematically represented in Figure 4.1 by the green semitransparent vessels surrounding the

<sup>1</sup>Linear alkylbenzene organic scintillator doped with 0.1% by mass of gadolinium.





**Figure 4.4:** Left: Schematic cross section of the dome and side region of the TPC and skin detectors, showing the position of the 2-inch PMTs at the bottom of the side skin volume (1) and on the dome looking outward (2). Right: Photograph of the inside walls of the inner cryostat vessel covered with PTFE reflectors. The ring of 20 2-inch PMTs looking upward to the side skin volume are visible at the bottom of the vessel. Images from Reference [A<sup>+</sup>20g].

cryostat. The setup of the acrylic tanks provides a thickness of GdLS of 61 cm on the side volumes of the outer detector, with the top and bottom sections providing a minimum thickness of 34.5 cm [A<sup>+</sup>20d]. An array of 120 8-inch Hamamatsu R5912 PMTs observe the outer detector from all sides, as well as provide a readout for Cherenkov light caused by muons crossing the water space. The PMT array is suspended in a Tyvek curtain in the water space at a distance of 115 cm from the outer wall of the acrylic tanks. The PMT readout is also schematically represented in Figure 4.1. The exterior surface of the outer cryostat vessel is covered with Tyvek reflectors to aid the light collection from the liquid scintillator veto [A<sup>+</sup>20g].


The average light collection efficiency over the entire outer detector volume is  $\sim 7\%$ , which corresponds to a light detection yield of 130 photoelectrons for a 1 MeV energy deposit in the liquid scintillator [M<sup>+</sup>17]. The energy threshold of the detector can be set as low as needed within the 3-fold PMT coincidence requirement, but the decay of  $^{14}\text{C}$  present in the organic scintillator, with a Q-value of 156 keV, can lead to false coincidences. For this reason, an energy threshold of 200 keV within a 500  $\mu\text{s}$  coincidence window is assumed for the WIMP search analysis, which corresponds to a minimal detection of 10 photoelectrons [A<sup>+</sup>20g]. The rate of events from  $^{14}\text{C}$  and other impurities in the GdLS is 5.9 Hz, but the largest contribution is from  $\gamma$ -rays from the rock walls of the laboratory at 31 Hz. The expected overall rate of events in the outer detector above threshold is  $\sim 51$  Hz [A<sup>+</sup>20g].

The outer detector will work as an active veto with the main goal of tagging neutrons that may produce WIMP-like signals in the TPC. A veto efficiency above 96.5% for neutrons that scatter once in the active volume is expected considering a coincidence window of 500  $\mu\text{s}$  before and after a nuclear recoil is recorded in the TPC. Gadolinium has a large neutron capture cross section and produces a distinct energy signature that can be easily identified: a neutron capture in  $^{155}\text{Gd}$  or  $^{157}\text{Gd}$  will produce a cascade of several  $\gamma$ -rays with total energy released of 7.9 MeV and 8.5 MeV, respectively. About 90% of thermal neutrons will be captured by  $^{155}\text{Gd}$  or  $^{157}\text{Gd}$ , with the remaining 10% being captured by hydrogen in the organic scintillator and releasing a characteristic 2.2 MeV  $\gamma$ -ray [M<sup>+</sup>17, A<sup>+</sup>20g]. These  $\gamma$ -rays can either be directly detected in the

liquid scintillator or in any other sensitive volume of the detector (skin or TPC). Despite being primarily an active veto system, the outer detector is a multi-purpose instrument that will be used in dedicated studies. The outer detector will make precise measurements of environmental neutron fluxes,  $\gamma$ -rays from the cavern rock and the cosmogenic muon flux at the laboratory.

## 4.4 Background Modelling and Simulations

Despite the effective background mitigation strategies implemented by LZ, like the usage of ultra-pure construction materials and the implementation of passive and active shielding, some backgrounds will undoubtedly be observed. The scientific output of LZ is linked to the ability to measure and model its backgrounds with high precision. The background model of LZ is built using the most recent radioactive assays of materials and measurements of environmental sources [A<sup>+</sup>20d, A<sup>+</sup>20b, A<sup>+</sup>20h] along with dedicated simulations [A<sup>+</sup>21] of the radioactive sources and detector materials. Figure 4.5 shows the simulated spectra of the main single-scatter ER background sources in LZ, normalized to their expected levels, in the inner 5.6 tonnes of the TPC and after applying the skin and outer detector anti-coincidence requirements.



**Figure 4.5:** Combined simulated spectrum of the main ER backgrounds in the 5.6 tonnes fiducial volume considered in the WIMP search analysis. The different spectra are for single scatter events with no coincident signal in the skin nor the outer detector. The spectra are taken directly from simulation and do not include any detector effects. Figure from Reference [A<sup>+</sup>20d].

The background model accounts for: trace radioactivity of  $^{40}\text{K}$ ,  $^{60}\text{Co}$ ,  $^{232}\text{Th}$  and  $^{238}\text{U}$  in detector materials; surface contamination from radon daughter plate-out (the long-lived  $^{210}\text{Pb}$  and progeny) and dust during the assembly of the TPC system; xenon contaminants from dust and surface radon emanation ( $^{220}\text{Rn}$  and  $^{222}\text{Rn}$ ), internal  $^{39}\text{Ar}$  and  $^{85}\text{Kr}$ ; natural radioactivity of the laboratory rock; radiogenic and muon-induced neutrons; two-neutrino double beta decay of  $^{136}\text{Xe}$ ; neutrino interactions from diffuse supernovae, atmospheric neutrinos and solar neutrinos from the pp,  $^7\text{Be}$ ,  $^{13}\text{N}$ ,  $^{15}\text{O}$ ,  $^8\text{B}$  and hep reactions. An overview of the backgrounds of LZ and

a detailed description of the background model for the WIMP search is presented in Reference [A<sup>+</sup>20d]. The background model for the <sup>136</sup>Xe neutrinoless double beta decay search is presented in Section 5.3.

The skin and outer detector systems, together with the LZ TPC itself, will be used to measure the backgrounds with high precision. The outer detector in particular will provide a detailed description of the external backgrounds from both high-energy  $\gamma$ -rays and external neutrons. The large density of liquid xenon ( $\sim 2.9 \text{ g cm}^{-3}$ ) leads to a self-shielding effect of the active target to external backgrounds. Most of the energy depositions from external radiation sources will occur in the outermost layers of the liquid xenon volume due to its strong  $\gamma$ -ray attenuation, resulting in the much-desired low-background levels at the center of the TPC where the sensitivity to the signal is larger. The outermost regions of the TPC will be used to fit the background model for external sources together with the veto detectors to fully characterize the backgrounds observed. On the other hand, neutrino interactions and the decay of sources that are mixed in the xenon, like <sup>85</sup>Kr, <sup>220</sup>Rn, <sup>222</sup>Rn and the neutrinoless double beta decay of <sup>136</sup>Xe, will be observed as a uniform background in the active volume of the TPC.

Simulations of cosmogenic and radiogenic backgrounds are performed with BACCARAT [A<sup>+</sup>21], a framework based on GEANT4 that evolved from the LUXSim simulation package developed for the LUX experiment [A<sup>+</sup>12a]. BACCARAT contains a detailed description of the full LZ detector geometry, from PMT cables to cryostat screws and bolts. The geometry and material definitions match the engineering drawings.

BACCARAT tracks particles using GEANT4 and characterizes their interactions in the sensitive volumes of the detector. The energy depositions in the liquid xenon volumes obtained from BACCARAT are then converted to scintillation and ionization yields using the Noble Element Simulation Technique (NEST) package [SBK<sup>+</sup>11, SFTT13].

For building background models for the different sensitivity analyses of LZ, the S1 and S2 signals are directly estimated using NEST, which allows for large statistics datasets to be generated [A<sup>+</sup>21]. However, these simulations rely on parametrizations of the detector response using averaged light and charge yields and do not contain time-based information of the S1 and S2 signals.

In order to generate S1 and S2 signals with PMT photon timing information that mimics the real detector response, a more complete simulation chain is used. In this chain, NEST is referenced in BACCARAT in order to calculate the raw scintillation photons and ionization electrons for interactions in the xenon volumes, and that information is then parsed into BACCARAT for tracking [A<sup>+</sup>21]. The scintillation photons produced by interactions in the outer detector are generated using a modified version of the Geant4 G4Scintillation code and are modeled according to the key optical properties of the GdLS. The scintillation photons form both the primary scintillation in xenon and the GdLS are propagated using optical tracking (ray-tracing) in GEANT4 [A<sup>+</sup>21]. Ionization electrons transported in the TPC into the extraction region are converted to a number of electroluminescence photons by NEST, using the extraction field strength and electron extraction efficiency [A<sup>+</sup>20d]. These photons, however, are too numerous

to be propagated efficiently in the same manner as the ones from primary scintillation, and are therefore parametrized using detector-averaged quantities, such as the gains  $g_1$  and  $g_2$ , that are estimated from optical simulations of VUV photons emitted uniformly in the xenon volumes [A<sup>+</sup>21]. The timing of the photons as they arrive at the window volumes of the PMTs is recorded and is then converted to waveforms by the Detector Electronics Response (DER) package. The DER simulates the response of the front-end electronics and considers the filtering and shaping effects of the data acquisition process. The final products of the DER are mock timelines with realistic waveforms of the simulated events that have an identical format to the output of the data acquisition of LZ [A<sup>+</sup>21]. The processing of this simulated data, and of real data once available, is performed by the LZ analysis programme (LZap) described in detail in Chapter 6.

# Chapter 5

## Sensitivity of a $0\nu\beta\beta$ decay search on $^{136}\text{Xe}$ with LZ

The low-radioactivity environment of the LZ detector system required for dark matter searches also enables the search for rare and unobserved physics phenomena, such as the neutrinoless double beta decay ( $0\nu\beta\beta$ ) of certain xenon isotopes, neutrino interactions with atomic nuclei or the interaction of Axions or Axion-like particles (ALPs), that would otherwise be shadowed by background events [M<sup>+</sup>17].

LZ is particularly well positioned to study the neutrinoless double beta decay mode of  $^{136}\text{Xe}$ , taking advantage of the natural abundance of this isotope in natural xenon. While most  $0\nu\beta\beta$  experiments use sources enriched in the isotope of interest, so as to increase the fraction of the relevant isotope and to decrease the passive non-source material, the large active volume of LZ with 7 tonnes of natural xenon yields 623 kg of  $^{136}\text{Xe}$ , a mass comparable to other world-leading  $0\nu\beta\beta$  experiments such as nEXO [A<sup>+</sup>18g], NEXT [MA<sup>+</sup>16], KamLAND-Zen [G<sup>+</sup>16] or PandaX-III [C<sup>+</sup>17a].

A detailed report of the expected LZ sensitivity to this decay is provided in this Chapter, as well as a sensitivity projection for a possible subsequent dedicated exposure using 90% enrichment of  $^{136}\text{Xe}$ . This Chapter follows the work co-authored and published in Reference [A<sup>+</sup>20c] by the author.

### 5.1 Assumptions About Detector Performance

Several assumptions are made about critical performance parameters of the LZ detector, such as energy and position resolution, multiple scatter discrimination, signal saturation and veto

efficiency, that will only be known when LZ starts collecting data. These will be discussed in the following sections.

### 5.1.1 Energy Resolution at $Q_{\beta\beta}$

The ability to distinguish the excess of  $0\nu\beta\beta$  decay events at the  $Q_{\beta\beta}$  from the endpoint  $2\nu\beta\beta$  events and other nearby backgrounds is strongly affected by the energy resolution ( $\sigma/E$ ), which in turn depends on the light collection and electron extraction efficiencies, and on the single electron amplification, *i.e.*, the number of detected photons per electron extracted into the gas region.

As discussed in Chapter 4 (Section 4.1.2), an interaction in the forward field region of the TPC of LZ will produce scintillation photons and ionization electrons, that in turn are responsible for generating the measured S1 and S2 signals, respectively. These light and charge quanta can fluctuate through recombination effects. The energy of an ER interaction can be estimated by combining both the S1 and S2 signals through the relation,

$$E = W \left( \frac{S1}{g1} + \frac{S2}{g2} \right), \quad (5.1)$$

where  $W = 13.7$  eV is the work function of liquid xenon, *i.e.*, the average energy needed to produce a single exciton or electron-ion pair [Pla61], and  $g1$ ,  $g2$  are the detector dependent gain factors determined directly from data using calibration sources of known energies. The baseline assumption for the photon detection efficiency ( $g1$ ) of LZ is 7.5%, averaged over the active volume [M<sup>+</sup>17]. However, more recent estimates suggest that  $g1=11.9\%$ , considering a 97.7% (85%) PTFE reflectivity of scintillation light in the liquid (gas) xenon phase, 100 m (500 m) scintillation absorption length in the liquid (gas), 27% average PMT efficiency<sup>1</sup>, 54% PMT photocathode coverage of the TPC surface area, and the geometric transparency and reflectivity of all grids [A<sup>+</sup>20d, M<sup>+</sup>17]. Most of these improvements come from the measured increase of PTFE reflectivity, which resulted of dedicated measurements [NLM<sup>+</sup>17]. The electron extraction efficiency at the extraction region, for an extraction field of 5.2 kV/cm (10.2 kV/cm) in the liquid (gas), is estimated at 95%, while the single electron amplification is estimated as 85 detected photons per extracted electron, resulting in  $g2=79$  [A<sup>+</sup>20d].

The LZ requirement for energy resolution at  $Q_{\beta\beta}$  is of 2.0% ( $\sigma/E$ ) [M<sup>+</sup>17] but this value is expected to be significantly better. Considering the previously mentioned detector parameters, the most recent NEST model [SBK<sup>+</sup>11, SFTT13] predicts that an energy resolution of 0.88% at  $Q_{\beta\beta}$  is achievable with LZ. More detailed studies including field modeling of the extraction region indicates that the energy resolution can be as low as 0.64% [Bai16], but these projections do not account for known detector effects due to photoionization on the gate grid, which may degrade the resolution [A<sup>+</sup>17c]. XENON1T demonstrated that an energy resolution of 0.79% at  $Q_{\beta\beta}$  is achievable with a tonne-scale dual-phase xenon detector [A<sup>+</sup>20j]. The drift field of XENON1T is 80 V/cm, significantly lower than the design 310 V/cm drift field of LZ [A<sup>+</sup>20d]

<sup>1</sup>Including QE and first dynode collection efficiency.

and the energy resolution is expected to improve with higher drift field as observed by EXO-200 [A<sup>+</sup>20i]. An energy resolution of 1.0% at  $Q_{\beta\beta}$  is assumed for the analysis presented henceforth. This value is expected to be conservative.


### 5.1.2 Position Resolution, Multi-site Interaction Identification and Minimal Vertex Separation

The ability to identify multiple scatter interactions typically associated with backgrounds is crucial for rare event searches. High-energy  $\gamma$ -rays, which are the main source of background in this search, have a high penetrative power and will often scatter more than once in the liquid xenon target, but sometimes these consecutive interactions are close to one another. The sensitivity to a  $0\nu\beta\beta$  decay signal is therefore dependant on the position resolution needed to identify a multiple scatter event. In the dual-phase TPC of LZ, the reconstruction of the radial position is determined using the S2 light pattern in the top PMT array, while the vertical position of the interaction is reconstructed using the timing between the S1 and S2 pulses for a given drift field. Therefore, the position resolution in the  $z$  coordinate is significantly better than the radial resolution, as will be addressed below.

The diffusion of the drifted electrons results in wider S2 signals from interactions at the bottom of the TPC when compared to S2 pulses from shallow events, which affects the time resolution of the S2 signal and therefore the position resolution in  $z$ . This is a known and well-understood Gaussian smearing that can be addressed, for instance, by applying a drift-time dependent S2 width cut or pulse deconvolution. The pulse timing information in LZ is of the order of tens of ns, and with a drift velocity of  $\leq 2$  mm/ $\mu$ s, the corresponding spacial resolution is expected to be sub-millimetric.

The full-width at half-maximum of an S2 signal from a shallow event in LZ will be  $\sim 1$   $\mu$ s, meaning that a separation of multiple interactions separated by 2 mm in  $z$  should be possible. Previous work assumed multiple scatter events could be rejected down to 3 mm separations in  $z$  [B<sup>+</sup>14], and the same is assumed here to accommodate the effect of diffusion. Moreover, the expected track length of the two electrons in a  $0\nu\beta\beta$ -decay in  $^{136}\text{Xe}$  is on average smaller than 3 mm even accounting for the emission of Bremsstrahlung photons, so rejecting events with vertex separation in depth less than 3 mm would result in the rejection of some fraction of signal events and a loss in sensitivity. Figure 5.1 shows an example of simulated ionization tracks in liquid xenon caused by two electrons emitted in opposite directions with half the energy of the  $Q_{\beta\beta}$  of  $^{136}\text{Xe}$ .

Vertex separation in the  $xy$  plane is not used to reject multiple scatter events in this analysis. Preliminary studies show that a vertex separation in the  $xy$  plane of 3 cm, a very conservative value for this type of detector [S<sup>+</sup>11, A<sup>+</sup>18e, Sol20], would improve background rejection by less than 4% in the innermost volume considered on this analysis.



**Figure 5.1:** Simulated ionization yield on liquid xenon of an event with two electrons emitted back-to-back in the  $z$  direction, with the initial energy of each electron equal to half the  $Q_{\beta\beta}$  of  $^{136}\text{Xe}$ . The red arrows represent the initial direction of emission of the two electrons. Credit to Andrey Solovov (M.S. at LIP-Coimbra).

### 5.1.3 Impact of S2 Signal Saturation

The energy range of this  $0\nu\beta\beta$  decay search is far greater than that of the main dark matter search goal of LZ. The expected 2458 keV energy deposition of a  $^{136}\text{Xe}$   $0\nu\beta\beta$  decay signal will lead to some saturation in the S2 signal due to limited dynamic range of the ADC and/or to capacitor depletion in the PMT biasing circuit [Fah14, A<sup>+</sup>20c].

Preliminary simulation results on light and charge yields, combined with the PMT voltage divider design, suggest that the S1 signal produced by a  $^{136}\text{Xe}$   $0\nu\beta\beta$  decay in the forward field region of the TPC will not lead to saturation. Likewise, the S2 signal will not saturate any channels in the bottom array but, depending on the drift time and the resulting level of electron diffusion before S2 production, around 7 to 22 top array channels are expected to saturate the ADCs. Nonlinearity effects caused by capacitor depletion following very large signals, which affects only the top PMTs in the S2 signal, will be corrected during data processing [A<sup>+</sup>20c].

Due to the saturation effects on the top array, using the full S2 signal to reconstruct the energy of the event would yield poor results. Fortunately, a large fraction of the S2 light will be detected on the bottom PMT array where no saturation is expected nor nonlinear behaviour (even at  $0\nu\beta\beta$  decay energies). For an S2 signal, approximately 79 photons detected (phd) per extracted electron are expected, with 27 of those photons being detected in the bottom array, on average [A<sup>+</sup>20d]. This ratio of top-bottom light collection is fairly consistent at higher energies since it is dominated by geometric factors such as the total reflective surface area on the top PMT array. Since events in the  $0\nu\beta\beta$  decay energy region will produce copious amounts of S2 light, using the S2 signal on the bottom PMT array alone for reconstructing the event energy will not



compromise the energy resolution.

The top array is also used for  $xy$  position reconstruction using the S2 light distribution amongst the PMTs. Having a good radial position reconstruction is crucial for rejecting radioactive backgrounds which are higher in rate on the outer regions of the detector. The effects of saturation, if left unchecked, will degrade the reconstruction by underestimating the light collection in the saturated channels, leading to reconstruction errors and potentially biasing the radial position of events near the wall inwards. Preliminary studies, that include the effects of S2 saturation, demonstrate that LZ will achieve an  $xy$  position resolution of 0.5 cm or better for interactions with deposited energies above 1.8 MeV within a radial distance of 68.8 cm from the center of the detector and across the full drift length of the TPC. Near the center of the detector, the resolution will be 0.2 cm or better [A<sup>+</sup>20c].

The effects of ADC saturation on the radial position reconstruction can be managed at the level of the reconstruction algorithms. It has been demonstrated in Mercury [S<sup>+</sup>11, A<sup>+</sup>18e], the position reconstruction algorithm used in LZ (see Section 6.1.1.2 for additional details), that excluding several saturated channels will not degrade the  $xy$  position resolution significantly [A<sup>+</sup>18e]. It can be shown that excluding up to 100 channels in the S2 signal produced by simulated 2.5 MeV energy depositions near the central region of the TPC would still yield an estimated  $xy$  position resolution below 2 mm FWHM with Mercury [Sol20].

#### 5.1.4 Veto Systems

The liquid xenon skin and outer detector systems, described in detail in Chapter 4 (Sections 4.2 and 4.3, respectively), will provide excellent discrimination capabilities for external radioactive backgrounds.

The position of the liquid xenon skin veto detector, adjacent to the TPC and encompassing the liquid xenon active volume and reverse field region entirely, will provide an excellent discrimination of incoming gamma backgrounds from radioactive contamination of detector materials outside the TPC. Any event which deposits more than 100 keV in the liquid xenon skin in coincidence with a candidate  $0\nu\beta\beta$  decay event in the TPC will be vetoed. The coincidence window considered for the liquid xenon skin veto is 800  $\mu$ s long.

The outer detector system will be crucial to mitigate the high-energy  $\gamma$ -rays produced in the rock surrounding the laboratory. Since the event triggering is performed in the TPC, the only factors that limit the energy threshold that can be set for a veto condition in the outer detector are its light collection and the decay rate of  $^{14}\text{C}$  in the liquid scintillator ( $\beta$ -decay with a 156 keV endpoint and an expected rate of  $\sim 5$  Hz [A<sup>+</sup>20g]) which can lead to false positives and thus impact the overall efficiency. The average light collection efficiency over the entire outer detector volume is estimated to be  $\sim 7\%$ , providing an average light yield of about 130 photoelectrons for a 1 MeV energy deposit in the liquid scintillator [M<sup>+</sup>17]. For dark matter searches, the energy threshold of the outer detector is set at 200 keV within a 500  $\mu$ s coincidence window to ensure a veto efficiency above 95% for neutrons that scatter once in the TPC. Since only  $\gamma$ -rays present a

relevant background in this  $0\nu\beta\beta$  analysis the coincidence window of the outer detector can be significantly reduced. The prompt coincidence window for the outer detector veto is set as  $1\ \mu\text{s}$  for this analysis, allowing the reduction of the rate of accidental coincidences due to  $^{14}\text{C}$  decays in coincidence with a potential  $0\nu\beta\beta$  event in the TPC to a negligible level and the reduction of the energy threshold from 200 keV down to 100 keV. Therefore, any energy deposition of more than 100 keV in the outer detector within a  $1\ \mu\text{s}$  coincidence window from a TPC event will be excluded.

## 5.2 Detector Calibration

The measurement of the detector response to known radioactive sources and subsequent calibration in both energy and position is a crucial step. The calibration plan for LZ predicts regularly deployment of both internal and external sources that will be used to calibrate the detector in a wide range of energies.

Given the size of LZ and the effect of xenon self-shielding, external radioactive sources have a very low efficiency to produce interactions in the xenon bulk, especially for low energies. To overcome this difficulty, most of the sources will be injected directly and mixed with the liquid xenon. The choice to inject calibration sources is also motivated by the benefits of having calibration information on the entire active volume of the TPC. All internal calibration sources decay away with short half-life or can be removed relatively easily from the xenon by the purification system, ensuring that the calibration campaigns minimally impact the scientific performance of the detector [ $A^+16\text{d}$ ,  $A^+18\text{d}$ ].

Most of the internal sources are aimed at low-energy calibration and deposit less than 200 keV per interaction, such as  $^3\text{H}$ ,  $^{14}\text{C}$ ,  $^{83\text{m}}\text{Kr}$  and  $^{131\text{m}}\text{Xe}$  [ $A^+20\text{g}$ ].  $^{220}\text{Rn}$  will also be injected and will be used to study the response of the detector at higher energies [ $A^+20\text{g}$ ]. The  $^{220}\text{Rn}$  decay chain<sup>2</sup> produces daughters that are short-lived, with the longest half-life being that of  $^{212}\text{Pb}$  at 10.6 hours. The different daughters also provide a decent range of decay energies. In particular, the energy spectrum of the  $^{208}\text{Tl}$  daughter has discontinuities at 3.2 MeV, 3.5 MeV, and 3.7 MeV created by the coincident  $\beta$  and  $\gamma$  decays (see Appendix A for a complete description of the  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains).

There is no benefit in having a high-energy ( $\sim\text{MeV}$ ) calibration uniformly through the xenon for position calibration. The position dependencies of the detector response can be handled by lower energy internal calibration sources quite robustly and remain valid for any energy regime. However, the calibration at higher energies using  $^{220}\text{Rn}$  can be used to characterize the effects of saturation and its radial dependence. Complementing the internal  $^{220}\text{Rn}$  source, an external  $^{228}\text{Th}$   $\gamma$ -ray source will also be deployed that, at the very least, will help calibrate the outer regions of the detector at the MeV energy scale. Since this is an external source, only  $\gamma$ -rays will reach the TPC, mostly the 2615 keV line from  $^{208}\text{Tl}$ . This  $\gamma$ -ray line is close to the  $0\nu\beta\beta$  decay energy and will be used to estimate the energy resolution close to  $Q_{\beta\beta}$ . Furthermore, this line

---

<sup>2</sup>From the late stages of the Thorium series, *i.e.*, the decay chain of  $^{232}\text{Th}$ .

will constitute one of the main backgrounds for this search, as explained in detail on Section 5.3, so having a source that can recreate this very important background is extremely useful. The  $^{228}\text{Th}$  calibration source will also be used to study the response of the skin detector in more detail.

Finally, the backgrounds measured once LZ starts collecting data in its first science run will also be used to characterize the detector response. The expected high rate of events in the outer-most regions of the TPC caused by known background sources with identifiable energy signatures can be matched with detailed simulations to provide an important cross-check to high-energy calibrations.

### 5.3 Background Model

The background model developed for this analysis includes contributions from decay products of radioactive sources in detector construction materials, natural radioactivity from the laboratory rock walls, internal radioactive contaminants in the xenon, neutron-activated xenon isotopes, the two-neutrino double beta decay of  $^{136}\text{Xe}$  and neutrino-induced backgrounds. The model was constructed using the most recent radioactivity assays and extensive detector simulations [A<sup>+</sup>20d, A<sup>+</sup>20b, A<sup>+</sup>20h, A<sup>+</sup>21].

The Monte Carlo simulations of radioactive contamination in detector components and the cavern rock were generated using BACCARAT (see Section 4.4). The simulations include the entire TPC, liquid xenon skin, outer detector, water tank and cavern rock. Dedicated simulation techniques are used to model the muon flux at the Davis laboratory and  $\gamma$ -ray backgrounds from the cavern rock, since these sources are strongly attenuated in the rock overburden and in the many shielding layers of LZ, respectively. More information on the latter simulations can be found on Section 5.3.2.

Simulations of each background source are performed independently to better characterize the background and the respective response of the detector. To ensure that the background analyses are statistically significant, each simulated component of the backgrounds model is generated with statistics corresponding to tens to hundreds of thousands of days, significantly exceeding the run time of the experiment, which is expected to be 1000 days.

The backgrounds will be measured with high accuracy once the detector begins taking data, thus allowing further validation of the simulation results. The data collected with the outer detector and skin systems will also be used to characterize backgrounds from external sources in more detail, taking advantage of their peripheral positioning in relation to the TPC. The outer detector system will be notably important for the characterization of backgrounds produced by  $\gamma$ -rays from cavern rock. All data will be fit to the simulations in order to best constrain the different backgrounds observed.

Table 5.1 summarizes the main background contributions for this  $0\nu\beta\beta$  search. The backgrounds are characterized in an inner volume with cylindrical shape and dimensions  $r < 39$  cm and

**Table 5.1:** Summary table of the masses, activities and estimated background counts in the  $\pm 1\sigma$  ROI and inner 967 kg mass, for a 1000 day run, considering 1.0% energy resolution at Q-value and 0.3 cm multiple scatter rejection along  $z$  (see text for details). The veto cuts discussed in Section 5.1.4 are also applied to the background estimates presented here.

Item	Mass (kg)	$^{238}\text{U}$ -late (mBq/kg)	Counts from $^{238}\text{U}$	$^{232}\text{Th}$ -late (mBq/kg)	Counts from $^{232}\text{Th}$	Total Counts
TPC PMTs	91.9	3.22	2.95	1.61	0.10	3.05
TPC PMT bases	2.80	75.9	1.52	33.1	0.03	1.55
TPC PMT structures	166	1.60	2.65	1.06	0.12	2.77
TPC PMT cables	88.7	4.31	1.44	0.82	0.19	1.63
Skin PMTs and bases	8.59	46.0	0.75	14.9	0.02	0.78
PTFE walls	184	0.04	0.39	0.01	0.00	0.39
TPC sensors	5.02	5.82	1.19	1.88	0.00	1.19
Field grids and holders	89.1	2.63	0.62	1.46	0.11	0.73
Field-cage resistors	0.06	1350	2.63	2010	0.03	2.65
Field-cage rings	93.0	0.35 <sup>†</sup>	0.82	0.24 <sup>†</sup>	0.00	0.82
Ti cryostat vessel	2590	0.08 <sup>†</sup>	1.30	0.22 <sup>†</sup>	0.20	1.49
Cryostat insulation	13.8	11.1 <sup>†</sup>	0.90	7.79 <sup>†</sup>	0.04	0.94
Outer detector system	22900	4.71 <sup>†</sup>	1.70	3.73 <sup>†</sup>	1.08	2.79
Other components	438	1.83	2.10	1.65	0.31	2.41
Det. components subtotal	-	-	21.0	-	2.32	23.3
Cavern walls	-	29000.00	3.21	12500.00	8.41	11.6
Neutron-induced $^{137}\text{Xe}$	-	-	-	-	-	0.28*
Internal $^{222}\text{Rn}$	-	-	-	-	-	0.45*
$^{136}\text{Xe}$ $2\nu\beta\beta$	-	-	-	-	-	0.01 <sup>†</sup>
$^8\text{B}$ solar neutrinos	-	-	-	-	-	0.02
Total	-	-	24.2	-	10.7	35.6

<sup>†</sup> Upper limit

\* preliminary estimate

$^{238}\text{U}$ -late chain is  $^{226}\text{Ra}$  and after.  $^{232}\text{Th}$ -late chain is  $^{224}\text{Ra}$  and after

$26 < z < 96$  cm, containing  $\sim 967$  kg of liquid xenon. This inner volume was optimized using a Feldman-Cousins cut-and-count analysis. The values of the radial and depth cuts were optimized in tandem in order to maximize the median 90% confidence level sensitivity to the half-life of  $^{136}\text{Xe}$   $0\nu\beta\beta$  decay. The sensitivity increases with the amount of  $^{136}\text{Xe}$  present within the control volume and is limited by the background rates. There is a trade-off between total xenon mass and background levels, ensuring that an optimal volume exists where the sensitivity is maximal. The background rates on Table 5.1 are estimated using this inner volume as it represents the most sensitive region of the detector. The actual sensitivity analysis for this decay search is performed using a profile-likelihood method that utilises an larger fiducial volume, as discussed in Section 5.4.

Since the expected  $0\nu\beta\beta$  decay signal is a monoenergetic peak, the region-of-interest (ROI) for this search is defined as a symmetric energy window around  $Q_{\beta\beta}$ , which is used to characterize the backgrounds in the most sensitive energy region. It is convenient to define the ROI in relation

to the energy resolution of the detector at this energy (considered here to be 1% ( $\sigma/E$ ) at  $Q_{\beta\beta}$ , see Section 5.1). A  $\pm 1\sigma$  energy window around  $Q_{\beta\beta}$  was used, or  $2433.3 < E_{dep} < 2482.4$  keV, resulting in a signal acceptance of 68.2%. The background rates on Table 5.1 are estimated for this ROI. The energy range used for the sensitivity analysis is extended from 2000 keV to 2700 keV in order to model the backgrounds more precisely, as discussed in Section 5.4.


The largest overall contributors to the backgrounds model of this search are the detector construction materials, composing approximately 65% of the total expected background. But the largest single component contributor is the laboratory rock, responsible for almost  $1/3$  of the total counts. The uncertainty in the estimated backgrounds is dominated by the uncertainties in the measured flux of  $^{238}\text{U}$  and  $^{232}\text{Th}$   $\gamma$ -rays from the cavern rock, at approximately 50% and 25%, respectively [ $A^{+20b}$ ]. These flux measurements were performed *in-situ* with a sodium iodide detector in different locations of the Davis campus. Despite the large uncertainty on this background component, it is estimated that a  $1\sigma$  increase in the total rock gamma background would only result in a 5% decrease of the sensitivity (see Section 5.4).

Only a handful of processes can generate backgrounds at the energy scale of the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$ . Neutrons do not pose a direct background at this energy scale since they will most likely scatter multiple times in the xenon target before thermalizing or being absorbed. Both  $\beta$  and  $\alpha$  particles can only generate backgrounds if they are emitted from the radioactive decay of impurities suspended in the liquid xenon bulk. Only  $\gamma$ -rays from the decay of radioactive contaminants in detector materials have the penetrative power to reach the liquid xenon bulk and produce an energy deposition at the energy scale of  $Q_{\beta\beta}$ .

The most important  $\gamma$ -ray lines near  $Q_{\beta\beta}$  are the 2614.5 keV from  $^{208}\text{Tl}$  decay in the  $^{232}\text{Th}$  decay chain and the 2447.7 keV  $\gamma$ -ray line from  $^{214}\text{Bi}$  in the  $^{238}\text{U}$  decay chain.  $^{214}\text{Bi}$  decays into  $^{214}\text{Po}$  via  $\beta$  emission with a half-life of  $19.9 \pm 0.4$  minutes (branching ratio of 99.979%) and Q-value of 3272 keV (see Appendix A). The 2447.7 keV  $\gamma$ -ray line is emitted along with the  $\beta$  with a 1.5% branching ratio. Since this  $\gamma$ -ray line is only 10 keV away from  $Q_{\beta\beta}$  it cannot be separated from the expected  $0\nu\beta\beta$  decay signal with the energy resolution of LZ (here assumed to be 1.0%), making it a very problematic background for this search.  $^{208}\text{Tl}$  is the progeny of  $^{212}\text{Bi}$  that decays by  $\alpha$  emission with a branching ratio of 35.9%.  $^{208}\text{Tl}$  itself decays into the stable  $^{208}\text{Pb}$  via  $\beta$  emission with a half-life of  $3.053 \pm 0.004$  minutes and a Q-value of 5001 keV. The  $\beta$  emission is always accompanied by the emission of a  $\gamma$ -ray of 2614.5 keV, about 160 keV higher in energy than  $Q_{\beta\beta}$ , that can undergo Compton scattering before reaching the xenon and produce a background for this search. The outer detector and skin vetoes will play an important role in tagging the 2614.5 keV  $^{208}\text{Tl}$  line, which is always emitted with another  $\gamma$ -ray of at least 583 keV. Another potential radiogenic background is caused by the sum of the 1173.2 keV and 1332.5 keV  $\gamma$ -ray lines from  $^{60}\text{Co}$  that could result in an energy deposition of 2505.7 keV, only  $\sim 50$  keV away from  $Q_{\beta\beta}$ . However, simulations indicate that a good rejection of multiple scatter events will render this background negligible. Nevertheless, the effect of  $^{60}\text{Co}$  in the cathode grid was simulated and included in the final sensitivity analysis, for completeness.

Figure 5.2 shows the spacial distribution of simulated background events in the ROI versus  $z$  and radius squared. Due to the large density of liquid xenon, and consequently its self shielding

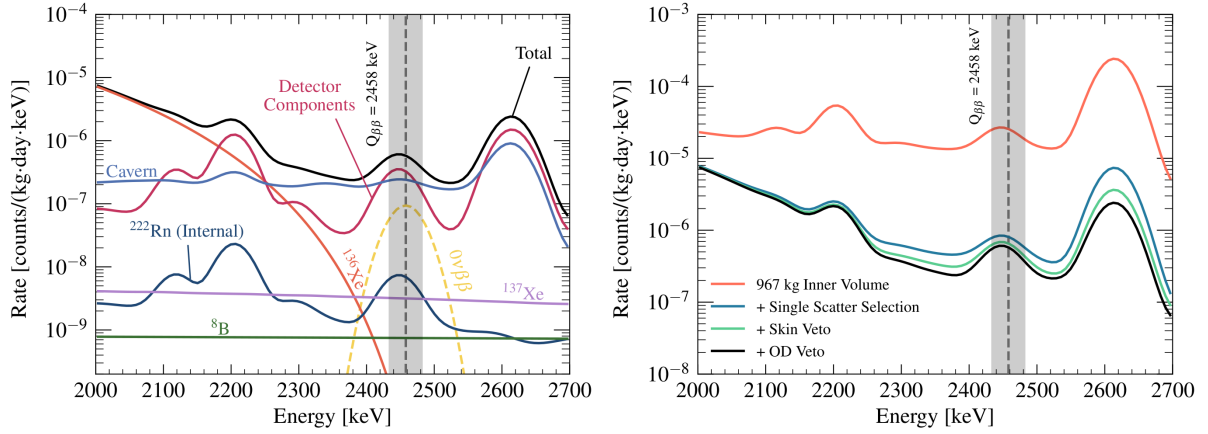
properties, the innermost region of the detector has a much lower background rate than the periphery. It is also visible that the background rates are higher at the top than at the bottom of the active volume, as the bottom PMTs are shielded by the liquid xenon in the reverse field region and also due to coincidences in the skin dome. The selection of an innermost region of the TPC optimised for this search results in the rejection of a large fraction of backgrounds at the expense of a large fraction of the available mass. The left-hand side plot of Figure 5.3 displays the total background spectrum within the 967 kg volume and for an exposure of 1000 days after all analysis cuts are applied (solid black line) and the contributions from individual background components (solid coloured lines). The dashed yellow line illustrates the expected spectrum of the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  for a half-life of  $1.06 \times 10^{26}$  years. The right-hand side plot of Figure 5.3 displays the impact of the successive selection cuts used in this analysis on the background spectrum in the inner 967 kg volume. The solid black line represents the total background spectrum within the 967 kg volume and for an exposure of 1000 days, the same total spectrum in the left-hand side plot of Figure 5.3. The “single scatter” selection provides the strongest background rejection for  $\gamma$ -rays of these energies. However, it cannot exclude events from the  $2\nu\beta\beta$  decay of  $^{136}\text{Xe}$  that are expected to produce single scatters similar to the expected  $0\nu\beta\beta$  decay signal, resulting in the apparent loss in background rejection efficiency of this analysis cut at lower energies.



**Figure 5.2:** Background event rate in the active region and in the  $\pm 1\sigma$  energy ROI as a function of  $r^2$  and  $z$ . The dashed black rectangle represents the inner 967 kg volume where LZ is most sensitive to the  $0\nu\beta\beta$  decay, while the larger dashed white rectangle represents the extended fiducial volume used on the profile likelihood analysis (see Section 5.4). Plots from Reference [A<sup>+</sup>20c].

### 5.3.1 Detector Components

As mentioned previously, the full LZ detector geometry, with every component of the TPC, skin, outer detector and auxiliary systems, is included in the detector simulations and in the background model. The first half of Table 5.1 is dedicated to the most important detector components, highlighting the main contributions to the background. Some of the items in Table 5.1 correspond to groupings of individual detector components and materials that are simulated individually. The masses and averaged activities of these components are also presented, and



**Figure 5.3:** Background energy spectrum in the inner 967 kg volume: contributions from the main background sources (left) and effect of successive selection cuts used in the analysis (right). More details regarding the analysis can be found in the text. The gray vertical band around  $Q_{\beta\beta}$  represents the  $\pm 1\sigma$  ROI used for this search. The  $^{60}\text{Co}$ ,  $^{238}\text{U}$ -late chain and  $^{232}\text{Th}$ -late chain backgrounds from the detector components are combined into a single curve on the left plot but are treated independently in the sensitivity analysis. The dashed yellow line in the left plot represents the expected signal spectrum for  $^{136}\text{Xe}$   $0\nu\beta\beta$  decay, considering a half-life of  $1.06 \times 10^{26}$  years (see Section 5.4), and is not included in the total spectrum. The spectra are smeared using the energy resolution function of LUX [ $A^+17c$ ], scaled to be 1.0% at  $Q_{\beta\beta}$ . Plots from Reference [ $A^+20c$ ].

divided into the contributions of the  $^{238}\text{U}$ -late and  $^{232}\text{Th}$ -late decay chains. These categories refer only to the late chain activity, from  $^{226}\text{Ra}$  and  $^{224}\text{Ra}$  onward respectively, as several samples measured by LZ were found to be out of secular equilibrium [ $M^+17$ ]. Simulations of the  $^{238}\text{U}$ -late and  $^{232}\text{Th}$ -late decay chains, as well as the decays of  $^{60}\text{Co}$ , are generated for each individual detector component using dedicated decay generators that correctly model the full decay chains and produce the correct branching ratios of all particle products. The contributions of  $^{60}\text{Co}$  decay to the background model are included on the final sensitivity analysis, as discussed on Section 5.4, but are excluded from Table 5.1 due to being negligible in the inner 967 kg volume. All construction materials have been screened for radioactivity in an aggressive radioassay campaign [ $A^+20h$ ]. Some contamination values for detector materials are measured upper limits, the most relevant of which will be re-screened to improve the accuracy of the model.

The top contributors to the total background rate amongst detector components are the TPC PMTs and related systems (PMT bases, support structures and cabling), the ceramic field-cage resistors, the cryostat vessel, and the outer detector system. Despite the outstanding radiopurity of the Hamamatsu R11410-22 PMT model, the relatively high mass ( $\sim 100$  kg) of the 494 used in LZ and their proximity to the active volume results in a contribution of nearly 15% the total background counts from detector components. The titanium used to manufacture the cryostat vessel has the lowest measured  $^{238}\text{U}$ ,  $^{232}\text{Th}$ ,  $^{60}\text{Co}$  and  $^{40}\text{K}$  activities reported to date [ $A^+17f$ ,  $A^+20h$ ]. Nevertheless, this radiopurity is compensated by the 2.6 tonne mass, located near the active volume, resulting in a sizeable background rate. It is worth noting that the titanium itself contributes with only half the total background counts associated to the cryostat, the remaining originate from the aluminium and helicoflex seals, nitrile elastomer and stainless

steel bolts in the vessel flanges.

The largest contribution to the background rate from the outer detector system is from the 1.1 tonne stainless steel structure that supports the acrylic tanks. The 4.3 tonne acrylic tanks, despite having low contamination levels, have a non-negligible contribution due to their large mass.

The ceramic resistors installed on the field-shaping rings of the TPC weigh less than 60 grams combined but have high contamination levels and are adjacent to the active volume, resulting in a contribution of more than 10% to the total background rate of the detector materials.

### 5.3.2 Davis Cavern Walls

$\gamma$ -rays from the laboratory rock are the largest single component contributor to the total background rate. These  $\gamma$ -rays are produced in the radioactive decay of natural rock contaminants, mainly  $^{40}\text{K}$  and the progeny of the  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains. The rock surrounding the Davis cavern is composed primarily of amphibolite and rhyolite [Les15, A<sup>+</sup>20b]. Amphibolite is a metamorphic rock with relatively low natural abundance of the aforementioned radioactive contaminants [Les15]. Some rhyolite intrusions are exposed in the walls of the Davis cavern and contain much higher radioactive contamination levels. Most of the floor of the laboratory is layed with 15 cm of low-radioactivity concrete, while the walls and ceiling of the cavern are lined with a layer of sprayed concrete (shotcrete) of average thickness 12.7 cm. Table 5.2 summarizes the activities of the main materials that are present in the walls of the Davis cavern [A<sup>+</sup>20b]. The floor concrete and shotcrete have similar measured activities for  $^{40}\text{K}$ ,  $^{238}\text{U}$  and  $^{232}\text{Th}$ , while the rhyolite displays roughly 5 times more activity of all radioactive species. Some gravel present below the water tank has an activity of 1.7 Bq/kg of  $^{232}\text{Th}$  and of 26.3 Bq/kg of  $^{238}\text{U}$ .

**Table 5.2:** Assayed activities of the rock and lining materials of the Davis cavern walls, taken from Reference [A<sup>+</sup>20b].

Sample	$^{40}\text{K}$ (Bq/kg)	$^{238}\text{U}$ (Bq/kg)	$^{232}\text{Th}$ (Bq/kg)
Homestake	297	2.7	1.3
Rhyolite	1291	108	44
Concrete	381	27	13
Shotcrete	272	23	12
Gravel	35.0	26.3	1.7

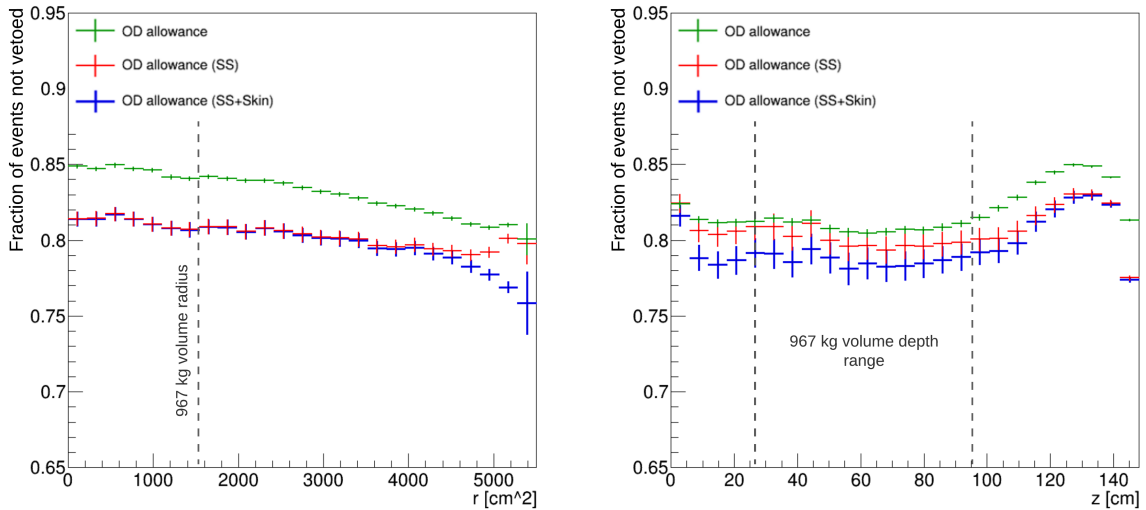
Previous measurements of the  $\gamma$ -ray flux from the rock at different SURF facilities demonstrated that the flux can vary significantly depending on the variation in the geological formations [MZTG10], prompting a direct measurement campaign on the Davis cavern in order to characterize this background for LZ. To that end, recent *in-situ* measurements were performed in several locations around the Davis cavern with a sodium iodide detector [A<sup>+</sup>20b]. These measurements indicate an average activity of 29.0 Bq/kg of  $^{238}\text{U}$  and of 12.5 Bq/kg of  $^{232}\text{Th}$  in the surrounding cavern rock. The uncertainties associated with these measurements are estimated



at approximately 50% and 25%, respectively, and dominate the uncertainties in the background model.

The water tank surrounding the LZ detector systems has a height of 591 cm and a radius of 381 cm, providing at least 70 cm of shielding to external  $\gamma$ -rays in any direction [A<sup>+</sup>20d]. Furthermore, an inverted pyramid shield build from 6 octagonal steel plates of 5 cm thickness is embeded between the concrete and the bottom of the water tank to provide additional passive shielding to  $\gamma$ -rays coming from the rock and gravel below. An additional 2.7 m steel shield with 8 cm thickness will be added to the top of the water tank to provide additional coverage of  $\gamma$ -rays coming from the top of the detector. No additional shielding is assumed to be on the sides of the water tank.

The outer detector system that encompasses the LZ cryostat will be crucial to mitigate the effect of  $\gamma$ -rays from the Davis cavern rock, providing both passive and active shielding to external  $\gamma$ -rays and neutrons. The outer detector vessels contain 17.3 tonnes of gadolinium-loaded liquid scintillator (GdLS) with a minimum thickness of 60 cm on the sides (see Figure 4.1 and Section 4.3). The scintillation light is detected by 120 Hamamatsu R5912 8 inch PMTs mounted outside of the acrylic tanks and suspended in the water space. With a light yield of about 130 photoelectrons for a 1 MeV energy deposit in the liquid scintillator [M<sup>+</sup>17] the outer detector will be able to measure the flux of  $\gamma$ -rays from the cavern walls with high precision. Figure 5.4 shows the impact of the outer detector in the rejection of events from the 2614.5 keV line from  $^{208}\text{Tl}$  in the rock walls. Roughly 20% of the background events produced by this source in the ROI, which is only 160 keV away from  $Q_{\beta\beta}$ , can be tagged using the outer detector by setting the energy threshold at 100 keV.



**Figure 5.4:** Fraction of background events from  $^{208}\text{Tl}$  ( $^{232}\text{Th}$  progeny) from the Davis cavern walls that are not vetoed by the outer detector system (assuming a 100 keV threshold) for different radial positions (left) and depths (right) of the interactions. The effect of the outer detector veto alone is represented in green, while the combination of this veto with the single scatter selection is displayed in red. The combination of both skin and outer detector vetoes and single scatter selection is represented in blue.

### 5.3.2.1 Simulation strategy for external $\gamma$ -ray sources


The background from the Davis cavern walls was estimated using dedicated simulations based on BACCARAT but slightly modified to account for the strong attenuation of sources outside the water tank [A<sup>+</sup>21]. The simulations include both veto systems, as well as the steel shield above the water tank and the bottom pyramid shield below that are not part of the remaining simulations for internal detector components. A 30 cm thick rock shell with the chemical composition of a Homestake rock sample was included around the water tank volume to simulate the Davis cavern walls [A<sup>+</sup>20b]. The full  $^{238}\text{U}$  and  $^{232}\text{Th}$  chains were simulated as contaminants in the rock shell. The contaminants are assumed to be in secular equilibrium.

Due to the extreme attenuation of external  $\gamma$ -rays on the top and bottom steel shields, water tank, veto systems and detector materials, only a very small fraction of the simulated events reach the TPC, resulting in simulations that are very inefficient [A<sup>+</sup>21]. In order to reach the statistics equivalent to 10 runs of 1000 days each for the background of the 2614.5 keV line from  $^{208}\text{Tl}$ , an estimated  $10^{17}$  full decay-chains of  $^{232}\text{Th}$  would have to be simulated, which is impractical to say the least. The same event biasing technique used in References [M<sup>+</sup>17, A<sup>+</sup>20b, A<sup>+</sup>21, Woo18, A<sup>+</sup>21] was used in this work to boost the simulations and provide enough statistics to model these backgrounds.

In essence, the event biasing technique tracks all  $\gamma$ -rays produced on the rock shell and records the positions, directions and energies of the  $\gamma$ -rays that reach a set of pre-defined nested cylindrical surfaces encompassing the LZ detector [A<sup>+</sup>21]. A succession of boosting steps, each corresponding to one of the surfaces, uses the particle information stored on the surface to simulate the particles onward with the same momentum but with a multiplicative factor in order to increase statistics. Figure 5.5 shows the approximated position of the surfaces used in each boosting stage relatively to the LZ detector system.

Due to the geometry of the boosting stages and the fact that the surface of the final stage is within the outer detector volume, the outer detector veto efficiency cannot be accounted for in the final simulation stage. This is not an issue for the low-energy analysis of LZ like the WIMP search analysis. Since the final simulation stage tracks all  $\gamma$ -rays until their absorption, the full OD veto efficiency is preserved for low energy searches where the outer detector will only be used to veto  $\gamma$ -rays that deposited little energy in the TPC via Compton scattering and deposit some energy in the veto systems afterwards. Conversely, for high-energy searches the outer detector will be used to veto incoming  $\gamma$ -rays that scatter in the scintillator and proceed to the TPC where they deposit most of their energy. As mentioned previously, this veto efficiency is of special importance for this  $0\nu\beta\beta$  decay search in order to mitigate the background from the 2614.5 keV line from  $^{208}\text{Tl}$  in the rock.

The outer detector veto efficiency was estimated by tracing back the information of  $\gamma$ -rays that produced a background in the TPC to the penultimate stage that strategically includes the full outer detector system. Any energy deposited in the scintillator by the  $\gamma$ -rays that produced a background is recovered and used to estimate the veto efficiency. The results on Figure 5.4 were obtained using this method. This estimated efficiency is then applied in the analysis after the




**Figure 5.5:** Approximated spacial positioning of the cylindrical surfaces used in the boosting stages of the rock gamma simulations. Each stage records the information of any incoming particle that crosses its boundary. The next iteration of the simulation will use the information recorded at the previous stage to propagate the particles with boosted statistics, recording once more the information of the particles that reach the boundary of the next stage. The statistics are boosted by an overall factor of  $10^8$ ,  $100\times$  on each stage.

background estimates.

### 5.3.2.2 High-energy $\gamma$ -rays from the cavern walls

The measurements of  $\gamma$ -ray flux in the Davis Cavern were limited to energies below 3 MeV [A<sup>+</sup>20b]. However, the simulations used to estimate the background from these  $\gamma$ -rays predict a non-negligible background rate from  $\gamma$ -rays with energies above 3 MeV. Figure 5.6 shows the spectrum of  $\gamma$ -rays from the cavern rock that reach the TPC for both the  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains as estimated from the simulation. The high-energy  $\gamma$ -rays are produced on both decay chains in similar ratios and seem to follow the same spectral shape, which could be an indication that the same physical process is responsible for producing these  $\gamma$ -rays.

The origin of the simulated high-energy  $\gamma$ -rays was traced back to the “AlphaInelastic” physics process in Geant4.10.3, that at the time used the *QGSP\_BIC\_HP* and *G4EmLivermorePhysics* physics lists for hadronic and electromagnetic processes, respectively [A<sup>+</sup>03]. Early simulations using Geant4.9.5 displayed a much lower flux of these high-energy  $\gamma$ -rays. Both the  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains have several  $\alpha$  emissions with energies up to 8.8 MeV that may produce these high energy  $\gamma$ -rays in  $(\alpha, \gamma)$  reactions in the rock on oxygen and silicon [TZMC17], the most abundant elements in the cavern rock at SURF, primarily composed of  $\text{SiO}_2$ ,  $\text{Al}_2\text{O}_3$ ,  $\text{FeO}$  and water [A<sup>+</sup>20b]. However, no quantitative comparison was performed between the simulated



**Figure 5.6:** Spectrum of the energy depositions in the TPC from  $\gamma$ -rays produced in the cavern walls by  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains (red and blue, respectively). Both spectra were produced with the same initial statistics. The flux of  $\gamma$ -rays with energies above 3 MeV is non-negligible due to their larger penetrative power. Despite the predominance of the  $^{232}\text{Th}$  spectrum at lower energies, the  $\gamma$ -ray flux above 3 MeV in the TPC is similar for both sources.

fluxes and  $(\alpha, \gamma)$  reactions in the rock nor could be found in the literature. Despite their high penetrative power, the flux of high-energy  $\gamma$ -rays in the TPC is expected to be 2 to 3 orders of magnitude lower than the radiogenic  $\gamma$ -rays from the  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains. This estimate, taken directly from simulation, results in a background rate from these  $\gamma$ -rays of less than one count in the inner 967 kg volume for a 1000 day exposure.

Since the overall background rate from high-energy  $\gamma$ -rays is much smaller than the one from radiogenic  $\gamma$ -rays, these were not excluded from the analysis, which is expected to be a conservative approach. The flux of these high energy  $\gamma$ -rays from the Davis Cavern walls will be measured in detail when LZ starts collecting data.

### 5.3.3 Internal Radon

Radon mixed in the liquid xenon bulk is the leading source of background at low energies. Most of the radon that finds its way into the liquid bulk emanates from detector materials and residual dust in the internal surfaces of the TPC. There are four naturally occurring Radon isotopes produced in the decay chains of long lived heavy elements, of which the two that are relevant for rare event searches are  $^{222}\text{Rn}$  produced in the  $^{238}\text{U}$  decay chain, and  $^{220}\text{Rn}$  produced in the  $^{232}\text{Th}$  decay chain. These radon isotopes are relatively short-lived, with half-lives of 3.82 days and 55.6 seconds, respectively. LZ has a requirement of  $<2.0 \mu\text{Bq/kg}$  of  $^{222}\text{Rn}$ , equivalent to 14 mBq in the active xenon.

Due to the long diffusion lengths of radon in solid materials, and the long lifetime of their progenitor isotopes, radon tends to accumulate in underground spaces with limited air flow. Radon levels are constantly monitored at several locations in the Homestake mine where LZ is installed, to ensure that there is no excessive radon build up in the air. Some charged radon-

daughters are expected to plate out onto material surfaces of detector components. To avoid contaminating the detector materials, all TPC assembly stages are performed at the surface in a radon-reduced and dust-reduced environment. During assembly and installation underground, the TPC and cryostat are stored in controlled atmospheres regularly purged with radon-reduced nitrogen to limit radon daughter plating.

The  $^{214}\text{Bi}$  daughters of the  $^{222}\text{Rn}$ -chain decays via  $\beta$  emission with Q-value of 3.27 MeV. The dominant background induced by this decay is the  $\beta$  emission without any accompanying  $\gamma$ -ray emission, here named a “naked- $\beta$ ” decay, with a branching ratio of 19.1%. This is the most problematic background from any radon-induced events in the bulk. Since this is a background completely contained within the liquid bulk, the skin and outer detector systems are ineffective at vetoing these events. At 1.0% energy resolution, 0.5% of  $^{214}\text{Bi}$  naked- $\beta$  decays will result in a single-scatter event in the ROI. Fortunately, this background can be vetoed by tagging the 7.8 MeV  $\alpha$  decay of the short-lived  $^{214}\text{Po}$  daughter of  $^{214}\text{Bi}$ , with a 163.6  $\mu\text{s}$  half-life. A rejection of more than 99.99% of the internal  $^{214}\text{Bi}$  decays can be achieved by excluding all events with an  $\alpha$  within a time window of 2.5 ms after the initial decay. This technique, named “BiPo tagging”, reduces the background from this source to around 0.03 background events per tonne in the ROI.

The 2448 keV  $\gamma$ -ray line of  $^{214}\text{Bi}$ , with 1.5% branching ratio, cannot be separated from the  $0\nu\beta\beta$  ROI by energy resolution alone. It can, however, be vetoed by detecting the coincident  $\beta$  emission for events in the xenon bulk. The low energy threshold of LZ (50% efficiency at 1.5 keV $_{ee}$ ) ensures that >99.97% of decays will be rejected by coincidence with the  $\beta$ , making this  $\gamma$  background negligible for  $^{214}\text{Bi}$  in the liquid bulk.

Some of the  $^{222}\text{Rn}$  daughters produced in the liquid bulk are positively charged and will drift to the cathode, due to the electric field and liquid circulation currents in the active xenon volume, or will be captured on the PTFE walls of the TPC. This effect was observed in EXO-200 [A<sup>+</sup>15e], where the internal  $^{214}\text{Bi}$  activity in their fiducial volume was found to be only 11.6% of the  $^{222}\text{Rn}$  activity, and also in LUX [A<sup>+</sup>15d, Bra14], where the measured activity of  $^{218}\text{Po}$  is  $\sim 20\%$  lower than that of  $^{222}\text{Rn}$  despite being neighbouring decays and secular equilibrium being reached very early in the science run. A large fraction of the  $\alpha$  emissions from  $^{210}\text{Po}$  were also directly observed coming from the PTFE walls and cathode grid on LUX. Once the  $^{222}\text{Rn}$  daughters are plated in one of the surfaces of the TPC, their decay can produce backgrounds that cannot be vetoed using a coincident decay. Considering that particle emissions are isotropic on the decays of these trapped daughters, there is a roughly 25% chance of both the  $\beta$  particle from  $^{214}\text{Bi}$  and the  $\alpha$  particle from  $^{214}\text{Po}$  being absorbed by the PTFE walls or by the cathode grid wires. In this scenario, a 2448 keV  $\gamma$ -ray from the decay of the plated out  $^{214}\text{Bi}$  that interacts in the liquid xenon cannot be vetoed. This is the main source of radon-related backgrounds for this search. Events in which the  $\beta$  decay goes into the active volume but the  $\alpha$  particle is absorbed are excluded by their position.

The same simulations used to predict the backgrounds from  $^{214}\text{Bi}$  in the  $^{238}\text{U}$  chain on the cathode grid were also used to predict this background, but accounting for the 25% of events where both the  $\beta$  and  $\alpha$  are absorbed by the walls, as well as the fraction of mass inside the

active region. To conservatively estimate this background, it is assumed that all the  $^{214}\text{Bi}$  decays in the cathode, since the plating in the PTFE walls is sub-dominant. Under this assumption, around 0.41 events in the inner 967 kg volume and in 1000 days are predicted for this source. This value can be reduced further by the detection of the recoil signal from  $^{210}\text{Pb}$  after the  $\alpha$  emission.

The total contribution of internal radon to the background rate is 0.44 counts per 1000 days in the ROI and in the inner 967 kg fiducial volume. Note that this estimate is conservative. First, the calculations of both backgrounds from radon daughters (cathode plating and decays in the xenon bulk) assume the total  $^{222}\text{Rn}$  activity, essentially duplicating the effect of the decay of  $^{214}\text{Bi}$ . Second, these calculations do not assume any removal of radon or its daughters by the circulation system.

### 5.3.4 Internal $^{137}\text{Xe}$

The isotope  $^{137}\text{Xe}$  is a  $\beta$ -emitter with a Q-value of 4173 keV and a half-life of 3.8 minutes. This  $\beta$ -decay spectrum overlaps with the  $0\nu\beta\beta$  ROI, and 67% of the decays do not have an accompanying  $\gamma$  emission, resulting in a “naked”  $\beta$ -decay that cannot be vetoed. This isotope is not present in natural xenon due to its short half-life, but it can be produced by the absorption of radiogenic and muon-induced neutrons by  $^{136}\text{Xe}$  through the reaction  $^{136}\text{Xe}(n,\gamma)^{137}\text{Xe}$ .

Of all the natural isotopes of xenon that can be produced by neutron absorption, only  $^{137}\text{Xe}$  poses as a potential background at the energy range of this  $0\nu\beta\beta$  decay search.

#### 5.3.4.1 Muon-induced neutron activation of $^{136}\text{Xe}$

Neutrons can be produced in electromagnetic, hadronic or nuclear cascades initiated by muons, or by muon-induced spallation of a nucleus [ $A^{+20d}$ ,  $A^{+15d}$ ]. The energies of these neutrons extend to the GeV range, making them harder to moderate and/or absorb when compared to the less energetic radiogenic neutrons ( $\leq 10$  MeV). With a mean free path in liquid xenon much longer than the size of the detector, muon-induced neutrons can travel far from the muon track that originated them and possibly evade the muon veto system of LZ. Fortunately, the flux of muon-induced neutrons tends to be significantly lower than that of radiogenic neutrons in deep underground facilities.

This cosmogenic background component was estimated using a conservative calculation, taking advantage of the measured background rates on similar experiments. EXO-200, located at the Waste Isolation Pilot Plant (WIPP) in New Mexico (USA), measured  $338_{-93}^{+132}$   $^{136}\text{Xe}(n,\gamma)^{137}\text{Xe}$  captures per year by fitting data in coincidence with hits in their muon veto. The background rate measured by EXO-200 is  $5.1_{-1.4}^{+2.0}$  events per year in a fiducial volume containing 76.5 kg of  $^{136}\text{Xe}$  [ $A^{+15e}$ ], or  $\sim 183$  ROI events per tonne of  $^{136}\text{Xe}$  over a 1000 day exposure [ $A^{+16e}$ ]. This rate implies that around 1.5% of neutron captures on  $^{136}\text{Xe}$  produce a background event in their ROI.

Comparatively to EXO-200, LZ has three advantages that reduce the muon-induced  $^{137}\text{Xe}$  background: the lower muon flux in the Davis cavern, the lower abundance of  $^{136}\text{Xe}$ , and the veto capabilities of the outer detector and skin systems.

The muon flux in the Davis cavern is calculated to be  $6.2 \times 10^{-9} \text{ cm}^{-2} \text{ s}^{-1}$  [M<sup>+</sup>17, A<sup>+</sup>21], a value nearly 100 times lower than at WIPP ( $4.0 \times 10^{-7} \text{ cm}^{-2} \text{ s}^{-1}$ ). The calculated flux agrees within 20% with the measurements carried out with the veto system of the Davis experiment [CDL<sup>+</sup>83] and the veto system of the Majorana demonstrator [A<sup>+</sup>17a] (see [A<sup>+</sup>21] for further discussion). This reduction in muon flux results in an overall decrease in  $^{137}\text{Xe}$  production from muon-induced neutrons by the same factor.

The low abundance of  $^{136}\text{Xe}$  in natural xenon compared to the 80% enriched EXO-200 xenon also contributes to the reduction of this background. With only 8.9% of  $^{136}\text{Xe}$  in natural xenon, the  $^{137}\text{Xe}$  activation in LZ is expected to be decreased by an order of magnitude from isotopic abundance alone for the same total mass.

Considering only the muon flux reduction by a factor of 100 and the abundance of  $^{136}\text{Xe}$  in natural xenon, the estimated background rate from muon-induced neutron capture on  $^{136}\text{Xe}$  is 0.16 events in the ROI and inner 967 kg volume of natural xenon over the 1000 day exposure. Furthermore, the presence of other isotopes of xenon with a higher neutron capture cross section than  $^{136}\text{Xe}$  will also contribute to reducing this background by more than an order of magnitude. Nearly half of natural xenon is  $^{129}\text{Xe}$  or  $^{131}\text{Xe}$ , each of which have neutron capture cross sections more than five times larger than  $^{136}\text{Xe}$  [NHW<sup>+</sup>07].

Finally, LZ will take advantage of its large active xenon mass and active veto systems that will provide excellent muon veto ability as well as an expected high efficiency at tagging neutron capture cascade  $\gamma$ -rays produced in the  $^{136}\text{Xe}(n,\gamma)^{137}\text{Xe}$  reaction. By looking back over several  $^{137}\text{Xe}$  half-lives for these signatures, potential  $^{137}\text{Xe}$  background may be vetoed as in EXO but with higher efficiency. The efficiency of tagging a muon that crosses any active volume of the LZ detector is expected to be close to 100% [A<sup>+</sup>21], while a neutron produced by a peripheral muon track would scatter several times inside the active xenon volume. XENON1T estimates a veto efficiency of 99.5% for muon tracks in their water tank from Cherenkov light alone [A<sup>+</sup>14c]. LZ is expected to reach a muon veto efficiency higher than 99%.

Taking into account all the assumptions mentioned above, the estimated background rate from muon-induced neutron capture on  $^{136}\text{Xe}$  becomes  $\ll 0.01$  events in the ROI and inner 967 kg volume of natural xenon over the 1000 day exposure, and is considered negligible for this search.

#### 5.3.4.2 Radiogenic neutron activation of $^{136}\text{Xe}$

$^{137}\text{Xe}$  can also be produced via thermal neutron capture on  $^{136}\text{Xe}$ . The flux of thermal neutrons within the LZ xenon space is strongly suppressed by the water tank and Gadolinium-loaded liquid scintillator on the outer detector system. However, the xenon purification system is located outside of the water tank and it is therefore not shielded from the thermal neutron

flux, measured to be  $\phi_n = 1.7 \times 10^{-6} \text{ cm}^{-2} \text{ s}^{-1}$  within the Davis cavern [B<sup>+</sup>15]. There are roughly 70 kg of unshielded xenon in the purification system at any given time, and given the 4000 kg/day xenon purification rate, 10 kg of xenon are delivered to LZ over the 3.8 minute half-life of  $^{137}\text{Xe}$ . Considering a thermal neutron capture cross-section  $\sigma_n = 0.26$  barn [A<sup>+</sup>16f] and the thermal neutron flux  $\phi_n$  mentioned above, the activation rate in the xenon purification system is given by:

$$R_{\text{Xe}137} = \sigma_n \phi_n \eta_{\text{Xe}136} \frac{m_{\text{Xe}} N_A}{M_{\text{Xe}136}} = 1.1 \text{ atoms/day},$$

where  $\eta_{\text{Xe}136}$  is the natural abundance of  $^{136}\text{Xe}$ ,  $m_{\text{Xe}}$  is the mass of xenon in the purification system and circulation pipes,  $M_{\text{Xe}136}$  is the molar mass of  $^{136}\text{Xe}$  and  $N_A$  is the Avogadro constant. Some of the  $^{137}\text{Xe}$  will decay within the purification system. The equilibrium concentration of  $^{137}\text{Xe}$  is given by  $R_{\text{Xe}137} T_{1/2} / \ln 2 = 0.06$   $^{137}\text{Xe}$  atoms per tonne of natural xenon (in the 70 kg of the circulation system).

Considering the 4 tonne/day purification rate and assuming that around half of the  $^{137}\text{Xe}$  atoms will either decay before reaching the TPC or will end up in the xenon skin volume or RFR, only 0.13 atoms/day will reach the active volume of LZ. Furthermore, considering that these atoms are uniformly distributed in the TPC – which is very conservative due to the slow mixing of liquid xenon in LZ and the half-life of  $^{137}\text{Xe}$  – only around 1/7 of the total atoms will end up in the inner 967 kg fiducial volume. Finally, assuming the 1.5% of  $^{137}\text{Xe}$  decays that fall within the  $\pm 1\sigma$  ROI in EXO [A<sup>+</sup>16e, A<sup>+</sup>15e], and conservatively assuming that this fraction also applies to LZ despite the superior energy resolution, the estimated background rate from radiogenic neutron capture on  $^{136}\text{Xe}$  is estimated to be 0.28 events in the ROI and inner 967 kg volume over the 1000 day exposure.

This background can be mitigated with the installation of a neutron shield around the xenon purification system. A reduction of thermal neutron flux by two orders of magnitude is achievable with the addition of a layer of 20 cm of HDPE around the xenon systems outside the water tank [Viv10, TRK10, LRK<sup>+</sup>06]. This additional shielding is not relevant for dark matter searches, but it would be required for a dedicated  $0\nu\beta\beta$  decay search using  $^{136}\text{Xe}$  enriched xenon in LZ, as discussed at the end of this Chapter.

### 5.3.5 Physics Backgrounds

One of the defining features of the multi-purpose physics laboratory that is LZ is its sensitivity to several physics interactions of significant scientific interest. However, some of these interesting phenomena pose as backgrounds<sup>3</sup> for this particular  $0\nu\beta\beta$  decay search. The two main physics backgrounds, which are discussed below, are neutrino-induced interactions from solar, supernova and atmospheric neutrinos (see Figure 3.2), and the very own two-neutrino decay mode of  $^{136}\text{Xe}$  (see Figure 3.10).

<sup>3</sup>As particle and astroparticle physics evolves in pursuit of the solutions to deeper and deeper questions, the once sought out *new physics* steadily converts into our *Standard Models*. Richard Feynman is quoted as saying “*yesterday’s discovery is today’s calibration*”, and to which Val Telegdi replies “*...and tomorrow’s background!*”.



### 5.3.5.1 Neutrino-electron elastic scattering

Solar neutrinos can scatter elastically with electrons from the liquid xenon target, and those from the  ${}^8\text{B}$  decay in particular can do so with enough energy to produce a signal in the  $Q_{\beta\beta}$  range. *hep* neutrinos also extend to higher energies, but their flux is approximately three orders of magnitude lower than that of  ${}^8\text{B}$  neutrinos and is therefore not considered here. The same argument is used for supernova and atmospheric neutrinos, despite their average higher energies. The spectrum of  ${}^8\text{B}$  neutrinos extends up to 16.36 MeV and can transfer up to 15.87 MeV to an electron in the liquid xenon [RBFFS14, Bah87]. A fraction of these events will produce a background in the  $0\nu\beta\beta$  decay ROI. The elastic scattering differential cross sections are calculated at tree-level for pure Z boson exchange, *i.e.*, neutral-currents (NC), and for a combination of Z and W boson exchange, *i.e.*, neutral-currents and charged-currents (NC+CC), following the procedure described in References [Bah87] and [MP03]. The four-fermion effective interactions for both NC and NC+CC interactions are represented in Equations 5.2 and 5.3, respectively.

$$\bar{\nu}_l e \longrightarrow \bar{\nu}_l e \quad (\text{NC}) \quad (l = \mu \text{ or } \tau) \quad (5.2)$$


$$\bar{\nu}_e e \longrightarrow \bar{\nu}_e e \quad (\text{NC+CC}), \quad (5.3)$$

Here  $\bar{\nu}$  represents either  $\nu$  or  $\bar{\nu}$ . The cross section for different neutrino flavours and for both interaction types are calculated separately. The total  ${}^8\text{B}$  neutrino flux considered in this estimate is  $5.79 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$ , corresponding to the largest value in literature and the worst case scenario (BP04) [BSB05]. The survival probability for  ${}^8\text{B}$  electron neutrinos, averaged over the entire energy spectrum, is  $P_{ee} = 0.33 \pm 0.02$  [Vis17], and the complement is used when calculating for muon and tau neutrinos<sup>4</sup>. Figure 5.7 shows the calculated differential rates for both pure Z boson exchange ( $\mu$  and  $\tau$  neutrinos, right plot) and a mixture of Z and W boson exchange (electron neutrinos, left plot). The expected combined background rate from neutrino-electron elastic scattering of solar neutrinos is  $<0.01$  events in the ROI per tonne per year, or 0.02 events over 1000 live days and in the inner 967 kg volume.

### 5.3.5.2 Charged-current neutrino-nucleus interactions

${}^8\text{B}$  solar neutrinos can also produce backgrounds in the  $0\nu\beta\beta$  decay energy region by the neutrino capture process  $\nu_e + {}^A\text{Xe} \longrightarrow e^- + {}^A\text{Cs}$ , both by the prompt emission of an energetic electron and also in the subsequent decay of the produced caesium nuclei. The prompt electron emitted following the neutrino capture in any of the xenon isotopes is expected to have an energy far greater than  $Q_{\beta\beta}$ , resulting in a negligible fraction of events in the ROI, while the prompt de-excitation of the Cs nuclei will likely produce several  $\gamma$ -rays and thus be rejected as multiple scatters with high efficiency [EE17, A<sup>+</sup>18g, EE14, PSY19].

<sup>4</sup>The survival probability for electron neutrinos considered in the original analysis [A<sup>+</sup>20c] is  $P_{ee} = 0.543$ , that corresponds to the averaged probability for the total solar neutrino spectrum [B<sup>+</sup>14].



**Figure 5.7:** Differential background rates for solar  $^8\text{B}$  neutrino-electron elastic scattering in the  $\pm 1\sigma$  ROI (1.0% energy resolution at  $Q_{\beta\beta}$ ), considering a neutrino flux of  $5.79 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$ . Left: mixture of Z and W boson exchange (NC+CC). Right: pure Z boson exchange (NC).

Table 5.3 summarizes the decay information of the caesium isotopes produced in the neutrino capture by different xenon isotopes. The caesium isotopes that decay through  $\beta^+$  emission will produce a clear signature with the energy deposition by the  $\beta^+$  itself and the 511 keV  $\gamma$ -rays from  $\beta^+$  annihilation and will thus be vetoed with high efficiency.  $^{131}\text{Cs}$ ,  $^{134}\text{Cs}$  and  $^{136}\text{Cs}$  decay through  $\beta^-$  emission, but the first two isotopes have a low Q-value and will not produce a background in the ROI of this search. Therefore, the only possible background from neutrino capture on xenon comes from the  $\beta$  decay of  $^{136}\text{Cs}$  into  $^{136}\text{Ba}$  with a half-life of 13.16 days and  $Q=2548.2 \text{ keV}$ . The  $\beta$  emission is always accompanied by at least two  $\gamma$ -rays that will most likely produce a multiple scatter in the xenon bulk.

