UNIVERSIDADE D
COIMBRA

Francisco Patrão Henriques Zagalo

DEVELOPMENT AND TESTING OF AN MRI-COMPATIBLE IMMOBILIZATION DEVICE FOR HEAD AND NECK IMAGING

Dissertação no âmbito do Mestrado em Física Médica orientada pelo Professor Doutor Antero José Pena Afonso de Abrunhosa, pelo Professor Doutor Paulo Jorge Ribeiro da Fonte e pela Professora Doutora Susete Teresa Gaspar do Fetal e apresentada ao Departamento de Física da Faculdade de Ciências e Tecnologia da Universidade de Coimbra.

Setembro de 2022

FACULDADE DE CIÊNCIAS E TECNOLOGIA
UNIVERSIDADE D COIMBRA

Departamento de Física

DEVELOPMENT AND TESTING OF AN MRI-COMPATIBLE IMMOBILIZATION DEVICE FOR HEAD AND NECK IMAGING

Dissertação no âmbito do Mestrado em Física Médica orientada pelo Professor Doutor Antero José Pena Afonso de Abrunhosa, pelo Professor Doutor Paulo Jorge Ribeiro da Fonte e pela Professora Doutora Susete Teresa Gaspar do Fetal e apresentada ao Departamento de Física da Faculdade de Ciências e Tecnologia da Universidade de Coimbra.

Francisco Patrão Henriques Zagalo

Coimbra

2022

Acknowledgments

This dissertation was the result of several months of hard work through which I received the help, support, motivation and friendship from several people. In the next lines, I will remember all those who made this work possible and lead me up to this moment.

I thank to the Institute of Nuclear Sciences Applied to Health (ICNAS) of the University of Coimbra, in the persons of the Directors of board, having accepted and supported me as a MSc student, helping in every way possible the completion of this work, ceding the many hours of utilization of the Magnetic Resonance, and financing the materials and production needed for the project.

I thank to the Laboratory of Instrumentation and Experimental Particle Physics of Coimbra (LIP-Coimbra), in the persons of the Directors of board, the generous time and resources ceded to produce all the parts involved in this project's main component, without which this work wouldn't have been completed.

I thank to Professor Antero José Pena Afonso de Abrunhosa, the opportunity given to work under his wise guidance, his availability, support and friendship. It was a pleasure and an honour to work under his supervision.

I thank to Professor Paulo Jorge Ribeiro da Fonte and Professor Susete Teresa Gaspar do Fetal, the opportunity to work under their guidance. I sincerely thank all the time and work devoted, the help, support and knowledge shared with me during these months, and the friendship shown at all moments. It was an honour and a pleasure to work under their supervision.

I thank to Professor Maria Isabel Silva Ferreira Lopes, the opportunity given to embrace this project, the work and time devoted as coordinator of the MSc in Medical Physics, helping whenever needed, and the friendship shown at all moments.

I thank to Sónia Gomes Afonso, MRI Technologist, all the help and support given obtaining the Magnetic Resonance images, crucial to this work, and all the patience and friendship shown at all moments.

I thank to Luís Alberto Vieira Lopes Lopes, all the help and supported given designing and producing the immobilization device, without which this work would not have been completed, and all the patience and friendship shown at all moments.

I thank to all my friends of the MSc in Medical Physics and the Physics department, for the friendship, help and support shown at all moments during these two years, without which this work would have been a lot harder to complete.

I thank to my childhood friends, André Andrade, Gonçalo Gamboa, João Gonçalves, Anselmo Batista, Júlio Ferreira, and my musically gifted friends Diogo Ferraro, Francisco Cidade and Luís Carvalho, all the moments and memories of pure friendship I will keep forever, and all the motivation, support and help devoted at all times during the time of this work.

I thank my girlfriend Tânia Maria Fonseca Amaral, all the love, friendship, help, motivation, support, patience, and much more, devoted at all times, in the good and especially in the bad moments.

I thank my family, especially my mother, my father and my brother, all the love and support, for always believing in me, giving everything possible and impossible to enable me to succeed in life.

Resumo

O principal objetivo da presente Dissertação é o estudo e desenvolvimento de um dispositivo de imobilização sub-milimétrica que impede/ previne o movimento da cabeça em exames de Tomografia por Emissão de Positrões (PET) e Imagem por Ressonância Magnética (MRI) da cabeça e pescoço. Este projeto foi desenvolvido no Instituto de Ciências Nucleares Aplicadas à Saúde (ICNAS) da Universidade de Coimbra, em parceria com o Laboratório de Instrumentação e Física Experimental de Partículas (LIP) Coimbra, no âmbito do projecto *High Resolution Brain PET* (HiRezBrainPET).

A imagem PET com longos tempos de aquisição é suscetível a artefactos de movimento que podem comprometer a qualidade da imagem e levar a uma má interpretação da mesma. Tornou-se claro que, para evitar estes erros num *scanner* cerebral PET de alta resolução, as técnicas de imobilização utilizadas atualmente não são suficientes, uma vez que movimentos maiores do que a resolução da imagem do sistema resultam em desfoque das imagens, sendo necessária a imobilização sub-milimétrica da cabeça.

Através do estudo do estado da arte dos dispositivos e técnicas de imobilização da cabeça, foi escolhido que o método a desenvolver seria a imobilização da maxila através de uma boqueira ligada a uma estrutura externa, um método frequentemente utilizado em radioterapia. O dispositivo de imobilização foi desenvolvido para exames de ressonância magnética da cabeça e pescoço, que podem fornecer facilmente informação fiável sobre o movimento, sem exposição à radiação ionizante, sendo que a mesma técnica também pode ser utilizada com o scanner HiRezBrainPET.

Foram realizados testes digitais para aplicar e optimizar dois algoritmos: o algoritmo SURF e o algoritmo Block Matching, e foi desenvolvido software baseado neles para analisar as imagens de ressonância magnética e detetar qualquer movimento com exatidão e precisão. Para calibrar e testar este software para deslocamentos translacionais e rotacionais, foi desenvolvido um fantoma, de forma a simular o melhor possível as imagens da cabeça a analisar, juntamente com um dispositivo para induzir movimento, de forma a produzir translações e rotações conhecidas.

O dispositivo de imobilização foi projetado em *Computer-Aided Design* (CAD) e construído nas instalações do LIP Coimbra. Foi produzido num polímero POM-C devido às suas propriedades físicas, particularmente a sua compatibilidade em MRI, resistência, rigidez, temperatura de fusão, e custo monetário.

Os testes finais ao dispositivo de imobilização foram realizados no sistema de Ressonância Magnética do ICNAS. Os resultados de imobilização finais consistem na análise de movimento de três sequências de movimentos diferentes com: o voluntário em repouso, o voluntário a exercer movimentos grandes e o voluntário a exercer movimentos pequenos (aplicando 50% da força exercida nos movimentos grandes). Todas as sequências de movimentos foram realizadas com e sem o dispositivo de imobilização, e adquiridas no plano sagital e no plano axial. Cada sequência foi adquirida durante 2 minutos e 26 segundos.

Para as sequências de movimentos forçados pequenos e grandes, a diminuição de movimentos com o uso do dispositivo imobilizador foi mais notória do que nas sequências de repouso. Todas as rotações, tanto no plano sagital como axial, foram reduzidas em cerca de 89%. No plano sagital, os movimentos de deriva translacional da cabeça apresentaram reduções de 81% e 88%, com a maior diminuição a fazer-se notar no eixo YY da imagem. No plano axial, estas reduções de movimento foram de 88% e 90%, desta vez com a maior diminuição a fazer-se notar no eixo XX da imagem. Esta análise de dados permite concluir que este sistema imobilizador restringe fortemente todos os movimentos da cabeça do voluntário.

Para as sequências de repouso, em ambos os planos, o movimento de deriva angular da cabeça foi reduzido para cerca de metade do seu valor sem imobilização e os valores dos erros aleatórios foram reduzidos aproximadamente 70% no plano sagital e 86% no plano axial. No plano sagital, os movimentos de deriva translacional da cabeça foram reduzidos em cerca de 36% no eixo XX e 85% no eixo YY da imagem, com os erros aleatórios a diminuirem aproximadamente de 200 µm para valores perto de 40 µm, no eixo XX e YY. No plano axial, os movimentos de deriva translacional da cabeça foram reduzidos em cerca de 32% no eixo XX e 58% no eixo YY da imagem, com os erros aleatórios a diminuirem aproximadamente de 40 µm para valores perto de 20 µm, no eixo XX e YY da imagem.

Assim, parece ficar claro que este sistema de imobilização é capaz de realizar imobilização submilimétrica em exames à cabeça e pescoço.

Palavras-Chave: Immobilização da cabeça; Imagem por Ressonância Magnética; Tomografia por Emissão de Positrões; Processamento de imagem; Fantoma de movimento; Testes experimentais.

Abstract

The present Dissertation's main goal is the study and development of a sub-millimetric immobilization device to prevent head movement in Positron Emission Tomography (PET) and Magnetic Resonance Imaging (MRI) head or brain exams. This project was developed in *Instituto de Ciências Nucleares Aplicadas à Saúde* (ICNAS) of *Universidade de Coimbra*, in partnership with *Laboratório de Instrumentação e Física Experimental de Partículas* (LIP) Coimbra, within the High Resolution Brain PET project (HiRezBrainPET).

PET imaging with long acquisition times is susceptible to movement artefacts that can compromise image quality and lead to image misinterpretation. It became clear that to avoid these errors in a high-resolution brain PET, the current immobilization techniques used weren't enough, as movement bigger than the image resolution of the system results in image blurring, and sub-millimetric immobilization of the head was needed.

Through the study of the state-of-the-art head immobilization devices and techniques, it was chosen that the method to follow and develop would be the immobilization of the maxilla through a mouthpiece connected to an external frame, a method often used in radiotherapy. The immobilization device was developed for head and neck MRI exams, which can provide us reliable movement information easily, without exposure to ionizing radiation, but the same technique can be also used with the HiRezBrainPET scanner.

Digital tests were created to apply and optimize two algorithms: the SURF algorithm and the Block Matching algorithm, and software was developed based on these to analyse MRI images and detect any movement with accuracy and precision. To calibrate and test this software for translation and rotational displacements, a head-like phantom was developed, together with a motion inducing device, to produce known translations and rotations.

The immobilization device was projected in Computer-Aided Design (CAD) and made in the LIP Coimbra facilities. It was produced in POM-C polymer due to its physical properties, particularly its MRI compatibility, strength, rigidness, melting temperature, and monetary cost.

The final tests to the immobilization device were performed in the MRI system at ICNAS. The final immobilization results consist in the movement analyses of three different movement sequences performed with: the volunteer at rest, the volunteer applying big periodical movements and the volunteer applying small periodical movements (with 50% of the force applied in the big movements). Every movement sequence was performed with and without the immobilization device and imaged in the sagittal and axial planes. Each sequence was imaged for 2 minutes and 26 seconds.

For the forced small and big movements sequences, the decrease in movement was more noticeable than in the rest sequences. Rotations both in the sagittal and the axial plane, in both sequences, were reduced approximately 89%. In the sagittal plane, the head translational movements presented a reduction of 81% to 88%, with the greater decreases showing for the YY image axis. In the axial plane,

these reductions in movement went from 88% to 90%, but the greater decreases were seen in the XX image axis. This data analysis allows the conclusion that this immobilization system strongly constrains every movement of the volunteer's head.

For the rest sequences, in both planes, angular head drift was reduced to approximately half of its non-immobilized values and the correspondent random errors were reduced in approximately 70% in the sagittal plane and in 86% in the axial plane. In the sagittal plane, head translational drifts were reduced in about 36% in the XX axis, and 85% in the YY axis and random errors decreased from approximately 200 μm to values around 40 μm in the XX and YY axis. In the axial plane, head translational drifts were reduced in about 32% in the XX axis, and in about 58% in the YY axis. Random errors also decreased from approximately 40 μm to values around 20 μm in the XX and YY axis.

Therefore, it seems clear that this immobilization system is capable of performing sub-millimetric immobilization in head and neck exams.

Keywords: Head immobilization; Magnetic Resonance Imaging; Positron Emission Tomography; Image processing; Movement phantom; Experimental tests.

Table of Contents

Table of Figures

Table of Tables

List of Acronyms

Acronym	Meaning
BRW	Brown-Robert-Wells
BM	Block Matching
CAD	Computed Aided Design
CT	Computed Tomography
CRW	Cosman-Robert-Wells
CS	Consensus Set
DP	Direction Processing
FBP	Filtered Back Projection
FOV	Field Of view
FWHM	Full Width at Half Maximum
GTC	Gill-Thomas-Cosman
HP	Histogram Processing
ICNAS	*Instituto de Ciências Nucleares Aplicadas à Saúde*
IGRT	Image-guided Radiotherapy
keV	Kiloelectron Volt
LGK	Leksell's Gamma Knife
LIP	*Laboratório de Instrumentação e Física Experimental de Partículas*
LOR	Line Of Response
MAD	Mean Absolute Difference
ME	Motion Estimation
MR	Magnetic Ressonance
MRI	Magnetic Ressonance Imaging

Acronym	Meaning
MSAC	M-estimator SAmple Consensus
MSE	Mean Square Error
MSS	Minimum Sample Space
MV	Motion Vector
NP	Neighbourhood Processing
PET	Positron Emission Tomography
PMMA	Poly(methyl methacrylate)
POM-C	Polyoxymethylene Copolymer
RANSAC	RANdom SAmple Consensus
RF	Radio Frequency
RPC	Resistive Plate Chamber
SNR	Signal to Noise Ratio
STD	Standard Deviation
SURF	Speeded Up Robust Features
UC	*Universidade de Coimbra*
2D	Two-dimentions

Chapter 1
Introduction

Positron Emission Tomography (PET) is a medical imaging technique in which a radiopharmaceutical, with affinity to a certain biological process in study, is marked with radioactive nuclei that undergo β^+ decay (positron emission) and is administrated to a patient. The molecules in this radiopharmaceutical will bound in the organ or tissue being studied and, once there, will emit positrons from the same point in space for as long as their bound lasts.

The emitted positrons will then annihilate with the electrons of the surrounding medium (as the positron is the antiparticle of the electron), producing two photons that are emitted in almost opposite directions with approximately the same energy (~511 keV). PET scanners detect the pairs of photons produced and form a Line Of Response (LOR) between them, that should pass near the point of annihilation. By acquiring large quantities of these LOR's, PET scanners can reconstruct the spatial distribution of the annihilation points and consequently the distribution of the radiopharmaceutical in the organism. (Couceiro, 2014)

Given the static nature of the images obtained in PET, we realize that patient movement during a PET exam, such as head movement, rises problems in the veracity of the computed radionuclide's spatial distribution in the organism. Since PET exams are obtained in the many minutes range, if the tissue to which the radiopharmaceutical is bound is moving, positrons that should be emitted in the same point in space are emitted in different spatial locations. Then, a PET image of a moving body corresponds to the average motion of that body during the exam, which usually results in blurred images (Pépin, et al., 2014).

Spatial resolution is an important metric, often used to evaluate the performance of PET systems, essential to achieve clinical performance in areas like oncology, neurology, and cardiology. In more common language, spatial resolution refers to the ability of a system to distinguish two closely placed objects. It is normally measured using the Filtered Back Projection (FBP) algorithm to reconstruct point source scans, calculating then the Full Width at Half Maximum (FWHM) values along the different directions (Gong, et al., 2016).

This spatial resolution is what usually restricts the utilization of PET scanners for brain imaging. As the normal spatial resolution of these scanners is around 5-6 mm, other brain imaging techniques with spatial resolutions below 1 mm have been used in detriment of PET. A High-Resolution Brain PET scanner has been studied and developed to combat this paradigm. The HiRezBrainPET project, in which this project is included, developed a brain PET based on Resistive Plate Chambers (RPC) detectors (Fonte, et al., 2018) with the objective of achieving sub-millimetric spatial resolution. Besides the obvious advantages in resolving smaller brain structures, like the ones involved in neuropsychiatric

diseases, this property could have benefits in brain strokes diagnosis and consequent therapy orientation and in central nervous system tumour staging, allowing better radiotherapy and surgery planning. This system also aims to be used in high-resolution imaging of other organs, like the heart and the liver. One of its most distinctive qualities is the low cost of the RPC technology, which could allow for a rapid expansion.

So, if the spatial resolution of a brain PET system is in the range of the millimetre or below, any movement from the patient that approaches this resolution level can cause blurring artifacts in the images being obtained, diminishing their overall resolution and compromising tumour detection. That brings us to the main objective of this Dissertation, to develop an immobilization device that can avoid movements in the sub-millimetric range, maintaining the precision of high-resolution brain PET systems.

To test such a device in a PET scanner would be difficult, considering the radiative nature of the exam itself and the lack of reference scanners with the required resolution. Therefore, the immobilization device was designed and tested in an imaging system where movement is measured in an easier way, without the need to subject volunteers to radiation, Magnetic Resonance Imaging (MRI). Movement artefacts in MRI can likewise be noticeable and degrade image quality, especially in acquisitions with low Signal to Noise Ratio (SNR) or in imaging small objects (Havsteen, et al., 2017). In our case, we are aiming to obtain structural images, that provide us information of the brain's shape, size and structures, and present a strong SNR.

The Dissertation is divided into seven chapters, the present one being the first. The second chapter is dedicated to the state-of-the-art immobilization systems. The third chapter is devoted to the development of the immobilization device and the fourth and fifth ones to software development and testing. In the sixth chapter, the results of the tests of the immobilization device will be presented, and in the seventh one, they will be discussed to withdrawal the final conclusions.

Chapter 2
State of the Art Immobilization Systems

2.1 Conceptualization

Although there are some immobilization techniques and devices developed specifically for PET and MRI head and neck exams, the ones that offer more precision are designed and developed for use in radiotherapy. Radiotherapy is a therapy used mainly in the treatment of oncologic diseases, where ionizing radiation is delivered to a tissue to kill cancer cells, controlling and, hopefully, eradicating tumours. It is used in approximately 75% of head-and-neck cancer patients and it requires highly precise treatment plans, as head-and-neck tumours are usually situated next to vital organs and structures in the human body. These treatment plans are usually administrated in multiple doses of radiation called fractions, over a specific timespan, trying to minimize the dose given to the surrounding tissues at risk, while maintaining an adequate dose given to the tumour. In theory, these plans provide effective results, but they are highly dependent on accurate geometric dose delivery. (Bruijnen, et al., 2018)

Motion of the tumour and the surrounding organs during the delivery of a designed treatment plan may result in tumour under dosage and over-exposure of healthy tissues (Bertholet, et al., 2019). This motion is usually categorized into two different types: interfraction motion and intrafraction motion.

Interfraction motion is related to the day-to-day differences in the patient setup and patient's anatomy, such as weight loss and tumour growth or shrinkage, and is normally monitored using Image-guided Radiotherapy (IGRT) before the treatment delivery (Bertholet, et al., 2019). Intrafraction motion is classified as the motion of the targeted volume in a shorter time span, referring to the setup's accuracy during a single fraction of the treatment plan, (Hoskin, et al., 2008) and usually requires immobilization of the patient to be mitigated. This last type of motion is the one we are interested in when looking for immobilization devices and techniques applicable to PET and MRI imaging, given the non-fractionated nature of the exams we're aiming to perform, and is the only one we will be referring to when mentioning motion from this point on.

The geometrical deviations we define as motion are also divided into two categories: systematic errors, derived from the treatment's preparation, and the random errors, that appear from the treatment's execution (Herk, et al., 2000). A systematic error is defined as a deviation that occurs in the same direction and is of similar magnitude during all treatment's fractions, being the mean error over the course of the treatment when we consider an individual, and an indication of the spread of individual mean errors when we consider a population (Hoskin, et al., 2008). A random error is defined as a deviation that can vary in direction and magnitude for each delivered fraction, being the standard deviation (SD) of the measured errors throughout the treatment for an individual, and the mean of these

individual random errors for a population. The intrafraction motion deviations we are looking for during treatments and exams are included in this last type of error. (Hoskin, et al., 2008)

In the remaining of the present Chapter, the existing immobilization systems and techniques in use today will be presented and evaluated through their capability for stopping intrafraction motion. From the frameless to the invasive and non-invasive frame-based systems and techniques for immobilization of the head, we will compare them and discuss the best possible solution for the problem in our hands.

2.2 Frameless Immobilization Systems

2.2.1 Standard Thermoplastic Masks

When talking about head immobilization techniques and devices, we must promptly discuss thermoplastic masks. First, because they are the most widely used of all the devices and techniques: they are a non-invasive technique, relatively cheap and easy to use. Second, because they are incorporated in all sorts of frameless immobilization systems, as we will see in the paragraphs below.

As the name hints, thermoplastic masks are made of a thermoplastic material. This means that this material changes its mechanical properties with different temperatures. Most thermoplastic masks become malleable when submerged in hot water for some minutes and are then moulded over the patient's face by technologists, shaping it to the patient's facial contours, and clipped into the base plate under the patient's head. The most important areas to mould are the chin, nose and inter-cantal area, as these are the crucial points to immobilize in the face to avoid it from moving. When the material hardens, it also shrinks in outside circumference by about 2% of its initial value, and it's then removed from the patient's face and trimmed for extra comfort. These masks are not meant for single-use and are only replaced by a new one when they start to not fit tightly to the patient's head (If, for example, the patient lost weight, or the mask simply became loose). (Thornton, et al., 1991)

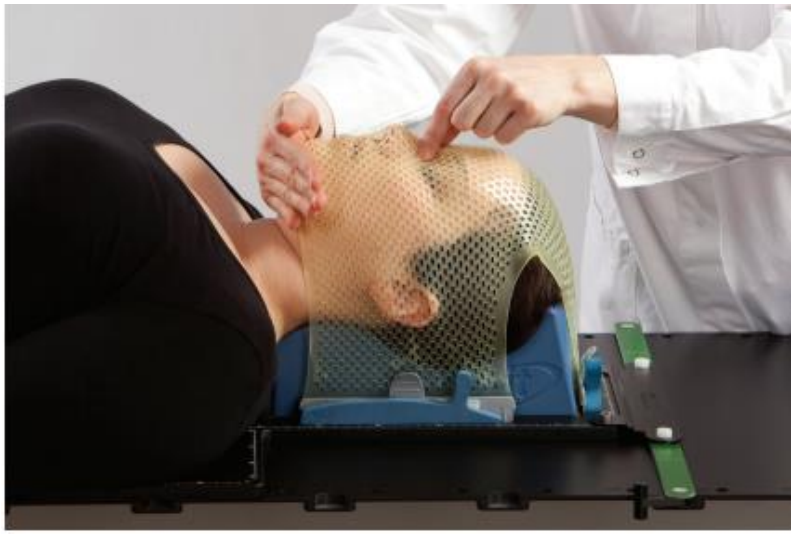

Figure 2.1- Thermoplastic mask and setup, from ORFIT Industries.

Thermoplastic masks are produced by numerous companies all over the world and come in a large variety of shapes and sizes, with multiple different mechanisms to connect to the base plate, different headrests, etc…. Here, we studied and compared the most notable brands and setups between them to reach a conclusion as for their overall immobilization precision.

The first comparison we are going to analyse is between systems from the same manufacturer. Orfit provides a white paper with a summary of the results of 2 comparable studies. In the first study, three different masks were compared: (1) Orfit HP Efficast micro-perforation 1.6 mm mask, (2) Orfit HP Efficast maxi perforation 2.0 mm and the (3) Orfit HP Nanor micro+ perforation 1.6 mm mask. The results of the translational and rotational intrafraction random error are shown in Table 2.1 (Mesch, et al., 2019). Note that pitch, yaw and roll rotations are rotations in turn of the x, y and z axes, respectively.

TABLE 2.1 - INTRAFRACTION TRANSLATIONAL AND ROTATIONAL RANDOM ERRORS, IN ALL AXES, OF THE THREE DIFFERENT SYSTEMS DEPICTED IN (MESCH, ET AL., 2019).

	Translation (mm)			Rotation (º)		
System	X (lat.)	Y (long.)	Z (vert.)	Pitch	Yaw	Roll
(1)	0.44	0.45	0.45	0.15	0.15	0.18
(2)	0.40	0.49	0.37	0.18	0.12	0.16
(3)	0.48	0.58	0.32	0.17	0.17	0.17

In the second study, different head supports were compared, all utilizing the same Orfit Hybrid mask, a reinforced mask for ultra-high precision. The results of the translational and rotational intrafraction random error are shown in Table 2.2. As these two studies are based on the same research method, a direct comparison between both is possible. (Meessen, et al., 2017)

TABLE 2.2 - INTRAFRACTION TRANSLATIONAL AND ROTATIONAL RANDOM ERRORS, IN ALL AXES, OF THE THREE DIFFERENT SYSTEMS DEPICTED IN (MEESSEN, ET AL., 2017).

	Translation (mm)			Rotation (º)		
System	X (lat.)	Y (long.)	Z (vert.)	Pitch	Yaw	Roll
(1)	0.24	0.25	0.18	0.24	0.12	0.26
(2)	0.27	0.31	0.16	0.18	0.16	0.22
(3)	0.36	0.24	0.31	0.27	0.18	0.26

In 2014 a study was published that compared three different head and neck immobilization systems. These three systems included: (1) Orfit AIO base plate, standard neck supports and a pre-cut five-point reinforced Efficast mask, fixed to the base plate with L-shaped profiles, (2) Q-Fix AccuFix Cantilever Board Featherline base plate with adjustable shoulder locks, Vacfix neck support and a U-Frame Aquaplast mask for the head, and (3) Aquaplast mask covering the head and shoulders, fixed to a Vacfix cushion at eight points with velcro strips. Each system was used for the treatment of 14 patients and was evaluated separately during a period of 12 months, one after the other. The translational and rotational intrafraction random error detected in these systems is shown in Table 2.3. (Hansen, et al., 2014)

TABLE 2.3 - AVERAGE INTRAFRACTION TRANSLATIONS AND ROTATIONS AND RESPECTIVE STANDARD DEVIATION, IN ALL AXES, FOR THE THREE DIFFERENT SYSTEMS DEPICTED IN (HANSEN, ET AL., 2014).

	Translation (mm)			Rotation (º)		
System	X (lat.)	Y (long.)	Z (vert.)	Pitch	Yaw	Roll
(1)	0.9 ± 0.1	1.0 ± 0.1	1.3 ± 0.1	0.7 ± 0.1	0.8 ± 0.04	0.5 ±0.04
(2)	1.7 ± 0.2	2.4 ± 0.2	2.1 ± 0.3	1.2 ± 0.1	1.1 ± 0.1	1.3 ± 0.1
(3)	1.4 ± 0.1	1.2 ± 0.1	1.8 ± 0.3	1.1 ± 0.1	1.0 ± 0.1	0.8 ± 0.04

Other studies involving standard thermoplastic masks show that the average intrafraction head motion of patients utilizing these is around 0.7 mm. (Babic, et al., 2018) (Ramakrishna, et al., 2010)

With the values shown until now from the most recent and notable standard thermoplastic masks and systems in the market, it's clear that values for their immobilization precision can present variations between institutions, research methods, fixation equipment used, etc... However, although the results do not converge to a single precision value, it's plausible to say that thermoplastic masks show a good promise when looking for sub-millimetric immobilization equipment.

Figure 2.2 - Standard, Hybrid and Open-face masks (respectively from left to right) Setups from ORFIT Industries.

2.2.2 Open-Face Thermoplastic Masks

Open-face masks are a type of thermoplastic mask designed to help claustrophobia suffering patients as well as improve comfort. These masks present an open space in the face region of the mask, as is shown in Figure 2.2, and rely mainly on the forehead and chin immobilization to secure the patient's head.

In a 2013 study, standard single-layer thermoplastic masks are compared to open-face masks. Forced movement was compared between both types of masks and was shown that the open-face masks allow, on average, a maximum lateral movement of 5.5 ± 2.5 mm and a maximum longitudinal movement of 4.2 ± 2.0 mm. For standard masks, the maximum allowed movement was 4.2 ± 3.0 mm lateral and 3.3 ± 1.5 mm longitudinal. With the open-face mask, the involuntary head motion was found to be 1.0 ± 0.6 mm with 0.4° ± 0.4° rotations in volunteers, and 0.8 ± 0.3 mm with 0.4° ± 0.2° rotations in patients during treatment. (Li, et al., 2013)

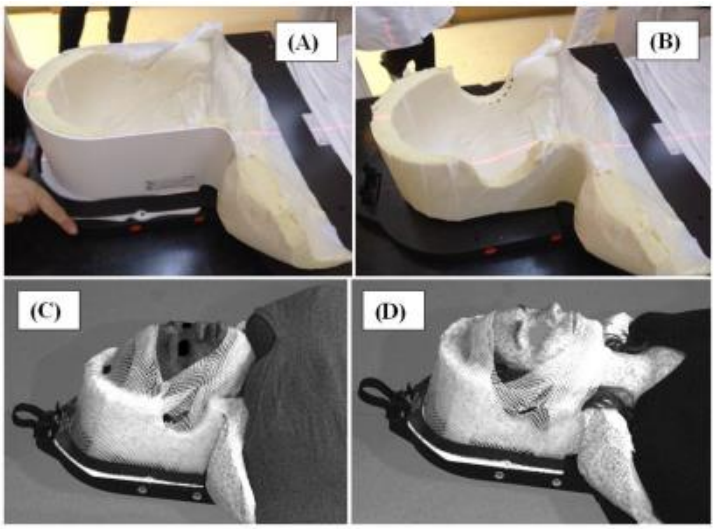

Figure 2.3 - Construction of the head mold in (A) and (B) for the use of the Freedom System in two volunteers (C) and (D). Image taken from (Li, et al., 2015).

