

*Machine Learning and Interactive Real-
Time Simulation For Training On Relevant
Total Hip Replacement Skills*

Thesis submitted for the Degree of Doctor of Philosophy

Mara Catalina Aguilera Canon

1 June 2021

Copyright Statement

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

A MIS PADRES

*por su apoyo y amor incondicional; y por enseñarme a luchar para
alcanzar todos mis sueños.*

TO MY PARENTS

*for their unconditional love and support, and for teaching me to fight for
my dreams.*

Acknowledgements

First, I would like to express my sincere gratitude to Bournemouth University for the studentship award, which made this research possible.

The completion of this thesis would not have been possible without the support of my supervisors. **Dr Hammadi Nait-Charif**, in whom I found a guide, a friend and an inspiration. **Dr Xiaosong Yang** and **Thomas W. Wainwright**, who provided patient guidance, enthusiastic encouragement and useful critiques during this research work.

I would like to express my very great appreciation to **Prof Robert Middleton**, for sharing their expertise in both traditional and robotic-assisted hip replacement. To the hip fellows visiting at ORI who provided feedback to my research progress. I am incredibly grateful for tremendous help during the first year of my research by **Kieran Gallagher**, **Paddy (Padmanabhan) Subramanian**, and **István Batta**. Thank you for providing in-depth and precise answers to my countless questions; for welcoming me into your surgical theatre and make sure I gained a clear understanding of the surgery. These special thanks are also extended to the staff of the **Royal Bournemouth and Christchurch Hospitals NHS Foundation Trust** and **Nuffield hospital**.

I must express my gratitude for their support and friendship to the whole Orthopaedic Research Institute (ORI) team, especially to **Tikki Immins**, **Shay Bahadori** and **Louise Burgess**. I am deeply grateful for their help.

I also want to thank my beloved friends **Alexander Andreevich Tereshin** and **Rantilini Samaratunga** who were of generous support in deliberating over our problems and findings, as well as providing a happy distraction to rest my mind outside of my research.

And finally, I would like to thank God Almighty for giving me the strength and knowledge to undertake this research and to persevere until completing it satisfactorily. Also, for blessing me with an unlimited source of love and support not only during this research but throughout my entire life; my mother **Elizabeth Cañon Vasquez** and my father **Salvador Aguilera Huerfano**.

Thank you all!

Contents

“First, think. Second, dream. Third, believe. And finally, dare.”

- Walt Disney

<i>Copyright Statement.....</i>	<i>1</i>
<i>Acknowledgements.....</i>	<i>5</i>
<i>List of Figures.....</i>	<i>9</i>
<i>List of Tables.....</i>	<i>13</i>
<i>Acronyms.....</i>	<i>1—14</i>
<i>Abstract.....</i>	<i>1—15</i>
1 Introduction.....	1—17
1.1 The hip joint.....	1—17
1.2 Total Hip Replacement (THR) surgery.....	1—18
1.3 Surgical training.....	1—23
1.4 Virtual and Augmented Reality simulators.....	1—24
1.5 Research Questions.....	1—27
1.6 Aims and objectives.....	1—28
1.7 Scope and Limitations.....	1—29
1.8 Thesis layout and remarks.....	1—30
2 Simulation-based training and learning theories.....	2-33
2.1 Serious games and surgical education.....	2-33
2.2 Simulation in orthopaedic training.....	2-39
2.3 Discussion.....	2-46
3 An augmented reality-based simulator for training on surgical navigation skills....	3-48
3.1 Robotic-assisted surgery and Imageless Navigation Systems for THR.....	3-48
3.2 Visual interface, second version.....	3-58
3.3 A communication protocol between Android devices and HoloLens.....	3-60
3.4 User study.....	3-62
3.5 Summary.....	3-64
4 The Finite Element Method (FEM) for physically-based simulation.	4-66
4.1 The Finite Element Method (FEM).....	4-67
4.2 Realtime FEM simulation.....	4-74
4.3 Bone reaming – tissue deformation.....	4-80
4.4 Summary.....	4-86
5 XFEM-based reaming simulation.	5-87
5.1 Mesh generation.....	5-88
5.2 Material model.....	5-92
5.3 Boundary conditions.....	5-95
5.4 Results.....	5-98
5.5 Summary.....	5-102

6	<i>Machine learning and volumetric deformation.</i>	103
6.1	Data-driven / Machine learning (ML) methods for mesh deformation	103
6.2	Machine learning on graph-structured data.	108
6.3	Summary.	115
7	<i>Machine Learning for reaming simulation.</i>	117
7.1	Learned 3D physically-based simulator.	118
7.2	Feature engineering and data cleaning.	119
7.3	Graph network model.	120
7.4	Modern GL – openGL for python	131
7.5	Summary.	133
8	<i>Conclusions and future work</i>	134
8.1	Future work	137
9	<i>References</i>	141

List of Figures

Figure 1 The hip joint.....	1—17
Figure 2 Total Hip Replacement. Reaming and acetabular implant insertion.....	1—19
Figure 3 Hip Implant.....	1—19
Figure 4 X-ray pelvis with hip joints and proximal 1/3rd of femur anteroposterior view showing templating for an uncemented total hip arthroplasty. From Bhaskar (2017).....	1—21
Figure 5 Stryker’s MAKO robot. Assisting robotic arm for orthopaedic arthroplasty procedures on hip and knee joints. Photography provided by ORI BU.....	1—22
Figure 6 Visual references of the progress during the reaming stage. a) Unreamed acetabulum. b) Initial ream, the labrum is completely removed and the true floor is exposed. c) Appearance of bleeding bone on the reamed surface) Reamer inside the acetabulum. Screenshots from TOTAL HIP REPLACEMENT - UNCEMENTED. [onlinevideo] (Orthopaedics, 2016) Orthopaedics, T., 2016. Available at: https://www.youtube.com/watch?v=sMBx3C8hhJg&t=217s [Accessed September 11, 2018].....	1—23
Figure 7 Taxonomy of knowledge (Krathwohl and Anderson, 2009).....	2-36
Figure 8 (Right) Learning outcomes method. Figure adapted from (Zigmont <i>et al.</i> , 2011). (Left) the Kolb’s Experiential Learning Cycle. (Kolb, 1984).....	2-37
Figure 9 Results of scores achieved by each simulator after the scoring stage. The value of 3 was taken as the minimum requirement for the simulator to achieve impact in learning.	2-44
Figure 10 (a) Beach stones; marker used for stage location. This resource was used under the image licence notice: michael clarke https://commons.wikimedia.org/wiki/File:Renvyle_beach_stones_(3585122829).jpg . Renvyle beach stones (3585122829)“, https://creativecommons.org/licenses/by-sa/2.0/legalcode . (b) Tools tracking marker to be folded around a cylindrical surface.....	3-53
Figure 11 Direction of interactable vectors. (a) Digital model of the infrared camera, \tilde{N}_{camera_view} represents the unit vector normal to the frontal face of the camera. (b) 3D model attached to the cylindrical tracked object, \tilde{N}_{tool} represents the normal of the mount surface and D_{tool} a normalized vector in direction of the cylinder’s axis. (c) Frame of reference fixed on the pelvis. \tilde{N}_{frame} represents the unit vector normal to the surface on which the 3D models of the reflective spheres are supported.....	3-55
Figure 12 (a) Mixed reality capture, surgical scene after setting up and anchoring the holographic content in space. The transparent- yellow parallelepiped gives the user a visual feedback about the location and orientation of the system tracking volume. (b) Mixed reality capture of the simulator.	3-56
Figure 13 (a) Planes of the body. This resource was used under the image licence notice: Connexions CC BY 4.0 (http://creativecommons.org/licenses/by/4.0/), via Wikimedia Commons. (b) Mixed reality capture, reamer orientation feedback. A ray is casted from the reamers tip in the direction of the tool, this ray changes its colour according to how close the user is to perform the procedure in an orientation close to the safe zone defined in the literature.....	3-58
Figure 14 3D model and time captures of the animation of a posterior approach.....	3-59
Figure 15 Screenshots of the app running on the device.....	3-59

Figure 16 Scheme of synchronization stage for HoloLens and Android Tablet.	3-60
Figure 17 Screenshots taken from the HoloLens first person view. (a) User interface to connect HoloLens to the mobile server. (b) Calibration stage using calibration marker. (c) Hologram projected based on received and corrected coordinates in the HoloLens.	3-61
Figure 18 User's test results.	3-63
Figure 19 Popular choices of 3D elements for volumetric discretization.	4-67
Figure 20 Stresses tensor in an infinitesimal element of a deformable body.	4-70
Figure 21 Mechanical material properties. (a) Stress-Strain curve of elasto-plastic material. (b) Lateral vs longitudinal deformations.	4-71
Figure 22 (a) Shows the overlying mesh of a hand model with 863 nodes of which 624 nodes lie on the surface. Displacements are determined for the visible nodes and an additional 100 nonvisible nodes that correspond to surface elements in order to allow real time stress/strain visualization. Adapted from (Berkley <i>et al.</i> , 2004)	4-77
Figure 23 (a) XFEM representation of discontinuities without retopologizing the main mesh. (b) Post-processed visualization of the mesh. Adapted from (Schoch <i>et al.</i> , 2013)	4-78
Figure 24 Enrichments on the central node. (Left) The node is only enriched by the first 2 functions. (right) All the elements that share a node are intersected by a single cut surfaces. However, the node is enriched 3 times as the cut creates 3 discontinuous regions in the elements connected by the node. Adapted from (Koschier <i>et al.</i> , 2017)	4-79
Figure 25 Nomenclature of chip formation model. Adapted from Piispanen (1948)	4-81
Figure 26 Bone tissue layer in the hip joint.	4-82
Figure 27 Scheme of the cortical bone structure. Figure adapted from Santiuste <i>et al.</i> (2014) and Hou <i>et al.</i> (2000)	4-85
Figure 28 Volume extracted from the geometrical mesh to create the tetrahedral refined mesh. In the picture the offset is intentional to show the fit between the surfaces.	5-88
Figure 29 Reamer sizes, frontal and back view of the cutting edges. Pictures provided by the Bournemouth University Orthopaedic Research Institute.	5-90
Figure 30 Acetabular Reamer. 3D model (a) and generated tetrahedral mesh (b). Orientation of the reamer in assembly during simulation (c).	5-91
Figure 31 Reamer design dimensions. Dimension in millimetres.	5-91
Figure 32 Stress-strain curve with progressive damage degradation. Adapted from [online] https://abaqus-docs.mit.edu/2017/English/SIMACAEMATRefMap/simamat-c-damageevolductile.htm	5-94
Figure 33 Loading and boundary conditions applied to the assembly	5-95
Figure 34 Protocol used to measure exerted axial force during reaming.	5-96
Figure 35 Samples of simulated scenarios with different inclination and antroversion angles.	5-97
Figure 36 Examples of some simulation under difference force directions, From left to right: Assembly view, results in bone tetrahedral mesh on frame 24 and 94.	5-98
Figure 37 Possible failure criteria.	5-99
Figure 38 Examples of results for reamer displacement.	5-99
Figure 39 PEEQ diffusion along the finite elements (tetrahedrons) during the 3 first frames	5-101
Figure 40 Types of graph structure data.	109

Figure 41 Nomenclature for the graph neural network. A node \mathbf{n}_i , is connected through edges \mathbf{en}_i with its neighbours \mathbf{m}_{ei}	110
Figure 42 Graph network block. Encode-Process decode in Sanchez-Gonzalez et al. (2018).Figure adapted from Battaglia et al. (2018).....	113
Figure 43 Flowchart of our 3D physically-based learned deformation renderer.....	118
Figure 44 Frequency distribution of stresses values.....	119
Figure 45 Feature vectors for each entity of our graph-shaped data.....	121
Figure 46 GN encoder-core-decoder model used.....	122
Figure 47 Comparison of optimization alternatives.....	124
Figure 48 Data exploration reamer displacement in some training examples vs number of removed tetrahedrons, the different series in the legend correspond to FEM results under different reamer orientations.....	125
Figure 49 Scatter plot for Predicted and Real updated α_i' values. Multilayerperceptron (MLP) vs the encoder-core-decoder GN model trained with our FEM data.....	128
Figure 50 Comparison of regression algorithms for prediction of reamer displacement. The estimations were evaluated from the test data.....	129
Figure 51 Results of resurfacing on curved surface. (a) Heatmap of ground Truth values on each tetrahedron (b) Heatmap representation of predicted values of accumulated stress on each tetrahedron on frame t	130
Figure 52 Render Pipeline.....	131
Figure 53 Geometrical mesh construction fitted to volumetric mesh. The assembly offset is intentional to show that perfect fit between the surfaces is possible.....	132
Figure 54 Acetabular cavity exposed in hip model AR headsets. For Augmented reality applications the background is set to black corresponding to the absence of projected light on top of the environment.....	133
Figure 55 A proposed layout of an ideal simulator for traditional THR.....	139

List of Tables

Table 1 Adapted from Stirling et al. (2014).....	2-39
Table 2 Assessment criteria with listed requirements for score assigning.....	2-42
Table 3 Assessment result per simulator. ☑ : The requirement is met. ☐ : The requirement is only met partially. -: The requirement was not met.....	2-43
Table 4 Johnson-Cook Plastic Hardening properties for cortical bone. Table adapted from (Baro and Deoghare, 2018).....	4-85
Table 5 Bone model meshing parameters	5-90
Table 6 Meshing properties for reamer 3D model.....	5-90
Table 7 Mechanical properties of the bone used to model its elastic behaviour.....	5-92
Table 8 Johnson-Cook Damage values.....	5-93
Table 9 Mechanical properties for AISI 420.....	5-95
Table 10 Accuracy (Acc) and Recall (Rcll) scores for selected removed elements in mesh after 1,30 and 60- step rollouts on the test dataset. Values in percentages.....	129

Acronyms

THR	Total Hip Replacement.
FEM	Finite Element Method.
VR / AR	Virtual Reality / Augmented Reality.
FPS	Frames per second.
MSE	Mean squared error.
GN	Graph Network.
PEEQ, ϵ_F	Plastic equivalent strain.
ML	Machine learning.
MLP	Multi-layer perceptron.
AI	Artificial Intelligence.
XR	Extended Reality.
GPU	Graphics Processing Unit.
VM	Von Mises Stress.
RNN	Recursive Neural Networks.
NN	Neural Networks.
CAE	Computer-Aided Engineering.