Estimates from the nEXO experiment predict a rate of 0.3 events per tonne of 80% enriched  $^{136}\text{Xe}$  per year, or 0.8 counts per tonne in a 1000 day exposure even if no single scatter selection is applied [A<sup>+</sup>18g]. Considering the full  $\beta+\gamma$  spectrum from the decay of  $^{136}\text{Cs}$ , about 5.7% of the decays will populate the ROI for an energy resolution of 2% FWHM [EE14] – very similar to the one considered in this analysis. Assuming the same event fraction and capture rate for LZ, a corresponding rate of 0.13 counts in the ROI for a 967 kg volume of natural xenon and an exposure of 1000 days is obtained. However, the decay of  $^{136}\text{Cs}$  always produces a cascade of  $\gamma$ -rays along with the  $\beta$  emission and is expected to be completely vetoed by multiple scatter rejection. A simulation of  $10^7$  decays of  $^{136}\text{Cs}$  in the bulk yielded no single scatter events within the ROI. Furthermore, some large fraction of  $^{136}\text{Cs}$  is expected to be removed from the liquid xenon bulk by the purification system over several circulation cycles within its 13.16 day half-life. This background is therefore considered negligible for this analysis.

### 5.3.5.3 $2\nu\beta\beta$ decay of $^{136}\text{Xe}$

As mentioned before, the two-neutrino decay mode of  $^{136}\text{Xe}$  shares the same Q-value as the neutrinoless decay mode,  $Q_{\beta\beta}$ , and as such it can produce a background in the energy ROI due to the finite energy resolution of the detector. These two decays also share the same position

**Table 5.3:** Backgrounds induced by neutrino capture on xenon isotopes.

Isotope	abundance [%]	$\sigma_{\text{CC}} [10^{-42} \text{ cm}^2]$	${}^A\text{Cs}$	$T_{1/2} ({}^A\text{Cs})$	$Q_{\beta} ({}^A\text{Cs})$	Decay mode (BR)
${}^{124}\text{Xe}$	0.09	-	${}^{124}\text{Cs}$	30.8 s	5.9	$\beta^+$ (100%)
${}^{126}\text{Xe}$	0.09	-	${}^{126}\text{Cs}$	1.68 min	4.8	$\beta^+$ (100%)
${}^{128}\text{Xe}$	1.92	1.22	${}^{128}\text{Cs}$	3.66 min	3.9	$\beta^+$ (100%)
${}^{129}\text{Xe}$	26.44	21.1	${}^{129}\text{Cs}$	32 h	1.2	$\beta^+$ (100%)
${}^{130}\text{Xe}$	4.08	2.15	${}^{130}\text{Cs}$	29 min	2.98	$\beta^+$ (100%)
${}^{131}\text{Xe}$	21.18	30.9	${}^{131}\text{Cs}$	9.7 d	0.35	$\beta^-$ (100%)
${}^{132}\text{Xe}$	26.89	3.63	${}^{132}\text{Cs}$	6.48 d	2.1	$\beta^+$ (100%)
${}^{134}\text{Xe}$	10.44	5.97	${}^{134}\text{Cs}$	2.1 a	2.1	$\beta^-$ (100%)
${}^{136}\text{Xe}$	8.87	12.1	${}^{136}\text{Cs}$	13.16 d	2.55	$\beta^-$ (100%)

distribution in the TPC, as they originate from the same isotope. Fortunately, the energy spectrum of the  $2\nu\beta\beta$  decay mode falls off sharply at the end point (see Figure 3.10), making this background less severe. Using the  $2\nu\beta\beta$  spectrum from Reference [KI12], for 1.0% energy resolution this results in  $6.9 \times 10^{-6}$  events per kg of  ${}^{136}\text{Xe}$  per year in the  $\pm 1\sigma$  ROI, or less than 0.01 background events in 1000 days and in the inner 967 kg volume. For the remaining analysis, the background caused by  $2\nu\beta\beta$  decay is simulated using DECAY0 [PTZ00]. This event generator produces the initial kinematics of the particles involved in both the  $2\nu\beta\beta$  and  $0\nu\beta\beta$  decay that are then simulated using BACCARAT.

## 5.4 Sensitivity Projection

The median expected sensitivity to  $0\nu\beta\beta$  decay is here defined as the median 90% confidence level (CL) upper limit that would be obtained by an ensemble of experiments with the expected background and no true signal [FC98]. A preliminary sensitivity study was performed using a Feldman-Cousins cut-and-count limit projection [RLC05] to determine the median expected sensitivity assuming 1000 days of detector live-time [Brá18]. The same event selection cuts used to characterize the backgrounds on Table 5.1 were used in this preliminary analysis:

1. The 967 kg inner volume depicted on Figure 5.2, with dimensions  $26 < z < 96$  cm and  $r < 39$  cm. This cylindrical volume corresponds to the most sensitive region of the TPC and yields 86 kg of  ${}^{136}\text{Xe}$ , corresponding to an exposure of 233 kg-years.
2. The ROI is defined as a  $1\sigma$  window on either side of  $Q_{\beta\beta}$ , corresponding to the energy interval  $2409 < E < 2507$  keV for an energy resolution ( $\sigma/E$ ) of 1% at  $Q_{\beta\beta}$ . The corresponding signal acceptance for this ROI definition is 68.2%.
3. Rejection of multiple scatter events by requiring that multiple interaction vertices are separated by less than 3 mm in the vertical direction. The  $0\nu\beta\beta$  decay signal acceptance for this selection alone is estimated to be 80% (see discussion below).
4. The vetoes exclude events with an energy deposition larger than 100 keV in the outer

detector (skin detector) within  $1\ \mu\text{s}$  ( $800\ \mu\text{s}$ ) of an interaction in the TPC.

Under these assumptions, and considering the background sources described above and summarized in Table 5.1, a 90% CL upper-limit of  $\mu_s = 10$  signal events is obtained. The corresponding half-life to this decay can then be calculated as

$$T_{1/2}^{0\nu} = \ln 2 \frac{m_{\text{Xe}} \eta_{\text{Xe}136} N_A}{M_{\text{Xe}136}} \frac{\epsilon}{\mu_s} t, \quad (5.4)$$

where  $m_{\text{Xe}} = 967\ \text{kg}$  is the mass of natural xenon in the inner volume,  $\eta_{\text{Xe}136}$  is the natural abundance of  $^{136}\text{Xe}$ ,  $M_{\text{Xe}136}$  is the molar mass of  $^{136}\text{Xe}$ ,  $N_A$  is the Avogadro constant and  $\epsilon \approx 0.5456$  is the signal efficiency from the ROI and SS selections. Using Equation 5.4, a median 90% confidence-level sensitivity to the  $0\nu\beta\beta$  decay half-life of  $^{136}\text{Xe}$  of  $4.1 \times 10^{25}$  years is obtained. This estimate does not include systematic uncertainties.

A more detailed sensitivity analysis was performed by constructing a multidimensional background model that accounts for each of the sources discussed in section 5.3. On this analysis, the 90% CL upper limit on the number of signal events is determined using the profile likelihood ratio (PLR) method, utilising the asymptotic one-sided profile likelihood test statistic [CCGV11]. This multidimensional background analysis takes advantage of the precise energy and 3-dimensional position reconstruction capabilities of the LZ detector to enhance the sensitivity far beyond what a simple cut-and-count analysis can achieve [ $A^{+20d}$ ,  $A^{+20c}$ ].

Each background is modelled with a probability density function (PDF)  $P(E, r^2, z)$  built using the energy deposited ( $E$ ), the depth ( $z$ ) and radial position ( $r^2$ ) of the interaction in the TPC. Due to the lack of statistics, the PDFs are approximated by combining the marginal energy and position distributions  $P(E, r^2, z) = P(E)P(r^2, z)$ , which has been verified to be a reasonable approximation for the energy range being considered (see discussion below). The energy and position distributions are obtained from extensive simulations of the contamination on detector materials,  $\gamma$ -rays from the laboratory rock and internal radon backgrounds. The internal  $^{137}\text{Xe}$ , neutrino-induced and  $2\nu\beta\beta$  decay backgrounds have uniform position distributions and are characterised by their energy spectra alone. The  $^{136}\text{Xe}$   $2\nu\beta\beta$  decay spectrum is obtained from Reference [KI12] and the  $^{137}\text{Xe}$   $\beta$ -decay spectrum is obtained from Reference [Mou17].

The expected signal from the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  is characterized by a uniform distribution of energy depositions in the TPC with an energy profile given by a Gaussian distribution centred at  $Q_{\beta\beta}$  and with standard deviation defined by the energy resolution of the detector, considered here to be 1% ( $\sigma/E$ ). The signal was simulated using DECAY0 [PTZ00] in order to compute the signal efficiency after applying the event selection criteria. It is estimated that 20% of  $0\nu\beta\beta$  decay events will be rejected by the 3 mm vertex separation threshold in  $z$  by producing a multiple scatter due to Bremsstrahlung emission, leading to a signal efficiency of 80%.

The signal and background PDFs used in the PLR analysis span over a wider range of energies and positions when compared to the event selections used in the cut-and-count analysis. This extended range of observables results in a better characterization of the backgrounds closer to the TPC walls and around the energy ROI, fitting the background rates more precisely and

leading to stronger constraints on the background rates in the innermost regions of the TPC. Furthermore, the full shape of the position distribution is expected to provide some additional discrimination power between some characteristically non-uniform background distributions and the uniform signal-like distributions.

The extended fiducial volume used in this analysis contains 5.6 tonnes of natural xenon, corresponding to a 1360 kg·years exposure of  $^{136}\text{Xe}$  assuming 1000 days of detector live-time. This cylindrical volume is represented in Figure 5.2 and its boundaries are defined at 4 cm from the TPC walls, 2 cm above the cathode grid and 13 cm below the gate grid. The dimensions of the fiducial volume were chosen in order to remove most of the backgrounds from the grids and the PTFE walls, mostly originated by  $\beta$ -emitting plated  $^{222}\text{Rn}$  daughters.

The energy range used in the PDFs is  $2000 < E < 2700$  keV, providing an extended coverage of the energy spectra of the backgrounds near  $Q_{\beta\beta}$ , most notably the 2615 keV  $\gamma$ -ray line of  $^{208}\text{Tl}$ , several  $\gamma$ -ray lines of  $^{214}\text{Bi}$  and the  $2\nu\beta\beta$  decay of  $^{136}\text{Xe}$ . Similarly to the position distributions, extending the energy distributions used in this analysis leads to a more precise fitting of the backgrounds near the ROI and to stronger constraints. All energy distributions considered in this analysis are smeared using the LUX energy resolution function [A<sup>+</sup>17c] that has been properly scaled to ensure a resolution of  $\sigma/E = 1\%$  at  $Q_{\beta\beta}$ .

The unbinned extended likelihood function  $L(\mu_s, \{\mu_b\})$ , defined in Equation 5.5, combines the signal PDF  $P_s$ , the background PDFs  $P_b^i$ , and the systematic uncertainties  $\sigma_b^i$  on the expected background rates  $a_b^i$  [A<sup>+</sup>20c],

$$L(\mu_s, \{\mu_b\}) = \left[ \mu_s P_s(E, r^2, z) + \sum_{i=1}^{n_b} \mu_b^i P_b^i(E, r^2, z) \right] \prod_{j=1}^{n_b} g(a_b^j, \sigma_b^j), \quad (5.5)$$

where  $\mu_s$  is the number of signal events (which explicitly depends on the decay half-life, see Equation 5.4) and  $\mu_b^i$  is the number of events for the  $i$ -th background source. These terms are the floating parameters of the function, while the background rates  $a_b^j$  and corresponding systematic uncertainties  $\sigma_b^j$  are treated as nuisance parameters combined in a set of Gaussian constraint terms  $g(a_b^j, \sigma_b^j)$ .

The relative systematic uncertainties on the background rates of the sources included in the likelihood are summarized on Table 5.4. All these uncertainties are expected to be conservative and will be estimated once these backgrounds are measured from acquired data. No other uncertainties are assumed in this analysis since they are expected to be subdominant.  $\gamma$ -rays from the cavern walls produce the largest single component contributor to the background model, and thus the uncertainties of the measured  $^{238}\text{U}$  and  $^{232}\text{Th}$  activities will have the largest impact on the sensitivity calculations. The uncertainties associated with the backgrounds from detector materials are conservatively set at 30%. The high uncertainty associated with the  $^{222}\text{Rn}$  background rate reflects the expected range determined from emanation measurements of individual materials and subsystems [A<sup>+</sup>20h, A<sup>+</sup>20d]. The internal  $^{137}\text{Xe}$  background and  $^{214}\text{Bi}$  cathode background rates are not known and are assigned a large uncertainty. On the other hand, the uncertainties for  $^{136}\text{Xe}$   $2\nu\beta\beta$  decay and  $^8\text{B}$  neutrino-induced backgrounds are constrained by

**Table 5.4:** Relative uncertainties on the background rates for the background components assumed in the PLR analysis [A<sup>+</sup>20c]. The uncertainties are treated in the unbinned extended likelihood function in Equation 5.5 as nuisance parameters.

Background	$\sigma/N$
$^{238}\text{U}$ (Detector)	30%
$^{232}\text{Th}$ (Detector)	30%
$^{60}\text{Co}$ (Detector)	30%
$^{238}\text{U}$ (Cavern)	50%
$^{232}\text{Th}$ (Cavern)	30%
$^{214}\text{Bi}$ (Cathode)	50%
$^{222}\text{Rn}$ (Internal)	50%
$^{137}\text{Xe}$ (Internal)	50%
$^{136}\text{Xe}$ $2\nu\beta\beta$	5%
$^8\text{B}$ solar $\nu$	5%


the measured half-life and uncertainties of the neutrino flux, respectively, resulting in lower uncertainties.

#### 5.4.1 Sensitivity Projection with Natural Abundance of $^{136}\text{Xe}$

Figure 5.8 shows the 90% CL sensitivity of LZ to the  $0\nu\beta\beta$  decay half-life of  $^{136}\text{Xe}$  as a function of the live-time of the detector. After 1000 days of data taking, LZ will reach a median sensitivity of  $1.06 \times 10^{26}$  years, a result comparable to the current best experimental limit from KamLAND-Zen [G<sup>+</sup>16]. This sensitivity corresponds to an upper limit on the number of signal events in the 5.6 tonne FV of  $\mu_s = 31.6$  (from Equation 5.4) above the expected background level for an exposure of 1000 days.

Figure 5.9a displays the impact of the energy resolution at  $Q_{\beta\beta}$  on the  $0\nu\beta\beta$  decay half-life sensitivity. The 1.0% energy resolution assumed for this analysis is conservative, as other dual-phase TPC detectors have demonstrated that a 0.8% resolution at these energies is achievable [A<sup>+</sup>20j]. A slight variation of the energy resolution within the displayed values would result in a variation of the sensitivity of only  $\pm 10\%$ . This is mainly due to the impossibility of rejecting the 2448 keV  $\gamma$ -ray line from  $^{214}\text{Bi}$  with energy resolution alone. However, the 2615 keV  $\gamma$ -ray line from  $^{208}\text{Tl}$  would impact the sensitivity significantly if the energy resolution was 2.0% or larger.

Figure 5.9b displays the half-life sensitivity as a function of the minimal vertical vertex separation. Multiple scatter events are assumed to be rejected based on  $z$  separation only. The multiple interactions at different depths in the TPC will produce multiple S2 pulses that are used to reconstruct the  $z$  vertex positions of the energy depositions. The 3 mm vertical vertex separation cut applied to the backgrounds will exclude the majority of the multiple scatter events from  $\gamma$ -rays but will also exclude a fraction of signal events due to Bremsstrahlung emission. At lower values of the vertex separation, a larger fraction of the signal is rejected and the sensitivity



**Figure 5.8:** LZ projected sensitivity to the half-life of the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  as a function of detector live time for a natural abundance of  $^{136}\text{Xe}$  (solid black line) and for a dedicated run with 90% isotopic enrichment (dashed black line). The light green shaded band represents a  $\pm 1\sigma$  statistical uncertainty on the sensitivity. The limits set by EXO-200 [A<sup>+</sup>19g] (orange full) and KamLAND-Zen [G<sup>+</sup>16] (purple full) are also shown, along with the respective projected sensitivities (dashed). Figure from Reference [A<sup>+</sup>20c].

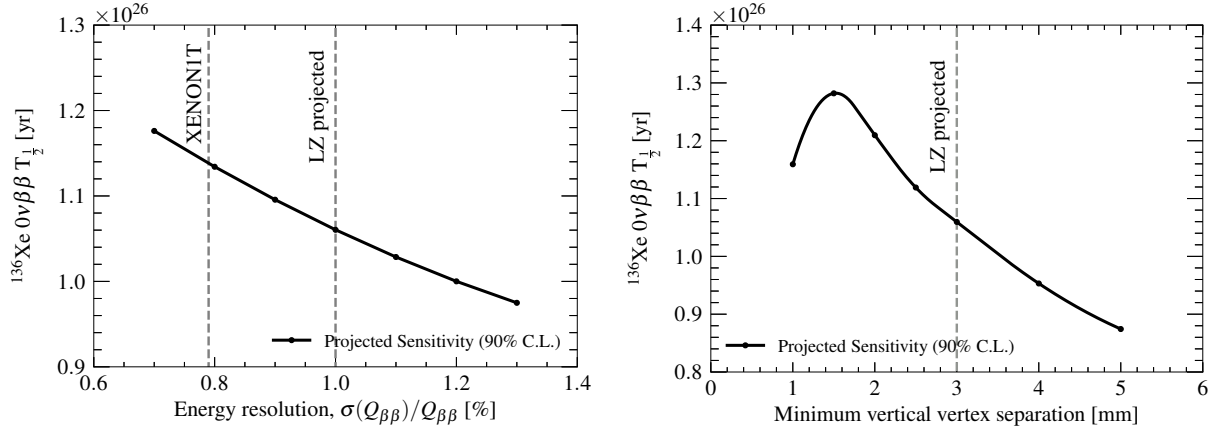
drops significantly. For larger  $z$  vertex separation values, more backgrounds will survive this cut and the sensitivity will also drop considerably.

Figure 5.10 shows the expected 90% exclusion sensitivity limits on  $\langle m_{\beta\beta} \rangle$ , estimated at 53–164 meV using Equation 3.54 for a half-life sensitivity of  $1.06 \times 10^{26}$  years. The uncertainty on the value of  $\langle m_{\beta\beta} \rangle$  is due to the range of values of the nuclear matrix elements  $\mathcal{M}^{0\nu}$  for different models. The upper and lower limits of  $\langle m_{\beta\beta} \rangle$  are obtained for the smallest and largest  $\mathcal{M}^{0\nu}$  values in literature, namely  $\mathcal{M}_{\text{QRPA}}^{0\nu} = 1.55$  [ME13] and  $\mathcal{M}_{\text{EDF}}^{0\nu} = 4.77$  [VRE13]. The value for the phase space factor considered here is  $G^{0\nu} = 3.84 \times 10^{-14} \text{ yr}^{-1}$  for an unquenched axial-vector coupling constant of  $g_A = 1.27$  [KI12]. The sensitivity to the normal and inverted neutrino mass hierarchy scenarios is also presented, with the results from this analysis almost reaching the upper bound of the inverted hierarchy parameter phase-space. The right-hand side of Figure 5.10 shows that  $^{136}\text{Xe}$  provides the strongest constraints on  $\langle m_{\beta\beta} \rangle$  so far [DPR19].

#### 5.4.2 Projection with 90% $^{136}\text{Xe}$ Enrichment

A dedicated  $0\nu\beta\beta$  decay search could be performed with LZ after the WIMP search run with no changes to the detector except for the deployment of 90%  $^{136}\text{Xe}$  enriched xenon into the TPC in order to increase the isotopic abundance by a factor of 10 [A<sup>+</sup>20c]. Considering the same 5.6 tonnes fiducial volume used in the analysis with natural abundance, a dedicated 1000 day run with 90% enrichment would yield an exposure of  $13.8 \times 10^3 \text{ kg}\cdot\text{years}$  of  $^{136}\text{Xe}$ . No additional improvements on detector performance over the main WIMP search run and non-enriched  $0\nu\beta\beta$  decay sensitivity analysis are considered here [A<sup>+</sup>20c].

The only changes in the background model for the dedicated 90%  $^{136}\text{Xe}$  run in comparison with



(a) LZ projected sensitivity as a function of the energy resolution at  $Q_{\beta\beta}$ .


(b) LZ projected sensitivity as a function of the minimum vertex separation in depth.

**Figure 5.9:** 90% CL sensitivity to the  $0\nu\beta\beta$  decay half-life of  $^{136}\text{Xe}$  as a function of the energy resolution at  $Q_{\beta\beta}$  (a) and the minimum vertex separation in depth (b), for a 1000 live-days run with LZ. The sensitivity for an energy resolution of 1% and vertical vertex separation of 3 mm, the values assumed in this analysis, are represented by the vertical dashed lines labelled “LZ projected”. For comparison, the projected LZ sensitivity assuming the energy resolution recently measured in XENON1T [A<sup>+</sup>20j] is also indicated on plot (a). Figure from Reference [A<sup>+</sup>20c].

the non-enriched scenario are the backgrounds that involve the enriched isotope itself, namely the neutron-induced  $^{137}\text{Xe}$  and the  $2\nu\beta\beta$  decay of  $^{136}\text{Xe}$ . The neutrino-induced backgrounds, mainly the charged-current neutrino-nucleus interactions with  $^{136}\text{Xe}$ , remain negligible even with 10 times more isotopic abundance.

The background rate associated with the  $2\nu\beta\beta$  decay of  $^{136}\text{Xe}$  was increased by a factor of 10 in this sensitivity analysis, but remains sub-dominant relatively to the Davis cavern backgrounds and radioactivity from detector materials, assumed here to be identical to the non-enriched scenario. Conversely, an increase of the neutron-induced  $^{137}\text{Xe}$  background by a factor of 10 would impact the sensitivity in a non-negligible way. In order to mitigate this background for a dedicated  $0\nu\beta\beta$  decay run, additional neutron shielding would have to be deployed around the xenon purification system and circulation lines outside the water tank. This shielding is not relevant for the WIMP search run since the activation of xenon isotopes does not constitute a relevant background for that run. An attenuation of the thermal neutron flux by two orders of magnitude is achievable with a layer of high-density polyethylene with a thickness of 20 cm [LRK<sup>+</sup>06, TRK10, Viv10], resulting in a reduction of the  $^{137}\text{Xe}$  background by the same amount. A 10-fold reduction of the thermal neutron flux from this shielding is considered in this sensitivity analysis as a conservative estimate, leading to a background rate of 0.28 events in the ROI and inner 967 kg volume over the 1000 day exposure, the same rate as in the non-enriched scenario. Even if this background is not properly mitigated with the addition of shielding, the estimated sensitivity would only decrease by 9%, a variation comparable to a  $+1\sigma$  increase in the background from the  $\gamma$ -rays from the Davis cavern rock. The muon-induced neutron absorption on  $^{136}\text{Xe}$  is considered to be negligible for both enrichment scenarios due to the same arguments





**Figure 5.10:** LZ projected exclusion sensitivity to  $\langle m_{\beta\beta} \rangle$  and the neutrino mass hierarchy as a function of the lightest neutrino mass. The width of the green band is the result of the uncertainty in the nuclear matrix elements [ME13, VRE13]. The red and blue contours represent the allowed parameter space for the inverted hierarchy and normal hierarchy neutrino mass scenarios, respectively, as well as the respective  $\pm 1\sigma$  contours (light colors) [A<sup>+</sup>18g]. This projection for LZ matches the current best experimental limit from KamLAND-Zen [G<sup>+</sup>16]. The inlay on the right-hand side of the plot displays the current best exclusion limits to  $\langle m_{\beta\beta} \rangle$  for different  $2\nu\beta\beta$  isotopes (solid arrows) and the projected limit of LZ for a dedicated run with 90% enriched  $^{136}\text{Xe}$  (see Section 5.4.2). The single pointed end of the arrows represents the lower-bound exclusion limit, while the length of the arrow represents the respective uncertainty. Figure from Reference [A<sup>+</sup>20c].

used in Section 5.3.4. Under the same assumptions used in the non-enriched scenario, the rate associated with the muon-induced background is estimated to be  $< 0.02$  events in the ROI and inner 967 kg volume for a 1000 day run. The overall  $^{137}\text{Xe}$  background rate was set at 0.3 events in the ROI and inner 967 kg volume for a 1000 day run with enriched xenon.

The 90% CL sensitivity of LZ to the  $0\nu\beta\beta$  decay half-life of  $^{136}\text{Xe}$  for a 1000 day run with enriched xenon is  $1.06 \times 10^{27}$  years, roughly 10 times higher than for the non-enriched scenario. The dashed black curve in Figure 5.8 represents this sensitivity as a function of detector live-time. The sensitivity to  $\langle m_{\beta\beta} \rangle$  would also improve by the same factor, covering the full parameter space for the inverted neutrino mass hierarchy.

## 5.5 Conclusions

Table 5.5 displays the summary of the results obtained in this work. The baseline assumptions presented in this analysis for the energy resolution, position reconstruction, background discrimination and  $^{136}\text{Xe}$  exposure of LZ indicate that a competitive search for the  $0\nu\beta\beta$  decay with a projected median 90% CL exclusion sensitivity of  $1.06 \times 10^{26}$  years for the half-life of  $^{136}\text{Xe}$ , and a sensitivity to  $\langle m_{\beta\beta} \rangle$  of 53–164 meV, comparable to the current best limits [G<sup>+</sup>16], is achievable while maintaining the main physics goal of searching for WIMP dark matter and without any modification to the nominal detector operations.

**Table 5.5:** Summary of the assumptions used and results obtained in this preliminary study of the sensitivity of LZ to the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$ .

Detector parameters	E resolution at $Q_{\beta\beta}$	1% $\sigma/E$
	ROI	$2433.3 < E < 2482.4$ keV ( $Q_{\beta\beta} \pm 1\sigma$ ) $2000.0 < E < 2700.0$ keV (PLR)
	Minimal vertex separation	3 mm in $z$ only
	Outer detector threshold	100 keV
	(coincidence window)	1 $\mu\text{s}$
	Skin detector threshold	100 keV
	(coincidence window)	800 $\mu\text{s}$
	967 kg inner volume	$r < 39$ cm and $26 < z < 96$ cm
	5.6 t fiducial volume (PLR)	$r < 68.8$ cm and $2 < z < 132.6$ cm
	Sensitivity ( $^{nat}\text{Xe}$ )	$^{136}\text{Xe}$ $0\nu\beta\beta$ $T_{1/2}$
$\langle m_{\beta\beta} \rangle$		53–164 meV (90% CL)
Sensitivity (90% $^{136}\text{Xe}$ )	$^{136}\text{Xe}$ $0\nu\beta\beta$ $T_{1/2}$	$1.06 \times 10^{27}$ years (90% CL)
	$\langle m_{\beta\beta} \rangle$	5.3–16.4 meV (90% CL)

A dedicated post WIMP search run with 90%  $^{136}\text{Xe}$  enrichment would be feasible with no improvements in detector parameters or operations. With proper mitigation of the  $^{137}\text{Xe}$  neutron-induced background, by placing additional shielding materials on the xenon purification system, LZ could reach a 90% CL sensitivity of  $1.06 \times 10^{27}$  years for the half-life of  $^{136}\text{Xe}$ , probing the complete parameter space for the inverted neutrino mass hierarchy.

A comprehensive background model was built using precise detector and physics simulations, radioactive assays of detector construction materials and *in situ*  $\gamma$ -ray flux measurements. The combined effect of using an extended phase-space in both energy and position with the profile likelihood analysis is an improvement of the sensitivity result by a factor of two when compared to a simple cut-and-count analysis. In order to improve the background estimates and associated uncertainties, an intensive background characterization campaign will be performed in the commissioning and first science run of LZ. During these stages the most relevant backgrounds will be measured with high precision with the TPC and outer detector system, most notably the contributions from internal radon and  $\gamma$ -rays from detector materials and cavern walls.

# Chapter 6

## Pulse Classification in LZ

The efficiency of a data analysis framework in identifying and interpreting the information recorded by a detector ultimately dictates the scientific performance of an experiment, especially one tailored to search for rare interactions. Different particle interactions in the detector will produce distinct event signatures that can be used to identify the physics processes that took place. For all events recorded in the detector, the analysis framework has to guarantee that all relevant information is available to the posterior physics analysis in order to distinguish the events of interest from the dominant backgrounds. Therefore, the overall goal of an analysis framework is to extract useful information from the features of the output signals of a detector. This is achieved by deploying specialized computational algorithms that process the raw data and retrieve a set of reduced quantities (RQs) that can then be used to fully describe the interactions on high-level physics analyses.

LZ will be the most sensitive detector for searching dark matter in the form of WIMPs. Moreover, it will be sensitive to several other interesting physical processes with a wide range of energies and distinct interaction mechanisms: neutrino-electron scattering, coherent neutrino-nucleus scattering, double beta decay, just to name a few. The signals resulting from these different interactions will need to be fully characterized in order to explore the total potential of LZ. For any triggered event, the data acquisition (DAQ) system records the voltage response of each PMT channel of the liquid xenon (LXe) TPC, LXe skin and outer detector (OD) systems, that is above a certain threshold, into a pulse timeline with 5 ms duration [LZ 16]. For the TPC and OD, the PMT signal is recorded in two separate gain stages, labeled high-gain (HG) and low-gain (LG), with a  $\times 10$  gain factor between the two [M<sup>+</sup>17, A<sup>+</sup>20g]. This allows LZ to be sensitive to low-energy physics with the higher gain data, *e.g.*, for WIMP search at the  $\mathcal{O}(10\text{keV})$  scale, while preserving the high energy response for physics beyond the WIMP search energy scale, *e.g.*, for <sup>136</sup>Xe neutrinoless double beta decay searches [A<sup>+</sup>20c]. The digitized event

timeline is the raw data of the detector and the input of the analysis framework, containing all the information about recorded events. Additional information regarding the conditions of the detector – field strength, gains, light collection maps, etc. – is stored in a database and accessed by the analysis framework to complement the required information about the detector state at the acquisition time.

The LZ experiment will produce copious amounts of event data along several years of operation. The data volume produced by the LZ detector can reach up to 1.4 PB per year, accounting for both WIMP search data and calibrations [M<sup>+</sup>17]. The analysis framework of LZ will have to balance efficiency with speed in order to return valuable information of the LZ science runs in a manageable time scale. The processing and storage of LZ data will be performed in two separate data centers: one based in the US and the other in the UK. The US data center (USDC) will use the infrastructure of the National Energy Research Scientific Computing (NERSC) center at LBNL, while the UK data center (UKDC) will use the GridPP infrastructure [A<sup>+</sup>20g] of the Imperial College London. This approach allows for distributed processing, creates storage redundancy and shares computational resources more efficiently among LZ collaborators. The USDC receives the raw data from the SURF surface staging computer and mirrors the data to the UKDC. Both data centers will have enough computing power and storage to house the raw data and run analysis and processing software in near real-time.

In this chapter, the inner workings of the analysis framework of LZ are discussed, from raw data to physics. The focus will be placed in the classification of pulses from the light signals produced within the detector, showcasing several tested algorithms that perform this task with great efficiency. These algorithms range from heuristic classifiers with simple decision chain logic to machine-learning algorithms as potential solutions to the limitations of *ad hoc* methods.

## 6.1 LZap - the LZ Analysis Programme

The goal of the analysis framework of LZ, dubbed *LZ Analysis Programme* (LZap), is to go from raw PMT data to a data format that can be used for physics analysis. The framework consists of a series of software modules, each containing a set of specialized algorithms, that process the raw data and return a complete characterization of an event in the form of a set of RQs. LZ will implement a similar modular structure as the one developed previously for the LUX experiment [M<sup>+</sup>17, A<sup>+</sup>20g]. The reason for this is twofold: the LUX analysis framework has been extensively tested and has proven to be successful at analysing the data collected, and the working principle and technology of the LZ detector will be a scaled-up copy of the LUX detector, making some of the features required in one analysis framework the same as the ones required for the other. Furthermore, the data to be analysed will be similar to the data collected by LUX, as both detectors use arrays of PMTs for collecting the primary and secondary scintillation of xenon in a dual-phase TPC. The main differences, aside from the absolute scale of the detectors and the larger number of light sensors used in LZ, will be the addition of a fully instrumented OD and skin volumes, that will play an important role in the event characterization and background mitigation.

The LZ detector is currently in the final stages of installation in the Davis cavern and is expected to begin operations by early 2021 [A<sup>+</sup>20d, A<sup>+</sup>20g]. Before real LZ data is available, the analysis framework modules are developed and tested using data from detector simulations [A<sup>+</sup>21]. The simulation data is obtained using a Geant4-based simulation package especially developed for the LZ detector and based on prior LUX simulation packages (see Section 4.4) [A<sup>+</sup>03, A<sup>+</sup>12a, A<sup>+</sup>21]. In addition to physics simulations that account for the full detector geometry, construction materials and nuclear and atomic physics processes, the simulated data will also include the detector electronics response (DER), with the inclusion of pulse shaping and filtering effects from acquisition electronics and realistic electronic noise, in order to produce a good representation of the real data.

To ensure that the analysis framework is ready for the first science run of LZ, a series of mock data challenges (MDCs) were planned with the goal of developing, tuning and producing analysis tools that could deliver high quality physics results from simulated data. These data challenges ran between mid 2017 and through mid 2020. The first MDC (MDC1) focused on providing the infrastructure and an end-to-end working framework. MDC2 focused on improving the quality of the simulation, generate reliable simulations of calibration sources, develop near-complete algorithms and produce limit/discovery results. The final MDC3 had the goal of fine-tune the processing framework, create advanced analysis and visualization tools, test the data centers, ramp up the realism of simulated data, generate salted<sup>1</sup> data and simulate the full physics commissioning of LZ, without Monte-Carlo truth (MCtruth) available to the analysers.

### 6.1.1 The Architecture of LZap

While the framework architecture of LUX was custom made, LZap is built within the GAUDI framework software package, an object-oriented software architecture that is used to build data processing applications for high-energy physics experiments [B<sup>+</sup>01]. The main advantage of using GAUDI is the architecture-centric nature of the software, where several instances of algorithm objects – called “modules”, for simplicity – can be instantiated and set to run in a specific order – processing chain – over a set of data objects. GAUDI also allows for the development of new data structures or analysis modules by collaborators that can be integrated into existing processing chains, ensuring that the analysis framework can be expanded and improved over the course of livetime of the experiment.

The overall data object structure handled by LZap is represented in Figure 6.1. For convenience, the data objects used by LZap are segmented into five “detector” structures, grouped by physical detector and gain mode: TPC high-gain (TPC-HG), TPC low-gain (TPC-LG), OD high-gain (OD-HG), OD low-gain (OD-LG) and Skin system. These structures contain the RQs for each of the physical detectors. Furthermore, a new “detector” data structure is created by joining the TPC-HG and TPC-LG data taking into account channel saturation in the HG data and

---

<sup>1</sup>Fake signal events added to the data stream to mitigate biases in the analysis, as an alternative to blinding the signal region.

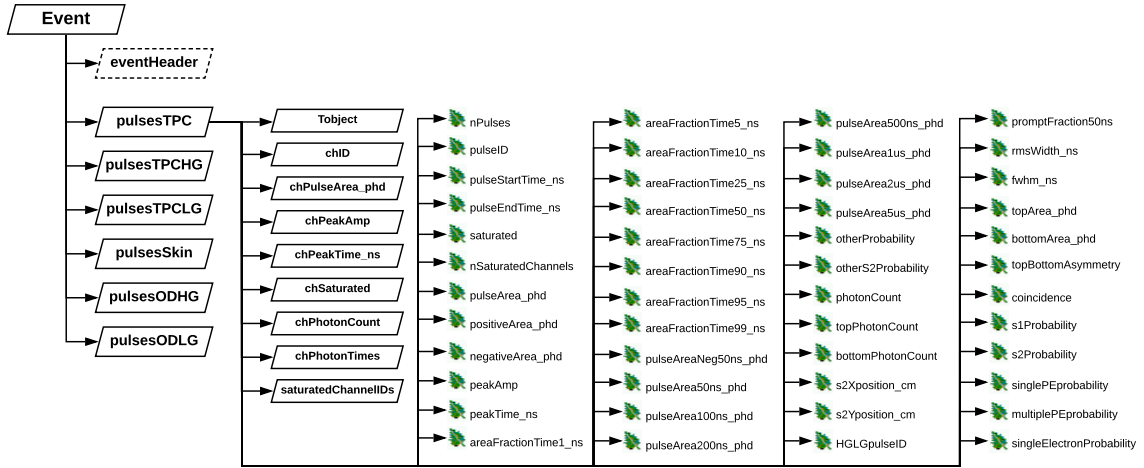


Figure 6.1: Core data structure created by LZap processing.

selecting the corresponding LG data<sup>2</sup>, thus folding the two gains into one “mixed TPC” data structure, often just referred to as simply “TPC”, that remains consistent at both low energy and high energy events.

The typical timeline of a triggered event in LZ is comprised of several digitized waveforms originated from the readout of the PMTs. The waveforms are digitized at 100M samples per second (resulting in a sample size of 10 ns) with a 14-bit resolution and a 2 V dynamic range [M<sup>+</sup>17]. The term waveform will be used as a general term to describe the raw data in the form of a time series voltage response, either per channel or over the sum of all channels. In LZ, data is only recorded in each channel when the voltage response of a given PMT is above a certain threshold [LZ 16]. This is called *pulse-only digitization* (POD) and the event timeline will be a succession of PODs positioned relative to one another with the correct acquisition times. These PODs may contain several distinct structures that are merged, adjacent or piled up in time. Several detector signals that are coincident in time may correspond to a single POD, and a POD may contain inner structures that are not necessarily univalent with respect to the original detector signals that were coincident in time.

In the context of the event timeline, a pulse is defined as a unique substructure in a POD that can be topologically distinguishable and whose boundaries in time (pulse start and end time) can be defined. A POD may contain several pulses, and the pulses in general will represent individual signals in the detector with a determined physical origin, if these are separable in time.

Figure 6.2 represents the data flow of LZap from raw event data to final RQs. At the core of LZap are two main processing chains: the *PhotonDetection* chain and the *InteractionDetection* chain, that process raw data in a pulse-based analysis and in an event-based analysis, respectively. The

<sup>2</sup>Some pulse shaping corrections are performed to account for the differences in pulse response between the two gain channels.

processed data created and used by the analysis modules is handled within a dedicated package called *PhysicsDataModel*, containing all the classes that instantiate the data objects that store transient information produced and used by the different LZap modules. These data objects are available during run-time and are part of the Transient Event Store (TES), a shared memory mechanism handled by GAUDI that allows the analysis modules to access all the available information: outputs of preceding algorithms, physics inputs or detector conditions. The data is stored in tree structures like the ones shown in Figure 6.1. The detector conditions are handled by yet another dedicated package dubbed *ConditionsDataModel* [A<sup>+</sup>20g]. The relevant information within the TES is stored by the *RootWriters* package to disk, using the persistency service of GAUDI [B<sup>+</sup>01]. The LZap files produced contain the RQs that are used in for physics analysis.

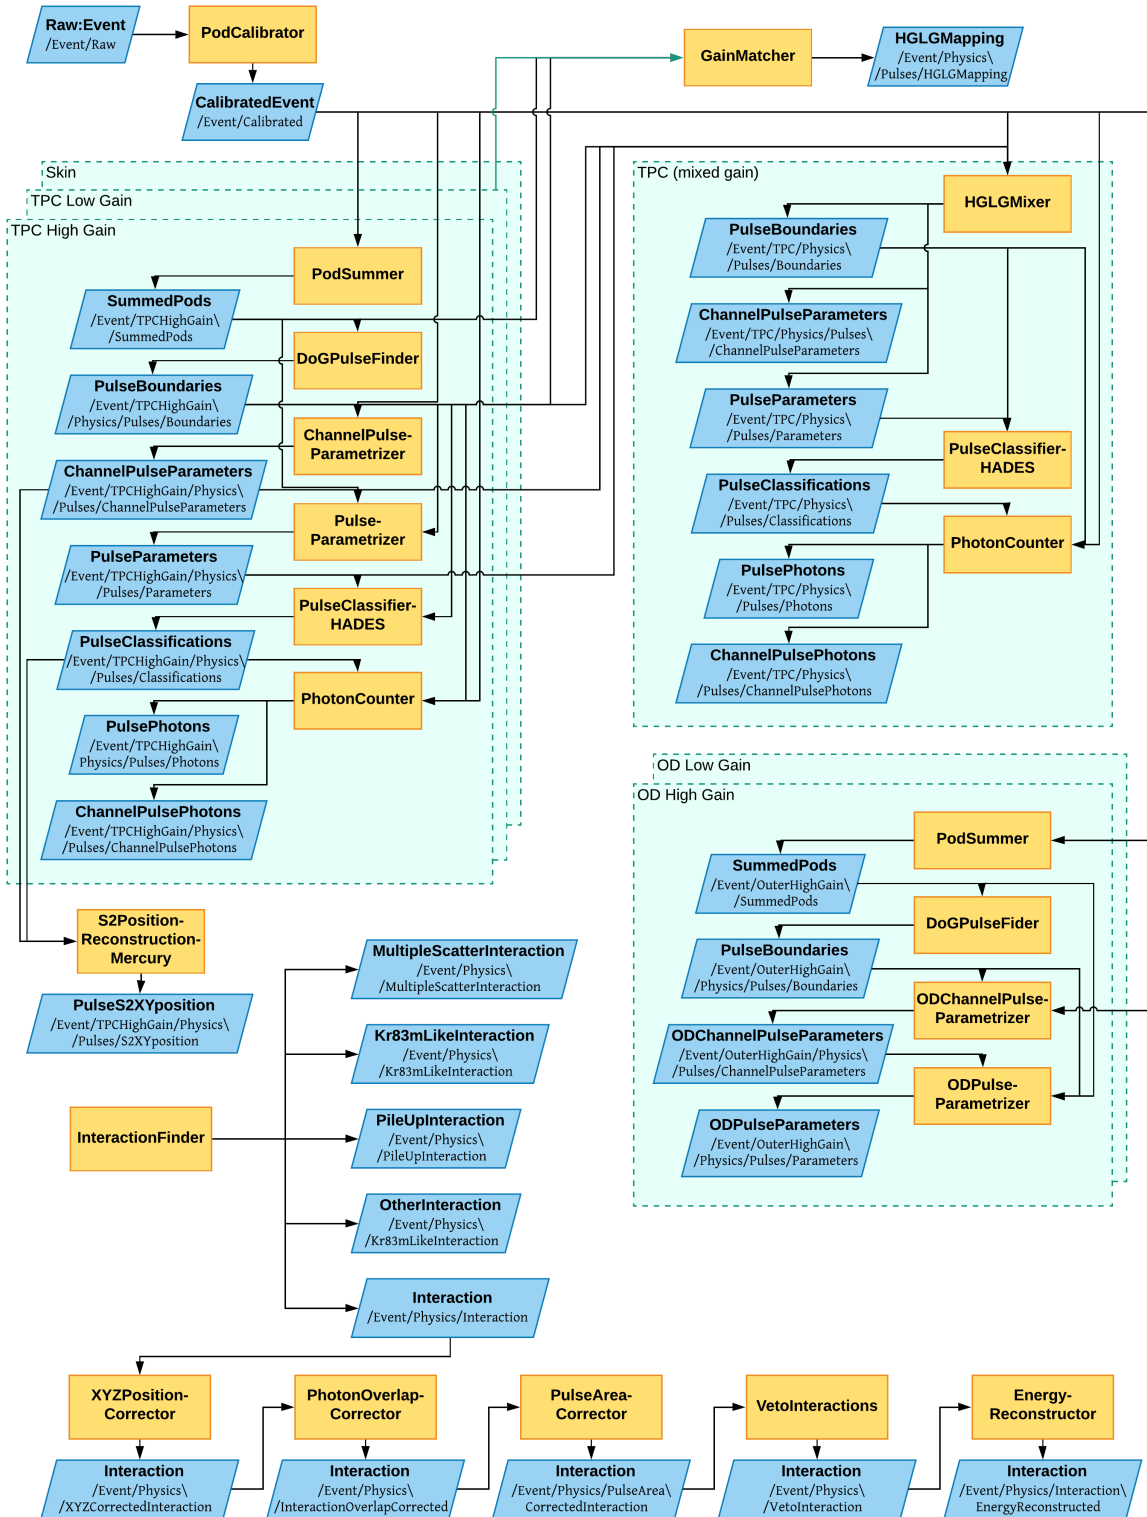
### 6.1.1.1 PhotonDetection processing chain

The *PhotonDetection* modules are responsible for the pulse-level processing. The first steps in the analysis framework of LZ are the identification, parametrization and classification of the individual pulses recorded by the detector using their distinctive features. The processing chain starts by taking the raw data and performing calibrations and baseline corrections to the PODs. Individual pulses are then isolated from the structures in the summed PODs, *i.e.*, the PODs summed over all channels of a given “detector”, by the pulse finder module. Each pulse is then parametrized in order to extract a series of RQs that can be used to characterize it (related to area, time, channel multiplicity, etc.), that are then used to classify the pulse into one of the various pulse classes considered in the analysis. At the end of the *PhotonDetection* chain a set of pulse data objects per detector per event and per channel is obtained, each with parameterization and classification RQs.

A schematic of the *PhotonDetection* chain in the context of LZap can be found in Figure 6.2, distributed across the several detector data structures. An brief overview of each *PhotonDetection* module is presented next.

**POD calibrator** (*PODCalibrator*) Converts the raw waveform data from ADC counts to units of photons detected (phd) per sample, performs time calibration of each channel and subtracts the baseline level of each POD, *i.e.*, levelling the POD waveform accounting for baseline fluctuations in the DAQ readout. The single photoelectron (SPE) size for HG and LG modes of the TPC and OD detectors is also used to scale the waveforms so that the gain factor is compensated for and the average SPE size is scaled to the proper unit of 1 photon detected (phd). Returns a *calibratedEvent* data object.

**POD summer** (*PODSummer*) Stacks all the PODs from all the channels in the same detector, adding the PODs that are overlapping or close in time into a single summed POD. Returns one or more *summedPOD* data objects per detector and per event.



**Figure 6.2:** Schematic of the data flow in LZap. Yellow rectangles represent algorithm objects (modules) and blue parallelograms represent data objects that are stored in the transient event store (TES). The aquamarine dashed rectangles encompassing some objects indicate the data flow for different detectors and their ensemble are a proxy to the *PhotonDetection* processing chain.



**Pulse finder** (*DoGPulseFinder*) Identifies pulse structures within summed PODs, returning the start and end times for each pulse structure identified. The current pulse finder algorithm in LZap is the *DoGPulseFinder*, that uses Difference of Gaussian (DoG) filters to obtain frequency-specific features in the signal data that can be used to identify transients in the waveforms. The frequencies are determined by the Gaussian widths and are tuned to the data structures that need to be partitioned. In LZap, two DoG filters are used, and are tuned to the average response of S1 and S2 pulses.

**HG-LG mixer** (*HGMixer*) Generates a new “mixed” summed waveform built with information from either the TPC-HG detector or the TPC-LG detector, toggling between the channels of the detectors when a channel in the HG is saturated. The toggling of pulse properties is only within the boundaries of each pulse.

**Channel pulse parametrizer** (*ChannelPulseParameterizer*) Calculates and returns pulse parameters (RQs) per channel, such as channel pulse height, channel pulse area, channel pulse width, among others, for the TPC-HG, TPC-LG, mixed TPC and Skin detectors.

**Pulse Parametrizer** (*PulseParameterizer*) Calculates and returns a more complete list of pulse parameters (RQs) using both the *summedPOD* and the individual channels per pulse, for each detector (TPC-HG, TPC-LG, mixed TPC and Skin) and for each event. The full list of pulse parameters used in LZap can be found in Appendix B.

**OD channel pulse parametrizer** (*ODChannelPulseParameterizer*) Similar module to the *ChannelPulseParameterizer*. Calculates and returns OD-specific parameters per pulse per channel.

**OD pulse parametrizer** (*ODPulseParameterizer*) Similar module to the *PulseParameterizer*. Calculates and returns OD-specific pulse parameters using mainly the *summedPOD* for each pulse.

**Pulse Classifier** (*PulseClassifierHADES*) Classifies each pulse as either S1, S2 or any of the considered classes (see discussion in Section 6.2.1), with the goal of ensuring the maximum classification efficiency possible. The module uses the geometrical pulse parameters calculated by the pulse parametrizer. Each pulse is classified individually and no correlation between pulses on the same event is considered.

A detailed description of the *PulseClassifierHADES* algorithm and in-depth discussion of the classification efforts in LZ can be found in Section 6.2 and is one of the main topics of this work.

**Photon counter** (*PhotonCounter*) Performs peak finding above a given threshold, by channel, and returns the spike photon count per pulse for each channel and for the *summedPOD*.

**Gain matcher** (*GainMatcher*) Maps the pulses in the HG detectors with corresponding pulses in the LG detectors. The algorithm tries to find the best subset of non-degenerated matches by maximizing the total sum of weights defined for each set of pulses overlapping in time.

### 6.1.1.2 InteractionDetection processing chain

The *InteractionDetection* modules perform the second main processing stage, where event-level information is obtained. The goal of this second stage is to characterize the event using the pulse information collected in the previous step by the *PhotonDetection* modules.

Using the information from the pulse-level analysis, S1 and S2 pulses from an event are paired in order to reconstruct the interaction vertices. The event is then classified as a *single scatter* (SS), *multiple scatter* (MS), *pileup event* (pileup), *krypton event* (Kr) or “*other*” if it does not fit the criteria of the previous categories. The XY-position reconstruction is performed using the PMT hit maps of pulses classified as S2 [S<sup>+</sup>11], while the depth of the interaction can be calculated using the time difference between the S2 signal and the S1 signal, that is proportional to the distance travelled by the drifted electrons in the liquid. Finally, the last algorithms apply corrections to the pulses and perform energy reconstruction, and store the data to disk. Once the detector begins operations, the detector parameters and PMT light responses used in the different corrections are obtained from calibration data after an initial analysis and processing, and the data is then reprocessed. A short overview of each *InteractionDetection* module is presented next.

**S2 position reconstruction** (*S2PositionReconstructorMercury*) Reconstructs the XY position of the interaction vertex from the light pattern of the associated S2 pulse. The *Mercury* algorithm [S<sup>+</sup>11, A<sup>+</sup>18e] uses light response functions (LRFs) of the PMTs in the top array to fit the PMT light pattern of an S2 pulse and estimate its position. The LRFs are stored in the *ConditionsDataModel* database.

**Interaction finder** (*InteractionFinder*) Uses the sequence of pulse classifications and some pulse parameters to classify the event into several categories: *single scatter* (SS) if the event contains only one valid S1 followed by a single valid S2; *multiple scatter* (MS) if the event contains only one S1 preceding multiple valid S2s; *pileup event* if several valid (n) S1s can be matched to several (N) S2s in a n×N mapping, as long as for each valid S2 there is a preceding valid S1; *krypton event* (Kr) if exactly two S1s with decreasing area precede a single S2; “*other*” if none of the previous conditions are met. The algorithm can also identify if some pulses are relevant to the interaction topology or if they can be ignored, such as in the case of a single

scatter event containing several pulses that are identified as S2 but none are true S2 pulses in the context of the event<sup>3</sup> (*e.g.*, multiple single electrons that overlap in time or S2 tails, see bottom plots of Figure 6.7).

**Veto interaction finder** (*VetoInteractionFinder*) Correlates pulses within the span of an event between the veto detectors and the TPC detector. Currently unused.

**XYZ position reconstruction** (*XYZPositionCorrector*) Reconstructs the depth ( $z$ -position) of the interaction from the calculated drift time. The electric field map used to convert the drift time into depth is stored in the *ConditionsDataModel* database. This module also corrects the  $xy$  positions of the S2 pulses obtained by the *S2PositionReconstructorMercury* module using the calculated interaction depth and considering the non-uniformity of the electric field.

**Photon overlap correction** (*PhotonOverlapCorrector*) Corrects for possible photon overlap in spike counting. Currently unused.

**Pulse area correction** (*PulseAreaCorrector*) Corrects the S1 and S2 pulse areas to account for detector effects on the charge and light yield of the interaction. For the S2 pulses, implements corrections to the  $z$  dependency due charge absorption by impurities in the LXe and  $xy$  variations due to non-uniform extraction field and liquid level. For the S1 pulses, implements corrections due to the  $z$  dependency caused by the photon absorption length at low xenon purity and  $xyz$  dependencies on, *e.g.*, light collection in the LXe bulk due to total internal reflection, solid angle between the event position and the PMT arrays, and PTFE reflectivity.

**Energy reconstruction** (*EnergyReconstructor*) Computes two energies, one assuming an ER interaction and one assuming a NR interaction. The ER combined energy is calculated using Equation 4.7. The NR energy is computed iteratively using the Lindhard model with biexcitonic quenching [LSS63, MYSH08].

In the end of the LZap processing chain, the data objects created by each module are stored to disk as RQ files (ROOT files). These RQs can then be used for physics analysis, providing a full representation of the events in the detector without forcing the analyser to use the raw waveform data from the detector. The full list of RQs calculated with LZap can be found in Appendix B.

---

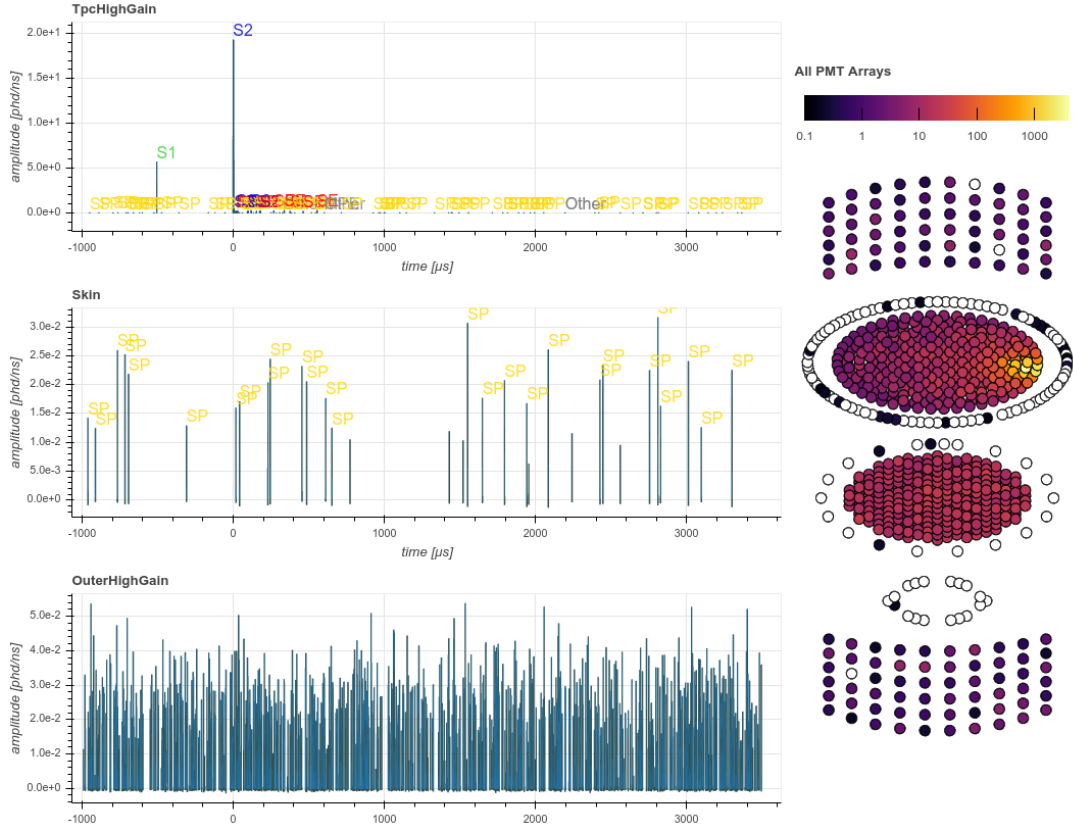
<sup>3</sup>As it will be discussed in depth on the *PulseClassifier* module Section, the classification is performed at the pulse level without any contextualization with the overall event or neighbouring pulses. This results in some pulses that are not relevant to the event being (correctly) classified as one of the many relevant classes. The *InteractionFinder* module is able to filter these pulses using the event context in order to correctly identify its topology.

A typical LZ event timeline and integrated PMT response is presented in Figure 6.3, taken from simulated data of the detector and electronics response. The three left-hand side plots show the event timelines for the TPC-HG (top left), Skin (middle left) and OD-HG (bottom left) detectors. The TPC-LG and OD-LG were omitted for simplicity. For this event, the Skin and OD-HG detectors recorded only PMT dark counts, with the dark count rate in the OD being significantly larger due to the PMTs being at room temperature. On the TPC-HG detector timeline the main S1 (green) and S2 (blue) pulses can be seen – the signature of an interaction in the active LXe volume. The S2 pulse is generally succeeded by some single electron (SE) emissions, seen on the TPC-HG as the SE pulses (red). The SPE pulses (yellow) also present in the TPC-HG are dark counts from the TPC PMTs. A spacial representation of the integrated PMT response for the event can be seen on the right side of Figure 6.3. The PMTs of the TPC (circular arrays at the center), OD (top and bottom rectangular arrays) and Skin (radial arrays adjacent to the TPC top and bottom array plus the smaller dome array at the bottom) are arranged to best convey their real spacial distributions.

## 6.2 Pulse Classification in LZap

The correct identification of interactions in LZ data is strongly dependant on the efficiency of the first modules in the chain, namely the *PulseFinder*, *PulseParametrizer* and *PulseClassifier*, the ones responsible for the identification and classification of pulse-level structures in an event. The science output of the data analysis requires that the processing framework correctly identifies the relevant pulses that make up a given event in the detector. In the context of a rare event search experiment such as LZ the pulse classification must strive for the maximum classification efficiency obtainable to guarantee that the science deliverables can be met. For example, a dark matter interaction on the LXe active target of LZ is expected to produce a low-energy nuclear recoil, that will appear in the event timeline as a single pair of S1 and S2 pulses. If the pulse classifier module fails to identify any of these pulses as being of their respective class, or if any other structure present in the timeline that is not one of these expected pulses is wrongly classified as either an S1 or an S2 pulse, the entire event can be misidentified, compromising the overall efficiency of the experiment. This is true for any type of event topology and for any energy scale, as there are many physics processes of interest in LZ extending throughout the entire energy range. Both the low and high energy interactions present challenges to the classification: at lower energies the pulse shapes may not be well defined due to the low light and charge yields, while at higher energies the large yields in both scintillation and ionization result in saturated pulses and busy timelines, with a plethora of structures that make identification and classification challenging.


Any pulse classification algorithm will have a strong dependence on the response of the algorithms used in the preceding modules. Minute changes to the pulse finder or pulse parametrizer algorithms often prompt an update of the classifier algorithm to match their new behaviour. Changes in the pulse finder behaviour, for example, have a profound impact on the efficiency of the classification module due to the addition of new pulse structures or differences in the



**Figure 6.3:** Example of an LZ event, taken from simulated data of the detector and electronics response. On the left-hand side, the event timelines for the TPC-HG (top left), Skin (middle left) and OD-HG (bottom left) detectors are shown. Pulses identified by the *PulseClassifier* are color-coded: S1 (green), S2 (blue), SE (red), SPE (yellow), MPE (gray), and Other (also gray). For a description of the pulse classes see Section 6.2.1. On the right-hand side, the spatial representation of the integrated PMT response for the full event is shown. The colors represent the signal intensity in the PMTs of the TPC (circular arrays at the center), OD (top and bottom rectangular arrays) and Skin (radial arrays adjacent to the TPC top and bottom array plus the smaller dome array at the bottom). Figure obtained using the official LZ Offline Event Viewer.

separation of overlaid waveforms. Developing a robust classification algorithm that can handle minute changes in the upstream modules is not a strict requirement for LZap but it is an advantageous quality of any pulse classifier algorithm. The following Sections will focus on trying to determine if such a pulse classification tool can be developed.

The two main topologies present in the data are the aforementioned S1 and S2 pulses. These are the critical signatures of an interaction and are therefore the main focus of the classification efforts. The topologies of these two pulse types are sufficiently distinct to allow for good discrimination between the two, even at the level of the individual pulse parameters. A simple algorithm designed to separate true S1 pulses from true S2 pulses in a dataset containing only the two would undoubtedly yield a high efficiency, but in the context of realistic data the challenge lies in having a pulse classifier that is tuned across all possible pulse parameter ranges. This challenge is the result of the algorithm having to perform the classification on a pulse-by-pulse



**Figure 6.4:** Pulse population distribution on the full phase-space of the pulse area (pA) and pulse length at 90% area (pL90) RQs (left) and top-bottom asymmetry (TBA) and pulse area RQs (right). Some degeneracy between the different populations is present and these 2-dimensional representations are merely a guideline for upcoming results. Not all isolated populations are labelled, only the main populations mentioned in Section 6.2.1 have been highlighted.

basis, without any insight on the energy scale of the event or information on the remaining pulses in the event, disregarding the context where that pulse is inserted. As an example, at different energy scales the pulse area phase-space of S1 and S2 pulses overlaps. However, by adding more pulse parameters to the analysis this issue becomes less prominent. Despite the radical differences between S1 and S2 pulses, the ambiguity of the energy scale of the event means that no *a-priori* assumption can be made about the pulse being classified. The problem only worsens when spurious pulses are considered. These can have similarities with either S1 pulses, S2 pulses or even share similarities with both. Furthermore, the class of S2 pulse is often attributed to a somewhat vast range of pulse topologies produced by S2 light, making this class very rich and diverse, as explained in Section 6.2.1. This extended definition of an S2 pulse undoubtedly complicates the classification efforts.

Figure 6.4 shows a rough labelling of the main pulse populations in the data. The goal of this figure is to act as a reference for the upcoming results. The populations are not fully separable in any 2-dimensional representation of pulse parameters but their general distributions can be intuitively inferred.

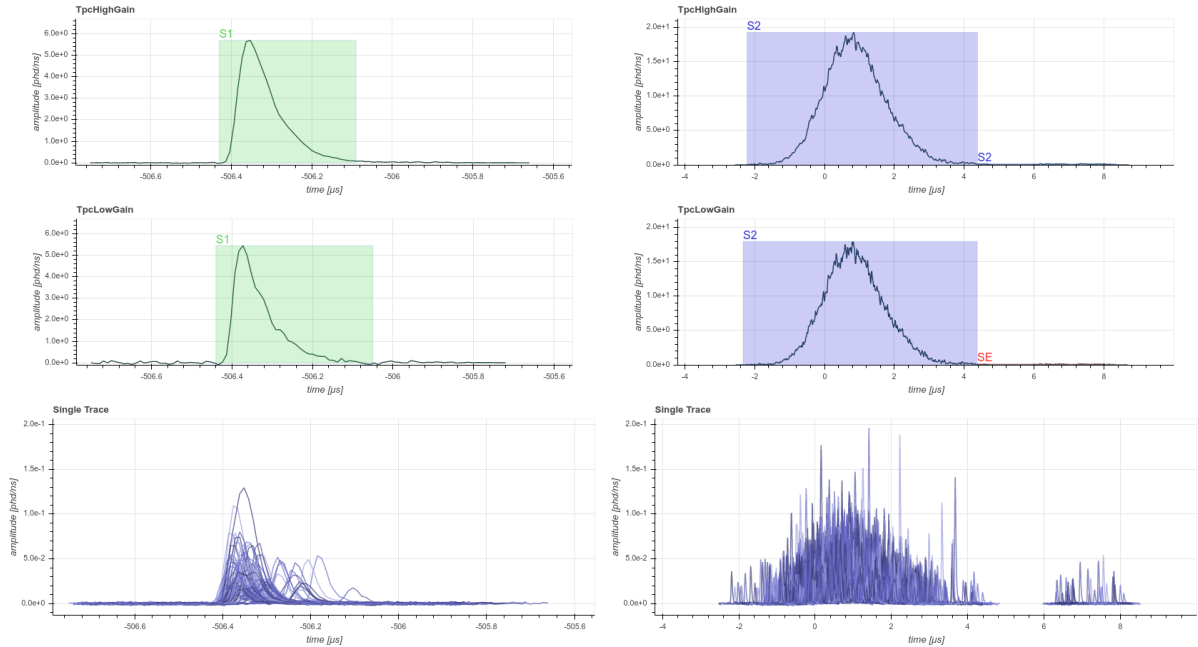
### 6.2.1 Main Pulse Topologies in LZ

A light signal collected by the PMTs will produce waveforms that can have only one of three possible origins in the LXe TPC of LZ:


- Primary scintillation light reaching the PMTs directly from the interaction site (S1 light).
- Electroluminescence created by the extraction of electrons into the gas phase (S2 light).
- Spurious signals from either the PMTs, the electronics or DAQ – the main causes being dark noise from thermionic emission in the PMT photocathode [Mey08], and PMT after-

pulsing originated by the delayed collection by the photocathode of residual ions excited by the moving charge within the PMT [A<sup>+</sup>19e, BDM<sup>+</sup>17].

Figure 6.5 shows an example of an S1 pulse (left) and an S2 pulse (right). The coloured boxes represent the pulse boundaries set by the *PulseFinder* algorithm. Due to the nature of scintillation light, the corresponding waveforms of the S1 signals are short, rising quickly ( $\sim 50$  ns) and falling exponentially ( $\sim 500$  ns), and are mostly recorded in the bottom array due to the internal reflections in the liquid-gas interface and on the walls of the detector. For a signal to be considered an S1 it must be registered in at least three individual PMTs (three-fold PMT hit coincidence requirement) in order to limit the rate of fake S1-only signals caused by accidental coincidences between multiple PMT dark counts [A<sup>+</sup>20d]. The waveforms of S2 signals have a longer time length and tend to have an almost symmetrical shape in time, dictated by the arrival times of the drifted electrons from the interaction site into the liquid-gas interface. Most of the S2 light is collected on the top PMT array and the light map it produces preserves the  $xy$  position of the extracted electrons, allowing for the reconstruction of the  $xy$  position of the interaction site – this can be seen in Figure 6.6 that shows the light signals in individual TPC PMTs for each of the pulses shown in Figure 6.5. Also in Figure 6.5, on the bottom right, a set of PMT waveforms can be seen after the main S2 pulse that resemble a SE pulse. This structure is classified in the TPC-HG as a secondary S2 because it exceeds the upper limit set on the area that a typical SE can have to be identified as one (see discussion below).



**Figure 6.5:** Typical waveforms of an S1 pulse (left, green) and S2 pulse (right, blue). These pulses are the ones present in Figure 6.3. The coloured boxes represent the pulse boundaries calculated by the *DoGPulseFinder*. The pulse classifications are obtained with the HADES algorithm (Section 6.2.3). The summed waveforms for both the S1 and S2 pulses in the TPC-HG (top) and TPC-LG (middle) detectors are presented, as well as the waveforms of 99 individual channels in the bottom array (bottom). Plots obtained with the official LZ Offline Event Viewer.



**Figure 6.6:** Light patterns in the top and bottom TPC PMT arrays for a typical S1 (left) and S2 (right) pulse – the pulses shown in Figures 6.3 and 6.5. Plots obtained with the official LZ Offline Event Viewer.

The lower limit of a pulse in LZ is the detection of a single photon by a PMT, producing a single photoelectron (SPE) signal. These pulses share the same basic shape as a typical low-energy S1 pulse but do not pass the S1 criteria of a three-fold PMT hit coincidence. The coincidence RQ is defined as the multiplicity of PMTs that recorded light within the boundaries of a pulse. Furthermore, if no pulse overlap is present, the identification and exclusion of isolated SPE pulses should be straightforward<sup>4</sup>. The unique case of an S1-like pulse that fails to meet the 3-fold coincidence criterion required to be considered an S1 and instead has coincidence equal to 2 is called a multiple photoelectron (MPE) pulse. The distinction between S1 and MPE pulses is exclusively dictated by the coincidence RQ mentioned above and the discrimination criterion can be implemented *a posteriori*. For this reason, in some analysis it is pertinent to consider S1 pulses and MPE pulses as belonging to the same S1-like family of pulses.

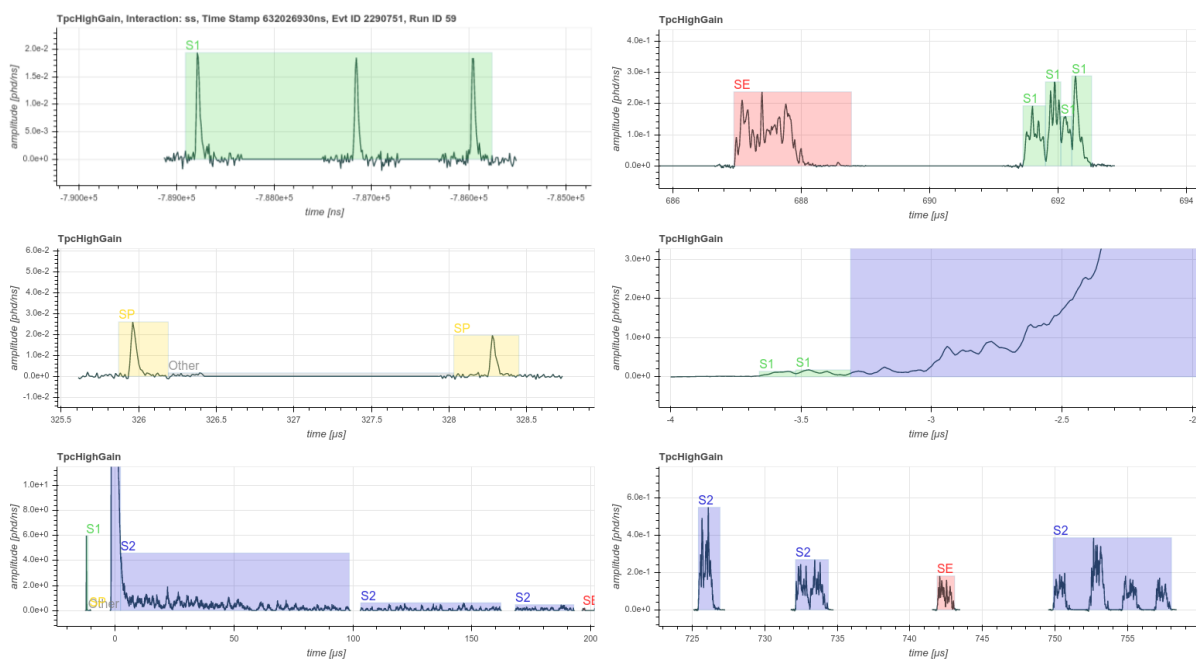
The lower limit of an S2 pulse is the average signal produced by the extraction of a single electron into the gas phase, which might happen sporadically or more frequently following a large S2 pulse [SK18, A<sup>+</sup>20f]. These single electron (SE) pulses are complex in shape but have a very characteristic pulse area due to the fixed strength of the extraction field, photon yield and light collection efficiency of the detector. Much like the coincidence criteria for valid S1 pulses, S2 pulses usually need to have several SE areas in order to be considered valid S2s, but this distinction is troublesome as the pileup of several uncorrelated SE pulses is common after a large S2 pulse. To allow for low-energy S2 analysis, the choice was made of merging SE pulses that are close in time as a single pulse that should not be classified as an SE but as either a

<sup>4</sup>See the discussion at the end of this section regarding artificial pulses and the problem with SPE coincidence.



multiple SE or an S2<sup>5</sup>. The weight of judging the merit of any S2 in the context of the full event falls to the *InteractionFinder* module in the *InteractionDetection* chain and ultimately to the person performing the analysis, allowing the *PulseClassifier* module to have some leeway on these particular cases.

Aside from the aforementioned pulse topologies, real LZ data will certainly have a large variety of pulses that will originate from artifacts in preceding processing modules, *e.g.*, partial pulses split by the pulse finder algorithm, or introduced by electronics, readout or PMT anomalous behaviour. For completeness, the main known pathological pulse structures found in the simulated data are presented in Figure 6.7.



**Figure 6.7:** Examples of pathological pulses in LZap at the *PulseClassifier* stage. Top-left: 2- and 3-fold coincident dark counts; Top-right: SE splits – *PulseFinder* issue; Middle-left: baselines between pulses with significant fluctuations; Middle right: S2 start ramps – *PulseFinder* issue; Bottom left: S2 tail structures; Bottom right: SE vertical sums (pile-up) and SE horizontal merging (E-train).

The presence of some of these structures is inherent to the data or stems from the behaviour of the detector itself, and cannot be completely avoided. A well-tuned pulse finder algorithm should consistently isolate these spurious structures and allow the classification algorithm to flag them. The best example of this situation is when more than one PMT dark count is registered in coincidence and is identified as a pulse. A 3-fold DC coincidence occurrence can be found in Figure 6.7 on the top left plot. Due to the nature of PMT dark counts, and considering their relatively high rate, occasionally having two or more in coincidence is inevitable. This is a topology of interest because to a basic pulse finder algorithm these structures can be wrongfully identified as an S1 pulse due to having fast peak-like structures within a short pulse duration and a 3-fold coincidence. The mitigation of these structures, and other alike, is therefore an

<sup>5</sup>Considering these pulses as true S2s is not strictly correct from the point of view of the final classification, but it is necessary to maintain the classification as general as possible on a pulse-level basis.

important task for the classification algorithm.

Although most separable waveforms will be identified as individual pulses, due to fluctuations in light collection timing, differences in scale of piled-up pulses, PMT saturation or digitization effects, some waveforms may be separated into more than one pulse in a non trivial way<sup>6</sup>. An example of particular interest is when a SE pulse is segmented by the pulse finder algorithm into spiked substructures due to these being close to the bandwidth associated with regular S1 pulses, triggering the S1 response of the *DoGPulseFinder* algorithm and leading it to perceive these features as if several S1 pulses were adjacent in time. This behaviour is represented in Figure 6.7 on the top right, where two SEs can be seen side by side, but the right-most one is shown to be split into several “artificial” pulses that get classified as S1. The reason for this classification is their short time length and spike-like appearance, coupled to the fact that the pulse classifier algorithm in LZap is designed to only look at each pulse structure and not its surrounding structures. This scenario poses a problem to the reconstruction of the event in many ways, the most obvious being when a single scatter event gains more S1 pulses and the *InteractionFinder* algorithm will likely tag it as a pileup event or in some rare cases as a Krypton event, contributing to the overall loss of efficiency of the analysis.

Fluctuations in the waveforms may also produce artificial structures that do not have a physical origin in the detector but end up being identified as pulses. These are often found adjacent to other pulse structures. The middle-left plot in Figure 6.7 shows an example of a baseline connecting two SPE pulses being isolated as a pulse. Some artificial structures can also be seen on the middle-right timeline of Figure 6.7, where over-splitting of the start of the S2 waveform results in sharp pulses that the classifier sees as more S1-like. Along with the SE-split structures seen in the top-right plot, this S2-ramp split represents a serious challenge to the pulse classification efforts (see previous comments).


The two bottom plots in Figure 6.7 are examples of S2 tails and electron “trains”, two very common structures that follow S2 pulses roughly above 100 keV [A<sup>+</sup>20f]. These structures are formed by the late extraction of drifted electrons from ionization in the interaction site, photoionization of chemical impurities in the xenon bulk, or the sporadic extraction of latent electrons in the liquid-gas interface, correlated to the S2 charge and light effects in the extraction region and field grids [SK18, A<sup>+</sup>20f]. Since these occurrences are more likely to happen closer in time to the main S2 pulse, the waveforms tend to overlap significantly immediately after the S2 pulse to form a long continuous tail, and get increasingly apart the further in time they occur, forming a trail of structures containing a varying number of piled-up SE waveforms. The bottom-right plot in Figure 6.7 shows four *summedPODs* containing, in order: vertical sum (pileup) of two SE waveforms, horizontal sum (chain) of two SE waveforms, a single SE pulse well isolated, and a large *summedPOD* with four SE pulses that get chained by the pulse finder for having a separation time under the set threshold<sup>7</sup>. These structures (S2 tails, SE pileup and chained SE pulses) are not exactly pathological since they are expected in a detector such as

<sup>6</sup>This is the main failure mode of the pulse finder but it is expected to be a very rare occurrence nonetheless.

<sup>7</sup>Chained SE pulses in the e-train structures is an expected effect and it is a by-product of tuning the pulse separation threshold at the SE scale to cluster afterpulsing with the respective SE pulse.

LZ. However, their shape and size can vary significantly, with some of these structures being indistinguishable from low energy S2 pulses. For that reason, these structures are classified as S2 pulses in order to preserve the overall efficiency of physics analyses at lower energies. The *InteractionFinder* module can then distinguish between these false S2 pulses and real low-energy S2 pulses using the context of the entire event.