In 2015, another study developed what they called the "Freedom System". This system consisted of a head mold and an open face mask made from a pre-cut thermoplastic material as we can see in Figure 2.3. This system was able to support the head, neck and shoulders of the patient through the head mold, which was marked with alignment lines for initial head alignment in the setup. (Li, et al., 2015)

The treatment started with the patient being initially positioned by matching the cast lines with room lasers on the couch extension board. The head position was manually adjusted with the mask off to minimize rotations and the mask was then placed and locked. The new residual rotations that resulted from this procedure were corrected using the Freedom rotational adjustments and standard couch

rotation. This system presented intrafraction translations and rotations of 0.6 ± 0.3 mm and 0.3° ± 0.1°, with a maximum range of forced head motion of 2.2 ± 1.0 mm. (Li, et al., 2015)

2.2.3 Other Frameless Immobilization Systems

Other frameless immobilization devices include simpler designs and concepts when comparing them to thermoplastic masks, relying mostly on patient cooperation to avoid large movements.

The most used non-thermoplastic frameless immobilization device is the classic "U-shaped" head holder, as we see in Figure 2.8, normally accompanied by some type of foam cushioning. A study from 2011 showed that for PET/CT and MRI exams with 20-30 minutes duration, the head holder couldn't restrict non-forced movements up to 4 mm (Li, et al., 2011). In 2010, another study developed a custom frameless and mask-less head mold to immobilize patients for radiotherapy, obtaining non-forced axial translations and rotations during 20 minutes of monitoring no greater than 1.2 mm and 0.5º (Cervino, et al., 2010).

Custom neck casts have also been used to immobilize patients but, without the help of any thermoplastic mask, it presented values of intrafraction motion of 2.0 ± 1.4 mm (Engelsman, et al., 2005).

2.3 Frame-based Immobilization Systems

Until now, we have only discussed frameless immobilization systems, systems that don't require a connection to a rigid frame to immobilize the patient. This is because frame-based systems, used in radiotherapy, take up a lot of space and don't usually fit inside PET or MRI systems, use materials that are not PET or MRI friendly and are much more expensive than frameless systems. Nevertheless, frame-based immobilization is the most precise form of immobilization, and it's divided into invasive and non-invasive immobilization.

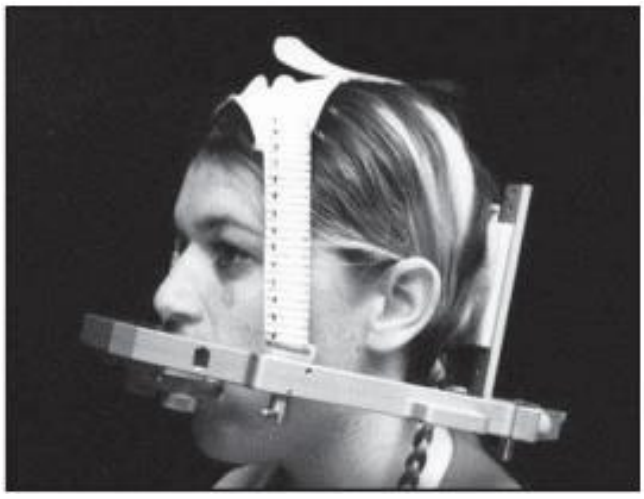

Figure 2.4 - Gill-Thomas-Cosman (GTC) Relocatable Stereotactic Frame. Image from (Chelvarajah, et al., 2004).

2.3.1 Invasive Frame-based Immobilization Systems

Invasive immobilization devices operate by attaching a rigid frame, that is usually a ring located on the lower portion of the head to limit interference with the treatment, with screws into the patient's cranium. Some of the most used, reliable, and stable invasive immobilization frames are the Cosman-Robert-Wells (CRW) frame and the Leskell frame. (Haffty, et al., 2018)

The CRW frame is a developed modification based on the original Brown-Robert-Wells (BRW) arc-radius design, utilizing its existing fixation and fiducial components to localize and verify target data, and needs to be applied to the skull by fixation pins, under general anaesthesia and endotracheal intubation (Pell, et al., 1991). The original BRW frame defines a probe insertion trajectory via the depiction of four different angles, positioning the point of a probe in the calculated 3D coordinates inside the brain, in the desired structure, through a small hole in the skull. Different probe insertion trajectories to the same target point require entirely different coordinates for each different trajectory. The CRW frame permits a more intuitive specification of the probe insertion trajectory by introducing a different arc-guidance system. (Khedr, et al., 2018), (Pell, et al., 1991), (Haffty, et al., 2018). The Leksell frame functions very closely to the CRW frame, although it is lighter and fully MR compatible. It is applicable in various treatment-planning systems, but it's mostly used in Leksell's Gamma Knife (LGK) system (used in radiosurgery), being the only frame in clinical use (Haffty, et al., 2018).

The BRW frame shows values of intrafraction motion of 0.4 ± 0.3 mm as the intrafraction motion measured in the CRW frame was 0.3 ± 0.21 mm (Babic, et al., 2018). Although this type of immobilization guarantees sub-millimetric immobilization precision, invasive methods are not viable in our area of research, as they are not practical, minimally comfortable, or logistically appliable.

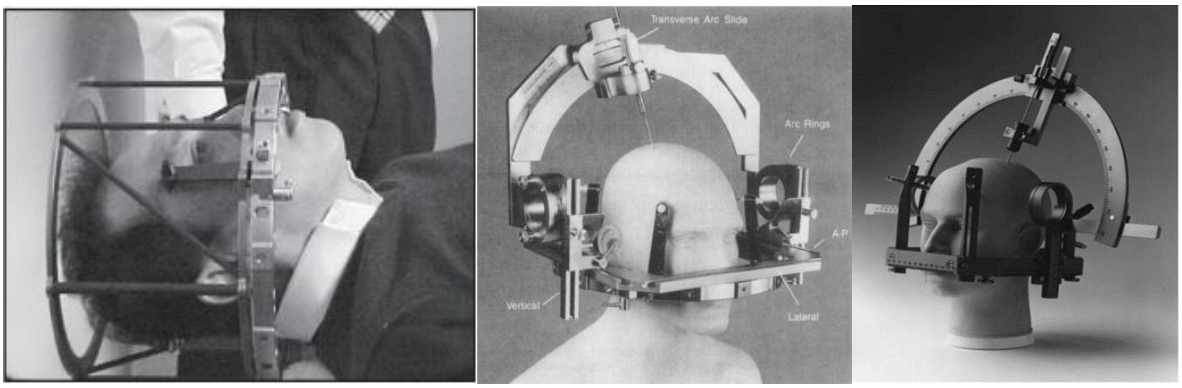

Figure 2.5 - BRW, CRW and Leksell frames, respectively from left to right. Images from (Chelvarajah, et al., 2004) and (Pell, et al., 1991).

2.3.2 Non-invasive Frame-based Immobilization Systems

Seeing the impressive precision of invasive frame-based immobilization systems, it's logical to attempt reproducing these same results in a non-invasive way, with improved comfort and logistics. For

that purpose, multiple non-invasive frame-based immobilization systems have been developed throughout the years, and the most effective and reliable ones in use today are covered in this section.

The first system we will discuss is the Gill-Thomas-Cosman (GTC) relocatable frame. The GTC relocatable frame consists of a standard head ring to which custom made dental and occipital impressions are attached, for stability and head support. These frames are more practical and comfortable than invasive fixed frames and are of quick release for cases of emergency. On the downside, the patients should have good dentition to be able to secure the mouthpieces. (Chelvarajah, et al., 2004)

In a 2017 study, this relocatable GTC frame was analysed and presented a mean intrafraction motion of 0.54 ± 0.76 mm, with pitch, roll and yaw rotations of -0.04º ± 0.35º, -0.01º ± 0.44 and -0.02º ± 0.48º, respectively. (Babic, et al., 2018)

A modified GTC (mGTC) frame was also developed, where the commercial occipital head supports were substituted with thin carbon concave support surfaces, to surround the base of the skull, shaped to hold and secure a patient-specific cushion. The first use of this mGTC frame was in 2002, and in 2011 a study was made to access its immobilization precision. This study concluded that, combined with implanted fiducials and an image-guidance system, this mGTC frame was capable of sub-millimetric precision of ± 0.66 mm, with rotational uncertainty no greater than 0.5º. (Winey, et al., 2012)

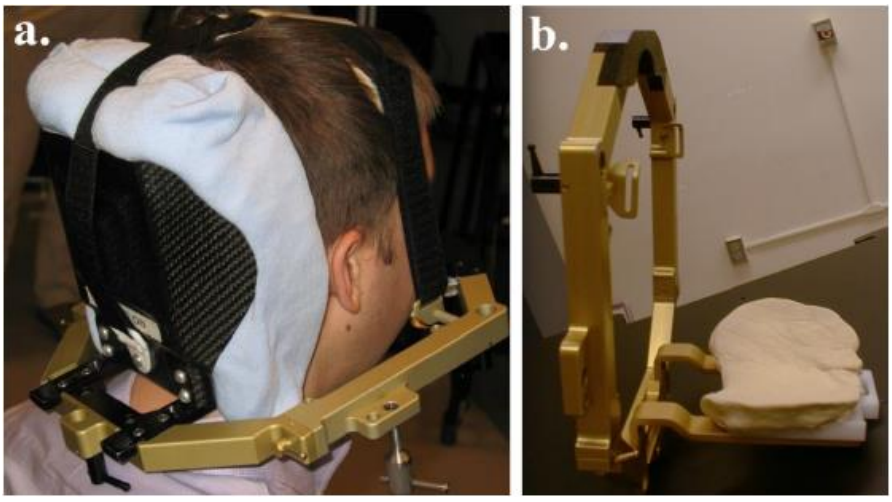

Figure 2.6 - A custom mGTC frame (a.) compared to a commercial standard GTC frame (b.). The black carbon support and the patient specific cushion can be seen in posterior part of the mGTC frame. Images were taken from (Winey, et al., 2012).

In another study, that analysed intra and interfraction motion for a variety of immobilization devices, the intrafraction patient motion recorded with only the mGTC frame was 0.9 ± 0.6 mm. (Engelsman, et al., 2005)

Another non-invasive frame-based immobilization technique is the immobilization of the patient's maxilla through the utilization of a patient-specific mouthpiece with incorporated mild vacuum suction on the upper palate. Vacuum suction is applied between the mouthpiece and the upper palate to

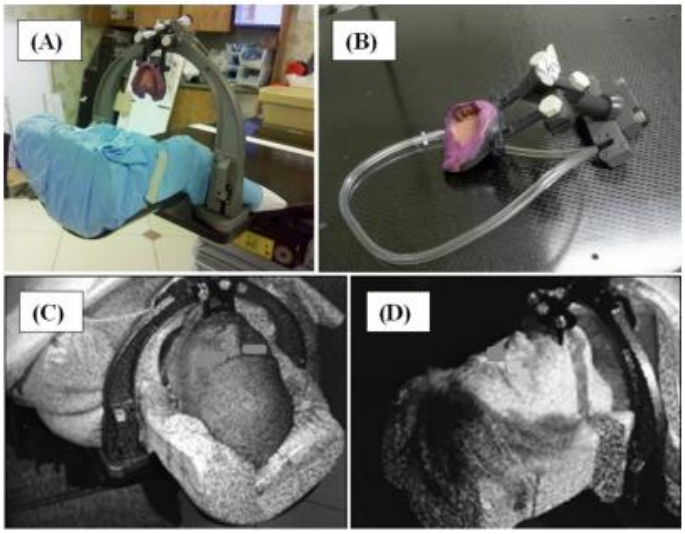

Figure 2.7 - Aktina Pinpoint System (A) and a custom made mouthpiece with vacuum suction (B). Two patients immobilized with the PinPoint system (C) and (D). Images were taken from (Li, et al., 2015).

assure that no biting is required to obtain tight contact. This mouthpiece is then connected to an external arc metal frame, through an adjustable rigid connecter, which in turn is fixed to a couch board locked on the treatment couch. (Li, et al., 2011)

There are multiple systems in the market utilizing this immobilization technique and different companies present different head supports, mouthpieces, frame designs, etc…, but the most studied between them is the PinPoint® system, by Aktina Medical, that we can see in Figure 2.7. In a 2011 study, the PinPoint immobilization system, with both the head mold and the mouthpiece with vacuum suction, showed the ability to restrict patients' intrafraction motion to the level of frame-based immobilization (Li, et al., 2011). The average head-motion magnitude of the patients was 0.3 ± 0.2 mm with rotations of 0.2º ± 0.1º (Li, et al., 2011). In another study made in 2017, the value of the average intrafraction motion measured with the PinPoint system was 0.45 ± 0.33 mm in stereotactic radiosurgery patients and 0.66 ± 0.33 mm in hypofractionated stereotactic radiotherapy patients. Pitch, roll and yaw rotations measured in this system were 0.07º ± 0.27º, 0.02º ± 0.29º, 0.03º ± 0.43º in hypofractionated stereotactic radiotherapy patients and 0.00º ± 0.24º, 0.03º ± 0.13º, 0.22º ± 0.27º in stereotactic radiosurgery patients, respectively. (Babic, et al., 2018)

2.4 Immobilization Techniques and Devices used in Brain PET and MRI

Since the conventional design of PET and MRI scanners feature a small-bore size, developing precise immobilization devices that can fit in this space becomes difficult. Immobilization devices and techniques commonly used in brain PET are rudimentary and, in most cases, unreliable. The great majority of the literature referring to brain imaging guidelines mentions the importance of patient immobilization, and the procedures to mitigate it normally include informing the patient about the need

to remain still and avoid voluntary motion of the head (Varrone, et al., 2009), (Waxman, et al., 2009). Exact positioning of the patient's head is usually attempted using external markers, such as cross-shaped laser marking systems some scanners have (Trembath, et al., 2015). Other techniques commonly used include the conventional head holders, that are mentioned in 2.2.3 page 8, foam inserts and elastic headbands to fix the head (Trembath, et al., 2015). We can see an example of that in Figure 2.8.

Figure 2.8 - Volunteer subject properly positioned in head holder using a laser system, cushions, chin strap, and forehead strap. Image taken from (Trembath, et al., 2015).

Nevertheless, in a 1991 study researching a head immobilization system for radiation simulation, a headrest, MRI insert, a rigid baseplate to receive the MRI insert and allow rigid mounting on the therapy couch, and a couch adapter were developed to position and support thermoplastic masks (which we'll discuss next) in MRI, Computed Tomography (CT) and PET systems. This system's schematics are shown in Figure 2.9 and its development is explained in the original article. The PET couch adapter was constructed of machined aluminium and supports the cranium at an extended distance from the couch maintaining the same indexing system used in MRI. (Thornton, et al., 1991)

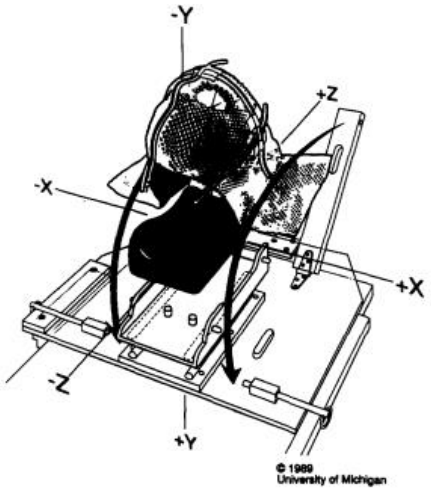

Figure 2.9 - Schematic of the mask immobilization system, from (Thornton, et al., 1991).

Three-dimensional motion analysis was performed on this device and the detailed procedure is described in the original article (Thornton, et al., 1991). The results were measured according to two different analyses: calculation of the average movement of individual anatomic reference points relative to their original locations, and determination of the centre of mass translation of all anatomic points relative to the centre of mass on the original frame. The average displacement of the centre of mass of the anatomical points from their original positions was 2.5 mm, with a standard deviation of 1.4 mm. The absolute difference in the positions of the anatomical markers was 3.8 mm, with a standard deviation of 1.3 mm. (Thornton, et al., 1991)

In current MRI systems, in addition to the limited bore size, there are also space restrictions forced by the head and neck radiofrequency (RF) coils (which were not mentioned in the previously referenced study). These RF coils used to perform head and neck imaging don't leave much room to spare inside them when a patient is placed in position, so all the devices developed to immobilize the patient's head have to either fit inside the coil or find a way around it.

Figure 2.10 - Pearl Technologies MRI positioning system.

Besides the usual foam cushioning, Pearl Technology developed a different immobilization solution to be used in head and neck MRI exams. This system consists of adjustable and inflatable padding between the patient's head and the RF coil, just as we see in Figure 2.10. The patient's head is first covered with a single-use sanitary cover and then the padding system is placed. The patient is placed inside the coil and the pads are pumped with air until a stable and comfortable position is reached.

A study was conducted to access the precision of this positioning aid, compared to the traditional foam cushioning. The results evaluated the movement permitted in the three different planes (transverse, frontal, sagittal) and the extent of cumulative positional change by the subjects during the sequences was compared for the two different positioning aids and is shown in Figure 2.11. (Eulenburg, et al., 2017)

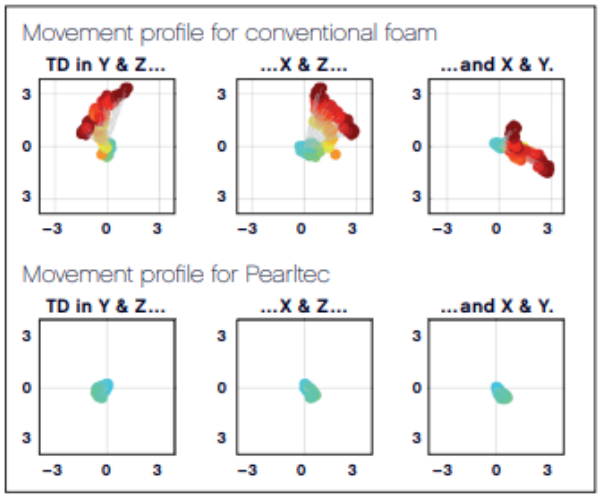

Figure 2.11 - Movement profiles of the patients using conventional foam cushioning (top) and Peartech (bottom). Taken from (Eulenburg, et al., 2017).

In 2019, another case study researching a practical immobilization solution for head and neck MR-only radiotherapy, where MR is the diagnostic image modality, developed a thermoplastic mask holder to fit inside a standard head and neck RF coil, as we can observe in Figure 2.12. In the movement restriction tests performed to this system was found that forced movements inside the RF coil during MRI acquisitions rarely exceeded 1.5 mm. (Mandija, et al., 2019)

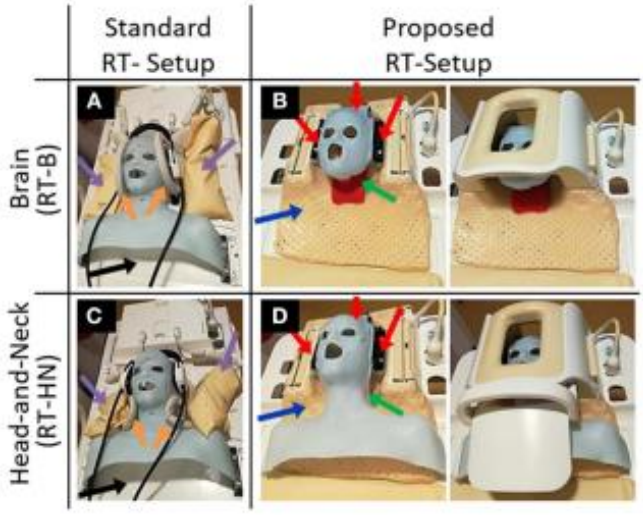

Figure 2.12 - Standard and proposed setups for MR imaging in treatment position for brain (A,B) and head-and-neck (C,D) radiotherapy patients. The blue arrows indicate the thermoplastic holder, the red arrows indicate the anchor points of the immobilization masks, and the green arrows indicate the neck support. Image and data from (Mandija, et al., 2019).

2.5 Overview

One specific type of the reviewed immobilization devices was followed and develop to create our own system, according to our problems and specifications. The MRI scanner we'll use to develop this case of study is the MAGNETOM Prisma Fit, with a Head/ Neck 64 Coil, depicted in Figure 2.13, from Siemens, that has a 60 cm bore diameter. From the previous analysed techniques and devices, we could only consider three approaches to obtain sub millimetric head and neck immobilization: the use of invasive frame-based techniques, the use of thermoplastic masks or the use of non-invasive frame-based techniques.

Figure 2.13 - Head/Neck 64 Coil, open (left) and closed (right), from Siemens.

The use of invasive techniques is right away excluded: besides the impracticality of the method itself, the pure size of the immobilization frames would hardly fit inside the head and neck coil and securing them outside the coil would be hard. To use thermoplastic masks, we would have to develop some type of thermoplastic mask holder like we saw in Figure 2.12 (Mandija, et al., 2019), but even if we could develop it for the RF coil we will use, there would be no guarantee the thermoplastic mask itself could assure sub millimetric immobilization (some studies show intrafraction motions of up to 1.5 mm, as shown in page 14). Finally, analysing the non-invasive frame based systems, since both the GTC and the mGTC frames need peripheral rigid frames, and we do not have that space inside the head and neck RF coil, we are left with a frame-based non-invasive system using an external fixed frame, similar to Aktina's PinPoint system, as the optimal system to explore and develop.

By itself, this external frame system also rises some problems that will be approached and explained in Chapter 3, along with the solutions devised to overcome them.

Chapter 3
Immobilization Device

3.1 The Concept

To develop a system like the one discussed in 2.5 for MRI, based on the devices currently in use, discussed in page 15, there are some problems we need to address. The first problem lies in the fact that the bore diameter of the MRI and the head and neck coil itself enforce space restrictions that don't usually appear in radiotherapy. To solve this, considering the specifications of the MRI system and RF coil available, we will have to mould the external frame to fit inside the bore diameter and around the head and neck coil, in a way that allows a mouthpiece to access the inside of the coil, specifically at the patient's mouth. The second problem lies in the materials to be used. As we are working on an MRI system, ferromagnetic materials are out of the question, and we must find a material that both presents the mechanical stiffness to firmly secure the patient's head and is MRI-compatible. Other changes have been made in relation to the original systems. One includes the mouthpiece, where the vacuum system will not be included. The fixations in and off the frame are also different, along with the mechanics used to move the mouthpiece to the correct location. These last changes were made both to simplify the production and the system itself to diminish its overall cost, while trying to maintain the precision reported in other devices.

The project of the system idealized was developed in 3D Computed Aided Design (CAD) software, specifically the Inventor® CAD software from Autodesk®. The system is divided into three main parts: the base plate, the external frame, and the mouthpiece. Each of these main blocks can be decomposed into different parts that will be explained later in this chapter. The external frame will be fixed to the base plate, where the patient will be lying, and the mouthpiece will then be connected to the external frame above the RF coil, right where the mouth entrance of the coil is. The schematics of this device will be described throughout this Chapter, starting with the MRI constituents, as the spatial references for the design.

3.2 Magnetic Resonance Model

A model of the relevant mechanical characteristics of the MRI scanner was developed as a spatial reference for the system being designed, to specify and account for the space restrictions we would encounter in the real MRI system. The model is divided into three components: the model of the MRI scanner, to account for the bore diameter; the model of the Head/Neck 64 coil, to account for the restrictions the external frame and mouthpiece had to overcome; and the MR table, to reference the base plate positioning and to take into account the space restriction reflected in the bore diameter implied by its positioning. The entire assembled model can be visualized in Figure 3.1.

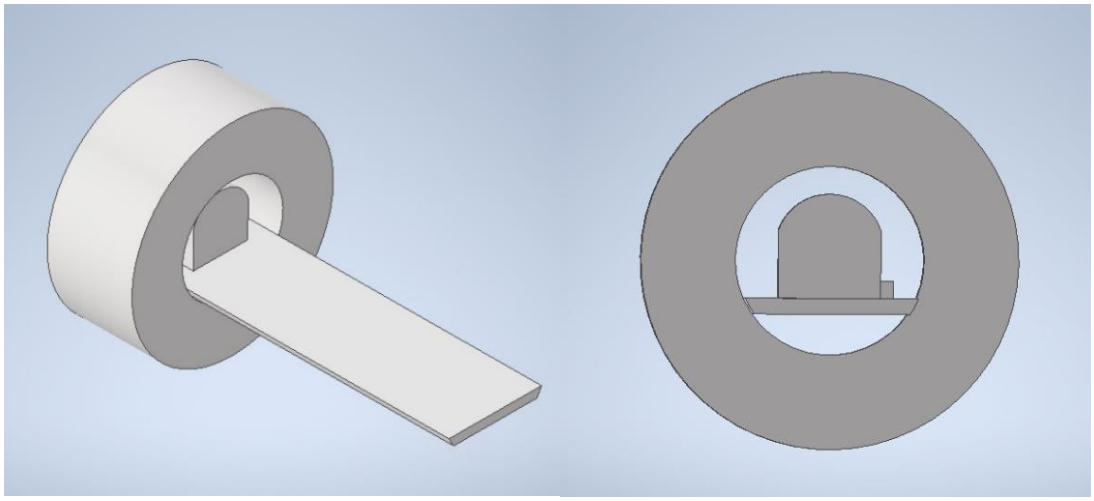

Figure 3.1 - The assembled model of the Magnetic Resonance system.

3.2.1 Magnetic Resonance Scanner

The MR model utilized in this project is based on Siemens's MAGNETOM Prisma Fit MRI scanner. The only important parameter to consider when designing this part was the 60 cm bore diameter, specified by Siemens. Therefore, the model is simply a hollow cylinder with the specified bore diameter and enough length to encompass the coil (50 cm), as shown in Figure 3.2.

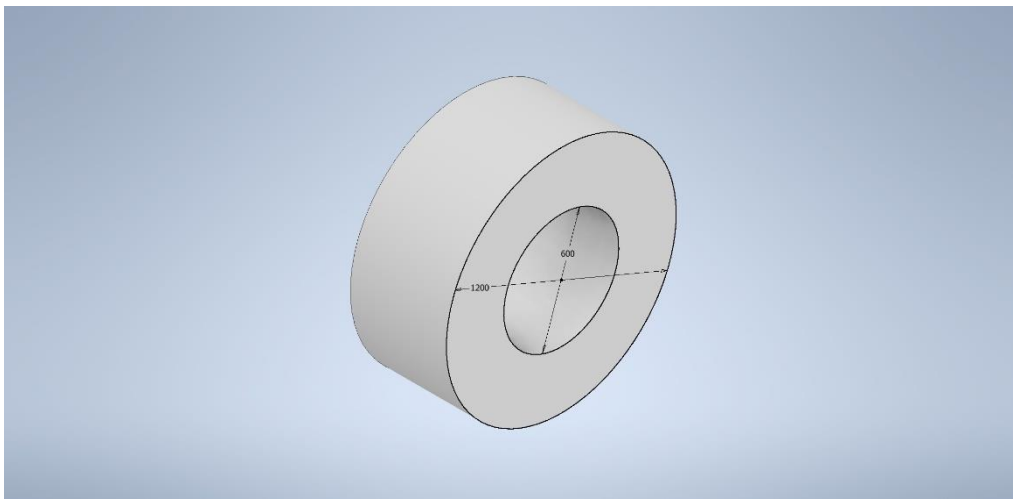

Figure 3.2 - The model of the Magnetic Resonance scanner.

3.2.2 Head/Neck 64 Coil

The process of modelling the head and neck coil was not so simple. Since the coil cannot be described as a simple solid like the MR scanner, as seen in Figure 2.13, we had to find a way to acquire its precise specifications. For this purpose, and since we could not get the specifications anywhere else, we submitted the head and neck coil to a Computerized Tomography (CT) scan. With this procedure, we could, in principle, precisely acquire every measure needed through the 3D Slicer software (Fedorov,

et al., 2012) and replicate it in CAD. This would assure the model functionality but given the streaking artefacts observed in the coil scan, like shown in Figure 3.3, together with other properties of the process, the coil had to be modelled differently and designed from scratch, however, based on the dimensions measured from the CT scan.

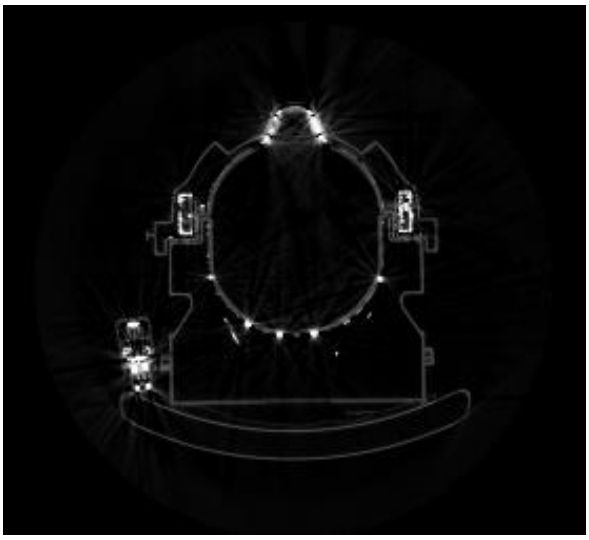

Figure 3.3 – Streaking artefacts in a plane of the CT scan of the Head and Neck 64 coil.