Abstract

Virtual Reality simulators have proven to be an excellent tool in the medical sector to help trainees mastering surgical abilities by providing them with unlimited training opportunities. Total Hip Replacement (THR) is a procedure that can benefit significantly from VR/AR training, given its non-reversible nature. From all the different steps required while performing a THR, doctors agree that a correct fitting of the acetabular component of the implant has the highest relevance to ensure successful outcomes. Acetabular reaming is the step during which the acetabulum is resurfaced and prepared to receive the acetabular implant. The success of this step is directly related to the success of fitting the acetabular component. Therefore, this thesis will focus on developing digital tools that can be used to assist the training of acetabular reaming.

Devices such as navigation systems and robotic arms have proven to improve the final accuracy of the procedure. However, surgeons must learn to adapt their instrument movements to be recognised by infrared cameras. When surgeons are initially introduced to these systems, surgical times can be extended up to 20 minutes, maximising surgical risks. Training opportunities are sparse, given the high investment required to purchase these devices. As a cheaper alternative, we developed an Augmented Reality (AR) alternative for training on the calibration of imageless navigation systems (INS). At the time, there were no alternative simulators that using head-mounted displays to train users into the steps to calibrate such systems. Our simulator replicates the presence of an infrared camera and its interaction with the reflecting markers located on the surgical tools. A group of 6 hip surgeons were invited to test the simulator. All of them expressed their satisfaction with the ease of use and attractiveness of the simulator as well as the similarity of interaction with the real procedure. The study confirmed that our simulator represents a cheaper and faster option to train multiple surgeons simultaneously in the use of Imageless Navigation Systems (INS) than learning exclusively on the surgical theatre.

Current reviews on simulators for orthopaedic surgical procedures lack objective metrics of assessment given a standard set of design requirements. Instead, most of them rely exclusively on the level of interaction and functionality provided. We propose a comparative assessment rubric based on three different evaluation criteria. Namely immersion, interaction fidelity, and applied learning theories. After our assessment, we found that none of the simulators available for THR provides an accurate interactive representation of resurfacing procedures such as acetabular reaming based on force inputs exerted by the user. This feature is indispensable for an orthopaedics simulator, given that hand-eye coordination skills are essential skills to be trained before performing non-reversible bone removal on real patients.

Based on the findings of our comparative assessment, we decided to develop a model to simulate the physically-based deformation expected during traditional acetabular reaming, given the user's interaction with a volumetric mesh. Current interactive deformation methods on high-resolution meshes are based on geometrical collision detection and do not consider the contribution of the materials' physical properties. By ignoring the effect of the material

mechanics and the force exerted by the user, they become inadequate for training on hand-eye coordination skills transferable to the surgical theatre. Volumetric meshes are preferred in surgical simulation to geometric ones, given that they are able to represent the internal evolution of deformable solids resulting from cutting and shearing operations. Existing numerical methods for representing linear and corotational FEM cuts can only maintain interactive framerates at a low resolution of the mesh. Therefore, we decided to train a machine-learning model to learn the continuum mechanic laws relevant to acetabular reaming and predict deformations at interactive framerates. To the best of our knowledge, no research has been done previously on training a machine learning model on non-elastic FEM data to achieve results at interactive framerates.

As training data, we used the results from XFEM simulations precomputed over 5000 frames for plastic deformations on tetrahedral meshes with 20406 elements each. We selected XFEM simulation as the physically-based deformation ground-truth given its accuracy and fast convergence to represent cuts, discontinuities and large strain rates. Our machine learning-based interactive model was trained following the Graph Neural Networks (GNN) blocks. GNNs were selected to learn on tetrahedral meshes as other supervised-learning architectures like the Multilayer perceptron (MLP), and Convolutional neural networks (CNN) are unable to learn the relationships between entities with an arbitrary number of neighbours. The learned simulator identifies the elements to be removed on each frame and describes the accumulated stress evolution in the whole machined piece. Using data generated from the results of XFEM allowed us to embed the effects of non-linearities in our interactive simulations without extra processing time. The trained model executed the prediction task using our tetrahedral mesh and unseen reamer orientations faster per frame than the time required to generate the training FEM dataset. Given an unseen orientation of the reamer, the trained GN model updates the value of accumulated stress on each of the 20406 tetrahedral elements that constitute our mesh during the prediction task. Once this value is updated, the tetrahedrons to be removed from the mesh are identified using a threshold condition. After using each single-frame output as input for the following prediction repeatedly for up to 60 iterations, our model can maintain an accuracy of up to 90.8% in identifying the status of each element given their value of accumulated stress. Finally, we demonstrate how the developed estimator can be easily connected to any game engine and included in developing a fully functional hip arthroplasty simulator.

Chapter 1

Introduction

This chapter describes the fundamentals of Total Hip Replacement, its relevance on improving human's life quality and how this procedure is performed. We describe how current technologies can be helpful to assist the training of orthopaedic surgeons efficiently and explain the motivation behind developing tools that can be used to assist such training. Finally, we present the research questions that will be addressed through the proposed objective and describe the summary of the content of chapters included in this thesis.

1.1 The hip joint.

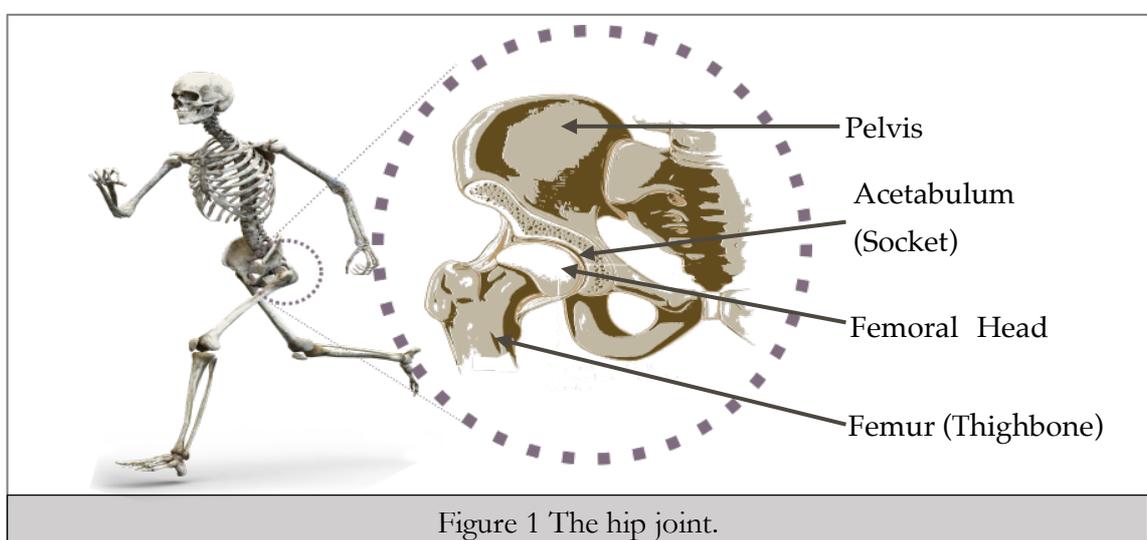


Figure 1 The hip joint.

The hip (Figure 1) is the joint where the thighbone (femur) meets the pelvis. The hip is known as a ball-type joint as it is formed by a semisphere (head) at the end of the femur that fits into a socket located on the pelvis (acetabulum). This assembly makes the hip very stable and provides a wide range of motion. The femoral head is coated with a layer of articular cartilage, which allows the femoral and pelvic bones to glide smoothly against each other. This cartilage absorbs shock forces on the joint. Therefore, it protects the surfaces of the femoral head and the acetabulum from excessive wear.

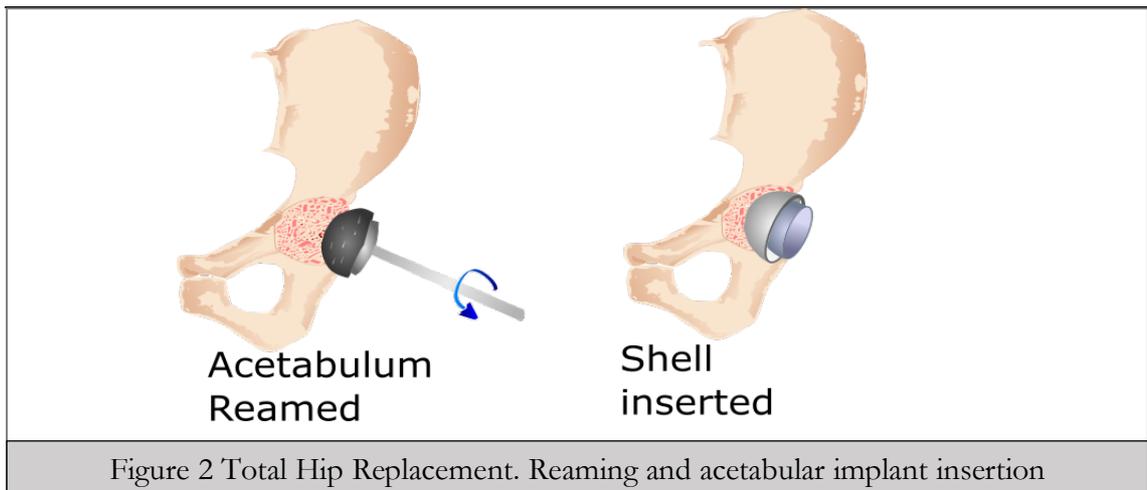
When the hip is healthy, much force is needed to injure it. However, sometimes overuse or falls can lead to severe hip injuries. In cases where a severe hip fracture occurs and the upper part of the thigh bone breaks, the entire joint must be replaced with an implant system. The most prevalent cause for requiring a hip replacement is severe Osteoarthritis, which results from excessive wear and tear on the joint caused by obesity, advancing age and disorders that cause unusual bone growth such as bone dysplasia. Hip dysplasia causes the hip socket to be too shallow to support the head of the femur. Most people diagnosed with hip dysplasia need replacement surgery to build better support for the hip as the wrong shape of the joint will cause the articular cartilage to wear out faster than one that has a more normal shape. Regardless of the cause for osteoarthritis, the cartilage covering the femoral head can degenerate so severely that the patients experience tremendous pain resulting from bone rubbing on bone, stiffness in the joint and difficulty walking (Wolford *et al.*, 2015).