Figure 6.8 shows an S1 pulse followed by four afterpulsing waveforms, with the characteristic  $\sim 1.5\mu\text{s}$  time after the main pulse, that are misclassified as S1s. The inlay of the figure shows only the individual channels with light in the time window of the afterpulse structures, where it is clear that these are single channel structures in coincidence. The classification of these structures as S1 pulses is a common false positive result<sup>8</sup>. However, on the perspective of a pulse classifier algorithm, these pulses with clear S1 shape and scale that happen to have coincidence larger than 1 are indistinguishable from real S1s or MPEs without context, making this failure mode irreducible with the current definition of coincidence. These artificial S1 pulses can be identified by the *InteractionFinder* module with high accuracy and are often inconsequential.



**Figure 6.8:** Summed PODs of an S1 pulse followed by baseline noise and four afterpulses. The two central afterpulses were overlaying and were merged into one pulse. The afterpulses were classified as S1s by having coincidence  $>1$  due to runaway baselines and pileup in different channels. The inlay shows the 4 channels that produced afterpulsing. Plots obtained with the official LZ Offline Event Viewer

A large portion of the work of the pulse classifier module is in identifying the origin of these pathological pulses, separating them from main pulse populations and labelling them accordingly. The pulse classification algorithm needs to maximize both the efficiency of identification of the main pulse topologies (S1, S2, SE, SPE) as well as a large efficiency in the identification of pathological pulse structures and the rejection of bad pulses, with particular concern in avoiding misclassification of bad pulses as the main pulse topologies. Here, pulses are considered to be “bad” if they contain any non-physical or clearly miscalculated RQ (negative total area, zero length or width, negative fractions, *etc.*) or if they originate from artifacts in preceding processing modules (pathologies), such as the isolated baselines or the split pulses mentioned previously. Some of these pulses do not correlate to any particular signal in the detector and should be classified as “Other”, which informs the following data processing stages and analyses that these pulses do not contribute to the physics of the event. It is, however, extremely important to monitor these pulses as they can indicate the existence of issues in the processing chain, data acquisition or the detector itself. Notice that not all the pathological pulses depicted in

<sup>8</sup>The probability of afterpulsing scales with the size of the preceding S1 pulse. The *InteractionFinder* module can be tuned to tag these misclassified afterpulses by comparing their size with the size of a preceding large S1.

Figure 6.7 should be classified as Other pulses, as in the case of S2 tails, SE pileup and chained SE pulses that should be classified as S2.

## 6.2.2 Classification Efficiency and Other Metrics of Performance

The measure of the performance of a pulse classification algorithm accounts for both the overall efficiency of classification, *i.e.*, the rate of successful classifications in a dataset, but also the individual rates of misclassification of each individual class, here named “failure modes” of the classifier. This distinction is important in a non-binary classification problem because the overall efficiency often shadows specific failure modes that may be more or less relevant to the performance of the algorithm. Figure 6.9 shows, on the left, a diagram displaying the relations between the true and the assigned classification, categorizing each relation according to the validity of the classification (true or false results) and the classification activity (positive or negative results), *i.e.*, the relevance of the pulse to the remaining algorithms on the processing framework. The Other pulse class is by default inactive and will be ignored by the rest of the processing chain unless specified otherwise, thus any classification to the Other category is negative. The diagram emphasises the importance of distinguishing between the different results of classification when discussing the efficiency of an algorithm, showing in red the quadrants that have severe impact in the overall efficiency of the LZ analysis.

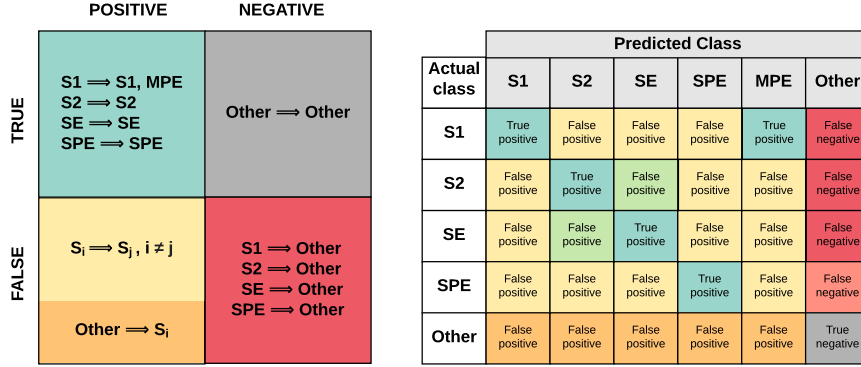
From Figure 6.9 the true positive (TP) results obtained by processing LZap data are defined as the number of occurrences with the following pairing of prediction and true class:  $\{S1, S1(MPE)\}$ ,  $\{S2, S2\}$ ,  $\{SE, SE\}$ ,  $\{SPE, SPE\}$ . The mixing of S2 pulses and SE pulses in the classifications is not impactful to the physics analysis, so sometimes it is useful to consider the pairing  $\{S2(SE), S2(SE)\}$  as part of the true positive results. These instances will be denoted accordingly. The true negative (TN) results are defined, in the same way, as the pair  $\{Other, Other\}$ . The false positive (FP) results correspond to any pairing of the type  $\{S_i, S_j \neq S_i\}_{S_i \neq Other}$ , meaning any prediction different than Other that is not the same as the corresponding class label. Finally the false negative (FN) results correspond to any pairing of the type  $\{Other, S\}_{S \neq Other}$ .

The top quadrants of the left schematic of Figure 6.9 represent the correct classifications (true positives and true negatives) that characterize the classical accuracy for a classifier, defined as:

$$acc = \frac{TP + TN}{N}, \quad (6.1)$$

with  $N = TP + TN + FP + FN$  the total number of occurrences. This is the classical definition of accuracy on a multi-class problem, where all diagonal elements are considered a success of classification and all non-diagonal elements are equally weighed as misclassifications.

The bottom quadrants of the left schematic of Figure 6.9 represent the misclassification “failure modes”: false positive and false negative results. The presence of false negative results effectively leads to information loss at the pulse level and misclassification of the overall event, which will directly impact the physics results, as mentioned previously. The false positive results can lead



**Figure 6.9:** Diagram of classification results (left) grouped by classification impact (positive or negative results) and validity of classification (true or false results) and the respective distribution over a generic confusion matrix (right). The colours represent the severity of the classification results to the ensuing analysis, from red (most severe) to green (not severe). The mixing between S2 pulses and SE pulses is tolerable since both represent the same process and are distinguished by area alone. False positive results are divided in misclassification of Other pulses and “mixing” of true pulses, as the repercussions of each case are different. The case where an SPE pulse is classified as Other bears no significant impact on the physics analysis.

to misclassification of the overall event too, but having positive/active classifications means that the *InteractionFinder* module will contextualize the pulse topology within the event and, if the failure mode is not critical within the event, possibly compensate for it or even actively correct it. The most problematic set of false positive results is the misclassification of Other pulses into active classes, highlighting the misclassification of SE splits into S1 pulses (top right example in Figure 6.7).

The positive predictive value, also called precision or “true efficiency”, is a very insightful metric for a classifier. It is defined as the ratio of true positives over all predicted positives (true positives and false positives) [IS17]. In the classification of LZap pulses, this means excluding the predictions that returned the Other class,  $\mathcal{I}(y_n \neq \text{Other})$ . This quantity is estimated using

$$PPV = \frac{TP}{TP + FP} \quad (6.2)$$

Conversely, the negative predictive value is a quantity that measures the ability to identify the true negatives over all the pulses that were identified as Other (predicted negatives), and is given by

$$NPV = \frac{TN}{TN + FN} \quad (6.3)$$

Achieving a high *PPV* is the first goal of the classifier. On the other hand, having a low *NPV* is not inconsequential, as it indicates that a large fraction of pulses are being misclassified as Other (FN). As mentioned previously, the FN misclassifications are a serious problem for the physics analysis, since in these cases the misclassification essentially removes one valid pulse from the analysis.

The performance of a non-binary classifier is often represented in the form of a confusion matrix, showcasing the partitioned efficiencies per class. An example of a confusion matrix relevant for the classification exercises on LZap can be found on the right-hand side of Figure 6.9. Both diagrams of Figure 6.9 are color-coded in order to highlight the correlations between the mixing in the results and the corresponding classifier efficiency.

### 6.2.3 Pulse Classifier HADES

The pulse classifier module implemented in LZap uses the *Heuristics Algorithm for Discrimination of Event Substructures*<sup>9</sup> (HADES), a heuristic decision tree that performs the classification at the pulse level using a set of RQs for each individual pulse. The module performs a sequence of decisions using several selection criteria tuned to provide the best classification accuracy, as explained below. HADES is a transparent algorithm designed to be the benchmark for more advanced pulse classification modules. It is designed to distinguish and categorize individual pulses within an event into six classifications:

- Primary scintillation (S1) - prompt scintillation released in an energy deposition in the LXe and detected in at least 3 different channels (3-fold coincidence).
- Electroluminescence (S2) - proportional scintillation generated by extracting drifting electrons into the gas phase of Xe.
- Single electron (SE) - scintillation in the gas phase created by a single electron extracted.
- Single photoelectron (SPE) - a single PMT signal, consequently with coincidence of 1.
- Multiple photoelectron (MPE) - an S1-like pulse with less than the required 3-fold coincidence level to be classified as an S1.
- Undetermined (Other) - pulses with “bad” parameters (accidentals like out-of-bounds, nonphysical values or miscalculations) or strange pathologies that excludes them from the remaining classes.

More details on the physical origin of these signals in the particular case of dual-phase TPCs such as LZ can be found in Section 4.1. As the different pulses expected to be observed have different features, one can separate distinct pulse populations in a given parameter-space with a certain efficiency using a set of parameterized functions. Combining a set of cuts across different pulse parameters can yield high classification efficiency. The selection criteria used by HADES are a set of functions fitted to the simulated data, after this is processed by the *PulseFinder* and *PulseParametrizer* modules. The selections aim at separating the different pulse populations with the highest efficiency across all of the phase-space they are defined in. Pulses that have the same physical origin in the detector will share common geometric traits and will naturally be grouped. A pulse population is defined as a set of pulses that are clustered together in a given phase-space by sharing similar pulse parameters.

The logical scheme of the HADES algorithm is presented in Figure 6.10. The pulse parameters

---

<sup>9</sup>Continuously integrated as part of the *PhotonDetection* package <https://gitlab.com/luxzeplin/lzap/PhotonDetection.git>

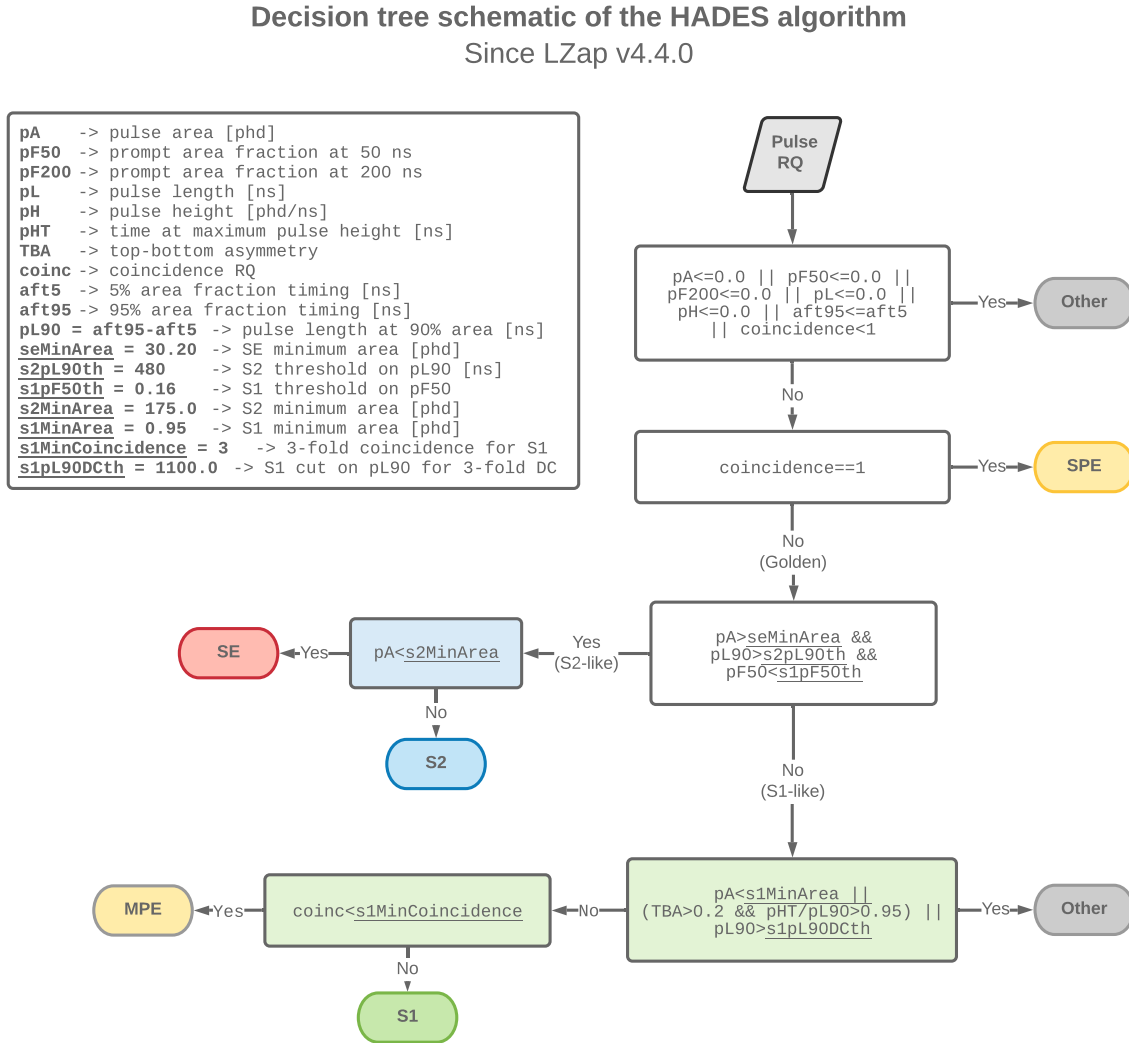
used are explained in detail in Appendix B. The algorithm starts by performing a series of “goodness” cuts to classify any pulses with non-physical properties, such as negative total area, negative pulse length or negative pulse amplitude as Other pulses. The next step is to sieve the single coincidence pulses by applying a requirement of more than one PMT hit in a given summed POD, classifying the pulses that fail this requirement as SPE<sup>10</sup> if coincidence=1 and as Other if coincidence<1. The remaining pulses that pass the first two steps (“goodness” cuts and coincidence>1) are called “Golden” pulses, and represent a population with all valid S1 pulses (MPEs included), S2 pulses and SE pulses. The next step implements the main cuts that will separate the pulses with S1-like properties (S1, MPE) from pulses with S2-like features (S2, SE, S2 tails, etc...). A combination of the total pulse area ( $pA$ ), area fraction within a 50 ns time window on the start of the pulse ( $pF50$ ) and pulse length at 90% pulse area ( $pL90$ ) prove to be highly efficient at separating S1-like pulses from S2-like pulses (see discussion below). Following the S2-like branch of the decision tree, the next cut implemented is tuned to separate SE pulses from S2 pulses using an area threshold, as SE pulses have a characteristic area that depends only on the conditions of the detector. The S1-like branching of the HADES decision tree will lead to another stage of “goodness” selections, where some pathologies on the S1 parameter phase-space are removed into the Other pulse class, while the remainder is separated into MPE pulses and S1 pulses based on the 3-fold coincidence requirement for an S1, `s1MinCoincidence`.

Figure 6.11 presents the selection functions or thresholds set by HADES and explained in the previous paragraph. The top panels show the conditions for the initial separation between S1-like and S2-like pulses, with the thresholds `seMinArea` = 30.2 phd, `s2pL90th` = 480 ns and `s1pF50th` = 0.16 represented by the coloured lines. Despite the initial separation being focused on S2-like pulses, the remaining populations on the S1-like phase-space, shown on the bottom left panel, are fairly pure. The second set of selections applied, namely the thresholds `s1MinArea` = 0.95 phd and `s1pL90DCth` = 1100 ns, separate baselines and 3-fold coincident dark counts from the main S1-like populations, respectively. The bottom right plot of Figure 6.11 shows the area threshold for separating S2 and SE pulses (red line) at `s2MinArea` = 175.0 phd. This threshold is set by finding the minimum value in the area distribution between SE pulses and pulses generated by two extracted electrons. However, the value currently being used was tuned several versions before the current LZap version considered in this work, and a new value tuned using the current LZap version is presented (orange line). The `s2MinArea` threshold depends on the single electron size, a very important parameter that depends on the conditions and performance of the detector (see discussion in Section 4.1.2), meaning that this threshold will have to be measured and adjusted to real data.

The pathologies sieved in the S1 branching of HADES are mostly 3-fold dark count coincidences (Figure 6.7 top left), floating baselines (Figure 6.7 center left) and S2 ramps (Figure 6.7 center right), although the efficiency of removal of the latter is low. The baselines can be removed by implementing an area threshold `s1MinArea` = 0.95 phd that ensures that for a 2-fold coincidence pulse (MPE) is still considered a good pulse, the individual SPE areas must be higher than

---

<sup>10</sup>The label “SPE” is a legacy of the time when no afterpulsing was simulated. The SPE class of pulses contains in fact all pulses with single coincidence, and are often noted as “single coincidence pulses” (SCPs).



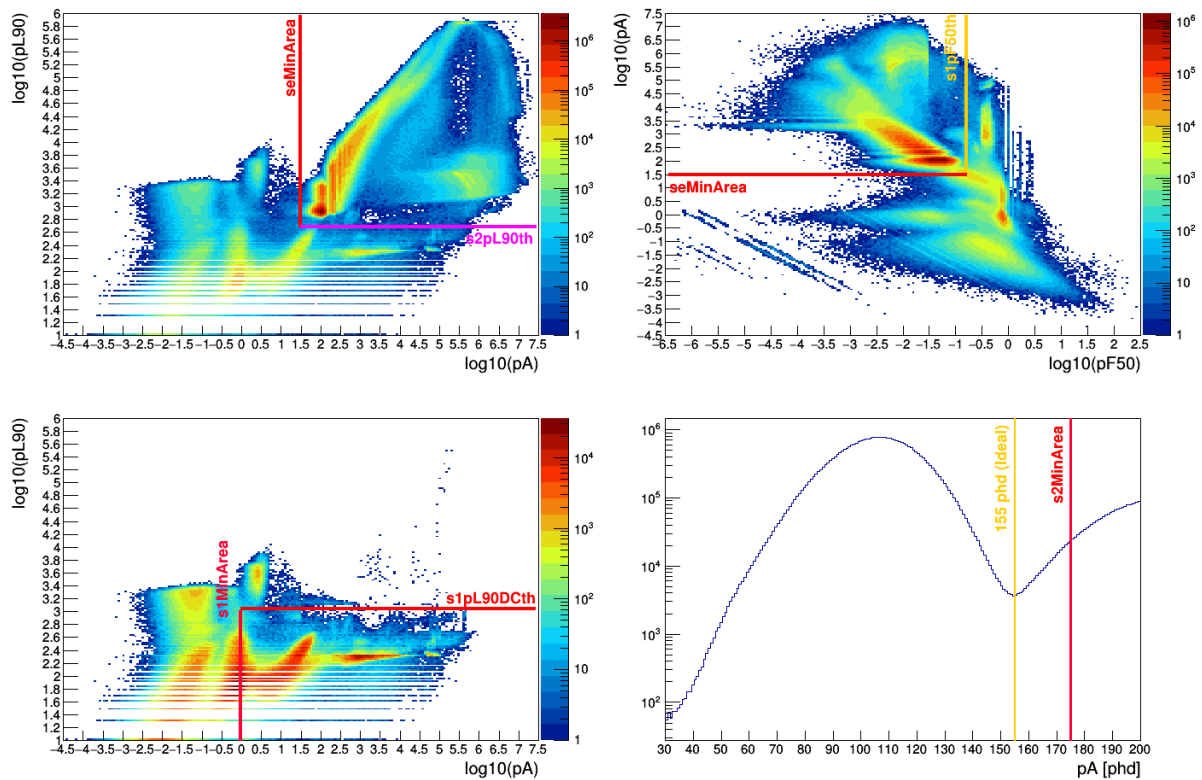
**Figure 6.10:** Diagram of the decision tree of the HADES algorithm for pulse classification in LZap. The depicted version is implemented in LZap since version 4.4.0.

0.455 phd on average<sup>11</sup>. The 3-fold dark count coincidence population, clearly seen in Figure 6.4 can be separated with a simple pulse length threshold on the S1 domain,  $s1pL90DCth = 1100.0$  ns. This pulse topology is well bounded in pulse length because of two main factors: from above by the ability of the pulse finder to separate individual dark counts that are a certain distance apart<sup>12</sup>, and from below due to the decreasing probability of having 3 pulses in coincidence in shorter time windows. The S2 ramps (Figure 6.7 center right) are partially removed by selecting pulses with peak time ( $pHT$ ), *i.e.*, the time at which the pulse reached peak amplitude ( $pH$ ), closer to the end of the pulse length, effectively tagging pulses whose shape

<sup>11</sup>The size of an SPE pulse depends on several factors and it is PMT-dependent, meaning that before an SPE calibration is performed to ensure  $pA(SPE) = 1$  phd, the SPE size can fluctuate significantly.

<sup>12</sup>If the individual DCs are too spread apart the pulse finder will be able to isolate them and thus the structure will not be considered a single pulse.





**Figure 6.11:** Depiction of HADES selections in the phase space of the pulse parameters. The top panels represent the first instance of S1-S2 separation as defined in the third decision node of the HADES diagram (Figure 6.10). The bottom left panel shows the separation of “bad” pulses from the remaining S1-like population. The bottom right panel depicts the separation of SE pulses from the S2 ensemble, showing the area threshold used in LZap for this analysis (red) and the optimized value for this dataset (orange).

is negatively skewed. However, this selection could also tag some SE pulses with  $pH$  upward fluctuations near the end of the pulse, but these occurrences were never observed in the data and are therefore believed to be extremely rare.

The S2-like side of the diagram does not include a selection for any pathological populations. The reason for this is that the criteria for separating S1-like and S2-like populations (third decision node in the HADES chain) is optimized for S2-like separation, ensuring that the S2-like populations are rather pure from the start.

### 6.2.3.1 Construction and tuning of HADES

HADES is a greedy algorithm, *i.e.*, it employs the decision criteria that maximizes classification accuracy at each decision step while still attempting to find the optimal global solution for the classification of all pulses. The ensemble of decision steps optimized individually gradually tend to a solution that approximates the optimal one, but it not guaranteed that it will reach this optimal solution. Nevertheless, the goal of HADES is to be a simple and intuitive algorithm,

and despite its simplicity it achieves a good classification performance, as discussed at the end of this section.

The pulse parameters (RQs) used in the selection functions in each decision step are carefully chosen for their ability to separate the pulse populations of interest. These choices were partly influenced by past experience, most notably from previous classification tools created for the LUX experiment. The selection functions are adjusted iteratively at each decision step and follow the flow of the data in the algorithm (top-to-bottom). The number of consecutive decision steps is limited to provide the minimum number of terminators that would ensure that each class was represented as an output. The reasons for this decision are twofold: the first reason is to avoid over-training of the algorithm to a given instance of the data, preserving the ability to generalize to possible variations of detector conditions and any minor changes on preceding processing stages (robustness), and the second reason is to have a minimal number of selection criteria and degrees of freedom that can be easily tuned. More insights on methods for building a decision tree are presented in Section 7.4.1.

In each iteration of the tuning process, a cross-validation is performed by observing the separation efficiency of pulse populations in the available transverse phase-space representations, *i.e.*, in the domain of the pulse parameters that are not included in the preceding decision node. If a new population is detected in one or several transverse representations that is distinctively disjoint from the target population, or if the target population displays non-trivial loss of pulses, the cut is readjusted and subsequently verified. All populations are sampled and handscanned<sup>13</sup> in the event viewer of LZ<sup>14</sup>. Some populations that are found to be isolated, *i.e.*, that present regular behaviour separated from the target populations or display unexpected features, are selected for handscanning with the event viewer in order to identify the pulse type. The pulse selection functions are tuned until no clear improvement is observed.

This straightforward way of modifying the selection criteria is one of the strengths of the HADES algorithm. The plasticity and simplicity of the selection functions means that these can be readily changed if the algorithms upstream are modified or detector conditions change, accounting for their response conveniently. Also for this reason, the functions are, most of the times, simple thresholds that have some physical or geometrical significance in the context they are used.

### 6.2.3.2 Overall efficiency of HADES

Despite its simplicity, HADES performs remarkably well at classifying true S1s, S2s, SEs and SPEs, with an estimated accuracy of 98.6% over all pulse types. The color-coded classes of the pulses presented in Figures 6.3, 6.5, 6.7 and 6.8 are from the results of the current HADES algorithm. Note that the cases displayed in Figure 6.7 were hand-picked as examples of pathological features in the data and should not be considered common. Some features in pulse populations

---

<sup>13</sup>Analysis of several events and their waveforms by eye by collaborators, with the goal of verifying the performance of the analysis algorithms and of the simulations.

<sup>14</sup>In reality, a dedicated event viewer created for the development of the *PulseFinder* algorithm was used during most of this analysis, due to severe delays in the delivery of the official event viewer of LZ.


separated by HADES are visible in Figure 6.11 and are the result of the algorithm not being fine-tuned for this dataset, due to the parallel development of HADES, the preceding algorithms in LZap and the simulation package. The version of HADES used to produce these results was tuned using an older version of simulated data that has been updated several times prior to the production of the data used in this work. The changes in the simulation were not significant enough to prompt a re-tuning of HADES but have destabilized the classification in some rare occasions. On the SE population on the bottom-right panel of Figure 6.11, it is noticeable that the `s2MinArea` threshold could be improved in order to minimize the number of pulses from the adjacent population (two merged SE pulses) being grouped with the SE population. Nevertheless this leakage of pulses that should be classified as S2 into the SE population is low and the impact on the physics analysis is limited to the characterization of the SE size, which can be easily corrected. The separation of spurious pulses present in the S1 populations is challenging, especially the SE splits shown in Figure 6.7 on the top-right. These occupy most of the same phase space as the real S1-like pulses and HADES is not tuned to separate them, meaning that they will mostly end on the S1 pile. The SE split population can be easily identified visually as a smeared Gaussian-like structure with average *TBA* of 0.2 and average area of 30 phd in the right plot in Figure 6.4.

The classification efficiency of HADES can be estimated directly by performing handscans of the data and reporting the occurrences where the algorithm is considered to have failed. The current version of the algorithm was tested extensively on several handscanning exercises, the latest performed at a collaboration-wide workshop covering a dataset of 750 events from simulated sources that included radioactive backgrounds, internal radon mixed in the LXe and neutrons from a deuterium-deuterium (DD) source. Table 6.1 summarizes the properties of this 750 event dataset after being processed with HADES. This dataset contained 98784 pulses found by the pulse finder algorithm. Of these pulses, 3329 (3.37%) were classified as Other pulses for having nonphysical RQs and 32860 (33.26%) were classified as SPE for having coincidence equal to 1. Of the remaining pulses, denoted as “Golden”, 6373 (9.87%) were classified as S1, 18862 (29.22%) were classified as S2 and 37043 (57.38%) were classified as SE. As it can be seen in the pie chart inscribed in Table 6.1, the “Golden” population is composed mostly of SE and S2 pulses, with only a small fraction of S1-like pulses. The S2 pulses are in fact mostly S2 tails, with only a fraction of those being real S2 pulses. The same can be said about the S1 pulses, where an estimate taken from the *InteractionFinder* results shows that there are less than 2 real S1 pulses per event.

The handscanning exercise concluded that no crippling classification issue was observed in the data and that the processed information was viable for physics analysis. These conclusions, despite favourable, do not return a clear metric of success of the algorithm. For that reason a dedicated handscan of 23 random events, corresponding to 2748 pulses, was performed in which all pulses were visualised individually. The results of that dedicated handscan can be found in Table 6.2. Of the 2748 pulses, HADES returned 39 detected failures, resulting in an estimated overall accuracy of  $acc = 98.6\%$  (Equation 6.1). However, some of the misclassified pulses are known issues that can be either flagged in the *InteractionFinder* module or simply ignored such as in the case of mixing between SE and S2 pulses. An example of the former is the SE splitting

**Table 6.1:** Metrics of the processing of 750 simulated events from multiple backgrounds,  $^{220}\text{Rn}$  decay in the LXe bulk and neutrons from a DD source. The values in the table were obtained through HADES directly. The pie chart on the right side shows the frequency of the different classes for all “Golden” pulses, defined as the pulses without bad RQs and with coincidence above 1. These results include misclassifications of HADES, particularly relevant for the fraction of S1 pulses in the data, which is largely inflated by the misclassification of SE splits.

	N pulses	event avg.	% total	% golden
All pulses	98784	131.7		
Bad pulses*	3329	4.4	3.37	
SCP	32860	43.8	33.26	
Golden	64561	86.1	65.36	
S1	6373	8.5	6.45	9.87
S2	18862	25.2	19.09	29.22
SE	37043	49.4	37.50	57.38
MPE	317	0.4	0.32	0.49
Bad S1-like*	1966	2.6	1.99	3.05
Bad S2-like*	0	0.00	0.0	0.00



\*Other pulses

**Table 6.2:** Confusion matrix of HADES results taken from the handscanning of 2748 pulses on 23 random events.

Actual class	Predicted class						Total	
	S1	S2	SE	SPE	MPE	Other		
S1 or MPE	23	0	0	0	1	1	25	0.9%
S2	0	494	3	0	0	0	497	18.1%
SE	0	1	1013	0	0	0	1014	36.9%
SPE	0	0	0	1089	5	7	1101	40.1%
Other	13	0	0	9	0	89	111	4.0%
<b>Total</b>	36	495	1016	1098	6	97	2748	

from the *PulseFinder* module represented in the top-right plot of Figure 6.7.

The positive predictive value,  $PPV$ , of HADES over this dataset is estimated using Equation 6.2 to be  $PPV = 98.83\%$ . If the mixing between S2 and SE pulses is ignored, the positive predictive value becomes  $PPV = 98.98\%$ . This mixing in HADES classification is due to fluctuations on SE size and fine tuning of the S2 area threshold. The mixing is fairly low and these cases will not hinder the physics analysis. On the dedicated handscan exercise, only one of the 1014 SE pulses was found to be misclassified as a S2 pulse. Conversely, only 3 S2 pulses were misclassified as SE. The SE pulses that are classified as S2 are simply pulses with area larger than the `s2MinArea` threshold, while virtually all S2 pulses that are classified as SE are two coincidental SE pulses with low area. As mentioned previously, the `s2MinArea` threshold used by HADES is not properly tuned due to the decrease of the SE size in this dataset, which leads to some leakage of small S2 pulses (two coincidental SE pulses with area under-fluctuations) into the SE population. The bottom right plot in Figure 6.11 shows the deviation from the optimal

separation value for the current version of the data. HADES is robust enough to handle these fluctuations to some extent, as in this case. Since e-trains like the ones presented in Figure 6.7 are correlated in size with the S2 pulses that precede them, the probability of having merged SE pulses increases with the size of the main S2 of the event. This means that the *InteractionFinder* module can easily veto false S2 pulses caused by SE merging or pileup by searching for a large S2 earlier in the event.

The negative predictive value of HADES over this data is estimated using Equation 6.3 to be  $NPV = 91.75\%$ . This result is dominated by the misclassification of SPE pulses as Other pulses (0.6%), which has no impact on the analysis in any meaningful way. The remaining cases are misclassifications of S1 pulses as Other pulses, which was observed only once in the handscan. This case, believed to represent the majority of the S1 misclassifications, is the result of the incorrect identification of a complicated structure between the S1 tail and afterpulsing. This failure mode can also be resolved by the *InteractionFinder* module.

For completeness, it is worth analysing the rate of Other pulses that are misclassified by HADES, *i.e.*, the fraction of Other pulses that are incorrectly classified as an active class. As mentioned in Section 6.2.2, this subset of FP results can severely impact the rest of the analysis and cause unintentional stress on the *InteractionFinder* module. This rate is calculated to be 19.8%, with the only contributions coming from the mixing of Other with S1 or SPE pulses. The major FP contribution is from SE splits (Other) being classified as S1 pulses (13 occurrences, 11.7% of Other pulses), as shown in Figure 6.7 and discussed in Section 6.2.1. This is the main issue in classification for the current version of LZap, and contributes to the decrease of the overall accuracy of HADES significantly. It is strictly a *PulseFinder* issue that could be potentially solved with a fine-tuning of the algorithm. However, real LZ data might yield other topologies that are as difficult to handle by the modules preceding the *PulseClassifier* and for this reason, in the context of the development of LZap, any pathology found in the simulated data for MDC3 should be thought of and treated as a real pathology would be. As mentioned before, this failure mode in particular can be identified by the *InteractionFinder* module by looking for tight clusters of pulses classified as S1 within a short time window and with characteristic area and height<sup>15</sup>, that are not preceded by a larger S1 and therefore not consistent with afterpulsing. This problem in particular is one of the major drives for exploring Machine Learning techniques for pulse classification, discussed in detail in Chapter 7.

The next contribution to the FP rate of Other pulses is the identification of bad pulses as SPE when these have coincidence of 1 (8.1% of total Other pulses). Shown in Figure 6.10, SPE pulses are selected by coincidence alone, which is not enough to separate PODs with a single channel baseline or isolated afterpulsing. Most of the cases of Other pulses with coincidence of 1 delimited by the *PulseFinder* module are from baselines in the tails of S1 pulses or within pulse structures as in the example in the left-middle panel of Figure 6.7. These structures are expected to have low area, tending to zero, but minute statistical fluctuations due to noise or imprecision on the pulse boundaries can easily increase the area of these baselines to that of

---

<sup>15</sup>The area of these SE splits is approximately 25 phd, which corresponds to a SE of area  $\sim 105$  phd being split into 4 or 5 structures, on average.

small SPE pulses. Adjusting the SPE selection on LZap is an exercise that must be studied in the future. These misclassifications are not critical in the context of the analysis and are only relevant for studies involving SPEs.

To improve the identification of SPE pulses the coincidence RQ needs to be modified to exclude contributions from empty channels. As mentioned previously, the coincidence RQ considers all channels (including empty ones) within the boundaries of the pulse, needing a more refined definition of what constitutes a valid channel POD to be considered for coincidence calculations<sup>16</sup>. The mixing between S1, MPE and SPE pulses will undoubtedly decrease when the coincidence RQ is refined. In a future version of HADES the selection of SPE pulses may also be performed after the separation between S1-like and S2-like pulses, since the coincidence RQ is only relevant in the S1-like domain.

As a final side note, the overall accuracy quoted above at 98.6% can be thought of as being the overall efficiency of the [*PulseFinder*, *PulseParametrizer*, *PulseClassifier*] system. The inclusion of pulses that have no actual correlation with physical processes in the detector but are artefacts created by data processing algorithms will inadvertently lower the effective accuracy of HADES, *i.e.*, the ability to correctly classify real pulse structures given that  $s \neq \text{Other}$ . If these pulses with dubious origins or created from pathological problems not related to classification are excluded from the accuracy calculations, like the ones represented in Figure 6.7, and the focus of the classification exercise is in successively separating S1-like pulses (S1, MPE, SPE) from S2-like pulses (S2, SE) the overall efficiency is, obviously, much higher.

A final method for estimating the efficiency of the algorithm is to analyse the performance of the *InteractionFinder* module and the physics results obtained at the end of MDC3. With the current version of the HADES algorithm, no major constraint was perceived on the physics output of the simulated MDC3 data, with all the deliverables being met. This indicates that HADES is not hindering the performance of the subsequent modules in LZap not the physics analysis itself.

### 6.2.3.3 Limitations of HADES and future improvements

Perhaps the most important remark to be made to any classifier constructed using simple boundary conditions as discriminants regards human bias. Every decision made by the developer while constructing the classifier algorithm is limited by the level of knowledge they possess about the data. On complex datasets such as the ones presented, it is natural that some pathological problems in the data are not identified correctly upon inspection via handscans. The central motivation for building an *ad-hoc* model is to easily achieve reasonable efficiency at identifying the most common data features.

HADES is meant to be simple and robust. The scope of HADES is limited to the separation of

---

<sup>16</sup>The channel POD does not end when the channel counts go to zero, there is a buffering window where samples are being recorded with no area or close to zero. This means that a baseline with  $pA \sim 0$  in one channel will be included in the coincidence calculation.

S1-like and S2-like pulses that form the basis of the analysis. For this reason the introduction of specialized selections of non-global pathologies is not desirable, as these can anchor HADES to a specific version of the data. Once the pathology is identified and solved in a new improved version of the data, HADES will have a legacy criterion that no longer fit its purpose and that can in fact be detrimental to the classification.

The characteristics that make HADES robust are also responsible for its lack of plasticity and adaptability to different data. When data changes enough to prompt a new version of HADES, this can be easily modified accordingly. However, it is not often clear when the pulse classification is affecting the scientific performance of the processing chain, or if the classification module is actually just identifying some unknown pathological behaviour in the data or from the preceding algorithms.

It is evident that HADES can be improved beyond its present version. These improvements should be implemented with the goal of yielding better results and improve the efficiency at partitioning the data in a general way. HADES must be able to handle realistic data, which will definitely present new challenges due to the expected increase in complexity when comparing with simulated data. Perhaps one of the best places to start looking for improvements for HADES, or even just techniques that may provide new insights on LZ data, is in machine learning methods. These tools can be used to interpret the data in an intrinsic way, that can lead to solutions to partitioning the data efficiently. Machine learning tools can provide some insight on the data that may lead to a better version of HADES, or at least one that can handle data in a more general way. Several machine learning solutions are explored in Chapter 7.

#### 6.2.4 Pulse Classifier COMPACT


The HADES classification algorithm was the second to be implemented in LZap during this work. The original pulse classifier algorithm, named *COMprehensive Pulse Analysis and Classification Tool*, or COMPACT<sup>17,18</sup>, used a set of two-dimensional probability density functions (PDFs) built from the distributions of S1-like pulses (S1, MPE, SPE) and S2-like pulses (S2, SE, S2 tails) to determine the most likely nature of a given pulse. The goal of COMPACT was to determine if a pulse was more S1-like or S2-like by comparing the parameters of the pulse to the PDF response at the parameter values. The pulse was given a score for each class, calculated from the responses for all PDFs considered for that class, that was then converted to a non-normalised probability vector. The pulse could then be assigned a categorical label by choosing the classification with the strongest score.

Figure 6.12 shows the original PDFs for S1-like and S2-like populations in the six considered 2-dimensional phase-space representations. The pulse parameters used to build the PDFs and classify the pulses are: top-bottom asymmetry (*TBA*), prompt fraction at 10 ns (pF10), height to length ratio (H2L), and a newly introduced pulse parameter named pulse asymmetry, defined

---

<sup>17</sup><https://gitlab.com/PauloBras/COMPACT.git>

<sup>18</sup>formerly at <https://lz-git.ua.edu/pbras/COMPACT>



**Figure 6.12:** COMPACT 2D KDE PDFs for all representations used, showing the S1-like distribution (blue) and S2-like distribution (red). The data used for building the PDFs was simulated for the first mock data challenge of LZ and did not include pulse-shaping effects from the detector or electronics response.

as

$$pS = \frac{aft50 - aft05}{aft95 - aft50}. \quad (6.4)$$

The choice of leaving out pulse area as a parameter was motivated by the strong correlation between this variable and the energy scale of the interaction that produced the pulses considered. Since the classifier had the goal of being energy-independent, pulse area was replaced by the pulse asymmetry parameter defined in Equation 6.4. This new parameter aimed to capture the skewness of a pulse by returning the ratio between the time it took the pulse to reach approximately half of its area and the time corresponding to the remaining area. Since S1 pulses have a fast rise time and slower decay time, while S2 pulses are mostly symmetrical, this variable is expected to provide some discrimination power. This can be clearly seen in Figure 6.12.

The pulse parameters are paired to produce 6 distinct 2-dimensional phase-space representations. The choice of bivariate PDFs is twofold: marginal PDFs have high degeneracy and limit the choice of good discrimination criteria, and higher-dimensional PDFs would suffer from the *curse of dimensionality* (see Section 7.2.1), where the lack of statistics would impede the production of smooth and general PDFs. Therefore, 2-dimensional PDFs are a good compromise between



simplicity and robustness.

The 2-dimensional PDFs are built using kernel density estimation (KDE), which is an alternative to binning. With KDE, each PDF is constructed by taking the sum of 2-dimensional Gaussian kernels centered at each data point of a given class. The dynamic range of the PDFs is optimized for the class and parameters being used in each parameter 2D representation. The effect of using a Gaussian kernel instead of simply binning the data within a grid is that statistical fluctuation in the PDFs are smoothed and minute features of the scale of the grid minimal separation can influence a wider space in the PDF. This avoids having sharp transients or discrete aberrations in the PDFs, which ensures that neighbouring pulses do not get widely different responses for the same PDF. It also helps in extrapolating low statistical regions in the PDF phase-space like population edges, that would otherwise suffer from statistical fluctuations. The result is a smooth PDF that is a good generalization of the data.

The COMPACT algorithm is summarised in Algorithm 1.

---

**Algorithm 1:** COMPACT algorithm

---

1. Get a pulse object  $x_i$  and retrieve parameters  $p \in [TBA, pF10, pS, H2L]$
2. For each pulse class  $k \in K = [S1\text{-like}, S2\text{-like}]$  do:
  - (a) For each of the 6 representations  $\{p, q\}$ , with  $p, q \in [TBA, pF10, pS, H2L]$  do:
    - i. Read  $[p(x_i), q(x_i)]$  and evaluate  $P_{\{p,q\}}(y_i = k) = PDF_{p,q}(x_i(p), x_i(q))$  by interpolating the PDF at the point  $[p(x_i), q(x_i)]$ .
    - ii. Convert  $P_{\{p,q\}}(y_i = k)$  into an instantaneous probability (relative frequency)
  - (b) Calculate the overall probability of the pulse  $x_i$  being of class  $y_i = k$ :

$$P(y_i = k) = \prod_{p \neq q} P_{\{p,q\}}(y_i = k)$$

3. Calculate the overall probability of the pulse  $x_i$  not being of class  $k$ :

$$P(y_i \neq k) = \sum_{h \neq k} P(y_i = h)$$


---

### 6.2.4.1 Results from COMPACT

The module was implemented in LZap early on 2016 and remained the primary pulse classification module until early 2017. The benchmarking of the module was performed using 1001 simulated pulses randomly selected from each class – S1-like includes S1 pulses as well as MPE pulses, S2-like includes S2 pulses, SE and S2 tails. These pulses were drawn from the same dataset used to build the PDFs and is not a cross-validation of the algorithm. For MDC1 data, COMPACT reached an S1 efficiency of 100%, with all 1001 pulses being correctly identified.

The S2 efficiency was calculated to be 99.60%, with 3 pulses being identified as S1-like and a fourth pulse being classified as Other. These results were limited by the available statistics at the time. An early independent test using LUX simulated data and  $10^4$  pulses, of which a fraction was reserved for out-of-sample testing, yielded an estimated S1 efficiency of 99.97% and an estimated S2 efficiency of 99.96%.

#### 6.2.4.2 Final remarks on COMPACT

The strengths of this method arise from the unbiased nature of the PDFs, which automatically account for all data features, the proven potential for identifying pathologies in the data and the ability to convert the PDF response into a probabilistic measure of the compatibility of the pulse with the classes considered.

Despite the success of the COMPACT method at classifying simulated pulses in the first mock data challenge, the inability to use the MC-truth at pulse-level to generate reliable PDFs for MDC2, as well as the overwhelming amount of irregular pulses following large S2s rendered this method obsolete.

The large number of irregular pulses that overpopulate most events, often several orders of magnitude more common than real S1 and S2 pulses, resulted in the contamination of the PDFs of the main pulse types to the point of shadowing any relevant pulse features. The spurious pulses are a symptom of the realism imposed in the simulated data and it is assumed that these issues will become more prevalent in real detector data.

The COMPACT algorithm is strongly affected by density changes within the PDFs. Consider the following scenario: a given class of pulses contains several distinct populations that are to be assigned the same label – the class label – and one of those populations happens to be more common than the remaining by some large factor. Any pulse from the less frequent populations will have a relatively weak response to the PDF of their respective class, when compared to a pulse of the same class but member of a more common population. Considering the importance of the populations to be the same when the PDFs are built inadvertently results in the strength of a given pulse to be dependent only on the density of pulses of that same class in the neighbouring regions of the phase-space it occupies. Since some populations have a larger impact on the overall result of the data analysis chain beyond classification, sampling each population equally in the construction of the PDFs leads to instabilities in the classification and suppression of relevant features.

This was the main reason that led to changing the pulse classifier algorithm to HADES, since the latter was able to handle these uncommon pulse populations in a robust and transparent way. As an example, real S2 pulses, S2 tails and electron trains permeate the data, but the SE and SE-sums in a typical event outnumber the real S2s by two to three orders of magnitude. This effect is clearly visible in Figure 6.11, where the density of pulses in the main S2 phase-space is much lower than the neighbouring S2-tail and e-train structures/populations. For COMPACT, the differences in density between pulses in the same broad class are imprinted into the PDFs for

that class, and the more frequent pulses end up shadowing the less frequent ones. The cut-based analysis of HADES allows for fine tuning of selection criteria that are unaffected by different population densities.

A possible revival of the COMPACT algorithm could be based on the production of PDFs using a multidimensional logistic regression algorithm that selects the entire phase-space regions where a class dominates over the other. This can be accomplished by simplifying the classification problem into a binary problem, considering only the generic S1-like and S2-like classes like the original COMPACT algorithm, and fitting a set of logistic functions to the data in order to best separate the two classes over a selected set of pulse parameters. This approach will be tested in the future.



# Chapter 7

## Machine Learning methods for Classification in LZap

This chapter focuses on the development of Machine Learning (ML) applications for classification of pulses in LZ data. The goal of this work is to find minimally-biased solutions for ML algorithms that can achieve high classification accuracy. Most of the effort will be directed to improving the HADES algorithm, the current *PulseClassifier* algorithm in LZap and explained in Section 6.2.3. The secondary objective of this work is to explore ML algorithms that may complement the heuristic models like HADES or eventually substitute them in the processing framework of LZ.

The ML models developed in this work are not required to outperform the heuristic algorithms currently available. The tools used to build and train a ML classifier can inform the decisions made when designing a specialized classification algorithm, assisting the development of heuristic algorithms such as HADES. As an example, the ability of some ML algorithms to rank the features of the data by discrimination potential will be explored in detail later in this chapter. These feature importance rankings are not the final goal of the ML algorithm, but rather a consequence of building and training the ML classification model, and a much welcomed one. Another advantage of ML methods is their ability to extract valuable information directly from the data, such as the most efficient way to partition it based on its labels or the detection of outliers.

One of the motivations for using ML techniques on the pulse data of LZ is to mitigate potential systematic problems of HADES. The main function of HADES is to be the standard classification algorithm for LZap and a benchmark to future classifiers. Being a simple, fully heuristic and greedy algorithm by design, HADES is very susceptible to human and data biases. It is possible and desirable to optimize HADES by exploring its limitations and the particularities of the data

itself with ML techniques.

Three ML methods were explored in this work: Section 7.3 focuses on *Gaussian mixture models* used for clustering analysis; Section 7.4 explores *random forests* as both a data mining tool and as a classification tool; Section 7.5 examines *artificial neural networks* and their implementation in a dedicated classification tool. These methods were chosen due to their flexibility, robustness, performance and generalizability, while also being readily available in open-source analysis packages: the neural networks were implemented using *Keras* [C<sup>+</sup>15] and the remaining methods were implemented using *scikit-learn* [P<sup>+</sup>11]. Some methods used the results obtained by HADES as an approximation to the real classification labels. A full discussion on the validity of this approach is provided in Section 7.5. The remaining models were developed with an unsupervised learning approach, with the use of the clustering algorithms mentioned above. More details on the implementations and training of each ML model can be found on their respective sections.

## 7.1 Introduction to Data Science

Science nowadays relies frequently in massive volumes of data. The ATLAS experiment at CERN produces around 300 TB of data per second. The LZ experiment will produce an excess of 1.4 PB in one year [M<sup>+</sup>17] and by the end of the expected 1000-day run it will have around 4 PB of raw data. The necessity of processing such massive volumes of complex data has given way to the field of *Big Data*, to new techniques of data science with focus on processing power with high performance computing (HPC), improvements on statistical treatment of data and highly-specialized data processing algorithms.

The advancements on the many fields of science naturally resulted in new problems with increasing complexity. As an example, the fields of genomics and computational biology have experienced large cross-disciplinary efforts (*Big Science*) to deal with increasingly complex problems that involve large quantities of complex data [HR13]. The Human Genome Project, to date the largest collaborative effort in biology and genomics, is an example of one of such large science problems that accompanied 15 years of computation evolution.

Particularly in modern physics, the abstraction of the formalism used to describe a model has steered physicists away from classical heuristic approaches to problem solving, with these being either too hard to derive or not-applicable at all. The ever-growing complexity of modern fields of study drives the constant improvement of science tools and detector systems, which result in an increase in volume, precision and complexity of experimental data. This resulted in the creation of techniques with the potential to solve a variety of challenging issues, especially more abstract ones. These solutions originated from computational techniques developed over the past decades, along with advanced statistical techniques.

None other set of tools has shown more potential for high-complexity problem solving than *artificial intelligence* systems – systems with the ability to encode a given environment and perform a set of actions that maximize its chance of successfully achieving its goals [Nil98, LH07,

RN95, PMG97]. It is undeniable that the groundwork set by Alan Turing on computer science led to the creation of the modern idea of a computing system. The pioneering work on early computer models and the theoretical and mathematical postulations credited to Turing were so ahead of its time that even today several models and postulations on programmable machine problems use the framework developed by Turing [CD77, Cv13, JWHT13]. Considered to be the father of modern computing science and artificial intelligence [Cv13], Turing went ahead of what the technology of the Second World War could deliver, and postulated over machines so advanced that even today are seen as hypothetical. Around the year 1950, Turing reflected on the idea that a machine could “think” [Tur50, Har08]. He imagined that a machine with enough computation power and with the ability to simulate any computational algorithm<sup>1</sup> was not restricted from reaching a state that could rival the complexity of a human mind. The philosophical and scientific implications of an “intelligent machine” led Turing to ponder over what would mean for a machine to have intelligence, which prompted the idea of the *imitation game*, also known as the Turing test: an exercise with the goal of determining if a machine possessed the ability to exhibit intelligent behaviour indistinguishable from that of a human [Har08]. The original Turing test was a simple setup where a judge would blindly communicate with a machine and a human being, and based on the conversation alone, would have to distinguish between the two. A machine that would prove indistinguishable from a human upon inquiry would be said to possess intelligence. The test is critiqued for not considering several scenarios such as a lack of response or the limitations of the interrogator, or for placing restrictions on the definitions of intelligent behaviour – humans are capable of unintelligent behaviour, and intelligent behaviour is not restricted to humans – and the core issue of limiting intelligence to manifest through speech. Ian McDonald critiqued the test by stating that “any A.I. smart enough to pass a Turing test is smart enough to know to fail it.” [Agg18]. Despite the original idea of the Turing test being widely regarded as insufficient or erroneous, the idea of testing the intelligence of an artificial system and the possible insights that such ideas might provide are as relevant today as they have ever been.

Arthur Samuel, an IBM researcher, developed a “Checkers-player” program in 1959 that was able to learn how to play checkers. Samuel saw games as the ideal starting point to study the ability of a computer to perform cognitive tasks and learn from experience [Sam59]. The Checkers-player used the position of the pieces on the board to determine the best possible move it could make to achieve victory. The move is chosen by analysing the possible moves up to a certain number of turns and using the minimax algorithm [Agg18]. Although a machine that could play a game of checkers is far from an intelligent system, the work of Samuel was pioneering and helped shape AI testing and development for the next decades.

In 1960, Donald Michie devised a simple physical system composed of 300 matchboxes called Machine Educable Noughts And Crosses Engine (MENACE) that could “play” a full game of noughts-and-crosses, or tic-tac-toe [Chi20]. The rules of the game, *i.e.*, the environment of the problem, were encoded in the architecture of the system: each matchbox represented the board of the game at a given state, and every possible state was accounted until the system or the

---

<sup>1</sup>A Turing-complete machine.

adversary reached victory or a draw. Furthermore, at each state the system had the freedom to “play” any valid move by randomly picking a coloured bead (each colour representing a valid move) taken from the matchbox by an operator, setting up the board for the adversary to play next. Michie designed MENACE to have the same number of coloured beads in each matchbox at the beginning, guaranteeing that the operator had equal chance of choosing each next possible move, but by playing several games against an opponent the system would adjust itself to favour the moves that provided the best results. It achieved this by adjusting the number of beads in each matchbox depending on the outcome of the game: if the model was successful in guaranteeing victory more beads of the colours that represented the winning set of moves were added to the respective boxes, but if the model failed those beads were removed from the matchboxes. This method of rewarding winning moves and punishing losing moves ensures that the losing moves become less likely to be picked again, while the winning moves have an increased chance of being selected. After some games the system will tend to the optimal state and will play a perfect game of tic-tac-toe, either winning or guaranteeing a draw.

Neither MENACE nor the Checkers-player program were intelligent systems, but rather highly specialised algorithms tuned for a specific task. They represent one of the first implementations of learning algorithms that could improve a system to the point where it could outperform a human in a cognitive task. MENACE serves as a simple and intuitive model of a machine that is capable of learning, even one made of matchboxes and beads. After the pioneering work of Michie and Samuel, combined with the invaluable work of Turing on theoretical computation and machine intelligence, AI had moved on to tackle more complex problems, mostly in the form of complex games that require some degree of proficiency and training from a human. In 1997, the at-the-time world chess champion Gary Kasparov was defeated by Deep Blue, an IBM computing system designed to read the chess board piece positions and predict up to 12 moves ahead of its opponent [Ert17]. In 2011, the IBM computing system Watson won a game of *Jeopardy!*, an American television quiz show, against two former champions [F<sup>+</sup>11, F<sup>+</sup>10, Ert17]. Watson understood the questions of the quiz, read by the host of the show, and formulated an hypothesis containing several possible answers, obtained by searching in a repository of information, and weighing the responses based on several criteria in order to chose the most likely answer. More recently, in 2016, the DeepMind Technologies AlphaGo computing program proved to be capable of besting top world players of the game of Go, one of the most complex board games. The AI not only performed at the level of its human adversaries, but also produced unexpected moves in-game that turned out to be advantageous in the long run. The remarkable move 37 of the second game of AlphaGo against Lee Sedol left most the top Go players and comentators completely baffled, with Fan Hui, three-time European Go champion, even calling it both “beautiful” and “not a human move” [Met20]. The AlphaGo program learned the rules and gameplay of Go without being exposed to the rules of the game explicitly, but rather through advanced deep learning techniques, recursive learning and by playing the game against alternative iterations of itself [Ert17].

AI research and development nowadays has shifted its focus towards human-machine interfaces, machine vision, machine-environment interaction, chatbots and the internet of things, with each of these topics being worthy of their own discussion. Society seems to be getting progressively



immersed in AI systems. Large communication corporations, device manufacturers and social media platforms are actively developing AI tools, such as device assistants that respond to voice commands from the user and interpret data on-the-fly. The advancements of AI technology have changed the way people communicate with their devices, control real world systems or even are exposed to content on the internet.

The advantages that AI systems can provide seem to be driving society to seek new solutions to problems past and future, obtained by a “mind” that can solve problems differently from a human mind. Whether such machine would be considered an intelligent system is strongly debated [LH07]. The general consensus amongst AI researchers seems to be that a system that can learn to interpret its own environment and act on it through learning has a certain degree of intelligence, but if the skills that led to the understanding of the tasks the system can perform are not generalizable and transferable to other environments and problems then the system is not an artificial general intelligence (AGI) – a machine with the potential to learn anything [Vos20].

### 7.1.1 Artificial Intelligence and Machine Learning

Arthur Samuel coined the phrase *Machine Learning* (ML) in 1952 to describe the techniques used to train his Checkers player program [Sam59, SB18]. The terms *machine learning* and *artificial intelligence* diverged, with the former no longer being synonym with methods for training the latter. Artificial intelligence nowadays refers to the study and development of computer systems able to perform tasks that require human intelligence, while machine learning describes how to construct self-learning algorithms capable of automatically improve with experience and that can ultimately be used to solve practical problems [Mit97].

Most of the ML framework available today arose from the study of finite systems with intrinsic high complexity, that allows them to have the plasticity to fit an enormous range of problems. These highly complex systems, born of the emergent property of simple units that form a connective group or ensemble, can be used to perform basic arithmetic and logical processing of data and try and solve a given problem. Left to chance, the probability of the system reaching a desired state that would allow for the problem to be solved would be inversely proportional to the total number of possible states of the system, and in most practical cases that probability would be close to zero. In order to help the system reach an optimal solution, a series of iterative tests can be devised, where the solution obtained at each iteration is evaluated under a certain metric of performance to inform the system if its current configuration has improved or what can be changed to improve it. The system can then be tuned towards a better state than the previous one, eventually approaching a close-to-optimal state. This process is called “training” and it is the cornerstone of machine learning techniques. The algorithm is said to have “learned” an interpretation of the problem by being modified to fit that problem against a desired solution. The potential of machine-learning techniques to learn and solve complex problems has made these methods quintessential in big data and modern science applications.

### 7.1.2 The Learning Process

The training of a ML algorithm can be accomplished in different ways, with the type of training having a deep impact on performance. Therefore, this is often the first design choice made when building a ML algorithm. The possible choices of the training settings revolve around three main attributes [Mit97]:

1. Learning through direct or indirect feedback. The learning process can be subdivided into three main categories: *supervised*, *unsupervised* and *reinforcement learning*. Supervised learning is a training method that uses labelled data to adjust the free parameters of the ML algorithm in order to fit its outputs to the data, using the data labels as target values. Unsupervised learning, on the other hand, does not depend on data labels, and instead the algorithm focuses on identifying relations between data objects, data trends and features. Supervised and unsupervised learning will be at the core of this work and will be discussed in detail below. Reinforcement learning is the method of training the algorithm by providing criticism on the quality of a solution, but not on how to improve it. The algorithm is expected to improve the solution by iteratively exploring the solution space [IS17].
2. Degree of control over the training data by the learner algorithm. Depending on the type of learner or the nature of the data, the training can be managed by the learner to some degree. On a supervised learning setup, the learner may divide the data into batches with limited samples or even limited attributes. Or the learner can generate new training data if the rules for doing so are available to it, and ask the supervisor for an opinion or try to compare the obtained response with different versions of itself.
3. Representativeness of the training data regarding future test examples. Whether the data used for training is representative of all the possible data examples that may be encountered in the future, after training. If not then the algorithm will likely not reach a state capable of handling some possible examples, *i.e.* will not be able to generalise, which will impact the accuracy of the model. Generalization is a very important property of any predictive or generative model.

#### Supervised learning

Supervised learning can be defined as the process of training an algorithm to predict or estimate an output by providing labelled input data, with the labels representing the target values for the output of the algorithm [JWHT13]. If the labelled data is categorical, *i.e.*, the object labels are one of a list of discrete numeric values or classes such as {`person`, `car`, `bicycle`} or {`alpaca`, `llama`, `huarizo`}, then the task is designated as *classification*. On the other hand, if the labels are numerical values and part of a broad number set, the task is described as *regression* [Bra16]. In summary, quantitative results call for a regression analysis, and qualitative results call for a classification analysis.

Generally speaking, the learning process involves updating the internal parameters of the model, whatever these may be, in a way that improves the relation between its output and the target values provided by the labels. This improvement is often quantified by some gain function that needs to be maximized or by a generalized error rate that should be minimized. The ultimate goal is for the algorithm to reach a state that can successfully predict the labels of the data and generalize to the set of all possible inputs [IS17].

To quantify the performance of a supervised learner, the simplest metric of performance is the accuracy of the model,  $acc$ , given by the fraction of correct predictions. In regression analysis, however, for a prediction to be considered correct it is often not required to be numerically equal to the expected result, since the output of the learning model maps to a continuous set. Instead, the most useful metrics of the performance of the model is the overall error rate,  $err$ , that can be given by the mean square errors of the results for each sample or any other standard error function [JWHT13]. The accuracy of the learner can then be obtained by complementarity,  $acc = 1 - err$  [IS17].

The metrics of performance defined in Section 6.2.2 can give an idea of the performance of the model but they do not inform of its ability to generalise to novelty data. Maximizing the accuracy over the training data, or, conversely, minimizing the training error may seem purely advantageous, but in reality a model with zero training error is most certainly overfit to the training data and it is not guaranteed to generalize properly [HTF09]. As the complexity of the model increases with training, it will gradually become more aware of complicated underlying structures of the training data. A learner is said to be overfitting when it starts to model these underlying structures that are not representative of the set of all possible inputs. To avoid overtraining of the model and thus overfitting to the training data, a common approach used throughout this work is to separate the available dataset into at least two subsets of data: a “training set” used to train the model and a “test set” used to monitor the performance of the model. The error rate obtained with the test set, also called the *generalization error*, is not only an indicator of the performance of the algorithm but it can also be used to monitor overfitting, and thus the ability of the model to generalize beyond the training data [HTF09]. If at some point during the learning process the model starts to overfit to the training set, the generalization error will begin to increase. At this point the learning can be halted, thus avoiding overtraining.

Some ML algorithms are more susceptible to overfitting than others. Some limit overfitting by implementing specialised training techniques, while others have inherent architectural advantages that suppress it to some extent. Ensemble methods, explained in detail in Section 7.4.1, are naturally resilient to overfitting since the model is constituted by several “weak” learners with limited learning scope (and thus limited overfitting) [Kle96]. Random forests, a type of ensemble method explored in this work (see Section 7.4), use *out-of-bag* sampling to minimize overfitting of the individual learners during training [HTF09]. Artificial neural networks, a ML model very susceptible to overfitting also explored in more detail in Section 7.5.1, generally use node input dropout as an indirect regularization method to prevent overfitting during training [Agg18, C<sup>+</sup>15].

## Unsupervised learning

Unsupervised learning describes a more data-centred approach where the goal is not to fit a model to the data but to extract information directly from it. This situation is definitely more challenging than in the supervised learning case. On an unsupervised learning implementation, the model uses the features of the data objects to extract relevant information like substructures, clusters of data or outliers, but no data labels are available to be set as target results [JWHT13]. Therefore, most unsupervised ML models are used to extract relationship or structural information about the dataset, using statistical, geometric or similarity criteria. Perhaps the most common families of unsupervised learning algorithms are those that perform clustering analysis, outlier detection or novelty detection. The information provided by these algorithms can help to understand the data model at a level not easily reached by heuristic methods, providing useful information about the multidimensional data distribution that can be used to inform higher level algorithms aimed at modelling the data or even be used to build generative models [Agg18].

In some fields of study, including particle and astroparticle physics, it is commonplace to generate models that simulate real physics processes. These simulations are extensively used as the basis to develop analysis tools aimed at processing real data from observatories and particle detectors. ML techniques used for data analysis tasks often rely in simulation data to train supervised ML models. However, in most situations there is no guarantee that the simulated data is representative of the real data, thus undermining the learning process by introducing some level of bias. This is especially concerning in rare searches where minute differences between simulation and reality can produce biased results and lead to wrong conclusions. Unsupervised methods can help transitioning from simulated data to real data by diagnosing the simulation models and performing data-driven sanity checks, or even better, avoiding using simulated data completely. The simulated data can be used to choose the best models and implementations, but the conclusions are derived from the real data directly<sup>2</sup>.

### 7.1.3 A Crisis of Reproducibility and the Role of ML in Science

Despite the potential of ML methods for problem solving and the successes of their application in modern science, the nature of most ML methods can make the interpretation of the results very challenging or downright impossible. A crisis of reproducibility of results in several fields of science can hinder scientific developments or undermine potential solutions to very important questions that are believed to have an optimal solution.

Several authors have verbalised concern with overuse of ML methods and similar “black-box”<sup>3</sup> approaches to problem solving [All20]. The concern strikes harder in fields where the systems under study have direct impact on society, such as medicine or economics. On fundamental sciences, and particularly in physics, the major concern relates to a variant of this problem: the

---

<sup>2</sup>However, it should be mentioned that using unsupervised techniques in this way does not eliminate all biases in the final results, as the information extracted by the ML models will most likely need to be interpreted or further analysed.

<sup>3</sup>Complex system whose operation is hidden or not readily understood.

inability of mitigating prior biases that condition the training of “black-box” algorithms and may lead to local solutions that do not reflect the optimal solution. This situation could lead to the usage of sub-optimal systems believed to have high precision, efficiency and confidence, to process data and draw conclusions from results that are in some way different from reality. A particularly relevant part of this argument concerns the training of ML algorithms with simulated data. The complexity of the ML algorithm can result in this picking up any tiny imperfection or biases of the simulated data that are not present in the real data, thus having an intrinsic misrepresentation of the target data that can be hard to mitigate. This issue is more prominent in systems trained in a supervised manner, where training data is generated artificially to mimic the real system, which is usually the case with simulated data.

Modelling reality often involves the simulation of stochastic processes, where realism is dependent on the quality of the analysis as well as the data simulated. Monte-Carlo methods are generally used in Physics and other areas of science to model such processes. The reliability of MCtruth<sup>4</sup> information in training algorithms meant to process real data is very debated. Especially when the goal of the analysis algorithms is the identification of sporadic features within the data, picking up the faintest imperfections between training and real data can compromise a physics result. This is the case for many state-of-the-art physics problems that revolve around rare event searches, as is the case of dark matter direct detection experiments or rare decay search experiments (for more details on these types of searches see Chapters 4 and 5, respectively).

Several methods can be implemented to mitigate biases generated from the mismatch of real and simulated data, starting by reducing the dependence on simulated data and performing the main development of analysis tools directly with real data. Other statistical methods can also be implemented to ensure that the risks of using imperfect data are minimized or at least understood.

In this work, it is recognised that these are important questions that need to be weighed when using “black box” models for analysis. The ML methods explored in this work were trained and tested with simulated data without MCtruth labels. As mentioned previously in Section 6.1, the mock data used to simulate the LZ detector does not contain MCtruth information at the level of the individual pulses, despite the information being available in the form of the generating process of the individual simulated photons that produce the pulses. The decision of not including the truth information at the pulse level had the goal of replicating the usage of real data for the development of a low-level processing framework.

## 7.2 Overview of the LZap Dataset

The algorithms developed in this work, including the heuristic algorithms HADES and COMPACT, were tested using a wide range of datasets, mostly chosen based on availability. This work spans over 4 years, having started in early 2016 and following most of the R&D phase of LZ,

---

<sup>4</sup>Monte-Carlo truth refers to the information of the underlying model that a given Monte-Carlo simulation uses to generate random-sampled data.

so it was inevitable that the data available changed significantly throughout this period. At the start of this work the only data available was from LUX simulations, that was used to develop most of the COMPACT algorithm and to test some ML approaches like neural networks (more details are provided at the end of this section). Starting in mid 2017, the mock data challenges (MDCs) set specific goals for the simulation and processing tools, resulting in a steady stream of even more realistic LZ simulated data that could be used to develop the classifier algorithms presented in this work.

The developments on simulation and data processing tools across all MDCs resulted in simulated data increasingly more complex and realistic. As discussed in Section 4.4, the simulations of the LZ detector are obtained using the GEANT4-based BACCARAT package and processed with the DER package, that recreates the effects from acquisition electronics [A<sup>+</sup>21]. The data is then processed by the LZap framework where the classification module resides, as explained in Section 6.1. The data used by the classifier algorithm is obtained from the *PulseFinder* and *PulseParametrizer* modules in the form of the reduced quantities (RQs) presented in Appendix B. These RQs are the attributes of the individual pulses (data objects) and will be referred to as “pulse parameters”.

The results presented in this Section focus on LZ simulated data obtained with LZap version 4.5.1<sup>5</sup>, released on the 8th of October 2019. This dataset will be designated as “LZap dataset” throughout this work. Any other versions of LZap and BACCARAT used in this work are explicitly mentioned. The data was obtained by selecting a random set of files processed by LZap 4.5.1 from simulated data corresponding to the LZ commissioning phase. This phase was chosen because it features several background and calibration sources that will be present in LZ, providing a reasonable variety of simulated interactions, event topologies and energies.

A pre-selection of pulses was made to exclude pulses with bad RQ values, such as negative total areas or non-positive pulse height and length. The criteria for excluding these instances are the same as the ones used in the HADES algorithm (see Section 6.2.3). All single coincidence pulses were also excluded from the data, as these are automatically identified in any analysis. This focuses the effort of the classification algorithms on the remaining pulses that are not so easily identifiable [Bra16]. It is important to notice that some contamination is expected due to the issues with the coincidence RQ calculation explained in Section 6.2.1. Based on the results of HADES, pulses with unit coincidence correspond to 42.7% of this entire dataset. After rejecting the single coincidence pulses, a total of  $10^6$  pulses were randomly selected from the entire dataset, which correspond to an average of 25402 distinct LZ events. The full dataset roughly corresponds to 10 minutes and 35 seconds of data taking (LZ expects an average event rate of 40 Hz, corresponding to an average of  $3.5 \times 10^6$  events per day [A<sup>+</sup>20g]).

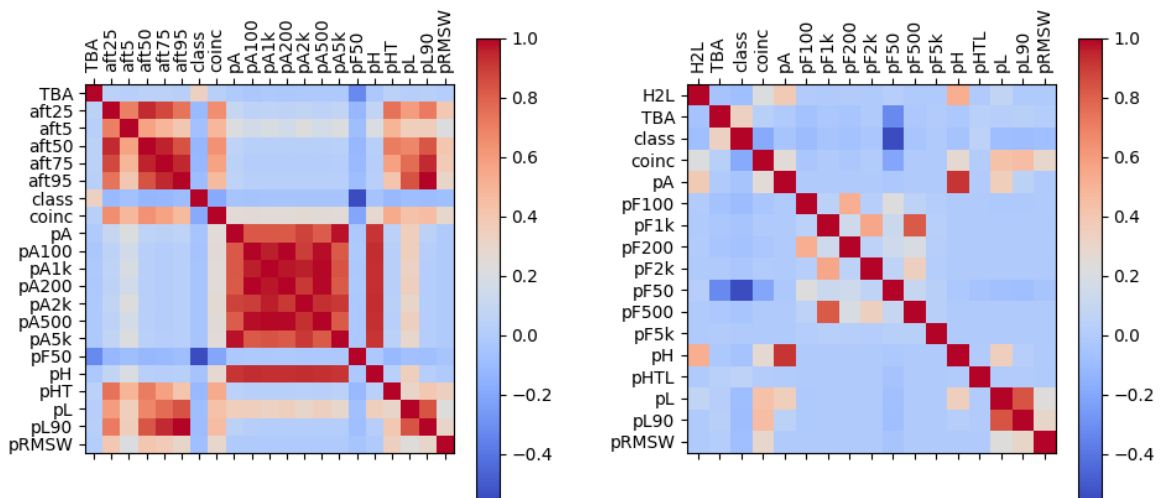
The original data processed with LZap contains 35 pulse-level RQs, of which 20 were selected for this analysis. This selection was based on past experience on pulse classification for the LUX experiment and on gained knowledge throughout this work. Along with the selected RQs, some bookkeeping information was also recorded for each pulse for traceability purposes: the

---

<sup>5</sup><http://teacher.pas.rochester.edu:8080/wiki/bin/view/Lz/LZapReleaseNotes>

simulation file name, LZap-processed file (RQ file) name, event ID on the raw file, event ID on the RQ file, and pulse ID. These variables allow the pulses selected on the different analyses to be visualised with the LZ event viewer, or the verification of individual pulse RQs if needed. Furthermore, the classification assigned to each pulse in the dataset by the HADES classifier algorithm, described in Section 6.2.3, is encoded as an integer variable and recorded as well. These specific labels from HADES will be used as prior classifications by some algorithms ahead, unless specified otherwise.

Figure 7.1 shows the correlation matrices of this dataset, for the raw RQs (left) and composite pulse parameters (right). The composite parameters are obtained by combining the pulse RQs in meaningful ways. More details on the pulse parameters can be found in Appendix B.



**Figure 7.1:** Correlation matrices for the non-normalized LZap dataset. The left plot shows the correlations between all pulse parameters (RQs) available in LZap. The right plot shows the correlations between the composite pulse parameters used extensively in this work.

From the correlation plots it is clear that some variables are strongly correlated, like the total pulse area ( $pA$ ) and any of the pulse areas integrated over different time windows ( $pA100$ ,  $pA200$ ,  $pA500$ ,  $pA1k$ ,  $pA2k$  or  $pA5k$ ). Using these strongly correlated variables together in any analysis would result in little information gain when compared to the usage of only one of the variables, with the added burden of having more variables to work with. However, these variables can be combined to obtain very useful parameters. In the case of pulse areas integrated over different time windows, these can be combined with the total pulse area to obtain the fraction of area within a given time window anchored at the start of the pulse, also called the prompt fraction at  $X$  ns, or  $pFX$ , with  $X = \{100, 200, 500, 1k, 2k, 5k\}$  ns in this particular case. These variables provide insight on the growth of the pulse early on, and are extremely useful to separate S1-like and S2-like pulses.

Of the 20 pulse parameters recorded for this dataset and presented in Figure 7.1, only a subset

is used to test and train the several algorithms explored. This constraint of the parameter space will help to mitigate some data-driven limitations detailed below.

### 7.2.1 Data-driven Limitations to ML Algorithms

There is no such thing as a perfect dataset. Different issues related to the intrinsic properties of the data can compromise the results of any data processing algorithm unless acted upon. The most common issues relate directly to erroneous values of the parameters of some data members, or the existence of data objects that were not expected in the dataset, *i.e.*, outliers. These are called “accidentals” in the data and require some level of preprocessing to identify and exclude them properly, if needed. Errors<sup>6</sup> while writing the data into file or poorly calculated variables are the main causes of these accidentals. An example related to the dataset used in this work is the flawed calculation of the coincidence RQ, mentioned previously.

Apart from accidentals, the dataset might not originally be in a state that is usable by most algorithms. Data inhomogeneities, non-linearity, density gradients, under-represented populations and parameter-space dimensionality<sup>7</sup> need to be accounted for, and not all algorithms respond to these issues equally. It is essential to perform a careful mitigation of the problems that may arise from these data features.

#### Data noise and parameter accidentals

Data noise often finds its way into a dataset, either through miscalculations, transcription errors, poor variable precision, typos or human error. Noise can be found on different levels in the dataset: in the values of parameters, in the object labels if available, or as data objects themselves [Bra16, Kub15].

Value noise can be generated when the data is being produced, handled or stored. There are three types of value noise: stochastic, systematic and arbitrary [Kub15]. Stochastic noise can be related to the randomness of any stochastic process of the system that produces the data, or by some expected variability of the values themselves. The LZ simulated data is obtained by simulating several stochastic processes related to particle and nuclear physics, such as nuclear decays and interaction of radiation with matter. Furthermore, the simulation of the electronics response of the LZ data acquisition will also introduce noise from the electronics that will be visible in the waveforms. No mitigation of these potential noise sources is performed in this work, and the potential failures<sup>8</sup> of the data generation tools are ignored, as they are outside of the scope of the classification efforts. Systematic noise is mainly introduced when a processing algorithm is poorly tuned, resulting in an atypical data trend. The *PulseFinder* and *PulseParametrizer* modules have known systematic failures already explained in the previous

---

<sup>6</sup>Colloquially referred to as “bugs” in the data.

<sup>7</sup>The dimensionality  $D$  of a dataset relates to the number of independent features that can be used to represent one element of the dataset (data object).

<sup>8</sup>The BACCARAT and DER packages are known to have issues that are being addressed.



Chapter, namely a systematic splitting of pulse substructures and biased RQ calculations, respectively. These accidentals are complicated to trace to a particular origin, but are expected to be negligible in this dataset.