The Head/Neck 64 coil model is then comprised of a rectangular shaped solid with an upper curvature correspondent to the one observed in the real coil, with a rectangular hole in the position of where the mouth hole should be and a protrusion 4 cm wide to simulate the coil connector. This model is represented in Figure 3.4.

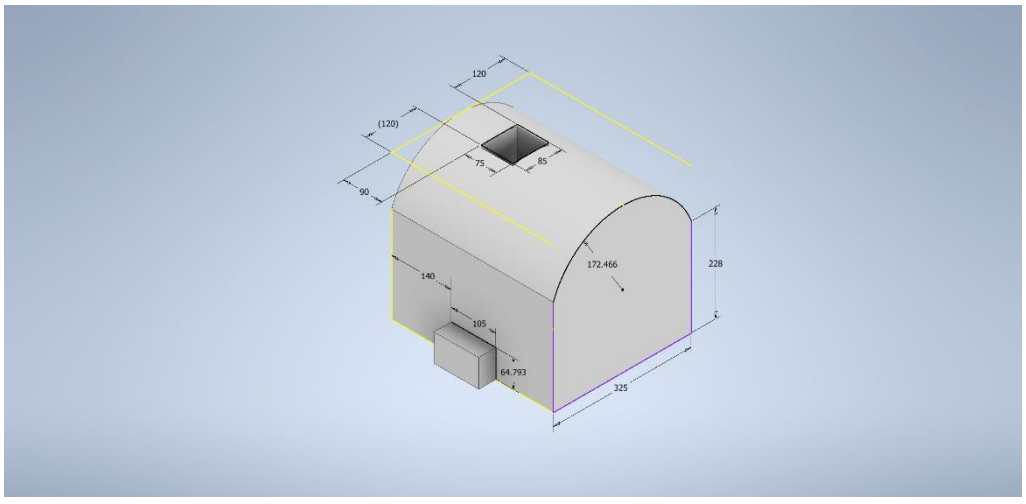

Figure 3.4 - The model of the Head and Neck 64 coil.

3.2.3 Magnetic Resonance Table

The MR table model is a simple one, standing just two details away from a perfect parallelepiped. The model of the table is a 205 cm by 55 cm rectangle with a small indentation on the upper face, shaped like the lower face of the MR coil model, and round lateral edges, to fit inside the MR scanner model in a specified position. The MR table model can be visualized in Figure 3.5.

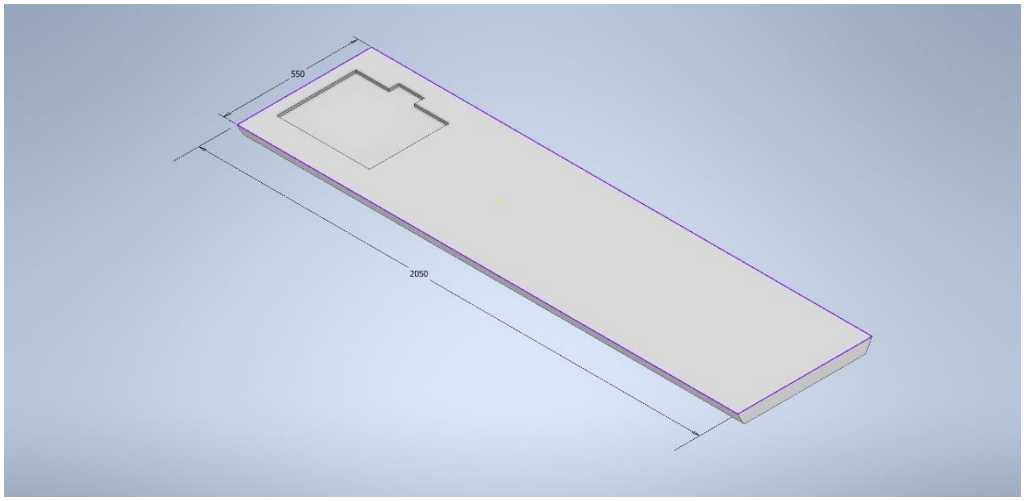

Figure 3.5 - The model of the Magnetic Resonance table.

3.3 Base Plate

The base plate's main purpose in this design is to firmly secure the external frame, providing a solid support. This base support is achieved by the force of the weight provided by the person being scanned when laid on top of the plate. It is composed of two parts, mainly because of practicality, and each part contains an opening for grabbing, making it easier to move. The main part of the base plate, the one that supports the frame, is moulded to the head and neck coil on its upper side, allowing the frame to be secured next to the coil, close to the mouth opening and avoiding contact between the frame and the patient shoulders. The external frame is connected to the base plate by two M8 butterfly wing

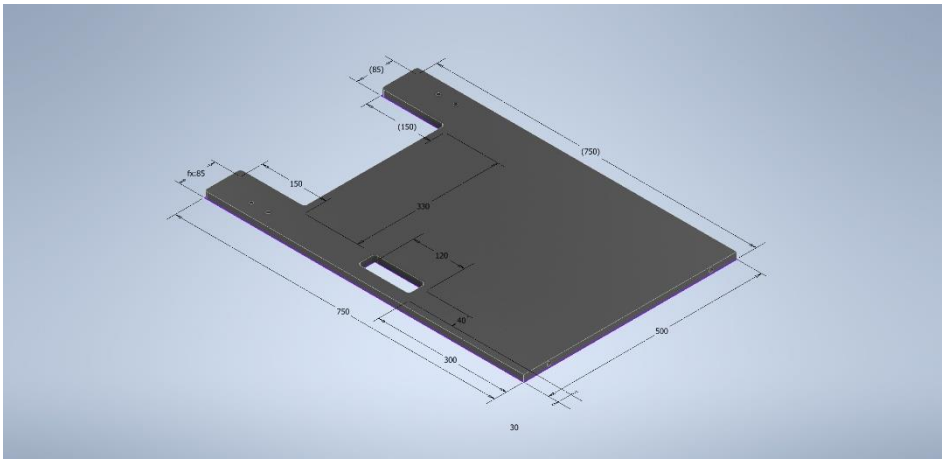

Figure 3.6 - The model of the main part of the base plate .

screws, to assure stability in the connection, positioned on the upper lateral sides, where the threaded holes are placed. The two parts of the base plate are connected through two cylindrical M10 insertions, seen in the lower part of the base plate. The base plate schematics and its details can be visualized in Figure 3.6 and Figure 3.7.

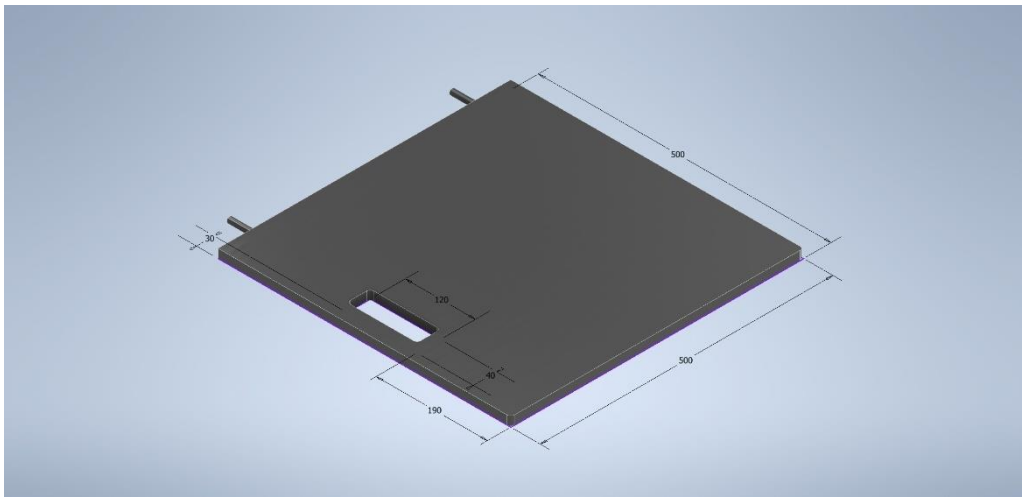

Figure 3.7 - The model of the secondary part of the base plate .

3.4 External Frame

The external frame is composed of five different parts, four of them being replicated to achieve symmetry, with its main objective being the positioning and tight holding of the mouthpiece to avoid movement from the patient. It consists of two frame bases, two frame arms, two upper screws and two lower screws, with one of this on each side of the base plate, and one central frame part (the only unique piece) that will hold the mouthpiece. With this frame design, the mouthpiece's position can be adjusted

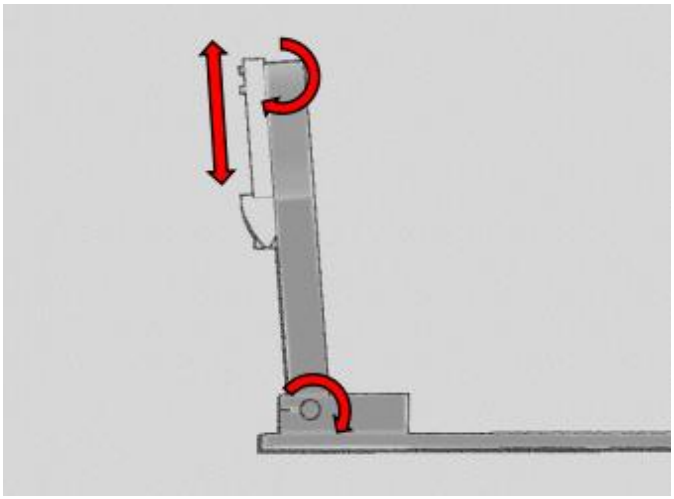

Figure 3.8 - Schematics of the external frame with its rotational and translational possible movements, pointed by red arrows.

at two different angles, 180 degrees on the lower one and 360 degrees on the higher one, and 85 cm in the vertical translational axis, as we see in Figure 3.8, allowing the adjustment to each different patient and, possibly, different MRI head and neck coils and setups.

3.4.1 The Frame's Base

The frame base consists of a solid parallelepiped 130 mm long, 40 mm wide and 40 mm tall. It has two M8 holes for the connection to the base plate through the insertion of two butterfly wing M8 screws, as mentioned prior. It also has a hole of 20 mm in diameter with the addition of a 2 mm gap that trespasses the solid sideways, for the connection to the frame's arm, and a third M8 whole placed precisely to secure it. The grip to the frame's arm is guaranteed through this mechanism, relying on the friction caused by the tightening of the M8 screw and it represents the first adjustable angle in the frame's system. We can see the mentioned details in Figure 3.9.

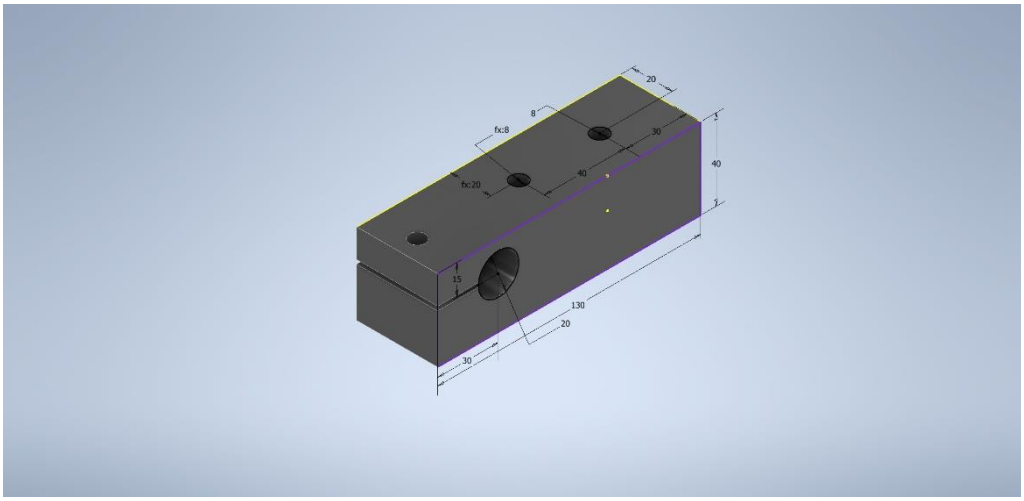

Figure 3.9 - The model of the external frame's base .

3.4.2 The Frame's Arm

The frame arm is a more detailed piece than the aforementioned. It is 345 mm tall, 40 mm long and 30 mm wide, presenting a straight section as well as a curved one. The straight section reaches an internal height of 225 mm and an external one of 235 mm. The curved section presents an internal arch of 172,5 mm and an external one of 202,5 mm, to match the curvature of the MRI head and neck coil. In the bottom section, the frame arm has an M10 hole crossing it sideways, where the lower screw (the connection between the arm and the base of the frame) will be inserted. In the top section, the frame arm has another M10 insertion, as we see in Figure 3.10, to hold the upper screw, the connection between the arm and the central part of the frame. This last connection is not a simple hole, it also presents a slope, complementary to the upper screw's shape, to provide extra stability. The whole scheme of the arm of the frame can be visualized in Figure 3.11.

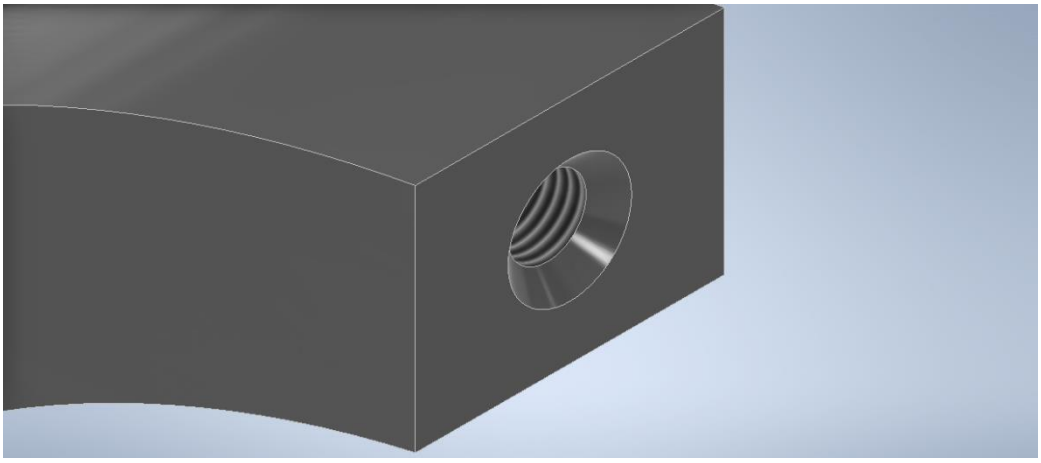

Figure 3.10 - Insertion for the connection between the frame's arm and its central piece, zoomed in.

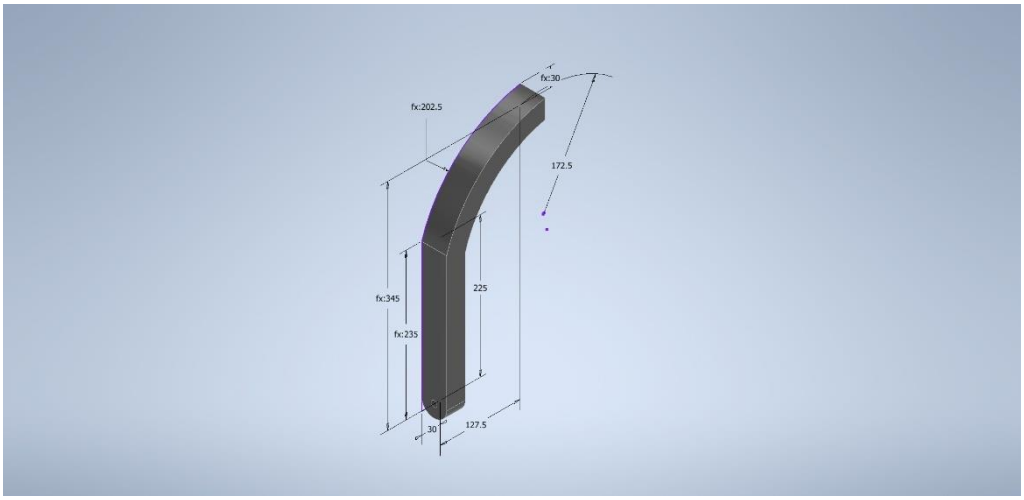

Figure 3.11 - The model of the external frame's arm.

3.4.3 The Frame's Central Piece

The central piece of this frame's design has multiple purposes. It is designed to hold and secure the mouthpiece, allowing its vertical translational movement, in one hand, and to rotate and secure it to the rest of the external frame in another, representing the second adjustable angle of the system. It is composed of a parallelepiped 83 mm long, 40 mm wide and 30 mm tall. It has a 12 mm diameter hole across its length together with a 2 mm gap in the same axis, to enable the piece to fit and rotate between the two frame arms and later secure it. On its upper side, it has two M8 holes, in the region of the aforementioned gap, where two M8 butterfly wing screws will be inserted to lock the piece in position. In its front part, the piece presents one cylindrical protrusion to position the mouthpiece and an M8 threaded hole where a butterfly wing screw will lock it in place. The cylindrical protrusion serves mostly

as orientation, to maintain the vertical direction of the mouthpiece's movements. The piece and its details can be observed in Figure 3.12.

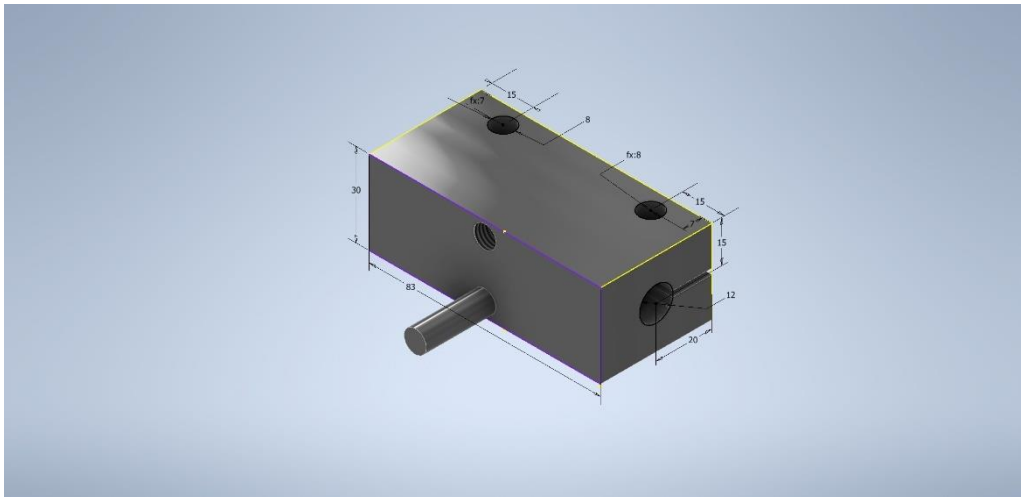

Figure 3.12 - The model of the frame's central piece .

3.4.4 The Lower Screw

The lower screws are the connections between the frame's arms and its bases. On the frame base's side, it is simply a solid 20 mm diameter cylinder, 40 mm long, to allow the rotation of the arm in relation to the fixed frame base. On the frame arm's side, it consists of a threaded M10 cylinder, 20 mm long, that represents the screw of the connection. The lower screw can be observed in Figure 3.13.

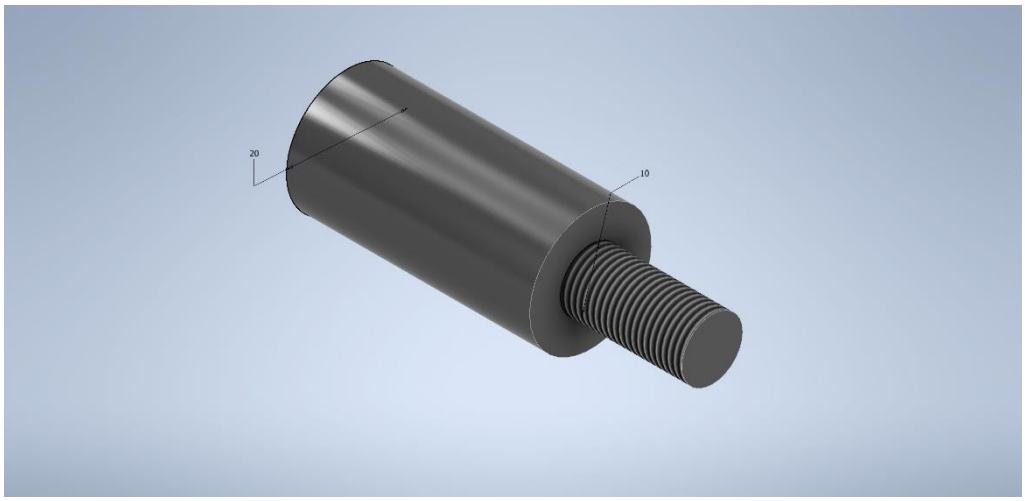

Figure 3.13 - The model of the lower frame screw.

3.4.5 The Upper Screw

The upper screws represent the connections between the frame's arms and its central piece. On the frame's arms side, this connection consists of an M8 threaded screw, 17 mm long, together with a cone-shaped part, rising another 3 mm, to assure greater stability in this connection. This first part of

the screw ends in a 1 mm tall step, that rises beyond the frame's arm insertion, to provide less friction with the central piece and thus, better movement in the connection. The second part of the upper screw, which is inserted into the central piece of the frame, consists of a simple 40 mm long cylinder with a 12 mm diameter. This whole piece and its details can be observed in Figure 3.14.

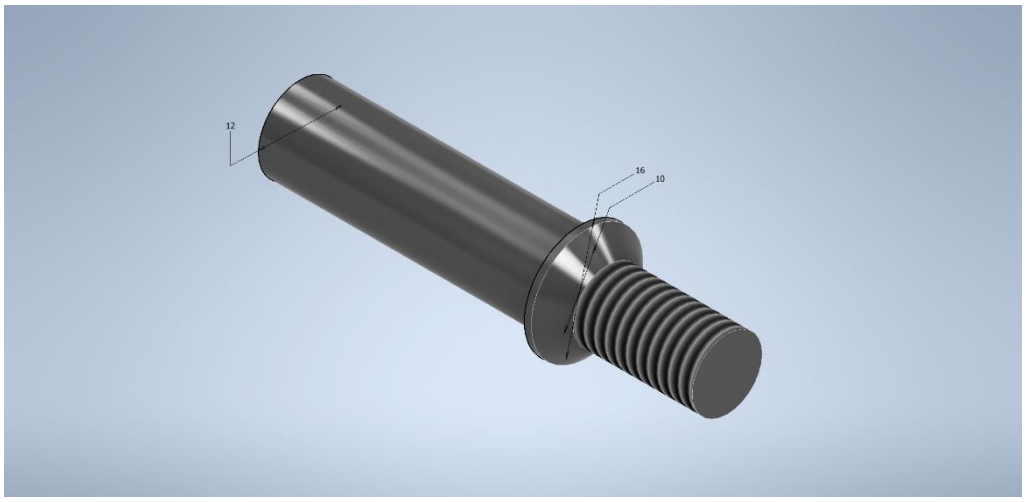

Figure 3.14 - The model of the upper frame screw.

3.5 Mouthpiece

The design of the mouthpiece applied in this system was based on the common dental impression trays used in dental clinics. The main modification added to that design was the junction of the upper and the lower tray in just one piece. To the main body of this mouthpiece was applied a parallelepiped 140 mm long with a 115 mm long cut down the middle, to provide a degree of movement liberty in the vertical axis of the system when positioning the mouthpiece. The upper part of the mouthpiece, which is in contact with the patient's maxilla, presents a protrusion, 11 mm tall 30 mm

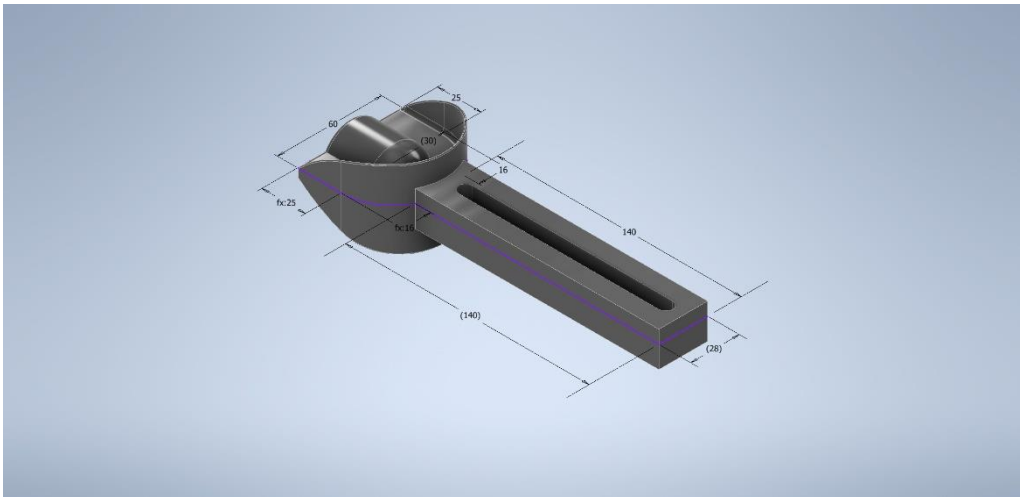

Figure 3.15 - The model of the mouthpiece designed, seen from the front .

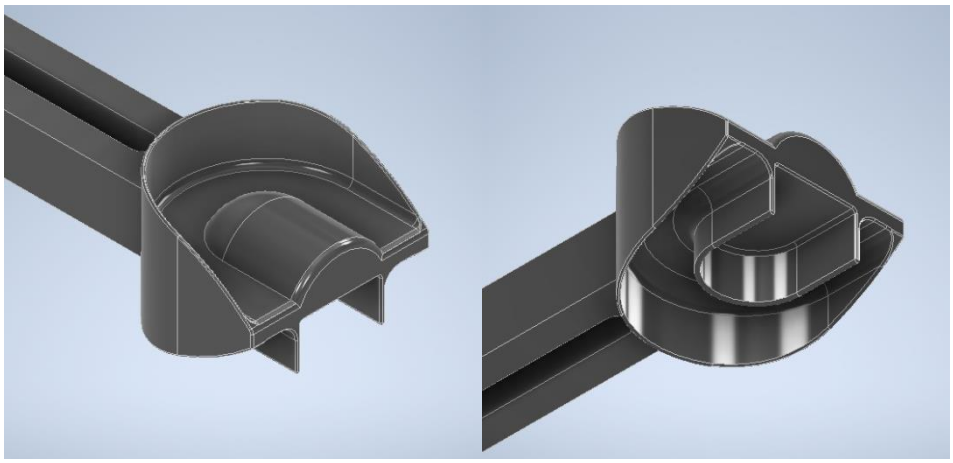

Figure 3.16 - The model of the mouthpiece designed seen from the back. The top part on the left and the lower part on the right.

wide and 41 mm long, rounded and positioned to accommodate the palate, with an 18,5 mm tall and 1 mm thick wall to stay between the teeth and gum and the lips of the patient. The lower part of the mouthpiece is similar to the first one. On the lower side, the protrusion was added to accommodate the patient's lower teeth, it is not as rounded as the upper one, due to the fact of not existing the palate, and has an extrusion within it to fit the patient's tongue. The outer walls of the mouthpiece were cut according to the usual design of the dental impression trays and all its corners were rounded to minimize discomfort. The mouthpiece designed and its details here mentioned are shown in Figure 3.15 and Figure 3.16.

3.6 Production of the Immobilization Device

After the completion of the design, the immobilization device's multiple parts were produced and assembled at *Laboratório de Instrumentação e Física Experimental de Partículas* (LIP) Coimbra. One of the main features we were looking for in the material to use in this device's production was MRI compatibility. The most MRI compatible material is usually PMMA (Poly(methyl methacrylate)) but, since we had the need of autoclaving the mouthpiece between usages, making it reusable for multiple patients, POM-C (Polyoxymethylene Copolymer) was the next best choice (Waplera, et al., 2014). Besides this, POM-C is ideal when it comes to replacing metal in mechanical machining. Some of its key features include strength and stiffness, toughness, creep resistance, fatigue resistance and dimensional stability, among others, all desirable characteristics for an immobilization project. Taking all these points into account, this was the material used to produce the immobilization device's.

Some minor alterations were made to the initial design during production in order to facilitate it. They include the separation of the body and the extension of the mouthpiece, being then connected through two locking pins, like we see in Figure 3.17, and the inclusion of another locking pin in the lower screw's connection to the frame's arms for improved stability, like we see in Figure 3.18. The

screws used to lock and unlock the device's position are made from Nylon Polyamide and are composed of a butterfly wing M8 nut and a threaded M8 rod, locked in place by a pin as we see in Figure 3.19.