1.2 Total Hip Replacement (THR) surgery.

1.2.1 Traditional surgery.

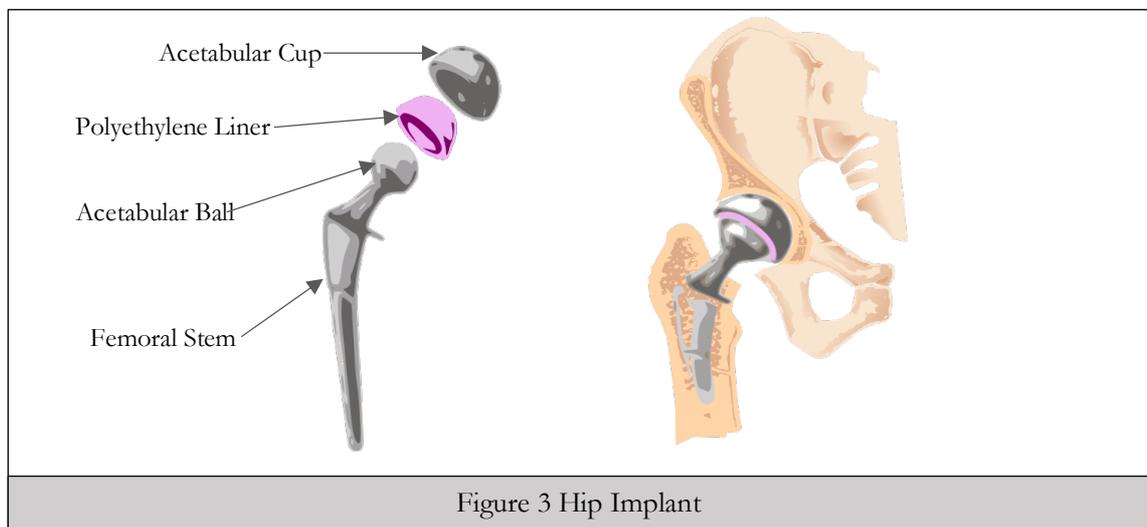
Each year over 96,000 Total Hip Replacement (THR) procedures are performed in England, Wales, Northern Ireland and the Isle of Man by both NHS and independent hospitals (National_Joint_Registry, 2018). Similarly, in the United States, more than 310,000 hip arthroplasties procedures are performed annually. Because of the ageing population, these numbers are expected to increase up to 205% annually (Wolford *et al.*, 2015). THR is considered to be one of the most successful and cost-effective surgeries in the orthopaedic field. However, up to 8% of the patients report discomfort related to several factors, including leg length discrepancies, dislocation¹, a high wear rate on the implant's surfaces

¹ Condition in which the head of the femur is forced out of its socket in the hip bone (pelvis).



that are gliding against each other, and other clinical complications (National_Joint_Register, 2018). These scenarios can lead to a more complicated and expensive procedure called *revision surgery*. In addition to the associated extra clinical costs and associated risk for being an invasive procedure, revision surgery decreases the patient's confidence in the procedure and increases their level of anxiety (Wolf *et al.*, 2012).

During surgery, surgeons restore the natural range of motion in the joint. Firstly, the surfaces on the acetabulum and femur must be prepared (resurfaced) to provide a perfect fit for the implant that will be inserted later (Figure 2). THR is a popular elective procedure as it relieves the pain caused by an arthritic condition (Bhaskar *et al.*, 2017). Figure 3 shows the main components of a complete hip implant system. The acetabular cup provides a convex gliding surface that allows the rotation of an acetabular ball with minimal friction. This sphere is attached to an elongated rigid part called the femoral stem responsible for transmitting the pelvis's mechanical load to the femur. During the entire procedure, the surgeons move the leg of the patient on several occasions. For example, to dislocate the hip joint, ease the



acetabular component's fitting stage and finally prepare the femoral canal and insert the femoral stem. Several authors such as Zinar & Schmalzried (2015) and Rojas et al. (2018) argue that the location and orientation of the acetabular component are of the highest importance when it comes to preventing the risks of dislocation; high wear on the polyethylene liner; and leg length discrepancies.

THR possesses qualities that differentiate it from other orthopaedic procedures such as total knee or shoulder replacement. Firstly, there is no other orthopaedic surgery in which the surgeon must move the patient's limbs as many times as in a total hip replacement. Secondly, surgeons have some limitations to manoeuvre their tools determined by the physical constraints of their selected approach². Finally, THR requires the surgeons to be “ambidextrous”, as applying the reaming and hammering force with either arm is equally relevant in the training path of a successful surgeon.

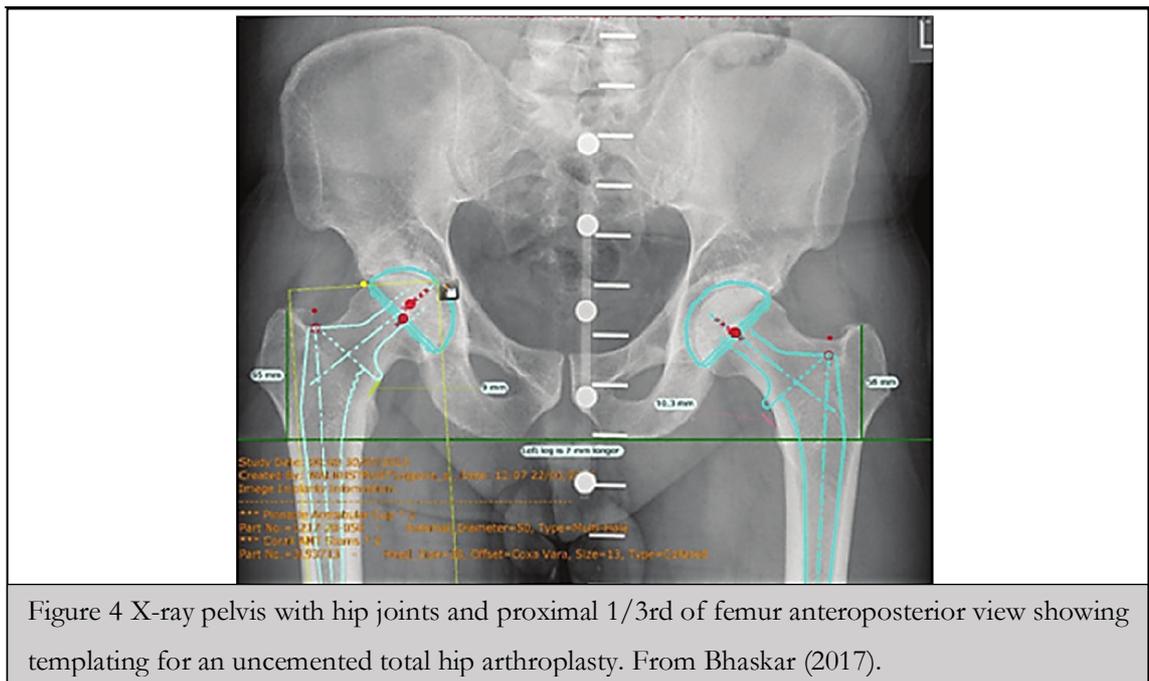
To implant an acetabular implant, surgeons must first prepare an adequate shape on the acetabulum. A uniform bleeding surface is required to ensure post-operative osseointegration³. Osseointegration is especially important when using systems that do not rely on surgical cement to achieve a stable fit between the implant and the bone. To do this, surgeons increase the size of the acetabular cavity using a hemispherical tool commonly known as a *reamer*. This step also referred to as reaming, progressively removes bone tissue until an adequate acetabular diameter and implantation depth is achieved. To ensure stable implantation of cementless systems, surgeons must achieve an interference fit between the pelvis and the cup. Such a fit is obtained by implanting a cup 1 or 2 mm bigger in diameter than the size of the last reamer head used.

Pre-operative surgical planning⁴ (Figure 4), helps surgeons to estimate the required implant size and type based on an X-ray image of the hip joint. Particularly the orientation of the cup edges relative to some anatomical landmarks on the pelvic bone provide a reference to the surgeons of the required reaming orientation.

² The technique chosen to enter the body to get to the site that needs to be operated on. For THR, the most common are the posterior approach, direct lateral approach, and direct anterior approach.

³ “A direct structural and functional connection between ordered, living bone and the surface of a load-carrying implant which is critical for implant stability”. (Parithimarkalaignan & Padmanabhan 2013)

⁴ Also known as *templating*.



While the surgeons are reaming, they need to use visual feedback to be aware of the progress of bone removal. The pelvic bone has three essential features the surgeon can observe. One is the disappearance of the acetabular labrum, which covers the acetabular *true floor*. The second one is the appearance of a bleeding bone on the whole surface of the cavity. Finally, the depth of the penetration from the last reamer diameter in use and how the reamer edges protrude the acetabulum bone compared with the initial template gives the surgeon an estimation on the future alignment of the implant the amount of bone removed.

1.2.2 Robotic Assisted surgery and Imageless Navigation Systems.

The accuracy of the reaming stage directly affects the stability of the implanted acetabular cup as it defines the geometry that will outline the fitting between the implant and the bone. Under-reaming causes difficulties during the cup insertion stage as well as undesired additional mechanical stresses in both the pelvic bone and the acetabular cup. On the other hand, over-reaming leads to increased micro-motion or traumatic disruption in the bone with loss of cement-less fixation. Unfortunately, it is hard to eliminate the error of the reamed surface due to different factors involved, such as the tool's cutting geometry, surgeon's technique and pelvic bone elasticity (Alexander, J. W., E. Kamaric, 1999). Devices such as navigation systems and robotic arms have proven to improve the final accuracy of the procedure (Schnurr *et al.*, 2009; Redmond *et al.*, 2016; Snijders *et al.*, 2017). Imageless Navigation Systems guide the surgeon's movements and give live feedback on the amount of removed bone tissue in real-time with an accuracy below 1mm (NDI Medical, 2017).



Figure 5 Stryker's MAKO robot. Assisting robotic arm for orthopaedic arthroplasty procedures on hip and knee joints. Photography provided by ORI BU

Robotic-assisted surgery removes human error and provides a more accurate and replicable cup orientation than manual implantation, thanks to its navigation system and powerful haptic units. After conducting a systematic review, Tilly (2016) defends that *“the use of robotic-assisted surgery can result in a greater range of motion in the hip, reduced soft-tissue damage, decreased bone-to-bone impingement, and enhanced stability”*. Even with the mentioned advantages, the learning curve involved in this new way of performing the surgery may delay its widespread use (Newman, 2014). The spatial calibration step on each surgery is essential to ensure the reliability of the navigation system. This step allows the robot to compute its relative positions to the patient and ensure that the implants' planned orientation and location are correct. Therefore, both trainees and already trained surgeons should learn to work without interfering with the tracking performed by the infrared cameras on both the patient and the robot.

Currently, surgeons start integrating the robot into their theatres supported by a team of technicians who guide them through the appropriate steps to perform during their firsts robotic-assisted surgeries. Consequently, the practical experience can only be obtained from direct interaction with the robot, which temporarily increases surgical times. An adequate training phase for this task is required since an incorrect spatial data acquisition would result in outcomes possibly less desirable than those typically obtained from traditional THR.

It is worth noting that the cost of acquiring a robotic assisting arm can be so high that it makes them unaffordable for many healthcare centres around the world.

1.3 Surgical training.

This research addresses the 14th item from the specialist training curriculum for trauma and orthopaedics training in the United Kingdom. This item states: “*Demonstrates familiarity and understanding of an acetabular preparation including trimming medially and at rim*” (Mitchell, 2017). Unfortunately, acetabulum reaming is not a reversible process as it involves material removal. Training on getting a stable result during the reaming stage and developing strong surgical skills requires repetition and rehearsal.