Regardless of their origin, variable noise will result in data objects whose attributes have unexpected values, either unrealistic (bad parameters) or simply outside the expected range (range outlier). The most common bad parameters found in the pre-processed data are non-positive total pulse area, area fractions outside the unit interval  $[0; 1]$  or other geometric variable with non-physical quantities. Most of these values are produced by parametrizing complex structures in the simulated waveforms like baseline noise or signal overshoot. A “range outlier”, on the other hand, is an instance of a parameter whose value is realistic but falls outside the expected parameter range. An example of these two types of variable noise can be found in the dataset in the form of the observed atypical values of the top-bottom asymmetry ( $TBA$ ) RQ of some S2 pulses. The populations produced by these parameter values can be observed in the right-hand plot of Figure 6.4. There are two discrete populations with  $TBA = \pm 1$  at higher pulse areas on either side of the main S2 population. Their origin was identified as a problem at assigning a value to the *topArea\_phd* and *bottomArea\_phd* RQs by the *PulseParametrizer* algorithm. This results in unpredictable RQ values that are often of the order of magnitude of the maximum range of a 16-bit floating point. This issue can be considered either a bad parameter or a range outlier problem. If one of these RQs used to calculate the  $TBA$  value has an arbitrarily large value, the  $TBA$  RQ will collapse into either  $+1$  or  $-1$  value. The *topArea\_phd* and *bottomArea\_phd* RQs are excluded from the dataset because the information contained in these variables can be retrieved from the  $TBA$  and *pulseArea\_phd* RQs.

It was mentioned in the previous Section that the HADES results were stored as label priors in this dataset. Despite the 98.6% overall accuracy of HADES, some misclassification of pulses is expected in this dataset. If the HADES labels are implemented in a supervised learning exercise, it is natural to think of these misclassifications as “class-label noise” [Kub15]. Methods for mitigating these misclassifications will be explored in detail whenever the HADES labels are used.

## Outliers

Data outliers are objects that are not associated with the underlying data model [Bra16]. These objects are also considered a type of noise, but at the level of the full dataset instead of the values of individual attributes, *i.e.*, range outliers. These outlier objects often integrate smoothly with the rest of the data and have well behaved attributes, or would otherwise be considered spurious objects caused by parameter accidentals. However, a large fraction of outliers is produced by less severe data noise or systemic data processing inaccuracies that cannot be immediately identified.

There are some identified outliers in the LZap data used in this work, namely SPE and MPE with incorrect coincidence values, and SE splits or other pulses with boundary issues, among others. The coincidence miscalculations, or rather the incorrect coincidence criterion, is a systemic problem of the *PulseParametrizer* module, while the over-splitting and inconsistent pulse

boundaries were a systemic problem of the *PulseFinder* as it was being adjusted.

The presence of SPE and MPE pulses with incorrect coincidence RQ values is not expected to impact the classification results significantly. The SPE pulses with coincidence greater than 1 are expected to be identified as S1-like pulses, due to the similarities in the RQs between these two classes. The same is true for the MPE pulses with coincidence greater than 2, that are expected to be organically assimilated into the S1 populations. Therefore, no outlier analysis is considered in this work with the specific goal of excluding these pulses from the dataset.

The pulse boundary issues and over-splitting by the *PulseFinder* module introduced some outlier populations in the dataset, one of the most problematic being the SE-splitting population that was studied extensively in this work. As explained before, the SE-splitting pulses produce an outlier population that is troublesome to remove with simple heuristic analyses, under the risk of diminishing the efficiency of classification of small S1 pulses. However, a choice was made to not exclude this population from the dataset. Instead the identification and classification of these pulses was set as one of the main goals of the classification algorithms being developed.

### Density, representativeness and sample biasing

Inhomogeneities in the data can sharply increase the complexity of the analysis performed by a ML algorithm. These inhomogeneities can be produced by density fluctuations in the parameter phase-space, differences in the abundances of data classes or data non-linearities, among other causes.

Some algorithms are very sensitive to local density gradients, especially when the computing of distances between data objects is involved [Agg15]. Most clustering algorithms use distance-based methods to infer the relation between data objects. This practice is tightly related to outlier detection algorithms as well. Both of these techniques are expected to produce a strong response to a density hotspot in the data. Conversely, less dense regions in the vicinity of a density hotspot will be “outvoted” and likely ignored or mistakenly associated with their denser counterparts, despite potentially containing distinct information. Most algorithms will focus on the stronger data trends and will be less sensitive to faint local features. The solution often involves implementing some sort of fine tuning method to these algorithms to fight this density shadowing effect. Algorithms that adjust to local density variations are expected to outperform those that do not [Agg15].

In LZ data, the tight parameter phase-space range of SE pulses tied with their expected high abundance produces a high density region clearly visible in Figure 6.4. This will influence the performance of some algorithms explored in this work, with the most noticeable case being the decision-tree-based algorithms. The high abundance of SE pulses is also tightly related to the problem of asymmetric representativeness of data classes, explained below.

The over-abundance of one or several object classes in the data when compared to less prominent species is another common example of inhomogeneity found in the data. If populations within a dataset have widely different representativeness, an algorithm may devote a disproportionate

amount of attention to the population that is more statistically significant, to the detriment of smaller valid populations that can be seen as outliers. This may be an advantage in some specific exercises, if the uncommon classes are not as important as the most common ones, but this often is not the case. Clustering and classification exercises often require the separation of less prominent data features that are as important as the most common ones. In ML branches like novelty detection or anomaly detection, the focus is on identifying the uncommon data subsets.

Several biasing methods can be implemented to mitigate the effects of density regions or data misrepresentation, of which the most common are *sample weighing* and *stratified sampling* [Agg15].

Assigning weights to individual data samples is a simple and robust way of highlighting some subsets of data in order to inform the algorithm of what data to focus on, and most ML algorithms can handle sample weights natively. The weights can be assigned heuristically by the analyser using some pre-selection of the relevant subsets of data, or they can be calculated using the intrinsic properties of the data. A very practical example is density-biasing, where sample weights are inversely proportional to the local sample density. This way, regions with higher density will not contribute as strongly to the final result. This effectively flattens the response of the algorithm in the full parameter phase-space. The data distributions presented in Figure 6.4 convey the extreme density variations present in the dataset being studied. The highest density region being that of the SE population, the most common pulse type in LZ data and with the most constrained parameter values. Several algorithms can be affected by the over-abundance of SE pulses in the data, often leading to incorrect conclusions. An instance of this issue is explored in detail when feature importance rankings are calculated using tree ensemble methods (see Section 7.4.1). It is important to point out that weights can be the source of severe bias in the final results. Proper management of weights is essential to ensure that the right information is obtained with the algorithms.

Stratified sampling is a method designed to compensate for misrepresented data by sampling some parts of the data more than others, either by selecting the relevant labels *a priori* or by selecting regions of the parameter phase-space that contain the most relevant subsets [Agg15]. No stratified sampling strategy was implemented in the work presented in this document. However, during the initial development stages of HADES at the MDC3, due to the low density and under-representation of S1 pulses with respect to S2 and SE pulses, a stratified sampling was performed to produce a working example dataset with roughly equal number of S1, S2 and SE pulses.

### Curse of dimensionality

It seems reasonable to assume that to solve a given problem involving a given dataset, the greater the number of discriminant features the easier it is to find a satisfactory solution. After all, if some degeneracy is still present in the solution, having additional discriminant features is expected to improve that result, even if only marginally. However, this is not entirely true from a computational standpoint. As the number of discriminant features increases so does the volume

of the parameter phase-space, often leading to the data becoming too sparse for the algorithm to extract useful information or find trends that are statistically significant. This is known as the *curse of dimensionality* [Kub15]. Any algorithm is expected to falter when the dimensionality of the feature space increases beyond a certain point, requiring considerably more computing time to reach a solution or even making the solution out of the reach of the algorithm entirely.

The “curse” in this context comes from the fact that the data density decreases exponentially with increasing dimensionality, placing a heavy burden on the amount of data required to properly represent a high-dimensional dataset. The quantity of data objects within must be enough to permeate the parameter phase space with enough density to allow a processing algorithm to operate properly, and the higher the dimensionality of the data the larger must the dataset be to ensure a fair data representation.

The *curse of dimensionality* is extremely important to take into account in clustering or classification exercises. The ability to partition the data into subsets that share similar features may be compromised if the number of features is large enough to impede the algorithm from determining the similarity between data objects. As more and more features are considered by the algorithms, the likelihood of two data objects sharing some of the features becomes lower and lower. Some of these algorithms are very sensitive to data density variations, as mentioned above, and often rely on intrinsic distances that must be estimated from the data. There are many clustering algorithms centred in the notion of “core distance” between data objects. These algorithms often evaluate the similarity of a data object to either their neighbouring objects or to a center point defined within the scope of the feature phase-space. But as discussed above, the intrinsic distances between data objects can easily overshoot the natural scope of a single feature. This may not seem damaging at first, but if the “core distance” used by one of these algorithms exceeds the natural scope of one or more data features in the dataset, then for all intents and purposes a marginal distribution of the dataset exists where most if not all data members are within said distance, thus making the partitioning more challenging.

To illustrate the problem, one can conceive a dataset that contains a finite number of data objects,  $\mathbf{x}_n, n = 1, 2, \dots, N$ , each with  $D$  independent features and with each feature value contained within the unit-length interval,  $[0; 1]$ . This  $D$ -dimensional dataset is a generic representation of any dataset that was preprocessed with a simple normalization of the range of the individual parameters, which is not expected to change any properties of the data since the parameters are assumed to be uncorrelated. The Euclidean distance,  $r$ , of a data point  $\mathbf{x} = \{x_1, x_2, \dots, x_D\}$  to a different data point  $\mathbf{x}' = \{x'_1, x'_2, \dots, x'_D\}$  over the  $D$ -dimensional feature phase-space is given by

$$r^2 = (x_1 - x'_1)^2 + (x_2 - x'_2)^2 + \dots + (x_D - x'_D)^2. \quad (7.1)$$

Since each feature is only defined within the unit-length, the maximum distance between points  $\mathbf{x}$  and  $\mathbf{x}'$  is  $r = \sqrt{D}$ . Notice that, despite each feature being limited to the interval  $[0; 1]$ , as  $D$  increases the range of distances between data objects is unbounded. For a 10-dimensional dataset like the one described above, the maximum distance within is  $r = \sqrt{10} \sim 3.1623$ ,

while a dataset with 100 independent features<sup>9</sup> could contain objects at an inner distance of  $r = \sqrt{100} = 10$ . With such unbounded distances between data members, it is expectable that any generic algorithm begins to struggle as the dimensionality increases.

The most common approach to mitigate the dimensionality issues is to pre-process the dataset and try to either reduce the number of parameters (feature selection) or find useful projections obtained by combinations of parameters (feature extraction). These exercises fall within the discipline of dimensionality reduction, whose main goal is to find the best compromise between information and dimensionality [Ert17]. Whenever the dataset used in this work is limited to a subset of the available parameters to accelerate learning, it is an example of dimensionality reduction via feature selection. This technique will be used extensively throughout this work, often paired with some method for ranking the data features by importance. More advanced techniques like hierarchical learning and distributed learning can also be implemented to help deep learning implementations in their learning process [Ert17].

### 7.3 Clustering Analysis of LZap Data

The underlying risks of training a ML classifier algorithm with non-representative data drives the usage of unsupervised learning methods. With these approaches, the data used for training is taken from the target dataset directly, avoiding misrepresentation of the problem upon fitting the ML model. By using real data it is guaranteed that all the available information and any possible data trends are present with no underlying hidden biases. Renouncing supervised methods implies that a new process must be introduced to extract information directly from the target data that can then be used by a classification algorithm. A particularly convenient set of tools used to process unlabelled datasets are clustering algorithms. These tools can be the first step to implement unsupervised learning to a specialized classification algorithm. Clustering is very often combined with classification efforts because it allows the data to be partitioned naturally due to intrinsic properties of the data members. For this reason, clustering of raw data is sometimes referred to as automatic classification. Combining clustering and classification methods results in a robust, minimally-biased solution for many problems involving unlabelled data.

Clustering algorithms are used to group data objects according to their proximity in their parameter phase-space [HTF09]. Data objects that have similar parameters can be grouped into an object cluster or population. The elements of a population obtained by clustering share common features that can be used to distinguish them from other data objects in different clusters. Different clustering methods can interpret the data in different ways, and the results from a given clustering algorithm are not generic. Choosing the correct type of clustering behaviour with respect to the data being processed is crucial to ensure that the clustering results are meaningful, as discussed in Section 7.2.1.

An initial consideration when attempting to cluster a dataset is to assess if the number of

---

<sup>9</sup>Quite mild when compared to some domains like text recognition and genomics.

distinct populations is known or if it can be inferred/estimated. Some clustering methods require an initial estimate of the number of clusters with which to group the data, while others can automatically estimate that information from the data [HTF09]. The presence of strong inhomogeneities in the dataset coupled to different data species can produce undesired behaviour in the clustering algorithm, as explained in detail in Section 7.2.1. Understanding the data and choosing the best clustering method for it is therefore a crucial step.

There are numerous methods for clustering data. Perhaps the most common algorithms are K-means, Gaussian Mixture Models (GMMs), Density-Based Spatial Clustering of Applications with Noise (DBSCAN), hierarchical clustering, amongst others [HTF09]. Each method presents advantages and disadvantages depending on the problem at hand [HTF09].

As an example, the simplest of the mentioned methods is probably the K-means algorithm [Llo82], that partitions data objects into  $K$  clusters depending on the squared Euclidean distance of each object to the centroid of each cluster, iteratively updating the centroid of the clusters and repeating the partitioning of the data objects at each iteration until no update can be made. This method requires an initial estimate of the number of clusters in the data, the hyperparameter  $K$ , which may have to be estimated with additional methods or by examining the data. It can also flounder easily if the data contains disproportional populations or a large number of outliers.

DBSCAN is another simple but powerful clustering algorithm that groups neighbouring data objects that are less than a parametric distance  $\epsilon$  apart from any other member of the cluster [EKSX96]: if a pre-determined number of objects is within  $\epsilon$  units of distance from the initial object, a new cluster is created and the objects are marked as its elements, otherwise the initial object is labelled as noise. The clustering stops when no new data object can be added to the current cluster, prompting the algorithm to seek a data object not yet visited and build a new cluster. The process is repeated until all data objects are either members of a cluster or are labelled as noise. DBSCAN has the advantage of performing an estimation of the number of clusters in the data directly, while also being able to identify outliers. However, the algorithm struggles with clusters with varying density, where the ideal values for the distance hyperparameter  $\epsilon$  and the minimum number of initial clustered objects will vary from cluster to cluster. It can also fall victim of the *curse of dimensionality*, where at higher dimensionality a decent value for  $\epsilon$  can be hard to estimate.

K-means and DBSCAN are two very conceptually different algorithms. The success of either algorithm is highly dependent on the properties of the target data, and the choice of which algorithm works best stems from a set of very unique driving factors that the user must consider or test. As discussed in Section 7.3.1, Gaussian Mixture Models provide an interesting set of advantages over K-means and DBSCAN, and will be the main focus of the cluster analysis performed in this work.

### 7.3.1 Gaussian Mixture Models

Mixture models are statistical tools often implemented for density estimation, but they can also be used for clustering and classification exercises [Bis06]. A mixture model can be defined as a multi-component probabilistic distribution, where a finite ensemble of distributions are assumed to originate the full observed data, and from which the unknown parameters of the individual distributions can be derived.

A Gaussian mixture can be defined as a linear superposition of  $K$  Gaussian distributions, in the form

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (7.2)$$

where  $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$  represents a normalized Gaussian distribution with the same dimensionality  $D$  as  $\mathbf{x}$  (the parameter space) and with mean vector  $\boldsymbol{\mu}_k$  and covariance matrix  $\boldsymbol{\Sigma}_k$ , with  $k = 1, 2, \dots, K$  [Bis06]. The mixing coefficients  $\pi_k$  are the weights of the component  $k$  of the mixture model. If the mixture model  $p(\mathbf{x})$  is set to represent a probability distribution, integrating over all normalized Gaussian components  $k$  yields that the mixing coefficients must satisfy

$$\sum_{k=1}^K \pi_k = 1. \quad (7.3)$$

Since both  $p(\mathbf{x})$  and  $\mathcal{N}$  must be non-negative, then it also follows that  $0 \leq \pi_k \leq 1$ , and the mixing coefficients  $\pi_k$  satisfy all the requirements to be probabilities.

The real observation, *i.e.*, the data, is assumed to have been generated by an underlying probabilistic model of the form of Equation 7.2, where each Gaussian component  $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$  is assumed to have generated a subset of elements in the data with probability  $\pi_k$ . Each Gaussian component can be assigned to a state of a  $K$ -dimensional binary random vector  $\mathbf{z}$ , of which only one particular element  $z_k$  is non-zero. This vector  $\mathbf{z}$  satisfies  $z_k \in \{0, 1\}$  and  $\sum_k z_k = 1$ . The marginal distribution  $p(\mathbf{z})$  can be associated with the mixing coefficients  $\pi_k$  such that  $p(z_k = 1) = \pi_k$ , and since the binary vector  $\mathbf{z}$  has  $K$  states for each element  $z_k = 1$ , it can be written in the form

$$p(\mathbf{z}) = \prod_{k=1}^K \pi_k^{z_k}. \quad (7.4)$$

Now it becomes clear that the conditional probability of observing  $\mathbf{x}$  given a particular state  $z_k = 1$ ,  $p(\mathbf{x} | z_k = 1)$ , can be interpreted as the probability of the data object represented by  $\mathbf{x}$  being generated by the  $k^{\text{th}}$  Gaussian component,  $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ . This can also be written as a parametrization of the binary vector  $\mathbf{z}$  as

$$p(\mathbf{x} | \mathbf{z}) = \prod_{k=1}^K \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)^{z_k}. \quad (7.5)$$

The joint distribution  $p(\mathbf{x}, \mathbf{z})$  can be obtained directly by combining Equations 7.4 and 7.5, while the original distribution  $p(\mathbf{x})$  can be obtained by marginalization of  $p(\mathbf{x}, \mathbf{z})$ . In fact, it is guaranteed that a distribution of observed variables can always be obtained by marginalization of an extended distribution that contains the observed variables plus some latent variables. Therefore, it is possible to describe a rather complex marginal distribution as a set of simpler joint distributions over the expanded variable space [Bis06]. In this sense,  $p(\mathbf{x})$  can be conveniently rewritten as

$$\begin{aligned} p(\mathbf{x}) &= \sum_{\mathbf{z}} p(\mathbf{x}, \mathbf{z}) \\ &= \sum_{\mathbf{z}} p(\mathbf{z})p(\mathbf{x}|\mathbf{z}), \end{aligned} \quad (7.6)$$

where the marginalization is performed with the sum over the discrete latent variables  $\mathbf{z}$ . It also follows from Equation 7.6 that for each  $\mathbf{x}_n$  there is a corresponding  $\mathbf{z}_n$ , with  $n = 1, 2, \dots, N$  the index of the data objects in a dataset with  $N$  elements.

Equation 7.6 can be rewritten in the same form as Equation 7.2 by using Equations 7.4 and 7.5, but this time with explicit latent variables  $\mathbf{z}$  included in the model. This result allows the estimation of the parameters of the Gaussian components to be obtained by manipulation of the joint distribution  $p(\mathbf{x}, \mathbf{z})$  instead of the marginal distribution of observed variables,  $p(\mathbf{x})$ .

Using Equations 7.2 and 7.6 together with Bayes' theorem [Bay63], the conditional probability  $p(z_k = 1|\mathbf{x})$  can be written in the form

$$\begin{aligned} p(z_k = 1|\mathbf{x}) &= \frac{p(z_k = 1)p(\mathbf{x}|z_k = 1)}{p(\mathbf{x})} \\ &= \frac{\pi_k \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}. \end{aligned} \quad (7.7)$$

The quantity in Equation 7.7 is the posterior probability of the  $k^{\text{th}}$  component of the model having generated the observed variable  $\mathbf{x}$ . It can be interpreted, quite intuitively, as a measure of the “response strength” of component  $k$  to the observation  $\mathbf{x}$ , being larger when the component  $k$  is compatible with  $\mathbf{x}$  and lower if otherwise. This quantity is also called the “responsibility” of the component  $k$  for explaining the observation  $\mathbf{x}$  [Bis06].

The quantities obtained provide the necessary framework to decompose the data into  $K$  components and to perform a maximum likelihood analysis that would allow for an optimal solution to be found. This solution can be obtained by finding the individual parameters of the Gaussian distributions and the respective mixing coefficients that maximize the log likelihood given by

$$\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{n=1}^N \ln \left( \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_n|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right), \quad (7.8)$$

where  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  is the full dataset of  $N$  data objects with  $D$  parameters. Likewise,



the latent variables form a  $N \times K$  matrix  $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N\}$ . However, the log likelihood is susceptible to singularities of the form  $\boldsymbol{\mu}_j = \mathbf{x}_n$ , *i.e.*, when the mean of a Gaussian component is equal to an observed variable. In the presence of such singularities, Equation 7.8 rapidly diverges when the variance of the  $j^{\text{th}}$  component goes to lower values, effectively collapsing the  $j^{\text{th}}$  Gaussian distribution into  $\mathbf{x}_n$ . Under these circumstances, no stable solution can be found. The divergent behaviour of the log likelihood can be mitigated by identifying the collapsing Gaussian component and assigning new values to the mean and covariance matrix [Bis06]. Singularities are not the only problem with using the log likelihood for Gaussian mixture models. Due to the explicit sum over the  $K$  components of the Gaussian mixture within the logarithm in Equation 7.8, finding a maximum of the log likelihood function often results in the absence of closed-form solutions. To avoid complications with the calculations of the log likelihood, a common approach to fitting a Gaussian mixture model to data is to apply the Expectation-Maximization (EM) algorithm [DLR77, MK07, Bis06] described below.

### Expectation-Maximization for GMMs

The goal of maximum likelihood estimation is to find a state of the system that maximizes the likelihood function, so it follows that the condition that must be met at the optimal state is equivalent to finding the root of the derivative of the likelihood function for a given parameter of the model [Bis06].

The derivatives of  $\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$  with respect to  $\boldsymbol{\mu}_k$  can be written as

$$\begin{aligned} \partial_{\boldsymbol{\mu}_k} (\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})) &= \sum_{n=1}^N \left( \frac{\partial_{\boldsymbol{\mu}_k} \left( \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} \right) \\ &= - \sum_{n=1}^N \left( \frac{\pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} \right) \boldsymbol{\Sigma}_k (\mathbf{x}_n - \boldsymbol{\mu}_k) \\ &= - \sum_{n=1}^N p(z_k = 1 | \mathbf{x}_n) \boldsymbol{\Sigma}_k (\mathbf{x}_n - \boldsymbol{\mu}_k), \end{aligned} \quad (7.9)$$

where Equation 7.7 was used in the last step, since the posterior probability, or “responsibility” of the component  $k$ , appears naturally in the derivative of the log likelihood. The expression for calculating the updated mean vector of component  $k$ ,  $\boldsymbol{\mu}_k^*$ , can now be obtained by setting Equation 7.9 to zero and rearranging

$$\boldsymbol{\mu}_k^* = \frac{1}{N_k} \sum_{n=1}^N p(z_k = 1 | \mathbf{x}_n) \mathbf{x}_n, \quad (7.10)$$

where

$$N_k = \sum_{n=1}^N p(z_k = 1 | \mathbf{x}_n), \quad (7.11)$$

can be thought of as an effective number of elements assigned to component  $k$ . This result is very

intuitive and can be interpreted as the weighted mean of the elements assigned to component  $k$ , with the weighting factor being the effective probability of the component  $k$  explaining the observation  $\mathbf{x}_n$ ,  $p(z_k = 1|\mathbf{x}_n)$ . The result in Equation 7.10 is only defined if  $\Sigma_k$  is non-singular, *i.e.*, if the Gaussian component  $k$  has not collapsed to a single point.

The same reasoning can be made for finding an expression for  $\Sigma_k$ . Setting the derivatives of  $\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$  with respect to  $\Sigma_k$  to zero returns

$$\Sigma_k^* = \frac{1}{N_k} \sum_{n=1}^N p(z_k = 1|\mathbf{x}_n) (\mathbf{x}_n - \boldsymbol{\mu}_k^*) (\mathbf{x}_n - \boldsymbol{\mu}_k^*)^T \quad (7.12)$$

which is of the same form as the expression for the covariance of a single Gaussian component, but again weighted globally by  $N_k$  and element-wise by the posterior probability  $p(z_k = 1|\mathbf{x}_n)$ .

Similarly, the derivatives of  $\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$  with respect to  $\pi_k$  can be written as

$$\begin{aligned} \partial_{\pi_k} \left[ \ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) + \lambda \left( \sum_{k=1}^K \pi_k - 1 \right) \right] &= \sum_{n=1}^N \partial_{\pi_k} \ln \left( \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right) + \lambda \\ &= \frac{\mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} + \lambda \\ &= \frac{1}{\pi_k} \sum_{n=1}^N p(z_k = 1|\mathbf{x}_n) + \lambda \end{aligned} \quad (7.13)$$

where the Lagrangian multiplier  $\lambda$  was introduced to ensure that the condition in Equation 7.3 is preserved. Setting these derivatives to zero and making use of Equation 7.11 and the condition in Equation 7.3, the value of  $\lambda = -N$  is obtained. Reintroducing this result in Equation 7.13 results in the expression for the updated  $\pi_k$ :

$$\pi_k^* = \frac{N_k}{N} \quad (7.14)$$

These equations for the Gaussian parameters and the mixing weights do not provide a direct solution for the mixture model given the observed data, but can be used to iteratively update the parameters of the model until a close-to-optimal solution is reached. This is the heart of the Expectation-Maximization algorithm for GMMs. The update of the posterior probabilities and the parameters of the model is performed in two separated instances, named the expectation step and the maximization step. On the expectation step, the posterior probabilities  $p(z_k = 1|\mathbf{x}_n)$  are calculated from the current parameters of the components of the mixture model using Equation 7.7. The maximization step then uses the new values of  $p(z_k = 1|\mathbf{x}_n)$  and Equations 7.10, 7.12 and 7.14 to estimate and update the parameters of the components.

The Expectation-Maximization algorithm for a Gaussian mixture model is described in Algorithm 2. It shares some similarities to the K-means algorithm. Both clustering methods display an iteration of two consecutive steps, a first one where data objects are distributed across the

existing components followed by a second step where the parameters of the components are tuned with respect to the elements associated with each component. The obvious difference between these two methods is that the assigning of a data object to a cluster is binary<sup>10</sup> in nature for the K-means algorithm, while for GMMs the assignment is probabilistic and continuous, obtained through the posterior probabilities  $p(z_k = 1|\mathbf{x}_n)$ . The K-means algorithm is, therefore, a non-probabilistic limit of the Expectation-Maximization applied to GMMs [Bis06].

---

**Algorithm 2:** Expectation-Maximization algorithm for Gaussian mixture models [Bis06]

---

1. Initialize the parameters of the  $K$  Gaussian components: mixing coefficients  $\pi_k$ , means  $\boldsymbol{\mu}_k$  and covariance matrices  $\boldsymbol{\Sigma}_k$ .
2. Compute the initial value of the log likelihood,

$$\ln p(\mathbf{X}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{n=1}^N \ln \left( \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right) \quad (7.8)$$

3. **Expectation step:** Calculate the posterior probabilities (“responsibilities”)

$$p(z_k = 1|\mathbf{x}_n) = \frac{\pi_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)} \quad (7.7)$$

4. **Maximization step:** Re-calculate the parameters of the Gaussian components using the “responsibilities”  $p(z_k = 1|\mathbf{x}_n)$

$$\boldsymbol{\mu}_k^* = \frac{1}{N_k} \sum_{n=1}^N p(z_k = 1|\mathbf{x}_n) \mathbf{x}_n \quad (7.10)$$

$$\boldsymbol{\Sigma}_k^* = \frac{1}{N_k} \sum_{n=1}^N p(z_k = 1|\mathbf{x}_n) (\mathbf{x}_n - \boldsymbol{\mu}_k^*)(\mathbf{x}_n - \boldsymbol{\mu}_k^*)^T \quad (7.12)$$

$$\pi_k^* = \frac{N_k}{N} \quad (7.14)$$

where

$$N_k = \sum_{n=1}^N p(z_k = 1|\mathbf{x}_n) \quad (7.11)$$

5. Re-calculate the log likelihood (Equation 7.8) with updated parameters and repeat from step 3 until either the parameters (Equations 7.10, 7.12 and 7.14) or the log likelihood converges.
- 

GMMs are often viewed as an overall more powerful algorithm than K-means, since they perform statistical inference on the data. However, it is worth noting that most GMM implementations perform a K-means cluster analysis over the dataset in order to obtain estimations of the initial

---

<sup>10</sup>Either the object is in the cluster or it is not in the cluster.

parameters of the mixture components that are better than random chance. This improves the computational speed of the GMM clustering, since K-means is usually some orders of magnitude faster than Expectation-Maximization.

### 7.3.2 Clustering of LZap Data using GMMs

The GMM implementation of *scikit-learn* [P<sup>+</sup>11] was used for the cluster analysis<sup>11,12</sup> of the LZap pulse data. The relevant hyperparameters for this implementation of the GMM algorithm are the number of components,  $K$ , the covariance type and the convergence threshold. The covariance type could be one of the following options: full, tied, diagonal and spherical. The full covariance was chosen as it produces a general covariance matrix for each component. The tied covariance type would set the same general covariance matrix to all components, while the diagonal and spherical covariance types would produce a diagonal covariance matrix and a simple variance for each component, respectively. The convergence threshold dictates the early stopping of the EM algorithm by monitoring the gain in the log likelihood at each step, and was set to its default value of  $1 \times 10^{-3}$ . It is also worth noting that this GMM implementation uses the K-means algorithm for initializing the mixture weights and the Gaussian parameters.

As explained in Sections 6.2.3 and 7.2, the classes SPE and MPE are ultimately separated by the coincidence RQ alone. The former class is considered to be extracted from the overall dataset with 100% efficiency by selecting pulses with coincidence equal to 1, while the latter class is separated from the S1-like population by selecting pulses with coincidence lower than `s1MinCoincidence`. For this reason, pulses with coincidence equal to 1 are excluded from this dataset and are not involved in the cluster analysis, while any MPE pulse is considered as S1-like and no *a-priori* separation is performed at this stage. It is important to recall that the coincidence RQ calculation is not ideal, as discussed in Section 6.2.1. The current coincidence definition leads to some real SPE pulses having coincidence greater than 1 and some MPE pulses having coincidence greater than or equal to `s1MinCoincidence`, resulting in some contamination of the dataset by these pulse classes. This, however, does not affect the final result, since the vast majority of SPE and MPE pulses resemble S1 pulses in the first place and should be grouped naturally by the clustering algorithm. For these reasons, the number of classes considered in this clustering exercise is  $S = 4$ , corresponding to the class labels  $\mathcal{S} = \{s_1, s_2, s_3, s_4\} = \{\text{S1, S2, SE, Other}\}$ .

The number of components of the mixture model was estimated using an implementation of the kernel-based algorithm described in Reference [SNB05] and was set at  $K = 67$ . This number is much larger than the number of categorical pulse classes being considered,  $S = 4$ . By inspecting the dataset, and indeed by understanding the nature of each pulse type, it is clear that each pulse class contains several distinct pulse topologies, as mentioned previously in Chapter 6.1. Figure 6.4 showcases the general way in which the different pulse populations are grouped into pulse classes. Since different sub-populations within a pulse class present different densities and

<sup>11</sup><https://gitlab.com/PauloBras/gmmclustering.git>

<sup>12</sup><https://gitlab.com/PauloBras/mlforpc.git>

representativity, it is not guaranteed that constraining the number of Gaussian components in the mixture model would result in a clustering of the data with greater significance or usefulness. Each class is too diversified to be singled out by any clustering algorithm, except for SE pulses which display a very consistent behaviour and whose parameters are tightly bound by the physics of the detector.

If the number of components,  $K$ , is smaller than the number of classes then it is guaranteed that at least one component  $k$  will contain more than one pulse population. A component  $k$  that contains objects from two or more classes could be further partitioned with some high efficiency, meaning that the clustering is not optimal. If  $K$  is equal or larger than the number of classes then it is more likely that each component  $k$  will contain only one species, assuming that the clustering efficiency is large. Therefore, it is preferable to have the GMM algorithm partition the pulse data beyond the number of categorical classes. There are two main arguments for this choice: The first is that over-partitioning allows for a deeper understanding of the differences between seemingly degenerated populations, possibly allowing for the detection of outlier populations or pathological structures that were not identified by developers. The second is that after the clustering is performed, the different components of the mixture model can be collapsed into the same categorical classes provided by HADES by simply checking the elements of each GMM component. The complexity of the classification problem is greatly reduced by converting the Gaussian components into class labels and thus reducing the dimensionality of the label phase-space. This is a form of hierarchical processing that greatly accelerates learning on the subsequent classification algorithms [Ert17].

However, it is important to notice that by manually assigning the pulse labels to the GMM components, the analyser is introducing biases into the results of the GMM algorithm, as some of the choices made in the assignment may not be ideal. The author believes that the benefits of working with pulse classes instead of GMM components outweigh the potential problems that would result of the biases introduced in the GMM results. Also, it is fair to assume that these biases would be minimal and that the contamination of the class phase-space that would result from them can be easily handled by the classification algorithms studied in this work. As with many aspects of cluster analysis, and ML in general, a heuristic approach is somewhat inevitable [HTF09].

### **RQs used to perform the cluster analysis**

For the clustering analysis of the LZap dataset, only a subset of the pulse parameters were used, namely pulse area ( $pA$ ), prompt fraction at 100 ns ( $pF100$ ), top-bottom asymmetry ( $TBA$ ), pulse length at 90% area ( $pL90$ ) and pulse height<sup>13</sup> ( $pH$ ). All these parameters are explained in detail in Appendix B. Reducing the number of parameters used by the clustering algorithm helps to prevent the *curse of dimensionality* (see Section 7.2.1). By providing the algorithm with only the parameters with greater discrimination power, it is ensured that the density of data over the full phase-space is improved without significantly compromising the ability to partition

---

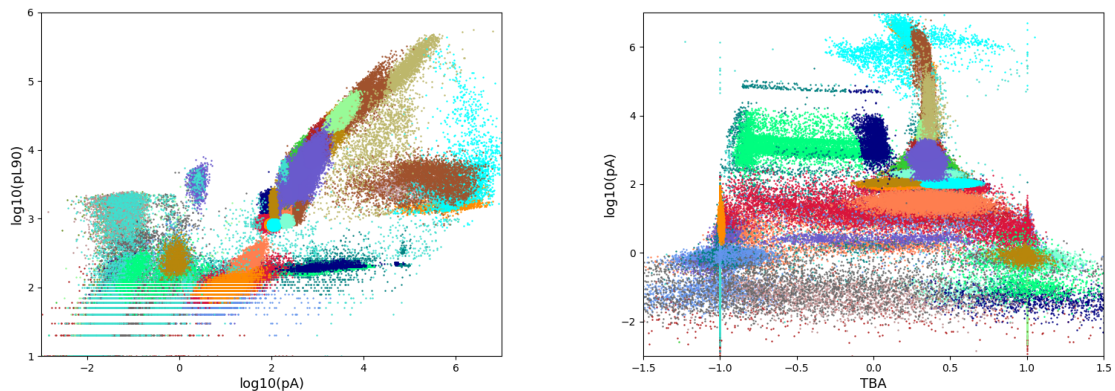
<sup>13</sup>The pulse height is the largest recorded peak voltage of a sample within the boundary limits of a waveform. It is, therefore, also called “peak amplitude”.

the data. Since the  $pA$ ,  $pL90$  and  $pH$  parameters span several orders of magnitude, the natural representation of these variables is in a logarithm scale. Therefore the following parameter space with  $D = 5$  pulse parameters was considered in the cluster analysis:

$$\mathbf{x} = \{\log_{10}(\mathbf{pA}), \mathbf{pF100}, \mathbf{TBA}, \log_{10}(\mathbf{pL90}), \log_{10}(\mathbf{pH})\}.$$

These pulse parameters were chosen based on their importance score (see Section 7.4) and form the simplest set of minimally-correlated geometrical RQs, with the only exception being  $pA$  and  $pH$  that are strongly correlated (see Figure 7.1). The  $pH$  parameter is expected to provide some discrimination ability to pulses with lower area.


Figure 7.2 displays the results of the GMM clustering of the LZap data, considering  $K = 67$  components. Despite the GMM algorithm providing probabilistic information, the presumed origin of each pulse with respect to the generation Gaussian component is converted into categorical information. This means that the pulse object  $\mathbf{x}_n$  can be associated to only one component  $k$  based on the largest value of  $p(z_k = 1|\mathbf{x}_n)$ .



**Figure 7.2:** Scatter plots of the GMM components for  $pA$  vs  $pL90$  pulse parameters (left) and  $TBA$  vs  $pA$  (right). The colors associated to each GMM component are cyclical and repeat for some components. The size and density of the 67 components vary significantly.

Several interesting pulse populations are highlighted in these plots. The first noticeable feature is the cyan population of S2-like pulses at high  $pA$  and trailing to higher  $pL90$  values, corresponding to GMM component  $k = 31$ , that clearly displays a pathological behaviour of  $TBA$  (range outliers), with fluctuations that are larger than expected for regular S2 pulses (see Figure 6.4). This suggests that the abnormal behaviour of the  $TBA$  parameter of these S2-like pulses is somehow correlated with their abnormal  $pL90$  behaviour. Another very relevant set of populations, that ultimately became the most useful result of this cluster analysis, are the SE split populations corresponding to components  $k = 15$  and  $k = 48$  (check Section 6.2.1 and Figure 6.7 for more details about SE splits). The ability to isolate these pulses from the neighbouring populations of SE and S1 pulses has proven to be quite challenging, as discussed in Section 6.2.3. Using this cluster analysis it is possible to separate these pulses from the remaining populations with high efficiency, leading to their parametrization and ultimately paving the way to find the

best discriminants available for classification purposes. GMM components  $k = 15$ ,  $k = 31$  and  $k = 48$  are displayed in Figure 7.3. Every component is checked for pathological features in these and other marginal distributions, and any strange population is handscanned to determine their constituents.




**Figure 7.3:** Scatter plots of the GMM components  $k = 31$  (cyan),  $k = 15$  (brown) and  $k = 48$  (orange), for  $pA$  vs  $pL90$  pulse parameters (left) and  $TBA$  vs  $pA$  (right). The component  $k = 31$  contains S2 pulses with abnormal behaviour and was selected for handscanning. Components  $k = \{15, 48\}$  were also selected for handscanning and were identified as SE splits.

### Generating new pulse class labels from the GMM results

These results provide the groundwork for generating training labels for subsequent classification algorithms. Each component will be assigned a categorical class based on the type of pulses that it contains. For determining the contents of each component, a series of tests were performed to the partitioned data: First, the distributions over the parameter phase-space for each component are plotted, like in Figures 7.2 and 7.3, which could hint to the analyser the nature of the pulses within; Second, the results from HADES are computed for each component, and since HADES has a high classification efficiency, the majority of the pulses within each component will be classified correctly; The third test, only performed if the first two tests provide insufficient information about the nature of the pulses, is to handscan a sample of pulses from each component.

Figure 7.4 shows the relative fraction of HADES labels for each GMM component, as well as the number of occurrences of each class. It is clear from the bottom plot that the components  $k = \{1; 5; 14; 17; 39\}$  have collapsed into singularities, having only one data element assigned to each. Component  $k = 4$  contains only 3 pulses, which may indicate that the component has also collapsed. Due to the lack of statistics, all these singularities are labelled as “Other” and are not expected to contribute significantly to the results of any algorithm that uses this data as training.

Components  $k = \{8; 11; 15; 16; 21; 22; 24; 30; 31; 34; 37; 38; 42; 44; 48; 50; 58; 59; 61; 62; 63; 65; 66\}$  contain either a significant mixture of HADES labels or have strange features in their parameter



**Figure 7.4:** Fraction of HADES labels per GMM component (top) and frequency of each HADES label per component (bottom). Error bars are Poisson in nature. Some mixture of HADES labels can be observed in most components, which means that either HADES is separating populations that the GMM clustering algorithm perceives as similar, or the GMM clustering is merging distinct populations that HADES can differentiate. Due to the simplicity of the HADES algorithm, the former hypothesis seems more likely. Arrows above the top plot mark the GMM components selected for handscanning.

distributions. These components were sampled and handscanned in order to determine the most likely class of pulses contained within. The number of pulses sampled varied from component to component, since some components are more well-behaved than others. Components that are constrained to a limited portion of the full scope of the parameter phase-space are expected to contain pulses that are very similar to one another, with little internal variability. For each handscanned component, only when a certain number of pulses with the same waveform topology are identified, without the presence of other pulse types, can the component be labelled. If two or more topologies that are not considered to be compatible with the same pulse class are present, *e.g.*, S1 pulses and S2 tails, the component is either labelled as “Other” or as the class with majority, depending on the impact of the impurities being introduced in the respective label.

After determining which pulse classes should be associated with each component, a mapping between the GMM component index and the pulse class labels can be constructed. Table 7.1 shows the mapping of components to pulse classes for this cluster analysis. The new labels assigned to each data object can now be used by other algorithms better suited for supervised learning, such as decision forests and forest ensembles. These methods will be studied in Section 7.4.

Table 7.2 shows the confusion matrix of the GMM results compared to the original HADES classifications for the entire LZap dataset. The large number of pulses seen by HADES as Other that are now labelled as S1(MPE) by the GMM algorithm are mostly afterpulses and two coincident dark counts that are purposely associated to the S1 population by the mapping in Table 7.1. These pulse topologies display great variability and permeate large portions of the parameter phase-space, leading HADES to misclassify a large fraction of these pulses. A similar



**Table 7.1:** Mapping from GMM component  $k$  to pulse class label. The labels marked with an asterisk were selected via handscanning of 20 or more random pulse samples from the respective component. The primary pulse type observed on each handscanned component is denoted in parenthesis: Afterpulsing (AP), two coincident dark counts (2fDC), three coincident dark counts (3fDC), SE splits, S2 splits, baselines and random mixtures of pulses (see Section 6.2.1 for details). The waveforms of the S2 pulses found in component  $k = 31$  do not display strange features, despite the odd behaviour of  $TBA$  and  $pL90$  (see text for details).

$k$	label	class $\in \mathcal{S}$ (type)	$k$	label	class $\in \mathcal{S}$ (type)	$k$	label	class $\in \mathcal{S}$ (type)
0	3	SE	23	2	S2	46	3	SE
1	4	Oth	24	1*	S1 (AP)	47	4	Oth
2	4	Oth	25	1	S1	48	4*	Oth (SE split)
3	2	S2	26	3	SE	49	2	S2
4	4	Oth	27	2	S2	50	1*	S1 (2fDC)
5	4	Oth	28	2	S2	51	3	SE
6	2	S2	29	2	S2	52	2	S2
7	3	SE	30	4*	Oth (S2 split)	53	4	Oth
8	1*	S1 (AP+2fDC)	31	2*	S2 (S2)	54	4	Oth
9	4	Oth	32	2	S2	55	2	S2
10	1	S1	33	2	S2	56	2	S2
11	1*	S1 (2fDC)	34	1*	S1 (2fDC)	57	2	S2
12	2	S2	35	4	Oth	58	1*	S1 (AP)
13	4	Oth	36	4	Oth	59	2*	S2 (S2 tail split)
14	4	Oth	37	4*	Oth (baseline)	60	1	S1
15	4*	Oth (SE split)	38	1*	S1 (AP)	61	4*	Oth (baseline)
16	1*	S1 (2fDC)	39	4	Oth	62	1*	S1 (2fDC)
17	4	Oth	40	4	Oth	63	1*	S1 (AP)
18	3	SE	41	2	S2	64	2	S2
19	2	S2	42	4*	Oth (S2/SE split)	65	4*	Oth (baseline)
20	4	Oth	43	3	SE	66	1*	S1 (2fDC)
21	4*	Oth (mixture)	44	4*	Oth (3fDC)			
22	1*	S1 (2fDC)	45	1	S1			

\* selected via handscan

problem has led to the labelling of several pulses as Other that were formerly seen as S1s or MPEs by HADES. As explained in Section 6.2.3, the S2-like pulse selection takes precedence over S1 selection, leading most spurious pulses to be in the S1-like phase-space after the first S1-S2 discrimination. The subsequent selection of true S1-like pulses is also affected by the relative small size of some S1-like topologies, that display great variability akin to spurious pulses and with which they share some of the phase-space. This leads to the compromise of finding selection criteria that select true S1-like pulses with high efficiency while excluding spurious pulses from the S1 classification. The GMM cluster is assumed to discern between these two pulse types with a larger efficiency than HADES, resulting in the differences displayed in Table 7.2.

**Table 7.2:** Confusion matrix of the GMM results compared with the results from HADES over the full dataset.

HADES class	GMM class				Total	
	S1(MPE)	S2	SE	Other		
S1(MPE)	39217	0	0	4545	43762	4.37%
S2	0	226577	14	194	226785	22.68%
SE	0	31694	642483	114	674291	67.43%
Other	20002	0	0	35160	55162	5.52%
<b>Total</b>	59219	258271	642497	40013	1000000	

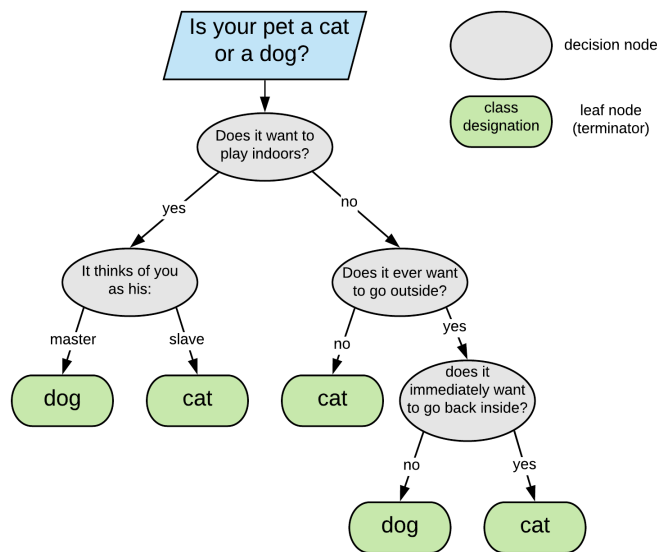
## 7.4 Random Forest Classifier

### 7.4.1 Decision Tree Models

Decision trees are flowchart-like structures that can be used for decision making, categorization and regression analysis [Tin95, Bre01, RYZ<sup>+</sup>05, HTF09]. They present themselves as a simple set of rules that control the flow of data and seek to arrange data objects in discrete categories. These algorithms are therefore very appealing to solve classification problems, with the advantages of being extremely fast and intuitively understood.

In their most general form, they are composed by a series of connected branches in sequence and by terminators of those branches that represent the different outputs of the tree, as depicted in Figure 7.5. In a typical decision tree the data flows from top to bottom, from the input node towards the tree terminators, passing through several branching nodes – where two or more data paths branch out from – where a decision is made to determine which branch the data will flow through next. The decisions in these branch nodes are tuned to separate the data into different subsets with similar features in the most efficient way. The endpoint of any given path along the tree will be reached by any data object whose properties conform with the decisions along the followed path. A subset of data objects that end in the same termination of the tree must have a similar set of properties.

A decision tree can be constructed using heuristic methods, *e.g.*, the HADES algorithm mentioned in Section 6.2.3, or it can be self-assembled and trained with data using mathematical models and tools [HTF09]. The latter variant of decision tree conforms with the notion of a machine learning algorithm. An efficient method for building a decision tree classifier, regardless of the dimensionality of the input data or the total number of object classes, is recursive partitioning. This method uses intrinsic metrics of the phase-space of the data to find the best split conditions of the different data classes. The most basic implementation of this method aims at finding the best hyperplane that separates any two data clusters that is perpendicular to the line that connects the center of mass of the clusters. This implementation, also called binary recursive partitioning or central-axis projection [Tin95], usually starts by splitting the two data classes whose centers of mass have the largest Euclidean distance. The choice of the separating hyperplane is often made by minimizing an impurity function, often a variant of



**Figure 7.5:** Schematic representation of a simple decision tree used for classification. The input data (blue) passes several decision nodes (gray) that evaluate one or more parameters of the data object and redirect it to the branch that represents the result of the test until a terminator node, or leaf node (green), is reached. Each leaf node has an associated class label that identifies the data objects that reach it, essentially grouping the data into classes of similar objects using their parameters as discriminants.

the sum of squared errors, where the error can be defined as the fraction of the data objects of one of the two classes that are on the opposite side of the hyperplane with respect to the majority (misclassification error). Other impurity criteria are often used for providing a better optimization of the decision step, like Gini impurity or deviance (cross-entropy) [HTF09]. After each partitioning of the data, the recursive algorithm repeats the same steps on the new subsets of data. The process can be repeated until all subsets contain only one data class or a single member. Since the goal is to optimize the separation of at least two clusters at each decision node, recursive partitioning is a greedy algorithm [HTF09].

If a decision tree uses only one data feature per decision node, like with the case of the tree depicted in Figure 7.5, it is called a binary decision tree. On such trees the data flow at each decision node is controlled by a single test point or threshold  $j$  on a single parameter  $d$ . The separation of data objects in the feature phase-space will simply follow hyperplanes perpendicular to the axis of a feature. This can be easily understood as a limitation on the choice of the best separation hyperplane for a multivariate complex dataset. Oblique decision trees are a more complex variant of binary decision trees where at each decision node a combination of features is used to guide the flow of data [Tin95]. Adjusting the hyperplanes used to separate the data can reduce the size of the tree and improve the purity of the data objects that go into each termination. These oblique trees are more general and naturally more compact than their binary counterparts for the same problem, but they can be harder to interpret [HTF09].

In a binary decision tree the partitioning of the data along a single parameter  $d$  with a threshold

$j$  results in a split of the parameter phase-space into two regions:  $R_1(d, j) = \{\mathbf{X} | \mathbf{X}_d \leq j\}$  and  $R_2(d, j) = \{\mathbf{X} | \mathbf{X}_d > j\}$ . The parameter  $d$  and separation threshold  $j$  are chosen in order to maximize the splitting. The tree is grown by applying these successive partitions until no more partitioning is possible or a stopping criterion is reached. The final tree will have  $M$  nodes corresponding to the same number of regions on the parameter phase-space,  $\{R_1, R_2, \dots, R_m, \dots, R_M\}$ , with  $m = 1, 2, \dots, M$  the index of nodes. Each region contains  $N_m = \#\{\mathbf{x}_n \in R_m\}$  elements within [HTF09]. The proportion of objects of class  $s$  observed in node  $m$  is given by

$$p_{ms} = \frac{1}{N_m} \sum_{\mathbf{x}_n \in R_m} I(y_n = s), \quad (7.15)$$

where  $I$  is the identity operator, that returns 1 if the statement is true and 0 if it is false, and  $y_n$  is the predicted result. This proportion can be calculated at any tree growing stage, not just using the final model. Using Equation 7.15, the majority class in node  $m$  can be obtained by  $\hat{s}(m) = \arg \max_s p_{ms}$ . But if each object within region  $R_m$  is classified to class  $s$  with associated probability  $p_{ms}$ , then the training error associated to tree node  $m$  is given by the Gini index, defined as

$$\sum_{s \neq s'} p_{ms} p_{ms'} = \sum_{s=1}^S p_{ms} (1 - p_{ms}). \quad (7.16)$$

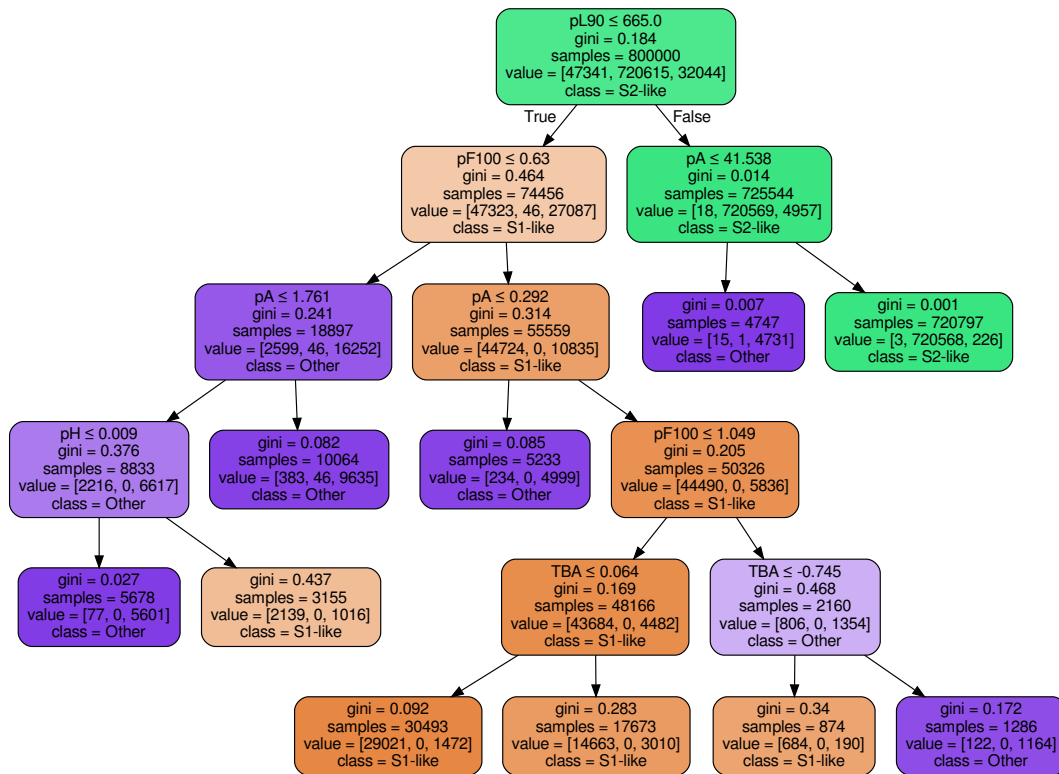
The Gini index is an intuitive solution to quantify the impurity at each node, with the additional advantage of being differentiable [HTF09].

A decision tree algorithm trained on data can be grown to any size and contain an arbitrarily large number of subdivisions of the data. Such trees are undesirable as they often overfit the data they were constructed with. Having a method that can build and train a tree to arbitrarily large complexity while maintaining high accuracy scores both on training and on cross-validation data is challenging if not impossible for most classification problems. Limiting the growth of the tree will ensure that the algorithm retains flexibility and can generalize to novelty data, but inhibiting the complexity of a tree will result in sub-optimal accuracy upon training [Tin95]. An option to limit tree size and reduce overfitting on training data is to “prune” the grown tree [HTF09]. Pruning is the removal of internal decision nodes (not terminal nodes) of a grown decision tree, often targeting nodes that have less power of separation of different classes. If a tree  $T_b$  is allowed to grow until each terminator either contains a single class or only one element<sup>14</sup>, pruning  $T_b$  returns a sub-tree  $T \subset T_b$  that will retain a high classification efficiency while reducing the overfit to the training data. The pruning can be guided using any impurity index, like the Gini index, to determine the nodes with less predictive power.

Figure 7.6 shows an example of a binary decision tree built for a classification exercise involving the S1-like, S2-like and Other pulse classes obtained with the GMM clustering analysis described in Section 7.3.1. The tree was not pruned and was build with a maximum depth of 5 consecutive decision nodes for simplicity.

---

<sup>14</sup>A fully grown tree will have a classification accuracy of 100% on the training data, but will most likely generalize poorly to novelty data.



**Figure 7.6:** Decision tree trained to classify LZap pulses into S1-like, S2-like and Other pulse classes. The S1-like class contains the S1 and MPE pulses while the S2-like class contains both the S2 pulses and the SE pulses clustered in the GMM analysis described in Section 7.3.1.

Despite the potential of single tree solutions for classification and regression problems, the major strength of decision trees is not in their individual complexity and accuracy but in their ability to work in large ensembles. An ensemble of  $B$  trees,  $\mathcal{T} = \{T_b\}_1^B$ , is conveniently called a “forest”, and the combined results of the individual weak classifiers can be monotonically improved so that the ensemble becomes a strong classifier. The tree displayed in Figure 7.6 is an example of a weak classifier that can be used as the basis for building these ensembles. Forest classifiers can be trained using weak learnability [Sch90], where each tree only needs to perform slightly better than random guessing for the overall accuracy of the forest to be large with low error and no loss of generalizability. The forest becomes a strong learner, meaning it can retain low error with high confidence for all classifications.

The two most important ideas for building and training a decision forest classifier are bootstrap aggregation and boosting, each leading to two of the most powerful tree ensemble models in ML, those being Random Forests (RFs) and Boosted Decision Trees (BDTs), respectively.

### 7.4.1.1 Bootstrap aggregation and random forests

Diversity of the weak models within the forest ensemble is what allows for both high accuracy and generalizability. If all weak classifiers in the forest ensemble have the same structure and act over the data in the exact same way, the ensemble will not be able to go beyond the limitations of the individual classifier. An elegant idea for building a decision forest is to create several subsets of training data chosen randomly and train each tree with a different subset or closed groupings of subsets. This technique is called bootstrap aggregation, or bagging, and ensures that the individual classifiers have a certain degree of diversity. When bagging is used, different classifiers will model the data in a different way, effectively reducing the variance of the forest.

An extension of the concept of bagging is the idea to not only sample the data objects that are fed to each classifier but to also sample the data features that are available to any given classifier (feature bagging). Several of these weak classifiers, built over a random subspace of the data features, form a Random Decision Forest [Tin95, Bre01, HTF09]. In a Random Forest model, the feature space is sampled at random in order to choose a subset of the total data parameters available for any given tree in the ensemble. Each tree is fully grown over a limited representation of the data, meaning that the individual classifiers cannot fully encode the data and ensuring that the ensemble is naturally resistant to overfitting. The results of a Random Forest are the combination of the results of each tree: for a regression model the predictions of the individual trees are averaged, while for a classification model the majority of the predictions is taken as the prediction of the model – essentially a vote. Algorithm 3 displays the generic procedure to build a random forest model and to predict the class of a data object.

A rule of thumb for the number of parameters selected when bagging is  $k = \sqrt{D}$ , with  $D$  being the dimensionality of the parameter space, and for the minimum number of samples in a node is `minNodeSplit` = 1, or conversely `minSamplesSplit` = 2, but these hyperparameters are model-dependant and should be optimized [Bre01, HTF09]. No constraint needs to be set on the size of each tree since these will be built in a marginal representation of the phase-space and thus will not easily overfit the data. However, this depth hyperparameter, `maxDepth`, should be tested nonetheless to ensure a good generalizability of the model.

Since each tree of the random forest is built using a bootstrapped sample of the dataset, the remaining data not used to build a given tree  $T_b$  can be used directly to estimate its performance. These samples are called *out-of-bag* (OOB) samples. For each observation  $\{\mathbf{x}_n, y_n\}$  the response of the random forest can be predicted by using the collective response of the individual trees where  $\{\mathbf{x}_n, y_n\}$  was OOB [JWHT13]. The overall classification error for the observations  $\{\mathbf{X}, \mathbf{Y}\}$  obtained from this OOB prediction is a valid estimate of the test error of the model.

Random Forests handle high-dimensional data very well and can maintain a high accuracy even when exposed to data with missing values. They also deal well with under-represented classes in the training data and are not overzealous with density regions on the phase-space of the features. However, in the presence of high dimensional data with a large fraction of noise/nuisance parameters, random forests may underperform due to the randomness in choosing the feature space for building the trees, especially for small forests [HTF09]. A tree grown in such

---

**Algorithm 3:** Random Forest for classification [HTF09]
 

---

1. For  $b = 1$  to  $B$  trees in the forest:
  - (a) Draw a bootstrap sample  $\mathbf{X}^*$  of size  $N^*$  from the training dataset  $\mathbf{X}$ , satisfying  $\mathbf{X}^* \subset \mathbf{X}$
  - (b) Grow a decision tree  $T_b$  with the bootstrapped data  $\mathbf{X}^*$ , by recursively repeating the following steps for each terminal node of the tree, until the minimum node size  $n_{min} \equiv \text{minNodeSize}$  is reached, or conversely if the number of samples in a terminal node is less than the minimum number of samples allowed to split a node,  $n < n_{split} \equiv \text{minSamplesSplit}$ :
    - i. Randomly select  $k$  parameters from the available  $D$  parameters
    - ii. Pick the best parameter among the  $k$  to split the data, using an impurity criterion to evaluate the split.
    - iii. Split the node into two daughter nodes.
2. Return the ensemble of trees  $\mathcal{T} = \{T_b\}_1^B$

To make a prediction for a data object  $\mathbf{x}$ :

1. Let  $\hat{f}_b(\mathbf{x})$  be the class prediction function of the  $b^{\text{th}}$  tree in the random forest ensemble, and  $\hat{f}_{model}(\mathbf{x})$  be the class predictor function of the full random forest model. Then:

$$\hat{f}_{model}(\mathbf{x}) = \text{majority vote}\{\hat{f}_b(\mathbf{x})\}_1^B$$


---

conditions has a small chance of selecting at random a relevant parameter in a decision node, worse even if the sampled features have a small representation of relevant variables. When working with data with large dimensionality it is often useful to rank the importance of the variables in the data beforehand, and choosing a subset of the features that ranked higher to avoid hindering of the efficiency of an individual tree by selecting at random less important variables. Feature importance ranking, or variable importance, is the measurement of the improvement in the splitting criterion attributed to each variable, accumulated over all the trees in a forest ensemble [HTF09]. Random Forest, as well as similar ensemble methods, can generate reliable estimations of the overall discrimination power of individual variables in data. This reliability is due to two main factors: bagging in the feature (variable) space ensures that variables with less discriminant power are given a chance if they are randomly sampled when stronger variables are not, and the averaging of the importance factors reduces the variance of the importance scoring [HTF09].

The variable importance can be calculated in different ways, the most basic estimation being to simply count the number of times a variable is selected by the individual trees and averaging the occurrence over the full forest. However, this simple method does not provide an assessment of the actual improvement that a variable brings to a single classifier, since two equally frequent variables might contribute in different ways to the end result. The Gini impurity index mentioned

above provides a more complete estimation of the importance of a variable by weighing the frequency of selection of the variable as a discriminant with the discrimination gain at each respective decision node.

Perhaps the best metric for determining the importance of a variable is the permutation accuracy importance [Bre01, SBZH07]. It assesses the relevance of a given variable for the overall accuracy of the classifier tree in an elegant way: after the classifier is built, randomly permuting some discriminant variables will disassociate the affected variable with the output of the tree, which is guaranteed to reduce the efficiency. Comparing the resulting accuracy with the one obtained for the intact tree yields a misclassification rate that can be interpreted as the overall effect of the variable on the accuracy, and thus its importance. Instead of doing permutations of variables and scoring the misclassification rate over all permutations, one can noise each variable independently and obtain a similar effect. Permuting variables is often preferred because it not only determines the global effect that each variable has on the performance of the tree but also accounts for any highly correlated variables in the data that provide limited improvements due to having viable alternatives.

Despite the appealing usefulness of the variable importance metric provided by Random Forest models, the resulting scores are often affected by strong biases that need to be carefully mitigated [SBZH07]. It is important to notice that these biases are contained in the feature importance ranking and are not in any direct way compromising the classifier model they are calculated from. The origin of these biases are essentially two: asymmetries in class representativeness in data and strong variable correlation. On the first issue, classes that are well represented will continue to be more prominent in subsets generated by bagging, and since they represent a larger portion of the data the decision criteria aimed at discriminating this classes might be viewed as stronger overall. The second issue involves having several variables that are strongly correlated in the data. If the random bagging of the variables when a tree is being built results said tree being exposed to two or more highly correlated variables, after one of those variables is chosen at random, the remaining variables will contribute marginally to the accuracy gain, as their discrimination power is already accounted for by the first correlated variable to be chosen. This last issue can be mitigated by implementing the aforementioned permutation accuracy importance test. For a possible mitigation of the first issue, one could compute the variable importance for each class by reducing the classification problem to a One vs All scenario. This technique does not give a general importance ranking, but the class-specific variable importance rankings provide insightful information nonetheless, especially in problems where the different classes are not equally important.

#### 7.4.1.2 Boosting and boosted decision trees

Boosting is a very useful idea for building ensembles of classifiers. The goal of boosting is to generate an ensemble of weak classifiers that are built and trained sequentially, so that at each new step  $b$  the new classifiers are trained with data weighted by the classifiers from the previous steps [HTF09]. The initial weights,  $\omega_{n,b=1}$ , are equal for all the  $n = 1, 2, \dots, N$  data objects



such that  $\omega_{n,b=1} = 1/N$ . The weights applied to the data at each boosting step  $b$  are related to the errors of the previous classifiers,  $T_{b-1}$ , with the weights of data that has been misclassified being increased while the weights of the correctly classified data being decreased. This forces the younger classifiers to focus on the mistakes of previous classifiers, assuring that at each boosting step  $b$  the classification accuracy improves. In the end of all boosting steps  $B$ , the boosted tree ensemble – more often called simply Boosted Decision Tree (BDT) – outputs not only the results of the final boosted tree but the weighed results of all trees in the ensemble, with the weight  $\alpha_b$  being calculated for each tree based on its performance at each boosting stage  $b$ .

Despite sharing some basic properties with Random Forests, Boosted Decision Trees are built and operate with a very different paradigm. Like Random Forests, the initial requirement set on the weak classifiers of the ensemble is that these need only to perform better than random chance at solving the problem, but while Random Forests use the strength of many weak classifiers, Boosted Decision Trees aim at improving the weak classifiers by sequentially applying the weak results onto the data and modifying it for the next classifier.

The most popular boosting algorithm is called *AdaBoost*<sup>15</sup> [Sch99]. For a given classification problem, *AdaBoost* uses the results of a given classifier,  $T(\mathbf{X})$ , that acts over a vector of data variables  $\mathbf{X} = \mathbf{x}_n$ ,  $n = 1, 2, \dots, N$ , producing a vector of results  $Y = y_n \in \{-1, 1\}$  (two-class problem) to modify the weights  $\omega_n$  at each boosting step  $b$ . At the first step,  $b = 1$ , the weights are uniform for all  $\mathbf{X} = \mathbf{x}_n$  data objects:  $\omega_n(b = 1) = 1/N$ . Over the following steps, from  $b = 1$  to  $B$ , a classifier  $T_b(\mathbf{x}_n)$  is fitted to the training data using the current weights  $\omega_n$  and the total error rate of that classifier is calculated using:

$$err_b = \frac{\sum_{n=1}^N \omega_n I(y_n \neq T_b(\mathbf{x}_n))}{\sum_{n=1}^N \omega_n}. \quad (7.17)$$

The weight  $\alpha_b$  for the classifier  $T_b$  is obtained using the error rate from Equation 7.17,  $\alpha_b = \log((1 - err_b)/err_b)$ . Finally the weights  $\omega_n$  are updated to  $\omega_n \leftarrow \omega_n \cdot \exp[\alpha_b \cdot I(y_n \neq T_b(\mathbf{x}_n))]$ . The process is then repeated for all  $B$  boosting stages [HTF09].

The output of the ensemble obtained using *AdaBoost* is given by

$$T(\mathbf{x}_n) = \text{sign} \left( \sum_{b=1}^M \alpha_b T_b(\mathbf{x}_n) \right). \quad (7.18)$$

The *AdaBoost* algorithm, taken from Reference [HTF09], is summarised in Algorithm 4.

*Adaboost* is valid only for a two-class classification problem. For a problem involving more than two classes *AdaBoost* can quickly diverge if  $err_b > 0.5$ , resulting in  $\alpha_b < 0$  and the weights  $\omega_n$  are updated in the opposite direction. The generalization of *AdaBoost* to handle multiple classes usually involves transforming the multi-class problem into several two-class problems. This can be cumbersome for larger class multiplicity and does not yield results as

---

<sup>15</sup>From “Adaptive Boosting”.

---

**Algorithm 4:** AdaBoost Algorithm

---

1. Initialize the weights  $\omega_n = 1/N$ ,  $n = 1, 2, \dots, N$ .
2. For  $b = 1$  to  $B$ :
  - (a) Fit a classifier  $T_b(\mathbf{x}_n)$  to the training data using weights  $\omega_n$ .
  - (b) Calculate:

$$err_b = \frac{\sum_{i=1}^N \omega_n I(y_n \neq T_b(\mathbf{x}_n))}{\sum_{n=1}^N \omega_n},$$

- (c) Calculate  $\alpha_b = \log((1 - err_b)/err_b)$ .
  - (d) Update individual weights as  $\omega_n \leftarrow \omega_n \cdot \exp[\alpha_b \cdot I(y_n \neq T_b(\mathbf{x}_n))]$ ,  $n = 1, 2, \dots, N$ .
  - (e) Renormalize  $\omega_n$ .

3. Return:

$$T(\mathbf{x}) = \text{sign} \left( \sum_{b=1}^B \alpha_b T_b(\mathbf{x}) \right).$$


---

good as alternative methods. A proposed alternative method that aims at modifying *AdaBoost* to handle multiple classes is the “Stagewise Additive Modeling using a Multi-class Exponential loss function” (SAMME) [ZRZH06]. By identifying that the boosting weight  $\alpha_b$  can be modified to become  $\alpha_b = \log((1 - err_b)/err_b) + \log(S - 1)$ , where  $S$  is the total number of classes, SAMME maintains the original premise of only requiring that a given weak learner performs better than random guessing. Here the results from a classifier  $T_b(\mathbf{x}_n) = y_n$  are part of the group of classes  $\{1, 2, \dots, s, \dots, S\}$ , where  $s$  is one of the  $S$  class labels, Notice that when  $S = 2$  SAMME reverts to the original *AdaBoost*. The final result of SAMME can be expressed as

$$T(\mathbf{x}) = \arg \max_s \left( \sum_{b=1}^B \alpha_b I(T_b(\mathbf{x}) = s) \right), \quad (7.19)$$

where  $\arg \max$  stands for the argument of the maximum. The SAMME algorithm steps, taken from Reference [ZRZH06], are summarised in Algorithm 5.

The large potential of BDTs and RFs for data mining is proven by their extensive use within the particle physics community. The usage of BDTs is a well established practice in the field, with several particle physics experiments and even some DM direct detection experiments adopting these methods instead of the classical analysis [A<sup>+</sup>16g, RYZ<sup>+</sup>05, WQY<sup>+</sup>15, K<sup>+</sup>16, A<sup>+</sup>15b]. This motivates their implementation and testing within the data analysis framework of LZ. Viability tests will be performed in the future to assess if these machine-learning techniques can also be

---

**Algorithm 5:** SAMME Algorithm

---

1. Initialize the weights  $\omega_n = 1/N$ ,  $n = 1, 2, \dots, N$ .
2. For  $b = 1$  to  $B$ :
  - (a) Fit a classifier  $T_b(\mathbf{x}_n)$  to the training data using weights  $\omega_n$ .
  - (b) Calculate:

$$err_b = \frac{\sum_{n=1}^N \omega_n I(y_n \neq T_b(\mathbf{x}_n))}{\sum_{n=1}^N \omega_n},$$

- (c) Calculate  $\alpha_b = \log((1 - err_b)/err_b) + \log(S - 1)$ .
  - (d) Update individual weights as  $\omega_n \leftarrow \omega_n \cdot \exp[\alpha_b \cdot I(y_n \neq T_b(\mathbf{x}_n))]$ ,  $n = 1, 2, \dots, N$ .
  - (e) Renormalize  $\omega_n$ .

3. Return:

$$T(\mathbf{x}) = \arg \max_s \left( \sum_{b=1}^B \alpha_b I(T_b(\mathbf{x}) = s) \right).$$


---

used in LZ for pulse identification<sup>16</sup> and event classification. These ideas will be explored in parallel with heuristics algorithms such as HADES, described in Section 6.2.3.

Despite the more widely usage of BDTs in physics analysis, it has been established that the performance of random forests can be very similar to boosting in several problems, with the added advantages of being simpler to train and tune [HTF09]. BDTs have the invaluable ability to model data at different levels of complexity, with the transition from data generalization to fine-tuning being gradually encoded into the sequence of boosted learners. This is very useful for novelty searches or in analyses where the expected target signal is a finely tuned feature within the larger model. For the classification analysis of LZap, however, the primary goals are not related to evaluating the compatibility of the model with specific features in the data but on providing insightful information that can be used by the HADES algorithm or any future classifier model. Exploring both RFs and BDTs can provide complementary information, but the main focus of this work will be on random forests.

The next Sections describe the implementation of the aforementioned ML methods applied to the task of classification of pulse structures in simulated LZ events.

### 7.4.2 The RFClassifier Pulse Classification Tool

The simplicity and robustness of random forests motivated the development of the *RFClassifier*<sup>17,18</sup> pulse classification tool presented in this Section. This tool has been developed using

---

<sup>16</sup>The identification and correlation of pulses within the context of the event.

<sup>17</sup><https://gitlab.com/PauloBras/rfclassifier.git>

<sup>18</sup><https://gitlab.com/PauloBras/mlforpc.git>

the *scikit-learn* implementation of the *RandomForestClassifier* model [P<sup>+</sup>11]. The main goals of this classifier are:

1. To provide a deeper understanding of bad pulse populations overlapping with the main expected populations of S1-like and S2-like pulses, and, if possible, determine the most efficient way of separating these two types of pulses from the data,
2. To determine which RQs, parametric thresholds and sequences of selection criteria yield the most efficient partitioning of LZap pulse data.

The results obtained in the clustering analysis of the LZap data with the GMM model contribute to the fulfilment of the first goal. The RFClassifier is expected to find the best strategy to separate the population of SE split pulses from the S1 population since the former was already successfully isolated during the GMM clustering analysis. The second goal can be achieved using the feature importance ranking capabilities of the random forest model that will be explored below.

### Input conditioning

The LZap dataset described in Section 7.2 was used to train and benchmark the RFClassifier model. The pulse parameters selected for benchmarking the model were pulse area ( $pA$ ), prompt fractions at 50 ns ( $pF50$ ), 100 ns ( $pF100$ ), 200 ns ( $pF200$ ) and 1  $\mu$ s ( $pF1k$ ), top-bottom asymmetry ( $TBA$ ), pulse length at 90% area ( $pL90$ ), pulse height ( $pH$ ), pulse height time scaled to pulse length ( $pHTL$ ), and pulse RMS width ( $pRMSW$ ). All RQs are described in detail in Appendix B. The full parameter space with  $D = 10$  pulse parameters used to build the RFClassifier model is then

$$\mathbf{x} = \{\mathbf{pA}, \mathbf{pF50}, \mathbf{pF100}, \mathbf{pF200}, \mathbf{pF1k}, \mathbf{TBA}, \mathbf{pL90}, \mathbf{pH}, \mathbf{pHTL}, \mathbf{pRMSW}\}.$$

Since the decision tree methods partition the data at simple thresholds on a given parameter, no parameters need to be converted to a logarithmic representation.

Despite being highly correlated, the four prompt fraction parameters,  $pF50$ ,  $pF100$ ,  $pF200$  and  $pF1k$ , were included in the input data to determine which one has the strongest discriminative power. These variables are calculated using time windows at different scales, and each will have a stronger response to pulse structures with compatible time scales. Using the feature importance ranking provided by the RFClassifier, it is possible to determine which of these scale-sensitive parameters is the best at separating different pulse classes.

The GMM results obtained in Section 7.3.1 and displayed in Table 7.1 were used as class labels during training. No additional selection of data was performed and all classes are considered to have the same importance. It is worth noting that some classes are more common than others in this dataset, namely SE pulses. Bagging may induce some class bias if some classes are more represented than others in a multi-class dataset. This is especially damaging if bootstrapping is done without replacement, which is the case. However, the asymmetry on the abundances of


the different species is not too severe and all classes are assumed to be sufficiently represented (see Table 6.1 for some idea on the relative abundance of the different pulse classes). Therefore, this effect is expected to be minimal for this analysis.

### Tuning the relevant hyperparameters

The *scikit-learn* implementation of the *RandomForestClassifier* model [P<sup>+</sup>11] supports the control of several hyperparameters, including the number of estimators in the model, the node split criterion, the max depth of the individual trees, the minimum number of samples for splitting a node, the minimum number of samples required to be at a leaf node, the maximum number of parameters to consider when looking for the best split and the application of bagging.