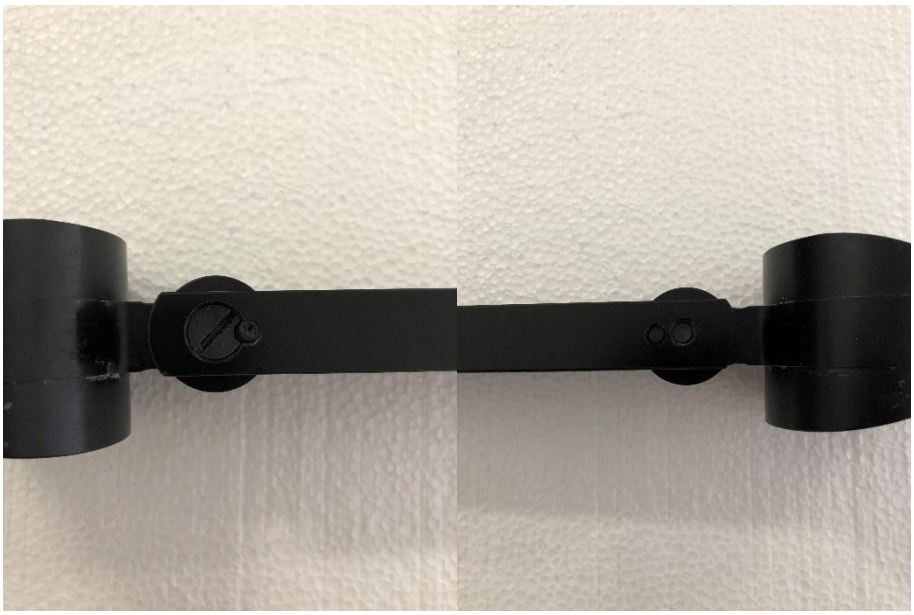

Figure 3.17 - Modifications introduced in the mouthpiece in the production process. Separation of the body and the extension of the mouthpiece, and connection through two locking pins

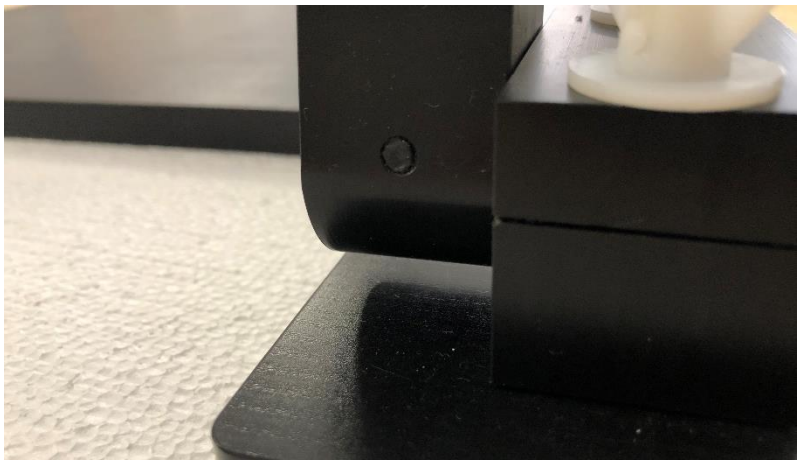

Figure 3.18 - Modifications introduced in the lower frame's arm connection in the production process. Inclusion of a locking pin in the lower screw's connection to the frame's arms to improve stability.

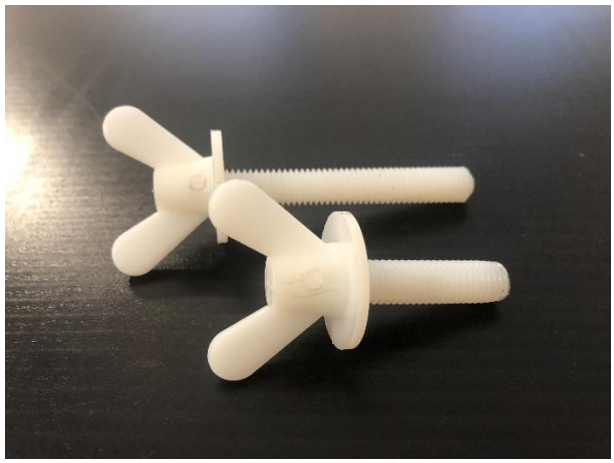

Figure 3.19 - Butterfly wing screws 3 cm long (in the front) and 5cm long (in the back).

In Figure 3.20 and Figure 3.21 we can observe a view of the final assembly with the Head and Neck coil and the produced immobilization device (since digital equipment is not allowed in the MR room, we could not obtain images of the actual process).

Figure 3.20 - Axial view for the preview of the final assembly between the immobilization device and the Head and Neck 64 coil.

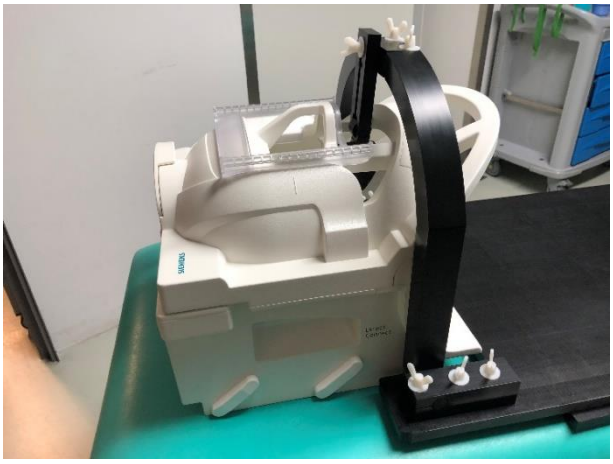

Figure 3.21 - Final assembly of the immobilization device and the Head and Neck 64 coil.

3.7 Patient Positioning

The positioning of the patient in the immobilization device was thought out to be simple, efficient, and of fast execution and it is exemplified in Figure 3.22. To begin, the immobilization device is placed on the MR table and the head and neck coil's top part is removed. The external frame is then rotated backwards until it reaches the top of the opened coil (A). The patient then lies on the system's base plate and rests his head inside the head and neck coil comfortably. Once this procedure is finished, the external frame is rotated forward until it reaches the patient's chest (B). At this point, the top part of the head and neck coil is set into place, closing the patient's head inside. After this, the external frame is rotated backwards until it reaches the mouth opening of the coil (C), and the mouthpiece is placed in the patient's mouth (D). The immobilization system is then adjusted to the patient's specific position and is locked together with the mouthpiece. Once the system is locked, the examination can begin.

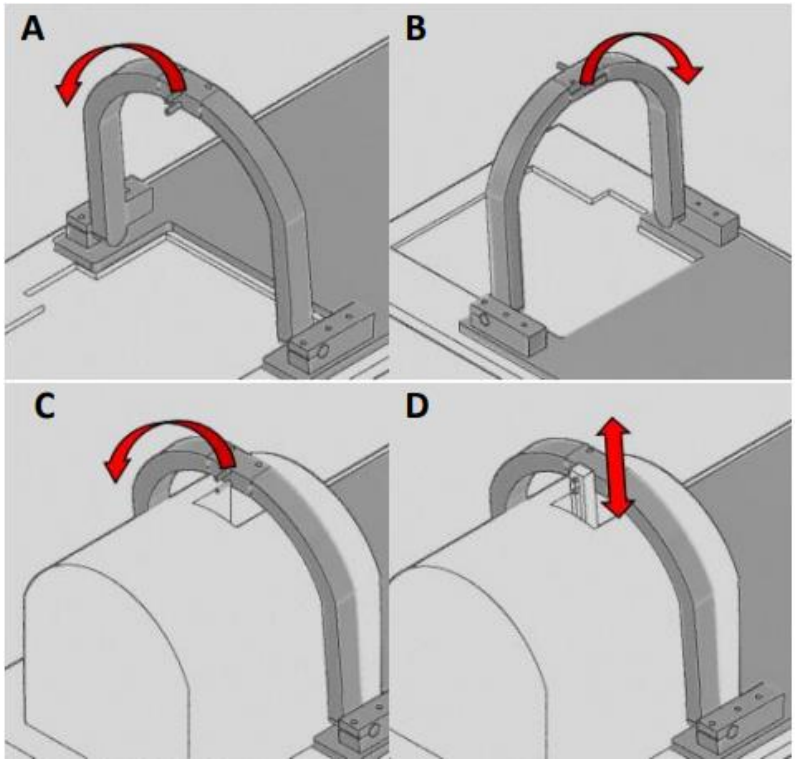

Figure 3.22 - Schematics for the patient's positioning process. (A) The external frame is rotated backwards until it reaches the top of the opened coil. (B) The external frame is rotated forward until it reaches the patient's chest. (C) The external frame is rotated backwards until it reaches the mouth opening of the coil. (D) The mouthpiece is placed in the patient's mouth and secured

Chapter 4
Software Developed to Analyse MRI Images

To access the existence of movement inside the MRI system, while the tests were undergoing, software had to be developed to analyse the images from the MRI scans and accurately calculate the translations and rotations present in those images. For that purpose, after some preliminary tests, two different algorithms were implemented, developed, and compared between them: the Block Matching algorithm and the SURF (Speeded Up Robust Features) algorithm. In this Chapter we'll present the different algorithms utilized and a general view of the final software. All of the algorithms were implemented in the software MATLAB®, developed by MathWorks Inc (MATLAB, 2021), as it has great computational power and is one of the best software to use in image processing.

4.1 Block Matching

Motion, in physics, is the change with time of the position or orientation of a body in reference to some position that is pre-recognized and assumed invariant and it is usually specified in a motion vector (MV) that comprises magnitude and direction. Motion estimation (ME) can be described as the process that obtains the motion vectors and specifies the transformation of a two-dimensional (2D) image frame into the next frame in an image sequence or video. There are different approaches to Motion Estimation that can be classified into direct, or pixel-based, and indirect, or feature-based, methods. While in direct methods MVs are calculated through the analysis and comparison of pixels between subsequent frames, in indirect methods these MVs are calculated with the analysis and comparison of features like colour tone and edge information, using statistical functions locally or globally. The Block Matching (BM) algorithm we present here is a pixel-based method. (Philip, et al., 2014)

The BM algorithm is the most popular method for motion estimation of local motion in an image sequence and its underlying supposition is that the patterns corresponding to objects and background in a frame move within the frame to form the corresponding objects on the next frame of the sequence. The basic idea behind the BM algorithm, represented in Figure 4.1, is the division of the image into a matrix of blocks that are compared to the corresponding block and its adjacent neighbours in the subsequent frame until a best-matching block is found. This search for the best-matching block in the subsequent frame is constrained to a specified $M \times M$ pixel search window, that can vary in size. The larger the motion, the larger should be the search window, but the larger the search window is, more computationally expensive the process becomes. Different BM algorithms differ in matching criteria, search method and the determination of the block size. (Tekalp, 1995) (Barjatya, 2004)

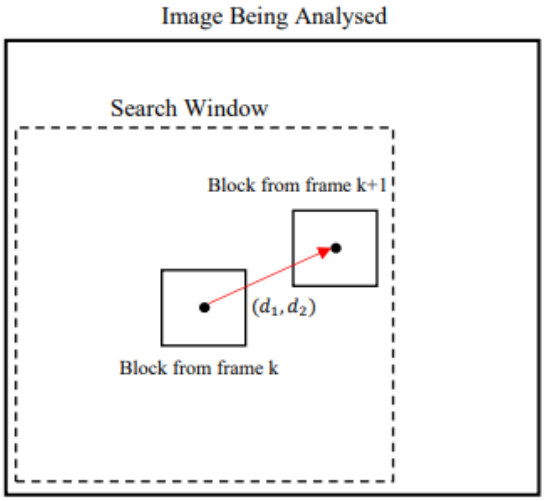

Figure 4.1 - Block Matching Algorithm Principles

The matching between corresponding blocks in different frames is based on the output of cost functions. The output of these cost functions is a numerical indication of the amount of mismatch between the blocks that are compared. The found block with the minimum dissimilarity, which minimizes the output of the cost function, is the one that matches the closest to the current block. There are various types of cost functions, but the most popular ones are the Mean Square Error (MSE) and the Mean Absolute Difference (MAD), that are given by equations (4. 1) and (4. 2).

$$MSE\ (d_1, d_2) = \frac{1}{N} \sum_{(n_1,n_2)\in B} [s(n_1, n_2, k) - s(n_1 + d_1, n_2 + d_2, k + 1]^2 \qquad (4.1)$$

$$MAD\ (d_1, d_2) = \frac{1}{N} \sum_{(n_1,n_2)\in B} |s(n_1, n_2, k) - s(n_1 + d_1, n_2 + d_2, k + 1)| \qquad (4.2)$$

where B represents the N-by-N block, for a set of candidate motion vectors (d_1, d_2) , (n_1, n_2) represent the pixels of the reference block (in frame k) and $(n_1 + d_1, n_2 + d_2)$ the pixels of the observed block (in frame k+1), with $s(n, n, k)$ being the intensity of the image in a specific pixel. The estimate of the motion vector is taken to be the value of (d_1, d_2) that minimizes the equation. (Tekalp, 1995) (Barjatya, 2004)

The search for the optimal correspondence block requires optimizing the matching criterions seen above to all possible MV candidates at each pixel (n_1, n_2) and this is accomplished by the usage of a method called Full-Search. This method applies the matching criterion to all possible values of (d_1, d_2), for each pixel in the block, which is computationally very expensive but normally reduced with the introduction of the search window, being the most effective type of search method. However, in most of the cases suboptimal solutions are used to reduce the computational expenses. Some of those methods include the Three-Step search or the Cross-Section search, amongst many others. (Tekalp, 1995) (Barjatya, 2004)

When it comes to the determination of the block and search window size, it is usually empirically adjusted, visually or quantitatively, until satisfying results are obtained, except in occasions where the object subject to the motion estimation is known. In this last case, the search window can be estimated using the expected object velocity. (Massanes, et al., 2011)

In the software developed for this project, the BM algorithm was utilized with the MSE matching criteria and the Full-Search method, with varying block and window sizes, adjusted to the overall image size that varied between different MRI acquisitions.

4.1.1 Block Matching's Data Processing

Unfortunately, due to unavoidable noisy grey tone fluctuations, the BM algorithm generates many wrong vectors over static blocks located on the background of the image (Stefano, et al., 1999). This problem was solved by applying a mask to the MV matrix generated by the BM algorithm, zeroing any background noise vectors. The mask applied was the binarization of the image being analysed, with the threshold value being adjusted manually. Another problem brought by these grey tone fluctuations is that they happen inside the image itself, and to solve this is a bit more complicated.

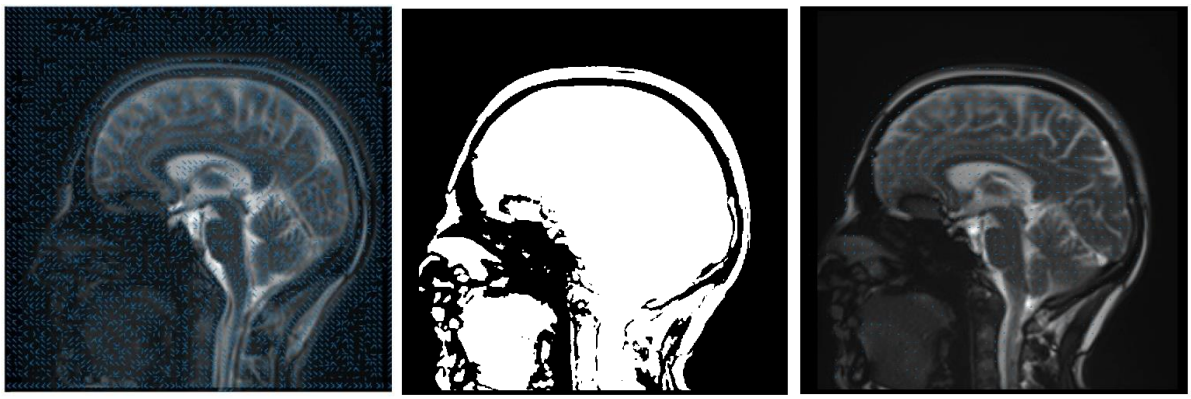

Figure 4.2 – Process of background noise removal, with the multiplication of the motion vector matrix (first image from the left) by a binary mask (middle image). Image at the right represents the result of the process.

To summarize, until now we have applied the BM algorithm to our sequence of images scanned by MRI and realized that in the MV matrix obtained, there are lots of noisy background MVs. We apply a mask to be able to neutralize this background noise, but we still have noisy MVs inside the image we're analysing that we must try to remove. Since this is a complex process, three processing algorithms were used and compared between them to solve this problem.

The first one, called Direction Processing (DP), relied on the principle that uniform movement in the image results in a MV matrix where the vectors should all be pointing in the same direction. The idea proposed was to analyse the vertical and horizontal components of every vector in the matrix and count the number of left (plus sign) and right (minus sign), and up (plus sign) and down (minus sign) vectors. Comparing the number of vectors in each state we can assume an overall direction of movement, eliminating every MV that doesn't follow it. The idea of this system is represented in Figure 4.3, where

the vectors surrounded by the circles would have been removed for not being in correspondence with the main movement direction. This processing algorithm was still considered in Chapter 5, but since, in reality, we would deal with head rotations, where the principle of a unanimous direction of movement is not applicable, it was only used when accessing movement in Regions Of Interest (ROIs).

The second processing algorithm developed was called Neighbourhood Processing (NP). This algorithm is based on the principle presented in the last paragraph but relied on the fact that closely related MVs should all have close intensities between them, and if a MV is not in correspondence with its neighbours, then it should be a noise vector, ultimately being removed. The MV matrix was then analysed vector by vector, eliminating the ones considered outside the tolerance interval of the neighbour's average intensity. Different neighbourhood sizes (3x3, 5x5, etc…) and tolerance values were applied to reach an optimal solution.

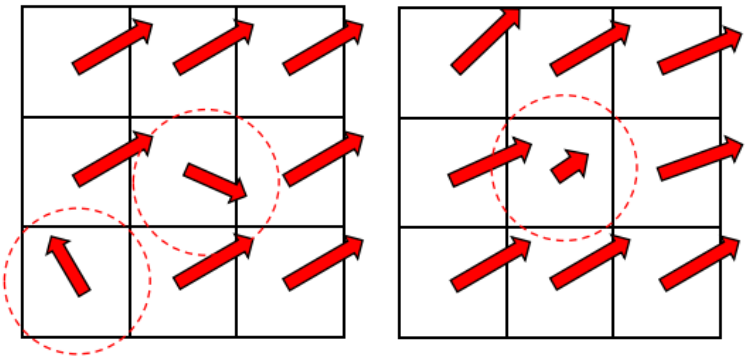

Figure 4.3 - Direction (on the left) and Neighbourhood (on the right) Processing principles visualized. The circled vectors would have been removed from the data.

Finally, a last processing algorithm, that we called Histogram Processing (HP), was applied to the BM algorithm's data. The principle here is basically the same as in NP, but it's applied in a larger scale. In this algorithm, all the MVs were collected from the matrix (in exception of the zeros of the background) and displayed by intensity in a histogram. This histogram was later analysed and fitted with a probability density function, from which we retrieved information about which vectors to accept and which to remove. There were various fits applied to the histogram (Kernel, Gamma, etc...) but the one found to be more reliable was the Gaussian fit, or Normal fit, from which we retrieved the coordinate's value for the maximum point and accepted the vectors which intensities lied within a tolerance interval of this value.

A comparison between these three processing algorithms, to assess the vantages and disadvantages of each one, was made during the testing of the BM algorithm and the processing algorithm chosen, who performed the best, was Histogram Processing.

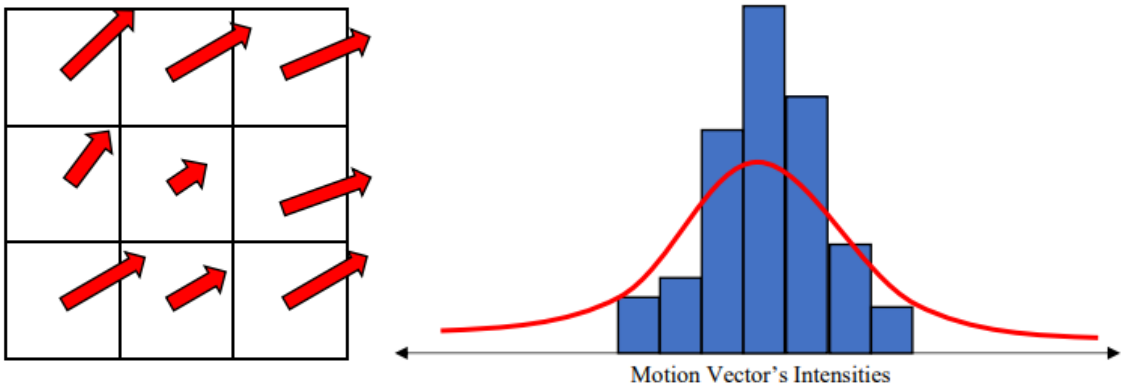

Figure 4.4 - Histogram Processing principle visualized. The vector intensities from the motion vector matrix on the left are represented in the histogram on the right, where a fitting curve was later applied.

4.2 SURF Algorithm

The SURF method (Speeded Up Robust Features) is an image comparison algorithm, based on scale and rotation-invariant interest point detection and description, which comprises three main steps. Firstly, is the detection of interest points of the image, defined as salient features from a scale-invariant representation, such as corners, blobs, or T-junctions, selected from different places in the image. The most valuable property of an interest point detector is its repeatability, the capability of finding the same interest point in different images, under different viewing conditions. The second step consists in building orientation invariant descriptors for every interest point, i.e. feature vectors that should represent the neighbourhood of each point detected. Desirable properties of a descriptor rely on its distinctiveness and robustness to noise, detection errors and geometric and photometric deformations. Lastly, the third step consists in matching the descriptor vectors between the images and is usually performed using distance or similarity metrics. (Bay, et al., 2008) (Oyallon, et al., 2015) (Parekh, 2021)

Several interest point detectors have been proposed. Between them, the Hessian-based methods (based on the Hessian matrix) seem to be the most stable and repeatable. The usage of the determinant of the Hessian matrix instead of its trace is also more advantageous, as it fires less on unwanted structures. The Hessian matrix is defined as $H(x,\sigma)$, like we see in equation (4. 3), where $x = (x, y)$ is a point in the image I, with σ being the scale, and where $L_{xx}(x,\sigma)$ is the convolution of the second derivative of the Gaussian function, $\frac{\partial^2}{\partial_x^2} g(\sigma)$, with the image I, at point **x** (the same goes for $L_{xy}(x,\sigma)$ and $L_{yy}(x,\sigma)$). (Bay, et al., 2008)

$$H(\boldsymbol{x},\boldsymbol{\sigma}) = \begin{bmatrix} L_{xx}(x,\sigma) & L_{xy}(x,\sigma) \\ L_{xy}(x,\sigma) & L_{yy}(x,\sigma) \end{bmatrix} \quad (4.3)$$

An approximation of Gaussian filters is used through box filters, like we see in Figure 4.5, that can be evaluated rapidly with the use of integral images. (Bay, et al., 2008)

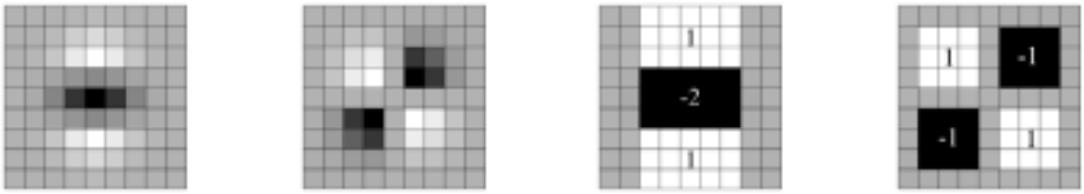

Figure 4.5 - The discretised and cropped Gaussian second-order partial derivatives in y-direction, on the first image from the left, and in the xy-direction, on the second image from the left. The approximation of the Gaussian second-order partial derivative in the y-direction using box filters, on the second image from the right, and in the xy-direction, on the first image from the right. The grey regions are equal to zero. From (Bay, et al., 2008).

The maximum of the determinant of the Hessian matrix operator is used to select interest point candidates throughout the image, that are then validated if the response is above a given threshold, with their location and scale being refined using quadratic fitting. Usually, a few hundred interest points are detected in a megapixel image. (Oyallon, et al., 2015)

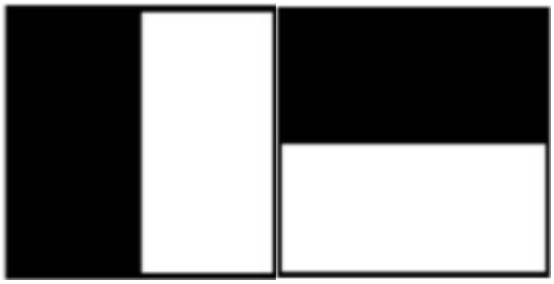

Figure 4.6 - Haar-wavelet horizontal (left) and vertical (right) types used for the SURF algorithm, From (Bay, et al., 2008).

To generate the descriptor, a square region is constructed and centred around the interest point, being oriented along the dominant orientation. It is calculated through the sum of Haar-wavelet responses in the horizontal and vertical directions, seen in Figure 4.6, in a circular neighbourhood of radius 6σ around the interest point, with σ being the scale of the interest point's detection. The direction which produces the largest sum of horizontal and vertical responses is the dominant one. The square region is then split into 4 x 4 square sub-regions, where the Haar-wavelet responses are again calculated, horizontal and vertically. If *dx* and *dy* denote the wavelet responses along the horizontal and vertical directions, then the descriptor vector for each sub-region is given by F, as in equation (4. 4). (Parekh, 2021) (Bay, et al., 2008)

$$F = \left(\sum dx, \sum dy, \sum |dx|, \sum |dy|\right) \tag{4.4}$$

Finally, to perform the image matching task, the local descriptors from several images are matched and compared. The matching is often based on a distance between the vectors, usually the Mahalanobis or Euclidean distance, and can be achieved using various techniques, like the similarity threshold or the nearest neighbour ratio. In the SURF algorithm the nearest neighbour technique is the

most used, where an interest point in the test image is compared to an interest point in the reference image by calculating the Euclidean distance between their descriptor vectors and a matching pair is detected, if its distance is closer than 0.7 times the distance of the second nearest neighbour. (Bay, et al., 2008)

4.2.1 Estimate Image Transformation

To assess the final transformation between consequent frames of an image sequence, i.e., to find the rotation and translation given from one frame to another, we need to extract the transformation matrix based on the matched interest points computed by the SURF algorithm. This is done using the statistically robust M-estimator SAmple Consensus (MSAC) algorithm, which is a variant of the RANdom SAmple Consensus (RANSAC) algorithm.

The RANSAC algorithm is used to estimate the robust transformation function between image frames. It is considered a robust estimation algorithm because its estimations are tolerant in the presence of outliers, points that fall outside the main group of data. Most algorithms using robust estimation gather the maximum amount of the data to provide an initial solution while RANSAC, on the other hand, uses the minimum data set possible and then tries to enlarge this data set with more consistent data. It consists mostly of two steps: Hypothesis and Test, and to be able to fit a transformation model to the data, the algorithm needs multiple sets of information beforehand: the minimum quantity of points required to match the model, also referred to as Minimum Sample Space (MSS), the minimum quantity of iterations to perform, the threshold value that differentiates the outliers from the inliers and the size of the data set, that corresponds to the end of the iterations. (Kulkarni, et al., 2017)

The Hypothesis step aims to fit a model to the data, able to minimize the impact of outliers. It begins with the selection of a MSS from the data set, and a transformation model is estimated based on these selected points (hypothesis). All of the remaining data set is tested to match this transformation model. If one or both selected points present themselves as outliers, then the model will not fit the rest of data, the algorithm will skip it and randomly pick another set of points to test another output model. We can visualize better this procedure by looking at Figure 4.7. Once the transformation model is defined the Testing step begins. Here, the algorithm iteratively checks which points in the entire data set are consistent with the transformation model, examining the distance between a specific point and the hypothetical transformation model and comparing it to the specified threshold value. If that specific point is within threshold, then it is classified as an inlier. The estimated model is correct when it reaches enough points classified as inliers. The most effective group of observations selected from the whole data set (the inliers) is known as the consensus set (CS). If the threshold value used in the RANSAC algorithm is considered very high, the robust estimation will be affected, as most of the points would be considered inliers. The MSAC algorithm was then introduced to solve this problem, evaluating the quality of the CS by calculating its likelihood. (Kulkarni, et al., 2017)

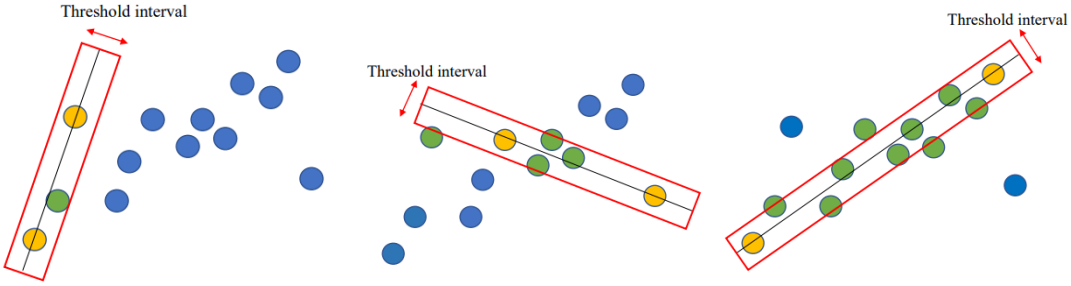

Figure 4.7 - Example of a RANSAC linear model acceptance. The MSS points are yellow and the accepted inliers, green. The accepted linear model would be the one on the right, where most inliers were accepted. The blue points represent the outliers.