For many years surgical training has been based on an apprenticeship model highly dependent on excellent trainer–trainee communication. However, gaining surgical experience became harder for starting trainees since the implementation of the European Working Time Directive (EWTD), which came into force in 2003 to ensure well-rested workers (European Parliament, 2003). The EWTD limits the number of working hours to 48 per week, and it demands employers to provide 11 hours of rest to their employees on any 24 hours and a minimum 20-minute rest break where the working day is longer than 6

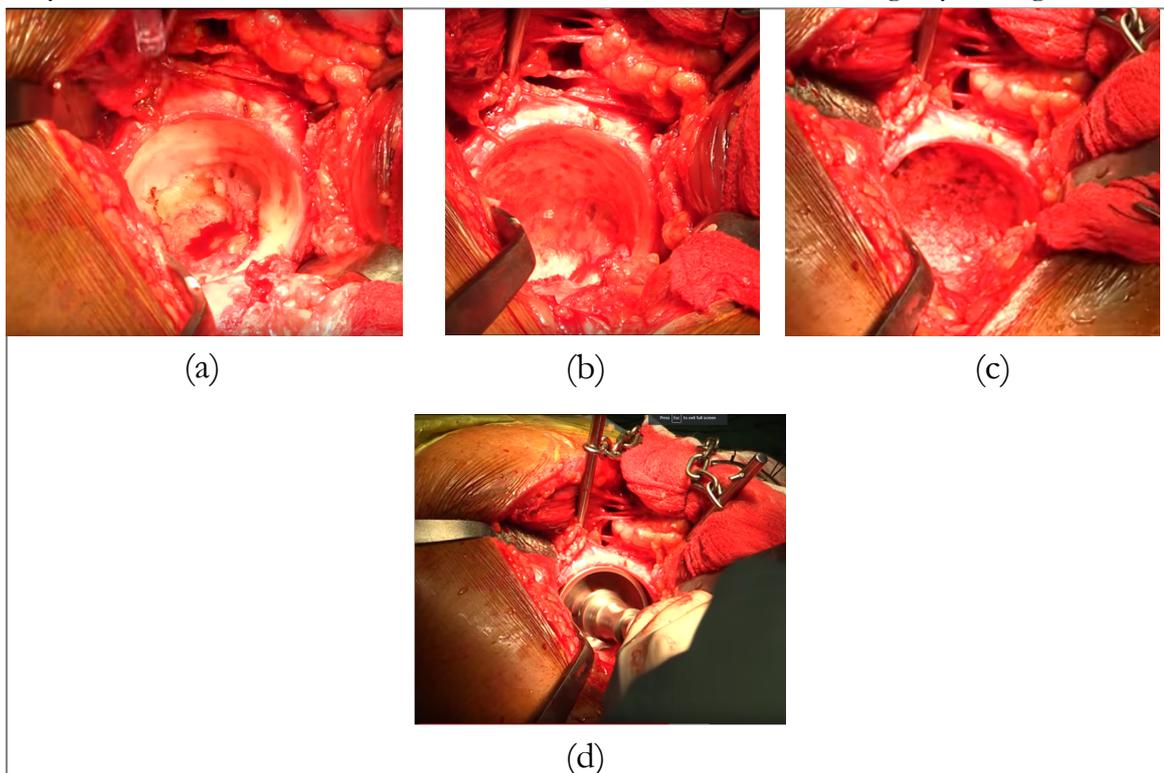


Figure 6 Visual references of the progress during the reaming stage. a) Unreamed acetabulum. b) Initial ream, the labrum is completely removed and the true floor is exposed. c) Appearance of bleeding bone on the reamed surface) Reamer inside the acetabulum. Screenshots from TOTAL HIP REPLACEMENT - UNCEMENTED. [onlinevideo] (Orthopaedics, 2016) Orthopaedics, T., 2016. Available at: <https://www.youtube.com/watch?v=sMBx3C8hhJg&t=217s> [Accessed September 11, 2018].

hours. A complete Total Hip Replacement surgery, including patient preparation, can take around 2 hours. Since trainees take longer to complete a whole procedure than experienced consultants, the total number of operations that residents have the opportunity to perform by themselves have dropped significantly. As a result of this, reducing the number of hours that medical staff is allowed to work has encouraged simulation-based training to progressively position itself as a tool of high importance in surgical training (Kotsis and Chung, 2013). Simulated training avoids patients risk and allows the trainees to get more profound knowledge of in-theatre skills such as spatial understanding and how to react to anatomy variation.

In their report "To err is human," the institute of medicine of the United States of America has suggested using *simulation whenever possible* (Kohn *et al.*, 2000). In fact, the impact of computer simulation in orthopaedic training and related literature has been growing exponentially over the last 12 years (Vaughan *et al.*, 2016). Virtual reality simulators represent an excellent tool for grasping abilities through repetition compared to other training methods such as single-use cadaver and or plastic bone (Seymour *et al.*, 2002; Aggarwal *et al.*, 2007). Also, Virtual and Augmented Reality (AR & VR) can emulate the adverse consequences from common mistakes that trainees make during learning (Gardner *et al.*, 2015). This allows the trainees to err in a controlled and safe digital environment that allows for better structuring of their knowledge.

1.4 Virtual and Augmented Reality simulators.

Virtual Reality aims to replace reality with an entirely computer-generated three-dimensional environment that can interact in a seemingly natural way. On the other hand, Augmented Reality overlays the computer-generated content into the user's visual field. This is, Augmented reality complements the natural user's vision instead of trying to replace it. Together, Virtual and Augmented reality form the notion of extended reality (XR). The realism of XR experiences perceived by the users is commonly assessed by its level of *immersion*. A proposed assessment of immersion for VR and AR simulators for orthopaedic training is presented in section 2.2.1. It is worth noting that XR experiences are not limited

to run only on stereoscopic⁵ devices. Under the right conditions, computer monitors and hand-held devices can also host these experiences and provide an adequate level of immersion.

Flat computer monitors or hand-held devices do not provide the visual immersion required to simulate open surgery procedures. During the real THR surgery, surgeons move their head around the incision to utilize their visibility of the exposed tissue. The incision made (Figure 6) must be big enough to have good tissue exposure and provide the surgeons with an area of adequate size to manipulate their tools comfortably. However, incisions too big require longer post-operative recovery times. The continuous change from a visual perspective is an essential practice in surgery that can only be replicated in simulation by using head-mounted displays. These displays can update their rendered image based on the tracked 6 DOF⁶ of the user's head and are available for both VR or AR experiences. For that reason, they improve the immersion during the simulation. In particular, AR headsets are recommended as a safer option as they allow for visualization of the real environment while procedures involving high impact and reaming forces are being trained.

Additionally, a simulator can provide haptic immersion. This term refers to the ability of the system to provide tactile and force feedback based on the user's interaction. A limitation of currently available open surgery simulators is implementing commercial haptic units to provide force feedback. The technical specifications of these type of devices usually report a range of motion smaller than the one required in THR surgery. The handles, usually in the shape of a haptic pen, are too small and must be attached to plastic models matching the tools' shape to provide tactile realism. Also, the maximum force exerted as resistance to movement has a magnitude below 40N (3D Systems, 2017). These hardware limitations make the human-machine interaction unnatural when trying to mimic joint replacement procedures such as THR.

Instead of using mechanical haptic units, physical objects can be used to deliver tactile feedback to the user. This approach is known as *passive haptics*. Passive haptics helps overcome the technical limitations of haptic units and provide truthful experience immersion

⁵ Stereoscopic rendering creates a sense of 3D depth by rendering the scene from 2 different view points with slightly different angles, just like human eyes. Common examples are gaming head-mounted headsets which have 2 displays, one for each eye.

⁶ 6 Degrees of Freedom. Namely, 3D position coordinates and 3D orientation angles.

once aligned in space with the digital representations of each object (Fotouhi, 2018). In a simulator for THR, such as the one proposed by Fotouhi (2018), a passive haptic model⁷ can be mounted on a supporting structure capable of resisting the forces applied to replicate the reaming and impacting stages of the surgery safely. The spatial alignment ensures that the digital content rendered by the XR headset is perceived to have the same position, orientation and scale as the passive haptic model. A haptic arm does not constrain the user's interaction. Instead, the movements and exerted force are monitored in real-time by adequate sensors such as a triaxial force sensor and a 3D position tracker.

Simulating the calibration step for robotic-assisted surgery does not require simulating complex tissue deformations as its central aim is to guide the surgeon in acquiring the anatomical landmarks needed for spatial calibration. Therefore, such a simulator can be hosted entirely on VR/AR headset. However, the same approach is not valid if we want to provide training opportunities for traditional THR. As mentioned in previous sections, acetabular reaming is an essential step to determine the success of THR. Reaming is not a reversible process, and consequently, surgeons must be skilful enough to handle the tools inside the patient's hip to avoid any possible complication. To be able to provide unlimited attempts inside the simulation, the “reamer” used during simulation should not induce any sign of wear on the surface of the passive haptics model. Therefore, the simulator must take advantage of the spatial alignment to display the progress of the reaming operation on the surface of the CGI model of the pelvis. To ensure the acquisition of relevant reaming skills, the estimated amount of bone removed during simulation must be similar to the one expected during real surgery.

Standard techniques used to simulate resurfacing operations, such as reaming, can be performed either in the geometrical or the volumetric domain. These techniques usually model the resurfaced mesh by generating new faces, seeding vertices (Bruyns *et al.*, 2002; Niu and Leu, 2007), voxelized boolean operations (Jang *et al.*, 2000; Wu *et al.*, 2009) or planar volumetric primitive subdivisions (Turini *et al.*, 2006; Sifakis *et al.*, 2007). Consequently, the evaluation of material removed is focused mainly on the collision detection of bounding volumes, which can be computationally expensive when requiring complex tool's geometries. For both the volumetric and geometric domain, mesh resurfacing operations are commonly

⁷ For example a 3D printed model or a synthetic bone commonly used for single-use training (Hetaimish, 2016).

dependant on the tool's swept path and collision detection. Therefore, these methods neglect the effects on the resulting surface of any changes in the cutting force magnitude and direction.

On the other hand, a physically-based⁸ simulation would compute the displacement of the cutting edges as a result of the fractured elements described by the primitive's stresses evolution given an applied cutting force. Numerical computer-aided simulation of machining operations on brittle⁹ materials (e.g. turning, milling, drilling and reaming) has been used to help cutting tool designers understand the dynamic behaviour in the interface between the tools and the resurfaced material (Ng and Aspinwall, 2002; Schermann *et al.*, 2006; Liu *et al.*, 2013). Depending on the model used, these simulations can be highly accurate. Numerical methods such as the finite element method (FEM) are adequate tools to calculate accurate plastic deformations. The FEM subdivides the deformable objects into smaller areas of interest called *finite elements* to approximate the solution of a system of differential equations over the volume domain. The finner the discretization of the domain, the more accurate and stable is the solution, but also the longer it takes for the solver to yield a solution of the system. Therefore, when a real-time analysis is desired, it is necessary to minimize the number of elements and constrain the model to a linear behaviour. Such simplifications usually compromise accuracy and make them non-adequate for medical applications (Gillies and Bourmpos, 2003). We propose to use machine learning on FEM data computed offline to train an accurate estimator of the amount of bone removed based on an interactive input faster than traditional numerical simulations.

1.5 Research Questions

This research aims to answer the following questions:

Q1 Which learning theories must be considered during the design process of a simulator for training on THR skills?

Q2 How to provide an alternative training opportunity outside of the surgical theatre to allows the user to acquire the relevant skills for the calibration step of robotic

⁸ The numerical solution of a problem that includes the physics models.

⁹ According materials science a brittle material is a material that breaks with small elastic deformation and almost no plastic deformation (Beer, 2011). The bone behaves as a brittle material.

assistants and imageless navigation systems?

Q3 How to obtain a deformation module capable of estimating the amount of removed bone interactively by taking the force applied on the reamer and the status of the volumetric mesh as its inputs?

1.6 Aims and objectives

Training strategies from gamification theories are only helpful for VR simulation-based training when these strategies add value to the construction of knowledge and enable a smooth transition of acquired skills to real-life applications. Therefore, this research aims to find the relevant learning theories around simulation-based training in the literature and identify how these should be included in a surgical simulator. An understanding of the learning directives for experience-based learning will give us a clearer view of the technical requirements and features that must be included in our XR simulator for THR.

To minimize human error in surgery, healthcare centres invest in Robotic Assistants and imageless navigation systems. Ideally, the learning curve for these devices should happen without assigning one of them exclusively for training. Consequently, we propose a prototype to train surgeons on the basic skills needed during the calibration stage of an imageless navigation system and assistive robotic arms. With this simulator, users will be able to learn to manipulate their tools without interfering with the continuous infrared tracking of both the patient and the robot. Furthermore, this prototype will also provide us with feedback from the orthopaedic surgeons about both comfort and functionality of the selected hardware and the interaction planned.

For cases where a system that ensures the correct location of the acetabular component is not affordable, it is required to assist the training of reaming skills. Under-reaming can cause difficulties during the cup insertion stage as well as undesired additional mechanical stresses in both the pelvic bone and the acetabular cup. On the other hand, over-reaming can lead to increased micro-motion or traumatic disruption in the bone with loss of cementless fixation. Another goal of the proposed research is to create a simulation able to execute efficiently and estimate the reduction on the acetabulum volume during simulation of the unassisted reaming procedure. This research will present a set of techniques that together will build a learned simulator of the changes triggered on a volumetric mesh based on the orientation and position of the reamer, as well as its exerted force. In other words, the model

should be able to respond to forces applied in different orientations on the acetabulum and calculate the amount of material removed. Our learned simulator module will be able to estimate the outcomes of the reaming procedure on both the left and right hip. This will be easy to include in the development pipeline of any other VR/AR simulator.

The objectives of this thesis are:

- To identify the learning theories that outline the design process of a simulator for training on THR skills and, therefore, ensure the learning impact of a simulator.
- To identify the current state of the art of surgical simulators for orthopaedic training and assess their available features against the requirements raised by the identified learning theories for simulation-based learning.
- To provide an alternative training opportunity outside of the surgical theatre that allows the user to acquire the relevant skills for the calibration step of robotic assistants and imageless navigation systems
- To replicate the acetabular reaming operation in a computer-generated environment.
 - To develop a FEM simulation of the reaming procedure under different reaming conditions.
 - To train a Machine Learning (ML) model capable of estimating the removed material on a refined area of a volumetric mesh faster than traditional FEM methods.
 - To train a model able to output the centre of the reamer, orientation and position according to the current bone volume distribution and force applied on a specific frame.