The hyperparameters that were tuned when building the RFClassifier were the number of tree estimators (`nTrees`), maximum depth of each tree (`maxDepth`) and the minimum number of samples for splitting a node (`minSamplesSplit`). The Gini index was selected as the split criterion, the maximum number of randomized parameters used to find the best split was set to the integer square root of the total number of parameters, and sample bootstrapping was enabled. The remaining hyperparameters of the model not mentioned here were set as default<sup>19</sup>.

Figure 7.7 displays the overall test accuracy as a function of the number of estimators, `nTrees`. The decrease in accuracy at higher forest size may be attributed to statistical fluctuations, and above `nTrees = 53` estimators the accuracy does not change significantly. The optimal forest size was determined to be `nTrees = 101`.




**Figure 7.7:** Impact of forest size (number of estimators) in the classification accuracy over the test set. Each accuracy value is the mean of the final accuracy of three independent training sessions with the same model architecture. Each tree is allowed to grow freely (`maxDepth=None`), and the minimum number of samples required to split a branch is set to 2 (`minSamplesSplit=2`).

The left plot in Figure 7.8 shows the performance of the RFClassifier as a function of the `maxDepth` hyperparameter, that controls the maximum number of branches each individual

<sup>19</sup>As in version 0.22.2 of *scikit-learn*.

tree is allowed to have, *i.e.*, the depth of each tree. Setting this variable as *None* will allow each tree to grow unbounded until other stopping criterion is met. At lower depth values the constraints on growth limit the performance of the model, as expected. Increasing the depth of each learner allows the model to grow in complexity and focus on finer details in the data. However, increasing the depth of the tree learners beyond 15 branches will have little impact on accuracy, independently of the value of the `minSamplesSplit` hyperparameter. This increasing complexity is likely not required. A maximum depth of around 15 branches for this model would probably result in enough complexity to handle the problem with high accuracy. Since the impact of depth on training time was determined to be minimal, the value for this hyperparameter was set to `maxDepth = None` (fully grown trees).

The right plot in Figure 7.8 shows the variation of the accuracy of the `RFClassifier` with the minimum number of samples required for a branch to be split, `minSamplesSplit`. This hyperparameter regulates the minimum number of samples required to convert a node into a decision node, therefore splitting the respective branch. The impact of this hyperparameter on accuracy is quite small, with noticeable changes being observed only for values beyond 100 samples per node. Beyond this value the accuracy slightly decreases as expected, since the model can no longer separate efficiently the final branches of the tree, that likely contain a mixture of classes. The value chosen for this hyperparameter was `minSamplesSplit = 2`.



**Figure 7.8:** Mean accuracy of the RF model for different values of `maxDepth` (left) and `minSamplesSplit` (right) hyperparameters, for a model with  $N_{trees} = 101$  trees. Each accuracy value is the mean of the final accuracy of three independent training sessions with the same model architecture.

## Training and results

The final model is composed of 101 learners with no limitation of growth, `maxDepth = None`, and minimum number of samples to split a branch set at `minSamplesSplit = 2`. The LZap dataset was divided into a training set and test set with 80–20% splitting ratio –  $8 \times 10^5$  samples for training and  $2 \times 10^5$  for validation.

Table 7.3 displays the confusion matrix of the `RFClassifier` for the test set of the LZap dataset,

using the GMM results as the “actual class” labels. The overall validation accuracy of the RFClassifier model, obtained using Equation 6.1, is  $acc = 99.37\%$ , higher than the estimated accuracy of the HADES algorithm obtained from handscans ( $acc_{\text{HADES}} = 98.6\%$ ). Considering only the classification of S1-like and S2-like pulses, *i.e.*, not considering the mixing of S2 and SE pulse labels to be a misclassification, the validation accuracy is  $acc^{\text{S1S2}} = 99.67\%$ .

**Table 7.3:** Confusion matrix of the RFClassifier results over the test dataset, using the GMM results as the truth labels.

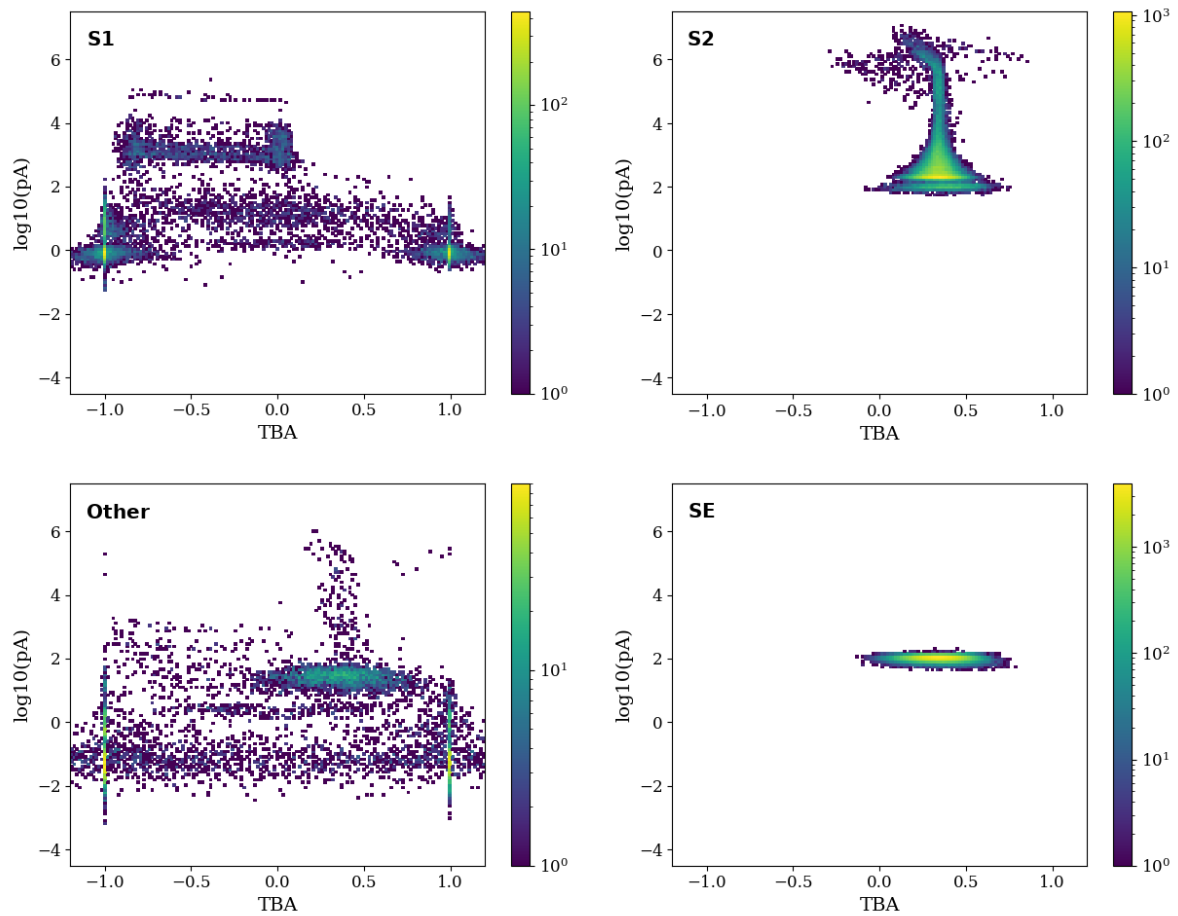
GMM class	RFClassifier Predicted class				Total	
	S1(MPE)	S2	SE	Other		
S1(MPE)	11646	0	0	247	11893	5.9%
S2	0	51114	351	4	51469	25.7%
SE	0	249	128275	3	128527	64.3%
Other	386	11	3	7711	8111	4.1%
<b>Total</b>	12032	51374	128629	7965	200000	

The main failure mode present in these results is the classification of Other pulses as S1 pulses, mainly from SE splits. Conversely, some S1 pulses are being identified as Other pulses by the RFClassifier, maybe hinting at some level of label mixing in the GMM results. However, comparing this misclassification results with the ones from the HADES algorithms shows that the RFClassifier has a higher efficiency at separating SE splits from S1 pulses and vice-versa. From the results presented in Table 6.2, the fraction of Other pulses that HADES classified as S1 pulses is  $\sim 12\%$  and the fraction of S1 pulses that were identified as Other is  $\sim 4\%$ <sup>20</sup>. The same fractions calculated from the results of the RFClassifier are  $\sim 5\%$  and  $\sim 2\%$ , respectively.

These improvements also resulted in a positive predictive value for the RFClassifier of  $PPV = 99.48\%$ , also higher than the result from HADES ( $PPV_{\text{HADES}} = 98.83\%$ ). If the mixing of S2 and SE pulses is not considered a misclassification, this value becomes  $PPV^{\text{S1S2}} = 99.79\%$ , again higher than the result from HADES ( $PPV_{\text{HADES}}^{\text{S1S2}} = 98.98\%$ ). This classification accuracy is not as high as desirable but it is an excellent result considering the limitations imposed by the data at this stage of development of LZap. It is apparent that the RFClassifier model surpassed the HADES algorithm in terms of performance over the same dataset. A direct comparison between the HADES algorithm, the RFClassifier and the remaining ML algorithms tested in this work can be found in Section 7.6.

Figures 7.9 and 7.10 display the results of the predictions of the RFClassifier model for the test dataset, represented in the marginal distributions  $\log_{10}(pA)$  vs  $\log_{10}(pL90)$  and  $TBA$  vs  $\log_{10}(pA)$ , respectively. These distributions can be compared to the plots in Figure 6.4, that displays the full distribution of all pulse populations. A detailed description of the main populations is also given in Figure 6.4 to aid with the interpretation of the results. The overall results obtained with the RFClassifier are very satisfactory and prove that random forests can provide a viable option for the classification of pulses in LZ data, or at least as a dedicated algorithm for assisting the main pulse classifier module.

<sup>20</sup>This value is obtained from a single misclassification occurrence, that bears little statistical significance.

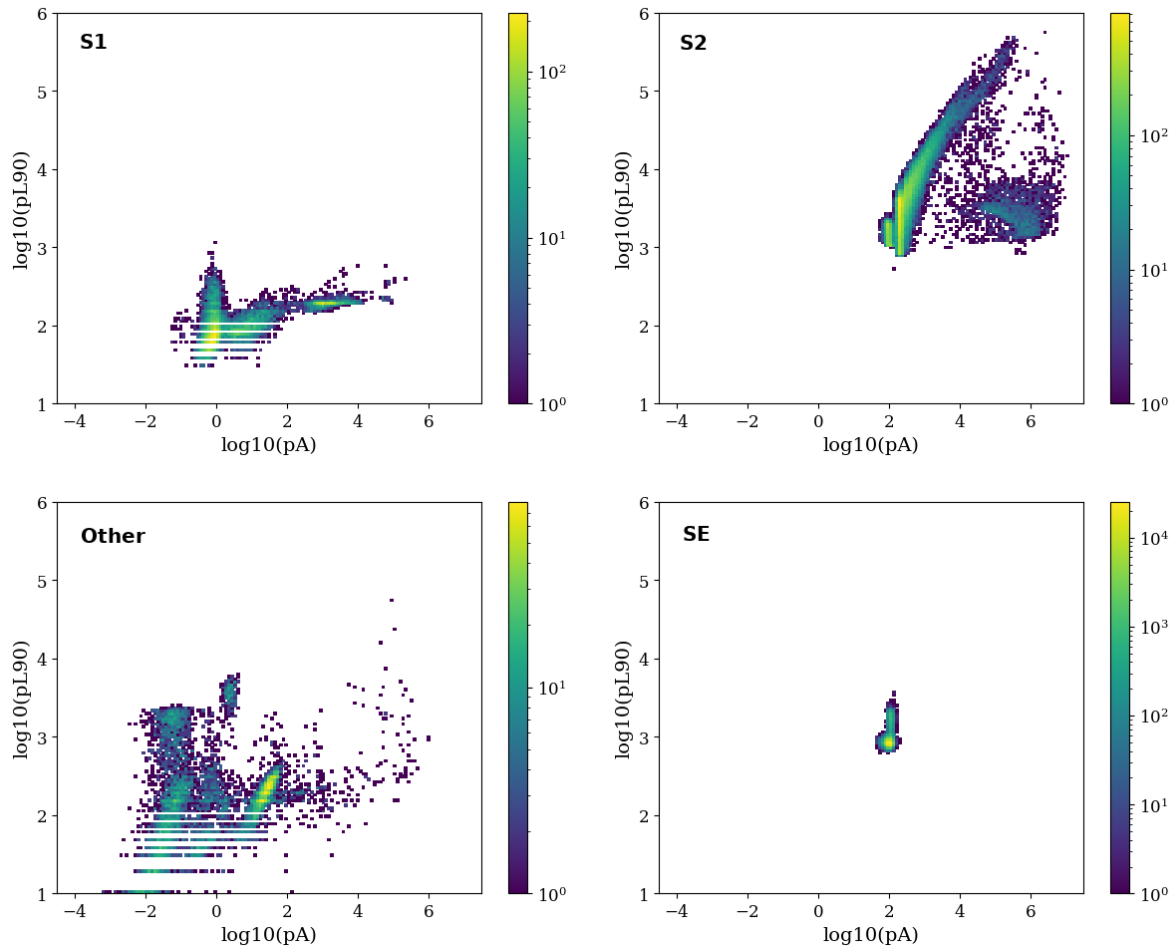


**Figure 7.9:** Distribution of the pulse populations in the marginal distribution  $TBA$  vs  $pA$  after being processed by the RFClassifier. The top-left plot displays the distribution of pulses classified as S1. The top right plot displays the distribution of S2 pulses, while the bottom plots display the Other population on the left and the SE population on the right.

The population of SE split pulses is clearly visible at the center of the bottom-left plot of Figure 7.9, that displays the distribution of the pulses classified as Other by the RFClassifier. This module successfully tagged most of the SE split pulses as Other, leaving only a small number of these spurious pulses in the distribution of pulses classified as S1, displayed on the top-left plot of Figure 7.9. As mentioned above, this ability to successfully classify these bad pulses and remove them from complicated overlaps with good pulses is a strong motivation for exploring random forests as alternatives to heuristic classifiers in LZap (see Sections 6.1 and 6.2 for more details).

The S2 and SE distributions present a high purity level, with no extraneous populations visible in any of the marginal distributions. The only oddity present in the S2 and SE plots is the sharing of pulses with area compatible with SE pulses but with larger pulse length. These pulses are likely single electrons that have either an afterpulsing or a coincident dark count that prompted the *PulseFinder* algorithm to extend the pulse boundaries beyond the limits of the SE waveform.



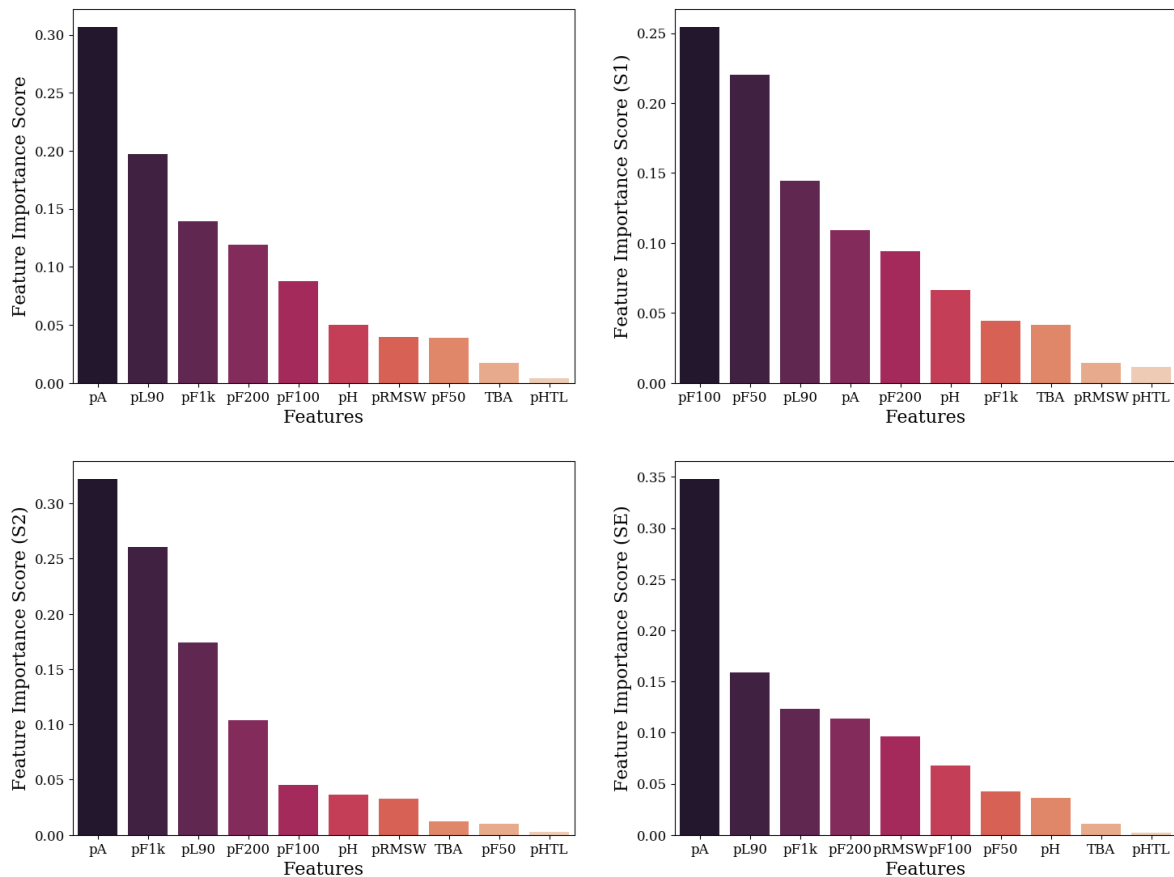


**Figure 7.10:** Distribution of the pulse populations in the marginal distribution  $pA$  vs  $pL90$  after being processed by the RFClassifier. The top-left plot displays the distribution of pulses classified as S1. The top right plot displays the distribution of S2 pulses, while the bottom plots display the Other population on the left and the SE population on the right.

These two sub-populations are likely to be the only contributors to the mixing between SE and S2 pulses. The oddity comes from the fact that the GMM results do not display this separation of the SE population into the main SE pulses and the ones with extended boundaries. The RFClassifier might be seeing these SE pulses with extended boundaries as more closely related to S2 pulses than to the average SE pulse. The pulses classified as SE can still provide a value for the SE size, a very important parameter of the LZ detector, since the only apparent distinction made by the RFClassifier seems to be correlated with pulse length and not area. Furthermore, the real SE pulses that are classified as S2 pulses can be easily recovered by reverting their classification and re-selecting them using an area threshold. Therefore, this distinction between different SE topologies can be easily mitigated using existing techniques and is inconsequential to the data analysis of LZ.

### 7.4.2.1 Feature importance ranking

The overall feature importance ranking presented in the top-left plot of Figure 7.11 was obtained using the trained RFClassifier model.



**Figure 7.11:** Feature importance ranking obtained with Random Forest classifiers. The top left plot shows the overall importance ranking for the final multi-class RFClassifier model. The top-right, bottom-left and bottom-right plots show the importance ranking for separating S1, S2 and SE pulses, respectively, obtained by three separate RFClassifier models trained in a “ONE-vs-ALL” scenario.

The overall model seems to favour pulse area and pulse length and several prompt fractions as the most discriminant features, while the top-bottom asymmetry appears near the bottom of the feature importance score. The *TBA* RQ has always been widely used for S1 and S2 discrimination due to the asymmetrical distribution of light between the top and bottom PMT arrays for these two pulse types. As it can be seen in the right plot of Figure 6.4, S1 pulses tend to have lower *TBA* due to the effect of internal reflection of the S1 light in the liquid-gas interface, while S2 light will be mostly collected at the top since it is produced in the gas phase of the TPC. This seemingly strong separation, while not entirely useful on its own, seems to be disregarded by the RFClassifier.

The results from the feature importance score are easily influenced by the asymmetric represen-


tation of class labels, especially on multi-class problems such as this. The feature importance ranking will be biased towards the parameters that partition the most common pulse classes, in this case the SE pulses that dominate the dataset. The focus on separating SE pulses resulted in the pulse area RQ being considered the strongest discriminant, as the main distinction between SE pulses and S2 pulses, the second most common pulse class, is their area. To fight this bias the multi-class classification problem was converted into three independent binary classification problems for each main pulse class. The same RF model architecture was trained on the LZap dataset, but in each implementation the training labels were collapsed into only two classes: S1-vs-{S2+SE+Other}, S2-vs-{S1+SE+Other} and SE-vs-{S1+S2+Other}. This is referred to as a One-vs-All approach and allows the feature importance ranking to provide information for each highlighted pulse class semi-independently.

The results of this analysis are displayed in the top-right and bottom plots of Figure 7.11. From these results, it is clear that the rankings change significantly for each highlighted pulse class, as expected. The rankings for S2 and SE separation are very similar, with larger emphasis on pulse area, prompt fraction within  $1 \mu\text{s}$  and pulse length. The rankings for S1 separation, however, are radically different, with the strongest parameters being the prompt fractions within 50 ns and 100 ns, and pulse length. The  $pF100$  parameter is especially strong at separating S1 pulses because the 100 ns window used in this parameter is very compatible with the typical S1 width. The top-bottom asymmetry is also more highly valued in the S1 separation than in S2 or SE separation. These results are in agreement with the knowledge collected while building and tuning the HADES and COMPACT algorithms, with HADES focusing strongly on a combination of prompt fraction, pulse area and pulse length thresholds for separating S1-like and S2-like pulses.

The feature importance score is also strongly dependent on correlations between pulse parameters. The parameter bootstrapping process for training a random forest classifier will randomly limit the pulse parameters that are available for each tree in the model, often leading to different trees having access to different but highly correlated parameters, *e.g.*, the prompt fraction RQs. This can lead to the false notion of the importance of these individual features by the tree ensemble. It is crucial to take this bias into account in order to properly determine the best overall parameters for classification of LZap data. The permutation importance score, explained at the beginning of this section, is a more reliable feature importance scoring method than the variable importance ranking presented in Figure 7.11 [Bre01, SBZH07]. Permutation importance is less sensitive to highly correlated variables and thus it might provide a sense of which prompt fraction scale is more appropriate to partition the overall data.

Figure 7.12 shows the permutation importance score for each composite RQ considered in the analysis. The scores were obtained using the fully trained RFClassifier model and can be directly compared to the feature importance scores in the top-left plot of Figure 7.11.

The pulse area and pulse length at 90% area are still considered the best overall discriminants, but unlike the results obtained with the feature importance ranking, the top-bottom asymmetry is now viewed as an overall stronger parameter, as expected. The results from the permutation importance scoring validate the initial assumptions made when building the first pulse classifica-



**Figure 7.12:** Permutation importance ranking obtained with the RFClassifier.

tion algorithms. These results seem to agree with the choices made when building the HADES algorithm since it focuses strongly on the four pulse parameters ranked the highest.

The  $pF100$  parameter scored the highest among the prompt fraction RQs. This result prompted a restructuring and tuning of the HADES algorithm in order to substitute the  $pF50$  RQ being used. This and other proposed modifications of the HADES algorithm will be discussed at the end of this Chapter.

The  $pRMSW$  and  $pHTL$  parameters scored the lowest on both the feature importance and permutation importance ranking studies. This does not mean that these parameters, or others with lower scores, are not useful for pulse classification. The  $pHTL$  parameter was used in HADES to exclude pulses with very strong negative skewness in the S1-like population. Despite being used in niche applications, all parameters can serve a purpose in this analysis.

Determining that HADES is already using the best discriminants among the provided RQs may seem like a dull result, but this is a very important confirmation. The goal of a classification module in the LZap chain is to have the best classification accuracy, and to satisfy that goal it must use the best discriminants. Furthermore, identifying the best discriminants is only the first step in a multi-stage process that also involves selecting the partitioning logic and tuning of the selection thresholds.

### On using HADES labels for training

Before the final design of the RFClassifier was implemented and trained using the GMM clustering results, the classifications provided by the HADES algorithm were used as approximations for the labels of the dataset. This initial choice provided a simple methodology for testing the random forest implementation and potentially yield useful information about ways to improve the HADES algorithm. The first results obtained with this method were in part satisfactory, and some information regarding parameter importance was obtained, but with little impact on the architecture of HADES.

It immediately became clear that using labels provided by a decision tree algorithm like HADES would result in severe overfitting of the random forest model. Since HADES uses simple thresholds on a limited range of parameters to partition the data, the individual trees of the forest were able to easily identify these thresholds, essentially leading to a rigid fit of the results of HADES by the RF model. This may seem like a training success at first, since the model learned how to partition the data like HADES. However, the main goal of this task is to generalize the classification results beyond HADES, while providing information on how to improve it.

Notice that in the results mentioned above the model was not overfitting to the training data but to the full dataset generated with HADES. Both the training and generalization errors were fairly low, meaning that the model was not generalizing but was instead mimicking HADES. Considering that the HADES algorithm is known to have systematic misclassification issues, it was clear that the model was also learning its inaccuracies. The clustering analysis described in Section 7.3 was implemented with the goal of obtaining the label information directly from the dataset and, by doing so, allow the random forest algorithm to generalize beyond HADES.


## 7.5 Neural Network Ensemble Classifier

Several pulse classification tools based on artificial neural networks were developed to evaluate their performance within LZap. The main goal of these tools is to either work as alternative algorithms to HADES in a dedicated LZap module or as auxiliary tools that would complement it. A brief explanation of neural networks is given in Section 7.5.1.

### 7.5.1 Artificial Neural Networks

A neural network is a collective of interconnected units (neurons) capable of receiving, processing and communicating information with each other [Agg18]. Biological neurons are cells that form interconnected structures with each other, connected via specialized cell structures called axons and dendrites. The connecting regions between neurons are called synapses and the strength of those connections can change in response to stimuli. This variability allows the system to adapt to the stimulations accordingly, in what is essentially the learning process.

Artificial neural networks are statistical systems aimed at recreating the structures present in biological brains in order to perform complex computing tasks. Throughout this work the term neural network (NN) will be used to refer to the artificial counterpart, unless mentioned otherwise. The connections between the computational units that form a NN are described by weights that can be modified in response to stimuli in the same way as the synaptic responses in the biological units. The output of a neuron is given by the weighed sum of its inputs passed through an activation function, where the weights are the connection strengths between neurons. Figure 7.13 illustrates the response of a single neuron connected to several inputs  $x_d$  that can be data parameters or the outputs of other neurons. The inputs to a neuron can also be constant responses unrelated to the data, also called bias neurons [Agg18].



**Figure 7.13:** Diagram of an artificial neuron, also called a perceptron. This particular NN architecture is also called a single layer perceptron since it contains only a single computational layer ( $\phi$ , see Equation 7.25) at the output node.

Multivariate data can be fed to an interconnected network of these artificial neurons and propagated through it to compute a function. The input data flows through the network and is modified by the weights and activation functions at each neuron until it reaches the output nodes. These network systems can be trained to solve complex problems by allowing the strength of the connections between neurons to be adjusted while exposing the network to sets of input-output pairs (supervised learning), by allowing the system to find commonalities in the data without external feedback (unsupervised learning) or by a reward-based approach (reinforced learning). The training of a NN is performed by adjusting the weights in response to the prediction error of the model. By iteratively updating the weights of the connections the neural network can encode the data. This work will focus on supervised and semi-supervised learning approaches explained in detail in the following Sections.

Due to the many ways the network and its connections can be arranged, a system with a relatively small number of neurons can display a large complexity. This structural complexity alone makes these systems highly flexible. Even though the individual neurons are very simple computing units, the full NN architecture displays the emergent property of high computational power. Furthermore, for the same network architecture the variability of the strength of the connections can produce a very large number of possible states, each state being a given combination of weights [Agg18]. Each of these states of a single NN will map the output space in a different way, and the sheer number of states available will often guarantee that a particular mapping able to fit the data model exists (provided that the response of the individual neurons is nonlinear, as will be explained below).

The most common NN architecture is a *feed-forward* network, where successive layers of computing neurons are fully interconnected with the preceding and succeeding layers. Other types of architectures are *convolutional neural networks* (CNNs) and *recurrent networks* [Agg18], not discussed in this work. In feed-forward networks, the data is computed at each successive layer, flowing from the input to the output layer. The input nodes of a neural network generally do not perform any level of computation over the data. For that reason this layer is often not counted in the total number of layers of a network.

A network with a single computational layer is called a *single-layer perceptron* (SLP), and its generic architecture is displayed in Figure 7.13. A SLP is composed of a single computational layer with many input nodes linked to a single output node. The output of a SLP is a weighed

linear function of the inputs. Considering a labelled dataset  $(\mathbf{X}, \mathbf{y})$ , with  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  the set of  $N$  data objects with  $D$  parameters each, and  $\mathbf{y} = \{y_1, y_2, \dots, y_N\}$  the respective real value labels, the output computed by a SLP with  $D$  input nodes and a single output node is given by

$$\hat{y} = \Omega \cdot \mathbf{x} = \sum_{d=1}^D \omega_d x_d, \quad (7.20)$$

where  $\Omega = \{\omega_1, \omega_2, \dots, \omega_D\}$  are the weights of the  $D$  connections between each input node and the output node. Having the target labels set as real values implies that the SLP will perform a regression analysis. However, a SLP regressor or any other NN architecture can function as a binary classifier by “pretending” that the class labels are real values [Agg18].

The computed output  $\hat{y}_n$  is compared to the target label  $y_n$  and a loss function can be defined to evaluate the result. Considering a loss function of the least-squares form,  $\mathcal{L}(y, \hat{y}) = (y - \hat{y})^2$ , where the squared term is the error of the prediction, the gradient of  $\mathcal{L}(y, \hat{y})$  with respect to the weights  $\Omega = \omega_d$  (the only free parameters of the SLP) can be written as

$$\frac{\partial \mathcal{L}_d}{\partial \omega_d} = \frac{\partial}{\partial \omega_d} \left( y - \sum_{d=1}^D \omega_d x_d \right)^2 = 2(y - \hat{y}) x_d, \quad (7.21)$$

with  $\mathcal{L}_d$  representing the partial loss at connection  $d$  as a result of weight  $\omega_d$  [Agg18]. The weights can be adjusted in the direction of lower gradients in order to minimize the loss function. This method is called *stochastic gradient descent* (SGD) and is used to minimize the error in the result by updating the weights of the model. It is referred as stochastic because the training algorithm cycles through the data objects in some random order. The updated weights  $\hat{\omega}_d$  can be obtained using Equation 7.22.

$$\begin{aligned} \hat{\omega}_d &= \omega_d + \Delta \omega_d = \omega_d - \frac{\partial \mathcal{L}_d}{\partial \omega_d} \\ &= \omega_d - 2(y - \hat{y}) x_d \end{aligned} \quad (7.22)$$

The update to the weights can be buffered by a constant  $\alpha$ , called the *learning rate*, that can be tuned to control the rate of change of the weights. Equation 7.22 can be generalized for the full set of weights as

$$\hat{\Omega} = \Omega + \alpha (y - \hat{y}) \mathbf{X} \quad (7.23)$$

with  $\hat{\Omega}$  the updated weights after evaluating the inputs  $\mathbf{X}$ . Generally, the weights are not updated after the evaluation of the entire dataset in order to avoid excessive computation. Instead, the update can be performed after evaluating a small batch,  $N'$ , of the inputs, since a loss function can be defined as a linearly separable sum of the individual loss functions for each individual training data point,

$$\mathcal{L} = \sum_{n=1}^{N'} \mathcal{L}_i \quad (7.24)$$

Smaller batches lead to more frequent updates but to a general loss in precision, while larger

batches result in less frequent but on average more accurate updates to  $\Omega$ . This approximation of gradient descent by averaging the gradient over a batch of data is referred to as *mini-batch stochastic gradient descent* [Agg18]. Each batch will result in one update of the weights, and once all training data has been passed through the model (without replacement), the training is said to have completed an *epoch*<sup>21</sup>. The learning rate and batch size are two important hyperparameters of SLPs and other NN architectures, and can greatly influence the speed of training and variance of the final state. A more detailed discussion of these two parameters is presented later in this Section for the generic example of a multi-layer feed-forward network.

### Activation functions

So far the SLP showcased is assumed to have a linear activation function, *i.e.*, the output of the SLP is proportional to the sum of the weighed inputs, with the proportionality constant equal to 1 in this particular case. However, it can be shown that any feed-forward network that exclusively uses linear activation functions in the computational layers can be collapsed into a SLP that computes a linear model of the inputs, regardless of the number of hidden layers [HTF09, Agg18]. The identity function was usually used in multi-class classification problems, but the limitations imposed on the model forced the implementation of other nonlinear functions and led to the popularity of the *softmax* function [HTF09].

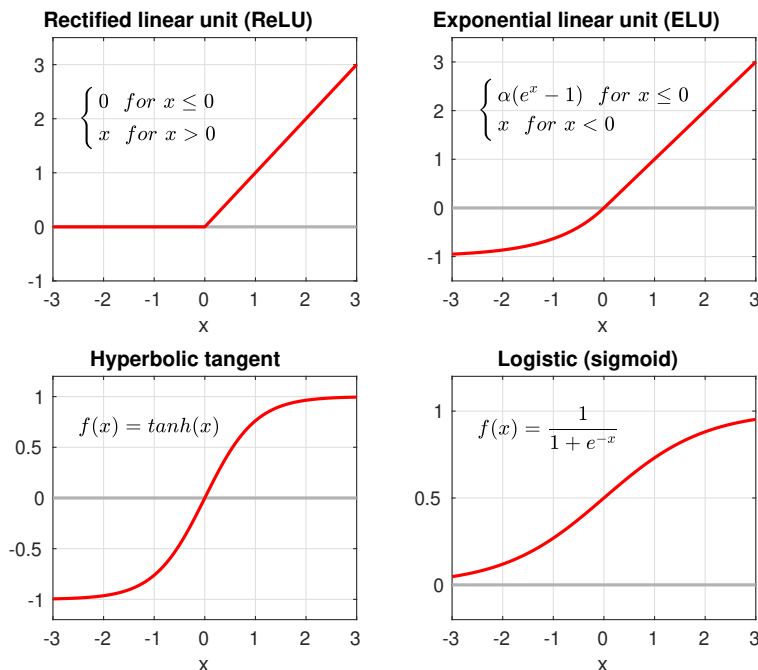
In most applications the desirable activation functions of a SLP or any other architecture are continuously-differentiable monotonic nonlinear functions. It is worth mentioning that a piecewise function can be used if the derivative is defined in each respective sub-domain. Moreover, the usage of piecewise activation functions can be extremely useful to simplify the calculation of the gradient of the loss function. The nonlinearity of the activation functions allows any NN system to transform the input space into any linearly separable space [Agg18]. Furthermore, if the activation functions are also squashing functions, *i.e.*, bounded from below and above, any network with at least one computational layer can approximate to some degree any multidimensional function, providing that the computational layer contains a sufficient number of units. Feed-forward networks of this type are known as “universal approximators” [HSW89].

Figure 7.14 shows some examples of common activation functions. The rectified linear unit (ReLU) is an example of a piecewise function that is widely used in NN applications due to the simplification of calculations it imposes during gradient descent [Agg18]. The derivative of the ReLU function can be defined in each of the sub-domains, being zero for  $x \leq 0$  and one for  $x > 0$ . However, this simplicity can lead to severe problems, most notably the “death” of neurons. A neuron using a ReLU activation can become unresponsive if it expects non-negative inputs but somehow its weights are initialized as negative, or vice versa. If the learning rate is too large then the fluctuations of the weights when training can also become negative due to large corrective leaps. When a ReLU unit becomes inactive to all inputs the gradient cannot be propagated to preceding layers and the weights cannot be updated further [Agg18]. Several

---

<sup>21</sup>An epoch is simply a term to describe the full pass of training data through the network and it is only used for logging and benchmarking purposes, having no special function rather than inform that all the data available has been processed.





**Figure 7.14:** Examples of common activation functions [Agg18, CUH15].

modifications to the ReLU activation can prevent this dying effect. The exponential linear unit (ELU), shown in the top-left plot of Figure 7.14, is also a piecewise function like ReLU but instead of setting the response of the neuron to zero for any  $x \leq 0$  it assigns the term  $\gamma(e^x - 1)$  that is bounded from below at

$$\lim_{x \rightarrow -\infty} \text{ELU}(x, \gamma) = -\gamma.$$

The derivative on the sub-domain  $x \leq 0$  is also trivial to compute, being simply defined as  $\text{ELU}(x, \gamma) + \gamma$ . Having a non-zero gradient in all the domain of the ELU prevents the unit from becoming inactive and ensures that the gradient can be propagated.

The “dying ReLU problem” is related to a broader issue of vanishing gradients. This effect can be described as the successive attenuation of the gradient of the loss function as it is propagated through the network [Agg18]. It is most notable when using an activation function whose derivative maxes-out below unity, like in the case of the sigmoid activation<sup>22</sup> represented in the bottom-left plot of Figure 7.14.

The output of the SLP model now with a nonlinear activation function  $\phi$  on the output node is given by:

$$\hat{y} = \phi(\Omega \cdot \mathbf{x}) = \phi\left(\sum_{d=1}^D \omega_d x_d\right) \quad (7.25)$$

<sup>22</sup>The maximum value of the first derivative of the logistic (sigmoid) function is 0.25, meaning that successive gradient descent steps across  $k$  computational layers will result in an attenuation of the gradient by a factor of at least  $0.25^k$ .

With a non-linear activation function, the gradient of the loss function needs to be modified accordingly. The gradient of  $L(y, \hat{y})$  with respect to the weights  $\omega_d$  can be decomposed by the chain rule as

$$\frac{\partial \mathcal{L}_d}{\partial \omega_d} = \frac{\partial \mathcal{L}_d}{\partial y} \frac{\partial y}{\partial \chi_d} \frac{\partial \chi_d}{\partial \omega_d}, \quad (7.26)$$

where the term  $\chi_d$  represents the input of the neuron before the activation function is applied, *i.e.*, the weighed sum of the inputs  $x_d$ ,

$$\chi_d = \sum_{d=1}^D \omega_d x_d.$$

The first term of Equation 7.26 is the partial derivative of the error function with respect to the output of the SLP, and can be written as

$$\frac{\partial \mathcal{L}_d}{\partial y} = \frac{\partial}{\partial y} (y - \hat{y})^2 = 2(y - \hat{y}) \quad (7.27)$$

The second term of Equation 7.26 is the partial derivative of the output of the SLP with respect to the neuron input  $\chi_d$ , which is simply the derivative of the activation function  $\phi$  of the output neuron,

$$\frac{\partial y}{\partial \chi_d} = \frac{\partial}{\partial \chi_d} \phi(\chi_d) = \phi' \left( \sum_{d=1}^D \omega_d x_d \right). \quad (7.28)$$

Finally, the third term of Equation 7.26 is the derivative of the input of the neuron with respect to the weights  $\omega_d$ . Since the input is simply the linear weighed sum of the inputs, this derivative can be simplified to

$$\frac{\partial \chi_d}{\partial \omega_d} = \frac{\partial}{\partial \omega_d} \left( \sum_{j=1}^D \omega_j x_j \right) = x_d. \quad (7.29)$$

Equation 7.26 can then be rewritten for the SLP as

$$\frac{\partial \mathcal{L}_d}{\partial \omega_d} = 2(y - \hat{y}) \phi' \left( \sum_{d=1}^D \omega_d x_d \right) x_d \quad (7.30)$$

The result in Equation 7.22 now becomes

$$\begin{aligned} \hat{\omega}_d &= \omega_d - \alpha \frac{\partial \mathcal{L}_d}{\partial \omega_d} \\ &= \omega_d - 2\alpha (y - \hat{y}) \phi' \left( \sum_{d=1}^D \omega_d x_d \right) x_d, \end{aligned} \quad (7.31)$$

which is fairly similar to the original result only with the additional correction from the non-linearity of the activation function. By taking a closer look at this result it becomes quite

apparent why the vanishing gradient problem can lead to the suppression of the weights of multi-layer networks and the eventual inability of the system to learn: If the derivative of the activation function becomes smaller than unity, and that results in the overall update term also being strictly smaller than one, then successive updates of the weights will result in successively smaller values. Furthermore, in architectures with more than one computational layer (explained below), propagating the gradient of the loss function backwards into the network will result in the successive attenuation of the weights further up the network. Avoiding that the gradient either vanishes or “explodes” to higher values is often managed by carefully choosing the best activation for the computational nodes of the network<sup>23</sup> or by implementing some sort of regularization of the gradient via adaptive learning rates, conjugate gradient methods or *batch normalization* [Agg18]. The batch normalization method can be described as the implementation of normalization layers either before the neuron activation (pre-activation normalization) or at the output of the unit (post-activation normalization) in order to regulate the magnitude of the gradients over a batch of training data [Agg18].

### Multi-layer perceptron and generic feed-forward networks


The concept of gradient descent used for training a SLP can be exported to more complex NNs with more than one computational layer. Figure 7.15 shows a schematic representation of a multi-layer perceptron (MLP). This architecture retains a high complexity while being manageable from a computational standpoint. The layers between the input and output layers are called “hidden” layers, since they do not communicate directly with the “outside” of the network [HTF09]. This will be the architecture explored in this work. The number of hidden layers, *i.e.*, the depth of the network, can be arbitrarily large, as well as the number of units in each layer. However, having more layers or units per layer will increase the training time very rapidly whilst not guaranteeing better results after a certain optimal network size. The specific architecture of the network should be optimized for the target problem to avoid excessive computation on one side and ensure performance on the other.

The input units in the generic MLP like the one depicted in Figure 7.15 only transmit the input variables into the computational layers without modifying their values. For convenience, let the nodes of the input layer be denoted as  $X_d$  for the  $d^{\text{th}}$  place in the layer, that will carry the  $x_d$  data feature of each  $D$ -dimensional data object  $\mathbf{x}_n \in \mathbf{X}$ , with  $n = 1, 2, \dots, N$ . The number of input nodes is therefore equal to the dimensionality of the dataset plus some bias neuron that might be included. In this deduction of the training process of the MLP the bias neurons are substituted by an input variable with constant positive value that is considered to be part of the input data, for simplicity. This is a common approach to biasing that is equivalent to having a dedicated bias neuron [Agg18].

The output of the hidden units in the first and second hidden layers are denoted as  $h_p^1$  and  $h_q^2$ , respectively, with  $p = 1, 2, \dots, P$  and  $q = 1, 2, \dots, Q$  the indices of each hidden unit in the respective layer. In a fully connected feed-forward model the full output of a layer is the

---

<sup>23</sup>The derivative of the ReLU activation is always 1 for any positive argument, therefore being resistant to either vanishing or exploding gradients.



**Figure 7.15:** Schematic representation of a feed-forward neural network, with  $D$  input nodes, 2 hidden layers with  $P$  and  $Q$  units each, and  $K$  output nodes. The connection between nodes have variable weights represented by the matrix  $\Omega = \Omega^l$ , with  $l = 1, 2, 3$  the index of the  $L = 3$  computational layers. This architecture is often called a *feedforward multi-layer perceptron*.

input of the following layer, so it is convenient to define the full set of outputs for each layer as  $H^1 = \{h_{1,1}^1, h_{1,2}^1, \dots, h_{1,P}^1\}$  and  $H^2 = \{h_{2,1}^2, h_{2,2}^2, \dots, h_{2,Q}^2\}$ . Here the superscript indices indicate the hidden layer number, and increase in the direction of the output layer.

The response of the first hidden layer is given by the linear combination of the inputs of each node  $p$  modified by the activation function  $\phi()$ ,

$$h_p^1 = \phi \left( (\omega_p^1)^T \cdot \mathbf{x} \right) = \phi \left( \sum_{d=1}^D (\omega_{pd}^1)^T \cdot x_d \right) \quad (7.32)$$

where  $\omega_p^1$  is the  $D \times P$  matrix of weights connecting the unit  $p$  of the first hidden layer with the  $D$  nodes of the input layer. The output  $h_p^1$  will be fed into the next layer with  $Q$  hidden units. For simplicity, and usually good practice, the same activation function  $\phi()$  is assumed for all hidden nodes. The output of the second hidden layer can then be written, much like the previous result, as

$$h_q^2 = \phi \left( (\omega_q^2)^T \cdot H^1 \right) = \phi \left( \sum_{p=1}^P (\omega_{qp}^2)^T \cdot h_p^1 \right) \quad (7.33)$$

The similarity between Equations 7.32 and 7.33 is not surprising. What a fully connected layer

does in a feed-forward model is sum the weighed inputs of that layer and then apply a function to that sum, so the results for the outputs of any hidden layer at any depth can simply be written as a function of the outputs of the preceding layers. At each computational layer, the outputs can be viewed as a new set of features derived from the initial data features at the input layer [HTF09]. This is essentially a chaining of outputs of successive computational layers, as expected from a feed-forward MLP architecture.

The number of units in the output layer usually depends on the type of task being performed. For regression, usually a single output node is used to compute the predicted value of the function being fitted by the model. For classification tasks, such as the ones concerning this work, the number of output nodes is usually set as the total number of class labels in the data,  $K$ . The output layer will receive the calculated features of the last hidden layer as its input, denoted by  $H^2 = \{h_1^2, h_2^2, \dots, h_q^2, \dots, h_Q^2\}$  in this example, with  $q$  the index of the units of the final hidden layer, and consequently the index of the new set of features produced at that final layer. The weighted sum of the features  $H^2$  at the  $k^{\text{th}}$  unit of the output layer is then defined as

$$T_k = \omega_k^T \cdot H^2 = \sum_{q=1}^Q (\omega_{kq}^2)^T \cdot h_q^2 \quad (7.34)$$

where  $\omega_{kq}$  denotes the weight of the connection between the  $q^{\text{th}}$  node of the preceding layer with the  $k^{\text{th}}$  node of the output layer.

Notice that at this stage the weighted sums have not yet been subjected to any transformation at the output nodes. This transformation should be treated differently from the ones provided by the activation functions at the hidden units, since it will convert the final encoding of the features at the final hidden layer into a set of outputs that should encode the predicted class in some way. There are different strategies to convert the multiple outputs of the output layer into the final class prediction. As an example, one can simply choose the argument of the maximum value at the output as the predicted class,  $\hat{y} = \arg \max_k \{T\}$ . An interesting approach that is widely used in multi-class classification is to implement the softmax function, defined as

$$\sigma_k(T) = \frac{e^{T_k}}{\sum_{j=1}^K e^{T_j}} \quad (7.35)$$

as the final processing step. This way each output node can be interpreted as a probability for each of the classes  $K$ , since the output of the softmax function over all output nodes sums to unity, by definition. The output of the network can now be written as a function of the final set of modified features  $T$

$$\hat{y}_k = \sigma_k(T). \quad (7.36)$$

Since in this setup the output takes the form of a vector with each element corresponding to a target class, the training labels need to be converted to categorical vectors  $y_k$  that encode the class labels in their indices.

An interesting property can be derived from this result: considering that the model being

described computes a function  $f(\mathbf{X}) = \hat{y}$ , with  $\hat{y}$  the prediction of the model, then this prediction can also be computed by another function that takes as arguments the modified features obtained at each computational layer,  $H^1, H^2, \text{etc.}$ , such that [HTF09]

$$f(\mathbf{X}) = g(H^1) = h(H^2) = \sigma_k(T) = \hat{y}. \quad (7.37)$$

The functions  $g, h$  can be interpreted as the functions computed by the networks obtained by truncating the original network  $f$  at the first and second hidden layers, respectively. This procedure is repeatable for any number of hidden layers and will be the basis of the *backpropagation* algorithm used to apply gradient descent when training the model, explained below.

The output of the model can be rewritten as

$$\begin{aligned} \hat{y} &= \sigma(T) = \sigma\left((\Omega^3)^T \cdot H^2\right) \\ &= \sigma\left((\Omega^3)^T \cdot \left(\phi\left((\Omega^2)^T \cdot H^1\right)\right)\right) \\ &= \sigma\left((\Omega^3)^T \cdot \left(\phi\left((\Omega^2)^T \cdot \left(\phi\left((\Omega^1)^T \cdot \mathbf{x}\right)\right)\right)\right)\right) \end{aligned} \quad (7.38)$$

The generalization of Equation 7.38 for a MLP with  $L$  computational layers can be expressed as:

$$\hat{y} = \sigma\left((\Omega^L)^T \cdot \phi\left((\Omega^{L-1})^T \dots \phi\left((\Omega^l)^T \dots \phi\left((\Omega^1)^T \mathbf{x}\right) \dots\right) \dots\right)\right) \quad (7.39)$$

## Backpropagation

Equations 7.32 and 7.33 can be rewritten to include the weighted sum of the inputs of the respective layers, defined as  $I^1 = (\Omega^1)^T \cdot \mathbf{x}$  and  $I^2 = (\Omega^2)^T \cdot H_1$ , becoming  $H^1 = \phi(I^1)$  and  $H^2 = \phi(I^2)$ , respectively. The gradient of a generic loss function  $\mathcal{L}(y_n, f(\mathbf{x}_n))$  with respect to the weights of the model can be expanded, applying the chain rule, in the following manner:

$$\frac{\partial \mathcal{L}}{\partial \Omega} = \frac{\partial \mathcal{L}}{\partial H^L} \cdot \frac{\partial H^L}{\partial I^L} \cdot \frac{\partial I^L}{\partial H^{L-1}} \cdot \frac{\partial H^{L-1}}{\partial I^{L-1}} \cdot \frac{\partial I^{L-1}}{\partial H^{L-2}} \dots \frac{\partial H^1}{\partial I^1} \cdot \frac{\partial I^1}{\partial \Omega} \quad (7.40)$$

where  $\Omega$  is the complete matrix of weights of the model. The following solutions can be found to the inner terms of the chain:

$$\frac{\partial H^l}{\partial I^l} = \phi'(I^l) \quad (7.41)$$

$$\frac{\partial I^l}{\partial H^{l-1}} = (\Omega^l)^T \quad (7.42)$$

The final partial derivative yields  $\mathbf{x}$ . Equation 7.40 can now be simplified as

$$\frac{\partial \mathcal{L}}{\partial \Omega} = \frac{\partial \mathcal{L}}{\partial H^L} \cdot \phi'(I^L) \cdot (\Omega^L)^T \cdot \phi'(I^{L-1}) \cdot (\Omega^{L-1})^T \dots \phi'(I^1) \cdot \mathbf{x} \quad (7.43)$$

Similarly to the result in Equation 7.37, the gradient can also be expanded into a succession of terms  $\delta^l$ , interpreted as the losses at the layer  $l$  [Agg18]. These terms are defined as

$$\delta^l = \phi'(I^l) \cdot \Omega^{l+1} \dots \Omega^{L-1} \cdot \phi'(I^{L-1}) \cdot \Omega^L \cdot \phi'(I^L) \cdot \nabla_{\Omega^L} \mathcal{L} \quad (7.44)$$

Notice that the ordering of the matrix multiplication was inverted, and the weight matrices were transposed accordingly.

The gradient of the loss with respect to the weights of the layer  $l$  can be written as

$$\nabla_{\Omega^l} \mathcal{L} = \delta^l \cdot (H^{l-1})^T \quad (7.45)$$

The values of the elements of  $\delta^l$  can be calculated recursively with the following expression:

$$\delta^{l-1} = \phi'(I^{l-1}) \cdot (\Omega^l)^T \cdot \delta^l \quad (7.46)$$

The evaluation of the loss terms  $\delta^l$  is performed from the output layer of the network towards the input, with the preceding loss terms being derived from the ones in the following computational layer. This means that the errors are propagated backwards in the network, and thus the term *backpropagation* for this method.

The backpropagation algorithm can be implemented in a two-step process: first, a batch of input data is fed forward into the current state of the network and the predictions  $\hat{y}_k$  are calculated; then the errors are computed using the predictions  $\hat{y}_k$  and the corresponding training labels  $y_k$ , and are backpropagated via  $\delta^l$ , with the gradients being computed to update the weights accordingly. These steps are depicted in Algorithm 6. A detailed explanation of the backpropagation algorithm can be found in [Agg18]

The loss function typically used for regression is constructed using the aforementioned sum-of-squares error, but for classification the most common choice is the deviance or cross-entropy, defined as:

$$\mathcal{L}(\Omega) = - \sum_{k=1}^K y_k \log \hat{y}_k = - \sum_{k=1}^K y_k \log f_k(\mathbf{x}) \quad (7.49)$$

A MLP that uses softmax activation at the output layer and cross-entropy loss in optimization is equivalent to a linear logistic regression model [HTF09].

### Optimizers and parameter-specific learning rates

When using stochastic gradient descent, the weights of the NN may not be updated in an optimal direction towards lower gradients, since fluctuations in the gradient calculations, local minima or large leaps may result in erratic movement down the potential. This erratic descent results in successive updates that often partially cancel each other and reduce the overall step towards the optimal state, hindering the learning process. Some optimization methods can be implemented to the learning process in order to improve the consistency of the updates to the weights and

---

**Algorithm 6:** Stochastic gradient descent with backpropagation

---

1. **Forward pass:** From a randomly selected batch of labelled training data, compute  $\hat{y}_n = f(\mathbf{x}_n)$  with the current weights.
2. Calculate the total loss  $\mathcal{L}(y_n, \hat{y}_n)$  associated with the prediction  $\hat{y}_n$  and target label  $y_n$ .
3. Calculate  $\delta^L$  for the final layer,  $L$ , using

$$\delta^L = \sigma'(T) \cdot \nabla_T \mathcal{L} \quad (7.47)$$

4. **Backward pass:** For each hidden layer  $l = L, L - 1, \dots, 2, 1$ , starting from the last:
  - (a) Move up the hidden layers and compute  $\delta^{l-1}$  using the result  $\delta^l$  calculated on the previous layer

$$\delta^{l-1} = \phi'(I^{l-1}) \cdot (\Omega^l)^T \cdot \delta^l$$

- (b) Compute the gradient at layer  $l$  using

$$\nabla_{\Omega^l} \mathcal{L} = \delta^l \cdot (H^{l-1})^T$$

- (c) repeat from (a) until the first hidden layer is reached.

5. Update the weights  $\Omega^l$  using

$$\hat{\Omega}^l = \Omega^l - \alpha \left( \nabla_{\Omega^l} \mathcal{L} \right) \quad (7.48)$$

6. Repeat from step 1 until the loss  $\mathcal{L}$  converges or, alternatively, until the loss calculated over a test set,  $\mathcal{L}_{test}$ , begins increasing over some training epochs (early stopping)
- 

accelerate training. The most common techniques are learning rate decay, momentum-based descent and parameter-specific learning rates [Agg18].

As mentioned above, the learning rate  $\alpha$  is a common hyperparameter that regularizes training by controlling the size of the updates to the weights. A smaller learning rate will make training slower but avoids big leaps in the gradient descent that may be highly inefficient, while a larger rate will accelerate training but will possibly overshoot the best solution. A desirable solution would be to limit the size of the steps as the training reaches the optimal state, which can be achieved by gradually decreasing the value of the learning rate. This is known as “learning rate decay” and is implemented by modifying the value of the initial learning rate  $\alpha_0$  at each training epoch  $t$ , *e.g.*,

$$\begin{aligned} \alpha_t &= \alpha_0 e^{-\kappa \cdot t} && \text{(exponential decay)} \\ \alpha_t &= \frac{\alpha_0}{1 + \kappa \cdot t} && \text{(inverse decay)} \end{aligned}$$

with  $\kappa$  being the parameter that controls the rate of decay.

Another common practice for optimizing gradient descent is to add a recursive term to the



gradient vector that smooths and averages the fluctuations:

$$\hat{\mathbf{V}} = \beta \mathbf{V} - \alpha \frac{\partial \mathcal{L}}{\partial \Omega} \quad (7.50)$$

$$\hat{\Omega} = \Omega + \hat{\mathbf{V}} \quad (7.51)$$

where  $\mathbf{V}$  is the current gradient vector of the loss function,  $\hat{\mathbf{V}}$  is the updated gradient vector,  $\Omega$  is the vector of weights at the current state,  $\hat{\Omega}$  is the updated vector of weights,  $\alpha$  is the aforementioned learning rate regularization factor and  $\beta$  is the *momentum* parameter that performs an exponential smoothing of the gradient vector. This smoothing is the result of favouring the components of  $\mathbf{V}$  that are in a consistent direction over many steps by attenuating the components that tend to change direction often, resulting in accelerated learning. A common issue of this algorithm is the overshooting of the targeted result since the final steps are enhanced in a particular direction by the momentum [Agg18].

The momentum approach is applied uniformly to every weight being updated in order to smooth the overall direction of the gradient descent vector  $\mathbf{V}$ . This method can be substituted by a different approach that involves modifying the learning rate for different weights, and penalizing weight updates with large oscillations caused by large partial derivatives. This is also a gradient regularizing technique but centred around smoothing individual gradient components, making it a more targeted method. Some common optimizer methods are *AdaGrad* [DHS11], *RMSprop* [Agg18, Hin12], *Adam* [KB14], among many others derived from these.

The *AdaGrad* optimization method [DHS11] adjusts the learning rate parameter at each step of gradient descent with a term  $A_i$  that is the summed square of the partial derivative with respect to each parameter  $\omega_i$ ,

$$\hat{A}_i = A_i + \left( \frac{\partial \mathcal{L}}{\partial \omega_i} \right)^2 \quad (\text{AdaGrad}) \quad (7.52)$$

$$\hat{\omega}_i = \omega_i - \frac{\alpha}{\sqrt{\hat{A}_i}} \left( \frac{\partial \mathcal{L}}{\partial \omega_i} \right).$$

This term benefits changes along a particular  $\omega_i$  direction on the gradient if those changes are consistently small during training, while penalising the changes in directions that have been experiencing larger leaps. The major drawback of *AdaGrad* is the eventual stalling of the training as  $A_i$  keeps increasing, causing updates to stop often before the optimal state is reached. Furthermore, even if the updates to a misbehaving component become stable and point in the right direction, that component will still be penalised due to the older behaviour since  $A_i$  takes all the history of the squared magnitude of the gradient into account [Agg18].

A solution to the limitations of the *AdaGrad* optimization could be to update  $A_i$  with averages instead of aggregation. The *RMSprop* algorithm [Agg18, Hin12] modifies the calculation of  $A_i$

by introducing a decay parameter  $\rho \in [0, 1]$  such that  $A_i$  becomes a running average of the type:

$$\begin{aligned}\hat{A}_i &= \rho A_i + (1 - \rho) \left( \frac{\partial \mathcal{L}}{\partial \omega_i} \right)^2 && (RMSprop) \\ \hat{\omega}_i &= \omega_i - \frac{\alpha}{\sqrt{\hat{A}_i}} \left( \frac{\partial \mathcal{L}}{\partial \omega_i} \right).\end{aligned}\tag{7.53}$$

This optimization retains the advantages of *AdaGrad* but now the impact of older updates to the weights will exponentially decrease with each successive update. However, since  $A_i$  must be initialized to zero the first gradient descent steps will be biased.

The *Adam* optimizer [KB14] combines the running average term of *RMSprop* (Equation 7.53) with a modified linear term  $B_i$  similar to the *momentum* method in Equation 7.51, but in this case that term also exponentially decays with each successive update:

$$\begin{aligned}\hat{B}_i &= \lambda B_i + (1 - \lambda) \left( \frac{\partial \mathcal{L}}{\partial \omega_i} \right) && (Adam) \\ \hat{\omega}_i &= \omega_i - \frac{\alpha_t}{\sqrt{\hat{A}_i}} \hat{B}_i.\end{aligned}\tag{7.54}$$

The term  $B_i$  performs a smoothing to the modified learning rate  $\alpha_t$ , that is dependent on the two decay rates  $\rho$  and  $\lambda$ , such that

$$\alpha_t = \alpha \left( \frac{\sqrt{1 - \rho^t}}{1 - \lambda^t} \right)\tag{7.55}$$

With the usage of the first-order and second-order parameters  $A_i$  and  $B_i$ , the *Adam* algorithm incorporates many of the techniques mentioned previously for optimization, making it an extremely popular optimiser [Agg18]. Other common optimizers like *Nadam* [Doz16] or *Adamax* [Cho15] are variations of these algorithms explained here.

### Overfitting and regularization

Neural networks are very prone to overfitting due to having many degrees-of-freedom, in the form of the connection weights, to fit the data. The careful mitigation of overfitting can be accomplished by different methods, two of the most common are *early stopping* of the training and the implementation of *dropout* [Agg18].

As mentioned in Section 7.1.2, the dataset used for training is usually divided into a training set and a generalization set of data, the latter used to monitor the progress of training in an independent way from the data that the model is using to learn. Neural networks are no exception for this methodology. By monitoring the loss function over the generalization dataset, or an equivalent error rate, it is possible to determine whether the model is starting to lose generalization power and beginning to overfit the training data, indicated by a rise in the error

rate over the generalization set. When this happens gradient descent is stopped and training ceases, leaving the model in its current state. This is the generic principle of *early stopping*. The premature stopping will most likely result in a sub-optimal state of the model, but it ensures that the loss of generalization is minimal.

The *dropout* method shares some similarities with the ensemble methods explained in Section 7.4.1, namely Random Forests, and consists in randomly disrupting some nodes of the network on each training batch. This disruption is essentially equivalent to set the input of that node to zero, meaning that that particular feature will not be propagated through the network. This is often implemented by introducing *dropout* layers within the hidden layers of the model, that will set some node inputs to zero at random with a certain probability  $p$ , effectively limiting the phase-space available for the remaining of the network [C<sup>+</sup>15, Agg18].

### Motivations for neural networks and NN ensemble models

There are many different subspecies of pulses in LZ simulated data, despite being generated by the same processes, as discussed in Section 6.2.1 and visible in Figures 6.4 and 6.7. As an example, the primary scintillation of xenon produced by an energy deposition (the S1 signal) can generate many different pulse topologies, some with widely different implications for the physics analysis of LZ than others. When LZ starts collecting data it is likely that even more families of pulses will be generated due to the stochastic nature of data collection, detector physics, electronics, and so many other processes that are not accounted for or not well described by simulations.

Most of the classification models explored in this work provide hard labels<sup>24</sup> for the predicted classes of the processed pulse data, with the only exception being the COMPACT algorithm explained in Section 6.2.4. There is a strong case to be made to obtain probabilistic information from the classification models instead of hard classifications, since a probabilistic result can provide more information about the nature of the pulse and the degree of ambiguity of the prediction. Furthermore, it might help to distinguish between the main type of pulses and the secondary ones, *e.g.*, separating real S2s from S2 tails and SE pileup, all classified as “S2” by a hard labelling algorithm such as HADES (see Section 6.2.3). The RQ structure of LZap is prepared to handle probabilistic values for each pulse class available instead of hard labels. This was implemented in order to take advantage of the probabilistic nature of the COMPACT algorithm, but it was never fully used due to the downfall of the algorithm (see Section 6.2.4).

The ML-based ensemble methods explored in Section 7.4 have the ability to provide probabilistic information by exploring the variance in the results from all the ensemble of classifiers. However, this is not a native property of these methods, unlike the neural network models explained above. A NN classifier can implement probabilistic prediction on a classification task natively by having multiple output nodes, with each being related to the probability of a pulse object being of a certain class. The class labels of the LZap dataset can be converted to categorical vectors of size  $K$  of the form  $y = \{y_1, y_2, \dots, y_K\}$ , with only the  $k^{\text{th}}$  component corresponding to the class label

<sup>24</sup>A discrete numerical or categorical label.

of the pulse being equal to 1 and all the others set to zero. This vector satisfies the conditions  $y_k \in \{0, 1\}$  and  $\sum_k y_k = 1$ , thus being a valid probabilistic object. A probabilistic vector provides more information about the nature of the pulses, but a hard label can be retrieved from the categorical vector if needed. Using the predicted categorical vector  $\hat{y}$ , the most probable pulse class can be assigned to the pulse by

$$\hat{y}_{\text{label}} = \arg \max_k \hat{y}. \quad (7.56)$$

Even though the results from a NN classifier may be hard to interpret and scrutinize, since NNs are a form of “blackbox” algorithm with high internal complexity, they provide a mean of exploration of the full potential of the analysis framework of LZ when using probabilistic information at the pulse-level classification.

Several multi-class NN classification tools were developed and tested during this work, where a single NN was used to perform multi-class classification over the LZap dataset. However, it was established that an ensemble of MLP binomial logistic regressors<sup>25</sup>, each trained in a One-vs-All configuration, returned better overall results than a single NN classifier model. This method is presented in Section 7.5.2. In a One-vs-All problem, each of the MLPs focuses on learning how to distinguish a single class from the remaining (specialization). However, this methodology deviates from a full probabilistic analysis, reducing the problem to an ensemble of bivariate classifications with limited information on the mixing probabilities between different classes. Nevertheless, this may be a generous trade-off between a full description of the pulse object in a probabilistic class space and an improved generalization ability, accuracy, and outlier detection capabilities.

Ensembles of neural network ensembles have been around for at least three decades [HS90, ZWT02, AM07, Tao19], but the implementations are often based on generalizations of boosting or bagging strategies applied to NN models or focus solely on deep learning applications. In this work, the focus is to find alternative solutions to the limitations of shallow NN implementations on multi class classification problems, based on the “strength of many” paradigm surrounding ensembles of weak classification models. The general description of this approach to a NN ensemble is presented below. To keep the same label notation used throughout this work, the effective number of classes with a dedicated NN classifier in the ensemble is described by the variable  $K'$ , meaning that the ensemble model is composed of  $K'$  NN classifiers. The prediction of the ensemble of  $K'$  NN classifiers is represented by a function  $\hat{y} = f(\mathbf{x})$ , with  $\hat{y} = \{\hat{y}_1; \hat{y}_2; \dots; \hat{y}_{k'}; \dots; \hat{y}_{K'}\}$ . The terms  $\hat{y}_{k'}$  are the output vectors of the  $k^{\text{th}}$  NN classifier in the ensemble, labelled here as  $\text{NN}_{k'}$ . These terms can be defined as independent functions such that  $\hat{y}_{k'} = f_{k'}(\mathbf{x})$ , with  $f_{k'}$  being the function fitted by  $\text{NN}_{k'}$ .

The terms  $\hat{y}_{k'}$  are two-vectors of the form  $\hat{y}_{k'} = (\epsilon_{k'}, \tilde{\epsilon}_{k'})$ . The component  $\hat{y}_{k'}^1 = \epsilon_{k'}$  representing the *response* of the network to the designated pulse class label  $k'$ , and the term  $\hat{y}_{k'}^2 = \tilde{\epsilon}_{k'}$  representing the response for all the other classes, here described as the *anti-response* of the

---

<sup>25</sup>A logistic regressor is a probabilistic model that classifies the instances with a probability instead of a hard label.

network to the class label  $k'$ . Each output vector  $\hat{y}_{k'}$  can be normalized to obtain a forced probability vector, *e.g.*, by using the sigmoid activation function in the output layers of each MLP classifier, or left with no restriction regarding the sum of its components. The idea behind not normalizing the output vectors is to perform an unrestricted fit to the provided training labels, in practice allowing the norm of the vector to vary and possibly leading to the identification of extreme novelty data that is not accounted for in the training set (generalization). In this way, it may be possible to extract additional information regarding the confidence of the NN classifier in the result by looking at the magnitude of the response vector.

The categorical pulse classes considered for this problem are  $\mathcal{S} = s_k = \{s_1, s_2, s_3, s_4\} = \{S1, S2, SE, Other\}$ , like many other methods explored in this work. And much like the rest of the methods explored, the S2-SE distinction is to draw attention to the separation of smaller S2-like pulses from S1-like pulses. However, since Other pulses are not of interest to the analysis, they do not need to be learned explicitly by the model. For that reason, an ensemble of only three NN classifiers is used, each specialized in S1, S2 and SE identification. For the current model, the number of NN classifiers is  $K' = 3$  for the classes S1, S2 and SE. This model was named *TriNet Classifier*. More details regarding the specific treatment of the results of the TriNet Classifier are provided below.

### 7.5.2 TriNet Classifier

The *TriNet*<sup>26,27</sup> pulse classification tool was first developed to evaluate the viability of using ML methods within LZap. This is a different approach to classification and data analysis than the decision tree-based ensemble methods showcased previously. Being based on neural networks, the TriNet Classifier is a black-box method that provides little information about data partitioning efficiency or parameter importance. Instead, the strength of this method resides in the ability to learn from impure data and still build a strong specialized classifier. The learning strategy is to have the TriNet model train over the results from HADES (Section 6.2.3), and providing that it can generalize naturally, outperform it.

#### Input preprocessing

The TriNet Classifier model was trained and tested using the LZap dataset described in Section 7.2. Using the knowledge obtained with the models developed in Section 7.4, the following list of pulse parameters were selected for classification: pulse area ( $pA$ ), prompt fractions at 50 ns ( $pF50$ ), 100 ns ( $pF100$ ), 200 ns ( $pF200$ ) and 1  $\mu$ s ( $pF1k$ ), top-bottom asymmetry ( $TBA$ ), pulse length at 90% area ( $pL90$ ), pulse height ( $pH$ ), and height to length ratio ( $H2L$ ).

It is important to determine if the input data is in a state that is apt to be handled by the model before training begins. The behaviour and range of the input values should be carefully controlled considering the limitations of the model architecture and hyperparameters, in order

<sup>26</sup><https://gitlab.com/PauloBras/trinetclassifier.git>

<sup>27</sup><https://gitlab.com/PauloBras/mlforpc.git>


to best reveal the nuances of the data. A practical example would be to choose a logarithmic representation of a variable whose features span several orders of magnitude. In other cases it is wise to scale or shift the variables in order to better suit the behaviour of the computational units. Having values and activation functions with scopes that are not compatible can result in loss of information and severe problems like the neuron “death” mentioned before, that results from having neuron activations limited to positive inputs, *e.g.*, *ReLU* activation, handling non-positive values.

There are two main forms of preprocessing input variables: mean-centering and normalization [Agg18].

- Mean-centering the data can help remove some bias effects and can be particularly important if the neuron activations have the largest gradient at the origin, otherwise the training progress becomes slow and the gradients are more susceptible to diverge, since the neurons are more frequently compensating the weights on one side of the domain of the activation functions.
- Variables with domains established *a priori*, or with a value range limited by definitions of other inherent processes are desirable over unbounded variables since they guarantee that novelty data will not exceed the expected range for that input. Bounded variables are trivial to normalize in order to ensure that the inputs are within the effective range of the domain of the activation function of the computational unit. On the other hand, normalisation of unbounded variables may result on range problems and divergences in the calculation of the gradients.

An example of variables with well-defined mathematical limits are the *TBA* RQ and the prompt fraction family of RQs, with *TBA* being restricted to the range  $[-1, 1]$  and the prompt fraction variables, as fractions, being restricted to the range  $[0, 1]$  (see Appendix B for more details). However, even these RQs suffer from computational miscalculations and some exceptions might fall out of the natural range of these parameters. In the case of the prompt fraction RQs, these miscalculations are the result of the time window that integrates the prompt area at the start of the pulse being defined outside of the pulse boundaries, which may result in larger integrated areas within the prompt window that the total area of the pulse. Nevertheless, the *TBA* RQ and the prompt fraction family of RQs are mainly contained within their nominal domain, with only a very small leakage of values outside. The values of pulse parameters that fall outside are often within an order of magnitude of the upper or lower limits, which is not expected to impact the classification of these pulses.

Figure 7.16 displays the distributions of RQs that were conditioned for training the NN models. The *pA* and *pL90* parameters were scaled with the average area and 90% length of the SE pulses, obtained using the results from the HADES classifier, resulting in the transformation  $pA \Leftarrow pA/(80.0 \text{ phd})$  and  $pL90 \Leftarrow pL90/(1000.0 \text{ ns})$ . This unorthodox choice of scaling is more advantageous than simply normalizing to the mean because they provide a natural middle point for the S1-like and S2-like phase-spaces, since the SE pulses are naturally at the lower bound of the S2-like phase-space and are very well defined in the data.



**Figure 7.16:** Stacked distribution of the preprocessed  $pA$ ,  $pL90$ ,  $pH$  and  $H2L$  RQs for each class label used to tune and train the TriNet Classifier model. The pulse class labels were obtained by the HADES classifier. The  $\log_{10} H2L$  parameter is mostly negative, and mean-centering its value distribution would not improve its discrimination ability significantly. Therefore, it was left as it is despite the risk of impacting the performance of the network.

The  $pH$  and  $H2L$  RQs, plus the aforementioned  $pA$  and  $pL90$  RQs were projected to a logarithmic representation to avoid large dynamic ranges, since these parameters are unbounded and span several orders of magnitude. This way, it is guaranteed that the values of these variables on novelty data will not fall too far from the ranges of known data. The  $H2L$  parameter is strongly asymmetric, with some classes appearing only at very small values. For this reason, a logarithmic representation was inevitable. The remaining parameters are already in a form that conforms with the requirements of the model.

The full parameter space with  $D = 9$  pulse parameters used to build the TriNet Classifier model is then defined as:

$$\mathbf{x} = \{ \log_{10}(\mathbf{pA}/80.0), \mathbf{pF50}, \mathbf{pF100}, \mathbf{pF200}, \mathbf{pF1k}, \dots \\ \mathbf{TBA}, \log_{10}(\mathbf{pL90}/1000.0), \log_{10}(\mathbf{pH}), \log_{10}(\mathbf{H2L}) \}$$

### Tuning of relevant hyperparameters

Preparing a NN model for training often requires making informed choices beforehand on many aspects of the model, most notably of its hyperparameters. These tunable parameters can be separated into two types: training hyperparameters that control and affect the training behaviour, and model hyperparameters that describe the properties of the model itself. The most relevant training hyperparameters are often considered to be the learning rate, the batch size and the number of training iterations (*epochs*), while the most important model hyperparameters are the number of hidden layers, the number of hidden units per hidden layer, and unit non-linearity (activation function) [Agg18, HTF09, Ben12].


The tuning of these parameters is usually performed by monitoring the training efficiency and overall results of the model for different combinations of values, ultimately choosing the values that return the best averaged performance. The performance parameters used in this work to tune the hyperparameters of the model are the generalization accuracy, loss over the generalization data and training time. Early stopping was implemented on the training of all models presented in this work. The training of the models was programmed to stop if the generalization accuracy did not improve over 10 consecutive training epochs (early stopping). The model is recorded at the state it was at the beginning of the final 10 consecutive training epochs.

The hyperparameters were tuned using a multi-class classifier model consisting of a single NN in order to have consistent performance measurements and not be dependent on the benchmarking of the ensemble model itself. It is assumed that each of the individual models of the TriNet Classifier ensemble will perform at least on the same level than the single-network multi-class classifier model benchmarked here, for the same architecture and hyperparameters. This assumption is justified by the same paradigm of the other ensemble methods explored in this work, that also applies to NN ensembles: The combination of weaker learners results in an overall stronger learner. By tuning each individual learner in the ensemble in the same way as one would tune a more complex model performing a more complex task, *i.e.*, a multi-class classification, it is ensured that the ensemble model will perform that same complex task as efficiently as the single dedicated learner.

As mentioned before, a NN model with nonlinear squashing functions as activations in at least a single hidden layer is a type of universal function approximator, providing that the number of hidden units is sufficient [HSW89]. This nonlinearity can be achieved by using several different activation functions, like the ones displayed in Figure 7.14. The four activation functions mentioned previously were tested in tandem with some gradient descend optimization algorithms, namely standard stochastic gradient descent with momentum (*SGD*), *RMSprop*, *Adam*, *Nadam* and *Adamax*. The results of this benchmarking of pairs of activation and optimizer are displayed in Figure 7.17.

All models with different combinations of activations and optimizers achieved a high averaged accuracy score, with the only exception being the model with sigmoid activation and standard stochastic gradient descent with momentum (*SGD*). This exception is easy to understand, and is an expected result. As mentioned previously, the sigmoid activation can lead to vanishing





**Figure 7.17:** Validation accuracy for different combinations of activation functions and optimizers for a network with 2 hidden layers, 42 hidden units total, and using a learning rate  $\alpha = 0.001$  and batch size  $n = 256$  on training. The results are averaged over the last half of training, where the accuracy and loss were observed to be stable.

gradients since its derivative is strictly lower than 0.25. This vanishing gradient problem results in stalled learning and suppression of weights, decreasing the overall accuracy of the model and the learning time. A model experiencing vanishing gradients when learning using gradient descent can only reach an optimal state if it performs a large number of small updates, which slows the learning process significantly [Agg18].