The transformation matrix is then derived from the CS, mapping the inliers from one frame to another. In the case of images, affine transformation matrixes are the most usual approach in 2D spaces and consist of six parameters, four of the parameters form the linear part, while the remaining two specify the vertical and horizontal translation components. Being (x, y) the coordinates of a point in the reference image and (u, v) the ones in the current image, the affine transformation is represented by the matrix A, like we see in equation (4. 5), where a, b, d and e are the elements of the linear part, and c and f the elements of the translation vector. (Dung, et al., 2013)

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = A \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, where\ A = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \quad (4.5)$$

This matrix is usually re-fitted as an easy scale-rotation-translation transform (s-R-t Transform) and the final parameters consist of one scale factor (s), one angle (Θ) and two translation parameters (t_x and t_y) like we see below in equation (4. 6). (Kulkarni, et al., 2017)

$$A = \begin{bmatrix} s * cos(\theta) & s * -sin(\theta) & 0 \\ s * sin(\theta) & s * cos(\theta) & 0 \\ t_x & t_y & 1 \end{bmatrix} \quad (4.6)$$

However, since the image sequences we are analysing in this work are all collected from the same spatial plan, the scale parameters can be neglected and a rigid transformation matrix can be used, providing better results. The final rigid transformation matrix A can now be written like we see below in equation (4. 7). Finally, from this transformation matrix, it's easy to understand how the rotation and translation between frames can be retrieved.

$$A = \begin{bmatrix} cos(\theta) & -sin(\theta) & 0 \\ sin(\theta) & cos(\theta) & 0 \\ t_x & t_y & 1 \end{bmatrix} \quad (4.7)$$

4.3 General Software Overview

With the different algorithm approaches selected, the main program was developed in MATLAB, to process the images retrieved from the MRI scans, apply the algorithms, and analyse their results. The images from the MRI scans can be loaded into the software in two different ways, as each one presented some benefits: directly, through the *dicomread* function MATLAB provides, and indirectly, using images pre-processed by syngo®, a software developed by Siemens for clinical imaging.

When applying the motion estimation algorithms, two approaches were utilized: one where the image sequence frames are compared in sequential form (the first with the second, the second with the third, etc..), requiring a cumulative sum of the final data to obtain the movement profile, and another one where the image sequence frames are always compared to the same frame, chosen manually (usually a frame from the middle of the sequence is chosen for better stability). The motion estimation algorithms are then applied between the two selected frames, process that repeats itself for every image sequence frame until the end of the sequence is reached. Once reached the end of the image sequence, we are left with three variables with the size of the image sequence, representing the horizontal and vertical translations in pixels, and the computed rotations in degrees for every frame comparison. To convert the computed translations into millimetres, we must convert the data from pixel to the wanted calibration. While for the images loaded through syngo® the *pixel* to *mm* conversion is done manually, using the FOV and image size information, loading the images through *dicomread* allowed the use of the *dicominfo* function, that provides the spacing between pixels directly. With the final data converted, the movement profiles are shown in millimetre (or degree) over time or frame number and can be interpreted and analysed posteriorly. The entire process of the software can be visualized in Figure 4.8.

When comparing translational movement profiles from both algorithms it is expected to observe some discrepancies, caused by the intrinsic basic principles of the algorithms. Since SURF works by identifying interest points, the translations obtained will be referent to those matched interest points. As for Block Matching, the movement is obtained through the MV matrix, referent to the hole image, divided into the block sizes analysed.This fact most often implies that BM movement estimation presents itself as a more conservative aproach when compared to SURF estimation. In the next Chapter we will test and calibrate the algorithms and software, to acess, assure and compare the veracity of the movement profiles aquired.

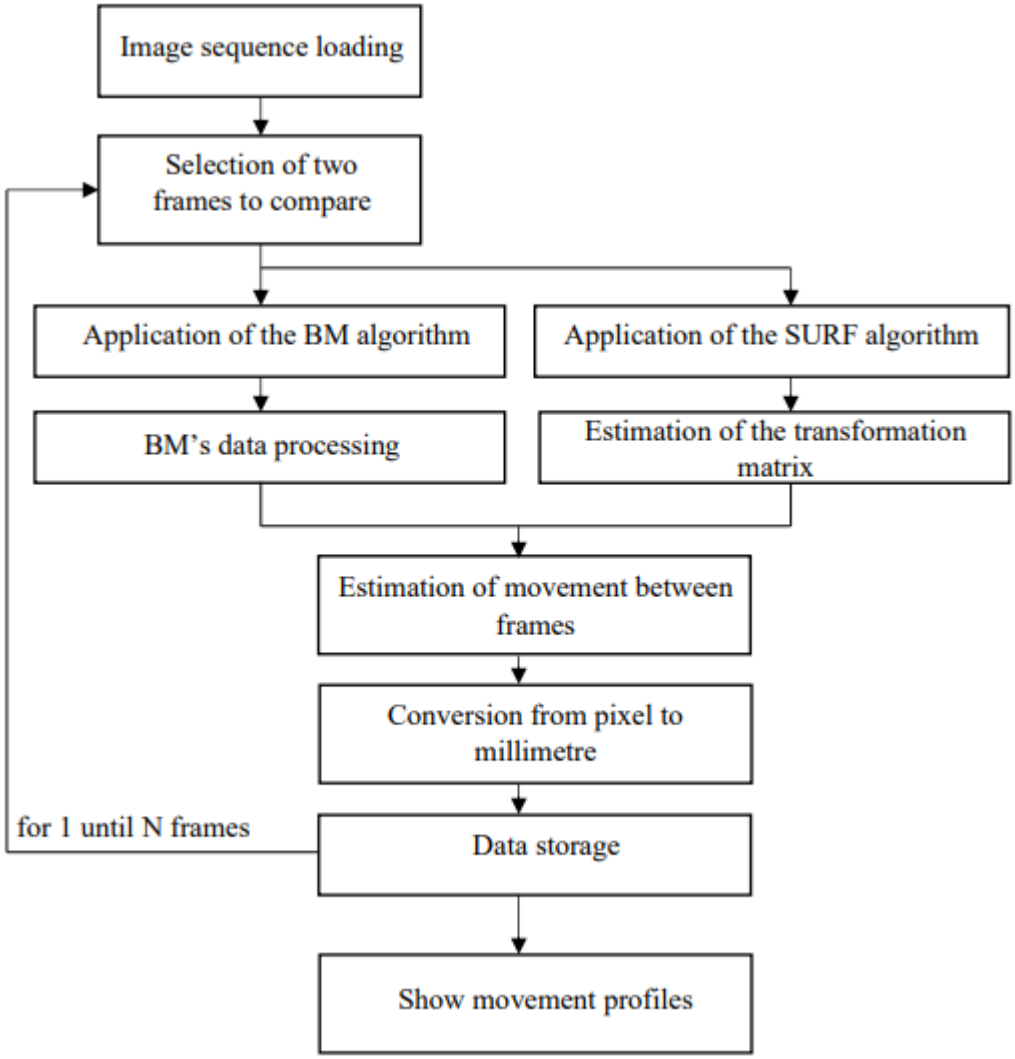

Figure 4.8 - Flow chart of the software developed to analyse MRI's images.

4.3.1 MATLAB functions and specifications

In this section we will describe the main MATLAB funtions used and its parameters, to better comprehend the software developed and its key points. The functions adressed will be referent to the BM and SURF algorithms, and are available when using the Computer Vision Toolbox provided by MATLAB.

The BM algorithm's utilization is extremely simplified when developing a MATLAB based software. It is based only on the *vision.BlockMatcher* class, that returns an object able to perform the motion estimation between two image frames by moving a block of pixels over a search region. To this function we must specify the wanted object's capabilities and its specifications. There is an extensive list of changeable properties when creating this block matching object, from which we emphasise the *ReferenceFrameSource*, the *SearchMethod*, the *BlockSize*, the *MatchCriteria* and the *OutputValue*. The *ReferenceFrameSource* must be set to *InputPort* to allow to provide a second image frame as a reference to the one being analysed and, like that, compare the motion between the two. As explained in 4.1, the

SearchMethod must be set to *Exhaustive*, and the *MatchCriteria* to *MES*, which are the default values. The *BlockSize*, as explained before, is set to a specific value depending on the characteristics of the images being analysied (mainly their size), with the default value being a 17x17 pixel region. Lastly, the *OutputValue* chosen was the *Horizontal and vertical components in complex form*, instead of *Magnitude-Squared*, to allow independent anylises of the horizontal and vertical components of the motion vectors.

The two image frames we're trying to compare are then served to the object returned by the *vision.BlockMatcher* function as inputs, allowing the return of the motion vector matrix that estimates the movement between the two.

MATLAB's SURF algorithm is equally of simple utilization, being based on four different functions: *detectSURFFeatures*, *extractFeatures*, *matchFeatures* and *estimateGeometricTransform2D*. The *detectSURFFeatures* function implements the SURF algorithm to find blob-like features and returns an object containing the point information about the features detected in the 2-D grayscale input image. The function has a variety of changeable properties from which the only one changed during our software development was the *MetricThreshold*, also known as the strongest feature threshold. This value had to be adjusted to the image characteristics, like the image size, since for example, to detect smaller features, we had to decrease it. This function is applied to both frames being compared, leaving us with two sets of detected feature points from the different frames. The next function applied is the *extractFeatures*. The *extractFeatures* function returns the vector descriptors and their corresponding location for each detected feature in each analysed frame. Its input arguments were the detected feature points and the corresponding frame, and its properties were set as default. This feature vectors from both frames are then passed on to the *matchFeatures* function to compare them and return the indices of the corresponding features between the two input sets, returned as a P-by-2 matrix of P number of indices. The matched points are then extracted using this matched features matrix and passed on to the *estimateGeometricTransform2D* function to estimate the final transformation between the two frames. This function receives the matched points of the two images and the output transformation type as arguments and returns the final geometric transformation matrix. In our software, the output transformation type was set to *rigid*, as explained in 4.2.1, to obtain the best estimation possible.

Chapter 5
Software Testing and Calibration

In this Chapter, the procedures implemented to assure the veracity of the displacement and rotation results, obtained from the motion estimation software, will be exposed and explained. To achieve this, multiple tests were conceived. To test the intrinsic limits of the algorithms we performed digital tests by manipulating the images to be processed beforehand, applying translations and rotations digitally, and then analyzing the results provided by the software. Later it was developed a movement phantom, capable of performing predictable displacements, to produce standard test images for the testing and calibration of the software. The phantom was developed to be MRI compatible and low-cost, with a high precision. The images obtained from the phantom tests were processed by the software, its data was analyzed, and the software's parameters were adjusted to ensure a good correspondence between real and computed displacements and rotations.

5.1 Digital Tests

The digital tests on the algorithm were divided into multiple steps. To perform the translational digital tests, we took a head and neck MR image and shifted it on the XX axis to simulate translations, firstly of one-by-one pixel, until twenty translations were reached, and then of subpixel lengths until image sequences of the same twenty translations were created. These subpixel translations were achieved through oversizing the original image, by increasing the number of pixels on each axis by a known factor (2, 3, 5, and 7), using image interpolation. A pixel-by-pixel shift was applied to these oversized images that were then undersized to their original size, creating the subpixel length (1/2, 1/3, 1/5, and 1/7 of a pixel) translation sequences. To perform the rotation digital tests a similar principle was put into use, we applied known rotations with different steps (2, 1, 0.5, 0.2, 0.1, and 0.05 degrees) to the original head and neck MR image, until we had nine frames for each image sequence, varying the maximum rotation angle reached for each one of the different sequences.

The limits of the algorithms were analyzed through the standard deviation of the residuals in the linear fit of the computed translations, in the XX axis, the standard deviation of the computed translations in the YY axis, and the standard deviation of the residuals in the linear fit of the computed rotations. To present the achieved translation results in millimeters we made the correspondence between the FOV of the MR images and their size in pixels, obtaining 0.329 mm for each pixel length in the image.

5.1.1 Block Matching Algorithm Results

As mentioned beforehand, the limits of the BM algorithm are here presented as the standard deviation of the residuals in the linear fit of the computed translations in the XX axis, which should have

a slope equal or near the value of the pixel, since all translations should be equal between them, and the standard deviation of the computed translations in the YY axis, that should be equal or near zero, since the induced movements are just in the XX axis. The limits of the BM algorithm are resumed in Table 5.1 for translational steps of 1, 1/2, 1/3, 1/5, and 1/7 of a pixel's length.

TABLE 5.1 - LIMITS OF THE ALGORITHM FOR THE TRANSLATIONS COMPUTED IN THE DIGITAL TESTS OF THE BM ALGORITHM, FOR THE XX AND YY AXIS FOR DIFFERENT TRANSLATIONAL STEPS.

	Translational Step (fraction of a pixel)				
	1	1/2	1/3	1/5	1/7
STD of XX Residuals	5.7 µm	58.54 µm	85.11 µm	85.19 µm	84.99 µm
STD YY	1.28 µm	18.55 µm	15.99 µm	19.14 µm	14.55 µm

As we can observe from Figure 5.1 to Figure 5.2, the BM algorithm detects almost perfectly the induced digital translations of one-pixel steps. From then on it is clear that this algorithm is not able to detect translations lower than one pixel length. Although we can see that the algorithm detects motion before the translation has reached the first pixel, and from then on only the translations of one pixel are detected.

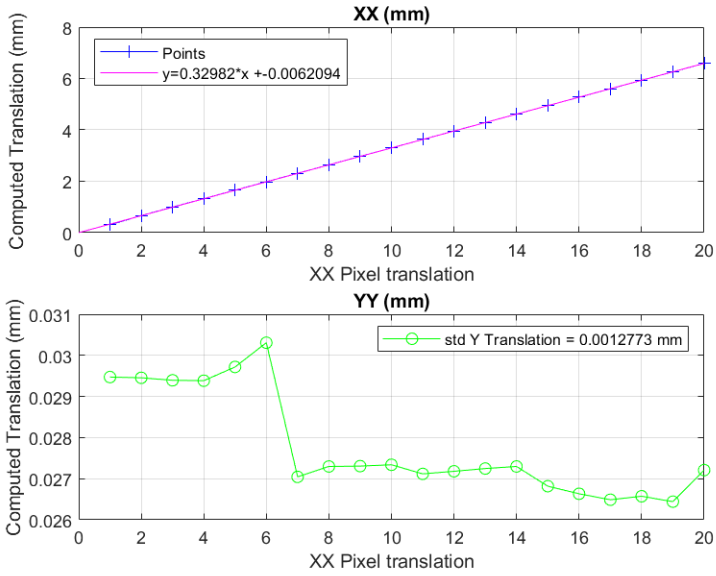

Figure 5.1 - Computed translations for the XX and YY axis, for each image frame, with 1 pixel horizontal translational steps, using the BM algorithm..

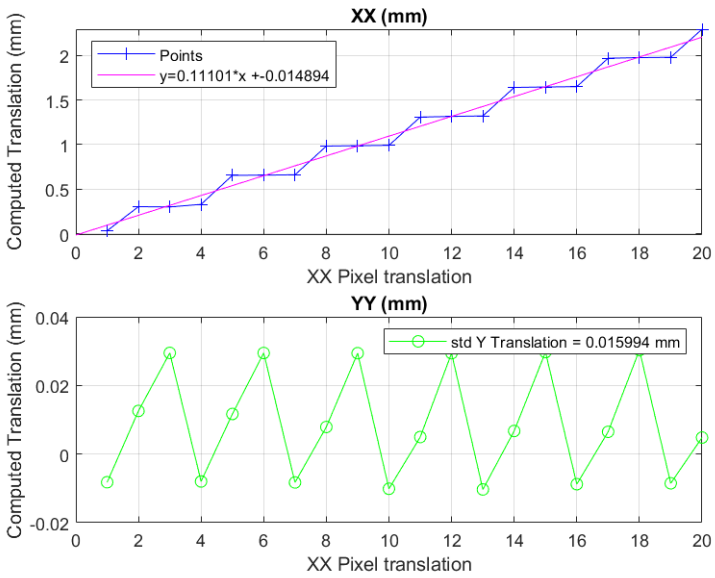

Figure 5.3 - Computed translations for the XX and YY axis, for each image frame, with 1/3 pixel horizontal translational steps, using the BM algorithm.

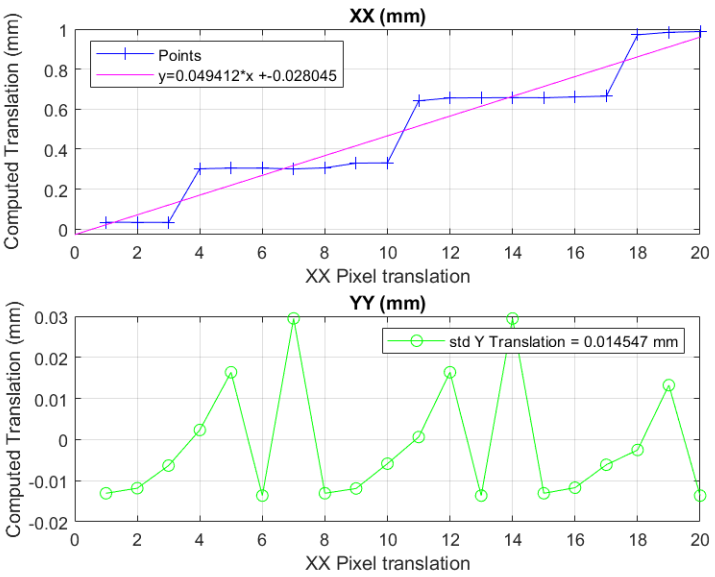

Figure 5.2 - Computed translations for the XX and YY axis, for each image frame, with 1/7 pixel horizontal translational steps, using the BM algorithm.

5.1.2 SURF Algorithm Results

The limits of the SURF algorithm are here presented in the same way as above, with the addition of the results of the rotational digital tests, presented as the standard deviation of the residuals in the linear fit of the computed rotations. The limits of the SURF algorithm are resumed in Table 5.2 for translational steps of 1, 1/2, 1/3, 1/5, and 1/7 of a pixel's length, and in Table 5.3 for rotational steps of 2, 1, 0.5, 0.2, 0.1, and 0.05 degrees.

TABLE 5.2 - LIMITS OF THE ALGORITHM FOR THE TRANSLATIONS COMPUTED IN THE DIGITAL TESTS OF THE SURF ALGORITHM, FOR THE XX AND YY AXIS FOR DIFFERENT TRANSLATIONAL STEPS.

	Translational Step (fraction of a pixel)				
	1	1/2	1/3	1/5	1/7
STD of XX Residuals	8.33 µm	7.67 µm	11.29 µm	9.93 µm	11.9 µm
STD YY	9.8 µm	7.73 µm	12.19 µm	11.51 µm	12.72 µm

TABLE 5.3 - LIMITS OF THE ALGORITHM FOR THE ROTATIONS COMPUTED IN THE DIGITAL TESTS OF THE SURF ALGORITHM, FOR DIFFERENT ROTATIONAL STEPS.

Rotational Steps	Maximum Angle	Residuals of STD
2º	18º	0.025º
1º	9º	0.011º
0.5º	4.5º	0.005º
0.2º	1.8º	0.006º
0.1º	0.9º	0.006º
0.05º	0.45º	0.033º

As we can observe from Figure 5.4 to Figure 5.6, the SURF algorithm also detects almost perfectly the induced digital translations of one-pixel steps. From then on, unlike the BM algorithm, this algorithm shows a greater capability of detecting sub-pixel translations, as it shows better linearity and limit values.

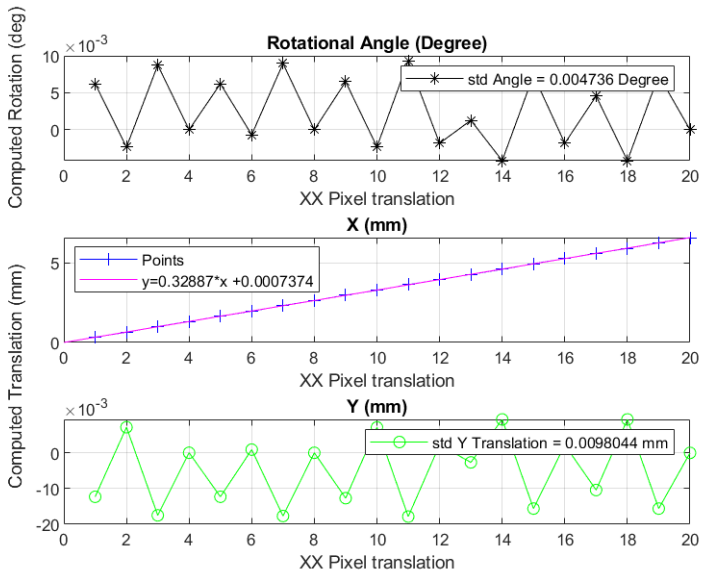

Figure 5.4 - Computed translations for the XX and YY axis and rotations, for each image frame, with 1 pixel horizontal translational steps, using the SURF algorithm.

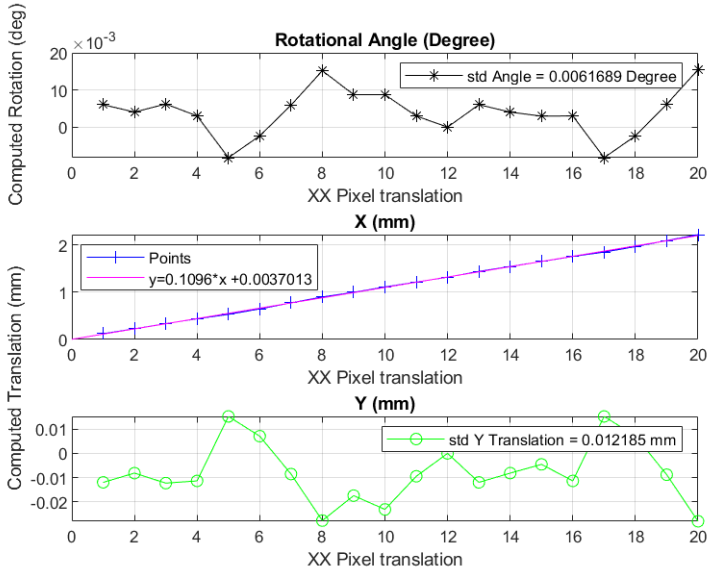

Figure 5.5 - Computed translations for the XX and YY axis and rotations, for each image frame, with 1/3 pixel horizontal translational steps, using the SURF algorithm.

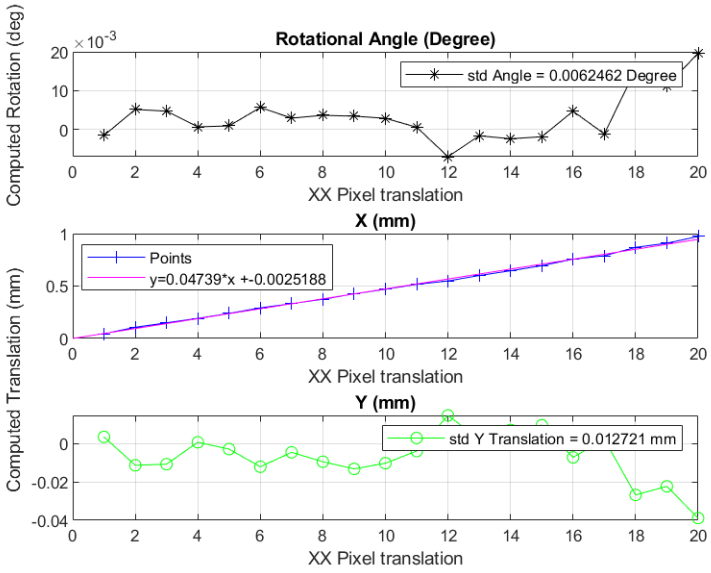

Figure 5.6 - Computed translations for the XX and YY axis and rotations, for each image frame, with 1/7 pixel horizontal translational steps, using the SURF algorithm.

As we can observe from Figure 5.7 to Figure 5.9, the SURF algorithm also detects almost perfectly the induced digital rotations until a certain scale factor. The precision of the algorithm is remarkable until the rotational steps around 0,1º, which is perfect for the type of rotations we are looking for in real life tests. From then on, we see a decline in the algorithm's precision, although the values of the computed rotations are still quite impressive for the size of the scale of the rotational steps.

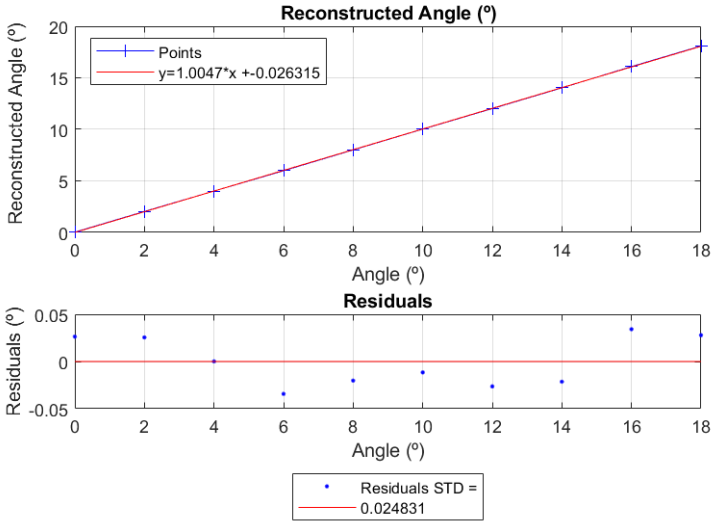

Figure 5.7 - Computed rotations and respective residuals with 2 degrees rotational steps, using the SURF algorithm.

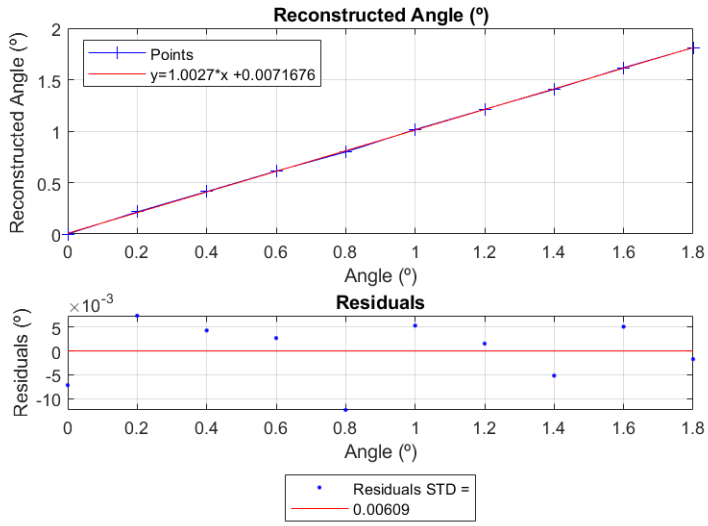

Figure 5.8 - Computed rotations and respective residuals with 0.2 degrees rotational steps, using the SURF algorithm.

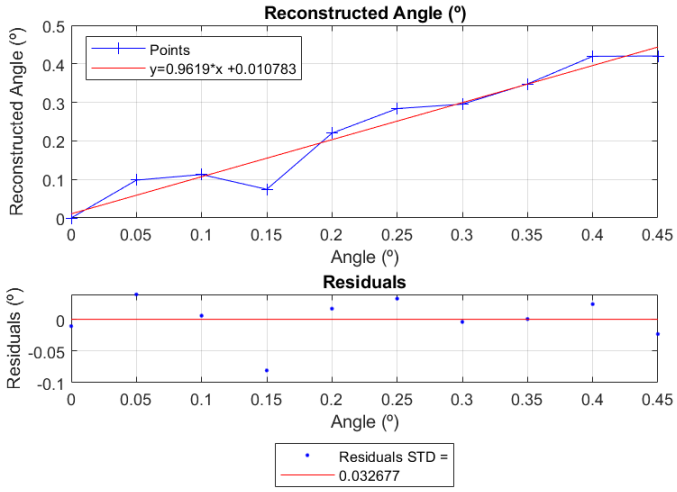

Figure 5.9 - Computed rotations and respective residuals with 0.05 degrees rotational steps, using the SURF algorithm.

5.2 The Movement Phantom

The concept of the movement phantom developed can be divided into three main constituents: the phantom itself, which must be visible in MRI, the movement-inducing mechanism and its parts, and the support for these components inside the imaging system.

The development of the phantom considered the need for MRI compatibility, as well as the final MR image characteristics that would be produced when imaged. The components of the phantom visible in MRI were thought out to simulate as possible (within our low-cost possibilities) the real images the software would have to process in a human head and neck imaging test. The easiest and cheapest way encountered was by filling water balloons with MRI contrasting liquids and placing them alternately inside a plastic container. For this experiment, we utilized four balloons inside a 60 by 60 mm plastic recipient, containing, alternately, mineral oil (liquid Vaseline) and water (two balloons with each). Mineral oil is a convenient liquid material for phantoms due to its relaxation time, low dielectric constant, high proton density, and low cost, giving a bright signal in T1-weighted MRI images (Gach, 2019). To contrast with the mineral oil's bright signal, regular water was chosen as the second material.

To create induced motion in the standard image sequence to be processed by the software, we needed to move the phantom inside the MR system. This is quite a difficult task to achieve while imaging, considering the need for MRI compatibility of the system. The simplest and easiest low-cost solution found was through inducing motion in the phantom between image scans, and not while imaging. It was designed a square base for the phantom to sit on, with a central plate inside, 2 mm shorter on each side of the base plate, as we see in Figure 5.10. These side gaps between the central plate and the bottom plate allow the insertion of plastic spacers with a known thickness (1 mm, for example), added between each image scan, to create motion in the final image sequence. To secure the phantom to the middle piece, velcro strips were placed in four different points of the central plate.