1.7 Scope and Limitations

The following considerations will be taken throughout this research.

- There are several paths a surgeon can choose to get a suitable view of the hip joint. In the surgical field, these are known as *exposures*. They differ from muscles chosen to incise to even the type of tools required. This project will consider only the **posterior approach** as it is the most popular one in the area where this research was conducted.
- Our finite element model will be built based on mechanical properties and material models used in the literature to model bone cutting processes that have been

validated with experimental data.

1.8 Thesis layout and remarks

The next chapter (Chapter 2) investigates the learning theories and gamification strategies identified as relevant for the design process of simulation-based training experiences. We will also propose a list of ideal requirements for an XR simulator in orthopaedics that the industry should be motivated towards and use these requirements to evaluate currently available simulators. The comparative assessments presented in Chapter 2 will help the reader to understand the context of this research based on the gap found in the available training opportunities for THR.

Devices such as navigation systems and robotic arms have proven to improve the final accuracy of the procedure. However, surgeons must learn to adapt their instrument movements to be recognised by infrared cameras. Chapter 3 describes the development and users study case from our AR simulator to help surgeons acquire infrared-based navigation skills before operating with them in reality (Aguilera-Canon *et al.*, 2019). Additionally, the development of an alternative communication protocol between an AR headset and an external tracking device is presented. Such protocol is helpful for headset devices whose computing capacity must be reserved for graphical computations or to enable advanced accurate position tracking systems.

Given that the cost of acquiring a robotic assisting arm can be so high that it makes it unaffordable for many healthcare centres around the world, we must also provide training opportunities to ensure the correct positioning of the acetabular component. Experimental data would be the ideal alternative of data to train a machine learning predictor on the resurfacing effects of the reaming procedure on the bone. However, this would require us to perform and measure the results of reaming experiments over many bone pieces. It is impossible to find several bone samples with identical mechanical behaviour when being reamed. This is because the bone density and its microstructure are different in every used instance. Therefore, this would introduce undesired variance to our training data.

Consequently, we must consider numerical simulations as a source of data to represent the evolution of the machined bone. Chapter 4 describes the fundamentals to define a FEM simulation of reaming procedures on bone. We will introduce the concept of material constitutive equations, which describe the relationships that govern the material's strain-

stress curve for all points in a continuum solid. We also provide an overview of current FEM models used for interactive VR surgical simulators. Finally, we present constitutive equations models found in the literature to describe the mechanical behaviour of bone.

Chapter 5 explains the methodology, material models, boundary conditions, loading conditions and set interactions used to generate the training data from a finite element simulation. The rationale behind some problem simplifications required due to hardware constraints is also explained in this chapter.

Chapter 6 presents a review of data-driven techniques that incorporate the accuracy of numerical simulations at faster computation times. It is usually not possible to represent FEM meshes in the Euclidean domain because of the variable size and orientation of its elements. Furthermore, in a volumetric FEM tetrahedral mesh, the number of neighbours connected to each element forming the volumetric mesh depends on the location of the element in the mesh. Therefore, a graph is the most appropriate data structure to keep the geometric topology and element-to-element relationships of a tetrahedral FEM mesh. This chapter explains the fundamentals of graph-structured data and learning in the graph domain.

Chapter 7 discusses how the results from the FEM simulation described in Chapter 5 were translated into the graph domain from their original tetrahedral mesh projection. We will discuss the rationale behind the features selection and incorporate the self-attention label that will allow us to *neglect* the presence of removed elements while keeping the graph size static. We will explain the learning strategy used to train the ML model to approximate the volumetric change of the reamed body through time. Finally, we demonstrate how the developed estimator can be easily connected to any game engine and included in developing a fully functional hip arthroplasty simulator.

Finally, Chapter 8 concludes this document with an analysis of the contributions from this thesis and proposals of future research directions that can improve the generalization accuracy of our trained machine learning model.

Chapter 2

Simulation-based training and learning theories

This chapter contextualises our research objectives given the current trends for adult and experience-based learning. Current reviews on simulators for orthopaedic surgical procedures lack objective metrics of assessment given a standard set of design requirements. Instead, most of them rely exclusively on the level of interaction and the features and functionality provided. Therefore, we constructed a list of the design requirements relevant to orthopaedic training simulation. Finally, we present a comparative assessment tool to review the current VR simulators for training on orthopaedic surgical skills.

2.1 Serious games and surgical education

2.1.1 *A review of the state of surgical education*

Since 1890, the training of resident surgeons has been ruled by the Halsteadian method commonly summarised with the saying “See One, Do One, Teach One”. Initially, the model was developed with the novel purpose of training both surgeons and high qualified mentors. However, several authors agreed that the uncertainty over whether trainees can perform a safe procedure at their first attempt represents a risk to the patient's wellbeing. (Vozenilek *et al.*, 2004; Rohrich, 2006) . Mason & Strike (2003) surveyed three different hospitals in the United Kingdom, asking doctors and nurses residents about their apprenticeship experience in eight practical procedures. They found out that “*Some 42% per cent of doctors felt inadequately trained to perform a practical procedure safely when first performing it alone*”. Even if this situation changed after 2005 after the inclusion of the reform Modernising Medical Careers (MMC)

in the UK, some other developing countries where such foundation training system does not exist are still implementing the Halsteadian model as their primary methodology.

The current UK training system is a more patient-protective one than the Halsteadian model. However, trainees continue to live a “pupil” role in it. After the culmination of the medical school, the apprentices could start observing or, in some cases, even assisting surgeries until their Foundation training stage. Then, after completing the Core Training (CT1-CT2) stage, they are expected to have been gradually introduced to the existing hip approaches and perform a full hemiarthroplasty. Finally, all candidates take part in a five-year orthopaedic surgical curriculum, during which they learn from an experienced consultant. This consultant has complete autonomy to decide how many procedures the trainee must observe before performing a whole surgical procedure or just a part of it (Syed *et al.*, 2009). The progress of each trainee towards their qualification goals is limited not only by the consultant’s criteria but also by the training opportunities available. For the particular case of total hip replacement, the NHS demands 40 successful surgeries performed at the end of the ST8 stage (McAlinden and Dougherty, 2014).

2.1.2 Adult learning theories and VR/AR simulation

A trainer-centred education, such as the Halsteadian model, is not entirely aligned with the current adults learning theories because these theories are more learner-based than problem-centred. Following the continuous development of technologies in both the medical and academic field, authors like Alimisis & Zoulias (2013) suggest that a new learner-centred approach that supports self-directed training is required. In these new learning scenarios, the goals should be clearly stated at the beginning of the simulation so that the evaluation can change from time-based to criterion-based. This section describes the learning theories found in the literature that justify the claims of Alimisis and Zoulias, and it will explain how they are relevant in the content planning of virtual simulation-based training experience.

In 1984, Knowles originally introduced the discipline that studies how adults learn, known as andragogy (Knowles, 1984) . His essay proposes that adults prefer to learn in a self-directed way as they are strongly guided by the motivation to acquire a new concept and its immediate application into their everyday problems. The previous knowledge available in the learner is also highly relevant in andragogy as adults need to create analogous thoughts around the new concepts and relate them to existing ones for the knowledge to stay on a

long-term basis (Clapper, 2010). These mental connections with previous concepts are relevant during the design of simulation-based training for orthopaedic skills since users should possess some basic concepts on orthopaedic surgery to build mental relationships between the theory and the practice.

Simulation-based learning satisfies Knowles' theories as it is learner-centred, facilitates the imitation of a realistic problem experience and allows the users to interact in a controlled virtual environment. Learning through past experiences can be achieved during serious games. In them, learning can be divided into levels that follow the order of steps that determine how the knowledge should be constructed. The construction of knowledge can also be enhanced by introducing the defined goals at the beginning of the interaction (Baby *et al.*, 2016). Serious games are considered inside the scope of serious games as long as the entertainment provided in the virtual environment keeps the learners' motivation towards practice.

Some authors believed that Knowles's approach should be complemented with a subjective component as not all adults learn in the same way. A completely learner-centred approach must be aware of the individuals and how their cultural context and social environment can influence their way to assimilate and internalise any new knowledge (Clapper, 2010; Merriam *et al.*, 2012). For example, time constraints can have an impact on an adult's learning motivation. Therefore, when planning training programs that deliver simulation, it is relevant to make the training accessible. For example, health professionals prefer a meaningful and short introduction or pre-training sessions when interacting with VR/AR devices (Clapper, 2010). Hence it is essential to design interaction mechanisms as user-friendly and intuitive as possible to keep the simulation attractive for all types of users regardless of their previous technical or computer-gaming experience.

Once a trainee has achieved sufficient proficiency in a task, improving until reaching an expert-level performance can involve different conditions on the learners and the way the knowledge is delivered. In an essay on mastering an ability through practice, Ericsson (2012) suggests that learners must be motivated to improve beyond autonomous performance while practising. Therefore, the aim of the training must change from being able to solve basic everyday problems to reach a higher level in every attempt. Unlike common beliefs, Ericsson argues that age and amount of experience do not automatically translate into the exceptional domain of a task. To achieve expertise, the individuals must address the training with clear and defined improvement goals; they also need to be provided with thorough feedback and

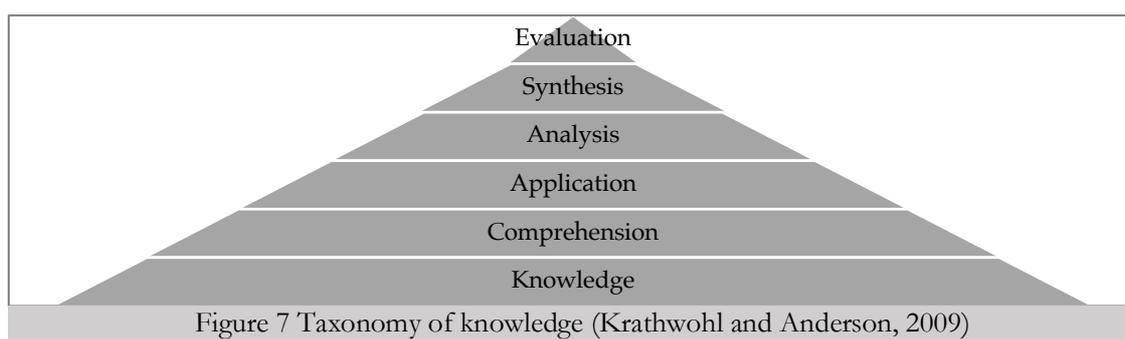
unlimited repetition opportunities. In orthopaedic surgery, achieving expertise beyond the common fundamental problems can be challenging since unique cases are not continuously available for interaction. VR/AR simulation provides an opportunity for deliberate practice, as unusual clinical cases can be replicated and allow the surgeons to improve their expertise with unlimited goal-based simulations. Deliberate practice and rare clinical cases are outside of this project's scope as they require higher orthopaedic expertise by the researcher and patient data collection.

2.1.3 *Constructivism and experience-based learning.*

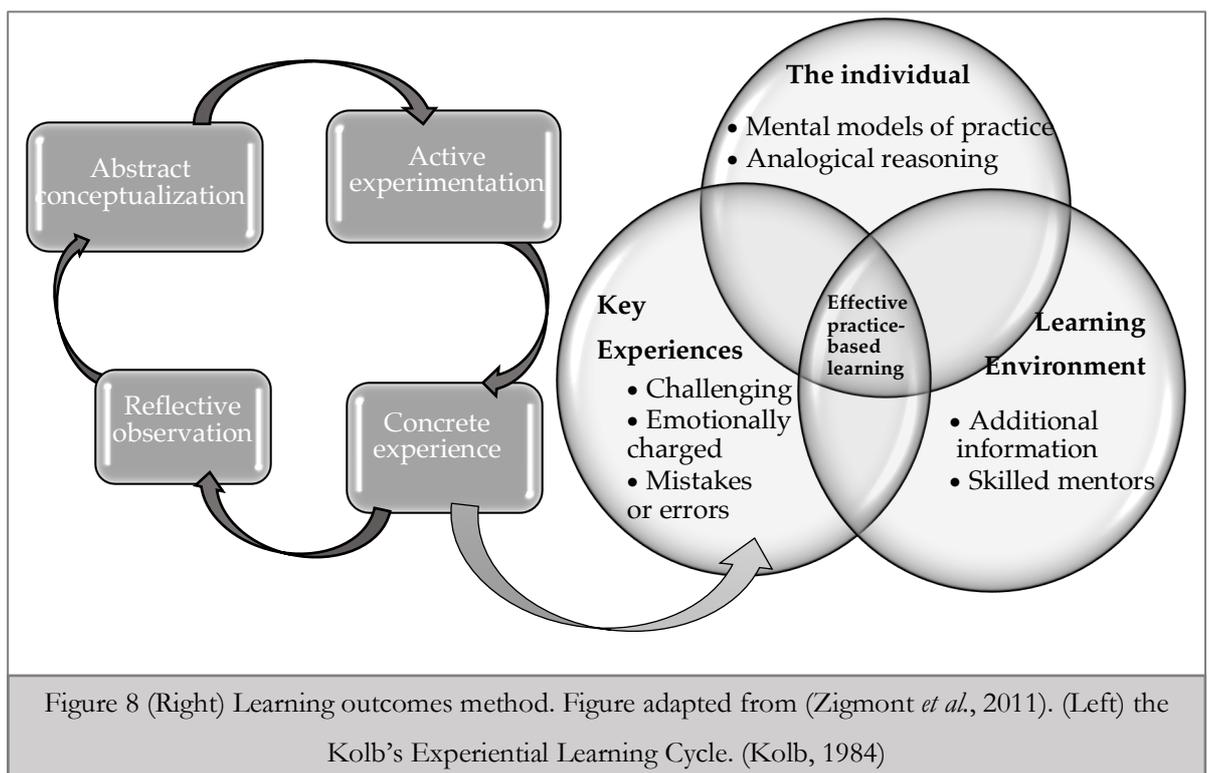
Constructivism is supported on the revised taxonomy of knowledge from Bloom (1965) (Krathwohl and Anderson, 2009) (Figure 7). The pyramid indicates that knowledge and comprehension are just the primary states of a complete learning path. The practice must be used to allow the users an opportunity to scale until the synthesis stage, rather than only deliver knowledge foundations to the pyramid (Lainema, 2008; Zigmont *et al.*, 2011). In order to do so, the hands-on learning experiences should be tools for the construction of the knowledge around applied situations, provide high-level feedback and ensure that the content is coherent with the relevant theoretical background (Wang, 2011).