By analysing the averaged accuracy score and loss of all the models, except the one using sigmoid activation and SGD optimizer, there is no clear advantage of one combination of activation and optimizer over another. Since accuracy seems not to be an issue for these models, the best combination of activation and optimizer to be used by the TriNet Classifier was chosen based on training performance. Table 7.4 shows the training time, in epochs, for the different models with unique combinations of activations and optimizers. The maximum training time allowed was set to 500 epochs. There are considerable differences on training time for different hyperparameter combinations. The largest effect on learning speed is attributed to the choice of optimizer, with the activations apparently having little effect over training performance. The *RMSprop* optimizer seems to provide the best training performance. All models using the SGD optimizer displayed slow learning, exceeding the maximum training time without stabilizing.

Since the activation choice will not impact the training significantly, the ELU activation was chosen for the TriNet Classifier model. This activation provided the best performance on training time using the *RMSprop* optimizer, on par with the ReLU activation. This result is expected

**Table 7.4:** Averaged training time (epochs) for different combinations of activation functions and optimizers, for the same model used to produce the results displayed in Figure 7.17. Numbers in red indicate the combinations that did not stabilize within 500 epochs, the maximum allowed number of cycles through the entire dataset considered in this benchmarking exercise. The value marked with † corresponds to the combination of activation and optimizer that was chosen for the TriNet Classifier.

	SGD	RMSprop	Adam	Adamax	Nadam
ReLU	500	30	35	100	250
Tanh	500	45	70	90	150
Sig	500	70	45	70	120
ELU	500	30 <sup>†</sup>	80	45	250
<b>average</b>	500	43.75	57.5	76.25	180

since both are linear activation functions that are known to significantly accelerate learning with gradient descent (by having simple derivatives). Furthermore, the ELU activation has the advantage over its rectifier counterpart of keeping the neurons active even when the neuron inputs are negative. Since the dataset contains some composite parameters with negative values, the ReLU activation could lead to some neuron “death”.


The ELU activation is not a squashing function, and therefore it limits the type of transformations that the model can perform to the data. The fact that all functions explored here – nonlinear, squashing and linear – are equally decent activations for this particular classification task is an indicator that the pulse classes in the LZap dataset are most likely linearly separable within the parameters available. This fact is not surprising since the HADES classifier and the RFClassifier also developed in this work can achieve high classification efficiencies despite both being linear models.

The depth of the network and the arrangement of the hidden units has a significant impact in performance. On one hand the increased complexity of a larger model results in a greater potential to process data, but on the other hand it makes training and benchmarking far more challenging. Deeper networks are more likely to suffer from vanishing or exploding gradient problems, and are more dependant of learning rate optimization algorithms. Furthermore, deeper networks often take longer time to converge to a stable state [Agg18].

An interesting property of NN models in general is that a model with a greater depth, *i.e.*, a larger number of hidden layers, does not need as many hidden units to achieve the same evaluation performance as a shallower model [Agg18]. This is due to the higher complexity of the composite variables available at the latter layers, achieved by the successive transformations performed by the upper layers. Furthermore, an increased depth can regularize the training process and lead to less overfitting, since the later layers are limited by the non-trivial transformations performed over the data by the initial layers. The composite variables available at later layers will generally present less biases and a smoother behaviour due to averaging and generalization effects. These two properties have contributed to the recent generalized appeal of deep network models and the preference for deep learning approaches on data mining applications [Agg18].

Figure 7.18 displays the performance results of models with different hidden unit configurations.

The two hyperparameters being tested here are the depth of the model, *i.e.*, the number of hidden layers,  $L$ , and the number of units in each hidden layer,  $P$ . For simplicity, all models considered in this benchmarking test and in the remaining study have the same number of units in all hidden layers. From this study, there is no evidence to suggest that a decreasing, increasing, convex or concave distribution of hidden units per layer is superior to a configuration with constant hidden unit number on all hidden layers in terms of performance for this classification task.



**Figure 7.18:** Averaged validation loss (bottom), validation accuracy (middle) and training time (top), for different combinations of NN depth and hidden unit number per layer. All models considered here were benchmarked using ELU activation, *RMSprop* optimization and learning rate  $\alpha = 0.001$  and batch size  $n = 256$ .

The averaged accuracy and loss scores do not vary significantly for the different architectures considered here. The fluctuations observed are most likely statistical in nature. The training time is also fairly consistent for all NN configurations, with the only exception being a generalized loss of training performance for the models with  $P = 5$  hidden units per layer. This performance loss is most likely related to an insufficient number of computational units required to encode the data. In this model, the number of units in the first hidden layer ( $P = 5$ ) is smaller than the number of input units of the model ( $D = 9$ ), which inadvertently leads to a collapse of the input phase space in only 5 composite variables and a slight loss in discrimination power.

A small trend seems to suggest that training appears to slow down for higher number of units per hidden node, which at some point was expected to happen because of the increasing free parameters (weights) of the fully-connected model. The trend seems to suggest that a layer size of  $P = 17$  is the best compromise between lack of computational power and training performance. However, the results for the training time display large statistical fluctuations and no definite

conclusions can be drawn from these results.

With the limited information available, it was decided that the choice of the depth and layer size should be made in order to have the simplest architecture possible, favouring depth over layer size and with a layer size larger than the dimensionality of the data. For these reasons, an architecture with depth of  $L = 3$  and  $P = 17$  hidden units per layer was selected for the TriNet Classifier.


The learning rate  $\alpha$  and the mini-batch size  $n$  are two very important hyperparameters that manage the learning process. A smaller learning rate will reduce the size of the updates to the weights and leads to slower training progress, while a larger learning rate can result in erratic updates that can easily overshoot the optimal state of the model and lead to a severe decrease in accuracy.

The batch size controls the number of data samples used in the forward pass of the gradient descent step, that in turn controls the frequency of the updates to the model. The loss calculated over a batch of data will have lower variance than one calculated over a single sample, and thus the updates tend to be more stable for larger batches.

Figure 7.19 displays the performance of the NN model, constructed with the previously tuned hyperparameters, trained with different learning rates and batch sizes. The bottom plot of this Figure shows the logarithm of the validation loss in order to be able to compare the disparity between the results for smaller and larger learning rates.

As expected, the learning rate significantly impacts the training time. Smaller values of  $\alpha$  result in smaller updates to the model and a generalized increase in training time. However, since the model is improved slowly and steadily, small values of  $\alpha$  generally result in greater validation accuracy. Conversely, larger values of  $\alpha$  result in bigger updates to the weights of the model, leading to erratic results: All models trained with learning rates  $\alpha > 0.1$  failed to reach a stable state, resulting in the underwhelming accuracies and large losses observed. The value for the learning rate was chosen to be  $\alpha = 0.001$  since it provides a decent compromise between training performance and model stability. This is, interestingly enough, the recommended value for the learning rate for the *RMSprop* optimizer algorithm according to the Keras documentation [C<sup>+</sup>15].

Similarly to the learning rate results, the batch size also affects the final performance of the model. The loss calculated over larger batch sizes tend to be more stable, which in turn results in a more controlled gradient descent. This stability aids the model to reach a better final state, ultimately improving the accuracy and decreasing the loss. The batch size does not appear to influence the overall training time significantly, which is somewhat unexpected. The training time displayed in Figures 7.17, 7.18 and 7.19 is defined as the number of training epochs, *i.e.*, the number of full passes over the entire training dataset, and does not represent real training time. The batch size controls the number of data samples used to perform an update to the model, meaning that each batch size dictates the number of updates per epoch. Larger batch sizes imply longer times for each update but fewer updates in a training epoch, leading to a null net effect on real training time but not on the number of training epochs. Nevertheless, this



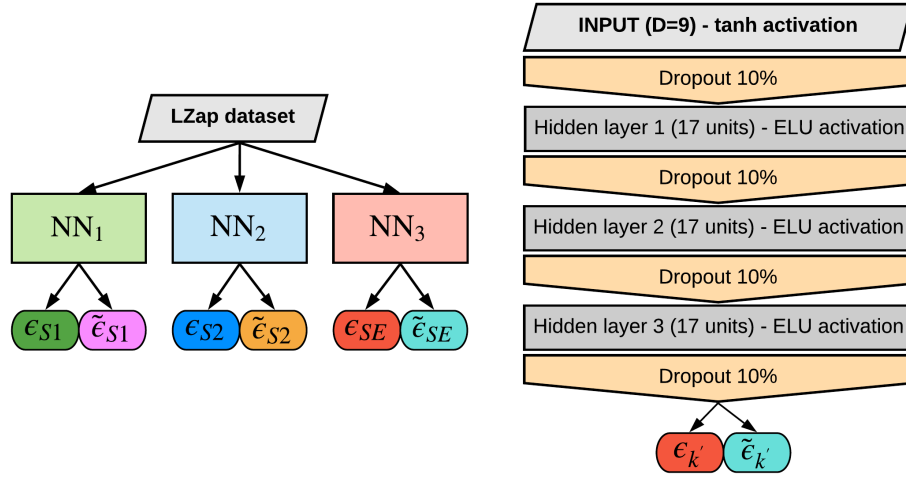
**Figure 7.19:** Averaged validation loss (bottom), validation accuracy (middle) and training time (top), for different combinations of the learning rate and batch size hyperparameters. All models considered here were benchmarked using ELU activation, *RMSprop* optimization,  $L = 3$  hidden layers and  $P = 11$  hidden units per layer.

effect is expected to be small and could easily be shadowed by statistical fluctuations caused by the random initialization of the parameters of the model. The batch size was set at  $n = 128$  samples for the training process of the TriNet Classifier.

### Training and results

The TriNet Classifier architecture is represented in the left side of Figure 7.20. The ensemble model is composed of three independently trained NN classifiers schematically depicted in the right side of Figure 7.20, each specialized in one of the three main pulse classes: S1, S2 and SE pulses. Each NN classifier has  $L = 3$  fully connected hidden layers, with  $P = 17$  hidden units each for a total of  $h = 51$  hidden units. The number of trainable parameters  $\Omega$  in each  $NN_{k'}$  classifiers is 1124. All hidden units possess a ELU activation. The input layer and all hidden layers have an associated dropout layer that will randomly shut down neurons with a 10% probability per neuron in each training batch.

The training of each NN classifier is performed in a One-vs-All scenario using the LZap dataset (see Section 7.2). The training and generalization subsets of data were divided 4:1, with  $8 \times 10^5$  samples for the training dataset and  $2 \times 10^5$  samples for the validation dataset. The training is optimized using the *RMSprop* algorithm with an initial learning rate  $\alpha = 0.001$  and batch size  $n = 128$  samples. The training is monitored using the accuracy score of each model over



**Figure 7.20:** Simplified schematic of the TriNet Classifier model (left). The model is composed of an ensemble of three NN models (right) each trained to separate the three main pulse classes in the data: S1, S2 and SE pulses. The outputs of each NN classifier,  $\hat{y}_{k'} = \{\epsilon_{k'}, \tilde{\epsilon}_{k'}\}$ , represent the response and anti-response of the classifier to its respective class.

the validation dataset, with an early stopping of the training if the validation accuracy does not improve for 10 consecutive epochs. To avoid overfitting, the model is saved at the state it was at the beginning of the 10 final training epochs.

The class labels used to train the TriNet Classifier were obtained using the HADES algorithm (Figure 6.10) since no truth information is available from the simulated data. As mentioned in Section 6.2.3, the overall efficiency of the HADES classifier, estimated by handscanning LZ simulated data, is  $\epsilon_{acc,HADES} = 98.6\%$ , meaning that these labels are not “pure”, *i.e.*, there is some mixture between the data labels and the real classes of the pulses in the data. This impurity of the labels reflects the systematic imperfections of the HADES algorithm and are expected to influence the performance of the TriNet Classifier in some way. However, this choice of using the class labels obtained with HADES is justified. According to the *Classification Without Labels* (CWoLA) paradigm, training a classifier model using impure data labels, especially an ensemble model like the TriNet Classifier, can return a classifier that is optimal at distinguishing between the true labels [MNT17]. The mixture of LZap pulse labels can be represented by the following expressions:

$$p_{S'_1}(\mathbf{x}) = \rho_1 p_{S_1}(\mathbf{x}) + (1 - \rho_1) p_{S_2}(\mathbf{x}) \quad (7.57)$$

$$p_{S'_2}(\mathbf{x}) = \rho_2 p_{S_1}(\mathbf{x}) + (1 - \rho_2) p_{S_2}(\mathbf{x}), \quad (7.58)$$

where  $p_{S'_i}(\mathbf{x})$  represents the probability distribution of the mixed class  $S'_i$ , that is composed of a fraction  $\rho_i$  of  $p_{S_1}(\mathbf{x})$  and  $(1 - \rho_i)$  of  $p_{S_2}(\mathbf{x})$ , the probability distributions of the true class labels  $S_1$  and  $S_2$ , respectively. Here the classes  $S_1$  and  $S_2$  are merely illustrative. Considering that the optimal classifier at distinguishing the true labels  $S_1$  and  $S_2$  is given by the likelihood ratio  $L_{S_1/S_2}(\mathbf{x}) = p_{S_1}(\mathbf{x})/p_{S_2}(\mathbf{x})$  [NP92], then the optimal classifier for distinguishing the mixed

labels  $S'_1$  and  $S'_2$  can be expanded using Equations 7.57 and 7.58 as

$$L_{S'_1/S'_2} = \frac{p_{S'_1}(\mathbf{x})}{p_{S'_2}(\mathbf{x})} = \frac{\rho_1 p_{S_1}(\mathbf{x}) + (1 - \rho_1) p_{S_2}(\mathbf{x})}{\rho_2 p_{S_1}(\mathbf{x}) + (1 - \rho_2) p_{S_2}(\mathbf{x})} = \frac{\rho_1 L_{S_1/S_2} + (1 - \rho_1)}{\rho_2 L_{S_1/S_2} + (1 - \rho_2)}, \quad (7.59)$$

which is simply a rescaling of the likelihood  $L_{S_1/S_2}$  for  $\rho_1 > \rho_2$ . Therefore,  $L_{S_1/S_2}$  and  $L_{S'_1/S'_2}$  represent the same classifier. If  $L_{S'_1/S'_2}$  represents an optimal classifier for separating  $S'_1$  and  $S'_2$ , then it is also optimal at separating  $S_1$  and  $S_2$ , for  $\rho_1 > \rho_2$ .

This paradigm applies to each of the classifiers of the TriNet ensemble, where each classifier is trained with the mixed samples obtained by HADES for each of the respective pulse classes. Despite the labels provided by HADES having a small fraction of mixed classifications, the classifiers trained with such labels should be able to approximate the optimal classifiers at distinguishing their respective true classes from the remaining classes.

To take full advantage of the CWoLA paradigm, the SE split population was manually removed from the remaining S1 populations using basic parametric selections. These selections are not optimized but they will exclude the majority of the SE split pulses from the main S1 populations. Even if the exclusion efficiency for the SE split pulses is not ideal, or more importantly if some real S1 pulses end up being removed from the training data, the CWoLA paradigm indicates that the model should be able learn to generalize these crude selections and optimize the separation of these bad pulses from the S1 populations.

The training labels are converted to categorical vectors  $y_{k'}$  for each of the  $\text{NN}_{k'}$  classifier, *e.g.*, when training the S1 NN classifier ( $\text{NN}_1$ ) all pulses classified as S1 by HADES are identified with the label vector  $y_1 = (1, 0)$ , while all the other pulses not classified as S1 by HADES are identified as  $y_1 = (0, 1)$ . The same labelling scheme is performed for the S2 and SE pulse classes when training the respective NN classifiers.

The output of the TriNet Classifier is given by the ensemble of the individual outputs of each NN classifier,  $\hat{y} = \{\hat{y}_1; \hat{y}_2; \hat{y}_3\}$ , with  $\hat{y}_{k'} = (\epsilon_{k'}, \tilde{\epsilon}_{k'})$  the output of the individual NN classifier assigned to class label  $k'$ . For simplicity, the output  $\hat{y}$  is explicitly written as a matrix of the type

$$\hat{y} = \begin{bmatrix} \epsilon_1 & \tilde{\epsilon}_1 \\ \epsilon_2 & \tilde{\epsilon}_2 \\ \epsilon_3 & \tilde{\epsilon}_3 \end{bmatrix}$$

with the row vectors representing the individual outputs of the  $\text{NN}_{k'}$  classifiers and the column vectors representing the overall response and anti-response of the ensemble, hereby defined by  $\epsilon = (\epsilon_1, \epsilon_2, \epsilon_3)^T$  and  $\tilde{\epsilon} = (\tilde{\epsilon}_1, \tilde{\epsilon}_2, \tilde{\epsilon}_3)^T$ , respectively.

As mentioned previously, the output vectors of each NN classifier are not forced to be probabilistic in nature. However, due to the binary nature of the One-vs-All training, the individual output vectors  $\hat{y}_{k'}$  will behave very similarly to a probabilistic entity. This is a natural response by the classifiers  $\text{NN}_{k'}$  that are essentially forced to see any given pulse as being similar to its designated pulse class  $s_{k'}$  or not similar to it. This results in a trade-off between response

and anti-response that will naturally mimic a probabilistic vector, meaning that  $\epsilon_{k'} + \tilde{\epsilon}_{k'} \approx 1$ . Nonetheless, and as stated before, there is no constrain on the output vector to have its components sum to unity, meaning that this is an emergent property of the model. Since the sum of the elements of each individual output  $\hat{y}_{k'}$  is approximately unity, the sum of all elements of the TriNet Classifier output is approximately 3, *i.e.*,

$$\sum_{k'=1}^{K'} (\epsilon_{k'} + \tilde{\epsilon}_{k'}) \approx K'. \quad (7.60)$$

For each of the pulse classes present in the data, the output  $\hat{y}$  of the TriNet Classifier ensemble is expected to be asymptotically equivalent to

$$f(\mathbf{x}|y=1) \sim \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}; f(\mathbf{x}|y=2) \sim \begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}; f(\mathbf{x}|y=3) \sim \begin{bmatrix} 0 & 1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad (7.61)$$

with the training label  $y = k'$  representing the pulse class  $s_{k'} \in \{S1, S2, SE\}$ .

The responses of the individual classifiers of the TriNet ensemble to an a pulse that resembles neither an S1 nor an S2 nor an SE pulse, *i.e.*, an Other pulse, are expected to be small, *i.e.*, the output of the TriNet Classifier model associated to an Other pulses is expected to be asymptotically equivalent to

$$f(\mathbf{x}|y \neq \{1, 2, 3\}) \sim \begin{bmatrix} 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}. \quad (7.62)$$


However, only in an extreme case would  $\epsilon$  be a non-zero vector. In practice it is expected that at least one of the responses  $\epsilon_{k'}$  would be non-zero. Furthermore, since Other pulses may have different degrees of resemblance to the other pulse classes, one of the components of  $\epsilon$  is expected to be greater than the other two. Nevertheless, the sum of responses should be smaller than that of good pulses, which is an indication that none of the classifiers in the ensemble “recognized” the pulse.

It is convenient to quantify the global strength of the responses  $\epsilon$  and anti-responses  $\tilde{\epsilon}$  in order to evaluate if there is any type of ambiguity on the predictions. The *confidence* on the result,  $\Gamma_\epsilon$ , can be expressed by the sum of the elements of  $\epsilon$ , and an equivalent quantity can be defined for  $\tilde{\epsilon}$ , designated by  $\Gamma_{\tilde{\epsilon}}$ .

$$\Gamma_\epsilon = \sum_{k'=1}^{K'} \epsilon_{k'} \quad (\text{confidence}) \quad (7.63)$$

$$\Gamma_{\tilde{\epsilon}} = \sum_{k'=1}^{K'} \tilde{\epsilon}_{k'} \quad (\text{confidence complement}) \quad (7.64)$$





**Figure 7.21:** Relationship between the confidence  $\Gamma_{\epsilon}$  and complement  $\Gamma_{\tilde{\epsilon}}$  for the TriNet Classifier (top-left), and the responses and anti-responses of the S1-vs-All  $\text{NN}_1$  classifier (top-right), S2-vs-All  $\text{NN}_2$  classifier (bottom-left), and SE-vs-All  $\text{NN}_3$  classifier (bottom-right).

The parameter  $\Gamma_{\tilde{\epsilon}}$  can be seen as the complementary to the confidence score  $\Gamma_{\epsilon}$ , since the result in Equation 7.60 implies that  $\Gamma_{\tilde{\epsilon}} \approx K' - \Gamma_{\epsilon}$ .

Figure 7.21 shows the relations between the responses and anti responses for the full TriNet Classifier (top-left) and for each individual  $\text{NN}_{k'}$  classifier. Here the response and anti-response of the full ensemble are the confidence  $\Gamma_{\epsilon}$  and complement  $\Gamma_{\tilde{\epsilon}}$ . The plots were obtained by processing the full LZap dataset.

The complementarity relationship between the pairs of quantities displayed in all plots is very evident. It is also evident that, despite their strong anti-correlation, the output vectors  $\hat{y}_{k'}$  are not probability vectors since the sum of their components only approximates unity from below. It is clear that using a probabilistic approach on each individual  $\text{NN}_{k'}$  classifier would yield very similar results. The individual outputs of the  $\text{NN}_{k'}$  classifiers are highly polarized, with most pulses returning either  $\epsilon_{k'} = 1$  or  $\tilde{\epsilon}_{k'} = 1$ . However, there are a few pulses that are assigned intermediate values for these responses, meaning that the individual classifier has some degree of confusion on the classification of those pulses.

From the top-left plot of Figure 7.21 it is possible to infer that the values of the confidence parameter are bound from above and below  $\Gamma_\epsilon \in [0, 2]$ . From the examples of expected outputs presented in Equation 7.61, the expected value of the confidence parameter for a pulse that is seen as a valid S1, S2 or SE pulse is  $\Gamma_\epsilon = 1$ , which corresponds to the case where only one classifier has displayed a strong response to the pulse. That is indeed the most common value observed in the top-left plot of Figure 7.21. When  $\Gamma_\epsilon > 1$  (and by extent  $\Gamma_{\tilde{\epsilon}} < 2$ ) the pulse has generated a strong response in two classifiers, indicating that there is some level of confusion regarding the nature of the pulse<sup>28</sup>. Finally, when  $\Gamma_\epsilon < 1$  (and by extent  $\Gamma_{\tilde{\epsilon}} > 2$ ) the pulse has generated a weak response in all  $\text{NN}_{k'}$  classifiers, which is an indication that the pulse is abnormal and, at some threshold of  $\Gamma_\epsilon$ , it should be tagged as an Other pulse. There are no instances of  $\Gamma_\epsilon > 2$ , which indicates that there are no pulses in this dataset that are seen as valid S1, S2 or SE pulses simultaneously.

There is a chance, albeit small, that a pulse will produce a response of the type  $\epsilon \simeq (1/3, 1/3, 1/3)$ , where the individual responses  $\epsilon_{k'}$  are weak but the confidence is equal to that of a unequivocal classification. Such a result would be hard to distinguish using only  $\Gamma_\epsilon$ . Fortunately, it is easy to show that, for all pulses processed, at least one of the components of the response vector is zero, guaranteeing that there is no case in the LZap dataset where all three  $\text{NN}_{k'}$  classifiers produce a non-zero response. This observation is discussed below. This is another emergent property of the TriNet Classifier that was not programmed *a priori*, but is welcome.

In order to make the results of the TriNet Classifier more intuitive, as well as to check if there is any instance where all three  $\text{NN}_{k'}$  classifiers have non-zero response, the values of  $\epsilon$  and  $\tilde{\epsilon}$  can be represented in a 2-dimensional Cartesian plane using a ternary plot. The elements of the response vectors are converted into 2-dimensional Cartesian coordinates by the following transformations:


$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \frac{1}{2} \frac{2\epsilon_2 + \epsilon_3}{\epsilon_1 + \epsilon_2 + \epsilon_3} \\ \frac{\sqrt{3}}{2} \frac{\epsilon_3}{\epsilon_1 + \epsilon_2 + \epsilon_3} \end{bmatrix} \quad (7.65)$$

The ternary plot for the anti-responses is obtained using the same transformations in Equation 7.65 but replacing  $\epsilon$  by  $\tilde{\epsilon}$ . Figure 7.22 displays the ternary plots for the responses (left) and anti-responses (right) of the TriNet Classifier. The bottom-left vertex of the (right)left ternary plot, located at  $(x, y) = (0, 0)$ , is associated with the (anti-)response of the  $\text{NN}_1$  classifier. The bottom-right vertex of the (right)left ternary plot, located at  $(x, y) = (1, 0)$ , is associated with the (anti-)response of the  $\text{NN}_2$  classifier. Finally, the top vertex of the (right)left ternary plot, located at  $(x, y) = (\sqrt{3}/2, 1/2)$ , is associated with the (anti-)response of the  $\text{NN}_3$  classifier.

The pulses that the TriNet Classifier identifies as more SE-like are expected to be closer to the top vertex of the left plot of Figure 7.22, while the pulses identified as more S1-like and S2-like will be closer to the bottom-left and bottom-right vertices, respectively. As an example, a response of the type  $\epsilon^T = (1, 0, 0)$  is associated with an S1-like pulse, and the equivalent point in the ternary plot of the responses would be located at the coordinates  $(x, y) = (0, 0)$ , the

---

<sup>28</sup>Overconfidence often leads to failure.



**Figure 7.22:** Ternary plots of the overall response  $\epsilon$  (left) and anti-response  $\hat{\epsilon}$  (right) of the TriNet Classifier.

bottom-left vertex of the ternary plot. It is easy to extrapolate that when the response vector contains an element which is significantly larger than the remaining two elements,  $\epsilon_{k'} \gg \epsilon_{j \neq k'}$ , the corresponding coordinates approximate the vertex coordinates of the ternary plot associated with the response of  $\text{NN}_{k'}$ .

From the left plot of Figure 7.22, it is clear that most pulses are placed in the vertices of the ternary plot, which indicates that the TriNet Classifier returns an unequivocal classification for most pulses. However, there are some pulses that are not assigned to any of the vertices but fall in between two vertices, on the edges of the ternary plot. The regions between each vertex correspond to pulses with mixed responses on different NN classifiers. This mixture is the result of two NN classifiers having a non-zero response to the same pulse, meaning that this pulse has some resemblance with the two respective pulse classes. As an example, a response vector of the type  $\epsilon^T = (0, 1, 1)$  is associated with a pulse perceived as both S2-like and SE-like by the TriNet Classifier, and has the corresponding transformed coordinates  $(x, y) \simeq (0.75, 0.43)$ , coinciding with the middle point between the top and bottom-right vertices of the ternary plot associated with the response of  $\text{NN}_3$  and  $\text{NN}_2$ , respectively. It is also trivial to extrapolate that any response  $\epsilon$  that contains a term that is significantly smaller than the remaining two,  $\epsilon_{k'} \ll \epsilon_{j \neq k'}$ , will be positioned in the region between the two vertices of the ternary plot associated with the responses of both classifiers  $\text{NN}_{j \neq k'}$ ,  $\forall k', j \in K'$ , *i.e.*, the furthest away from the vertex associated with  $\text{NN}_{k'}$ .

A significant amount of crossover between  $\epsilon_2$  and  $\epsilon_3$  is clearly visible in the left plot of Figure 7.22, which is the result of the close relation between S2 and SE pulse types. The opposite is visible in the region between  $\epsilon_1$  and  $\epsilon_2$ , where significantly less crossover is observed, as expected. The region between  $\epsilon_1$  and  $\epsilon_3$  presents some crossover as well since S1 and SE pulse populations share part of their parameter phase-space.

Since the ternary plot of the responses returns only the outline of the triangle, at least one of the components of  $\epsilon$  is zero for all pulses processed. This means that the only degree of mixing

of classifications is between any two classifiers at maximum, with the third always returning no response. This is an indication that all three different classes can be fully separated by the TriNet Classifier.

The analysis of the response ternary plot alone is not enough to provide a decent classification, since it does not provide an efficient way to identify Other pulses. These pulses are expected to return low responses  $\epsilon_{k'}$  for all the  $\text{NN}_{k'}$  classifiers, since each classifier is expected to view Other pulses as distinct from their designated class  $s_{k'}$ . Furthermore, the expected response associated with an Other pulses, *i.e.*,  $\epsilon_{k'} = 0, \forall k'$ , cannot be represented in a ternary plot due to the denominator in Equation 7.65 going to zero. Such a weak response  $\epsilon$  associated with a Other pulse will be transformed, under Equation 7.65, in the same way as a strong response associated with a valid pulse due to the summation term in the denominators of the Cartesian components. The transformation given by Equation 7.65 effectively returns a normalized response that has no information on the magnitude of  $\epsilon$ , *i.e.*, its confidence  $\Gamma_\epsilon$ . In order to reduce the fraction of false positives, a classification criterion should account for both the confidence  $\Gamma_\epsilon$  and the individual values of  $\epsilon$ , with the former providing discrimination of bad pulses and the latter evaluating the most likely class that the pulse belongs to.

The ternary plot of the anti-responses (right side of Figure 7.22) contains information that is complementary, but not reciprocal, to the information in the ternary plot of the responses (left side of Figure 7.22). The vertices of the plot of the anti-responses are associated with mixed classifications, while unequivocal classifications are associated to the points between vertices. As an example, the typical anti-response vector associated with an S1 pulse is  $\tilde{\epsilon}^T = (0, 1, 1)$ , meaning that both the  $\text{NN}_2$  and  $\text{NN}_3$  classifiers saw the pulse as not being similar to either an S2 or an SE pulse. This result translates to a position in the ternary plot of the anti-responses of  $(x, y) \simeq (0.75, 0.43)$  that is the middle point between the vertices associated with the anti-responses of  $\text{NN}_2$  and  $\text{NN}_3$ . So the expected anti-response  $\tilde{\epsilon}$  for a pulse of class  $s_{k'}$  will be represented at the antipodal of the vertex associated with the anti-response of  $\text{NN}_{k'}$ .

Other pulses will likely produce a low confidence,  $\Gamma_\epsilon < 1$ , which implies that  $\Gamma_{\tilde{\epsilon}} \gtrsim 2$ . Since  $\tilde{\epsilon} \in [0, 1]$ , it is guaranteed that an Other pulse will have every element  $\tilde{\epsilon}_{k'} \neq 0$ , and so Other pulses will appear towards the centroid of the ternary plot and further from the edges. In the most extreme case where  $\tilde{\epsilon} = (1, 1, 1)$ , the Cartesian coordinates in the ternary plot are  $(x, y) \simeq (0.5, 0.29)$ , which coincides with the centroid of the ternary plot.

These cases can be clearly observed in the right plot of Figure 7.22 in three different regions. The first is in the region closer to the vertex associated with  $\tilde{\epsilon}_1$ , where a substantial number of scattered points fill the inward space adjacent to this vertex. These occurrences are caused by pulses seen as both SE-like and S2-like, which are expected since these two pulse topologies are very closely related and form a continuous population. The second and third regions are in the form of the straight lines that connect the centroid with the in-between points that are antipodal to the vertices associated with  $\tilde{\epsilon}_1$  and  $\tilde{\epsilon}_3$ . These cases are related to pulses that slightly resemble S1-like and SE-like pulses, respectively. The S1-like pulses in these regions are mostly

SPE and MPE pulses with coincidence<sup>29</sup> larger than 2 that the TriNet Classifier is somehow distinguishing from the remaining S1 pulses.

In summary, the ternary plots showcase the incredible versatility of having several independent outputs in an ensemble model, and indicate that the TriNet Classifier can be used to classify LZ pulses with high accuracy. It is also apparent that the TriNet Classifier has the potential to identify outlier pulses by using the anti-responses, being able to distinguish between anomalous pulses associated to different classes. The ternary plots in Figure 7.22, although not essential tools for analysis, provide a strong visual queue to all this information.

### Defining a suitable classification metric using $\epsilon$ and $\tilde{\epsilon}$

In practice, one can use both  $\epsilon$  and  $\tilde{\epsilon}$  to evaluate the most likely class, the degree of mixing between the classifications and even to which class is an Other pulse most closely related to. This is a powerful accomplishment that expands the range of analyses that this model can perform in low-level pulse processing. The assignment of hard pulse classes to the processed pulses can be as simple as setting a threshold to each of the responses of the individual classifiers  $NN_{k'}$  and checking the confidence factor  $\Gamma_\epsilon$ . Since the TriNet Classifier seems to favour a particular response to a label  $k'$  with a low level of ambiguity, *i.e.*,  $\exists \epsilon_{k'} \geq \epsilon_{j \neq k'}, \forall k' \in \{1, 2, 3\}$ , selecting all pulses with  $\Gamma_\epsilon \approx 1$  and the response of  $NN_{k'}$  larger than a given threshold will return mostly pulses that are of the class  $s_{k'}$ . This approach would return a hard classification label similar to the HADES algorithm.

If, however, one wished to use a probabilistic approach and expand the classification problem beyond what HADES can provide, the TriNet output can be converted into a probabilistic vector using a simple set of rules like the ones described in Algorithm 7. Despite the output of the TriNet Classifier not being readily associated with a probabilistic result, by using the response vector  $\epsilon$  (as an intermediate conditional probability entity) and the confidence score  $\Gamma_\epsilon$  a probabilistic formulation can be constructed.

The formulation presented in Algorithm 7 is the simplest arrangement of the responses of the ensemble into a meaningful probability vector. It is important to recognise that the probabilities extracted from the NN ensembles need to account for spurious pulses that do not fit into the primary classes used to train the  $NN_{k'}$  classifiers, *e.g.*, the Other pulses in the LZap dataset. Since there is no representation of spurious pulses in training, the probability vector should have  $K' + 1$  number of terms, one for each  $K'$  primary class plus the remainder probability assigned to the remaining classes. In the case of the TriNet classifier, the  $K' = 3$  primary classes will return  $K' + 1 = 4$  probability results. This probabilistic approach is more versatile than a categorical classification, and can be converted to one in post analysis if needed. It is worth pointing out that the model predictive power will be diminished by converting the 6 total terms of the responses and anti-responses into 4 probabilistic values with Algorithm 7, but this information loss is compensated by having a result that is easier to interpret.

---

<sup>29</sup>See the discussion at the end of Section 6.2.1 regarding artificial pulses and the problem with SPE coincidence.

**Algorithm 7:** TriNet probability classification

1. From the TriNet output  $\hat{y} = \{\hat{y}_1; \hat{y}_2; \hat{y}_3\}$ , with  $\hat{y}_{k'} = (\epsilon_{k'}, \tilde{\epsilon}_{k'})$ , take the first and second column vectors :

$$\begin{aligned}\epsilon &= (\epsilon_1, \epsilon_2, \epsilon_3)^T \quad (\text{response}) \\ \tilde{\epsilon} &= (\tilde{\epsilon}_1, \tilde{\epsilon}_2, \tilde{\epsilon}_3)^T \quad (\text{anti-response})\end{aligned}$$

2. Calculate the confidence  $\Gamma_\epsilon = \sum_{k'=1}^{K'} \epsilon_{k'}$ .
3. Construct the probability vector

$$p(k) = (p(k' = 1), p(k' = 2), p(k' = 3), p(k' \neq \{1, 2, 3\}))$$

by assigning probabilities to each class:

- (a) If  $\Gamma_\epsilon \leq 1$  (no classification mixture):

$$\begin{aligned}p(k \in K') &= \epsilon_{k'} \\ p(k = 4) &= 1 - \Gamma_\epsilon \quad (\text{Other})\end{aligned}$$

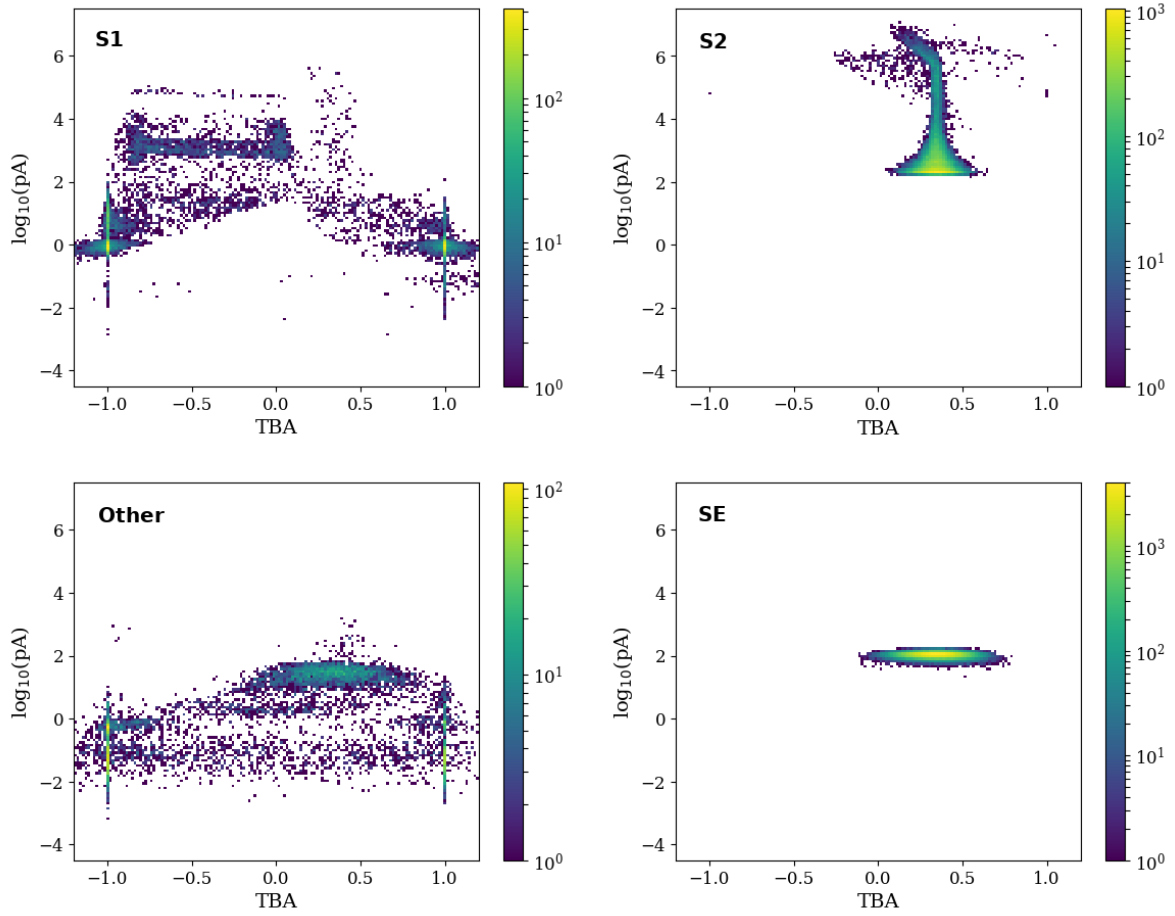
- (b) If  $\Gamma_\epsilon > 1$  (some classification mixture):

$$\begin{aligned}p(k = 4) &= \frac{1}{2}(\Gamma_\epsilon - 1) \quad (\text{Other}) \\ p(k \in K') &= \epsilon_{k'} \frac{1 - p(k = 4)}{\Gamma_\epsilon}\end{aligned}$$

In order to represent the pulse populations for different classes obtained with the TriNet Classifier, the  $\hat{y}$  output, and in turn the probabilistic vector of the pulse classes, are converted to categorical class labels. Figures 7.23 and 7.24 show some examples of classifications obtained by setting simple thresholds to the elements of the probability vector obtained using Algorithm 7.

Most pulses belonging to the  $s_k$  class will produce a value of  $p_k \approx 1^-$ , with just a small fraction going below 95%. But since there is little correlation between probability components, pulses belonging to one of the main pulse classes represented in the ensemble model can have relatively low  $p_k$  values and still be unambiguously associated to their respective class. The selection cuts applied to the probability vector  $p$  in this example are quite loose, with the probability thresholds for each pulse class being as low as 23%. More restricted rules can be devised in order to extract the pulses of each given class. Such rules can be tuned to the specific classification requirements of the processing framework.

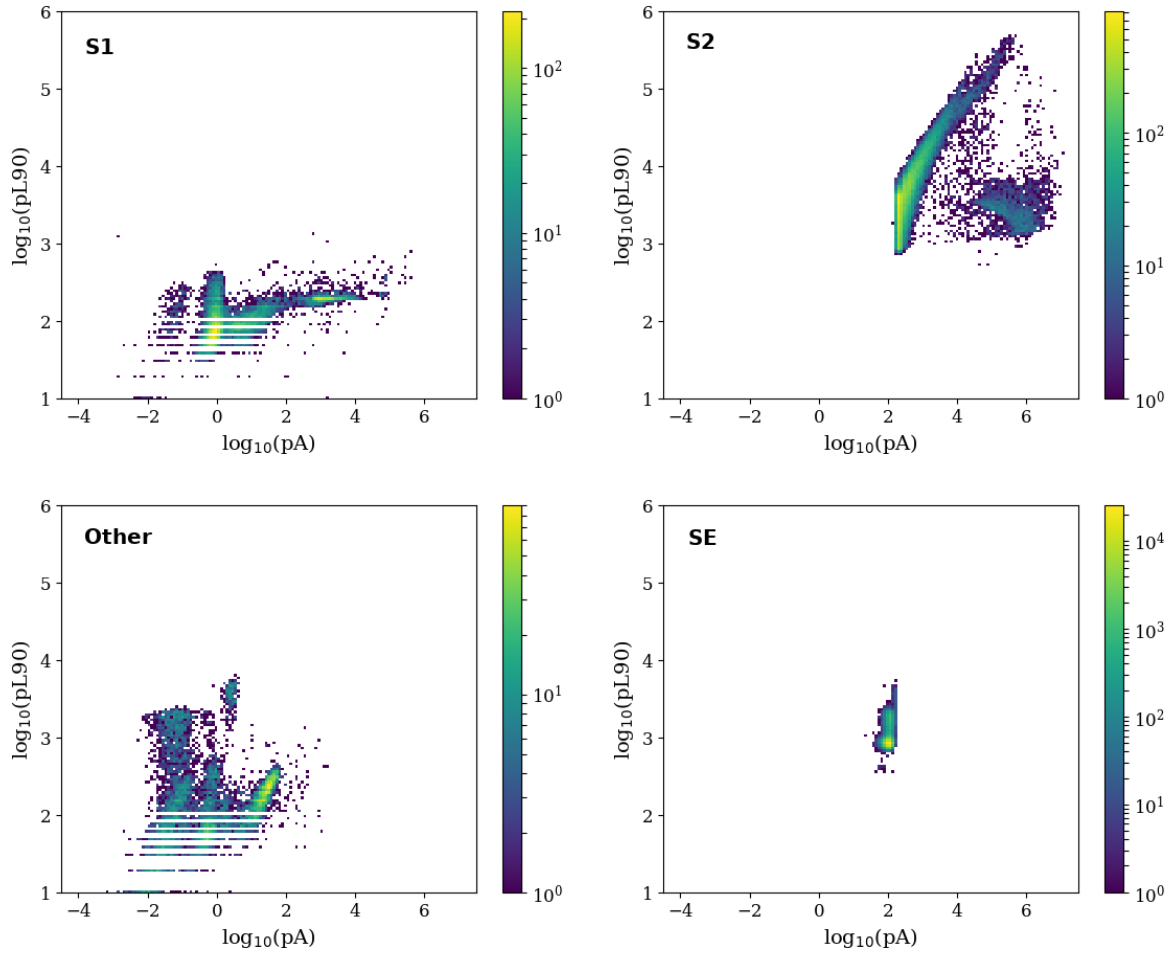
The S2 population seems to be well defined and no clear outlier populations are observed. There is a clear separation between SE and S2 pulses that seems to follow the area trend naturally. The S1 population appears to be well identified, with only two small outlier populations being



**Figure 7.23:** Distribution of the pulse populations in the marginal distribution  $TBA$  vs  $pA$  after being processed by the TriNet Classifier. Top-left: pulses classified as S1 using the selection  $p(k = 1) > 0.23$ . Top right: S2 pulses selected using  $p(k = 2) > 0.25$ . Bottom-right: SE population selected using  $p(k = 3) > 0.25$ . Bottom-left: distribution of the remaining pulses, *i.e.*,  $p(k = 1) < 0.23 \wedge p(k = 2) < 0.25 \wedge p(k = 3) < 0.25$ , classified as Other.

observed: the first one can be found in the top-left plot of Figure 7.23 at  $pA > 1000.0$  phd and  $TBA > 0.2$ , and the second one can be found in the top-left plot of Figure 7.24 for values of  $pA < 0.1$  phd. Some SPE and MPE pulses (with coincidence RQ larger than 2) seem to have been classified as Other, too. This is not entirely bad and seems to indicate that the TriNet Classifier can be used to separate the populations of MPE and SPE pulses with coincidence larger than 1 with some efficiency, despite not having any knowledge of the coincidence RQ. However, these pulses with miscalculated coincidence RQ are associated to the S1 class in training, meaning that they will be counted as misclassifications. This makes the quantification of the real performance of the TriNet Classifier quite complicated.

Tables 7.5 and 7.6 compare the results of the TriNet Classifier with the labels obtained with HADES and with the GMM clustering method developed in Section 7.3.1, respectively. Being able to generalize beyond the training labels is the main accomplishment of the TriNet Classifier,



**Figure 7.24:** Distribution of the pulse populations in the marginal distribution  $pA$  vs  $pL90$  after being processed by the TriNet Classifier. The distributions of S1-like (top-left), S2-like (top-right), SE-like (bottom-right), and Other (bottom-left) pulses were obtained using the same selections as in Figure 7.23.

but it also makes the comparison between the results and the training labels more complicated. After all, the metrics of performance computed from misclassification results are agnostic to the model being closer or further from reality than the training data is. For this reason, the comparison of the results with the GMM results obtained in Section 7.3.1 seems more appropriate, since these two techniques were developed completely independently of one another.

The overall validation accuracy of the TriNet Classifier model (Equation 6.1) calculated using the labels from the GMM clustering is  $acc(\text{GMM}) = 95.56\%$ , lower than the estimated accuracy of the HADES algorithm ( $acc_{\text{HADES}} = 98.6\%$ ). The positive predictive value for the TriNet Classifier is calculated using the same GMM labels as  $PPV(\text{GMM}) = 96.19\%$ , also significantly lower than the result from HADES ( $PPV_{\text{HADES}} = 98.83\%$ ). This decrease in efficiency is due to the TriNet displaying some degeneracy in distinguishing between S2 and SE pulses. Considering only the classification of S1-like and S2-like pulses, *i.e.*, not considering the mixing of S2 and SE pulse labels to be a misclassification, the validation accuracy becomes  $acc^{\text{S1S2}}(\text{GMM}) = 98.56\%$ ,



**Table 7.5:** Confusion matrix of the TriNet Classifier results compared with the HADES results, over the entire LZap dataset.

HADES class	TriNet Predicted class				Total	
	S1(MPE)	S2	SE	Other		
S1(MPE)	8122	0	0	549	8671	4.3%
S2	1	45296	13	1	45311	22.7%
SE	0	446	134449	2	134897	67.4%
Other	3385	0	19	7717	11121	5.6%
<b>Total</b>	11508	45742	134481	8269	200000	

**Table 7.6:** Confusion matrix of the TriNet Classifier results compared with the GMM class labels obtained in Section 7.3.1, over the entire LZap dataset.

GMM class	TriNet Predicted class				Total	
	S1(MPE)	S2	SE	Other		
S1(MPE)	10257	0	0	1576	11833	5.9%
S2	0	45717	5992	0	51709	25.9%
SE	0	7	128450	0	128457	64.2%
Other	1251	18	39	6693	8001	4.0%
<b>Total</b>	11508	45742	134481	8269	200000	

and the positive predictive value becomes  $PPV^{S1S2}(\text{GMM}) = 99.32\%$ . The performance of the TriNet Classifier on these conditions is on par with the HADES algorithm. A comparison between the models developed in this work can be found in Section 7.6.

Similarly to the results of the RFClassifier, the main failure mode is the classification of Other pulses as S1 pulses, but in this case the pulses are most likely small baselines or SPE and MPE pulses that were mislabelled in the GMM clustering analysis. The TriNet Classifier has a higher efficiency at separating SE splits from S1 pulses, and vice-versa, when compared to HADES.

One of the strongest results obtained with the TriNet Classifier is the natural identification of spurious pulses without any preconceived notion of what a bad pulse should look like. The bottom-left plots of Figures 7.23 and 7.24 contain some known populations of bad pulses that were not represented in the training data with dedicated labels. The SE split population that the HADES algorithm could not separate effectively seems to be isolated with high efficiency, despite the crude removal of these pulses from the S1 classification, as mentioned above. This is a clear indication that the TriNet Classifier was able to generalise beyond HADES, which is an indication that the CWoLA paradigm is valid in this context.

It is worth mentioning that since a linear activation function was used in this ensemble model, *i.e.*, the ELU activation, the input space as well as the subsequent combined parameter phase-spaces can only be partitioned into convex regions. This effect might lead to some loss in generalizability. A different approach to the design of the classifier model presented here might lead to better classification results in the future. But despite this limitation, the current version

of the TriNet Classifier demonstrated here it is still able to provide excellent classification results.

### 7.5.2.1 Alternative methods of training and model implementation

There are many different ways of implementing a NN classifier, and more so an ensemble of NN classifiers like the TriNet Classifier model. The choices presented in this work were made with the goal of exploring this technique in its most unrestricted form and should not be taken as being the most optimized, in the sense that not all possibilities were explored. In the light of the results obtained in this work, some changes to the implementation and training of the TriNet Classifier are being considered for a future study.

The first and most simple change would be to explore a probabilistic implementation of each NN classifier model. The results presented above seem to indicate that the model is probabilistic in nature and, due to the binary nature of the One-vs-All training of the individual classifiers, it tends to return quasi-probabilistic entities as outputs. An exploration of this change would in principle provide similar results to those obtained in this work, with the possibility of providing more stable outputs and consistent classifications.

Another important change that should be explored is the implementation of the classification models using nonlinear activation functions that would allow a more complex transformation of the input space and most likely a better generalization of the results. The choice of using the ELU activation in this work was based primarily on the training speed it provided, a strong incentive given the limited time that the author had to develop, implement and test this model. A future study on this aspect of the model would most certainly result in a better classification tool.

An alternative method of implementing the TriNet classifier using outlier detection techniques in tandem with the NN classifiers was explored briefly along the development of this algorithm. The goal was to use outlier detection methods, namely *isolation forests* [LTZ08], to skim the data first, and then let the TriNet Classifier generalize the pulse class populations via CWoLA. To avoid having less represented populations tagged as outliers, the pulses were weighed using a simple logistic regression model that used the selection from HADES. This alternative implementation of the model provided similar classification results when compared to the one explored in this work but the ability to identify outliers was subpar. The main reason for the lack of classification power for outliers is most likely related to the forced generalization that was demanded of the model: by using the isolation forest results to skim the data, a large fraction of pulses of the  $s_{k'}$  class were labelled as not belonging to  $k'$ , which pushed the individual  $NN_{k'}$  classifiers to classify good and bad pulses alike more strongly towards a class  $k'$ . This is an interesting result that, despite not providing a strong coverage to the classification of spurious pulses, had a stronger classification score for the learned classes  $s_{k'}$ . However, this is not the desired behaviour for a classifier that is expected to encounter a large percentage of spurious pulses. The idea was abandoned but it might be explored in the future, where a more careful tuning of the outlier detection method may solve the initial issues observed.

Since this work was completed, the quality of simulated data improved significantly. A new study of this technique with the most recent LZap data, expected to better approximate real LZ data, would be enlightening. Furthermore, the LZap processing algorithms that precede the classification were also extensively modified, with the most recent versions having improved the quality of the RQs significantly. These changes would undoubtedly improve the classification accuracy of the TriNet Classifier model.

## 7.6 Comparison Between Techniques

All pulses classification tools explored in this work proved to be able to handle, learn and predict with high efficiency the many pulse topologies present in the LZap data. They were chosen based on their ability to handle multi-class problems as well as provide insightful information about the data itself [IS17].

Table 7.7 compares the 3 models developed in this work: HADES, the RFClassifier and the TriNet Classifier. Based on these results, the RFClassifier seems to lead in terms of classification performance.

**Table 7.7:** Comparison of the metrics of performance for all the classifier models developed in this work. The results for S2-SE mixing are obtained by considering that a SE seen as an S2 and vice-versa is not a misclassification. The results from HADES were verified by handscan. The metrics for the RFClassifier and TriNet Classifier were obtained using the results of the GMM clustering as a proxy to the “true” class.

Metric (Eq.)	Without S2-SE mixing			With S2-SE mixing		
	HADES	RF+GMM Classifier	TriNet Classifier	HADES	RF+GMM Classifier	TriNet Classifier
Accuracy (6.1)	0.9858	0.9937	0.9556	0.9873	0.9967	0.9856
PPV (6.2)	0.9883	0.9948	0.9619	0.9898	0.9979	0.9932
NPV (6.3)	0.9175	0.9681	0.8094	0.9175	0.9681	0.8094

Random forests and boosted decision trees are a great solution for the classification efforts in LZap. These methods produce models that can go from pure blackbox to fully transparent depending on the level of control required for the analysis. They have a fast runtime and provide high classification efficiency even beyond HADES.

In the development stages of the RFClassifier several important tools were used to improve the classification of LZ pulse data, namely the development of clustering techniques and feature importance ranking tools that provide valuable information about the data itself. These two techniques will undoubtedly be used in real LZ data once it is available, as they provide the necessary information to go beyond simulations as the only source of training data for the development of the classification algorithms. The results from the RFClassifier and the aforementioned peripheral techniques were used to improve the HADES algorithm and helped to identify systemic issues with the simulated data that could have potentially impacted its performance.

The only downside to the RFClassifier is its limited generalization ability. The rigid nature of decision tree methods leads to limited flexibility when the input data is changed in some non-trivial ways, something that is expected to happen in a real experiment. Despite the combination of clustering analysis and tree ensemble methods being incredibly powerful for processing unchanging data, due to the time-consuming clustering analysis required to tune this method, updating the RFClassifier has proven to be inefficient. Nevertheless, when the data is stable this method has tremendous potential as a dedicated classification tool.

The TriNet Classifier model was not as developed nor explored as much as the remaining ML methods due to its greater complexity and also due to some time constraints related to the rest of this work. Despite the limited investment in this technique it provided excellent results and achieved all goals that were set from the start: It performs on par with HADES as a standalone pulse classification tool, and as an auxiliary tool it can be used to generalize the classification results and possibly identify systemic failures in the official classification algorithms implemented in LZap. The TriNet Classifier model was able to learn and generalize beyond HADES, which is the most important achievement of the model, indicating that the CWoLA paradigm can be used as an asset to NN-based classification models. The results provided by this model were obtained with minimal preprocessing of pulses or pulse populations, with the only exception being the non-optimized rejection of the SE split population.

The ability to provide probabilistic results is a sought out feature of the classification tools for LZap. Since the original classification algorithm, COMPACT, the LZap data structure is prepared to handle probabilistic information as classification, and the TriNet Classifier model is able to deliver that information. This also provides the TriNet Classifier with a higher classification flexibility than the remaining methods studied here.

Of all the ML tools available for pulse classification, random forests may be the most accessible model to be implemented in LZap, providing a viable alternative to heuristics algorithms such as HADES. The unsupervised learning approach taken with this implementation of the RFClassifier can be a viable alternative to the usage of simulations to train the classification algorithms, once LZ starts collecting data.

Despite the slight advantages of the RFClassifier, all models have contributed to the improvement of the classification of the LZ data processing framework, an effort that will impact the scientific performance of the experiment once it begins its scientific runs.

## 7.7 Final Remarks

The work presented here on ML classification tools was performed in parallel with the development and routine updates of the HADES classifier algorithm. As the default classification module in LZap, it is the responsibility of the author to ensure that it is continuously up to date. The continuous development of LZap modules during the mock data challenges, as well as the preparations for processing the initial data of LZ require constant maintenance of HADES. Some of the updates on the HADES algorithm were also prompted by the results obtained during the

development of the tools presented in this chapter.

The timing of this work was challenging, having coincided with the early to mid stages of the development of the LZ analysis tools, a phase that required a high level of engaging and time dedication. Most of the work presented in this Chapter and Chapter 6 focuses on the end result of an effort that spanned some 5 years, since 2015.

The tools developed in these classification studies are the groundwork for pulse classification in LZ, and will undoubtedly be an important factor on the performance of this dark matter experiment, not only in its dark matter search but also in the many other physics topics that will be studied.



# Chapter 8

## Conclusions

The PhD activities were performed at the *Laboratório de Instrumentação e Física Experimental de Partículas* (LIP) at the Physics Department of the University of Coimbra. All PhD activities are inserted in the LUX-ZEPLIN (LZ) collaboration of which the author is an active working member. The work plan of the PhD programme was completed successfully and all milestones have been reached. The work produced in this PhD is an integral part of the physics output of the LZ collaboration. The PhD roadmap focused on the development of analysis and data processing tools, namely the pulse classification module of the official LZ data processing framework, as well as the study of the sensitivity of LZ to the neutrinoless double beta decay of the isotope  $^{136}\text{Xe}$ , the second main physics goal of the experiment.

One of the goals of the PhD programme was the study of the sensitivity of the LZ detector to the neutrinoless double beta decay of the isotope Xe-136. The author is the lead researcher for this analysis since March 2017 and is a working member of the High-Energy ER working group that supervises this study. This is the second physics goal of the LZ experiment, that takes advantage of the fact that it naturally contains around 623 kg of  $^{136}\text{Xe}$ , comparable with dedicated  $0\nu\beta\beta$  decay experiments [A<sup>+</sup>12b, G<sup>+</sup>16]. A complete analysis of the main backgrounds for this search was performed, including the full description of the radioactive sources, their activities and how they interact within the LZ detector. An extensive analysis campaign, lasting over a year, was dedicated to characterizing the background caused by gamma-rays coming from the laboratory rock, originated from the decays of naturally occurring  $^{238}\text{U}$  and  $^{232}\text{Th}$  daughters. The extensive shielding of the LZ detector poses a challenge to the simulation of this background source. Using the results obtained by the LZ collaboration from measurements *in situ* [A<sup>+</sup>20b] and dedicated simulation techniques, this background was fully characterized and included in the analysis. The sensitivity projections for the LZ experiment to the half-life of the  $0\nu\beta\beta$  decay

of  $^{136}\text{Xe}$  are of  $1.06 \times 10^{26}$  years. Despite being an experiment tailored for dark matter searches, LZ can be a competitive detector for this decay, having a sensitivity comparable to that of the current best dedicated experiments [G<sup>+</sup>16]. After WIMP search, LZ could run a dedicated  $0\nu\beta\beta$  search with a 90% enriched mass of  $^{136}\text{Xe}$ . This scenario was also considered in this work and the corresponding sensitivity to the half-life of the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  was estimated to be  $1.06 \times 10^{27}$  years. This work culminated in the publication of a scientific paper by the collaboration, of which the author is a corresponding author [A<sup>+</sup>20c].

The author actively participated on the LZ analysis working group and on the LZ offline computing working group, where most of the work on data processing algorithms is managed. These working groups are responsible for the development of tools and algorithms to be used in the official LZ data processing framework (LZap). Several algorithms were developed and some implemented in the data analysis framework of the LZ experiment, aimed at the classification of pulse structures in LZ data that are then used to reconstruct the information of the interaction that originated the recorded event. It is therefore one of the most important steps in the data processing chain and most of the succeeding modules use its results in order to correctly perform their function. Two pulse classification algorithms, named COMPACT and HADES, were implemented in LZap and were the main classification modules in the framework since the start of the PhD task.

The *Comprehensive Pulse Analysis and Classification Tool*<sup>1</sup> (COMPACT) algorithm proved to have high classification efficiency, estimated at 100% at identifying S1 signals and 99.97% at identifying S2 signals on LZ data. Extensive testing of the classification module was performed by the author and by the LZ collaboration during the first and second Mock Data Challenges (MDC1 and MDC2). The algorithm was, unfortunately, unable to handle the complexity of the simulated data once it approached that of real data, prompting the development of a new classification algorithm (HADES). The reasons for the under-performance of COMPACT were studied in detail and are well understood. The lessons learned from COMPACT were useful in the development of new, more robust algorithms. An upgrade of the COMPACT algorithm is planned in the future as an alternative method for classification that may complement the classification module used in the data processing. Several advanced statistical techniques like logistic regression or even Machine Learning tools can be devised in order to solve the weaknesses displayed by COMPACT.

The *Heuristics Algorithm for Discrimination of Event Substructures*<sup>2</sup> (HADES) algorithm is a heuristic decision tree that performs the classification at the pulse level using several pulse parameters as discriminants and by performing a sequence of decisions using several selection criteria tuned to provide the best classification efficiency. HADES is a simple, robust and transparent algorithm designed to be modified on-the-fly and serve as the benchmark for more advanced pulse classification modules. Despite its simplicity, HADES performs remarkably well.

---

<sup>1</sup><https://gitlab.com/PauloBras/COMPACT.git>

<sup>2</sup>Continuously integrated as part of the *PhotonDetection* package <https://gitlab.com/luxzeplin/lzap/PhotonDetection.git>



The overall classification accuracy of HADES, estimated directly by performing handscans of the data, is estimated to be  $\sim 98.6\%$ . HADES will be maintained and upgraded as the default classifier algorithm for LZap. Machine Learning methods may help to improve HADES or at least aid in the tuning of this algorithm by providing new insights on LZ data.

The author is still the main developer of both these classification algorithms for the LZ experiment and an active member of the analysis and offline computing working groups to this day. He is also expected to manage some of the algorithms and processing software modules of the framework until the LZ detector starts taking data. Several presentations on official LZ working group meetings were also given in order to discuss the work developed during the PhD.

Several Machine Learning methods were studied for pulse classification in LZap<sup>3</sup>. Some methods were identified as having high potential for producing pulse classification algorithms that would improve the heuristic methods used in LZap. Of these methods, artificial neural networks (ANN) and random decision forests were chosen to process LZap data. Due to the lack of truth information in the simulated data (by design) the classification algorithms explored in this work resorted to unique ways of generating training labels, either by adopting the results of HADES or by performing clustering analysis using Gaussian mixture models (GMM). The implementations were provided by Keras [C<sup>+</sup>15] and scikit-learn [P<sup>+</sup>11].

Implementations of random forests (RFs) and boosted decision trees (BDTs) were tested in this work. While BDTs are widely used in high energy physics for their incredible classification power and versatility, RFs performed better in the context of LZ data and the work focused solely on these models. This resulted in the development of the *RFClassifier*<sup>4</sup> classification tool. Due to the nature of these methods, the results of HADES cannot be used as training labels, which was verified with MDC3 data. A new training strategy was implemented using clustering analysis that featured Gaussian mixture models<sup>5</sup> (GMMs), which successfully partitioned the data into several components that were subsequently labelled via handscans. The *RFClassifier* achieved a classification accuracy of 99.37%, estimated with both handscans and using the labels generated by the GMM algorithm. This is the highest accuracy achieved by any of the methods explored in this work. Furthermore, these methods have the ability of identifying the pulse parameters with higher discrimination power, a technique called “feature importance ranking”. This ranking provided precious information for other classification algorithms, having ultimately led to major improvements to the HADES algorithm.

The work on ANNs began in 2016 with the initial proof-of-concept algorithms produced and tested on LUX data and LZ simulated data from MDC1. The classification accuracy for MDC1 data was calculated to be over 99.5%. The success of this method prompted additional testing and during the year of 2018 and 2019 a dedicated classification algorithm was developed, called *TriNet*<sup>6</sup>, that could provide high classification efficiency for the highly complex simulated data of

<sup>3</sup><https://gitlab.com/PauloBras/mlforpc.git>

<sup>4</sup><https://gitlab.com/PauloBras/rfclassifier.git>

<sup>5</sup><https://gitlab.com/PauloBras/gmmclustering.git>

<sup>6</sup><https://gitlab.com/PauloBras/trinetclassifier.git>

MDC3. The algorithm is an ensemble of multilayer perceptrons (MLPs), each being specialized at identifying one specific pulse class. In this work, three MLPs were trained to identify S1, S2 and SE pulses using the results obtained with HADES as training labels. It was verified that the TriNet classifier successfully generalized over the imperfection of HADES, a result conjectured in the literature [MNT17]. The TriNet classifier achieved an overall classification accuracy of 95.56% compared to the results from the GMM clustering analysis. The most important result from the TriNet classifier is the ability to identify spurious pulses that could not be easily separated using conventional methods such as HADES. The insights gained from this algorithm led to the isolation of pathological pulse populations in LZap data, which were subsequently used to identify the best methods to separate them from the most relevant pulse populations. This knowledge is crucial for improving other classification algorithms like HADES and COMPACT.

The creation of the LZ Machine Learning working group, established in July 2019 and of which the author is a working member, aims at the implementation and testing of several advanced data analysis algorithms in LZap, including the TriNet and the RFClassifier models. The work performed on machine learning methods for pulse classification in dark matter direct detection experiments will result in the publication of a paper currently in preparation.

The author participated actively in the construction and research of the LZ detector, performing shifting duties for a total of 21 days, 13 of which were spent on the Davis Lab in the 4850 ft level of the Sanford Underground Research Facility (Lead, South Dakota, USA), where the LZ experiment is installed. On the Davis lab, the author participated in the maintenance and installation of the liquid nitrogen-cooled thermosyphons that control the detector temperature and minimize thermal gradients, the wiring of several peripheral sensors on the liquid xenon tower, xenon circulation control cabinets, and system recovery vessel (SRV) to the process loop controller (PLC) units and the testing of the communication between the underground control monitor and several sensors and actuators. The wiring, cabling and labelling of the SRV control panel was done exclusively by the author. The author received training in cryogenic safety, chemical hygiene and safety, electrical safety, pressure safety, high-voltage safety, compressed gas cylinder handling, oxygen deficiency and self-contained self-rescuer usage, as well as site-specific training as a requirement for working underground in the Davis Lab environment.

This doctoral program is aligned with the state-of-the-art of experimental particle and astroparticle physics, and provided the opportunity to develop valuable skills in software development and benchmarking, mostly centred in advanced data processing techniques and statistical tools like Machine Learning and Deep Learning.

# Bibliography

- [A<sup>+</sup>65] C. V. Achar et al. Detection of muons produced by cosmic ray neutrinos deep underground. *Phys. Lett.*, 18:196–199, 1965. doi:10.1016/0031-9163(65)90712-2.
- [A<sup>+</sup>93] P. Anselmann et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Nucl. Phys. Proc. Suppl.*, 31:117–124, 1993.
- [A<sup>+</sup>01] Q.R. Ahmad et al. Measurement of the rate of  $\nu_e + d \rightarrow p + p + e^-$  interactions produced by  $^8\text{B}$  solar neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87:071301, 2001. arXiv:nuc1-ex/0106015, doi:10.1103/PhysRevLett.87.071301.
- [A<sup>+</sup>02] Q.R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002. arXiv:nuc1-ex/0204008, doi:10.1103/PhysRevLett.89.011301.
- [A<sup>+</sup>03] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003. doi:{10.1016/S0168-9002(03)01368-8}.
- [A<sup>+</sup>05] G.J. Alner et al. First limits on nuclear recoil events from the {ZEPLIN} I galactic dark matter detector. *Astroparticle Physics*, 23(5):444 – 462, 2005. doi:{10.1016/j.astropartphys.2005.02.004}.
- [A<sup>+</sup>07a] D. Yu Akimov et al. The ZEPLIN-III dark matter detector: instrument design, manufacture and commissioning. *Astropart. Phys.*, 27:46–60, 2007. doi:{10.1016/j.astropartphys.2006.09.005}.
- [A<sup>+</sup>07b] G. J. Alner et al. First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two phase xenon detector for dark matter detection. *Astropart. Phys.*, 28:287–302, 2007. doi:{10.1016/j.astropartphys.2007.06.002}.

- [A<sup>+</sup>08] J. Angle et al. First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory. *Phys. Rev. Lett.*, 100, 2008. doi:{10.1103/PhysRevLett.100.021303}.
- [A<sup>+</sup>10] J. Argyriades et al. Measurement of the two neutrino double beta decay half-life of Zr-96 with the NEMO-3 detector. *Nucl. Phys. A*, 847:168–179, 2010. arXiv:0906.2694, doi:10.1016/j.nuclphysa.2010.07.009.
- [A<sup>+</sup>11a] N. Ackerman et al. Observation of Two-Neutrino Double-Beta Decay in <sup>136</sup>Xe with EXO-200. *Phys. Rev. Lett.*, 107:212501, 2011. arXiv:1108.4193, doi:10.1103/PhysRevLett.107.212501.
- [A<sup>+</sup>11b] E. Aprile et al. Design and Performance of the XENON10 Dark Matter Experiment. *Astropart. Phys.*, 34:679–698, 2011. doi:{10.1016/j.astropartphys.2011.01.006}.
- [A<sup>+</sup>12a] D.S. Akerib et al. LUXSim: A Component-Centric Approach to Low-Background Simulations. *Nucl. Instrum. Meth. A*, 675:63–77, 2012. arXiv:1111.2074, doi:10.1016/j.nima.2012.02.010.
- [A<sup>+</sup>12b] M. Auger et al. Search for Neutrinoless Double-Beta Decay in Xe136 with EXO-200. *Physical Review Letters*, 109(3), Jul 2012. doi:{10.1103/PhysRevLett.109.032505}.
- [A<sup>+</sup>13a] Georges Aad et al. Search for dark matter candidates and large extra dimensions in events with a photon and missing transverse momentum in *pp* collision data at  $\sqrt{s} = 7$  TeV with the ATLAS detector. *Phys. Rev. Lett.*, 110(1):011802, 2013. arXiv:1209.4625, doi:{10.1103/PhysRevLett.110.011802}.
- [A<sup>+</sup>13b] R. Agnese et al. Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *Phys. Rev. Lett.*, 111(25):251301, 2013. arXiv:1304.4279, doi:{10.1103/PhysRevLett.111.251301}.
- [A<sup>+</sup>13c] D.S. Akerib et al. The Large Underground Xenon (LUX) Experiment. *Nucl. Instrum. Meth. A*, 704:111–126, 2013. arXiv:1211.3788, doi:10.1016/j.nima.2012.11.135.
- [A<sup>+</sup>14a] D.S. Akerib et al. First results from the LUX dark matter experiment at the Sanford Underground Research Facility. *Phys. Rev. Lett.*, 112:091303, 2014. arXiv:1310.8214, doi:10.1103/PhysRevLett.112.091303.
- [A<sup>+</sup>14b] J. B. Albert et al. Improved measurement of the  $2\nu\beta\beta$  half-life of <sup>136</sup>Xe with the EXO-200 detector. *Phys. Rev. C*, 89:015502, 2014. doi:{10.1103/PhysRevC.89.015502}.
- [A<sup>+</sup>14c] E. Aprile et al. Conceptual design and simulation of a water Cherenkov muon veto for the XENON1T experiment. *JINST*, 9:P11006, 2014. doi:10.1088/1748-0221/9/11/P11006.

- [A<sup>+</sup>15a] Alexander Aab et al. The Pierre Auger Cosmic Ray Observatory. *Nucl. Instrum. Meth.*, A798:172–213, 2015. doi:{10.1016/j.nima.2015.06.058}.
- [A<sup>+</sup>15b] Georges Aad et al. Search for  $s$ -channel single top-quark production in proton–proton collisions at  $\sqrt{s} = 8$  TeV with the ATLAS detector. *Phys. Lett.*, B740:118–136, 2015. doi:{10.1016/j.physletb.2014.11.042}.
- [A<sup>+</sup>15c] M. Agostini et al. Results on  $\beta\beta$  decay with emission of two neutrinos or Majorons in  $^{76}\text{Ge}$  from GERDA Phase I. *Eur. Phys. J. C*, 75(9):416, 2015. arXiv:1501.02345, doi:10.1140/epjc/s10052-015-3627-y.
- [A<sup>+</sup>15d] D.S. Akerib et al. Radiogenic and Muon-Induced Backgrounds in the LUX Dark Matter Detector. *Astropart. Phys.*, 62:33–46, 2015. arXiv:1403.1299, doi:10.1016/j.astropartphys.2014.07.009.
- [A<sup>+</sup>15e] J. B. Albert et al. Measurements of the ion fraction and mobility of  $\alpha$ - and  $\beta$ -decay products in liquid xenon using the EXO-200 detector. *Phys. Rev. C*, 92:045504, 2015. doi:{10.1103/PhysRevC.92.045504}.
- [A<sup>+</sup>16a] J. Aalbers et al. DARWIN: towards the ultimate dark matter detector. *JCAP*, 11:017, 2016. arXiv:1606.07001, doi:10.1088/1475-7516/2016/11/017.
- [A<sup>+</sup>16b] D.S. Akerib et al. Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data. *Phys. Rev. Lett.*, 116(16):161301, 2016. arXiv:1512.03506, doi:10.1103/PhysRevLett.116.161301.
- [A<sup>+</sup>16c] D.S. Akerib et al. Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics. 8 2016. arXiv:1608.05381.
- [A<sup>+</sup>16d] D.S. Akerib et al. Tritium calibration of the LUX dark matter experiment. *Phys. Rev. D*, 93(7):072009, 2016. arXiv:1512.03133, doi:10.1103/PhysRevD.93.072009.
- [A<sup>+</sup>16e] J.B. Albert et al. Cosmogenic backgrounds to  $0\nu\beta\beta$  in EXO-200. *Journal of Cosmology and Astroparticle Physics*, 2016(04):029, 2016. doi:{10.1088/1475-7516/2016/04/029}.
- [A<sup>+</sup>16f] J.B. Albert et al. Measurement of neutron capture on  $^{136}\text{Xe}$ . *Phys. Rev. C*, 94(3):034617, 2016. arXiv:1605.05794, doi:10.1103/PhysRevC.94.034617.
- [A<sup>+</sup>16g] E. Armengaud et al. Constraints on low-mass WIMPs from the EDELWEISS-III dark matter search. *JCAP*, 5:019, May 2016. doi:10.1088/1475-7516/2016/05/019.
- [A<sup>+</sup>16h] R. Arnold et al. Measurement of the  $2\nu\beta\beta$  decay half-life of  $^{150}\text{Nd}$  and a search for  $0\nu\beta\beta$  decay processes with the full exposure from the NEMO-3 detector. *Phys. Rev. D*, 94(7):072003, 2016. arXiv:1606.08494, doi:10.1103/PhysRevD.94.072003.

- [A<sup>+</sup>16i] R. Arnold et al. Measurement of the double-beta decay half-life and search for the neutrinoless double-beta decay of  $^{48}\text{Ca}$  with the NEMO-3 detector. *Phys. Rev. D*, 93(11):112008, 2016. [arXiv:1604.01710](#), [doi:10.1103/PhysRevD.93.112008](#).
- [A<sup>+</sup>17a] N. Abgrall et al. Muon Flux Measurements at the Davis Campus of the Sanford Underground Research Facility with the Majorana Demonstrator Veto System. *Astropart. Phys.*, 93:70–75, 2017. [arXiv:1602.07742](#), [doi:10.1016/j.astropartphys.2017.01.013](#).
- [A<sup>+</sup>17b] R. Agnese et al. Projected Sensitivity of the SuperCDMS SNOLAB experiment. *Phys. Rev. D*, 95(8):082002, 2017. [arXiv:1610.00006](#), [doi:10.1103/PhysRevD.95.082002](#).
- [A<sup>+</sup>17c] D. S. Akerib et al. Signal yields, energy resolution, and recombination fluctuations in liquid xenon. *Phys. Rev. D*, 95:012008, Jan 2017. [doi:{10.1103/PhysRevD.95.012008}](#).
- [A<sup>+</sup>17d] D.S. Akerib et al. 3D Modeling of Electric Fields in the LUX Detector. *JINST*, 12(11):P11022, 2017. [arXiv:1709.00095](#), [doi:10.1088/1748-0221/12/11/P11022](#).
- [A<sup>+</sup>17e] D.S. Akerib et al. First Searches for Axions and Axionlike Particles with the LUX Experiment. *Phys. Rev. Lett.*, 118(26):261301, 2017. [arXiv:1704.02297](#), [doi:10.1103/PhysRevLett.118.261301](#).
- [A<sup>+</sup>17f] D.S. Akerib et al. Identification of radiopure titanium for the LZ dark matter experiment and future rare event searches. *Astroparticle Physics*, 96:1–10, 2017. [doi:{10.1016/j.astropartphys.2017.09.002}](#).
- [A<sup>+</sup>17g] D.S. Akerib et al. Limits on spin-dependent WIMP-nucleon cross section obtained from the complete LUX exposure. *Phys. Rev. Lett.*, 118(25):251302, 2017. [arXiv:1705.03380](#), [doi:10.1103/PhysRevLett.118.251302](#).
- [A<sup>+</sup>17h] D.S. Akerib et al. Results from a search for dark matter in the complete LUX exposure. *Phys. Rev. Lett.*, 118(2):021303, 2017. [arXiv:1608.07648](#), [doi:10.1103/PhysRevLett.118.021303](#).
- [A<sup>+</sup>17i] Shadab Alam et al. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *Mon. Not. Roy. Astron. Soc.*, 470(3):2617–2652, 2017. [arXiv:1607.03155](#), [doi:10.1093/mnras/stx721](#).
- [A<sup>+</sup>17j] V. Anastassopoulos et al. New CAST Limit on the Axion-Photon Interaction. *Nature Physics*, 13:584–590, 2017. [arXiv:1705.02290](#), [doi:10.1038/nphys4109](#).
- [A<sup>+</sup>17k] E. Aprile et al. First Dark Matter Search Results from the XENON1T Experiment. *Phys. Rev. Lett.*, 119(18):181301, 2017. [arXiv:1705.06655](#), [doi:10.1103/PhysRevLett.119.181301](#).