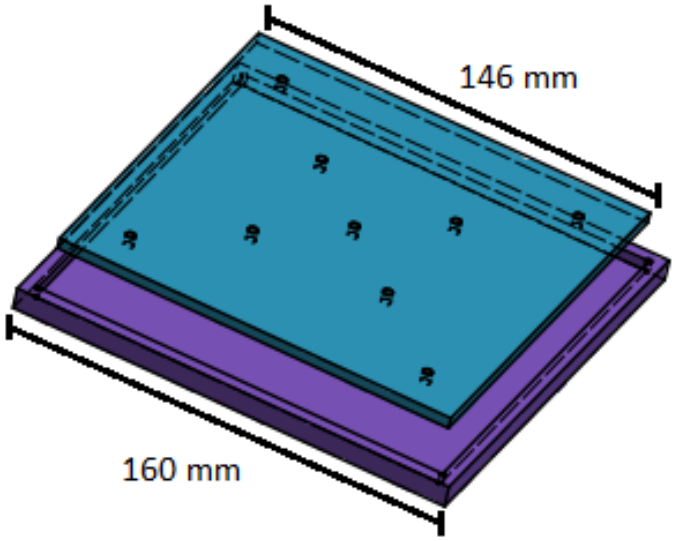

Figure 5.10 - Schematics of the motion inducing device developed. The square base of the device on purple and its movable central piece in blue.

To support this system inside the head and neck coil a mechanism was designed to compensate for the interior curvature present, usually used to support the patient's head and neck, and secure the phantom always in the same plane. This support mechanism also needed to be developed with MRI-

compatible materials. The low-cost solution found for this purpose was through the construction of a mold of the interior curvature of the coil, that would be flat on the top to secure the functional set-up of the phantom. This mold was constructed with a mixture of styrofoam and styrofoam glue placed inside a plastic bag to keep the coil clean. The mixture was made beforehand and, while it did not harden, was placed inside the RF coil, being pressed with a flat surface from above for a few minutes. With the mixture hardened, a flat styrofoam top was placed and secured to the mold, with velcro strips and wood pleats, to tightly secure the functional set-up of the phantom to the support. The whole system, positioned inside the open head and neck RF coil, can be seen in Figure 5.11. In Figure 5.13 we can see an imaged slice of the transversal plane of the phantom.

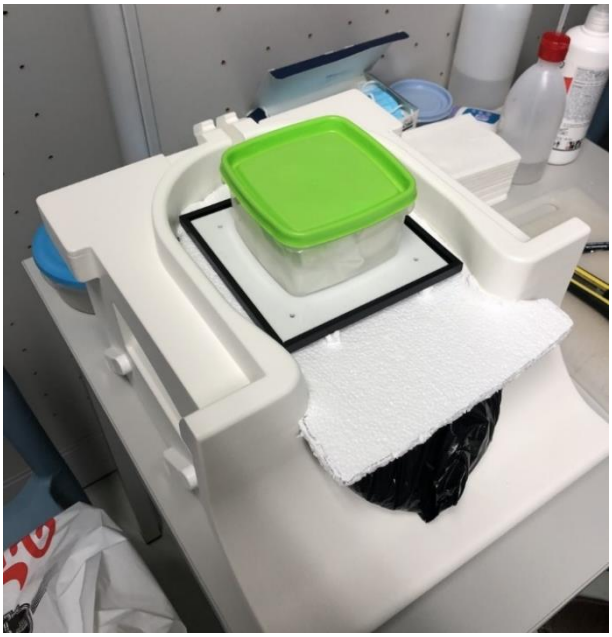

Figure 5.11 - Movement phantom and its set-up, placed inside the open head and neck RF coil, before being imaged.

To generate the image sequence with the simulated translational motion, the coil was placed inside the MRI system with the phantom system inside and the entirety of the spacers inserted into one side of the base plate, keeping it in place and tight, as seen in Figure 5.12. After the first image of the sequence was generated, a spacer would be removed from its original side and inserted on the opposite one, moving the phantom one spacer's thickness length to the side and maintaining the tightness of the central plate, avoiding unwanted movements. This procedure was repeated until all the spacers were positioned on the opposite side from where they were originally placed. To simulate rotational motion, the procedure was similar, differing on the position and length of the spacers. Here, smaller square "point-like" spacers, with the wanted thickness, were used. By placing them in the corners of the base plate we can create a degree of rotation like we see in Figure 5.12, adding rotational steps with the more spacers inserted.

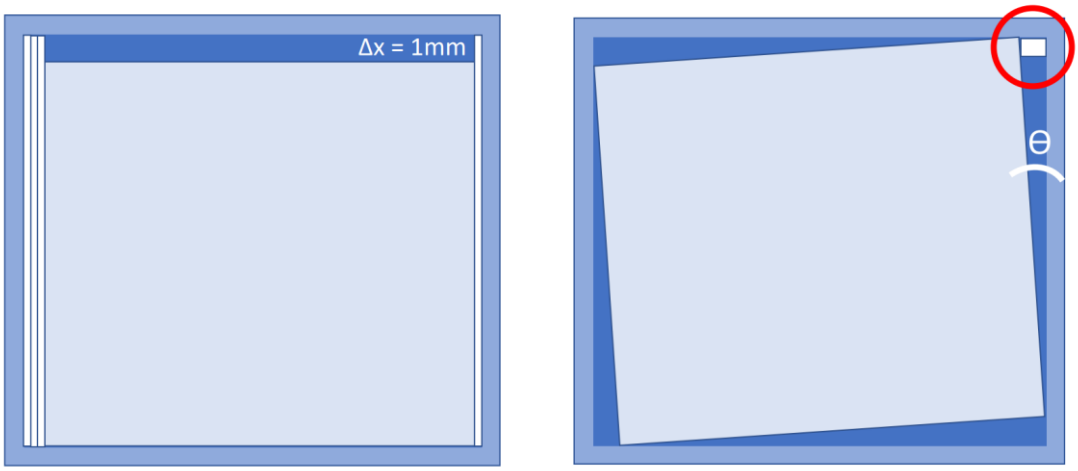

Figure 5.12 - Translation inducing procedure scheme, on the left, with the introduced spacers to create a 1mm translation. Rotation inducing procedure scheme, on the right, with the introduced spacer, circled in red, and the rotation angle Θ created.

Five image sequences were obtained to test the algorithms with the first sequence, composed of ten images, being generated with the phantom in the same position, to assess the algorithm's intrinsic noise and resolution. Two other sequences, named "BIG" translations and rotations, were obtained utilizing spacers of 1 mm, generating five images each with 0, 1, 2, 3, and 4 mm translations of the phantom, or the equivalent rotations (that only go until 3 mm). The last two sequences, named "SMALL" translations and rotations, utilized spacers with 0,1 mm thickness, also generating five images each. The expected rotations were calculated through the arcsine of the thickness of the spacer inserted, over the side length of the central plate of the motion inducing mechanism. The expected

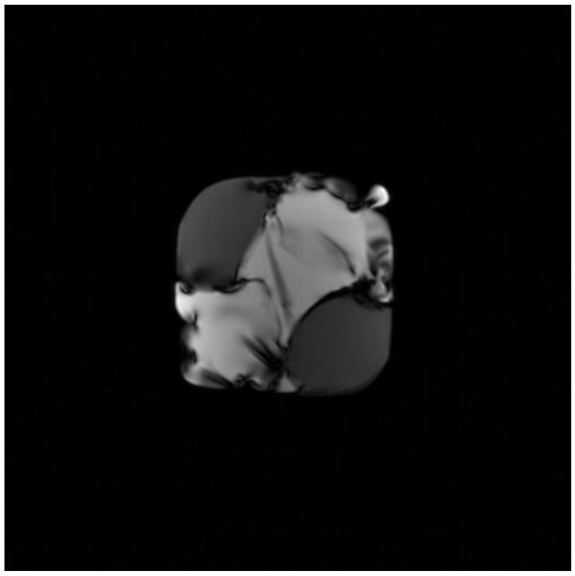

Figure 5.13 - Slice of the transversal plane of the movement phantom, imaged through Magnetic Resonance.

rotations calculated are 0º, 0.39515º, 0.79031º and 1.1855º for the BIG rotations sequence, and 0º, 0.0395º, 0.0790º, 0.1185º and 0.158º for the SMALL rotations sequence. All analyses of the image sequences involving translations or rotations were performed with the original sized images and with the original images oversized by a factor of 3, to eventually help the algorithms search for smaller scale movements.

5.2.1 Block Matching Algorithm Results

In Table 5.4 and Figure 5.14 we can analyze the mean position and the standard deviations obtained for the computed translations returned by the algorithm, for the ten frames of the phantom exactly in the same position and for the XX and the YY axis.

TABLE 5.4- COMPUTED STANDARD DEVIATION FOR THE IMAGE SEQUENCE OF THE PHANTOM IN THE SAME POSITION, FOR THE XX AND YY AXIS, USING THE BM ALGORITHM.

	Translational STD
XX Axis	47.3 µm
YY Axis	67.9 µm

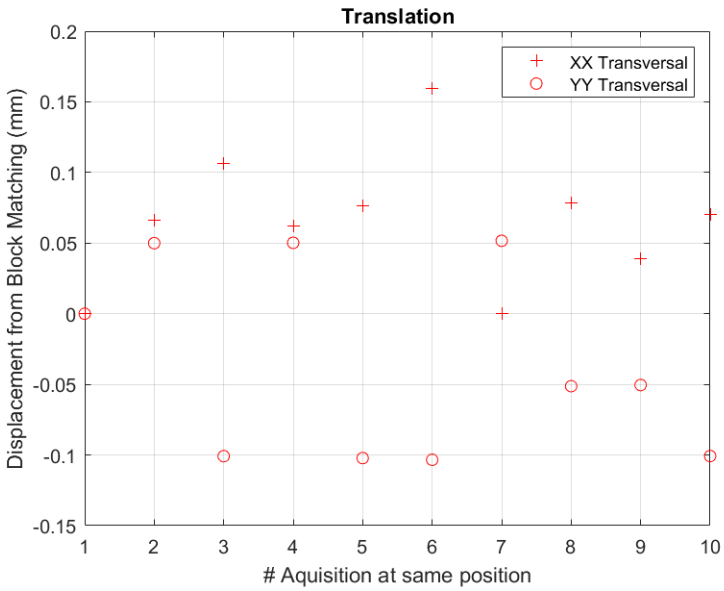

Figure 5.14 - Computed translations for each frame of the sequence, for the XX and YY axis, in the transversal plane of the phantom, using the BM algorithm.

For the BIG translations sequence, each of the five positions of the phantom were imaged in three different slices of the same transversal plane, ten times each, to generate statistics, so in the end, the sequence is made up of a hundred and fifty images, ten of each slice in each position. In Figure 5.15 we can observe the computed translations of the induced phantom movements for each slice imaged, with an added linear fit and respective mean position for each frame.

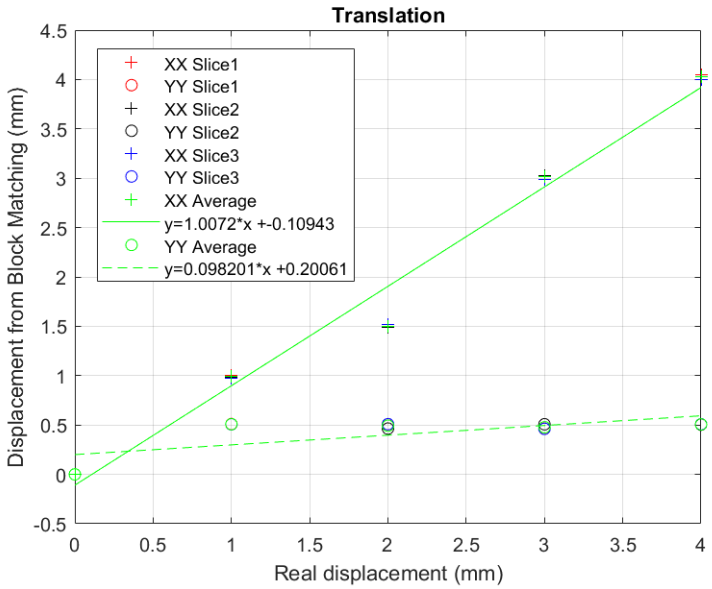

Figure 5.15 - Computed translations for the "BIG translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the BM algorithm.

TABLE 5.5 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION IN THE "BIG TRANSLATIONS" SEQUENCE, FOR THE XX AND YY AXIS, AND XX LINEAR FIT'S SLOPE, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE BM ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear fit slope in the XX axis	1.0072	0.9920
Mean STD for the 3 slices in the XX axis	19.9 µm	25 µm
Mean STD for the 3 slices in the YY axis	13.7 µm	5.3 µm

From the data collected, shown in Table 5.5, we can retrieve a linear fit slope for the XX axis of 1.0072 which corresponds to a computed average of 1.0072 mm per translation, with an average of the standard deviations of 19.9 µm for the XX axis and 13.7 µm for the YY axis, meeting the expectations. Although the value of the computed translation for the second image frame seems to be 0.5 mm too short, the computed slope of the linear fit, as well as the rest of the data shown, indicates the accuracy and precision of the computed translations given by the algorithm in this range of motion.

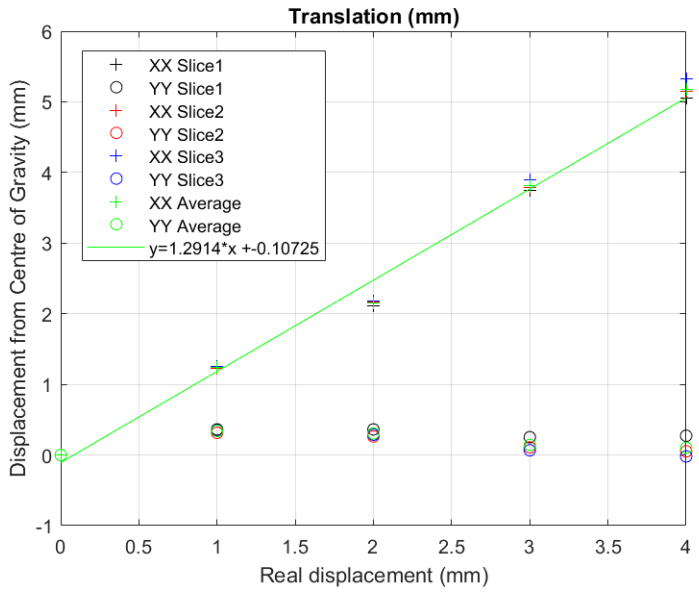

Figure 5.16 - Computed translations for the "BIG translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the CoG algorithm.

To clear the suspicions that the value observed in the translation to the second image frame was not an error in the algorithm's process, we ran the image sequence through a CoG (Centre of Gravity) algorithm, and the results are shown in Figure 5.16. This data clearly shows the same type of behavior for the second translation computed, leading us to believe that those images have an experimental error. We can then conclude that the algorithm is probably computing the correct translations. Figure 5.16 also helps to understand the fact that led us to abandon this CoG algorithm. Although it shows very good linearity, the precise values of the translations come up with significant errors, reflected in the calculated slope of 1.29.

For the SMALL translations sequence, each of the five positions of the phantom were imaged in three different slices of the same transversal plane as well, five times each, so in the end, the sequence is made up of seventy-five images, five of each slice in each position. In Figure 5.17 we can observe the computed translations of the induced phantom movements with 0.1 mm steps, with an added linear fit and respective mean positions for each frame. Analyzing the graphic shown it is clear that the algorithm cannot detect the sub-pixel step translations implied. Even though the algorithm detects that there is movement present in the image sequence (the detected movement always appears as the size of the pixel

length), it cannot specify the sub-pixel lengths in which it is moving, as we expected through the results of the digital tests performed earlier.

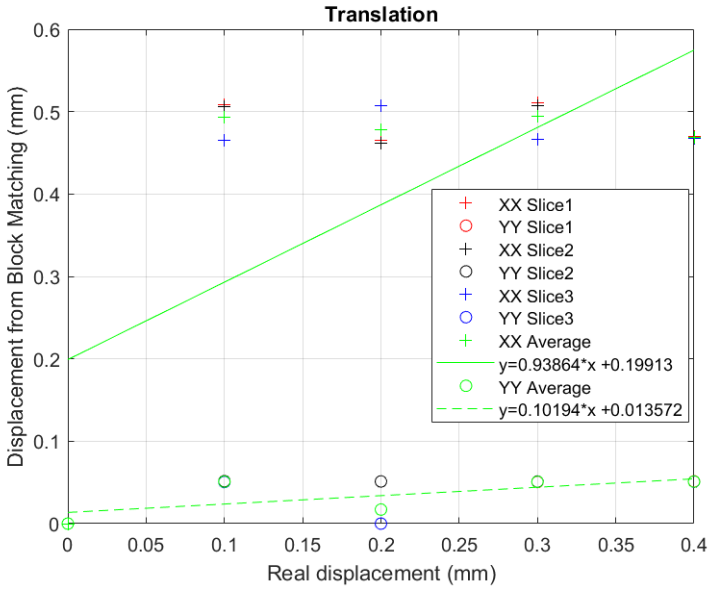

Figure 5.17 - Computed translations for the "SMALL translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the BM algorithm.

TABLE 5.6 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION IN THE "SMALL TRANSLATIONS" SEQUENCE, FOR THE XX AND YY AXIS, AND XX LINEAR FIT'S SLOPE, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE BM ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear fit slope in the XX axis	0.93864	1.0456
Mean STD for the 3 slices in the XX axis	18.97 µm	21.8 µm
Mean STD for the 3 slices in the YY axis	7.6 µm	3.5 µm

5.2.2 SURF Algorithm Results

In Table 5.7 and Figure 5.18 we can analyze the computed translations returned by the SURF algorithm for the ten frames of the phantom exactly in the same position, and the standard deviation of

the positions obtained for the XX and the YY axis, respectively. From this analysis we can calculate a mean standard deviation of 27.8 µm for the XX axis and 29.7 µm for the YY axis, as well as a rotational

TABLE 5.7 - COMPUTED TRANSLATIONAL STANDARD DEVIATION FOR THE IMAGE SEQUENCE OF THE PHANTOM IN THE SAME POSITION, FOR THE XX AND YY AXIS, AND ROTATIONAL STANDARD DEVIATION, USING THE SURF ALGORITHM.

	Standard Deviation
XX Axis	27.8 µm
YY Axis	29.7 µm
Angle	0.0118º

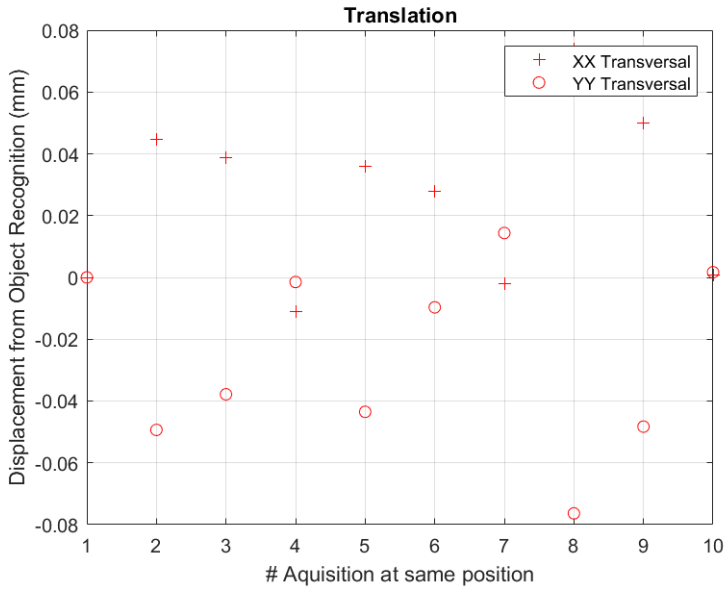

Figure 5.18 - Computed translations for each frame of the sequence, for the XX and YY axis, in the transversal plane of the phantom, using the SURF algorithm.

standard deviation of 0.0118 degrees.

The BIG translations sequence analyzed here is the same as described above, for the BM algorithm. In Figure 5.19 we can observe the computed translations of the induced phantom movements for each slice imaged, with an added linear fit and respective mean position for each frame.

From the data collected, seen in Table 5.8, we can retrieve a linear fit slope for the XX axis of 1.0237 which corresponds to a computed average of 1.0237 mm per translation, with an average of the standard deviations of 145.6 µm for the XX axis and 146.7 µm for the YY axis, meeting the expectations. Although the value of the computed translations does not always represent 1 mm, the computed slope of the linear fit, as well as the rest of the data shown, indicate the veracity of the computed translations given by the algorithm in this range of motion. It's also notable a slight decline in the mean computed XX position of the second translation, just as mentioned in page 55.

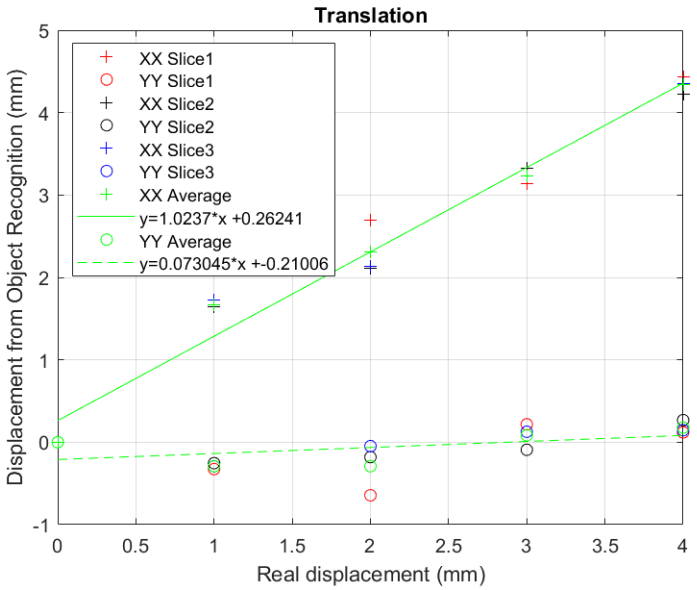

Figure 5.19 - Computed translations for the "BIG translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the SURF algorithm.

TABLE 5.8 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION, IN THE "BIG TRANSLATIONS" SEQUENCE, FOR THE XX AND YY AXIS, AND XX LINEAR FIT'S SLOPE, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE SURF ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear fit slope in the XX axis	1.0237	1.0245
Mean STD for the 3 slices in the XX axis	145.6 µm	190.2 µm
Mean STD for the 3 slices in the YY axis	146.7 µm	242.3 µm

The SMALL translations sequence analyzed here is also the same as the one described for the BM algorithm. In Figure 5.20 we can observe the computed translations of the induced phantom movements with 0.1 mm steps, with an added linear fit and respective mean position for each frame. Here, the results seem more promising that on the BM algorithm when analyzing the mean translational positions of the three different slices, showing a perception of increased movement between each step translation. Even though the computed translations are not in steps of 0.1 mm, the linearity of the data, shown in Table 5.9, presents a linear slope of 1.7497, corresponding to steps of approximately 0.2 mm between image frames.

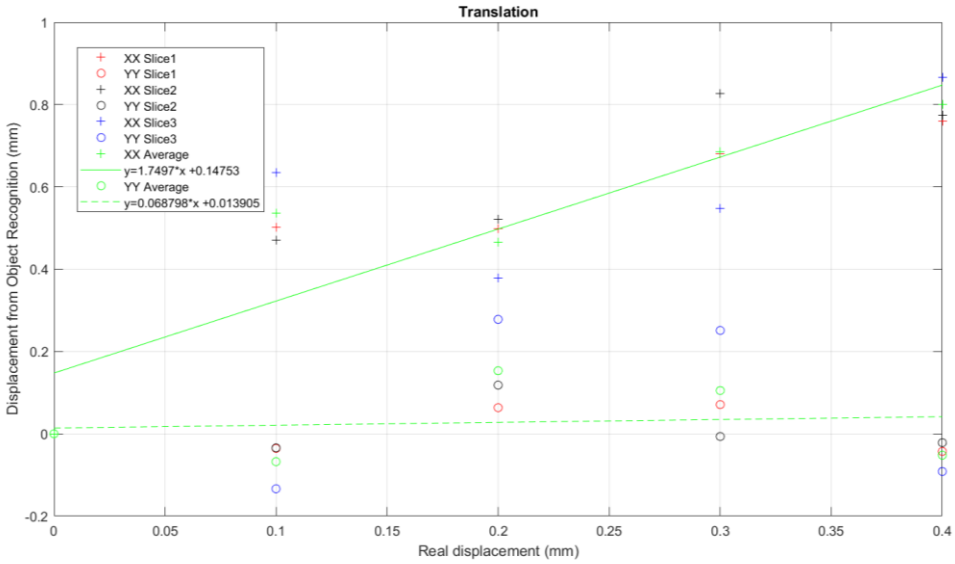

Figure 5.20 - Computed translations for the "SMALL translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the SURF algorithm.

TABLE 5.9 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION, IN THE "SMALL TRANSLATIONS" SEQUENCE, FOR THE XX AND YY AXIS, AND XX LINEAR FIT'S SLOPE, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE SURF ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear fit slope in the XX axis	1.7497	1.5387
Mean STD for the 3 slices in the XX axis	90.8 µm	59.4 µm
Mean STD for the 3 slices in the YY axis	84.2 µm	57.5 µm

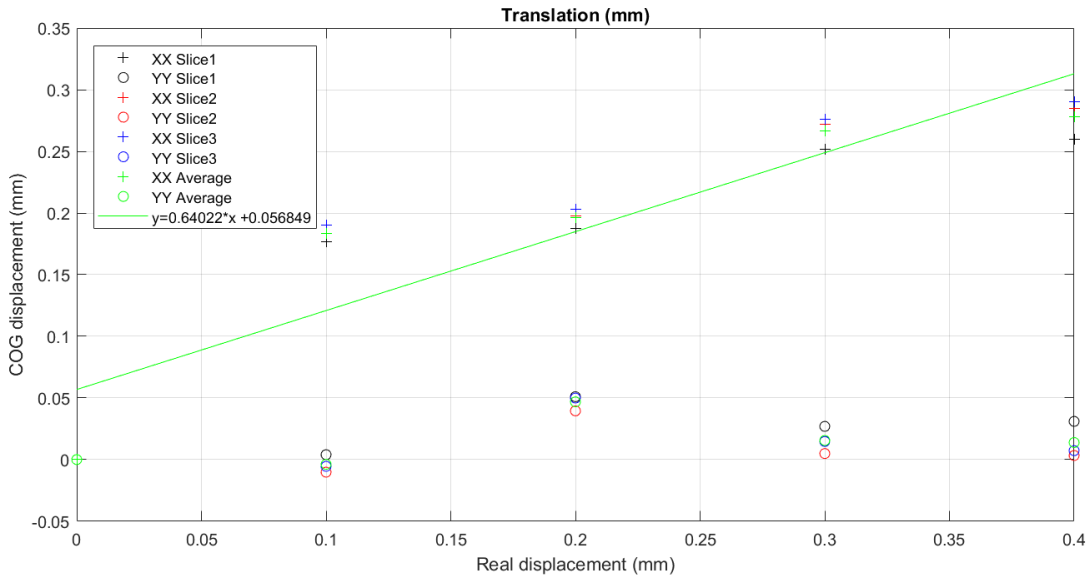

Figure 5.21 - Computed translations for the "SMALL translations" image sequence, for each imaged slice and for the XX and YY axis, and respective mean positions and linear fits, using the CoG algorithm.

We used again the Centre of Gravity algorithm to help confirm the linearity of the real translations created in the image sequence. As we can see in Figure 5.21, the computed rotations spread, calculated by this algorithm, matches the results given by the SURF algorithm, although the length of the translations are different, indicating that there is in fact a probable experimental error in the movement phantom's imaging (which is highly likely due to the small length of the translations).

For the BIG and SMALL rotations sequences, each of the five positions of the phantom were imaged in three different slices of the same transversal plane as well, five times each, leaving each sequence with seventy-five images, five of each slice in each position. Taking into account the small size of the rotations to be analyzed, the oversized image sequences were the ones who presented the best results. In Figure 5.22 and Table 5.10 we can analyze the computed rotations obtained for the BIG rotations sequence, oversized by a factor of 3, with an addition of a linear fit for each of the slices as well as a linear fit for the mean rotations calculated (to note that the values obtained are negative due to the anti-clockwise induced movement). With an average linear slope of 1.1352, we can observe that the rotational steps are not exactly 0.395º, but the proximity to this value, along with the good linearity leaves us to trust the rotations computed by the algorithm in this range of motion. In Figure 5.23 and Table 5.11 we can analyze the computed rotations obtained for the SMALL rotations sequence, oversized by a factor of 3, corresponding to an image oversizing factor of three (since it shows an improved performance in this specific case), with an addition of a linear fit for each of the slices as well as a linear fit for the mean rotations calculated. The results obtained for this sequence are worse than for the previous one in terms of linearity, but with an average linear slope of 1.4266, and considering the smallness of the induced rotations implied, we can be satisfied with the performance of the algorithm in such a minimal range of motion.