Healthcare education and serious gaming are well aligned with constructivist theories as well since they are experience-centred. The knowledge is expected to be constructed progressively during a problem-solving practice and not only delivered as in lecture-based learning (Alimisis and Zoulias, 2013). AR/VR scenarios provide an excellent tool for constructivist learning as long as the content of the simulation is aligned with an approved training curriculum and the tasks are perceived as authentic (Huang *et al.*, 2010).

Another relevant author in experience-based learning is Kolb (1984). He proposed the required features for meaningful learning activities. The Kolb's cycle of experimental learning (Figure 8) points out that the experience alone has not the entire influence on the



learning process. During this step, the individual can only identify the substantial knowledge gap and the type of mental model required to complete the exercise. Afterwards, in the observation stage, meaningful feedback occurs, and debate is carried around the results obtained. Several authors agree that this reflective observation is the most valuable stage of Kolb's learning cycle (Alimisis and Zoulias, 2013; Damewood, 2016). Later, during the stage of abstract conceptualisation, the learner can build mental connections to previous experiences and interiorise the new skillset. Finally, the step of active experimentation refers to a real-life opportunity to test the new concepts acquired. This situation is more meaningful when the occasion arises by itself as it obliges the learners to identify the problem and test their mental models. However, if the time until active experimentation is too long, the knowledge can get lost, and some of the stages of the cycle would have to be re-done.



A motivating assessment is the central pillar of a successful reflective observation, and it is preferred over a failure/success evaluation (Clapper, 2010). Adult learners prefer their feedback to be immediate, but it is also recommended to deliver a meaningful overall assessment at the end of the practice rather than specific step assessments during the interaction (Alimisis and Zoulias, 2013). Throughout the simulator development, results can be displayed to the user immediately after the interactive task has finished, allowing further

data analysis with the supervision of the experts. Moreover, artificially intelligent agents could be implemented to support the users during their analysis and enrich the value of the simulator without requiring additional human assets.

Similar to adult learning theories, one frequent critique of constructivism is done by the objectivist tradition. Objectivists authors defend that the experience value and its further reflection is subjected to different factors around the learners themselves and their environment (Lainema, 2008). Zigmont et al. (2011) integrated Knowles's theories with a closer focus on the individual and joined them with Kolb's learning cycle on experience-based training (See Figure 8). For Zigmont et al. (2011), the individuals will learn successfully when their motivation comes from a known problem, and there is a clear understanding of the benefits and application of the new skill in their professional reality. Students must also involve themselves without bias in the experience and be supported towards change and growth by their environment (Wang, 2011).

Several best practices regarding the multimedia gamification procedure were also found in the literature. Some authors agree that an immersive high fidelity experience, including replicating the expected sounds and visual surroundings, is important to allow learners to construct real-world analogies faster (Wang, 2011). On the contrary, others suggest that inducing several stimuli simultaneously (i.e. text instructions, audio, animations) can jeopardise the learners' ability to focus (Baby *et al.*, 2016). From the interaction perspective, Alimisis and Zoulias (2013) present arguments to emphasise the need to include extended decision trees inside the simulation content to allow users to "own" the simulated problem and test their decision making. However, we know that providing unlimited decision choices is not always feasible from the programming perspective as it involves extensive memory and developing time resources. Lastly, Salas et al. (2009) suggest that a good practice to carry out a performance measurement is by calculating it automatically and free of human subjectivities. It is also recommended to compute the relevant assessment values based on the measurement from different sensors and observation angles. This broader data spectrum allows the simulator's developers to provide multiple levels of analysis of the results of the interactions. Baby et al. (2016) complement this idea by mentioning that the feedback provided in serious games should be able to identify a user as expert or novice and track them as they progress in their practice.

2.2 Simulation in orthopaedic training

Particularity in the field of orthopaedics, the most popular simulators are cadaveric training, polymer-based stiff foam bone model and virtual reality. Table 1 describes a classification presented by Stirling et al. (2014) about the types of simulators used in orthopaedic training and their main characteristics.

Simulation model	Advantages	Disadvantages
Cadaveric simulation	<ul style="list-style-type: none"> • High Fidelity • Shown to develop transferable operative skills • Allows understanding of relevant clinical anatomy and surgical approaches 	<ul style="list-style-type: none"> • Expensive • Not easily accessible with specialist storage demands • Time-consuming preparation time • Relies on tissue donation • Risk of disease transmission • Lack of uniformity amongst specimens
Synthetic bone simulation	<ul style="list-style-type: none"> • Relatively inexpensive and portable • Widely available • Develop understanding and familiarity with orthopaedic instruments and equipment 	<ul style="list-style-type: none"> • Does not allow understanding of the influence of soft tissues • Lack of true haptic feedback • Single use
Cognitive simulation	<ul style="list-style-type: none"> • Potentially cost-free • Accessible on mobile devices • Point of care education 	<ul style="list-style-type: none"> • Limited evidence to support use in clinical training/improvement in technical procedural skills
Virtual / Augmented reality	<ul style="list-style-type: none"> • Able to record progress and assess motion analysis • Allows for development of hand-eye coordination and complex tool's manipulation • A wide range of procedures may be possible • Allows for scenario simulation 	<ul style="list-style-type: none"> • High initial setup costs • Available haptic devices can be below the feedback requirements for some orthopaedic procedures and therefore limit realism

Table 1 Adapted from Stirling et al. (2014)

Even after the evident progress in computer graphics and motion tracking technologies during the past seven years, several authors agree on the absence of VR technology application in the hip surgical field (Vaughan *et al.*, 2016). The first procedures with which the simulation industry dabbled in surgical training were those related to laparoscopy. Mabrey et al. (2010) attempted to identify the market and contextualise their developed knee arthroscopy simulator. At that time, laparoscopic VR simulators started to incorporate themselves in the training curriculum, and orthopaedic procedures were “far less common” to find in literature. Some authors have performed systematic reviews of VR simulations for orthopaedic training. Vaughan et al. (2016) present arguments to emphasise that surgical skills training for total hip replacement procedures has not enough

representation in the spectrum of virtual reality-based simulators. In addition, Li et al. (2017) mention some trending products as examples of how training environments lead the market of hip procedure simulators for technical skills such as trauma operation, drilling of femoral bone with haptic feedback, and positioning the elements for fractured bone fixation.

2.2.1 Comparative assessment of current relevant VR simulators for orthopaedic training

The purpose of this section is to describe and compare several VR simulators that claim to be suitable for orthopaedics training found in both the market and the literature. Current reviews on simulators for orthopaedic surgical procedures lack objective metrics of assessment given a standard set of design requirements. Instead, most of them rely exclusively on the level of interaction, the functionality and the features provided. The simulators were grouped into three different categories based on their hardware and interaction characteristics, namely cognitive simulations, arthroscopy VR simulators, and open surgery simulation.

The category of cognitive simulators includes all solutions that do not provide a natural human interaction with the digital content. This means that the users' interaction is limited to mouse click events or pressing buttons available on gaming controllers. The main aim of this type of simulators is to illustrate the processes involved in different orthopaedic procedures rather than provide a real hands-on experience. One work included in this category was the computer-hosted simulator developed by Blyth et al. (2008). This software guides the trainees through the steps of hip trauma surgery and lets them choose from a set of defined tools by clicking a button on the computer screen.

Another popular example is the application for smartphones Touch Surgery (TouchSurgery, 2017). This app illustrates through animations and slides the steps involved in a wide variety of procedures. For Total Hip Replacement, the user interaction is limited to move forward or backwards in the sequence of animations. A more immersive but not interactive enough solution is OssoVR surgery simulation (OssoVR, 2016), available for some arthroplasty procedures and developed for VR headsets. Even though the environment provided by OssoVR is visually immersive, its interaction is limited to the movement of the gaming controllers which handles are unnatural compared with the grip of a surgical tool. The lack of haptic feedback breaks the illusion of immersion and reduces training skills that are actually transferred to the surgical theatre.

As part of the second group, we decided to include two of the biggest market

representants of arthroscopy simulators due to their privileged position in the market—namely, Virtamed ArtroS (VirtaMed AG, 2014) and Arthromentor (3Dsystems, 2017). Arthroscopy differs from arthroplasty in several ways. For example, the procedures are less invasive as small incisions are required to insert the small probe camera and the tools manipulated by the surgeon based on the projection of the camera view on a screen. As the surgeon is looking at a monitor, the field of view during both the real and simulated experience is exactly the same. The layout of the surgery itself makes it easier to replicate in a simulated environment as the range of motion of the instruments is limited, and the movement of the tools is easy to track.

The last category includes simulators developed to train skills for open surgery. Unlike arthroscopy, open surgeries require a more extensive incision on the patient's tissue and an adjustment of the surgeon's angle of view to get an unobstructed exposure of the relevant bones, nerves and muscles during the operation. In this comparative study, all of the training tools identified used a haptic unit consisted of 1 or 2 arms to handle the interaction. TraumaVision, developed by Swemac (2017), is a simulator for hip fracture and trauma surgery that allows the surgeon to use an X-ray view of the patient while performing the fixation of a hip fracture. The simulator also allows the users to move the patient's legs and complete basic interactions triggered by buttons on the screen. At the end of the simulation, users receive a numerical assessment of their performance based on the ideal positions of the fixation elements. SimOrtho (OsSimTech, 2018) includes a 3D display and an additional screen for the learner to get an X-ray and patient representation view at the same time. The 3D display provides a depth effect to the experiences but has no feedback from the user's perspective. OsSimTech stands out among its competitors as it uses a personalised haptic accessory attachable to different tools, improving the range of motion and making the human-machine grip more realistic. Lastly, FundamentalVR (2018) released Fundamental Surgery, a visually immersive VR experience with a commercially average haptic unit that includes Total Hip Replacement training. An assessment stage is still under development, but feedback for the achieved angles of cup implantation is provided while performing the surgery. It is worth noting that both of the latter mentioned simulators are still under continuous development phase, with Fundamental Surgery being released in June 2018 and the last update with breaking hardware and software changes of SimOrtho released in October 2018. The scores assigned in this study are given according to the versions available to the date.

The criteria to perform the market state-of-art evaluation were divided into three

categories, as shown in Table 2. The immersion criteria refer to the similarity between the simulated task and the surgery in real-life regarding its visual representation and interactivity. This group of criteria is usually dependent on the hardware specifications as sometimes unique handles mimicking surgical instruments are required to accurately train specific hand-eye coordination skills. On the other hand, the “response of the rendered content to the interaction” evaluates the similarity of the triggered deformations on the digital content with the ones expected in real surgery. These requirements usually depend on the software design or the deformation algorithms included. Nevertheless, accurate deformations can also require considerable computational resources. To obtain a high score in the second group of criteria, simulators should implement deformations algorithms such as physical inputs in the real surgery should produce similar visual or haptic outputs as the one rendered in XR.

The final set of criteria evaluate the inclusion of some relevant learning theories in the content design of the different simulators. The training of decision-making skills ensures a learner-centred process. An ideal simulator requires some modularity to allow trainees to practice with different sets of tools and diverse cases that would enable deliberate practice. Furthermore, as explained in section 2.1.3, a valuable assessment report at the end of the training experience is essential to allow reflective observation during experience-based training.