- [A<sup>+</sup>18a] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints*, 7 2018. [arXiv:1807.06209](#).
- [A<sup>+</sup>18b] P. Agnes et al. Low-Mass Dark Matter Search with the DarkSide-50 Experiment. *Phys. Rev. Lett.*, 121(8):081307, 2018. [arXiv:1802.06994](#), [doi:10.1103/PhysRevLett.121.081307](#).
- [A<sup>+</sup>18c] M. Agostini et al. Improved Limit on Neutrinoless Double- $\beta$  Decay of  $^{76}\text{Ge}$  from GERDA Phase II. *Phys. Rev. Lett.*, 120(13):132503, 2018. [arXiv:1803.11100](#), [doi:10.1103/PhysRevLett.120.132503](#).
- [A<sup>+</sup>18d] D. S. Akerib et al. Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment. *Phys. Rev.*, D97(10):102008, 2018. [doi:10.1103/PhysRevD.97.102008](#).
- [A<sup>+</sup>18e] D. S. Akerib et al. Position Reconstruction in LUX. *JINST*, 13(02):P02001, 2018. [doi:10.1088/1748-0221/13/02/P02001](#).
- [A<sup>+</sup>18f] D.S. Akerib et al. Chromatographic separation of radioactive noble gases from xenon. *Astropart. Phys.*, 97:80–87, 2018. [arXiv:1605.03844](#), [doi:10.1016/j.astropartphys.2017.10.014](#).
- [A<sup>+</sup>18g] J. B. Albert et al. Sensitivity and Discovery Potential of nEXO to Neutrinoless Double Beta Decay. *Phys. Rev.*, C97(6):065503, 2018. [doi:10.1103/PhysRevC.97.065503](#).
- [A<sup>+</sup>18h] C. Alduino et al. First Results from CUORE: A Search for Lepton Number Violation via  $0\nu\beta\beta$  Decay of  $^{130}\text{Te}$ . *Phys. Rev. Lett.*, 120(13):132501, 2018. [arXiv:1710.07988](#), [doi:10.1103/PhysRevLett.120.132501](#).
- [A<sup>+</sup>18i] E. Aprile et al. Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys. Rev. Lett.*, 121(11):111302, 2018. [arXiv:1805.12562](#), [doi:10.1103/PhysRevLett.121.111302](#).
- [A<sup>+</sup>19a] Morad Aaboud et al. Combination of searches for invisible Higgs boson decays with the ATLAS experiment. *Phys. Rev. Lett.*, 122(23):231801, 2019. [arXiv:1904.05105](#), [doi:10.1103/PhysRevLett.122.231801](#).
- [A<sup>+</sup>19b] Morad Aaboud et al. Constraints on mediator-based dark matter and scalar dark energy models using  $\sqrt{s} = 13$  TeV  $pp$  collision data collected by the ATLAS detector. *JHEP*, 05:142, 2019. [arXiv:1903.01400](#), [doi:10.1007/JHEP05\(2019\)142](#).
- [A<sup>+</sup>19c] A.H. Abdelhameed et al. First results from the CRESST-III low-mass dark matter program. *Phys. Rev. D*, 100(10):102002, 2019. [arXiv:1904.00498](#), [doi:10.1103/PhysRevD.100.102002](#).
- [A<sup>+</sup>19d] M. Aker et al. Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN. *Phys. Rev. Lett.*, 123(22):221802, 2019. [arXiv:1909.06048](#), [doi:10.1103/PhysRevLett.123.221802](#).

- [A<sup>+</sup>19e] P. A. Amaudruz et al. In-situ characterization of the Hamamatsu R5912-HQE photomultiplier tubes used in the DEAP-3600 experiment. *Nucl. Instrum. Meth.*, A922:373–384, 2019. doi:10.1016/j.nima.2018.12.058.
- [A<sup>+</sup>19f] C. Amole et al. Dark Matter Search Results from the Complete Exposure of the PICO-60 C<sub>3</sub>F<sub>8</sub> Bubble Chamber. *Phys. Rev. D*, 100(2):022001, 2019. arXiv:1902.04031, doi:10.1103/PhysRevD.100.022001.
- [A<sup>+</sup>19g] G. Anton et al. Search for Neutrinoless Double-Beta Decay with the Complete EXO-200 Dataset. *ArXiv e-prints*, 2019. arXiv:1906.02723. arXiv:1906.02723.
- [A<sup>+</sup>19h] E. Aprile et al. Constraining the spin-dependent WIMP-nucleon cross sections with XENON1T. *Phys. Rev. Lett.*, 122(14):141301, 2019. arXiv:1902.03234, doi:10.1103/PhysRevLett.122.141301.
- [A<sup>+</sup>19i] E. Aprile et al. Observation of two-neutrino double electron capture in <sup>124</sup>Xe with XENON1T. *Nature*, 568(7753):532–535, 2019. arXiv:1904.11002, doi:10.1038/s41586-019-1124-4.
- [A<sup>+</sup>19j] O. Azzolini et al. Evidence of Single State Dominance in the Two-Neutrino Double- $\beta$  Decay of <sup>82</sup>Se with CUPID-0. *Phys. Rev. Lett.*, 123(26):262501, 2019. arXiv:1909.03397, doi:10.1103/PhysRevLett.123.262501.
- [A<sup>+</sup>20a] F. Agostini et al. Sensitivity of the DARWIN observatory to the neutrinoless double beta decay of <sup>136</sup>Xe. *arXiv e-prints*, 3 2020. arXiv:2003.13407.
- [A<sup>+</sup>20b] D. S. Akerib et al. Measurement of the Gamma Ray Background in the Davis Cavern at the Sanford Underground Research Facility. *Astropart. Phys.*, 116:102391, 2020. arXiv:1904.02112, doi:10.1016/j.astropartphys.2019.102391.
- [A<sup>+</sup>20c] D. S. Akerib et al. Projected sensitivity of the LUX-ZEPLIN experiment to the  $0\nu\beta\beta$  decay of <sup>136</sup>Xe. *Phys. Rev. C*, 102:014602, 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevC.102.014602>, doi:10.1103/PhysRevC.102.014602.
- [A<sup>+</sup>20d] D. S. Akerib et al. Projected WIMP sensitivity of the LUX-ZEPLIN dark matter experiment. *Phys. Rev. D*, 101:052002, 2020. URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.052002>, doi:{10.1103/PhysRevD.101.052002}.
- [A<sup>+</sup>20e] D.S. Akerib et al. Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers. 4 2020. arXiv:2004.06304.
- [A<sup>+</sup>20f] D.S. Akerib et al. Investigation of background electron emission in the LUX detector. 4 2020. arXiv:2004.07791.
- [A<sup>+</sup>20g] D.S. Akerib et al. The LUX-ZEPLIN (LZ) Experiment. *Nucl. Instrum. Meth. A*, 953:163047, 2020. arXiv:1910.09124, doi:10.1016/j.nima.2019.163047.
- [A<sup>+</sup>20h] D.S. Akerib et al. The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs. *arXiv e-prints*, 6 2020. arXiv:2006.02506.



- [A<sup>+</sup>20i] G. Anton et al. Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment. *Phys. Rev. C*, 101(6):065501, 2020. [arXiv:1908.04128](#), [doi:10.1103/PhysRevC.101.065501](#).
- [A<sup>+</sup>20j] E. Aprile et al. Energy resolution and linearity in the keV to MeV range measured in XENON1T. *arXiv 2003.03825*, 2020. [arXiv:2003.03825](#).
- [A<sup>+</sup>20k] E. Aprile et al. Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment. 7 2020. [arXiv:2007.08796](#).
- [A<sup>+</sup>20l] E. Armengaud et al. Precise measurement of  $2\nu\beta\beta$  decay of  $^{100}\text{Mo}$  with the CUPID-Mo detection technology. *Eur. Phys. J. C*, 80(7):674, 2020. [arXiv:1912.07272](#), [doi:10.1140/epjc/s10052-020-8203-4](#).
- [A<sup>+</sup>20m] E. Armengaud et al. The CUPID-Mo experiment for neutrinoless double-beta decay: performance and prospects. *Eur. Phys. J. C*, 80(1):44, 2020. [arXiv:1909.02994](#), [doi:10.1140/epjc/s10052-019-7578-6](#).
- [A<sup>+</sup>20n] Q. Arnaud et al. First germanium-based constraints on sub-MeV Dark Matter with the EDELWEISS experiment. *Phys. Rev. Lett.*, 125(14):141301, 2020. [arXiv:2003.01046](#), [doi:10.1103/PhysRevLett.125.141301](#).
- [A<sup>+</sup>20o] M. Askins et al. THEIA: an advanced optical neutrino detector. *Eur. Phys. J. C*, 80(5):416, 2020. [arXiv:1911.03501](#), [doi:10.1140/epjc/s10052-020-7977-8](#).
- [A<sup>+</sup>21] D.S. Akerib et al. Simulations of Events for the LUX-ZEPLIN (LZ) Dark Matter Experiment. *Astropart. Phys.*, 125:102480, 2021. [arXiv:2001.09363](#), [doi:10.1016/j.astropartphys.2020.102480](#).
- [ABD17] Matteo Agostini, Giovanni Benato, and Jason Detwiler. Discovery probability of next-generation neutrinoless double- $\beta$  decay experiments. *Phys. Rev. D*, 96(5):053001, 2017. [arXiv:1705.02996](#), [doi:10.1103/PhysRevD.96.053001](#).
- [AFZ06] G.W. Angus, B. Famaey, and H.S. Zhao. Can MOND take a bullet? Analytical comparisons of three versions of MOND beyond spherical symmetry. *Mon. Not. Roy. Astron. Soc.*, 371:138–146, 2006. URL: <http://adsabs.harvard.edu/abs/2006MNRAS.371..138A>, [doi:10.1111/j.1365-2966.2006.10668.x](#).
- [Agg15] Charu C. Aggarwal. *Data Mining: The Textbook*, volume 1, XXIX, 734. Springer International Publishing, 2015. [doi:{10.1007/978-3-319-14142-8}](#).
- [Agg18] Charu C. Aggarwal. *Neural Networks and Deep Learning*, volume 1, XXIII, 497. Springer International Publishing, 2018. [doi:{10.1007/978-3-319-94463-0}](#).
- [AH48] R.A. Alpher and R.C. Herman. On the Relative Abundance of the Elements. *Phys. Rev.*, 74:1737–1742, 1948. URL: <http://adsabs.harvard.edu/abs/1948PhRv...74.1737A>, [doi:10.1103/PhysRev.74.1737](#).

- [All20] Genevera Allen. *Can we trust scientific discoveries made using machine learning?*, 2019 (last accessed June 1, 2020). [eurekaalert.org/pub\\_releases/2019-02/ru-cwt021119.php](http://eurekaalert.org/pub_releases/2019-02/ru-cwt021119.php).
- [AM07] M. A. H. Akhand and Kazuyuki Murase. Neural Network Ensemble Training by Sequential Interaction. In *Artificial Neural Networks – ICANN 2007*, pages 98–108, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. doi:10.1007/978-3-540-74690-4\_11.
- [B<sup>+</sup>93] P. Benetti et al. Detection of energy deposition down to the keV region using liquid xenon scintillation. *Nucl. Instrum. Meth. A*, 327(1):203 – 206, 1993. doi:10.1016/0168-9002(93)91442-P.
- [B<sup>+</sup>01] G. Barrand et al. GAUDI - A software architecture and framework for building HEP data processing applications. *Comput. Phys. Commun.*, 140(1):45–55, 2001. URL: <http://www.sciencedirect.com/science/article/pii/S0010465501002545>, doi:10.1016/S0010-4655(01)00254-5.
- [B<sup>+</sup>03] R. Bernabei et al. Dark matter search. *Riv. Nuovo Cim.*, 26N1:1–73, 2003. arXiv:astro-ph/0307403. arXiv:astro-ph/0307403.
- [B<sup>+</sup>05] R. Barbier et al. R-parity violating supersymmetry. *Phys. Rept.*, 420:1–202, 2005. arXiv:hep-ph/0406039, doi:10.1016/j.physrep.2005.08.006.
- [B<sup>+</sup>06] C. A. Baker et al. Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Phys. Rev. Lett.*, 97:131801, 2006. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.97.131801>, doi:10.1103/PhysRevLett.97.131801.
- [B<sup>+</sup>12a] E. Behnke et al. First dark matter search results from a 4-kg CF<sub>3</sub>I bubble chamber operated in a deep underground site. *Physical Review D*, 86(5):052001, 2012. URL: <http://adsabs.harvard.edu/abs/2012PhRvD..86e2001B>, doi:10.1103/PhysRevD.86.052001.
- [B<sup>+</sup>12b] J. Beringer et al. Review of Particle Physics. *Phys. Rev. D*, 86:010001, 2012. URL: <http://link.aps.org/doi/10.1103/PhysRevD.86.010001>, doi:10.1103/PhysRevD.86.010001.
- [B<sup>+</sup>12c] J. Bovy et al. The Milky Way’s Circular-velocity Curve between 4 and 14 kpc from APOGEE data. *Astrophys. J.*, 759:131, 2012. URL: <http://adsabs.harvard.edu/abs/2012ApJ...759..131B>, arXiv:1209.0759, doi:10.1088/0004-637X/759/2/131.
- [B<sup>+</sup>13] R. Bernabei et al. Final model independent result of DAMA/LIBRA-phase1. *European Physical Journal C*, 73:2648, 2013. URL: <http://adsabs.harvard.edu/abs/2013EPJC...73.2648B>, doi:10.1140/epjc/s10052-013-2648-7.
- [B<sup>+</sup>14] L. Baudis et al. Neutrino physics with multi-ton scale liquid xenon detectors. *Journal Cosm Astropart Phys*, 2014(01):044, 2014. doi:{10.1088/1475-7516/2014/01/044}.

- [B<sup>+</sup>15] A. Best et al. Low energy neutron background in deep underground laboratories. *Nuclear Instrumentation and Methods A*, 812:1–6, 2015. doi:10.1016/j.nima.2015.12.034.
- [B<sup>+</sup>18] A. S. Barabash et al. Final results of the Aurora experiment to study  $2\beta$  decay of  $^{116}\text{Cd}$  with enriched  $^{116}\text{CdWO}_4$  crystal scintillators. *Phys. Rev. D*, 98(9):092007, 2018. arXiv:1811.06398, doi:10.1103/PhysRevD.98.092007.
- [B<sup>+</sup>20] T. Braine et al. Extended Search for the Invisible Axion with the Axion Dark Matter Experiment. *Phys. Rev. Lett.*, 124(10):101303, 2020. arXiv:1910.08638, doi:10.1103/PhysRevLett.124.101303.
- [Bah87] J.N. Bahcall. Neutrino-electron scattering and solar neutrino experiments. *Rev. Mod. Phys.*, 59:505–521, 1987. URL: <https://link.aps.org/doi/10.1103/RevModPhys.59.505>, doi:10.1103/RevModPhys.59.505.
- [Bai16] Adam Bailey. *Dark matter searches and study of electrode design in LUX and LZ*. PhD thesis, Imperial College London, London, England, 2016. doi:{10.25560/41878}.
- [Bar20] Alexander Barabash. Precise Half-Life Values for Two-Neutrino Double- $\beta$  Decay: 2020 Review. *Universe*, 6(10):159, 2020. arXiv:2009.14451, doi:10.3390/universe6100159.
- [Bay63] Thomas Bayes. An essay towards solving a problem in the doctrine of chances. by the late rev. mr. bayes, f. r. s. communicated by mr. price, in a letter to john canton, a. m. f. r. s. *Phil. Trans. R. Soc.*, (53):370–418, 1763. doi:10.1098/rstl.1763.0053.
- [BC10] D.R. Bes and O. Civitarese. Matrix elements for the ground-state to ground-state  $2\nu\beta^-\beta^-$  decay of Te isotopes in a hybrid model. *Phys. Rev. C*, 81:014315, 2010. doi:10.1103/PhysRevC.81.014315.
- [BCD<sup>+</sup>08] M.-M. Bé, V. Chisté, C. Dullieu, E. Browne, V. Chechev, N. Kuzmenko, F. Kondev, A. Luca, M. Galán, A. Pearce, and X. Huang. *Table of Radionuclides*, volume 4 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, Pavillon de Breteuil, F-92310 Sèvres, France, 2008. URL: [http://www.bipm.org/utils/common/pdf/monographieRI/Monographie\\_BIPM-5\\_Tables\\_Vol4.pdf](http://www.bipm.org/utils/common/pdf/monographieRI/Monographie_BIPM-5_Tables_Vol4.pdf).
- [BCD<sup>+</sup>13] M.-M. Bé, V. Chisté, C. Dullieu, X. Mougeot, V.P. Chechev, F.G. Kondev, A.L. Nichols, X. Huang, and B. Wang. *Table of Radionuclides*, volume 7 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, Pavillon de Breteuil, F-92310 Sèvres, France, 2013. URL: [http://www.bipm.org/utils/common/pdf/monographieRI/Monographie\\_BIPM-5\\_Tables\\_Vol7.pdf](http://www.bipm.org/utils/common/pdf/monographieRI/Monographie_BIPM-5_Tables_Vol7.pdf).
- [BCnG10] Gianfranco Bertone, David G Cerdeño, and Anne M. Green. *Particle Dark Matter: Observations, Models and Searches*, chapter 17, pages 347–369. Cambridge University Press, 2010. ISBN:9780521763684.

- [BD12] Philippe Brax and Anne-Christine Davis. Modified Gravity and the CMB. *Phys. Rev. D*, 85:023513, 2012. arXiv:1109.5862, doi:10.1103/PhysRevD.85.023513.
- [BDM<sup>+</sup>17] A. Butcher, L. Doria, J. Monroe, F. Retière, B. Smith, and J. Walding. A method for characterizing after-pulsing and dark noise of PMTs and SiPMs. *Nucl. Instrum. Meth.*, A875:87–91, 2017. doi:10.1016/j.nima.2017.08.035.
- [Ben12] Yoshua Bengio. Practical recommendations for gradient-based training of deep architectures. *arXiv e-prints [1206.5533]*, 2012. arXiv:1206.5533.
- [BFH10] Yang Bai, Patrick J. Fox, and Roni Harnik. The Tevatron at the Frontier of Dark Matter Direct Detection. *JHEP*, 12:048, 2010. arXiv:1005.3797, doi:10.1007/JHEP12(2010)048}.
- [Bis06] Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006. ISBN:978-0-387-31073-2.
- [BKS<sup>+</sup>09] Michael Boylan-Kolchin, Volker Springel, Simon D.M. White, Adrian Jenkins, and Gerard Lemson. Resolving cosmic structure formation with the Millennium-II Simulation. *Mon. Not. Roy. Astron. Soc.*, 398(3):1150–1164, 2009. doi:10.1111/j.1365-2966.2009.15191.x}.
- [Bla13] Pasquale Blasi. The Origin of Galactic Cosmic Rays. *Astron. Astrophys. Rev.*, 21:70, 2013. doi:10.1007/s00159-013-0070-7}.
- [BM07] J.R. Brownstein and J.W. Moffat. The Bullet Cluster 1E0657-558 evidence shows Modified Gravity in the absence of Dark Matter. *Mon. Not. Roy. Astron. Soc.*, 382:29–47, 2007. arXiv:astro-ph/0702146, doi:10.1111/j.1365-2966.2007.12275.x}.
- [BP04] John N. Bahcall and M.H. Pinsonneault. What do we (not) know theoretically about solar neutrino fluxes? *Phys. Rev. Lett.*, 92:121301, 2004. arXiv:astro-ph/0402114, doi:10.1103/PhysRevLett.92.121301}.
- [BPB01] John N. Bahcall, M.H. Pinsonneault, and Sarbani Basu. Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties. *Astrophys. J.*, 555:990–1012, 2001. arXiv:astro-ph/0010346, doi:10.1086/321493}.
- [Bra14] A.W. Bradley. *LUX Thermosyphon Cryogenics and Radon-related Backgrounds for the First WIMP Result*. PhD thesis, Case Western Reserve University, Department of Physics at Case Western Reserve University, Cleveland, Ohio, May 2014.
- [Bra16] Max Bramer. *Principles of Data Mining*, volume 3, XV, 526. Springer-Verlag London, 2016. doi:10.1007/978-1-4471-7307-6}.
- [Brá18] P. Brás. The 136xe neutrinoless double beta decay search with lz. In *Neutrino 2018 - XXVIII International Conference on Neutrino Physics and Astrophysics*. Zenodo, jun 2018. URL: <https://doi.org/10.5281/zenodo.1300887>, doi:10.5281/zenodo.1300887}.

- [Bre01] L. Breiman. Random Forests. *Machine Learning*, 45, 2001. doi:10.1023/A:1010933404324.
- [BSB05] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. New solar opacities, abundances, helioseismology, and neutrino fluxes. *Astrophys. J. Lett.*, 621:L85–L88, 2005. arXiv:astro-ph/0412440, doi:10.1086/428929.
- [BTZ14] N. Barros, J. Thurn, and K. Zuber. Double beta decay searches of  $^{134}\text{Xe}$ ,  $^{126}\text{Xe}$  and  $^{124}\text{Xe}$  with large scale Xe detectors. *Journal of Physics G Nuclear Physics*, 41(11), Nov 2014. doi:10.1088/0954-3899/41/11/115105.
- [C<sup>+</sup>15] François Chollet et al. Keras, 2015. <https://keras.io>.
- [C<sup>+</sup>17a] Xun Chen et al. PandaX-III: Searching for neutrinoless double beta decay with high pressure  $^{136}\text{Xe}$  gas time projection chambers. *Sci. China Phys. Mech. Astron.*, 60(6):061011, 2017. arXiv:1610.08883, doi:10.1007/s11433-017-9028-0.
- [C<sup>+</sup>17b] Xiangyi Cui et al. Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment. *Phys. Rev. Lett.*, 119(18):181302, 2017. arXiv:1708.06917, doi:10.1103/PhysRevLett.119.181302.
- [CA13] V. Chepel and H. Araújo. Liquid noble gas detectors for low energy particle physics. *JINST*, 8:R04001, 2013. doi:10.1088/1748-0221/8/04/R04001.
- [CBG<sup>+</sup>06] Douglas Clowe, Maruša Bradač, Anthony H. Gonzalez, Maxim Markevitch, Scott W. Randall, Christine Jones, and Dennis Zaritsky. A Direct Empirical Proof of the Existence of Dark Matter. *Astrophys. J.*, 648(2):L109, 2006. URL: <http://stacks.iop.org/1538-4357/648/i=2/a=L109>, doi:10.1086/508162.
- [CCGV11] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71:1554, 2011. [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)]. arXiv:1007.1727, doi:10.1140/epjc/s10052-011-1554-0.
- [CD77] B. E. Carpenter and R. W. Doran. The other Turing machine. *The Computer Journal*, 20(3):269–279, 1977. doi:{10.1093/comjnl/20.3.269}.
- [CDDJ<sup>+</sup>98] B.T. Cleveland, T. Daily, R. Davis Jr., J.R. Distel, K. Lande, C.K. Lee, P.S. Wildenhain, and J. Ullman. Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector. *The Astrophysical Journal*, 496(1):505, 1998.
- [CDL<sup>+</sup>83] M. L. Cherry, M. Deakyne, K. Lande, C. K. Lee, R. I. Steinberg, B. Cleveland, and E. J. Fenyes. Multiple muons in the Homestake underground detector. *Phys. Rev. D*, 27:1444–1447, Apr 1983. URL: <https://link.aps.org/doi/10.1103/PhysRevD.27.1444>, doi:10.1103/PhysRevD.27.1444.
- [Cha32] J. Chadwick. The Existence of a Neutron. *Proc. Roy. Soc. Lond. A*, 136(830):692–708, 1932. doi:10.1098/rspa.1932.0112.

- [Chi20] Oliver Child. *Menace: the Machine Educable Noughts And Crosses Engine*, 2016 (last accessed June 1, 2020). [chalkdustmagazine.com/features/menace-machine-educable-noughts-crosses-engine/](http://chalkdustmagazine.com/features/menace-machine-educable-noughts-crosses-engine/).
- [CHL15] Eric Carlson, Dan Hooper, and Tim Linden. Improving the Sensitivity of Gamma-Ray Telescopes to Dark Matter Annihilation in Dwarf Spheroidal Galaxies. *Phys. Rev.*, D91(6):061302, 2015. arXiv:1409.1572, doi:{10.1103/PhysRevD.91.061302}.
- [Cho15] Chollet, François and others. Adamax - Keras documentation, 2015. [keras.io/api/optimizers/adamax](https://keras.io/api/optimizers/adamax), last accessed on June 14 2020.
- [CK09] Robert Caldwell and Marc Kamionkowski. Cosmology: Dark matter and dark energy. *Nature*, 458:587–589, 2009. doi:{10.1038/458587a}.
- [Col15] Planck Collaboration. Planck mission picture gallery. <http://www.cosmos.esa.int/web/planck/picture-gallery>, 2015. last accessed 31 August 2020.
- [CPT92] S.M. Carroll, W.H. Press, and E.L. Turner. The cosmological constant. *Ann. Rev. Astron. Astrophys.*, 30:499–542, 1992. URL: <http://adsabs.harvard.edu/abs/1992ARA%26A..30..499C>, doi:10.1146/annurev.aa.30.090192.002435.
- [CRH<sup>+</sup>56] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the Free Neutrino: A Confirmation. *Science*, 124:103–104, Jul 1956. doi:{10.1126/science.124.3212.103}.
- [CUH15] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). *arXiv preprint arXiv:1511.07289*, 2015.
- [Cv13] S. Cooper and J. van Leeuwen. *Alan Turing: His Work and Impact*. Elsevier Science, 1st edition, 2013. ISBN: 9780123869807.
- [DAP15] A. De Angelis and M.J.M. Pimenta. *Introduction to Particle and Astroparticle Physics: Questions to the Universe*. Undergraduate Lecture Notes in Physics. Springer Milan, 2015. URL: <https://books.google.pt/books?id=J06GCgAAQBAJ>.
- [Det20] Jason Detwiler. Future Neutrinoless Double-Beta Decay Experiments. In *Neutrino2020*, 2020. URL: [https://indico.fnal.gov/event/43209/contributions/187827/attachments/130703/159511/20200701\\_Nu20\\_FutureNDBD\\_Detwiler.pdf](https://indico.fnal.gov/event/43209/contributions/187827/attachments/130703/159511/20200701_Nu20_FutureNDBD_Detwiler.pdf).
- [dG<sup>+</sup>13] A. de Gouvea et al. Working Group Report: Neutrinos. In *Community Summer Study 2013: Snowmass on the Mississippi*, 10 2013. arXiv:1310.4340.

- [DGG<sup>+</sup>62] G. Danby, J.M. Gaillard, K. Goulianos, L.M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962. doi:{10.1103/PhysRevLett.9.36}.
- [DHS11] John Duchi, Elad Hazan, and Yoram Singer. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011. [jmlr.org/papers/v12/duchi11a.html](http://jmlr.org/papers/v12/duchi11a.html).
- [DLR70] B. A. Dolgoshein, V. A. Lebedenko, and B.U. Rodionov. New Method of Registration of Ionizing-particle Tracks in Condensed Matter. *JETP Lett.*, 11:351, 1970.
- [DLR77] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977. URL: <http://www.jstor.org/stable/2984875>.
- [DMVV16] Stefano Dell’Oro, Simone Marcocci, Matteo Viel, and Francesco Vissani. Neutrinoless double beta decay: 2015 review. *Advances in High Energy Physics*, 2016. doi:{10.1155/2016/2162659}.
- [DNN08] Sacha Davidson, Enrico Nardi, and Yosef Nir. Leptogenesis. *Phys. Rept.*, 466:105–177, 2008. arXiv:0802.2962, doi:10.1016/j.physrep.2008.06.002.
- [Doz16] Timothy Dozat. Incorporating Nesterov Momentum into Adam. In , 2016.
- [DPR19] Michelle J. Dolinski, Alan W.P. Poon, and Werner Rodejohann. Neutrinoless Double-Beta Decay: Status and Prospects. *Annual Review of Nuclear and Particle Science*, 69(1):219–251, 2019. doi:{10.1146/annurev-nucl-101918-023407}.
- [EE14] H. Ejiri and S.R. Elliott. Charged current neutrino cross section for solar neutrinos, and background to  $\beta\beta(0\nu)$  experiments. *Phys. Rev. C*, 89(5):055501, 2014. arXiv:1309.7957, doi:10.1103/PhysRevC.89.055501.
- [EE17] S. R. Elliott and H. Ejiri. Solar Neutrinos as Background to Neutrinoless Double-beta Decay Experiments. *AIP Conf. Proc.*, 1894(1):020008, 2017. doi:10.1063/1.5007633.
- [EK SX96] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In , pages 226–231. AAAI Press, 1996.
- [End00] P.M. Endt. The Berkeley Laboratory Isotopes Project, May 2000. accessed 10 May 2015. URL: <http://ie.lbl.gov/education/isotopes.htm>.
- [Ert17] Wolfgang Ertel. *Introduction to Artificial Intelligence*, volume 2, XIV, 356. Springer International Publishing, 2017. doi:{10.1007/978-3-319-58487-4}.

- [ESW07] Daniel J. Eisenstein, Hee-jong Seo, and Martin J. White. On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter. *Astrophys. J.*, 664:660–674, 2007. [arXiv:astro-ph/0604361](#), doi:{10.1086/518755}.
- [EV02] S.R. Elliott and P. Vogel. Double beta decay. *Annual Review of Nuclear and Particle Science*, 52:115–151, 2002. doi:{10.1146/annurev.nucl.52.050102.090641}.
- [F<sup>+</sup>98a] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998. [arXiv:hep-ex/9807003](#), doi:10.1103/PhysRevLett.81.1562.
- [F<sup>+</sup>98b] Y. Fukuda et al. Measurements of the Solar Neutrino Flux from Super-Kamiokande’s First 300 Days. *Phys. Rev. Lett.*, 81, Aug 1998. doi:10.1103/PhysRevLett.81.1158.
- [F<sup>+</sup>01] Wendy L. Freedman et al. Final Results from the Hubble Space Telescope Key Project to Measure the Hubble Constant. *Astrophys. J.*, 553(1):47, 2001. URL: <http://stacks.iop.org/0004-637X/553/i=1/a=47>, doi:10.1086/320638.
- [F<sup>+</sup>02] S. Fukuda et al. Determination of solar neutrino oscillation parameters using 1496 days of Super-Kamiokande I data. *Phys. Lett.*, B539:179–187, 2002. doi:10.1016/S0370-2693(02)02090-7.
- [F<sup>+</sup>10] D. Ferrucci et al. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3):59–79, 2010. doi:{10.1609/aimag.v31i3.2303}.
- [F<sup>+</sup>11] Dave Ferrucci et al. Watson jeopardy! challenge, 2011. accessed on 2020/01/17. URL: [https://researcher.watson.ibm.com/researcher/view\\_page.php?id=6859](https://researcher.watson.ibm.com/researcher/view_page.php?id=6859).
- [F<sup>+</sup>20] Y. Fukuda et al. ZICOS - A new project for neutrinoless double beta decay using Zirconium complex in organic liquid scintillator. *Journal of Physics: Conference Series*, 1342:012093, 2020. URL: <https://doi.org/10.1088/1742-6596/1342/1/012093>, doi:{10.1088/1742-6596/1342/1/012093}.
- [Fah14] C.H. Faham. *Prototype, Surface Commissioning and Photomultiplier Tube Characterization for the Large Underground Xenon (LUX) Direct Dark Matter Search Experiment*. PhD thesis, Brown University, Physics department at Brown University, Providence, Rhode Island, May 2014.
- [FC98] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998. [arXiv:physics/9711021](#), doi:10.1103/PhysRevD.57.3873.
- [Fer34] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934. doi:10.1007/BF01351864.



- [Fis18] Vincent Fischer. Search for neutrinoless double-beta decay with SNO+. In *13th Conference on the Intersections of Particle and Nuclear Physics*, 9 2018. arXiv:1809.05986.
- [FMV14] N. Fornengo, L. Maccione, and A. Vittino. Constraints on particle dark matter from cosmic-ray antiprotons. *JCAP*, 1404(04):003, 2014. doi:10.1088/1475-7516/2014/04/003.
- [Fri99a] A. Friedmann. On the Curvature of Space. *General Relativity and Gravitation*, 31(12):1991–2000, 1999. URL: <http://dx.doi.org/10.1023/A%3A1026751225741>, doi:10.1023/A:1026751225741.
- [Fri99b] A. Friedmann. On the Possibility of a World with Constant Negative Curvature of Space. *General Relativity and Gravitation*, 31(12):2001–2008, 1999. URL: <http://dx.doi.org/10.1023/A%3A1026755309811>, doi:10.1023/A:1026755309811.
- [FRS<sup>+</sup>14] G. W. Fraser, A. M. Read, S. Sembay, J. A. Carter, and E. Schyns. Potential solar axion signatures in X-ray observations with the XMM–Newton observatory. *Mon. Not. Roy. Astron. Soc.*, 445(2):2146–2168, 2014. arXiv:1403.2436, doi:{10.1093/mnras/stu1865}.
- [FSB08] Ignacio Ferreras, Prasenjit Saha, and Scott Burles. Unveiling dark halos in lensing galaxies. *Mon. Not. Roy. Astron. Soc.*, 383:857, 2008. arXiv:0710.3159, doi:10.1111/j.1365-2966.2007.12606.x.
- [FSW05] Ignacio Ferreras, Prasenjit Saha, and Liliya L.R. Williams. Stellar and total mass in early-type lensing galaxies. *Astrophys. J. Lett.*, 623:L5–L8, 2005. arXiv:astro-ph/0503168, doi:10.1086/429995.
- [FT13] C. Fry and M. Thoennessen. Discovery of actinium, thorium, protactinium, and uranium isotopes. *Atomic Data and Nuclear Data Tables*, 99(3):345 – 364, 2013. doi:{10.1016/j.adt.2012.03.002}.
- [G<sup>+</sup>16] A. Gando et al. Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen. *Phys. Rev. Lett.*, 117(8):082503, 2016. [Addendum: Phys.Rev.Lett. 117, 109903 (2016)]. arXiv:1605.02889, doi:10.1103/PhysRevLett.117.082503.
- [Gai16] Thomas K. Gaisser. Atmospheric Neutrinos. *J. Phys. Conf. Ser.*, 718(5):052014, 2016. doi:10.1088/1742-6596/718/5/052014.
- [Gam48] G. Gamow. The Origin of Elements and the Separation of Galaxies. *Phys. Rev.*, 74:505–506, 1948. URL: <http://link.aps.org/doi/10.1103/PhysRev.74.505.2>, doi:10.1103/PhysRev.74.505.2.
- [Gan20] Yoshihito Gando. First results of KamLAND-Zen 800. *J. Phys. Conf. Ser.*, 1468(1):012142, 2020. doi:10.1088/1742-6596/1468/1/012142.

- [Gav01] V. N. Gavrin. Solar neutrino results from SAGE. *Nucl. Phys. Proc. Suppl.*, 91:36–43, 2001. doi:10.1016/S0920-5632(00)00920-8.
- [GGK<sup>+</sup>13] Yu.M. Gavriilyuk, A.M. Gangapshev, V.V. Kazalov, V.V. Kuzminov, S.I. Panasenko, and S.S. Ratkevich. Indications of  $2\nu 2K$  capture in  $^{78}\text{Kr}$ . *Phys. Rev. C*, 87(3):035501, 2013. doi:10.1103/PhysRevC.87.035501.
- [GGT95] Evalyn I. Gates, Geza Gyuk, and Michael S. Turner. The Local halo density. *Astrophys. J.*, 449:L123–L126, 1995. arXiv:astro-ph/9505039, doi:{10.1086/309652}.
- [GK07] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. OUP Oxford, 2007. ISBN:978-0-19-850871-7. URL: <https://books.google.pt/books?id=SdAcSwTR0CgC>.
- [Gra20] Christopher Grant. KamLAND-Zen and SNO+. In *Neutrino2020*, 2020. [https://indico.fnal.gov/event/43209/contributions/187836/attachments/130621/159305/CGrant\\_KLZandSNOplus\\_Neutrino2020\\_reduced.pdf](https://indico.fnal.gov/event/43209/contributions/187836/attachments/130621/159305/CGrant_KLZandSNOplus_Neutrino2020_reduced.pdf).
- [GTR<sup>+</sup>07] Raphael Gavazzi, Tommaso Treu, Jason D. Rhodes, Leon Ve Koopmans, Adam S. Bolton, Scott Burles, Richard Massey, and Leonidas A. Moustakas. The Sloan Lens ACS Survey. 4. The mass density profile of early-type galaxies out to 100 effective radii. *Astrophys. J.*, 667:176–190, 2007. arXiv:astro-ph/0701589, doi:10.1086/519237.
- [H<sup>+</sup>74] F.J. Hasert et al. Observation of Neutrino Like Interactions without Muon or Electron in the Gargamelle Neutrino Experiment. *Nucl. Phys. B*, 73:1–22, 1974. doi:10.1016/0550-3213(74)90038-8.
- [H<sup>+</sup>88] K.S. Hirata et al. Experimental Study of the Atmospheric Neutrino Flux. pages 235–246, 3 1988. doi:10.1016/0370-2693(88)91690-5.
- [H<sup>+</sup>91] K.S. Hirata et al. Real time, directional measurement of B-8 solar neutrinos in the Kamiokande-II detector. *Phys. Rev. D*, 44:2241, 1991. [Erratum: *Phys.Rev.D* 45, 2170 (1992)]. doi:10.1103/PhysRevD.44.2241.
- [Har08] Stevan Harnad. The Annotation Game: On Turing (1950) on Computing, Machinery, and Intelligence (PUBLISHED VERSION BOWDLERIZED). In R. Epstein, G. Roberts, and G. Beber, editors, *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*, pages 23–66. Springer, 2008. Chapter: 3 Commentary On: Turing, A.M. (1950) Computing Machinery and Intelligence. Mind 49 433-460 Address: Amsterdam.
- [HD02] Wayne Hu and Scott Dodelson. Cosmic Microwave Background Anisotropies. *Ann. Rev. Astron. Astrophys.*, 40:171–216, 2002. arXiv:astro-ph/0110414, doi:10.1146/annurev.astro.40.060401.093926.

- [Hel56] R. H. Helm. Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei. *Phys. Rev.*, 104:1466–1475, 1956. URL: <http://link.aps.org/doi/10.1103/PhysRev.104.1466>, doi:10.1103/PhysRev.104.1466.
- [HH00] Wick C. Haxton and Barry R. Holstein. Neutrino physics. *Am. J. Phys.*, 68:15–32, 2000. arXiv:hep-ph/9905257, doi:10.1119/1.19368.
- [Hin12] G. Hinton. Neural networks for machine learning. *Coursera Video*, 2012. [youtube.com/playlist?list=PLoRl3Ht4J0cdU872GhiYwf6jwrk\\_SNhz9](https://youtube.com/playlist?list=PLoRl3Ht4J0cdU872GhiYwf6jwrk_SNhz9), last accessed on June 14, 2020.
- [HN19] Jonathan Haefner and NEXT Collaboration Team. Recent results from NEXT-White, a demonstrator for the NEXT neutrinoless double beta decay experiment. In *APS April Meeting Abstracts*, volume 2019 of *APS Meeting Abstracts*, page D10.003, 2019.
- [HR13] Leroy Hood and Lee Rowen. The Human Genome Project: big science transforms biology and medicine. *Genome Medicine*, 5, 2013. doi:10.1186/gm483.
- [HS90] L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1001, 1990. doi:10.1109/34.58871.
- [HSW89] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. doi:10.1016/0893-6080(89)90020-8.
- [HTF09] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York, NY, 2nd edition, Feb 2009. ISBN 978-0-387-84858-7.
- [Hu08] Wayne Hu. Lecture Notes on CMB Theory: From Nucleosynthesis to Recombination. *ArXiv e-prints*, 2008. ArXiv:[0802.3688]. arXiv:0802.3688.
- [I<sup>+</sup>16] T. Iida et al. The CANDLES experiment for the study of Ca-48 double beta decay. *Nuclear and Particle Physics Proceedings*, 273-275:2633 – 2635, 2016. 37th International Conference on High Energy Physics (ICHEP). URL: <http://www.sciencedirect.com/science/article/pii/S2405601415009724>, doi:10.1016/j.nuclphysbps.2015.10.013.
- [Ian14] Aldo Ianni. Solar neutrinos and the solar model. *Physics of the Dark Universe*, 4:44 – 49, 2014. DARK TAUP2013. URL: <http://www.sciencedirect.com/science/article/pii/S2212686414000211>, doi:10.1016/j.dark.2014.06.002.
- [IM18] N. S Israel and J. W Moffat. The Train Wreck Cluster Abell 520 and the Bullet Cluster 1E0657-558 in a Generalized Theory of Gravitation. *Galaxies*, 6(2):41, 2018. arXiv:1606.09128, doi:10.3390/galaxies6020041.

- [IS17] Laura Igual and Santi Seguí. *Introduction to Data Science*, volume 1, XIV, 218. Springer International Publishing, 2017. doi:{10.1007/978-3-319-50017-1}.
- [Jak94] Karl Jakobs. The Physics results of the UA2 experiment at the Cern  $p\bar{p}$  collider. *Int. J. Mod. Phys. A*, 9:2903–2978, 1994. doi:10.1142/S0217751X94001163.
- [JKG96] Gerard Jungman, Marc Kamionkowski, and Kim Griest. Supersymmetric dark matter. *Phys. Rept.*, 267:195–373, 1996. arXiv:hep-ph/9506380, doi:{10.1016/0370-1573(95)00058-5}.
- [JWHT13] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning*, volume 1, XIV, 426. Springer-Verlag New York, 2013. doi:{10.1007/978-1-4614-7138-7}.
- [K<sup>+</sup>01] K. Kodama et al. Observation of tau neutrino interactions. *Phys. Lett.*, B504:218–224, 2001. doi:10.1016/S0370-2693(01)00307-0.
- [K<sup>+</sup>16] Vardan Khachatryan et al. Search for s channel single top quark production in pp collisions at  $\sqrt{s} = 7$  and 8 TeV. *arXiv e-prints*, 2016. arXiv:1603.02555.
- [Kan17] G. Kane. *Modern Elementary Particle Physics: Explaining and Extending the Standard Model*. Cambridge University Press, 2017. URL: [https://books.google.pt/books?id=Xt6\\_DwAAQBAJ](https://books.google.pt/books?id=Xt6_DwAAQBAJ).
- [KB14] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, 2014. URL: <https://ui.adsabs.harvard.edu/abs/2014arXiv1412.6980K>, arXiv:1412.6980.
- [KI12] J. Kotila and F. Iachello. Phase-space factors for double- $\beta$  decay. *Phys. Rev. C*, 85:034316, 2012. doi:10.1103/PhysRevC.85.034316.
- [KK98] Marc Kamionkowski and Ali Kinkhabwala. Galactic halo models and particle dark-matter detection. *Phys. Rev. D*, 57:3256–3263, 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevD.57.3256>, doi:10.1103/PhysRevD.57.3256.
- [KK04] Andriy Kurylov and Marc Kamionkowski. Generalized analysis of weakly interacting massive particle searches. *Phys. Rev.*, D69:063503, 2004. arXiv:hep-ph/0307185, doi:{10.1103/PhysRevD.69.063503}.
- [Kle96] E. M. Kleinberg. An overtraining-resistant stochastic modeling method for pattern recognition. *Ann. Statist.*, 24(6):2319–2349, 1996. doi:{10.1214/aos/1032181157}.
- [KM73] Makoto Kobayashi and Toshihide Maskawa. CP Violation in the Renormalizable Theory of Weak Interaction. *Prog. Theor. Phys.*, 49:652–657, 1973. doi:10.1143/PTP.49.652.

- [KSLB14] Prajwal Raj Kafle, Sanjib Sharma, Geraint F. Lewis, and Joss Bland-Hawthorn. On the Shoulders of Giants: Properties of the Stellar Halo and the Milky Way Mass Distribution. *Astrophys. J.*, 794(1):59, 2014. URL: <http://stacks.iop.org/0004-637X/794/i=1/a=59>, doi:10.1088/0004-637X/794/1/59.
- [Kub15] Miroslav Kubat. *An Introduction to Machine Learning*, volume 1, XIII, 291. Springer International Publishing, 2015. doi:{10.1007/978-3-319-20010-1}.
- [KYYK15] J. Koda, M. Yagi, H. Yamanoi, and Y. Komiyama. Approximately a Thousand Ultra-diffuse Galaxies in the Coma Cluster. *Astrophys. J.*, 807:L2, 2015. URL: <http://adsabs.harvard.edu/abs/2015ApJ...807L...2K>, arXiv:1506.01712, doi:10.1088/2041-8205/807/1/L2.
- [LC<sup>+</sup>19] J.M. López-Castaño et al. Recent Results of the Majorana Demonstrator Experiment. *AIP Conf. Proc.*, 2165(1):020018, 2019. arXiv:1909.07509, doi:10.1063/1.5130979.
- [LCG19] J.M. López-Castaño and Ian Guinn. Current status of LEGEND: Searching for Neutrinoless Double-Beta Decay in <sup>76</sup>Ge: Part II. *PoS, LeptonPhoton2019*:162, 2019. arXiv:1912.03308, doi:10.22323/1.367.0162.
- [Lee20] Moo Hyun Lee. AMoRE: A search for neutrinoless double-beta decay of 100Mo using low-temperature molybdenum-containing crystal detectors. *JINST*, 15(08):C08010, 2020. arXiv:2005.05567, doi:10.1088/1748-0221/15/08/C08010.
- [Lem31] Abbé G. Lemaître. A Homogeneous Universe of Constant Mass and Increasing Radius accounting for the Radial Velocity of Extra-galactic Nebulae. *Mon. Not. Roy. Astron. Soc.*, 91(5):483–490, 1931. URL: <http://mnras.oxfordjournals.org/content/91/5/483.short>, doi:10.1093/mnras/91.5.483.
- [Les15] K.T. Lesko. The Sanford Underground Research Facility at Homestake (SURF). *Physics Procedia*, 61:542–551, 2015. 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013. doi:{10.1016/j.phpro.2014.12.001}.
- [LH07] S. Legg and M. Hutter. A Collection of Definitions of Intelligence. *arXiv e-prints*, 2007. URL: <https://ui.adsabs.harvard.edu/abs/2007arXiv0706.3639L>, arXiv:0706.3639.
- [LKM<sup>+</sup>15] Brian Lenardo, Kareem Kazkaz, Aaron Manalaysay, Jeremy Mock, Matthew Szydakis, and Mani Tripathi. A Global Analysis of Light and Charge Yields in Liquid Xenon. *IEEE Trans. Nucl. Sci.*, 62(6):3387–3396, 2015. arXiv:1412.4417, doi:10.1109/TNS.2015.2481322.
- [Llo82] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982. doi:10.1109/TIT.1982.1056489.

- [LNST63] J Lindhard, V Nielsen, M Scharff, and P.V Thomsen. INTEGRAL EQUATIONS GOVERNING RADIATION EFFECTS. (NOTES ON ATOMIC COLLISIONS, III). *Kgl. Danske Videnskab., Selskab. Mat. Fys. Medd.*, 33(10), 1963.
- [LRK<sup>+</sup>06] R. Lemrani, M. Robinson, V. A. Kudryavtsev, M. De Jesus, G. Gerbier, and N. J. C. Spooner. Low energy neutron propagation in MCNPX and GEANT4. *Nucl. Instrum. Meth.*, A560:454–459, 2006. doi:10.1016/j.nima.2005.12.238.
- [LS61] J. Lindhard and M. Scharff. Energy Dissipation by Ions in the keV Region. *Phys. Rev.*, 124:128–130, 1961. URL: <https://link.aps.org/doi/10.1103/PhysRev.124.128>, doi:{10.1103/PhysRev.124.128}.
- [LS96] J.D. Lewin and P.F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astropart. Phys.*, 6:87–112, 1996. doi:10.1016/S0927-6505(96)00047-3.
- [LSS63] J. Lindhard, M. Scharff, and H.E. Schiøtt. Range concepts and heavy ion ranges (notes on atomic collisions, ii). *Kgl. Danske Videnskab. Selskab. Mat. Fys. Medd.*, 33(14), 1963. OSTI ID: 4153115.
- [LTZ08] F. T. Liu, K. M. Ting, and Z. Zhou. Isolation Forest. In , pages 413–422, 2008. doi:10.1109/ICDM.2008.17.
- [LZ 16] LZ Collaboraion. The LZap Reconstruction Package. unpublished, 2016.
- [LZ 20] LZ Collaboration. Projected sensitivity of the LUX-ZEPLIN experiment to search for the two-neutrino and neutrinoless double beta decay of <sup>134</sup>Xe. (in preparation), 2020.
- [M<sup>+</sup>17] B. J. Mount et al. LUX-ZEPLIN (LZ) Technical Design Report. *arXiv:1703.09144 (physics.ins-det)*, 2017. arXiv:1703.09144.
- [M<sup>+</sup>20] J.-P. Macquart et al. A census of baryons in the Universe from localized fast radio bursts. *Nature*, 581(7809):391–395, 2020. arXiv:2005.13161, doi:10.1038/s41586-020-2300-2.
- [MA<sup>+</sup>16] J. Martín-Albo et al. Sensitivity of NEXT-100 to Neutrinoless Double Beta Decay. *JHEP*, 05:159, 2016. arXiv:1511.09246, doi:10.1007/JHEP05(2016)159.
- [Mac17] Carla Macolino. Search for neutrinoless double-beta decay with the SuperNEMO demonstrator. *PoS, EPS-HEP2017:121*, 2017. doi:10.22323/1.314.0121.
- [Mal14] D.C. Malling. *Measurement and Analysis of WIMP Detection Backgrounds, and Characterization and Performance of the Large Underground Xenon Dark Matter Search Experiment*. PhD thesis, Physics Department at Brown University, Providence, Rhode Island, May 2014.

- [Mar18] David J.E. Marsh. Axions and ALPs: a very short introduction. In *13th Patras Workshop on Axions, WIMPs and WISPs*, pages 59–74, 2018. arXiv:1712.03018, doi:10.3204/DESY-PROC-2017-02/marsh\_david.
- [ME13] M. T. Mustonen and J. Engel. Large-scale calculations of the double- $\beta$  decay of  $^{76}\text{Ge}$ ,  $^{130}\text{Te}$ ,  $^{136}\text{Xe}$ , and  $^{150}\text{Nd}$  in the deformed self-consistent skyrme quasiparticle random-phase approximation. *Phys. Rev. C*, 87:064302, Jun 2013. doi:10.1103/PhysRevC.87.064302.
- [Mel99] Yannick Mellier. Probing the universe with weak lensing. *Ann. Rev. Astron. Astrophys.*, 37:127–189, 1999. arXiv:astro-ph/9812172, doi:{10.1146/annurev.astro.37.1.127}.
- [Met20] Cade Metz. *How Google’s AI Viewed the Move No Human Could Understand*, 2016 (last accessed June 1, 2020). [www.wired.com/2016/03/googles-ai-viewed-move-no-human-understand/](http://www.wired.com/2016/03/googles-ai-viewed-move-no-human-understand/).
- [Mey08] H.O. Meyer. Dark Rate of a Photomultiplier at Cryogenic Temperatures. *arXiv e-prints*, 5 2008. arXiv:0805.0771.
- [MHPK01] A. P. Meshik, C. M. Hohenberg, O. V. Pravdivtseva, and Ya. S. Kapusta. Weak decay of Ba-130 and Ba-132: Geochemical measurements. *Phys. Rev. C*, 64:035205, 2001. doi:10.1103/PhysRevC.64.035205.
- [Mit97] Tom M. Mitchell. *Machine Learning*. McGraw-Hill Science/Engineering/Math, Ohio, U.S.A, 1st edition, 1997. ISBN-13: 978-0070428072.
- [MK07] Geoffrey J McLachlan and Thriyambakam Krishnan. *The EM algorithm and extensions*, volume 382. John Wiley & Sons, 2007. doi:{10.1002/9780470191613}.
- [MNS62] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962. doi:10.1143/PTP.28.870.
- [MNT17] Eric M. Metodiev, Benjamin Nachman, and Jesse Thaler. Classification without labels: Learning from mixed samples in high energy physics. *JHEP*, 10:174, 2017. arXiv:1708.02949, doi:10.1007/JHEP10(2017)174.
- [Mou17] Xavier Mougeot. Betashape: A new code for improved analytical calculations of beta spectra. *EPJ Web of Conferences*, 146:12015, 01 2017. doi:10.1051/epjconf/201714612015.
- [MP03] William J. Marciano and Zohreh Parsa. Neutrino electron scattering theory. *J. Phys. G*, 29:2629–2645, 2003. arXiv:hep-ph/0403168, doi:10.1088/0954-3899/29/11/013.
- [MS85] S.P. Mikheyev and A.Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917, 1985.

- [MS16] M. Maltoni and A.Y. Smirnov. Solar neutrinos and neutrino physics. *Eur. Phys. J.*, A52(4):87, 2016. doi:10.1140/epja/i2016-16087-0.
- [MYSH08] D. M. Mei, Z. B. Yin, L. C. Stonehill, and A. Hime. A Model of Nuclear Recoil Scintillation Efficiency in Noble Liquids. *Astropart. Phys.*, 30:12–17, 2008. doi:10.1016/j.astropartphys.2008.06.001.
- [MZTG10] D.-M. Mei, C. Zhang, K. Thomas, and F. Gray. Early Results on Radioactive Background Characterization for Sanford Laboratory and DUSEL Experiments. *Astropart. Phys.*, 34:33–39, 2010. arXiv:0912.0211, doi:10.1016/j.astropartphys.2010.04.003.
- [N<sup>+</sup>20] Irene Nutini et al. The CUORE Detector and Results. *J. Low Temp. Phys.*, 199(1-2):519–528, 2020. doi:10.1007/s10909-020-02402-9.
- [NAS06] NASA. 1E 0657-56: NASA Finds Direct Proof of Dark Matter. <https://www.chandra.harvard.edu/photo/2006/1e0657/> - last accessed Sep. 1st 2020, 2006.
- [NHW<sup>+</sup>07] K. Ni, R. Hastly, T.M. Wongjirad, L. Kastens, A. Manzur, and D.N. McKinsey. Preparation of neutron-activated xenon for liquid xenon detector calibration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 582(2):569–574, 2007. doi:10.1016/j.nima.2007.08.180.
- [Nil98] Nils J. Nilsson. Artificial intelligence: A new synthesis. In *Artificial Intelligence: A New Synthesis*, pages xix – xxi. Morgan Kaufmann, Oxford, 1998. ISBN 978-1-55860-467-4. doi:10.1016/B978-0-08-049945-1.50004-6.
- [NLM<sup>+</sup>17] F. Neves, A. Lindote, A. Morozov, V. Solovov, C. Silva, P. Bras, J.P. Rodrigues, and M.I. Lopes. Measurement of the absolute reflectance of polytetrafluoroethylene (PTFE) immersed in liquid xenon. *JINST*, 12(01):P01017, 2017. arXiv:1612.07965, doi:10.1088/1748-0221/12/01/P01017.
- [NP92] J. Neyman and E.S. Pearson. *On the Problem of the Most Efficient Tests of Statistical Hypotheses*, volume In: Kotz S., Johnson N.L. (eds) Breakthroughs in Statistics. Springer Series in Statistics (Perspectives in Statistics). Springer, New York, 1992. doi:10.1007/978-1-4612-0919-5\_6.
- [Oor32] J.H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bulletin of the Astronomical Institutes of the Netherlands*, 6:249, 1932.
- [P<sup>+</sup>97] S. Perlmutter et al. Measurements of the cosmological parameters  $\Omega$  and  $\Lambda$  from the first 7 supernovae at  $z \geq 0.35$ . *Astrophys. J.*, 483:565, 1997. arXiv:astro-ph/9608192, doi:10.1086/304265}.
- [P<sup>+</sup>11] F. Pedregosa et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.



- [P<sup>+</sup>14] T. Piffl et al. The RAVE survey: the Galactic escape speed and the mass of the Milky Way. *Astron. Astrophys.*, 562:A91, 2014. URL: <http://adsabs.harvard.edu/abs/2014A%26A...562A..91P>, arXiv:1309.4293, doi:10.1051/0004-6361/201322531.
- [PA<sup>+</sup>15] Planck Collaboration, P.A.R. Ade, et al. Planck 2015 results. XIII. Cosmological parameters. *astro-ph.CO*, February 2015. arXiv:[1502.01589]. arXiv:1502.01582.
- [Pau] Wolfgang Pauli. Pauli letter collection: letter to Lise Meitner. Typed copy. URL: <https://cds.cern.ch/record/83282>.
- [PD<sup>+</sup>15] Nathalie Palanque-Delabrouille et al. Neutrino masses and cosmology with Lyman-alpha forest power spectrum. *JCAP*, 11:011, 2015. arXiv:1506.05976, doi:10.1088/1475-7516/2015/11/011.
- [Pla61] R.L. Platzman. Total ionization in gases by high-energy particles: An appraisal of our understanding. *The International Journal of Applied Radiation and Isotopes*, 10(2):116–127, 1961. URL: <http://www.sciencedirect.com/science/article/pii/0020708X61901089>, doi:10.1016/0020-708X(61)90108-9.
- [PMG97] D. Poole, A. Mackworth, and R. Goebel. *Computational Intelligence: A Logical Approach*. Oxford University Press, Inc., USA, 1997. ISBN:0195102703.
- [Pon58] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958.
- [Pon91] B. Pontecorvo. Inverse beta process. *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.*, 1:25–31, 1991.
- [PQ77a] R. D. Peccei and Helen R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D*, 16:1791–1797, Sep 1977. URL: <http://link.aps.org/doi/10.1103/PhysRevD.16.1791>, doi:{10.1103/PhysRevD.16.1791}.
- [PQ77b] R. D. Peccei and Helen R. Quinn. CP Conservation in the Presence of Pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.38.1440>, doi:{10.1103/PhysRevLett.38.1440}.
- [PSY19] P. Pirinen, J. Suhonen, and E. Ydrefors. Charged-current neutrino-nucleus scattering off Xe isotopes. *Phys. Rev. C*, 99:014320, 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevC.99.014320>, doi:10.1103/PhysRevC.99.014320.
- [PTZ00] O. A. Ponkratenko, V. I. Tretyak, and Yu. G. Zdesenko. Event generator DECAY4 for simulating double-beta processes and decays of radioactive nuclei. *Physics of Atomic Nuclei*, 63(7):1282–1287, Jul 2000. doi:10.1134/1.855784.
- [PW65] A.A. Penzias and R.W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *"Astrophys. J."*, 142:419–421, 1965. URL: <http://adsabs.harvard.edu/abs/1965ApJ...142..419P>, doi:10.1086/148307.

- [QV15] X. Qian and P. Vogel. Neutrino Mass Hierarchy. *Prog. Part. Nucl. Phys.*, 83:1–30, 2015. [arXiv:1505.01891](#), [doi:10.1016/j.pnpnp.2015.05.002](#).
- [R<sup>+</sup>10] Beth A. Reid et al. Cosmological constraints from the clustering of the Sloan Digital Sky Survey DR7 luminous red galaxies. *Mon. Not. Roy. Astron. Soc.*, 404(1):60–85, 2010. [doi:{10.1111/j.1365-2966.2010.16276.x}](#).
- [R<sup>+</sup>16] E. Richard et al. Measurements of the atmospheric neutrino flux by Super-Kamiokande: energy spectra, geomagnetic effects, and solar modulation. *Phys. Rev. D*, 94(5):052001, 2016. [arXiv:1510.08127](#), [doi:10.1103/PhysRevD.94.052001](#).
- [RBFFS14] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari. Complementarity of dark matter detectors in light of the neutrino background. *Phys. Rev. D*, 90(8):083510, 2014. [arXiv:1408.3581](#), [doi:10.1103/PhysRevD.90.083510](#).
- [RC97] F. Reines and C. Cowan. The Reines-Cowan experiments: Detecting the Poltergeist. *Los Alamos Sci.*, 25:4–27, 1997.
- [RCJ<sup>+</sup>65] F. Reines, M.F. Crouch, T.L. Jenkins, W.R. Kropp, H.S. Gurr, G.R. Smith, J.P.F. Sellschop, and B. Meyer. Evidence for high-energy cosmic ray neutrino interactions. *Phys. Rev. Lett.*, 15:429–433, 1965. [doi:10.1103/PhysRevLett.15.429](#).
- [RCY<sup>+</sup>19] Adam G. Riess, Stefano Casertano, Wenlong Yuan, Lucas M. Macri, and Dan Scolnic. Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond  $\Lambda$ CDM. *Astrophys. J.*, 876(1):85, 2019. [arXiv:1903.07603](#), [doi:10.3847/1538-4357/ab1422](#).
- [Rei96] F. Reines. The neutrino: From poltergeist to particle. *Rev. Mod. Phys.*, 68:317–327, 1996. [doi:10.1103/RevModPhys.68.317](#).
- [RLC05] Wolfgang A. Rolke, Angel M. Lopez, and Jan Conrad. Limits and confidence intervals in the presence of nuisance parameters. *Nucl. Instrum. Meth. A*, 551:493–503, 2005. [arXiv:physics/0403059](#), [doi:10.1016/j.nima.2005.05.068](#).
- [RN95] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, 3rd edition, 1995.
- [Rob35] H.P. Robertson. Kinematics and World-Structure. *Astrophys. J.*, 82(12):284, 1935. URL: <http://adsabs.harvard.edu/abs/1935ApJ...82..284R>, [doi:10.1086/143681](#).
- [Rod57] A. W. Rodgers. Radius Variation and Population Type of Cepheid Variables. *Mon. Not. Roy. Astron. Soc.*, 117(1):85–94, 1957. URL: <http://mnras.oxfordjournals.org/content/117/1/85.abstract>, [doi:10.1093/mnras/117.1.85](#).

- [RT15] Marco Roncadelli and Fabrizio Tavecchio. No axions from the Sun. *Mon. Not. Roy. Astron. Soc.*, 450(1):L26–L28, 2015. arXiv:1411.3297, doi:{10.1093/mnrasl/s1v040}.
- [RWMM07] M. Redshaw, E. Wingfield, J. McDaniel, and E. Myers. Mass and Double-Beta-Decay  $Q$ -Value of  $^{136}\text{Xe}$ . *Phys. Rev. Lett.*, 98:053003, 2007. doi:{10.1103/PhysRevLett.98.053003}.
- [Ryd16] B. Ryden. *Introduction to cosmology*. Cambridge University Press, 11 2016.
- [Ryd17] B. Ryden. *Introduction to Cosmology*. Cambridge University Press, 2017. URL: <https://books.google.pt/books?id=07WSDQAAQBAJ>.
- [RYZ<sup>+</sup>05] B.P. Roe, H.J. Yang, J. Zhu, Y. Liu, I. Stancu, and G. McGregor. Boosted decision trees, an alternative to artificial neural networks. *Nucl. Instrum. Meth.*, A543(2-3):577–584, 2005. doi:{10.1016/j.nima.2004.12.018}.
- [S<sup>+</sup>10] J. Silk et al. *Particle Dark Matter: Observations, Models and Searches*. Cambridge Univ. Press, Cambridge, 2010. doi:10.1017/CB09780511770739.
- [S<sup>+</sup>11] V.N. Solovov et al. Position Reconstruction in a Dual Phase Xenon Scintillation Detector. *accepted for publication in IEEE Trans. on Nuc. Sci*, 2011. doi:10.1109/TNS.2012.2221742.
- [S<sup>+</sup>12] N. Suzuki et al. The Hubble Space Telescope Cluster Supernova Survey. V. Improving the Dark-energy Constraints above  $z \lesssim 1$  and Building an Early-type-hosted Supernova Sample. *Astrophys. J.*, 746(1):85, 2012. URL: <http://stacks.iop.org/0004-637X/746/i=1/a=85>, doi:10.1088/0004-637X/746/1/85.
- [S<sup>+</sup>19] Albert M Sirunyan et al. Search for dark matter produced in association with a single top quark or a top quark pair in proton-proton collisions at  $\sqrt{s} = 13$  TeV. *JHEP*, 03:141, 2019. arXiv:1901.01553, doi:10.1007/JHEP03(2019)141.
- [S<sup>+</sup>20] Albert M Sirunyan et al. Search for dark matter particles produced in association with a Higgs boson in proton-proton collisions at  $\sqrt{s} = 13$  TeV. *JHEP*, 03:025, 2020. arXiv:1908.01713, doi:10.1007/JHEP03(2020)025.
- [Saa13] R. Saakyan. Two-Neutrino Double-Beta Decay. *Ann. Rev. Nucl. Part. Sci.*, 63:503–529, 2013. doi:10.1146/annurev-nucl-102711-094904.
- [Sak91] A.D. Sakharov. Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe. *Sov. Phys. Usp.*, 34(5):392–393, 1991. doi:10.1070/PU1991v034n05ABEH002497.
- [Sam59] A. Samuel. Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3:210–229, 1959.

- [SB18] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, Cambridge, Massachusetts, 2nd edition, 2018. ISBN-13: 978-0262193986.
- [SBK<sup>+</sup>11] M. Szydagis, N. Barry, K. Kazkaz, J. Mock, D. Stolp, M. Sweany, M. Tripathi, S. Uvarov, N. Walsh, and M. Woods. NEST: a comprehensive model for scintillation yield in liquid xenon. *Journal of Instrumentation*, 6(10):P10002, 2011. doi:10.1088/1748-0221/6/10/P10002.
- [SBZH07] C. Strobl, A-L. Boulesteix, A. Zeileis, and T. Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8, 2007. doi:10.1186/1471-2105-8-25.
- [Sca06] R. Scarpa. Modified Newtonian Dynamics, an Introductory Review. In E.J. Lerner and J.B. Almeida, editors, *First Crisis in Cosmology Conference*, volume 822 of *American Institute of Physics Conference Series*, pages 253–265, 2006. URL: <http://adsabs.harvard.edu/abs/2006AIPC...822..253S>, doi:10.1063/1.2189141.
- [Sch90] Robert E. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990. doi:10.1007/BF00116037.
- [Sch99] Robert E. Schapire. A Brief Introduction to Boosting. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'99*, page 1401–1406, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
- [Sch11] Richard W. Schnee. Introduction to dark matter experiments. In *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small*, pages 775–829, 2011. arXiv:1101.5205, doi:10.1142/9789814327183\_0014.
- [Sch15] Peter Schneider. *Extragalactic Astronomy and Cosmology: An Introduction*. Springer-Verlag Berlin Heidelberg, 2 edition, 2015. doi:10.1007/978-3-642-54083-7.
- [SD11] Peter Sorensen and Carl Eric Dahl. Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter. *Phys. Rev. D*, 83:063501, 2011. arXiv:1101.6080, doi:10.1103/PhysRevD.83.063501.
- [SDF02] Fedor Simkovic, Pavol Domin, and Amand Faessler. Neutrinoless double beta decay of Xe-134. *arXiv hep-ph/0204278*, 4 2002. URL: <https://arxiv.org/abs/hep-ph/0204278>, arXiv:hep-ph/0204278.
- [SFTT13] Matthew Szydagis, Adalyn Fyhrie, Daniel Thorngren, and Mani Tripathi. Enhancement of NEST Capabilities for Simulating Low-Energy Recoils in Liquid Xenon. *JINST*, 8:C10003, 2013. arXiv:1307.6601, doi:10.1088/1748-0221/8/10/C10003.

- [SGN<sup>+</sup>19] N. Solomey, C. Gimar, A. Nelsen, M. L. Buchele, and H. Meyer. Astrophysics and Technical Study of a Solar Neutrino Spacecraft NASA 2018 NIAC Phase-1 Grant Final Report. 2019.
- [SIK<sup>+</sup>13] Anse Slosar, Vid Irsic, David Kirkby, Stephen Bailey, Nicolas G. Busca, et al. Measurement of Baryon Acoustic Oscillations in the Lyman-alpha Forest Fluctuations in BOSS Data Release 9. *J. Cosmology and Astrop. Phys.*, 1304:026, 2013. arXiv:1301.3459, doi:10.1088/1475-7516/2013/04/026.
- [SK18] P. Sorensen and K. Kamdin. Two distinct components of the delayed single electron noise in liquid xenon emission detectors. *JINST*, 13(02):P02032, 2018. arXiv:1711.07025, doi:10.1088/1748-0221/13/02/P02032.
- [SNB05] Vikas Sindhvani, Partha Niyogi, and Mikhail Belkin. Beyond the point cloud: From transductive to semi-supervised learning. In , pages 824–831, 2005. doi: {10.1145/1102351.1102455}.
- [Sol20] V. Solovov. Effects of S2 saturation on  $xy$  position reconstruction of  $0\nu\beta\beta$  decay events with Mercury. Private communication, Oct. 2020.
- [SPVF99] F. Simkovic, G. Pantis, J.D. Vergados, and Amand Faessler. Additional nucleon current contributions to neutrinoless double beta decay. *Phys. Rev. C*, 60:055502, 1999. arXiv:hep-ph/9905509, doi:10.1103/PhysRevC.60.055502.
- [SSF96] J. Schwieger, F. Simkovic, and Amand Faessler. The Pauli principle, QRPA and the two neutrino double beta decay. *Nucl. Phys. A*, 600:179–192, 1996. arXiv:nucl-th/9602032, doi:10.1016/0375-9474(96)00024-3.
- [SV82] J. Schechter and J.W.F. Valle. Neutrinoless Double beta Decay in SU(2) x U(1) Theories. *Phys. Rev. D*, 25:2951, 1982. doi:10.1103/PhysRevD.25.2951.
- [SZ70] R.A. Sunyaev and Y.B. Zeldovich. Small-Scale Fluctuations of Relic Radiation. *Astrophysics and Space Science*, 7:3–19, 1970. URL: <http://adsabs.harvard.edu/abs/1970Ap%26SS...7....3S>, doi:10.1007/BF00653471.
- [T<sup>+</sup>06] Max Tegmark et al. Cosmological Constraints from the SDSS Luminous Red Galaxies. *Phys. Rev.*, D74:123507, 2006. arXiv:astro-ph/0608632, doi:{10.1103/PhysRevD.74.123507}.
- [T<sup>+</sup>07] P. Tisserand et al. Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *Astron. Astrophys.*, 469:387–404, 2007. arXiv:astro-ph/0607207, doi:{10.1051/0004-6361:20066017}.
- [T<sup>+</sup>18] M. Tanabashi et al. Review of Particle Physics. *Phys. Rev. D*, 98(3):030001, 2018. doi:10.1103/PhysRevD.98.030001.
- [Tao19] Sean Tao. Deep Neural Network Ensembles. *arXiv e-prints*, 2019. URL: <https://ui.adsabs.harvard.edu/abs/2019arXiv190405488T>, arXiv:1904.05488.

- [TEC91] A.L. Turkevich, T.E. Economou, and G.A. Cowan. Double beta decay of U-238. *Phys. Rev. Lett.*, 67:3211–3214, 1991. doi:10.1103/PhysRevLett.67.3211.
- [Tin95] Tin Kam Ho. Random decision forests. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 1, pages 278–282, 1995. doi:{10.1109/ICDAR.1995.598994}.
- [TRK10] V. Tomasello, M. Robinson, and V.A. Kudryavtsev. Radioactive background in a cryogenic dark matter experiment. *Astroparticle Physics*, 34(2):70 – 79, 2010. doi:10.1016/j.astropartphys.2010.05.005.
- [Tur50] Alan M. Turing. Computing Machinery and Intelligence. *Mind*, 59:433–60, 1950. doi:{10.1093/mind/LIX.236.433}.
- [TZMC17] A. Tiwari, C. Zhang, D. M. Mei, and P. Cushman. Observation of annual modulation induced by  $\gamma$  rays from  $(\alpha, \gamma)$  reactions at the Soudan Underground Laboratory. *Phys. Rev. C*, 96(4):044609, 2017. [Erratum: *Phys. Rev. C*98,no.1,019901(2018)]. doi:{10.1103/PhysRevC.96.044609, 10.1103/PhysRevC.98.019901}.
- [vBBS85] T.S. van Albada, J.N. Bahcall, K. Begeman, and R. Sancisi. Distribution of dark matter in the spiral galaxy NGC 3198. *Astrophys. J.*, 295:305–313, August 1985. URL: <http://adsabs.harvard.edu/abs/1985ApJ...295..305V>, doi: 10.1086/163375.
- [Vis17] Francesco Vissani. Joint analysis of Borexino and SNO solar neutrino data and reconstruction of the survival probability. *Nucl. Phys. Atom. Energy*, 18(4):303–312, 2017. arXiv:1709.05813, doi:10.15407/jnpae2017.04.303.
- [Viv10] L. Viveiros. *Optimization of Signal versus Background in Liquid Xe Detectors Used for Dark Matter Direct Detection Experiments*. PhD thesis, Physics department at Brown University, Providence, Rhode Island, May 2010.
- [Vos20] P. Voss. *Essentials of general intelligence: The direct path to AGI*, 2002 (last accessed June 1, 2020). [www.kurzweilai.net/essentials-of-general-intelligence-the-direct-path-to-agi](http://www.kurzweilai.net/essentials-of-general-intelligence-the-direct-path-to-agi).
- [VRE13] Nuria López Vaquero, Tomás R. Rodríguez, and J. Luis Egido. Shape and pairing fluctuation effects on neutrinoless double beta decay nuclear matrix elements. *Phys. Rev. Lett.*, 111:142501, Sep 2013. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.142501>, doi:10.1103/PhysRevLett.111.142501.
- [W<sup>+</sup>20] Qiuhong Wang et al. Results of Dark Matter Search using the Full PandaX-II Exposure. 7 2020. arXiv:2007.15469.
- [WHC<sup>+</sup>15] Wenting Wang, Jiaxin Han, Andrew P. Cooper, Shaun Cole, Carlos Frenk, and Ben Lowing. Estimating the dark matter halo mass of our Milky Way using dynamical tracers. *Astro-ph.GA*, July 2015. arXiv:1502.03477. URL: <http://arxiv.org/pdf/1502.03477v2.pdf>.