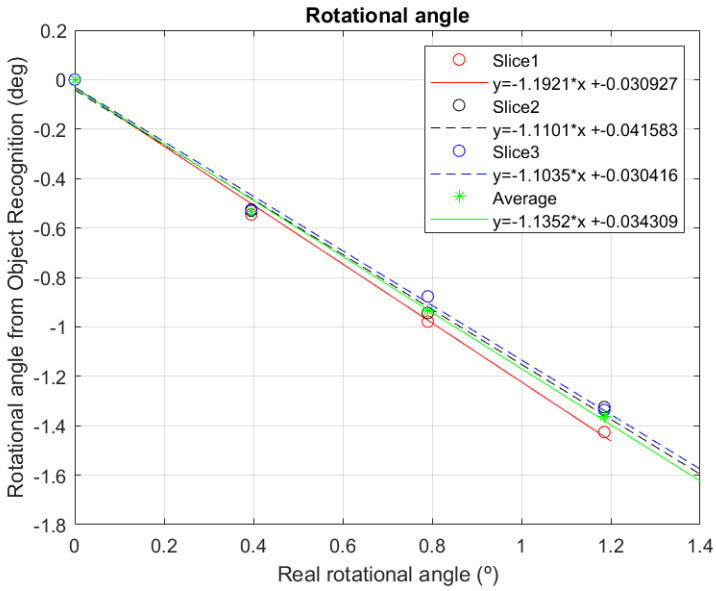

Figure 5.22 - Computed rotations for the "BIG rotations" image sequence, oversized by a factor of three, for each imaged slice, and respective mean positions and linear fits, using the SURF algorithm.

TABLE 5.10 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION, IN THE "BIG ROTATIONS" SEQUENCE, AND LINEAR FIT'S SLOPE OF THE MEAN POSITIONS, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE SURF ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear Fit slope	1.1363	1.1352
Mean angle STD for the 3 slice positions	0.044º	0.013º

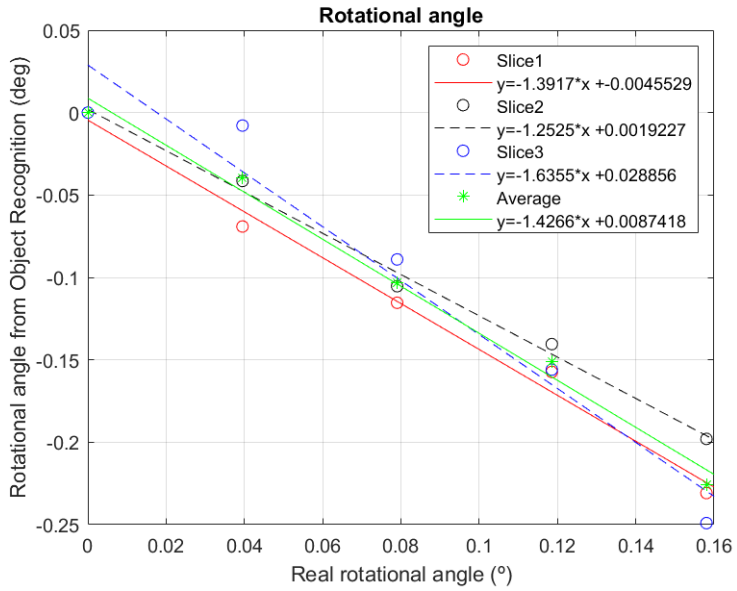

Figure 5.23 - Computed rotations for the "SMALL rotations" image sequence, oversized by a factor of three, for each imaged slice, and respective mean positions and linear fits, using the SURF algorithm.

TABLE 5.11 - COMPUTED MEAN STANDARD DEVIATIONS FOR THE THREE SLICES IN EACH POSITION, IN THE "SMALL ROTATIONS" SEQUENCE, AND LINEAR FIT'S SLOPE OF THE MEAN POSITIONS, FOR THE IMAGE OVERSIZING FACTOR OF ONE AND THREE, USING THE SURF ALGORITHM.

	Image Oversizing Factor	
	1	3
Linear fit slope	0.96202	1.4266
Mean angle STD for the 3 slice positions	0.017º	0.013º

5.3 Software Tests Conclusions

We can now conclude that both algorithms seem to be capable of accurately calculating translational and rotational movements in the millimeter range, which objectively is the goal of the software in this project. Analyzing Table 5.4 and Table 5.7 it becomes clear that the SURF algorithm presents less intrinsic noise associated to the standard deviation of the calculated translations, it shows a 20 µm lower STD in the XX axis and an approximately 40 µm lower STD in the YY axis. Although the BM algorithm does not recognize translational movements bellow the pixel length, it shows great accuracy and reliability above that threshold, even surpassing the SURF algorithm's performance in the millimeter range. For the "BIG" translations, the BM algorithm shows STDs in the order of approximately 20 µm while the SURF algorithm presents them in the order of 200 µm. The SURF algorithm still shows good results in the millimeter range, with the benefit of detecting sub-pixel translations, shown when analyzing translational steps of 0.1 mm. When analyzing the "SMALL" translations, looking at Figure 5.17 we can observe that, even though the BM algorithm does not differentiate the sub-pixel 0.1 mm translations, it is able to detect that some kind of translation has happened, setting its value in the length of the pixel. The rotational results obtained from the SURF algorithm also demonstrate great precision and reliability, providing plausible results even in a scale of 0.0395 degree steps. It's also worth mentioning that some of the errors detected in the computed rotations and translations may have been experimentally introduced, as observed in 5.2.1 and 5.2.2, with a higher probability of occurrence in the "SMALL" translations and rotations sequences.

Chapter 6
Immobilization Tests and Results

In this Chapter, we will present the results obtained from the in vivo tests performed to the immobilization device. The immobilization tests were realized in the Magnetic Resonance, at ICNAS, Coimbra, by submitting a volunteer to three main tests: one where the volunteer applied forced movements with the maximum possible force, another one where he applied movements with 50% of his maximum force, and a final one where he simply tried to be at rest, as still as possible. All three tests were performed four times, two acquisitions with and without the immobilization device, to allow the evaluation of the system's performance, and two acquisitions in the sagittal and the axial planes. In the sagittal plane was chosen a medial slice to image. In the axial plane was imaged a slice approximately in the middle of the brain. Twelve image sequences were created in total to analyze "Rest", "Small" and "Big" movements, with and without the immobilization device, in the sagittal and transversal plane of the system. All the sequences imaged were T1-weighted, with each frame being acquired every 0.4 seconds, resulting in a total of 340 frames and a time of 2 minutes and 26 seconds for each sequence to be imaged. The movements applied in each imaged plane were in concordance with that same plane's orientation, as we see in Figure 6.1. The image sequences obtained were then processed by our software, with the SURF and BM algorithms, to reach the data exposed bellow.

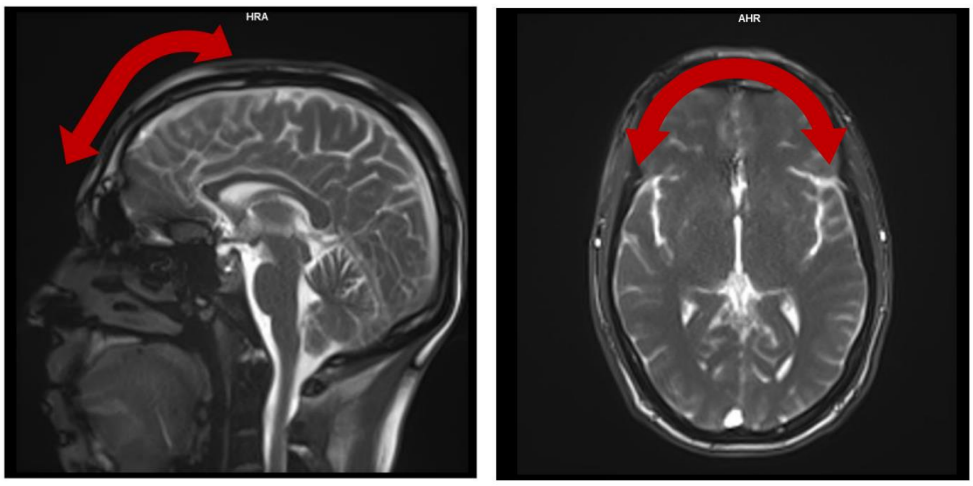

Figure 6.1 - Sagittal head motion, indicated by the red arrows, on the left, and Axial head motion, indicated by the red arrows, on the right.

6.1 SURF Algorithm Results

The results are quantified here as the standard deviation of the movement detected, without the drift of the head throughout the sequence, which was calculated through the smoothing curve of the original data and was considered by itself one of the most important points in the results, considering

the usual longer examination times utilized in PET. All the random errors are calculated by quadratically deconvoluting the intrinsic image noise of 0.0118º, 27.8 µm in XX and 29.7 µm in YY, determined before in Table 5.7, to the STDs obtained. The results obtained for the analyses of the data in the sagittal plane, at rest, without and with the immobilization system, are shown in Figure 6.2 and Figure 6.3, respectively, and are compiled in Table 6.1.

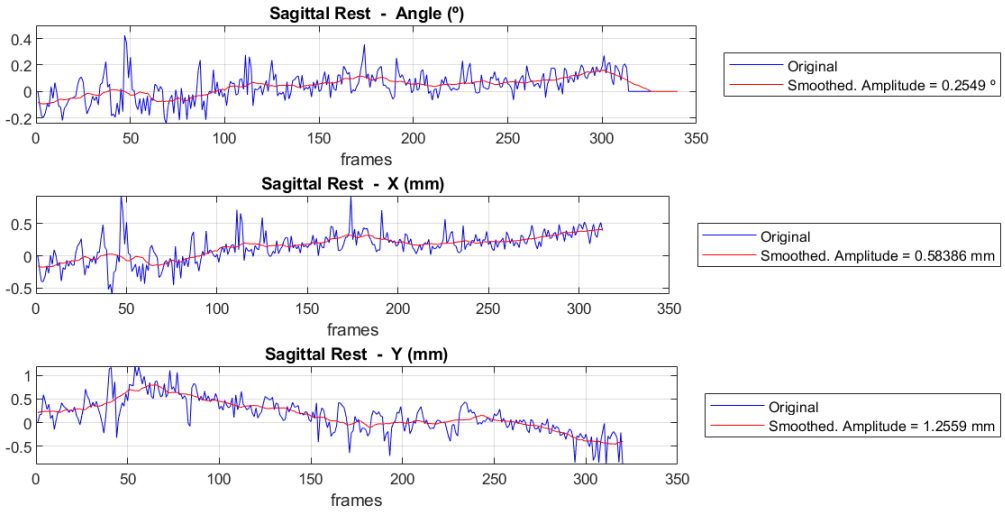

Figure 6.2 - Computed rotations and translations, in the XX and YY axis, for the sagittal rest sequence without the immobilization device. The smoothed amplitude indicates the head drift during the acquisition.

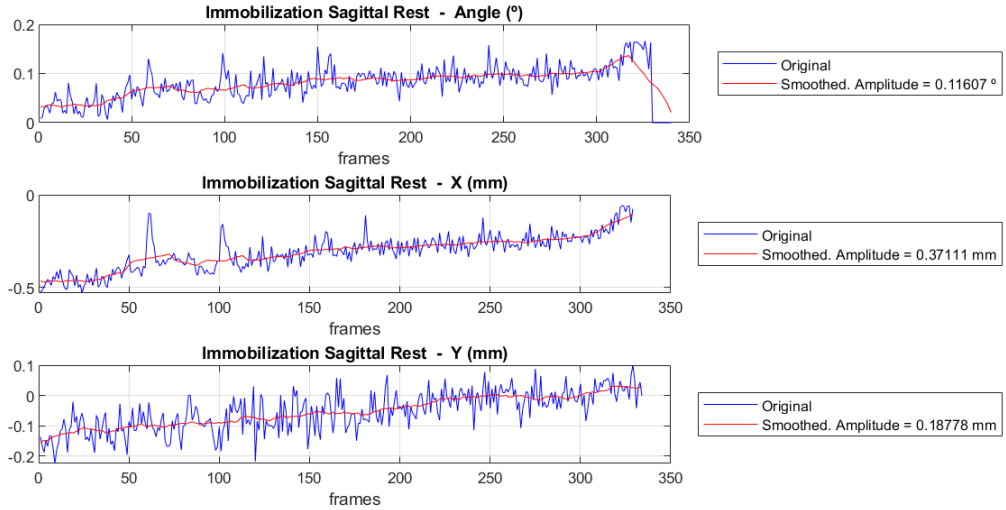

Figure 6.3 - Computed rotations and translations, in the XX and YY axis, for the sagittal rest sequence with the immobilization device. The smoothed amplitude indicates the head drift during the acquisition.

TABLE 6.1 - COMPUTED HEAD DRIFT AMPLITUDES AND STANDARD DEVIATIONS FOR THE SAGITTAL REST SEQUENCES, WITHOUT AND WITH THE IMMOBILIZATION DEVICE, USING THE SURF ALGORITHM.

	Without the immobilization device			With the immobilization device		
	Angle (º)	X (µm)	Y (µm)	Angle (º)	X (µm)	Y (µm)
Drift Amplitude	0.255	584	1256	0.116	371	188
STD (without the drift)	0.086	171	216	0.024	48	45
Deconvoluted STD	0.085	169	213	0.021	39	33

As we can see in these results, the usage of the immobilization device highly reduces both the drift amplitude and the standard deviation of the movements, reducing the angular drift to less than half. In the XX axis the drift amplitude decreased in about 36% and in the YY axis this decrease was even larger, going from 1256 µm to 188 µm. The random error, analyzed through the deconvoluted STDs, suffered an angular reduction of almost 80%, and decreased from approximately 200 µm to values around 40 µm in the XX and YY axis.

While analyzing the plotted data referent to the Rest sequence, without immobilization, regular frequency oscillations seemed to appear. By applying the Fast Fourier Transform (FFT) to this data, a distinguishable peak was found in the frequency's surrounding the 0.2 Hz, like we see in Figure 6.4.

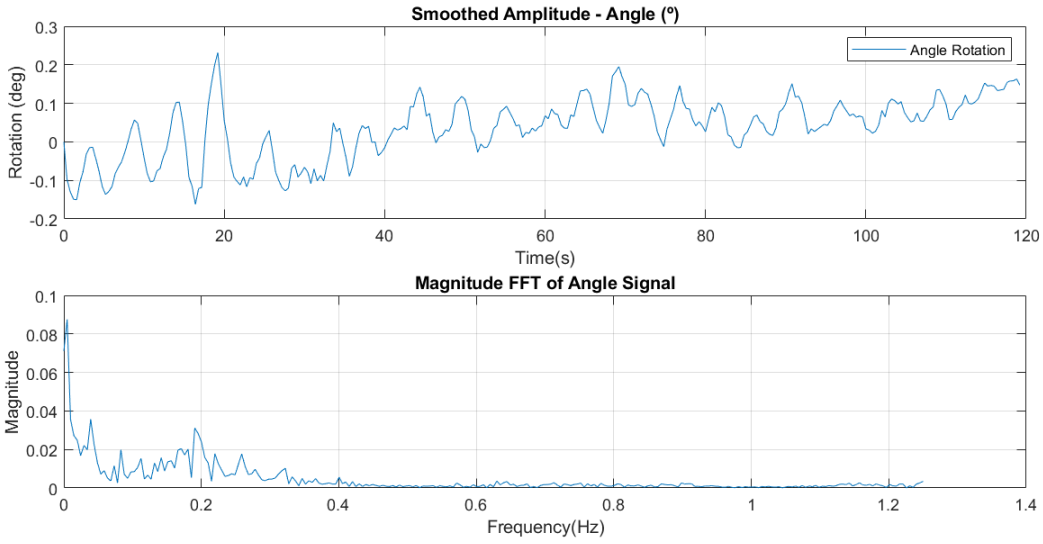

Figure 6.4 - Smoothed amplitude of the computed rotations for the sagittal rest sequence, without the immobilization device, and correspondent FFT analysis.

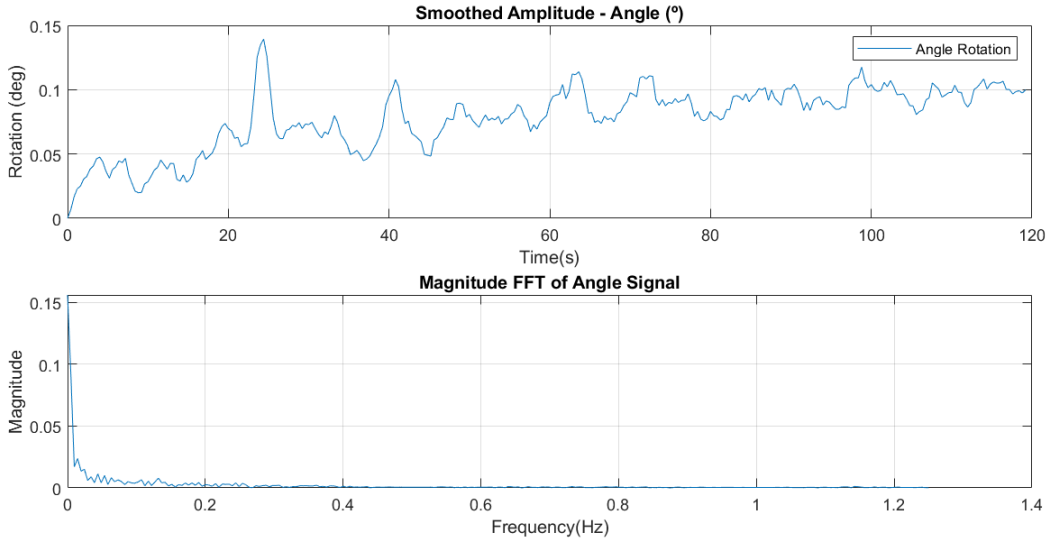

Figure 6.5 - Smoothed amplitude of the computed rotations for the sagittal rest sequence, with the immobilization device, and correspondent FFT analysis.

This frequency was found to be in the range of the respiratory frequency, which usually situates between 12 to 20 breaths per minute (Flenady, et al., 2017), with the lower level corresponding to 0.2 breaths per second (therefore 0.2 Hz), showcasing once again the remarkable sensitivity of the algorithm to small movements in the image sequences. The FFT was also applied to the data referent to the Rest sequence with the immobilization device, where the 0.2 Hz peak was no longer visible, as we see in Figure 6.5, uncovering another immobilization benefit of the system developed.

For the axial plane, at rest, the results obtained are shown in Figure 6.6 and Figure 6.7, and compiled in Table 6.2. The analysis's results are presented in the same way as for the sagittal plane. Analyzing these results, we can observe that the immobilization system reduced angular head drift amplitude to around half of their original value. In the XX axis the drift amplitude decreased in about

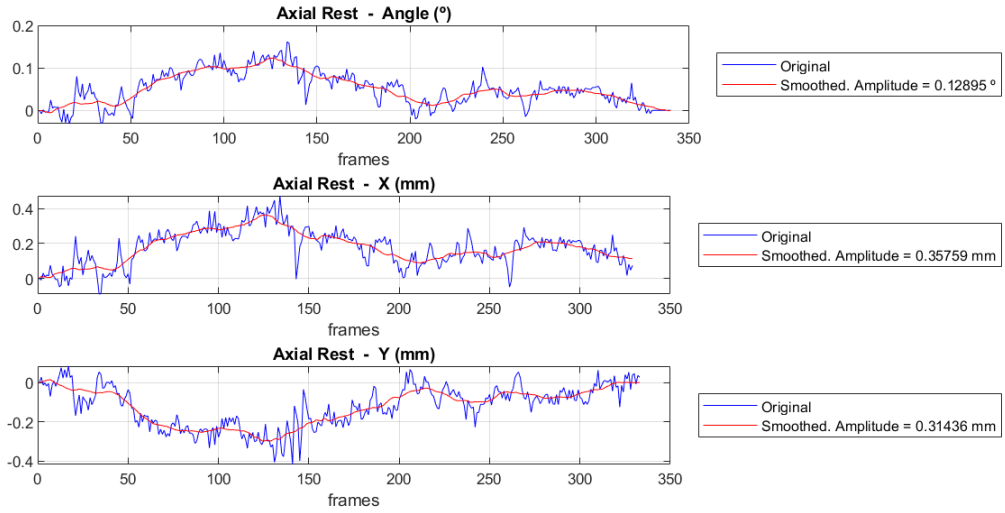

Figure 6.6 - Computed rotations and translations, in the XX and YY axis, for the axial rest sequence without the immobilization device. The smoothed amplitude indicates the head drift during the acquisition.

32% and in the YY axis this decrease was even larger, decreasing head drift in around 60%. The random error, analyzed through the STDs, decreased from approximately 40 µm to values around 20 µm in the XX and YY axis. A reduction in 86% is also notable in the angular STDs.

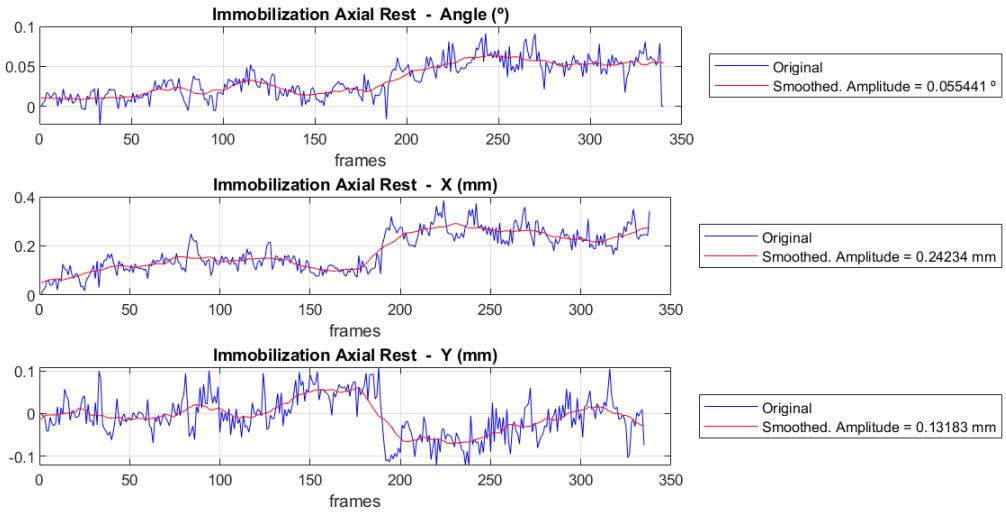

Figure 6.7 - Computed rotations and translations, in the XX and YY axis, for the axial rest sequence with the immobilization device. The smoothed amplitude indicates the head drift during the acquisition.

TABLE 6.2 - COMPUTED HEAD DRIFT AMPLITUDES AND STANDARD DEVIATIONS FOR THE AXIAL REST SEQUENCES, WITHOUT AND WITH THE IMMOBILIZATION DEVICE, USING THE SURF ALGORITHM.

	Without the immobilization device			With the immobilization device		
	Angle (º)	X (µm)	Y (µm)	Angle (º)	X (µm)	Y (µm)
Drift Amplitude	0.129	358	314	0.055	242	131
STD (without the drift)	0.020	55	50	0.012	36	34
Deconvoluted STD	0.016	47	40	0.002	23	17

When comparing the sagittal and axial immobilizations, the results show that the immobilization in the drift amplitude in angle and in the XX axis are very similar in both planes, with the drift amplitudes in the YY axis presenting greater decreases in the sagittal plane, even though the drift value calculated in the axial plane is still lower than in the sagittal plane. For the random errors, the results show greater

decreases overall in the sagittal plane, but the STD values calculated with immobilization are still lower in the axial plane.

For the movement ranges of the Small and Big sequences, seen from Figure 6.9 to Figure 6.15, and compiled in Table 6.3 and Table 6.4, for both planes, the head drift movement was not taken into account. There for, the results here are only quantified as the standard deviation of the movements detected. The goal of these forced movements sequences was mainly to test the immobilization device's limits.

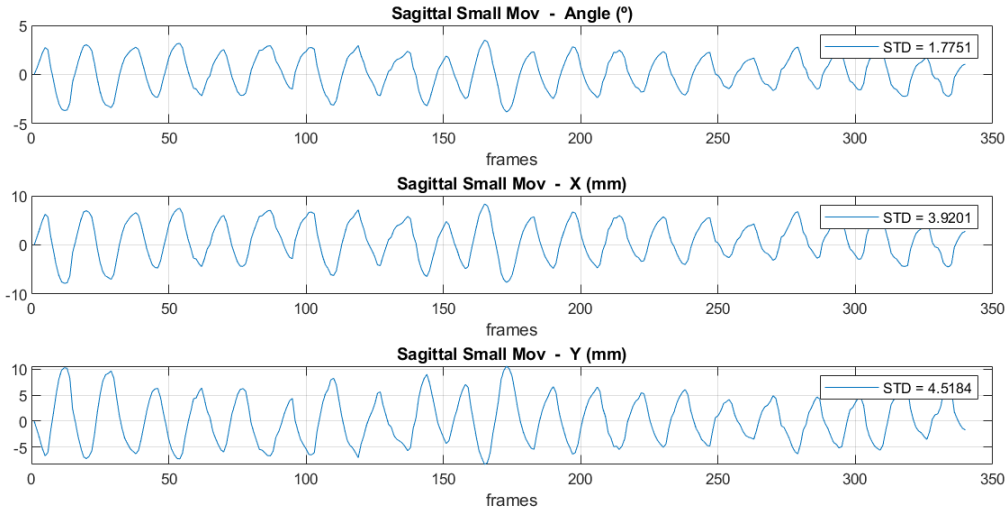

Figure 6.9 - Computed rotations and translations, in the XX and YY axis, for the sagittal small movements sequence without the immobilization device.

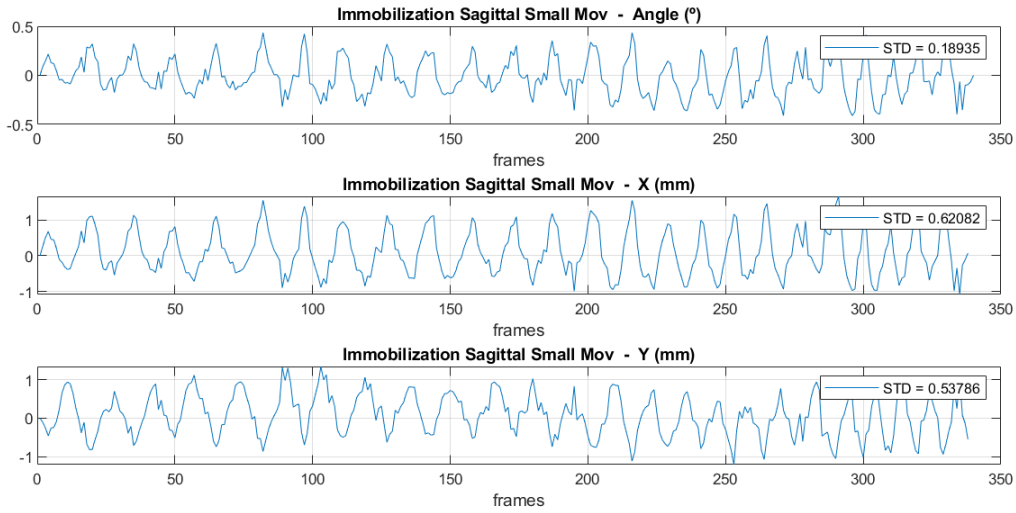

Figure 6.8 - Computed rotations and translations, in the XX and YY axis, for the sagittal small movements sequence with the immobilization device.

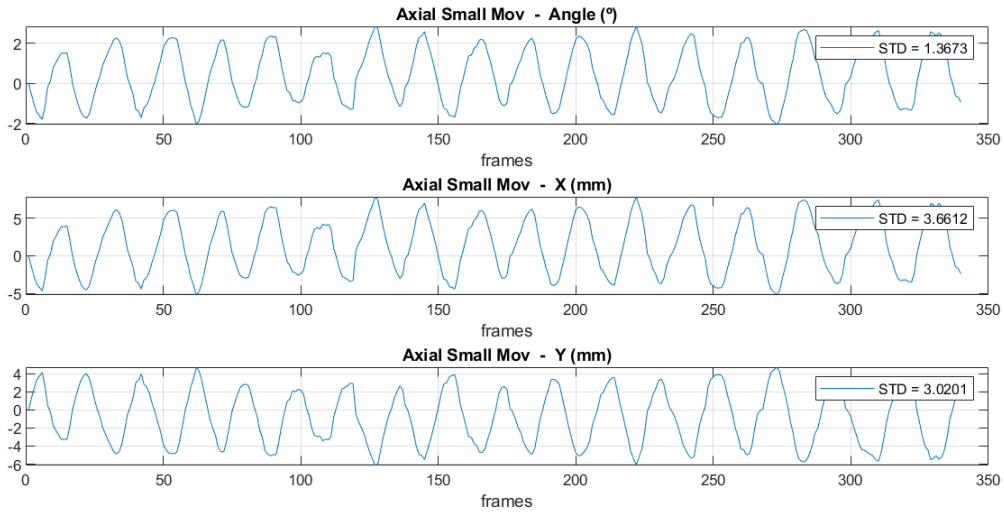

Figure 6.10 - Computed rotations and translations, in the XX and YY axis, for the axial small movements sequence without the immobilization device.