Immersion	Digital content response to real-time interaction	Applied learning theories
a. Field of view similar to the reality	d. Mesh deformation based on input forces and tool movements.	g. Decision making allowed
b. Realistic content rendering	e. Interactive kinematic models to allow change of patient position	h. Modular (different tools, and approaches to select from)
c. The interaction feels natural (haptics and tools)	f. Interaction with all rendered tissues	i. Detailed and valuable assessment report

Table 2 Assessment criteria with listed requirements for score assigning

The results of the assessment per simulator are listed in Table 3. In order to provide an objective quantitative evaluation for different types of simulators with different strengths, there were three requirements assigned to each criteria group expected to be met. The assignation of the value was performed as described below:

- 0 - No requirement met.
- 1 - One requirement partially addressed, 2 not included

- 2 – Two requirements partially addressed, 1 not included
- 3 - All requirements partially addressed
- 4 – Two requirements completely met, and 1 just partially included
- 5 - 3 requirements completely met in the content of the simulators

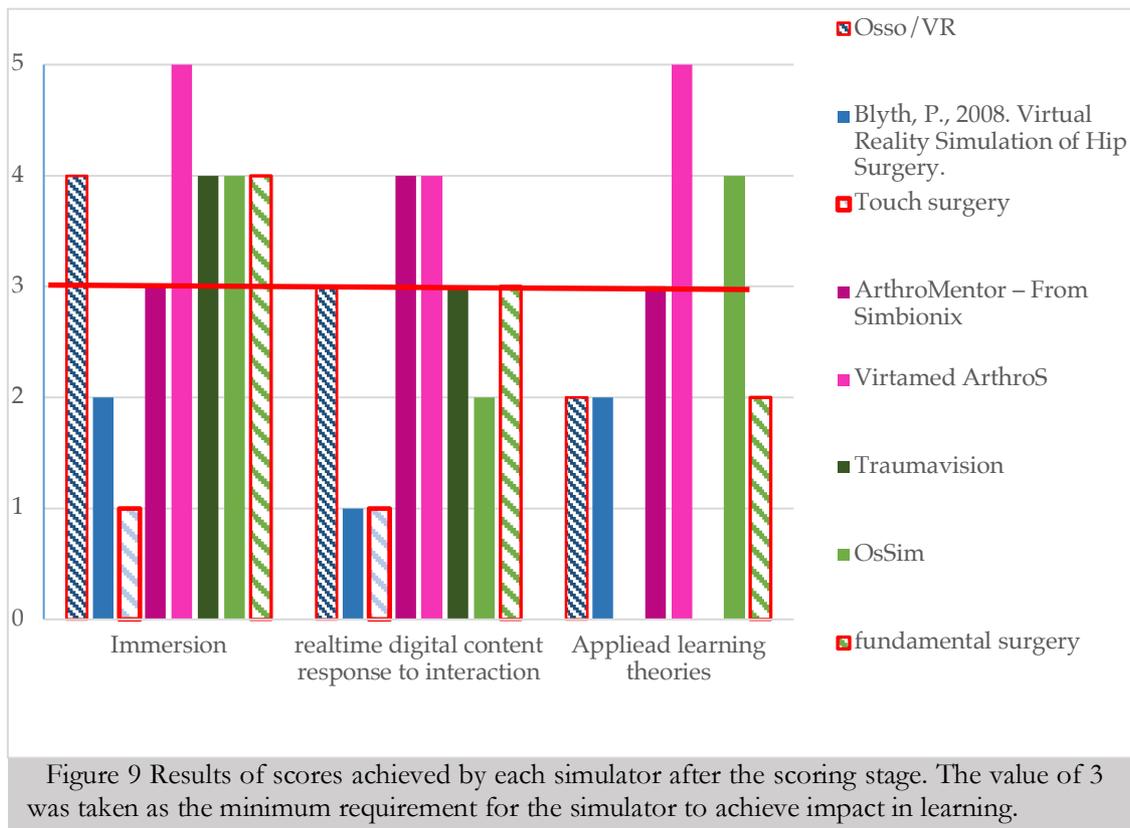
Simulator/requirements	a	b	c	d	e	f	g	h	i
OssoVR Osso	○	○	∩	-	∩	∩	∩	-	∩
Blyth P. Virtual Reality simulator	∩	∩	-	-	∩	-	∩	-	∩
Touch Surgery	∩	∩	-	-	∩	-	-	-	-
Simbionix - ArthroMentor	○	∩	∩	○	○	∩	∩	∩	∩
Virtamed ArthroS	○	○	○	○	○	○	○	○	∩
TraumaVision	○	∩	○	∩	∩	∩	-	-	∩
OsSim Tech	○	○	∩	-	∩	∩	∩	○	○
Fundamental Surgery	○	○	∩	∩	∩	∩	∩	-	∩

Table 3 Assessment result per simulator. ○: The requirement is met. ∩: The requirement is only met partially. -: The requirement was not met.

Figure 9 shows the scores assigned to each simulator according to the previously defined assessing criteria. The results are grouped by criteria and type of simulator. In the diagram, cognitive simulators are shaded using different tones of blue, arthroscopic ones using magenta tones and open surgery simulators using greens. Simulators that include a module for total hip replacement are highlighted with a red border and a pattern shading.

In the resulting assessment, cognitive simulators were awarded the lowest values from the gamification perspective, while open surgery simulators are closer to satisfy the minimum average value of 3 for all the assessed requirement. Cognitive simulators usually lack the application of learning theories or allowing the testing of the user’s decision-making skills. Arthroscopy simulators were assessed with the highest score in almost all categories. These simulators have received accreditation from schools of surgeons and are being used in training curricula already (Antonis *et al.*, 2019). As mentioned before, the nature of the procedure makes it easy to replicate, and there has been a significant amount of development around their validation and improving their learning impact (Morgan *et al.*, 2017).

Open surgery simulators are close to satisfying the minimum average value of 3 for all the assessed requirements. However, improvements are still required in order to include a meaningful feedback report in open surgery simulators. Detailed feedback is essential for surgeons to construct knowledge models rather than just succeeding by completing a task.



Open surgery simulators are close to meet the minimum score for all assessment categories, but there is still work to be done to allow decision-making and provide accurate haptic interaction.

2.2.2 Analysis and comparison with independent validation studies.

Several validation reports in the literature support the scores assigned in Figure 9. The validation of training simulators is defined as the measurement of their reliability, and depending on the features to be assessed, it can be referred to as face, construct or content validation (Schijven and Jakimowicz, 2005). A face, content and construct validation study of the hip module of the Virtamed ArthroS establish that both expert and novice surgeons consider the visual and haptic interaction of the simulator to be natural, and the majority believes the interaction is adequate to train the skills of hip arthroscopy (Antonis *et al.*, 2019). The validated simulation consisted of capturing several anatomical landmarks with the camera probe while minimising the tool's movement and the area of touched tissue. In that validation study, the acquired performance scores from novice and experts deferred when assessing the damage caused to the structures by manipulating the tools and the amount of distance travelled by the instruments. However, there was no significant difference between

the scores acquired by both groups when comparing their ability to expose predetermined points within the hip joint to the camera field. The face and content validation scores assigned by Antonis et al. (2019) for the Virtamed ArthroS are similar to the scores assigned in this assessment.

Another existing face and content validation study was carried out on a Total Knee Replacement simulation running on a version of the OsSimTech earlier than the one considered in this report (Newman *et al.*, 2018). In general, experienced surgeons did not wholly agree with the realism of the simulation. This result would be equivalent to a partially met requirement in our comparative study. The visual content of the simulation was significantly improved in version 3.2 released in 2019. However, the haptics of the tools and their range of motion were just partially addressed for the interaction to feel completely natural. The content validity in this report was assessed by the surgeons' opinion of the accuracy of the simulation. Moreover, all participants partially agreed that it was helpful for training on drilling, sawing and other skills (Newman *et al.*, 2018). An important observation was that experienced surgeons provided lower scores than novice ones to the content validity of the simulator.

Fundamental Surgery self-reported a pilot face validation study performed with surgeons on mixed training stages including medical students (Rainger, 2019). Participants found that the experience's realism is sufficient from the visual and haptic perspective when performing a total hip arthroplasty. However, unlike other validation studies reported in the literature, no consultants were recruited during this study. The results of the face validation study were reported as a joint percentage of trainees that agree and strongly agree with the validation statements. As there is no reported data to establish the frequency difference between both the *agreeing* and *strongly agreeing* used as validation metrics, it is impossible to directly compare this face validation study to the scores assigned in Figure 9.

Besides a realistic graphical interface, simulator interaction must feel natural and contain physically accurate reactions to the user inputs. Currently, a trade-off has to be made when enabling decision trees, physically accurate deformation and haptic feedback because of technological constraints. Also, the interaction with the tools must feel natural in terms of haptics and range of motion to help the user to see the relevance of the experience and keeps them motivated towards practice. Together with a good simulator design, simulated experiences must be aligned with the day-to-day life of healthcare professionals, keeping in mind their time constraints (Clapper, 2010; Wang, 2011). On younger generations, previous

gaming experience could give the trainees an undesired advantage during the simulation. Therefore, the design of the interaction interface must minimise the risk of technological bias by replicating the real-life range of motion, grip, and feedback of the surgical tools.

2.3 Discussion.

Different requirements in a simulator are needed at different stages of surgical training. In Table 2, the ones we propose as ideal for a simulator intended to be used after the trainee has a basic understanding of the procedure until achieving proficiency for independent performance. It is worth mentioning that for a fully visual immersive simulator, the more rendered tissues, the harder it is to provide interaction with all 3D content. This performance decrease should be considered when analysing the assessment reported in Figure 9. Although extensive interaction is vital for allowing decision making during training, the high amount of required collision detection would make the software complexity untractable from the hardware perspective.

Although there has been significant progress in the simulation of elastic deformation for medical procedures, orthopaedics simulators require algorithms to predict accurate plastic deformations based on the user force and movement input. More research is needed in this area as irreversible procedures can benefit massively from an accurate simulation opportunity. FundamentalVR tracks the orientation of the surgical tool during the reaming step, but there is no progressive evolution in the appearance of the acetabular surface during reaming. Instead, a change in the surface texture is triggered after the reamer has been inside the acetabular cavity for a certain time. Physically accurate interaction with all rendered tissues and a dense range of decision possibilities would be ideal for achieving expertise and evaluating problem-solving skills. However, it is worth recognising that the available technology has memory and performance limitations to meet all suggested comparison criteria. Though, the required computational resources can make the experience challenging to deliver at interactive framerates, especially if an accurate model of the mechanical properties of the bone is desired. An equally important feature that requires tremendous computational resources is providing opportunities to train the trainees' decision-making process, allowing them to make mistakes and displaying the appropriate consequences during the interaction. More research in complex data structure manipulation is required for future simulators to be able to provide such training possibilities.

Finally, Table 3 is evident that a physically-based mesh deformation based on the tool's forces and movements is a requirement that is still missing in many of the identified simulators. As a matter of fact, only arthroscopy simulators meet this requirement completely. However, elastic deformations, simulated for training on arthroscopy, are not comparable to the plastic resurfacing ones present in procedures such as acetabular reaming.

Chapter 3

An augmented reality-based simulator for training on surgical navigation skills.

Imageless navigation systems (INS) and Robotic Assisting Arms have been used in orthopaedics to improve the outcomes of several orthopaedic procedures such as THR. This chapter presents a Mixed Reality (MR) simulator that helps surgeons acquire infrared-based navigation skills before performing a real surgical operation. At the time of its development, to the best of our knowledge, there were no existing solutions using head-mounted displays to train users into the steps to calibrate such systems. The simulator was developed using Unity3D and executed on the Microsoft HoloLens headset. The device capabilities were enhanced using the PTC Vuforia engine to support the movement recognition of the calibration tools. Finally, a group of 7 hip surgeons were invited to try the application, expressing their satisfaction with all of the features evaluated. The latter allows us to affirm that our simulator represents a cheaper and faster option to train surgeons in INS than the current training methods.

3.1 Robotic-assisted surgery and Imageless Navigation Systems for THR .

Robotic-assisted surgery removes human error and provides a more accurate and replicable cup orientation than manual implantation, thanks to its navigation system and powerful haptic units. After conducting a systematic review, Tilly (2016) defends that after

performing a robotic-assisted surgery, the range of motion of the hip is broader, and the soft-tissue damage is minimised. Even with the mentioned advantages, the learning curve involved in this new way of performing the surgery may delay its widespread use (Newman, 2014). In real procedures, the infrared camera (IR) cast rays which are reflected at the surface of passive spheres attached in a specific position on the surface of a unique mount. To ensure that trackable objects are always visible to the camera, this mount has to be pointing in a similar direction as the camera view and remain inside its tracking volume. The spatial calibration step on each surgery is essential to ensure the reliability of the surgery. This step allows the robot to compute its relative positions to the patient and ensure that the implants' planned orientation and location are correct. Therefore, both trainees and expert surgeons should learn to work without interfering with the continuous tracking of both the patient and the robot.