- [Wil15] Jeanne R. Wilson. An Experimental Review of Solar Neutrinos. In *Topical Research Meeting on Prospects in Neutrino Physics (NuPhys2014) London, UK, December 15-17, 2014*, 2015. URL: [inspirehep.net/record/1360284/files/arXiv:1504.04281.pdf](https://inspirehep.net/record/1360284/files/arXiv:1504.04281.pdf).
- [Wol78] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys. Rev. D*, 17:2369–2374, 1978. doi:10.1103/PhysRevD.17.2369.
- [Woo18] D. Woodward. *Simulations of cosmic muons and background radiations for muon tomography and underground experiments*. PhD thesis, University of Sheffield, Department of Physics & Astronomy at University of Sheffield, Sheffield, United Kingdom, March 2018.
- [WQY<sup>+</sup>15] Yiwen Wen, Huilin Qu, Daneng Yang, Qi-shu Yan, Qiang Li, and Yajun Mao. Probing triple-W production and anomalous WWWW coupling at the CERN LHC and future  $\mathcal{O}(100)$  TeV proton-proton collider. *JHEP*, 03:025, 2015. doi:{10.1007/JHEP03(2015)025}.
- [WX16] Yifang Wang and Zhi-zhong Xing. Neutrino Masses and Flavor Oscillations. *Adv. Ser. Direct. High Energy Phys.*, 26:371–395, 2016. arXiv:1504.06155, doi:10.1142/9789814733519\_0019.
- [X<sup>+</sup>19] Jingkai Xia et al. PandaX-II Constraints on Spin-Dependent WIMP-Nucleon Effective Interactions. *Phys. Lett. B*, 792:193–198, 2019. arXiv:1807.01936, doi:10.1016/j.physletb.2019.02.043.
- [Z<sup>+</sup>20] P.A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020. doi:10.1093/ptep/ptaa104.
- [ZRZH06] Ji Zhu, Saharon Rosset, Hui Zou, and Trevor Hastie. Multi-class AdaBoost. *Statistics and its interface*, 2, 2006. doi:{10.4310/SII.2009.v2.n3.a8}.
- [Zub03] K. Zuber. *Neutrino Physics*. Series in High Energy Physics, Cosmology and Gravitation. CRC Press, 2003. URL: <https://books.google.pt/books?id=mvLKBQAAQBAJ>.
- [Zwi33] F. Zwicky. Die Rotshiebung von extragalaktischen Nebeln. *Helv. Phys. Acta*, 6:110–127, 1933.
- [Zwi37] F. Zwicky. On the Masses of Nebulae and of Clusters of Nebulae. *Astrophys. J.*, 86:217, October 1937. URL: <http://adsabs.harvard.edu/abs/1937ApJ...86..217Z>, doi:10.1086/143864.
- [ZWT02] Zhi-Hua Zhou, Jianxin Wu, and Wei Tang. Ensembling neural networks: Many could be better than all. *Artificial Intelligence*, 137(1):239 – 263, 2002. doi:10.1016/S0004-3702(02)00190-X.





# List of Figures

- 2.1 NGC3198 rotation curve, obtained by measuring the redshift of the 21 cm emission line of neutral hydrogen gas. The upper curve represents the best fit of the velocities to the data. The curves marked with “halo” and “disk” represent the mass density contributions of the dark matter halo and galactic disk (stars, gas and dust) obtained from this fit, respectively. Picture from [vBBS85]. . . . . 6
- 2.2 Combined data of type Ia supernovae from several surveys. The black line indicates the best fit to the cosmological parameters of the  $\Lambda$ CDM model. Figure from [S<sup>+</sup>12]. . . . . 8
- 2.3 Cosmic recipe of our universe according to the latest results from Planck [A<sup>+</sup>18a]. 9
- 2.4 Mollweide projection of the CMB map measured by the Planck satellite (2015 data), with monopole (mean temperature) and dipole (Doppler shift) contributions removed. Figure from [Col15]. . . . . 11
- 2.5 Illustrative sensitivity of the acoustic peaks of the CMB theoretical power spectrum to the baryon energy density  $\Omega_b h^2$  (left), total matter energy density  $\Omega_m h^2$  (center) and dark energy density  $\Omega_\Lambda h^2$  (right), as predicted by the  $\Lambda$ CDM model. The default cosmological parameters considered here are  $\Omega_k = 0$ ,  $\Omega_\Lambda = 0.65$ ,  $\Omega_b h^2 = 0.02$  and  $\Omega_m h^2 = 0.147$ , not the current best estimates. Figure adapted from Reference [HD02]. . . . . 12
- 2.6 Planck temperature power spectrum of the CMB, obtained using multipolar analysis over the CMB map. The blue line represents the  $\Lambda$ CDM model theoretical spectrum prediction that best fits the data. The first acoustic peak is clearly visible at multipole  $l \sim 200$ , followed by the secondary peaks and damping tail, at higher  $l$ . Figure from Reference [A<sup>+</sup>18a]. . . . . 13

- 2.7 Evolution of an acoustic wave in the primordial plasma (upper-left), showing the different species present and their behaviour (mass distribution) over different moments: baryons (named “gas”, blue), dark matter (black), photons (red), and neutrinos (green). The upper-right panel represents the decoupling of photons and baryons, where photons disperse and the baryonic peak stalls. Dark matter and baryonic matter then rearrange due to gravity, resulting in the merging of the respective curves as portrayed in the lower-left figure [ESW07]. The lower-right figure shows the baryon acoustic peak measured from the Lyman- $\alpha$  forest [SIK<sup>+</sup>13]. The peak in the correlation function is clearly visible at length scale  $r = 100 - 130 h^{-1}$  Mpc. . . . . 14
- 2.8 Predicted 95% CL bands of the primordial abundances of  ${}^4\text{He}$  (top pink), D (center blue),  ${}^3\text{He}$  (center red), and  ${}^7\text{Li}$  (bottom green) as a function of the baryon-to-photon ratio  $\eta$ . The boxes over each band indicate the observed range of values for the light element abundances. The vertical blue dashed band represents the 95% CL measurement of the baryon density from the CMB, and the vertical red dashed band indicates the 95% CL constraints on the baryon content from BBN considering the observed abundances of D and  ${}^4\text{He}$  only. Figure from [T<sup>+</sup>18]. . . 16
- 2.9 Comparison of the measured distribution of large-scale structures of galaxies from the Sloan Digital Sky Survey (left) and the Millennium-II simulations for a redshift of  $z = 6.2$  (top) and today (bottom), assuming a cosmological model with  $\Omega_m = 0.25$  and  $\Omega_\Lambda = 0.75$ . The side length of the Millennium-II boxes is  $100 \text{ Mpc } h^{-1}$  in comoving units. The same distance scale is represented in the left image by the red slice centred on the observer. . . . . 17
- 2.10 Image of the Bullet Cluster where the contributions of the intergalactic gas obtained from its X-ray emission (pink) and the majority of the matter in the cluster inferred from weak gravitational lensing (blue) are represented. Credit to: X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al [NAS06]. . . . . 18
- 2.11 Combined constraints of the cosmological parameters  $\Omega_m$  and  $\Omega_\Lambda$  from type Ia Supernovae, CMB and BAO from cosmological observations prior to 2012. Despite the outdated results displayed in the image, the same agreement on  $\Omega_m$ ,  $\Omega_\Lambda$  and curvature can be found in the most recent data. The constrains from CMB and BAO results have improved significantly due to the Planck measurements [A<sup>+</sup>18a]. The region marked with “No Big Bang” means that for that set of cosmological parameters the universe starts in a contracting state, resulting in a bouncing universe [CPT92]. Figure from [S<sup>+</sup>12]. . . . . 20
- 2.12 Limits on the spin-independent WIMP-nucleon (left) and spin-dependent WIMP-proton (right) scattering cross sections from the ATLAS experiment [A<sup>+</sup>19b]. . . 25

2.13 Projected sensitivity on the spin-independent WIMP-nucleon (left) and spin-dependent WIMP-neutron (right) scattering cross sections of LZ [A<sup>+</sup>20d]. Solid lines represent the most recent limits from direct detection experiments [A<sup>+</sup>17h, W<sup>+</sup>20, A<sup>+</sup>18i, A<sup>+</sup>17g, A<sup>+</sup>19f, A<sup>+</sup>19h, X<sup>+</sup>19]. . . . . 27

3.1 Detection of inverse  $\beta$ -decay on hydrogen (proton) in the reactor-neutrino experiment by Cowan and Reines [CRH<sup>+</sup>56]. The signature of this interaction is the delayed ( $\sim$ ms) observation of  $\gamma$ -rays from neutron capture that follows the detection of characteristic  $\gamma$ -rays from  $e^+$  annihilation. Figure from [RC97]. . . . . 35

3.2 The Solar neutrino flux as a function of the neutrino energy, highlighting the contributions from different nuclear fusion reactions [SGN<sup>+</sup>19]. . . . . 37

3.3 Solar  $^8\text{B}$  neutrino fluxes, showing the measured muon and tau neutrino flux ( $\phi(\nu_{\mu\tau})$ ) obtained from neutral current interactions in SNO versus the combined results from SNO and Super-Kamiokande for the electron neutrino flux ( $\phi(\nu_e)$ ) obtained from charged current interactions ( $\nu_e e$  elastic scattering). The diagonal bands represent the total flux predicted by the BPB01 solar model [BPB01] (dashed) and the total flux measured by SNO and Super-Kamiokande (full). An updated figure showing the most recent combined results can be found in Reference [Z<sup>+</sup>20]. Figure from Reference [A<sup>+</sup>01]. . . . . 38

3.4 The up-to-down (U-D)/(U+D) asymmetry as a function of momentum for single ring fully contained (FC) and partly-contained (PC) e-like (top) and  $\mu$ -like (bottom) events in Super-Kamiokande. FC events have both the neutrino interaction vertex and resulting particle tracks entirely within the fiducial volume, while PC events only have the former. The hatched bands represent the expected asymmetry under the assumption of no oscillations. The red dashed line on the bottom plot represents the expected asymmetry for  $\mu \leftrightarrow \tau$  oscillations considering maximal mixing. Figure from Reference [F<sup>+</sup>98a]. . . . . 39

3.5 Graphic representation of the normal (top) and inverted (bottom) neutrino mass hierarchies. Oscillation experiments are only sensitive to the differences of squared-masses  $\Delta m_{21}^2 = m_2^2 - m_1^2$  (from solar neutrino oscillations) and  $\Delta m_{32}^2 = m_3^2 - m_2^2$  (from atmospheric neutrino oscillations). Since  $\Delta m_{21}^2 \ll |\Delta m_{32}^2|$ , the squared-mass difference  $\Delta m_{32}^2$  can be redefined as  $\Delta m_{32}^2 \approx m_3^2 - (m_1^2 + m_2^2)/2$ . The sign of  $\Delta m_{32}^2$  is unknown, leading to the two possible mass hierarchies. The colors associated to each mass eigenstate ( $\nu_1, \nu_2, \nu_3$ ) represent the mixing of flavour eigenstates ( $\nu_e, \nu_\mu, \nu_\tau$ ), *i.e.*, probability of finding one of the flavour eigenstates if the neutrino is in a certain mass eigenstate. Figure adapted from Reference [DMVV16]. . . . . 43

- 3.6 Masses of the SM fermions. The neutrino masses presented consider the normal hierarchy scenario. The six orders of magnitude gap between the mass scale of neutrinos and the remaining fermions could indicate a new mass scale and a different mass generating mechanism for neutrinos. . . . . 47
- 3.7 Nuclear mass as a function of the atomic number  $Z$  for a  $2\nu\beta\beta$  decay candidate with  $A$  even (left) and  $A$  odd (right). For the even-even isotope (lower mass parabola on the left plot) the single  $\beta$ -decay is kinematically forbidden but the  $2\nu\beta\beta$  decay is not. For odd mass number nuclei (mass parabola on the right plot) the  $2\nu\beta\beta$  decay is strongly suppressed as the single  $\beta$ -decay would be dominant. Figure from Reference [DMVV16]. . . . . 51
- 3.8 First order diagram of the  $0\nu\beta\beta$  decay process mediated by the exchange of a light massive Majorana neutrino  $\nu_M$ . Figure from Reference [DMVV16]. . . . . 54
- 3.9 Predicted probability density distribution of  $\langle m_{\beta\beta} \rangle$  as a function of the lightest neutrino mass,  $m_l$ , for the normal hierarchy (left) and inverted hierarchy (right) assuming a flat prior. Solid contours indicate the allowed parameter phase-space, considering the  $3\sigma$  intervals of the neutrino oscillation parameters and assuming QRPA NMEs. Figure from Reference [ABD17]. . . . . 55
- 3.10 Illustration of the spectra of the summed electron kinetic energies ( $K_e$ ) scaled with the total energy of the decay ( $Q$ ) for the  $2\nu\beta\beta$  decay (dotted line) and  $0\nu\beta\beta$  decay (solid line) convolved with an energy resolution of 5% FWHM. The area of the curves is arbitrarily set for illustration purposes and does not represent the expected relative rates for these two decay modes. Figure from Reference [EV02]. 56
- 4.1 Schematic representation of the LZ detector. The active xenon TPC in the center (1) is surrounded by the acrylic tanks of the outer detector (2) that are observed by a cylindrical array of 120 8-inch PMTs (3). The full detector is within a water tank (4). The high-voltage umbilical (5) and the neutron source tube (6) protrude the instrumented volume of the outer detector from the sides. The main liquid xenon circulation conduit that connects the xenon purification system to the TPC is adjacent to the water tank (7). A detailed view of the innermost region of the detector is presented on the inlay in the right [A<sup>+</sup>20d, A<sup>+</sup>20g]. . . . . 61

- 4.2 Schematic representation of an interaction in the liquid xenon volume of a dual-phase xenon TPC. The energy deposited by a particle in the liquid xenon volume excites and ionizes the xenon medium, releasing scintillation photons and ionization electrons. The scintillation is promptly detected in the top and bottom PMT arrays as the S1 signal, while the electrons are drifted from the interaction site by an applied electric field to the liquid-gas interface, where a stronger electric field (extraction field) produced between the gate and anode grids extracts them into the gas phase, producing the strong electroluminescence light corresponding to the S2 signal. The time difference between the S1 and S2 signals is used to determine the depth of the interaction, while the S2 light pattern in the top array can be used to reconstruct the XY position of the interaction [M<sup>+</sup>17, A<sup>+</sup>20g]. . . . . 65
- 4.3 Electron recoil (top) and nuclear recoil (bottom) bands from LUX calibration data. The ER band was obtained from tritium calibration data, while the NR band was measured using a monoenergetic neutron source. The solid lines indicate the band medians, while the dashed lines represent the band contours that contain 80% of the events of the NR (red) and ER (blue) distributions [A<sup>+</sup>16d]. . . . . 69
- 4.4 Left: Schematic cross section of the dome and side region of the TPC and skin detectors, showing the position of the 2-inch PMTs at the bottom of the side skin volume (1) and on the dome looking outward (2). Right: Photograph of the inside walls of the inner cryostat vessel covered with PTFE reflectors. The ring of 20 2-inch PMTs looking upward to the side skin volume are visible at the bottom of the vessel. Images from Reference [A<sup>+</sup>20g]. . . . . 71
- 4.5 Combined simulated spectrum of the main ER backgrounds in the 5.6 tonnes fiducial volume considered in the WIMP search analysis. The different spectra are for single scatter events with no coincident signal in the skin nor the outer detector. The spectra are taken directly from simulation and do not include any detector effects. Figure from Reference [A<sup>+</sup>20d]. . . . . 72
- 5.1 Simulated ionization yield on liquid xenon of an event with two electrons emitted back-to-back in the  $z$  direction, with the initial energy of each electron equal to half the  $Q_{\beta\beta}$  of  $^{136}\text{Xe}$ . The red arrows represent the initial direction of emission of the two electrons. Credit to Andrey Solovov (M.S. at LIP-Coimbra). . . . . 78
- 5.2 Background event rate in the active region and in the  $\pm 1\sigma$  energy ROI as a function of  $r^2$  and  $z$ . The dashed black rectangle represents the inner 967 kg volume where LZ is most sensitive to the  $0\nu\beta\beta$  decay, while the larger dashed white rectangle represents the extended fiducial volume used on the profile likelihood analysis (see Section 5.4). Plots from Reference [A<sup>+</sup>20c]. . . . . 84

- 5.3 Background energy spectrum in the inner 967 kg volume: contributions from the main background sources (left) and effect of successive selection cuts used in the analysis (right). More details regarding the analysis can be found in the text. The gray vertical band around  $Q_{\beta\beta}$  represents the  $\pm 1\sigma$  ROI used for this search. The  $^{60}\text{Co}$ ,  $^{238}\text{U}$ -late chain and  $^{232}\text{Th}$ -late chain backgrounds from the detector components are combined into a single curve on the left plot but are treated independently in the sensitivity analysis. The dashed yellow line in the left plot represents the expected signal spectrum for  $^{136}\text{Xe}$   $0\nu\beta\beta$  decay, considering a half-life of  $1.06 \times 10^{26}$  years (see Section 5.4), and is not included in the total spectrum. The spectra are smeared using the energy resolution function of LUX [A<sup>+</sup>17c], scaled to be 1.0% at  $Q_{\beta\beta}$ . Plots from Reference [A<sup>+</sup>20c]. . . . . 85
- 5.4 Fraction of background events from  $^{208}\text{Tl}$  ( $^{232}\text{Th}$  progeny) from the Davis cavern walls that are not vetoed by the outer detector system (assuming a 100 keV threshold) for different radial positions (left) and depths (right) of the interactions. The effect of the outer detector veto alone is represented in green, while the combination of this veto with the single scatter selection is displayed in red. The combination of both skin and outer detector vetoes and single scatter selection is represented in blue. . . . . 87
- 5.5 Approximated spacial positioning of the cylindrical surfaces used in the boosting stages of the rock gamma simulations. Each stage records the information of any incoming particle that crosses its boundary. The next iteration of the simulation will use the information recorded at the previous stage to propagate the particles with boosted statistics, recording once more the information of the particles that reach the boundary of the next stage. The statistics are boosted by an overall factor of  $10^8$ ,  $100\times$  on each stage. . . . . 89
- 5.6 Spectrum of the energy depositions in the TPC from  $\gamma$ -rays produced in the cavern walls by  $^{238}\text{U}$  and  $^{232}\text{Th}$  decay chains (red and blue, respectively). Both spectra were produced with the same initial statistics. The flux of  $\gamma$ -rays with energies above 3 MeV is non-negligible due to their larger penetrative power. Despite the predominance of the  $^{232}\text{Th}$  spectrum at lower energies, the  $\gamma$ -ray flux above 3 MeV in the TPC is similar for both sources. . . . . 90
- 5.7 Differential background rates for solar  $^8\text{B}$  neutrino-electron elastic scattering in the  $\pm 1\sigma$  ROI (1.0% energy resolution at  $Q_{\beta\beta}$ ), considering a neutrino flux of  $5.79 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$ . Left: mixture of Z and W boson exchange (NC+CC). Right: pure Z boson exchange (NC). . . . . 96

5.8 LZ projected sensitivity to the half-life of the  $0\nu\beta\beta$  decay of  $^{136}\text{Xe}$  as a function of detector live time for a natural abundance of  $^{136}\text{Xe}$  (solid black line) and for a dedicated run with 90% isotopic enrichment (dashed black line). The light green shaded band represents a  $\pm 1\sigma$  statistical uncertainty on the sensitivity. The limits set by EXO-200 [A<sup>+</sup>19g] (orange full) and KamLAND-Zen [G<sup>+</sup>16] (purple full) are also shown, along with the respective projected sensitivities (dashed). Figure from Reference [A<sup>+</sup>20c]. . . . . 101

5.9 90% CL sensitivity to the  $0\nu\beta\beta$  decay half-life of  $^{136}\text{Xe}$  as a function of the energy resolution at  $Q_{\beta\beta}$  (a) and the minimum vertex separation in depth (b), for a 1000 live-days run with LZ. The sensitivity for an energy resolution of 1% and vertical vertex separation of 3 mm, the values assumed in this analysis, are represented by the vertical dashed lines labelled “LZ projected”. For comparison, the projected LZ sensitivity assuming the energy resolution recently measured in XENON1T [A<sup>+</sup>20j] is also indicated on plot (a). Figure from Reference [A<sup>+</sup>20c]. . . . . 102

5.10 LZ projected exclusion sensitivity to  $\langle m_{\beta\beta} \rangle$  and the neutrino mass hierarchy as a function of the lightest neutrino mass. The width of the green band is the result of the uncertainty in the nuclear matrix elements [ME13, VRE13]. The red and blue contours represent the allowed parameter space for the inverted hierarchy and normal hierarchy neutrino mass scenarios, respectively, as well as the respective  $\pm 1\sigma$  contours (light colors) [A<sup>+</sup>18g]. This projection for LZ matches the current best experimental limit from KamLAND-Zen [G<sup>+</sup>16]. The inlay on the right-hand side of the plot displays the current best exclusion limits to  $\langle m_{\beta\beta} \rangle$  for different  $2\nu\beta\beta$  isotopes (solid arrows) and the projected limit of LZ for a dedicated run with 90% enriched  $^{136}\text{Xe}$  (see Section 5.4.2). The single pointed end of the arrows represents the lower-bound exclusion limit, while the length of the arrow represents the respective uncertainty. Figure from Reference [A<sup>+</sup>20c]. . . . . 103

6.1 Core data structure created by LZap processing. . . . . 108

6.2 Schematic of the data flow in LZap. Yellow rectangles represent algorithm objects (modules) and blue parallelograms represent data objects that are stored in the transient event store (TES). The aquamarine dashed rectangles encompassing some objects indicate the data flow for different detectors and their ensemble are a proxy to the *PhotonDetection* processing chain. . . . . 110

- 6.3 Example of an LZ event, taken from simulated data of the detector and electronics response. On the left-hand side, the event timelines for the TPC-HG (top left), Skin (middle left) and OD-HG (bottom left) detectors are shown. Pulses identified by the *PulseClassifier* are color-coded: S1 (green), S2 (blue), SE (red), SPE (yellow), MPE (gray), and Other (also gray). For a description of the pulse classes see Section 6.2.1. On the right-hand side, the spacial representation of the integrated PMT response for the full event is shown. The colors represent the signal intensity in the PMTs of the TPC (circular arrays at the center), OD (top and bottom rectangular arrays) and Skin (radial arrays adjacent to the TPC top and bottom array plus the smaller dome array at the bottom). Figure obtained using the official LZ Offline Event Viewer. . . . . 115
- 6.4 Pulse population distribution on the full phase-space of the pulse area (pA) and pulse length at 90% area (pL90) RQs (left) and top-bottom asymmetry (TBA) and pulse area RQs (right). Some degeneracy between the different populations is present and these 2-dimensional representations are merely a guideline for upcoming results. Not all isolated populations are labelled, only the main populations mentioned in Section 6.2.1 have been highlighted. . . . . 116
- 6.5 Typical waveforms of an S1 pulse (left, green) and S2 pulse (right, blue). These pulses are the ones present in Figure 6.3. The coloured boxes represent the pulse boundaries calculated by the *DoGPulseFinder*. The pulse classifications are obtained with the HADES algorithm (Section 6.2.3). The summed waveforms for both the S1 and S2 pulses in the TPC-HG (top) and TPC-LG (middle) detectors are presented, as well as the waveforms of 99 individual channels in the bottom array (bottom). Plots obtained with the official LZ Offline Event Viewer. . . . . 117
- 6.6 Light patterns in the top and bottom TPC PMT arrays for a typical S1 (left) and S2 (right) pulse – the pulses shown in Figures 6.3 and 6.5. Plots obtained with the official LZ Offline Event Viewer. . . . . 118
- 6.7 Examples of pathological pulses in LZap at the *PulseClassifier* stage. Top-left: 2- and 3-fold coincident dark counts; Top-right: SE splits – *PulseFinder* issue; Middle-left: baselines between pulses with significant fluctuations; Middle right: S2 start ramps – *PulseFinder* issue; Bottom left: S2 tail structures; Bottom right: SE vertical sums (pile-up) and SE horizontal merging (E-train). . . . . 119
- 6.8 Summed PODs of an S1 pulse followed by baseline noise and four afterpulses. The two central afterpulses were overlaying and were merged into one pulse. The afterpulses were classified as S1s by having coincidence  $>1$  due to runaway baselines and pileup in different channels. The inlay shows the 4 channels that produced afterpulsing. Plots obtained with the official LZ Offline Event Viewer . 121



- 6.9 Diagram of classification results (left) grouped by classification impact (positive or negative results) and validity of classification (true or false results) and the respective distribution over a generic confusion matrix (right). The colours represent the severity of the classification results to the ensuing analysis, from red (most severe) to green (not severe). The mixing between S2 pulses and SE pulses is tolerable since both represent the same process and are distinguished by area alone. False positive results are divided in misclassification of Other pulses and “mixing” of true pulses, as the repercussions of each case are different. The case where an SPE pulse is classified as Other bears no significant impact on the physics analysis. . . . . 123
- 6.10 Diagram of the decision tree of the HADES algorithm for pulse classification in LZap. The depicted version is implemented in LZap since version 4.4.0. . . . . 126
- 6.11 Depiction of HADES selections in the phase space of the pulse parameters. The top panels represent the first instance of S1-S2 separation as defined in the third decision node of the HADES diagram (Figure 6.10). The bottom left panel shows the separation of “bad” pulses from the remaining S1-like population. The bottom right panel depicts the separation of SE pulses from the S2 ensemble, showing the area threshold used in LZap for this analysis (red) and the optimized value for this dataset (orange). . . . . 127
- 6.12 COMPACT 2D KDE PDFs for all representations used, showing the S1-like distribution (blue) and S2-like distribution (red). The data used for building the PDFs was simulated for the first mock data challenge of LZ and did not include pulse-shaping effects from the detector or electronics response. . . . . 134
- 7.1 Correlation matrices for the non-normalized LZap dataset. The left plot shows the correlations between all pulse parameters (RQs) available in LZap. The right plot shows the correlations between the composite pulse parameters used extensively in this work. . . . . 149
- 7.2 Scatter plots of the GMM components for  $pA$  vs  $pL90$  pulse parameters (left) and  $TBA$  vs  $pA$  (right). The colors associated to each GMM component are cyclical and repeat for some components. The size and density of the 67 components vary significantly. . . . . 164
- 7.3 Scatter plots of the GMM components  $k = 31$  (cyan),  $k = 15$  (brown) and  $k = 48$  (orange), for  $pA$  vs  $pL90$  pulse parameters (left) and  $TBA$  vs  $pA$  (right). The component  $k = 31$  contains S2 pulses with abnormal behaviour and was selected for handscanning. Components  $k = \{15, 48\}$  were also selected for handscanning and were identified as SE splits. . . . . 165

- 7.4 Fraction of HADES labels per GMM component (top) and frequency of each HADES label per component (bottom). Error bars are Poisson in nature. Some mixture of HADES labels can be observed in most components, which means that either HADES is separating populations that the GMM clustering algorithm perceives as similar, or the GMM clustering is merging distinct populations that HADES can differentiate. Due to the simplicity of the HADES algorithm, the former hypothesis seems more likely. Arrows above the top plot mark the GMM components selected for handscanning. . . . . 166
- 7.5 Schematic representation of a simple decision tree used for classification. The input data (blue) passes several decision nodes (gray) that evaluate one or more parameters of the data object and redirect it to the branch that represents the result of the test until a terminator node, or leaf node (green), is reached. Each leaf node has an associated class label that identifies the data objects that reach it, essentially grouping the data into classes of similar objects using their parameters as discriminants. . . . . 169
- 7.6 Decision tree trained to classify LZap pulses into S1-like, S2-like and Other pulse classes. The S1-like class contains the S1 and MPE pulses while the S2-like class contains both the S2 pulses and the SE pulses clustered in the GMM analysis described in Section 7.3.1. . . . . 171
- 7.7 Impact of forest size (number of estimators) in the classification accuracy over the test set. Each accuracy value is the mean of the final accuracy of three independent training sessions with the same model architecture. Each tree is allowed to grow freely (`maxDepth=None`), and the minimum number of samples required to split a branch is set to 2 (`minSamplesSplit=2`). . . . . 179
- 7.8 Mean accuracy of the RF model for different values of `maxDepth` (left) and `minSamplesSplit` (right) hyperparameters, for a model with  $N_{trees} = 101$  trees. Each accuracy value is the mean of the final accuracy of three independent training sessions with the same model architecture. . . . . 180
- 7.9 Distribution of the pulse populations in the marginal distribution  $TBA$  vs  $pA$  after being processed by the RFClassifier. The top-left plot displays the distribution of pulses classified as S1. The top right plot displays the distribution of S2 pulses, while the bottom plots display the Other population on the left and the SE population on the right. . . . . 182
- 7.10 Distribution of the pulse populations in the marginal distribution  $pA$  vs  $pL90$  after being processed by the RFClassifier. The top-left plot displays the distribution of pulses classified as S1. The top right plot displays the distribution of S2 pulses, while the bottom plots display the Other population on the left and the SE population on the right. . . . . 183

- 7.11 Feature importance ranking obtained with Random Forest classifiers. The top left plot shows the overall importance ranking for the final multi-class RFClassifier model. The top-right, bottom-left and bottom-right plots show the importance ranking for separating S1, S2 and SE pulses, respectively, obtained by three separate RFClassifier models trained in a “ONE-vs-ALL” scenario. . . . . 184
- 7.12 Permutation importance ranking obtained with the RFClassifier. . . . . 186
- 7.13 Diagram of an artificial neuron, also called a perceptron. This particular NN architecture is also called a single layer perceptron since it contains only a single computational layer ( $\phi$ , see Equation 7.25) at the output node. . . . . 188
- 7.14 Examples of common activation functions [Agg18, CUH15]. . . . . 191
- 7.15 Schematic representation of a feed-forward neural network, with  $D$  input nodes, 2 hidden layers with  $P$  and  $Q$  units each, and  $K$  output nodes. The connection between nodes have variable weights represented by the matrix  $\Omega = \Omega^l$ , with  $l = 1, 2, 3$  the index of the  $L = 3$  computational layers. This architecture is often called a *feedforward multi-layer perceptron*. . . . . 194
- 7.16 Stacked distribution of the preprocessed  $pA$ ,  $pL90$ ,  $pH$  and  $H2L$  RQs for each class label used to tune and train the TriNet Classifier model. The pulse class labels were obtained by the HADES classifier. The  $\log_{10} H2L$  parameter is mostly negative, and mean-centering its value distribution would not improve its discrimination ability significantly. Therefore, it was left as it is despite the risk of impacting the performance of the network. . . . . 205
- 7.17 Validation accuracy for different combinations of activation functions and optimizers for a network with 2 hidden layers, 42 hidden units total, and using a learning rate  $\alpha = 0.001$  and batch size  $n = 256$  on training. The results are averaged over the last half of training, where the accuracy and loss were observed to be stable. . . . . 207
- 7.18 Averaged validation loss (bottom), validation accuracy (middle) and training time (top), for different combinations of NN depth and hidden unit number per layer. All models considered here were benchmarked using ELU activation, *RMSprop* optimization and learning rate  $\alpha = 0.001$  and batch size  $n = 256$ . . . . . 209
- 7.19 Averaged validation loss (bottom), validation accuracy (middle) and training time (top), for different combinations of the learning rate and batch size hyperparameters. All models considered here were benchmarked using ELU activation, *RMSprop* optimization,  $L = 3$  hidden layers and  $P = 11$  hidden units per layer. . . . . 211

7.20	Simplified schematic of the TriNet Classifier model (left). The model is composed of an ensemble of three NN models (right) each trained to separate the three main pulse classes in the data: S1, S2 and SE pulses. The outputs of each NN classifier, $\hat{y}_{k'} = \{\epsilon_{k'}, \tilde{\epsilon}_{k'}\}$ , represent the response and anti-response of the classifier to its respective class. . . . .	212
7.21	Relationship between the confidence $\Gamma_\epsilon$ and complement $\Gamma_{\tilde{\epsilon}}$ for the TriNet Classifier (top-left), and the responses and anti-responses of the S1-vs-All $NN_1$ classifier (top-right), S2-vs-All $NN_2$ classifier (bottom-left), and SE-vs-All $NN_3$ classifier (bottom-right). . . . .	215
7.22	Ternary plots of the overall response $\epsilon$ (left) and anti-response $\hat{\epsilon}$ (right) of the TriNet Classifier. . . . .	217
7.23	Distribution of the pulse populations in the marginal distribution $TBA$ vs $pA$ after being processed by the TriNet Classifier. Top-left: pulses classified as S1 using the selection $p(k = 1) > 0.23$ . Top right: S2 pulses selected using $p(k = 2) > 0.25$ . Bottom-right: SE population selected using $p(k = 3) > 0.25$ . Bottom-left: distribution of the remaining pulses, <i>i.e.</i> , $p(k = 1) < 0.23 \wedge p(k = 2) < 0.25 \wedge p(k = 3) < 0.25$ , classified as Other. . . . .	221
7.24	Distribution of the pulse populations in the marginal distribution $pA$ vs $pL90$ after being processed by the TriNet Classifier. The distributions of S1-like (top-left), S2-like (top-right), SE-like (bottom-right), and Other (bottom-left) pulses were obtained using the same selections as in Figure 7.23. . . . .	222
A.1	[1/8] Decay scheme of $^{214}\text{Bi}$ [BCD <sup>+</sup> 08]. . . . .	282
A.9	Decay scheme of $^{208}\text{Tl}$ [BCD <sup>+</sup> 13]. . . . .	290
A.10	Uranium series diagram. The Uranium decay chain starts with the naturally-occurring isotope $^{238}\text{U}$ and terminates with the stable $^{206}\text{Pb}$ following a $4n + 2$ rule for the mass number [FT13]. The total energy release for this chain is 51.7 MeV. The isotope $^{222}\text{Rn}$ , with a half-life of $T_{hl} = 3.8$ d, is produced along this decay chain. Figure from [Mal14]. . . . .	291
A.11	From left to right: the Thorium and Actinium series. The Actinium decay chain starts with the naturally-occurring isotope $^{235}\text{U}$ and terminates with the stable $^{207}\text{Pb}$ following a $4n + 3$ rule for the mass number $A$ [FT13]. The total energy release for this chain is 46.4 MeV. The Thorium decay chain starts with the naturally-occurring isotope $^{232}\text{Th}$ and terminates with the stable $^{208}\text{Pb}$ following a $4n$ rule for the mass number. The total energy release for this chain is 42.6 MeV. The isotopes $^{219}\text{Rn}$ ( $T_{hl} = 4$ s) and $^{220}\text{Rn}$ ( $T_{hl} = 55$ s) are produced along the Actinium and Thorium decay chains, respectively. Figure from [Mal14]. . . . .	292

# List of Tables

2.1 Latest results from the Planck Collaboration for the cosmological parameters of the  $\Lambda$ CDM model. The values presented have 68% error bands and are obtained using a combined likelihood on the CMB temperature (TT) power spectrum, E-mode polarization (EE) power spectrum, high-multipole temperature-polarization cross-correlation (TE), low-multipole polarization (lowE), lensing and baryon acoustic oscillations (BAO) data [A<sup>+</sup>18a]. . . . . 9

3.1 The three generations of the left-handed lepton sector of the Standard Model (SM) and their the electromagnetic charge (**Q**). The left-handed leptons can be arranged in  $SU(2)_L$  doublets of the weak sector, while right-handed particles are  $SU(2)$  singlets and therefore do not undergo weak interactions. . . . . 35

3.2 Sources of neutrinos from the solar fusion cycle according to the Standard Solar Model [BSB05, GK07].  $\langle E_\nu \rangle$  is the average neutrino energy and  $E_\nu^{max}$  is the endpoint energy of the reaction. Notice that neutrinos with discrete energies are the result of two body interactions, while neutrinos with a continuous distribution of energies (for which only  $E_\nu^{max}$  is given) are the result of three body interactions. 36

3.3 Current best combined results of the 3-neutrino oscillation parameters [Z<sup>+</sup>20]. . . 44

3.4 Measured half-lives and Q-values of two-neutrino double beta decay processes ( $2\nu\beta\beta$ ,  $2\nu$ ECEC) of several isotopes. The  $2\nu$ ECEC decays of <sup>78</sup>Kr and <sup>124</sup>Xe are the longest decays ever observed. As a reference, the age of the Universe is estimated around  $13.8 \times 10^9$  years. The  $2\nu\beta\beta$  decay of <sup>238</sup>U and the  $2\nu$ ECEC decay of <sup>130</sup>Ba are observed indirectly through radiochemistry and geochemistry techniques, respectively. From Reference [Bar20]. . . . . 53

3.5	Experiments searching for the $0\nu\beta\beta$ decay process in different isotopes, listing the total available (active) isotopic mass and the sensitivity limits on the half-life of the $0\nu\beta\beta$ decay. Future experiments under commissioning or R&D stages have the expected sensitivity values marked as “projected” [Det20]. . . . .	58
4.1	Dimensions and masses of some relevant elements of the LZ detector. Table adapted from Reference [A <sup>+</sup> 20d]. . . . .	61
4.2	Physical properties of the four grid electrodes of the TPC of LZ: working voltage, diameter of each 304 stainless steel ultra-finish wire, pitch of the grid mesh and number of wires in the mesh [A <sup>+</sup> 20g]. . . . .	63
4.3	Projected detector performance parameters for the TPC of LZ. Indented entries indicate a dependency on preceding parameters (see text for details). Table from [A <sup>+</sup> 20d]. . . . .	68
5.1	Summary table of the masses, activities and estimated background counts in the $\pm 1\sigma$ ROI and inner 967 kg mass, for a 1000 day run, considering 1.0% energy resolution at Q-value and 0.3 cm multiple scatter rejection along $z$ (see text for details). The veto cuts discussed in Section 5.1.4 are also applied to the background estimates presented here. . . . .	82
5.2	Assayed activities of the rock and lining materials of the Davis cavern walls, taken from Reference [A <sup>+</sup> 20b]. . . . .	86
5.3	Backgrounds induced by neutrino capture on xenon isotopes. . . . .	97
5.4	Relative uncertainties on the background rates for the background components assumed in the PLR analysis [A <sup>+</sup> 20c]. The uncertainties are treated in the unbinned extended likelihood function in Equation 5.5 as nuisance parameters. . . . .	100
5.5	Summary of the assumptions used and results obtained in this preliminary study of the sensitivity of LZ to the $0\nu\beta\beta$ decay of $^{136}\text{Xe}$ . . . . .	104
6.1	Metrics of the processing of 750 simulated events from multiple backgrounds, $^{220}\text{Rn}$ decay in the LXe bulk and neutrons from a DD source. The values in the table were obtained through HADES directly. The pie chart on the right side shows the frequency of the different classes for all “Golden” pulses, defined as the pulses without bad RQs and with coincidence above 1. These results include misclassifications of HADES, particularly relevant for the fraction of S1 pulses in the data, which is largely inflated by the misclassification of SE splits. . . . .	130
6.2	Confusion matrix of HADES results taken from the handscanning of 2748 pulses on 23 random events. . . . .	130

7.1	Mapping from GMM component $k$ to pulse class label. The labels marked with an asterisk were selected via handscanning of 20 or more random pulse samples from the respective component. The primary pulse type observed on each handscanned component is denoted in parenthesis: Afterpulsing (AP), two coincident dark counts (2fDC), three coincident dark counts (3fDC), SE splits, S2 splits, baselines and random mixtures of pulses (see Section 6.2.1 for details). The waveforms of the S2 pulses found in component $k = 31$ do not display strange features, despite the odd behaviour of <i>TBA</i> and <i>pL90</i> (see text for details). . . . .	167
7.2	Confusion matrix of the GMM results compared with the results from HADES over the full dataset. . . . .	168
7.3	Confusion matrix of the RFClassifier results over the test dataset, using the GMM results as the truth labels. . . . .	181
7.4	Averaged training time (epochs) for different combinations of activation functions and optimizers, for the same model used to produce the results displayed in Figure 7.17. Numbers in red indicate the combinations that did not stabilize within 500 epochs, the maximum allowed number of cycles through the entire dataset considered in this benchmarking exercise. The value marked with † corresponds to the combination of activation and optimizer that was chosen for the TriNet Classifier. . . . .	208
7.5	Confusion matrix of the TriNet Classifier results compared with the HADES results, over the entire LZap dataset. . . . .	223
7.6	Confusion matrix of the TriNet Classifier results compared with the GMM class labels obtained in Section 7.3.1, over the entire LZap dataset. . . . .	223
7.7	Comparison of the metrics of performance for all the classifier models developed in this work. The results for S2-SE mixing are obtained by considering that a SE seen as an S2 and vice-versa is not a misclassification. The results from HADES were verified by handscan. The metrics for the RFClassifier and TriNet Classifier were obtained using the results of the GMM clustering as a proxy to the “true” class. . . . .	225
B.1	Event header RQs: <code>eventHeader.&lt;RQ&gt;</code> . . . . .	293
B.2	Pulse-level RQs for the mixed TPC and Skin detectors: <code>pulsesTPC.&lt;RQ&gt;</code> , <code>pulsesSkin.&lt;RQ&gt;</code> . The HG and LG detectors contain the same RQ structure as the mixed TPC detector: <code>pulsesTPCHG.&lt;RQ&gt;</code> , <code>pulsesTPCLG.&lt;RQ&gt;</code> . . . . .	295
B.3	Pulse-level composite parameters used exclusively within the <i>PulseClassifier</i> module, calculated from the original RQs used in LZap, for the mixed TPC and Skin detectors: <code>pulsesTPC.&lt;RQ&gt;</code> , <code>pulsesSkin.&lt;RQ&gt;</code> . . . . .	298

- C.1 Useful nomenclature of data-specific concepts. Uppercase letter refer to generic aspects of the variables, while lowercase letter refer to specific or observed variables or elements. Bold letters denote a matrix, often referring to a set of multidimensional data objects, but a lowercase bold letter refers to a single multidimensional data object. *E.g.*, a set of  $N$  data objects with  $D$  parameters can be named  $\mathbf{X}$ , with  $\mathbf{x}_n$  the  $1 \times D$   $n^{\text{th}}$  multidimensional data object of the  $N \times D$  variable  $\mathbf{X}$ . . 302

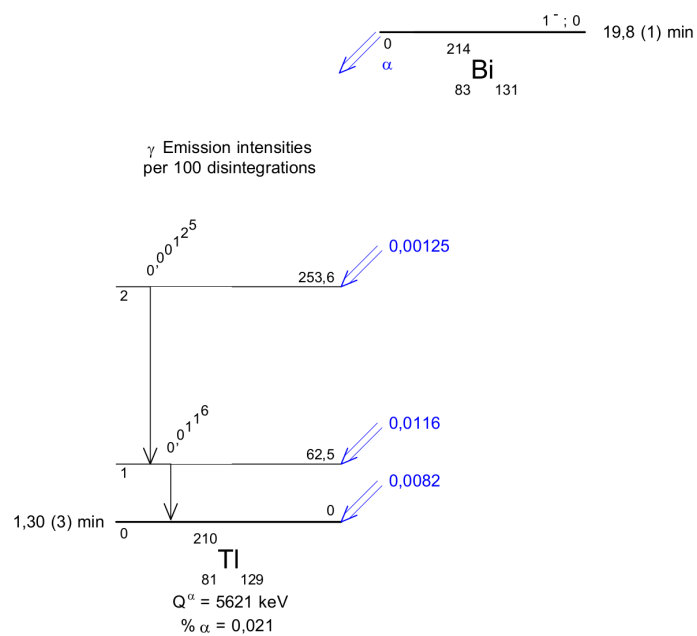


# Attachments



# Appendix A

## Decay Schemes and Decay chains



**Figure A.1:** [1/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08].

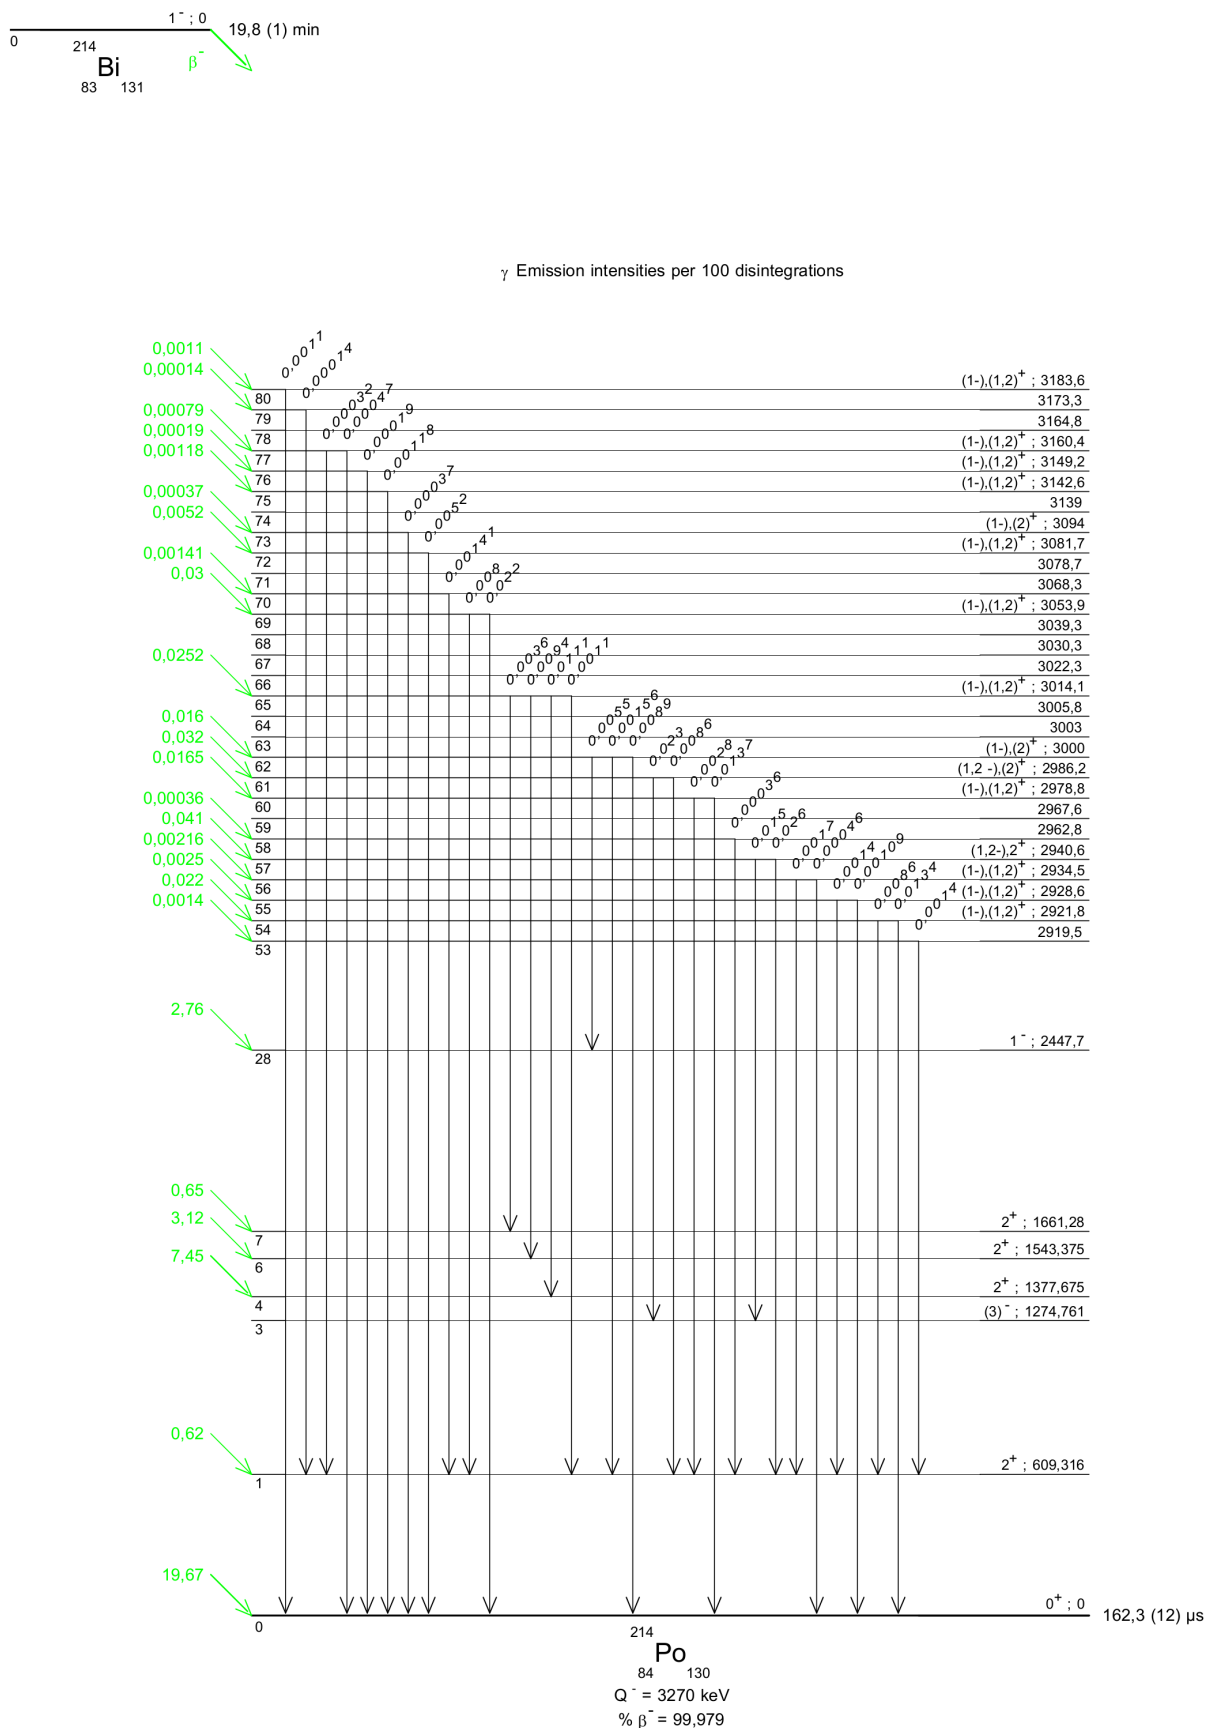


Figure A.1: [2/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).




Figure A.1: [3/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).




Figure A.1: [4/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).




Figure A.1: [5/8] Decay scheme of <sup>214</sup>Bi [BCD<sup>+</sup>08] (continuation from page 282).






Figure A.1: [6/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).




Figure A.1: [7/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).




Figure A.1: [8/8] Decay scheme of  $^{214}\text{Bi}$  [BCD<sup>+</sup>08] (continuation from page 282).


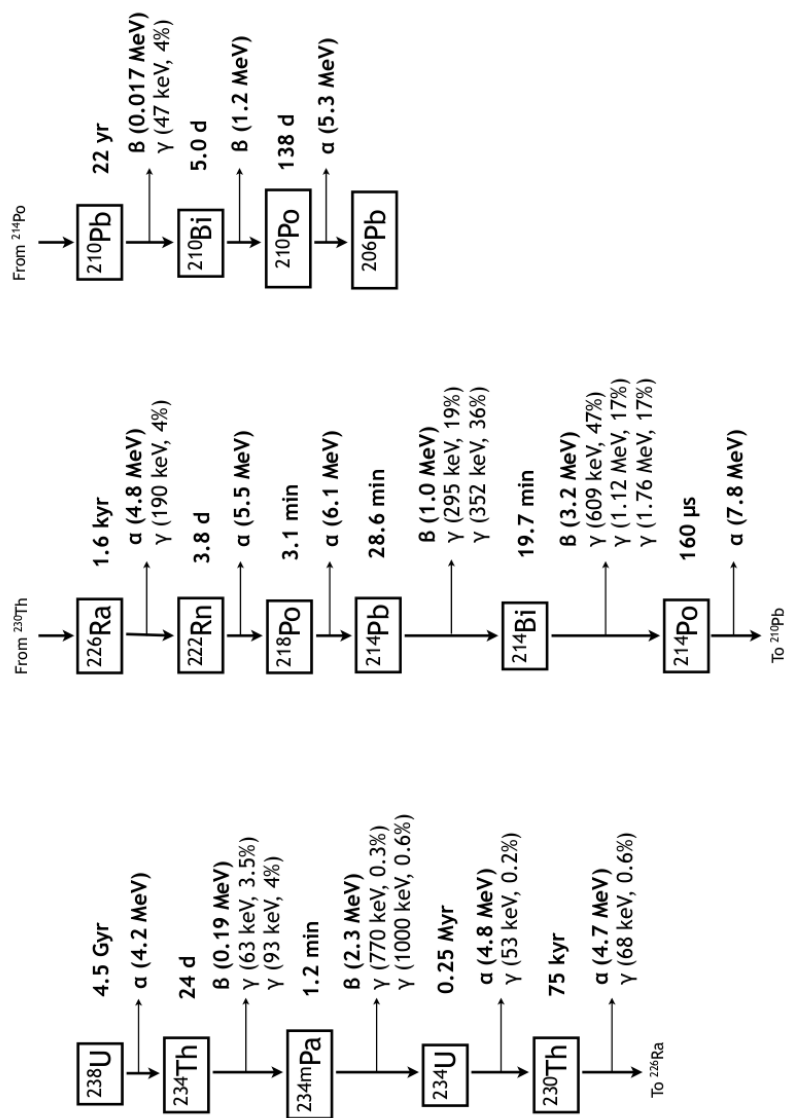
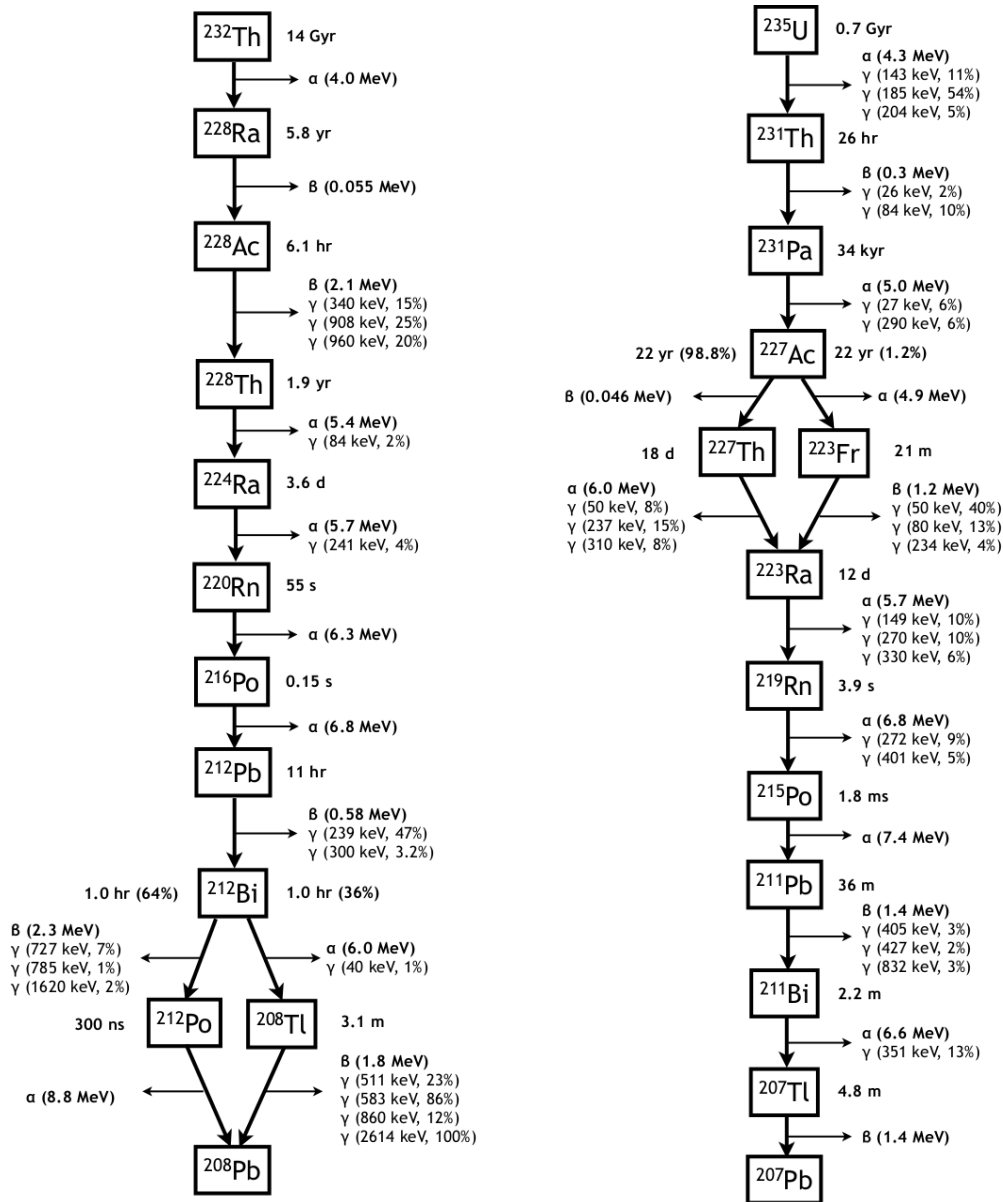


Figure A.9: Decay scheme of  $^{208}\text{Tl}$  [BCD<sup>+</sup>13].



**Figure A.10:** Uranium series diagram. The Uranium decay chain starts with the naturally-occurring isotope  $^{238}\text{U}$  and terminates with the stable  $^{206}\text{Pb}$  following a  $4n + 2$  rule for the mass number [FT13]. The total energy release for this chain is 51.7 MeV. The isotope  $^{222}\text{Rn}$ , with a half-life of  $T_{1/2} = 3.8$  d, is produced along this decay chain. Figure from [Mal14].



**Figure A.11:** From left to right: the Thorium and Actinium series. The Actinium decay chain starts with the naturally-occurring isotope  $^{235}\text{U}$  and terminates with the stable  $^{207}\text{Pb}$  following a  $4n + 3$  rule for the mass number  $A$  [FT13]. The total energy release for this chain is 46.4 MeV. The Thorium decay chain starts with the naturally-occurring isotope  $^{232}\text{Th}$  and terminates with the stable  $^{208}\text{Pb}$  following a  $4n$  rule for the mass number. The total energy release for this chain is 42.6 MeV. The isotopes  $^{219}\text{Rn}$  ( $T_{hl} = 4$  s) and  $^{220}\text{Rn}$  ( $T_{hl} = 55$  s) are produced along the Actinium and Thorium decay chains, respectively. Figure from [Mal14].

# Appendix B

## Reduced Quantities (RQ) of LZap data

All LZap RQs were taken from:

<https://luxzeplin.gitlab.io/docs/softwaredocs/analysis/analysiswithrqs/rqlist.html>.

Formerly at:

<https://lz-git.ua.edu/docs/softwaredocs/blob/master/docs/analysis/analysiswithrqs/rqlist.md>

**Table B.1:** Event header RQs: `eventHeader.<RQ>`

Type	RQ name	Description
unsigned long	runID	run ID, taken from <code>Global.runNumber</code> in raw data file
unsigned long	eventID	event ID, taken from <code>Event.globalEvent</code> in raw data file
string	rawFileName	raw data file that this event belongs to
unsigned long	triggerTimeStamp_s	seconds component of trigger time (POSIX)
unsigned long	triggerTimeStamp_ns	ns component of trigger time (ns past <code>triggerTimeStamp_s</code> )
int	triggerMultiplicity	the number of channels that were above threshold and contributing to the trigger
int	triggerType	the type of trigger that accepted this event
Continued on next page		

Table B.1 – continued from previous page

Type	RQ name	Description
unsigned long	sumPodStartTime_s	seconds component of start time of sum POD associated with this trigger; N.B. this is an Event Builder sum pod, NOT an LZap sum pod! (UNIX Epoch)
unsigned long	sumPodStartTime_ns	ns component of start time of sum POD (ns past sumPodStartTime_s)
unsigned int	sumPodSampleCount	number of samples in the sum POD; N.B. this is an Event Builder sum pod, NOT an LZap sum pod!
vector<float>	sumPodSamples	samples of the sum POD associated with this trigger; N.B. this is an Event Builder sum pod, NOT an LZap sum pod!
int	runType	
unsigned long	runStartTime_s	seconds component of run/acquisition start time (UNIX Epoch)
unsigned long	runStartTime_ns	ns component of run/acquisition start time (ns past runStartTime_s)
int	rawFileFormatVersion	Format version of the raw data file from which this event came
vector<int>	nPodsDC	the number of PODs in each data collector (14 DCs)
int	nPods	the total number of PODs in the event
vector<unsigned long>	bufferStartTime_s	seconds component of time when each buffer went live for the event (46 buffers; UNIX Epoch)
vector<unsigned long>	bufferStartTime_ns	ns component of time when each buffer went live for the event (46 buffers; ns past bufferStartTime_s)
vector<unsigned long>	bufferStopTime_s	seconds component of time when each buffer went dead for the event (46 buffers; UNIX Epoch)
vector<unsigned long>	bufferStopTime_ns	ns component of time when each buffer went dead for the event (46 buffers; ns past bufferEndTime_s)



**Table B.2:** Pulse-level RQs for the mixed TPC and Skin detectors: `pulsesTPC.<RQ>`, `pulsesSkin.<RQ>`. The HG and LG detectors contain the same RQ structure as the mixed TPC detector: `pulsesTPCHG.<RQ>`, `pulsesTPCLG.<RQ>`

Type	RQ name	Description
int	nPulses	Number of pulses found in this detector in this event
vector<int>	pulseID	<b>pID</b> ; ID of the pulse (time ordered)
vector<int>	pulseStartTime_ns	<b>pST</b> [ns]; Start time of pulse relative to event trigger time
vector<int>	pulseEndTime_ns	<b>pET</b> [ns]; End time of pulse relative to event trigger time
vector<bool>	saturated	Is at least one channel of the pulse saturated
vector<int>	nSaturatedChannels	Number of saturated channels
vector<vector<unsigned int>>	saturatedChannelIDs	ChannelIDs of the channels that are saturated
vector<float>	pulseArea_phd	<b>pA</b> [phd]; Total area of summed pod from pulse start to end
vector<float>	positiveArea_phd	<b>posA</b> [phd]; Total positive area of summed pod from pulse start to end
vector<float>	negativeArea_phd	<b>negA</b> [phd]; Total negative area of summed pod from pulse start to end
vector<float>	peakAmp	<b>pH</b> [phd/sample]; Max amplitude of summed pod within pulse
vector<int>	peakTime_ns	<b>pHT</b> [ns]; Time at which summed pod reaches maximum within pulse
vector<int>	areaFractionTime1_ns	<b>aft1</b> [ns]; Time at which summed pod reaches 1% of total area; ns relative to pulse start
vector<int>	areaFractionTime5_ns	<b>aft5</b> [ns]; Time at which summed pod reaches 5% of total area; ns relative to pulse start
vector<int>	areaFractionTime10_ns	<b>aft10</b> [ns]; Time at which summed pod reaches 10% of total area; ns relative to pulse start
vector<int>	areaFractionTime25_ns	<b>aft25</b> [ns]; Time at which summed pod reaches 25% of total area; ns relative to pulse start
vector<int>	areaFractionTime50_ns	<b>aft50</b> [ns]; Time at which summed pod reaches 50% of total area; ns relative to pulse start
Continued on next page		

Table B.2 – continued from previous page

Type	RQ name	Description
vector<int>	areaFractionTime75_ns	<b>aft75</b> [ns]; Time at which summed pod reaches 75% of total area; ns relative to pulse start
vector<int>	areaFractionTime90_ns	<b>aft90</b> [ns]; Time at which summed pod reaches 90% of total area; ns relative to pulse start
vector<int>	areaFractionTime95_ns	<b>aft95</b> [ns]; Time at which summed pod reaches 95% of total area; ns relative to pulse start
vector<int>	areaFractionTime99_ns	<b>aft99</b> [ns]; Time at which summed pod reaches 99% of total area; ns relative to pulse start
vector<float>	pulseAreaNeg50ns_phd	<b>pAN50</b> [phd]; Area in fixed integration window ranging from 50ns before <b>aft5</b> to <b>aft5</b>
vector<float>	pulseArea50ns_phd	<b>pA50</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 50ns after <b>aft5</b>
vector<float>	pulseArea100ns_phd	<b>pA100</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 100ns after <b>aft5</b>
vector<float>	pulseArea200ns_phd	<b>pA200</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 200ns after <b>aft5</b>
vector<float>	pulseArea500ns_phd	<b>pA500</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 500ns after <b>aft5</b>
vector<float>	pulseArea1us_phd	<b>pA1k</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 1us after <b>aft5</b>
vector<float>	pulseArea2us_phd	<b>pA2k</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 2us after <b>aft5</b>
vector<float>	pulseArea5us_phd	<b>pA5k</b> [phd]; Area in fixed integration window ranging from <b>aft5</b> to 5us after <b>aft5</b>
vector<float>	promptFraction50ns	<b>pF50</b> ; Fraction of summed pod area in 50ns window at start of pulse relative to total pulse area; 50ns window defined from 10ns before to 40ns after <b>aft5</b>
vector<int>	rmsWidth_ns	<b>pRMSW</b> [ns];
vector<int>	fwhm_ns	<b>pFWHM</b> [ns];
vector<float>	topArea_phd	<b>topA</b> [phd]; Total area in top array
vector<float>	bottomArea_phd	<b>botA</b> [phd]; Total area in bottom array

Continued on next page

Table B.2 – continued from previous page

Type	RQ name	Description
vector<float>	topBottomAsymmetry	<b>TBA</b> ; Ratio of total area of top PMTs vs. total area of bottom PMTs
vector<int>	coincidence	<b>coinc</b> ; Number of channels that had non-zero contribution to pulse
vector<float>	topCentroidX_cm	[cm]; centroid X position, calculated using top array only
vector<float>	topCentroidY_cm	[cm]; centroid Y position, calculated using top array only
vector<float>	bottomCentroidX_cm	[cm]; centroid X position, calculated using bottom array only
vector<float>	bottomCentroidY_cm	[cm]; centroid Y position, calculated using bottom array only
vector<float>	s1Probability	Probability that pulse is S1 (0 or 1)
vector<float>	s2Probability	Probability that pulse is S2 (0 or 1)
vector<float>	singlePEprobability	Probability that pulse is SPE (0 or 1)
vector<float>	multiplePEprobability	Probability that pulse is MPE (0 or 1)
vector<float>	singleElectronProbability	Probability that pulse is SE (0 or 1)
vector<float>	otherProbability	Probability that pulse is none of the above (0 or 1)
vector<float>	otherS2Probability	Probability that pulse is S2-like (has correct hit pattern) but non-S2 pulse shape (0 or 1)
vector<string>	classification	Classification of pulse given above probabilities
vector<int>	photonCount	Pulse photon count. Photon count is not evaluated for all pulses.
vector<int>	topPhotonCount	Pulse photon count, top array
vector<int>	bottomPhotonCount	Pulse photon count, bottom array
vector<float>	s2Xposition_cm	[cm]; X position of pulse; 999 if pulse is not classified as S2
vector<float>	s2Yposition_cm	[cm]; Y position of pulse; 999 if pulse is not classified as S2
vector<int>	s2XYrecStatus	Returned status of Mercury reconstruction
vector<int>	s2RecDof	Degrees of freedom in XY reconstruction
vector<float>	s2XYchi2	Chi2 from Mercury XY reconstruction
vector<float>	s2XXcov_cm	[cm <sup>2</sup> ]; XX component of covariance matrix from XY reconstruction; NB units are cm <sup>2</sup> , despite the variable name!

Continued on next page

Table B.2 – continued from previous page

Type	RQ name	Description
vector<float>	s2YYcov_cm	[cm <sup>2</sup> ]; YY component of covariance matrix from XY reconstruction; NB units are cm <sup>2</sup> , despite the variable name!
vector<float>	s2XYcov_cm	[cm <sup>2</sup> ]; XY component of covariance matrix from XY reconstruction; NB units are cm <sup>2</sup> , despite the variable name!
vector<float>	s2XYarea_phd	[phd]; LCE-corrected S2 area, as returned by Mercury
vector<int>	HGLGpulseID	LG pulse that best matches HG pulse. If HG pulse has no good matching LG pulse, entry is -1. Vice versa for LG.
vector<vector<int>>	chID	Channel ID for the corresponding channel in the chPulseArea_phd branch
vector<vector<float>>	chPulseArea_phd	chpA [phd]; Pulse area for each channel. If channel is not present in chID branch, then pulse area is zero for that channel.
vector<vector<float>>	chPeakAmp	chpH [phd/sample]; Pulse height for each channel.
vector<vector<int>>	chPeakTime_ns	chpHT [ns]; Time of peak for each channel; ns relative to pulse start time
vector<vector<bool>>	chSaturated	Saturation flag for each channel.
vector<vector<int>>	chPhotonCount	Pulse photon count for each channel. Photon count is not evaluated for all pulses.
vector<vector<vector<int>>>	chPhotonTimes	Vector of photon times (in ns relative to pulse start) for each channel.

**Table B.3:** Pulse-level composite parameters used exclusively within the *PulseClassifier* module, calculated from the original RQs used in LZap, for the mixed TPC and Skin detectors: `pulsesTPC.<RQ>`, `pulsesSkin.<RQ>`.

Type	RQ name	Description
vector<int>	pulseLength_ns	<b>pL</b> = pST - pET [ns]; Length of the pulse, from pulseStartTime_ns to pulseEndTime_ns.
vector<int>	pulseLength90_ns	<b>pL90</b> = aft95 - aft5 [ns]; Length of the pulse at 90% area, from aft5 to aft95.
Continued on next page		

Table B.3 – continued from previous page

Type	RQ name	Description
vector<float>	promptFraction100ns	<b>pF100</b> = $pA_{100}/pA$ ; Fraction of summed pod area in 100ns window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 100ns after.
vector<float>	promptFraction200ns	<b>pF200</b> = $pA_{200}/pA$ ; Fraction of summed pod area in 200ns window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 200ns after.
vector<float>	promptFraction500ns	<b>pF500</b> = $pA_{500}/pA$ ; Fraction of summed pod area in 500ns window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 500ns after.
vector<float>	promptFraction1kns	<b>pF1k</b> = $pA_{1k}/pA$ ; Fraction of summed pod area in 1 $\mu$ s window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 1 $\mu$ s after.
vector<float>	promptFraction2kns	<b>pF2k</b> = $pA_{2k}/pA$ ; Fraction of summed pod area in 2 $\mu$ s window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 2 $\mu$ s after.
vector<float>	promptFraction5kns	<b>pF5k</b> = $pA_{5k}/pA$ ; Fraction of summed pod area in 5 $\mu$ s window at start of pulse relative to total pulse area; window defined from <b>aft5</b> to 5 $\mu$ s after.
vector<float>	pulseHeightTimeLength	<b>pHTL</b> = $pHT/pL$ ; Ratio between <b>peakTime_ns</b> and <b>pL</b> . Proxy to pulse skewness.
vector<float>	heightToLength	<b>H2L</b> = $pH/pL_{90}$ ; Ratio between <b>peakAmp</b> and <b>pL90</b> .



# Appendix C

## Useful Nomenclature

### Definitions:

- feature: observable property of the dataset or individual data object
- attribute: independent variable that characterizes one feature of the data object
- parameter: a factor (numerical, categorical or logical) used to define a system or sets the conditions of its operation (hyperparameter). It can also be used to describe a data attribute<sup>1</sup>.
- label: an assigned discrete attribute that indicates the class to which a data object belongs to, often assigned by a predictive model (classifier).
- class: a set or category of things having some property or attribute in common and differentiated from others by kind, type, or quality.
- data object: a container for data values (attributes, labels, features) and whose properties can be used to identify it and describe it. An instance of a class type is a class object.
- dataset: a collection of data objects related by origin, creator process or environment. The properties of the dataset emerge from the properties of the individual data objects.

---

<sup>1</sup>This is the case for the LZap pulse RQs, that are referred to as “pulse parameters” throughout this work.

**Table C.1:** Useful nomenclature of data-specific concepts. Uppercase letter refer to generic aspects of the variables, while lowercase letter refer to specific or observed variables or elements. Bold letters denote a matrix, often referring to a set of multidimensional data objects, but a lowercase bold letter refers to a single multidimensional data object. *E.g.*, a set of  $N$  data objects with  $D$  parameters can be named  $\mathbf{X}$ , with  $\mathbf{x}_n$  the  $1 \times D$   $n^{\text{th}}$  multidimensional data object of the  $N \times D$  variable  $\mathbf{X}$ .

Symbol	Name	Description
$N, n$	Number of data objects	The number of data objects in the dataset. Index $n = 1, 2, \dots, N$ .
$D, d$	Number of parameters	The number of parameters on the data, a.k.a, the dimensionality of the dataset. Index $d = 1, 2, \dots, D$ .
$K, k$	Number of GMM components	The number of GMM components considered in the GMM model. Index $k = 1, 2, \dots, K$ .
$S, s$	Number of classes	The number of classes considered in the dataset. Index $s = 1, 2, \dots, S$ .
$M, m$	index of tree nodes	The index of the nodes of a treenumber of classes considered in the dataset. Index $s = 1, 2, \dots, S$ .
$\mathcal{S}$	Class categorical vector	Categorical vector of the $S$ classes: $\mathcal{S} = \{\text{class}_1, \text{class}_2, \dots, \text{class}_S\}$ . The indexing of $\mathcal{S}$ , written as $\mathcal{S}_i = \text{class}_i$ is a categorical, not a numerical value.
$\mathbf{X}$	Dataset	The full dataset of $N$ data objects with $D$ parameters – shape $[N \times D]$ . Alternative forms: $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n, \dots, \mathbf{x}_N\}$ , with $n = 1, 2, \dots, N$ the object/element index.
$\mathbf{x}_n$	Data object	The $n^{\text{th}}$ data object of the dataset $\mathbf{X}$ – shape $[1 \times D]$ . Alternatives: $\mathbf{x}_n = x_{n,1}, x_{n,2}, \dots, x_{n,d}, \dots, x_{n,D}$ , with $d = 1, 2, \dots, D$ the feature index.
$x_d$	parameter	The $d^{\text{th}}$ parameter of the data object $x$ – size 1. Alternatives: $x_d \in x_1, x_2, \dots, x_d, \dots, x_D$ , with $d = 1, 2, \dots, D$ the feature index.
$\mathbf{Y}$	Vector of results	The results obtained by the ML model after processing the dataset $\mathbf{X}$ with $N$ data objects – shape $[N \times 1]$ . Alternative forms: $\mathbf{Y} = \{y_1, y_2, \dots, y_n, \dots, y_N\}$ , with $n = 1, 2, \dots, N$ the object/element index.
Continued on next page		



Table C.1 – continued from previous page

Symbol	Name	Description
$y_n$	prediction result	The result of processing the $n^{\text{th}}$ data object ( $\mathbf{x}_n$ ) of the dataset $\mathbf{X}$ , <i>i.e.</i> , $y_n = \hat{f}(\mathbf{x}_n)$ where $\hat{f}$ is the predictor function of the model – shape [1]. If $\mathbf{x}_n$ is of class $\mathcal{S}_i$ , therefore having an arbitrary label value $S_i = i$ , the prediction is correct if $y_n == S_i$ .
$\mathbf{Z}$	Latent variable matrix	The latent variable matrix of a GMM model with $K$ components – shape $[N \times K]$ . Alternative forms: $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n, \dots, \mathbf{z}_N\}$ , with $n = 1, 2, \dots, N$ the object/element index.
$\mathbf{z}_n$	Latent variable	The $n^{\text{th}}$ latent variable of the GMM with $K$ components – shape $[1 \times K]$ . Alternatives: $\mathbf{z}_n = x_{n,1}, x_{n,2}, \dots, x_{n,k}, \dots, x_{n,K}$ , with $k = 1, 2, \dots, K$ the GMM component index.
$z_k$	Latent variable element	The $k^{\text{th}}$ element of the latent variable $z$ of the GMM with $K$ components – size 1. Alternatives: $z_k \in z_1, z_2, \dots, z_k, \dots, z_K$ , with $k = 1, 2, \dots, K$ the GMM component index.
$z_k$	Latent variable vector	The $k^{\text{th}}$ latent variable of the GMM with $K$ components – size $K$ .
$\hat{f}_{model}(\mathbf{x})$	predictor function of model	The predictor function $\mathbf{x} \rightarrow \hat{f}_{model}(\mathbf{x})$ of the model: $\hat{f}_{model} : \mathbf{X} \rightarrow \mathbf{Y}$ .
$\mathcal{T}$	Tree ensemble (forest)	An ensemble of $B$ tree classifiers, called a forest, $\mathcal{T} = \{T_b\}_1^B$ , with $b = 1, 2, \dots, B$ the Ensemble component index.
$T_b$	Tree	The $b^{\text{th}}$ tree classifier belonging to the ensemble $\mathcal{T} = \{T_b\}_1^B$ .