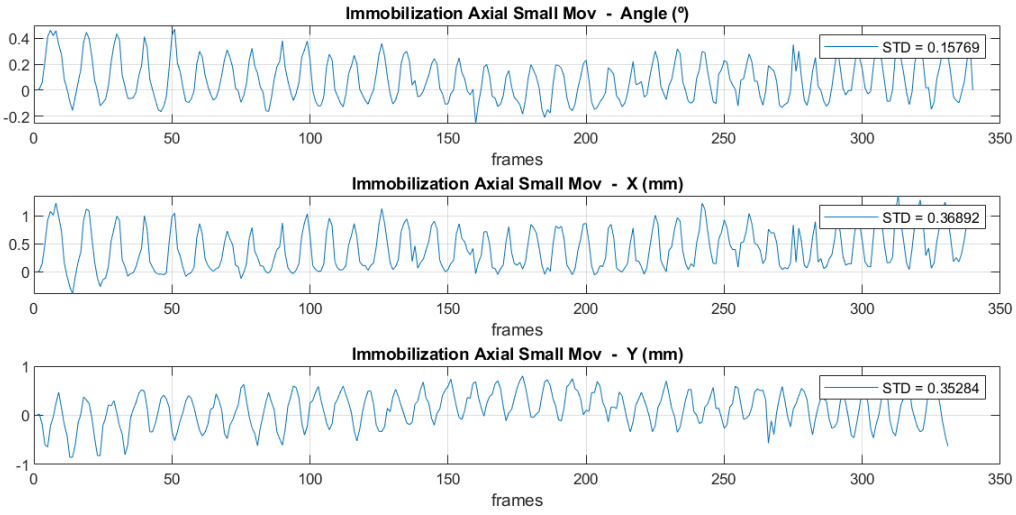

Figure 6.11 - Computed rotations and translations, in the XX and YY axis, for the axial small movements sequence with the immobilization device.

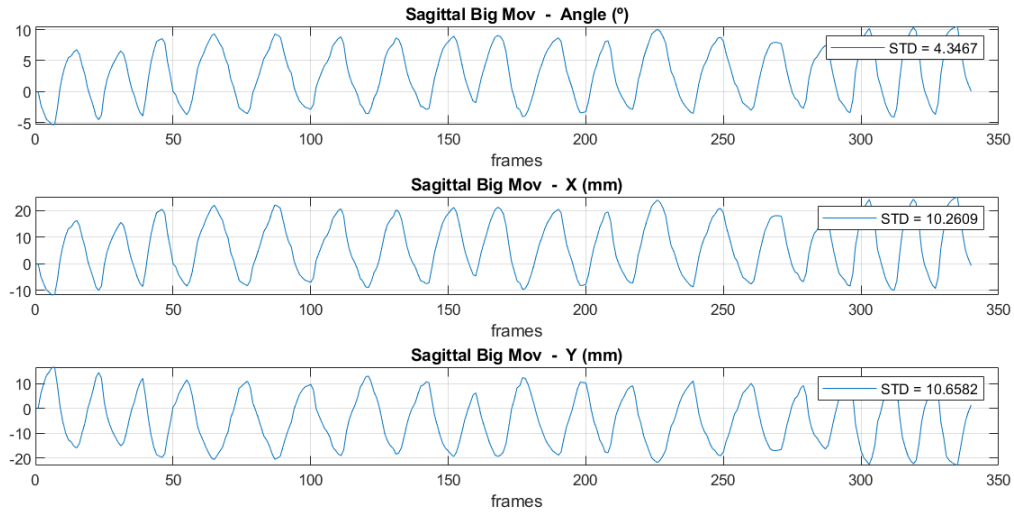

Figure 6.12 - Computed rotations and translations, in the XX and YY axis, for the sagittal big movements sequence without the immobilization device.

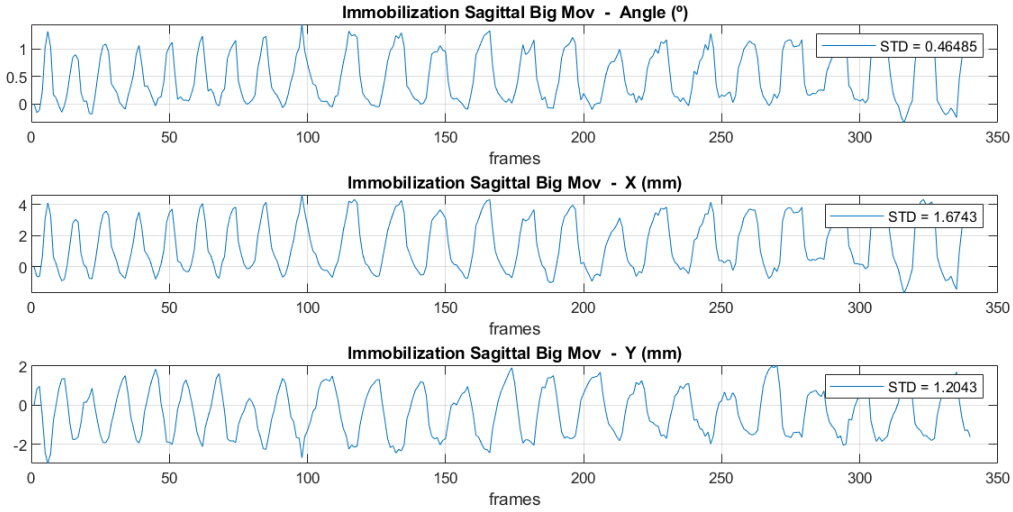

Figure 6.13 - Computed rotations and translations, in the XX and YY axis, for the sagittal big movements sequence with the immobilization device.

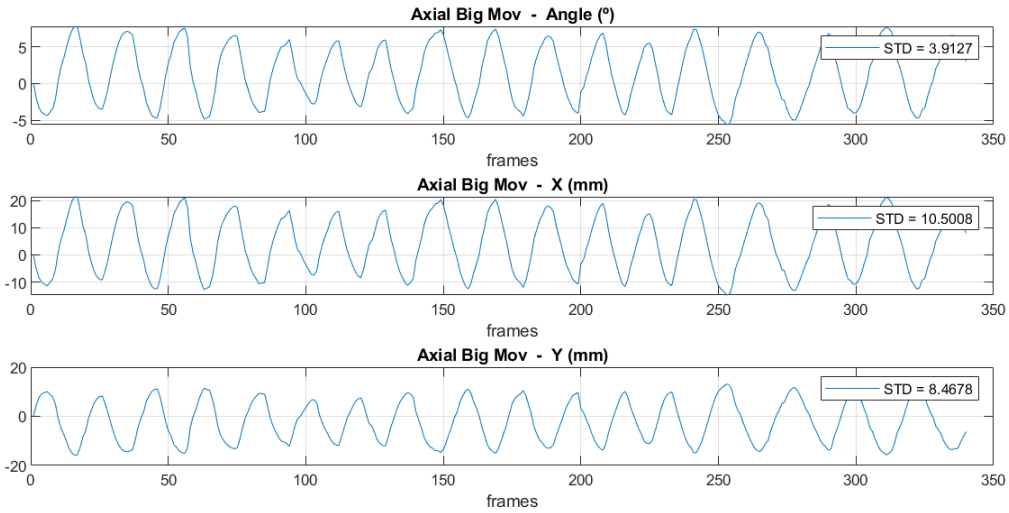

Figure 6.14 - Computed rotations and translations, in the XX and YY axis, for the axial big movements sequence without the immobilization device.

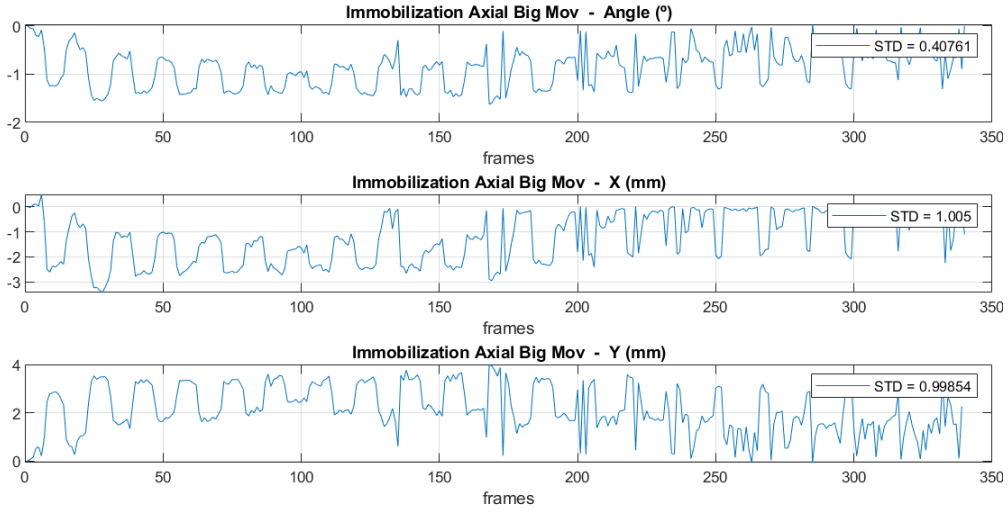

Figure 6.15 - Computed rotations and translations, in the XX and YY axis, for the axial big movements sequence with the immobilization device.

TABLE 6.3 - COMPUTED STANDARD DEVIATIONS FOR THE SAGITTAL SMALL AND BIG MOVEMENTS SEQUENCES, WITHOUT AND WITH THE IMMOBILIZATION DEVICE, USING THE SURF ALGORITHM.

	Without the immobilization device			With the immobilization device		
	STD Angle (º)	STD X (µm)	STD Y (µm)	STD Angle (º)	STD X (µm)	STD Y (µm)
Small Mov.	1.775	3290	4518	0.189	621	538
Big Mov.	4.347	10261	10658	0.465	1674	1204

TABLE 6.4 - COMPUTED STANDARD DEVIATIONS FOR THE AXIAL SMALL AND BIG MOVEMENTS SEQUENCES, WITHOUT AND WITH THE IMMOBILIZATION DEVICE, USING THE SURF ALGORITHM.

	Without the immobilization device			With the immobilization device		
	STD Angle (º)	STD X (µm)	STD Y (µm)	STD Angle (º)	STD X (µm)	STD Y (µm)
Small Mov.	1.37	3661	3020	0.16	369	352
Big Mov.	3.91	10500	8468	0.40	1005	998

From the data we can see that both the sagittal and the axial plane of the system present similar resistance to the maximum applied forces, although in the axial plane the allowed movement stops completely at a point, saturating the curves in the graphics of the Big movements sequences. For the Small movements, in the sagittal and axial planes, respectively, there is a reduction of 81% to 89% and 88% to 89% of the movements read through the standard deviations calculated. For the Big movements, in the sagittal and axial planes, respectively, there is a reduction of 83% to 89% and 89% to 90% of the translational movements, read through the standard deviations calculated. The rotations on both planes and both movement sequences are also reduced approximately 89% with the usage of the immobilization device.

6.2 Block Matching Algorithm Results

The results calculated with the BM matching algorithm were not considered from this point on because there were many contradictions with the results obtained in the movement phantom and digital tests, in comparison with the SURF algorithm. As we see in Figure 6.16 and Figure 6.17, as examples, the translations calculated for the Small and Big movements are not aligned with the ones obtained with the SURF algorithm, whose computed translations appear more realistic considering the type of movements performed in the imaging process.

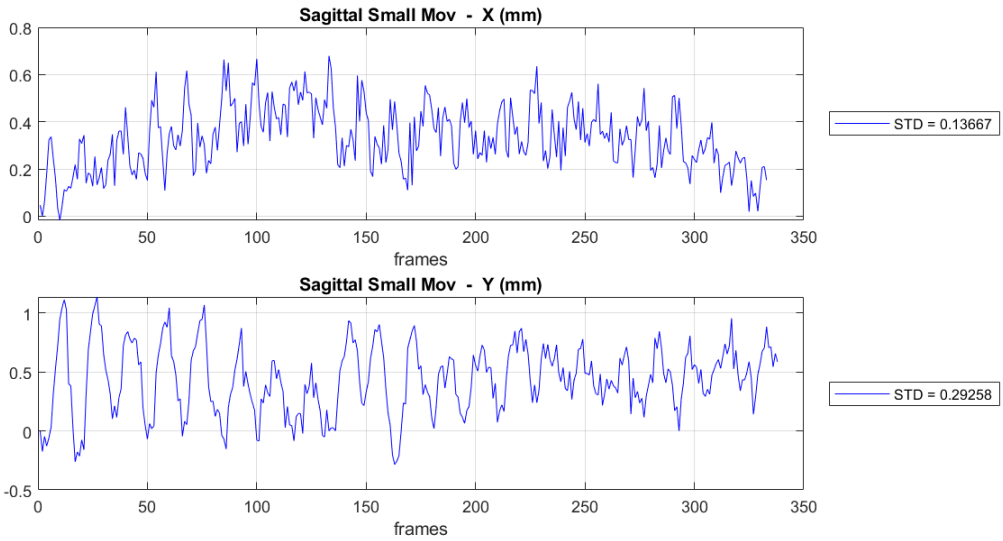

Figure 6.16 - Computed translations, in the XX and YY axis, for the sagittal small movements sequence without the immobilization device, using the BM algorithm.

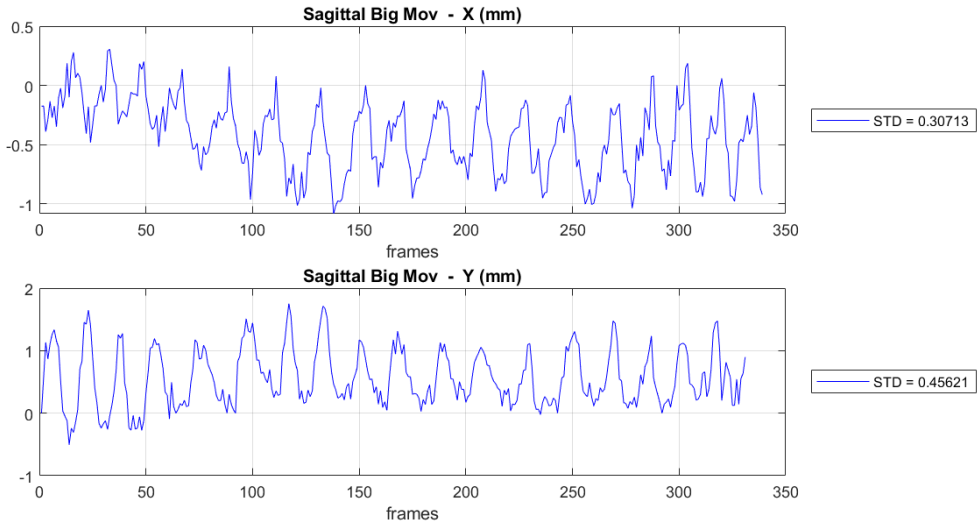

Figure 6.17 - Computed translations, in the XX and YY axis, for the axial big movements sequence without the immobilization device, using the BM algorithm.

There were many attempts to try and solve this problem in the scale of the translational amplitudes of the movements but with no solution reached, the algorithm was left aside. The hypothesis proposed for these results lies within the algorithm itself. Since the movements in these translational sequences are a lot bigger than the pixel length, the need to increase the maximum displacement value the algorithm can look for causes a significant increase in the noise of the movement vectors matrix, overlapping the histogram processing method used to eliminate it, consequently returning bad measures.

Chapter 7
Conclusions

In this work, non-invasive frame-based system capable of sub-millimetric immobilization of the head, for PET and MRI, was studied, developed and tested. The system was developed in Inventor® CAD software from Autodesk®, produced in the facilities of LIP Coimbra, and tested in the MAGNETOM Prisma Fit Magnetic Resonance, with a Head/ Neck 64 Coil, at ICNAS UC. The tests data were analysed by software developed in MATLAB®. To assess the accuracy of the movement detection, between other tests, a movement phantom was developed and imaged, as explained in Chapter 5. From there, an algorithmic intrinsic error of approximately 28 μm for the computed translations and 0.0118º for the rotations was calculated.

The final immobilization tests consisted in three different movement sequences performed with and without the immobilization device, imaged in the sagittal and axial planes. Two of the sequences, the "Big" and "Small" movements, consisted of the movement profile of the volunteer's head when applying his maximum force, and then 50% of his maximum force, respectively, to create periodical movements in the imaging plane's direction, with and without the immobilization device. The other sequence, also performed with and without the immobilization device, consisted simply of the movement profile of the volunteer's head when he is staying as still as possible, simulating a real MRI or PET head and neck exam. Each sequence was imaged for 2 minutes and 26 seconds.

From the results exposed in Chapter 6, for the Rest sequences, both in the sagittal and the axial planes, it's observable a clear reduction in the head drift during imaging and in the translational and rotational STDs. In both planes, angular head drift is reduced to approximately half of its non-immobilized values, from 0.255º to 0.116º in the sagittal plane, and from 0.129º to 0.055º in the axial plane. Random errors, represented as standard deviations were reduced in approximately 70% in the sagittal plane, decreasing from 0.086 degrees to 0.024 degrees. In the axial plane this reduction is even higher, of 86%, decreasing from 0.0161º to 0.002º.

For the sagittal plane, head translational drifts in the XX axis were reduced in about 36%, from 584 μm to 371 μm, and in the YY axis this decrease was even larger, going from 1256 μm to 188 μm (around 85% drift reduction). The random error, analyzed through the STDs decreased from approximately 200 μm to values around 40 μm in the XX and YY axis. In the axial plane, head translational drifts in the XX axis were reduced in about 32%, from 358 μm to 242 μm, and in the YY axis this decrease was also larger, going from 314 μm to 121 μm (around 58% drift reduction). The random error, analyzed through the STDs also decreased from approximately 40 μm to values around 20 μm in the XX and YY axis.

Both the head drift and the movement standard deviation results, obtained with the usage of the immobilization device, stay within the sub-millimetric range aimed. It is also notable that the head movements implied by the body's natural respiratory motion seem to be immobilized in these sequences.

For the forced Small and Big movements sequences, in the sagittal plane, the decrease in movement is more noticeable than in the Rest sequences. In the sagittal plane, the head translational movements presented a reduction of 81% to 88%, with the greater decreases showing for the YY image axis. In the axial plane, these reductions in movement went from 88% to 90%, but the greater decreases were seen in the XX image axis. It is also noticeable that the Big movements in the axial plane appeared saturated by the immobilization device. The rotations both in the sagittal and the axial plane, in both sequences, were reduced approximately 89%. With the use of the immobilization device, amplitudes of the Big forced movements decreased from approximately 30 mm to values around 4 mm, in the sagittal plane, and from values around 20 mm to approximately 3 mm, in the axial plane. For the Small forced movements, the movement amplitudes decreased from 15 mm to approximately 1.5 mm, in the sagittal plane, and from 10 mm to values around 1mm, in the axial plane. With these heavily forced movements, the worst translational standard deviation obtained was of approximately 1.6 mm. The results for the immobilization in the Small movements, however, show standard deviations of 500 / 600 μm.

Therefore, it seems clear that this immobilization system is capable of performing sub-millimetric immobilization in head and neck exams, bearing still space for improvements.

Future Work

Although this immobilization system proved itself capable of performing sub-millimetric immobilization, some aspects remain eligible for optimizations and improvements.

One big improvement would be patient comfort when utilizing the immobilization device. This would require added cushioning to the base plate, and minor changes in the mouthpiece design, including the development of different mouth size pieces to accommodate all types of patients. Some simple mechanisms could also be applied to this system to improve its overall capabilities like, for example, angle and metric indicators to save patient positions and improve the position repeatability between exams.

Finally, the immobilization capabilities of the device can also be improved. With the process of adapting this system from a MRI to a PET system, the possibility of utilizing stronger ferromagnetic materials rises. With this, the locking mechanisms, used to lock the device in position, could also be improved, changing the "low-cost" friction-based methods applied here to more reliable ones, with stronger materials and different locking principles.

Bibliography

Babic, Steven, et al. 2018. To frame or not to frame? Cone-beam CT-based analysis of head immobilization devices specific to linac-based stereotactic radiosurgery and radiotherapy. *Journal of Applied Clinical Medical Physics* . 2018, Vol. 19, 2.

Barjatya, Aroh. 2004. Block Matching Algorithms For Motion Estimation. 2004.

Bay, Herbert, et al. 2008. SURF: Speeded Up Robust Features. *Computer Vision and Image Understanding.* 2008, Vol. 110, 3.

Bertholet, Jenny, et al. 2019. Real-time intrafraction motion monitoring in external beam radiotherapy. *Physics in Medicine & Biology.* 2019, Vol. 64, 15TR01.

Bruijnen, Tom, et al. 2018. Intrafraction motion quantification and planning target volume margin determination of head-and-neck tumors using cine magnetic resonance imaging. *Radiotherapy and Oncology*. 2018, Vol. 130.

Cervino, Laura, et al. 2010. Frame-less and mask-less cranial stereotactic radiosurgery: A feasibility study. *Physics in Medicine and Biology.* 2010, Vol. 55, 7.

Chelvarajah, Radhini, et al. 2004. Cranial Immobilisation – Is There A Better Way? *Journal of Medical Radiation Sciences.* 2004, Vol. 51, 5.

Couceiro, M. 2014. *Simulation Study of a Wide Axial Field of View Positron Emission Tomography System Based on Resistive Plate Chamber Detectors.* Coimbra, Portugal : Universidade de Coimbra, Faculdade de Ciências e Tecnologia, Departamento de Física, 2014. (https://estudogeral.sib.uc.pt/handle/10316/23827.).

Dung, Lan-Rong, Huang, Chang-Min and Wu, Yin-Yi. 2013. Implementation of RANSAC Algorithm for Feature-Based Image Registration. *Journal of Computer and Communications.* 2013, Vol. 01, 06.

Engelsman, Martijn, et al. 2005. Intra- and interfractional patient motion for a variety of immobilization devices. *Medical Physics.* 2005, Vol. 32, 11.

Eulenburg, Peter zu and Maximilian, Ludwig. 2017. Head immobilization during fMRI with galvanic vestibular stimulation. *Pearl Technology*. 2017.

Fedorov, A., et al. 2012. 3D Slicer as an Image Computing Platform for the Quantitative Imaging Network . *Magnetic Resonance Imaging.* 30, 2012, Vol. 09, 1323-41.

Flenady, Tracy, Dwyer, Trudy Ann and Applegarth, Judith Anne. 2017. Accurate respiratory rates count: So should you! *Australasian Emergency Nursing Journal.* 2017, Vol. 20, 1.

Fonte, P., Peskov, V. and Abbrescia, M. 2018. *Resistive Gaseous Detectors: Designs, Performance, and Perspectives.* s.l. : Wiley, 2018. 978-3-527-34076-7.

Gach, H. Michael. 2019. Technical Note: T1 and T2 and Complex Permittivities of Mineral Oil, Silicone Oil, and Glycerol at 0.35, 1.5, and 3 T. *Medical Physics.* 2019, Vol. 46, 4.

Gong, Kuang, Cherry, Simon R and Qi, Jinyi. 2016. On the assessment of spatial resolution. *Physics in Medicine and Biology.* 2016, Vol. 61, 5.

Haffty, Bruce G. and Goyal, Sharad. 2018. *Precision Radiation Oncology.* 1. s.l. : Rutgers University Press, 2018. 0813585961.

Hansen, Christian Rønn, et al. 2014. Comparison of three immobilisation systems for radiation therapy in head and neck cancer. *Acta Oncologica.* 2014, Vol. 53, 3.

Havsteen, et al. 2017. Are Movement Artifacts In Magnetic Ressonance Imaging a Real Problem? - A Narrative Review. *Frontiers in Neurology.* 2017, Vol. 8, 232.

Herk, Marcel Van, et al. 2000. THE PROBABILITY OF CORRECT TARGET DOSAGE: DOSE-POPULATION HISTOGRAMS FOR DERIVING TREATMENT MARGINS IN RADIOTHERAPY. *Int. J. Radiation Oncology Biol. Phys.* 2000, Vol. 47, 4.

Hoskin, Peter, et al. 2008. *On target: ensuring geometric accuracy in radiotherapy.* RCR London UK : s.n., 2008. 1-905034-33-8.

Khedr, Ali, Alaminos Bouza, Armando and Brown, Russell. 2018. Use of the Brown-Roberts-Wells Stereotactic Frame in a Developing Country. *Cureus.* 2018, Vol. 10, 1.

Kulkarni, Shamsundar M., Bormane, D. S. and Nalbalwar, Sanjay laxmikant. 2017. RANSAC algorithm for matching inlier correspondences in video stabilisation. *International Journal of Signal and Imaging Systems Engineering.* 2017, Vol. 10, 4.

Li, G., et al. 2013. Migration from full-head mask to "open-face" mask for immobilization of patients with head and neck cancer. *Journal of Applied Clinical Medical Physics* . 2013, Vol. 14, 5.

Li, George, et al. 2015. Clinical experience with two frameless stereotactic radiosurgery (fSRS) systems using optical surface imaging for motion monitoring. *Journal of Applied Clinical Medical Physics.* 2015, Vol. 16, 4.

Li, George, et al. 2011. Correction of motion-induced misalignment in co-registered PET/CT and MRI (T1/T2/FLAIR) head images for stereotactic radiosurgery. *Journal of Applied Clinical Medical Physics.* 2011, Vol. 12, 1.

Li, George, et al. 2011. Motion monitoring for cranial frameless stereotactic radiosurgery using video-based three-dimensional optical surface imaging. *Medical Physics.* 2011, Vol. 38, 7.

Mandija, Stefano, et al. 2019. Brain and Head-and-Neck MRI in Immobilization Mask: A Practical Solution for MR-Only Radiotherapy. *Frontiers in Oncology.* 2019, Vol. 9.

Massanes, Francesc, Cadennes, Marie and Brankov, Jovan G. 2011. Compute-unified device architecture implementation of a block-matching algorithm for multiple graphical processing unit cards. *Journal of Electronic Imaging.* 2011, Vol. 20, 3.

MATLAB. 2021. *Version 9.11 (R2021b).* Natick, Massachusetts : The MathWorks Inc., 2021.

Meessen, S., et al. 2017. Immobilisation systems for brain treatment: are individual head supports needed for stable fixation? *Journal of the European Society for Radiotherapy and Oncology.* 2017, Vol. 123, 1.

Mesch, L., Essers, M. and Rozema, M. 2019. A clinical evaluation of the stability, patient comfort and ease in use of the new Nanor mask. *Journal of the European Society for Radiotherapy and Oncology*. 2019, Vol. 133, 1.

Oyallon, Edouard and Rabin, Julien. 2015. An Analysis of the SURF Method. *Image Processing On Line*. 2015, Vol. 5, 10.5201/ipol.2015.69.

Parekh, Ranjan. 2021. *2021 : Fundamentals of IMAGE, AUDIO, and VIDEO PROCESSING Using MATLAB®: With Applications To PATTERN RECOGNITION*. s.l. : CRC Press (Taylor and Francis group), 2021. 9781003019718.

Pell, Malcolm F and Thomas, David G. T. 1991. The initial experience with the Cosman-Roberts-Wells stereotactic system. *British Journal of Neurosurgery.* 1991, Vol. 5, 2.

Pépin, Audrey, et al. 2014. Management of respiratory motion in PET/computed tomography. *Nuclear medicine communications.* 2014, Vol. 35, 2.

Philip, Jobin T, et al. 2014. A Comparative Study of Block Matching and Optical Flow Motion Estimation Algorithms. *2014 Annual International Conference on Emerging Research Areas: Magnetics, Machines and Drives (AICERA/iCMMD).* 2014.

Ramakrishna, Naren, et al. 2010. A clinical comparison of patient setup and intra-fraction motion using frame-based radiosurgery versus a frameless image-guided radiosurgery system for intracranial lesions. *Radiotherapy and Oncology* . 2010, Vol. 95, 1.

Stefano, Luigi Di and Viarani, Enrico. 1999. Vehicle Detection and Tracking Using the Block Matching Algorithm. 1999.

Tekalp, A. Murat. 1995. *Digital Video Processing.* Upper Saddle River, NJ 07458 : Prentice Hall PTR, 1995. O-13-190075-7.

Thornton, A.F., et al. 1991. A head immobilization system for radiation simulation, CT, MRI, and PET imaging. *Medical dosimetry: official journal of the American Association of Medical Dosimetrists.* 1991, Vol. 16, 2.

Thornton, A.F., et al. 1991. Three-dimensional motion analysis of an improved head immobilization system for simulation, CT, MRI, and PET imaging. *Radiotherapy and Oncology.* 1991, Vol. 20, 4.

Trembath, LisaAnn, Newell, Maureen and Devous, Michael D. 2015. Technical Considerations for Brain Amyloid PET Imaging with F-18 Florbetapir. *Journal of Nuclear Medicine Technology.* 2015.

Varrone, Andrea, et al. 2009. EANM procedure guidelines for PET brain imaging using [18F]FDG, version 2. *European Journal of Nuclear Medicine and Molecular Imaging.* 2009, Vol. 36, 12.

Waplera, Matthias C, et al. 2014. Magnetic properties of materials for MR engineering, micro-MR and beyond. *Journal of Magnetic Resonance.* 242, 2014.

Waxman, A.D., Herholz, K. and Lewis, D.H. 2009. Society of nuclear medicine procedure guideline for FDG PET brain imaging. *Journal of Nuclear Medicine* . 2009.

Winey, Brian, et al. 2012. Immobilization precision of a modified GTC frame. *Journal of Applied Clinical Medical Physics.* 2012, Vol. 13, 3.