Like robotic-assisted surgery, imageless navigation systems (INS) provide an alternative to conventional methods in achieving a more accurate position of the cup implant intraoperatively than conventional methods (Schnurr *et al.*, 2009; Chang *et al.*, 2017; Snijders *et al.*, 2017). Unlike other image-based navigation systems, INS minimises the amount of radiation the patient is exposed to by avoiding the need of extra CT scans or X-Rays as the relevant surgical planes and axes are constructed through precise registration of bony landmarks (Chang *et al.*, 2017).

In order to provide both INS and Robotic Assisting Arms with the correct spatial data input and avoid further orientation mistakes intraoperatively, surgeons and trainees must master the skills needed for a proper calibration stage using the infrared tracking camera. Longer surgeries mean a more prolonged state under anaesthesia and a higher risk of infections for the patient. Unfortunately, while surgeons learn to adapt their instrument movements to be recognised by the infrared cameras, surgical times can be extended up to 20 minutes more (Schnurr *et al.*, 2009; Silvennoinen *et al.*, 2012). However, this amount of extra time needed for positioning the optical unit and point registration can be shortened up to 4.8 +/- 3.8 minutes after getting valuable experience with this kind of device (Thorey *et al.*, 2009).

Mixed and Augmented reality has been applied to the medical training field due to their standalone nature and spatial understanding capabilities. The latter allows the system to blend the pre-processed digital content into the physical world through spatial anchoring. Among some examples are CAE Healthcare (Healthcare, 2017) and Fundamental VR (VR,

2017), which have developed holographic interfaces to train medical staff in diagnostic ultrasound, anatomy and surgical approaches. However, to the date of development and to the best of our knowledge, there is no existing MR simulator to help surgeons train in the use INS without requiring high budget technologies such as the system itself.

3.1.1 Development

This section explains the decisions carried out in terms of the software and hardware selections of the assets used to develop the simulator.

3.1.1.1 Hardware selection

The spatial alignment ensures that the digital content rendered by the XR headset is perceived to have the same position, orientation and scale as a passive haptic model. Augmented reality (AR), was chosen over Virtual Reality (VR) as the users are still able to see their real environment through them. This feature makes AR applications easier to align with the required passive haptics and minimises the risk of user tripping over cables or other objects during the interaction.

The development of AR technologies has grown fast over the last five years. Currently, there is a variety of AR headsets available in the market, such as the Microsoft HoloLens® (1st and 2nd edition) (Microsoft, 2017), the MagicLeap® AR glasses (Magic Leap Inc., 2018) and the Epson Moverio. Furthermore, both of the world's main smartphone manufacturers have released machine vision libraries to provide AR experiences while using the information from a single RGB camera. These APIs are the Google's ARCore®(Google, 2018) and Apple ARKit®(Apple, 2018). The multiplatform nature available for AR experiences provides an opportunity for the development of collaborative experiences where trainees and instructors can immerse themselves in the same holographic environment and enable social interactions. Even though AR experiences running on tablets do not provide full interaction with the digital content, a joined immersion between instructors and trainees still represents a training advantage. In real surgical procedures, the number of spectators in a room is constrained to maintain air quality and decrease infection risks. Through a shared AR/VR experience, users can observe the steps followed by a trainer or a peer.

Head-mounted displays provide the head tracking capabilities required to simulate scenarios with limited tissue exposure, such as those in a THR posterior approach. These devices are able to update the displayed digital content according to the position and

orientation of the users' head while letting them interact with the digital content. Opinions in the literature about the use of current AR headsets are controversial and constantly changing. Even though some authors discourage the use of HoloLens because of their limited field of view (Silvennoinen, 2017), other authors encourage their use not only for medical training but also during surgery (Cui *et al.*, 2017; McNutt, 2017), arguing that the field of view limitation loses relevance for open surgery simulation as surgeons usually focus their sight into the incision area.

Our prototype was designed to train surgeons on the basic skills needed during the calibration stage of an imageless navigation system, namely the awareness of the infrared camera's tracking volume and how to manipulate surgical instruments while avoiding targets occlusion. The design phase ran on the HoloLens1, an MR headset with spatial understanding capabilities capable of rendering holograms in the user's field of view while tracking their movements in space. The Microsoft HoloLens has been used in the medical training field due to its mixed reality (MR) capabilities and standalone nature. Microsoft defines MR as the intersection between real-time computer processing, human interaction and spatial understanding (HoloLens, 2017). The latter allows the system to blend the pre-processed digital content into the physical world through spatial anchoring.

3.1.1.2 Software development.

The simulator was developed on the game engine Unity3D (Unity Technology, 2018). This platform was selected as Microsoft provides a development kit for the HoloLens that include helper methods, allowing us to use the information obtained from the tracking sensors RGB-D cameras to enhance the user's interactivity. Unlike standard gaming VR headsets, the Microsoft HoloLens 1 does not include gaming controllers to handle the user's interactions. Instead, the orientation of the user's head in space is used as an analogy for a 3D cursor capable of interacting with the rendered holograms. Additionally, the hand movements are tracked, and by performing a hand gesture, the user can "click" over the content as desired. In order to make the best use of the device features during the developing stage, we used the following assets from the HoloLens' Software Development Kit (SDK) (Microsoft-OpenSource, 2018):

- **GazeManager:** This asset reads the tracked position and orientation of the user's head based on the device's GPS and motion sensors. The user's gaze vector is defined as an imaginary line that points outwards from a point between the user's eyes, it is normal to the lenses of the headset and follows

the orientation of the head.

- **CursorManager:** This asset takes as its inputs the orientation of the user's head and the spatial understanding¹⁰ capabilities of the headset. It finds the optimal location of the cursor in the 3D space by evaluating the collision between the user's gaze and the rendered holograms.
- **Gesture Manager:** Once the cursor is located on top of a hologram, this machine vision asset recognises the "clicking" hand gesture performed by the user.
- **StandardFast shader:** A custom shader optimized to run in standalone¹¹ devices responsible for speeding up the rendering pipeline of the holograms.

In real procedures, the infrared camera (IR) cast rays which are reflected on passive spheres attached to the surface of the surgical instruments to determine their location with high accuracy. Therefore, it is important to be aware of the user's instrument manipulation during simulation to provide a valid training alternative. To do this, we have included in the simulation a computer vision application programming interface (API) that allowed us to measure the position of a particular object relative to the AR headset. The selected tool was the PTC Vuforia API (Vuforia, 2018), which estimates the position of objects by processing the visual information from the RGB camera of the device and looking for *image targets*. Vuforia can be easily be imported as an external asset in the Unity3D development environment. During the development phase, this allowed us to create image targets and include them as objects in the scene, which would obtain their position and orientation once these are estimated in space by Vuforia given the information received from the RGB camera.

The chosen image target must be set during the design stage of the AR application. Vuforia first detects feature points in this target and then uses them to compare them with the receiving frame from the camera in real-time¹². A feature point is defined as a “*sharp*,

¹⁰ A digital awareness of the environment achieved by a continuous 3D scanning of the device's surroundings.

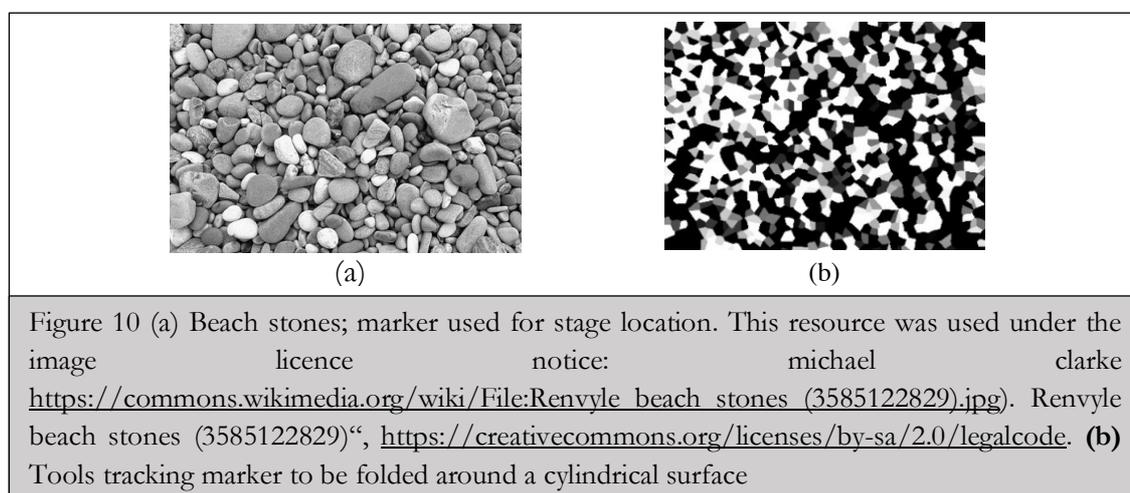
¹¹ For example a 3D printed model or a synthetic bone commonly used for single-use training (Hetaimish, 2016).

¹² A system provides results in real-time if it reacts to a steady flow of new information immediately and without interruption.

pointed, chiselled detail in the image” (Vuforia, 2018). Image targets can also be folded around a cylindrical or brick object. It is essential to mention that the tracking performance depends on the device's camera resolution, lighting conditions, and the percentage of the marker inside the RGB camera's field of view (Simonetti Ibañez and Paredes Figueras 2013).

Figure 10 shows the two raw images used as markers in our MR simulator. We ensured that both targets had a high density of feature points arranged in non-repetitive patterns over the picture. The image target in Figure 10(a) was chosen to help the users to locate the holograms representing the surgical room in the best suitable position for their space. This image file was designed to have a minimum of 20-30 feature points per 10% of the total area of the image to guarantee smooth tracking (Simonetti Ibañez and Paredes Figueras 2013). Vuforia provides developers with a Target Manager tool to assess the quality of their targets before including them in the AR applications. The density of point features in the picture was validated in the target manager, obtaining an assessment score of 4 out of 5 stars. Figure 10(a) was printed on an A4 size to ensure accurate orientation recognition at distances greater than 1m from the frontal RGB camera of the HoloLens 1.

The tracking of the tools' handles represents a more demanding challenge. This is because the dispersion and distribution of a high number of feature points can be indistinguishable on a small target image at a considerable distance from the RGB camera. Therefore, we design a target image with a density and disperse feature points around the image (Figure 10(b)). This marker ensures smooth and continuous detection and recognition, avoiding oversized targets. Given that a point feature is created in a point of high contrast,



the grey-scale¹³ values of the image's pixels are constrained to a range of 5 values (0,50,100,177,255). This was done to facilitate real-time feature detection. A predominantly black background was chosen for the marker. Simonetti Ibañez and Paredes Figueras (2013) proved that Vuforia API requires a difference of 50 units in the grey-scale values of adjacent pixels to recognise the change as a feature on dark backgrounds. Alternatively, the API would require a minimum difference of 60 units between lighter tones. In addition, the distribution of the features and their density inside the image were generated to avoid supersaturation when their printed version was located more than 30 cm away from the RGB camera. Figure 10(b) shows the final image target created for the tool's handle. This image file was assessed by the Vuforia Target Manager, obtaining a score of 5 out of 5 stars.

The image chosen for the tool was folded around a cylindrical object. After trial and error, cylindrical volumes were found to be the most suitable shape for smooth tracking of hand handled instruments. A flat image would require the camera to be facing it directly at all times, which would make the instrument manipulation unnatural. Additionally, a cuboid object showed difficulties in tracking the target's position when the camera was facing its edges. Instead, the cylindrical shape provides the camera with a continuous size of the projected area when the instruments are rotated around their symmetry axis.

3.1.2 Application description.

The location and orientation of the digital surgical scene were attached to the image target shown in Figure 10(a). This allowed the users to locate the simulation scenario in their preferred position by placing the image on top of a flat surface. Within the digital scene, a digital representation of a patient was laying on its left side with its hip joint previously dislocated as concordant with real surgery. When using an INS in real surgery, a fixed frame of reference is attached to the patient pelvis (Figure 11(c)). This allows the INS to measure the position and orientation of the tracked tools relative to the patient's body, even if this must be moved during surgery. We included a digital representation of this reference frame to introduce the concept of relative tracking to trainees.

¹³ The grey scale is a continuous measure of light intensity that covers a range from 0 to 255. In this scale black areas correspond to values closer to 0 and white areas to values closer to 255.

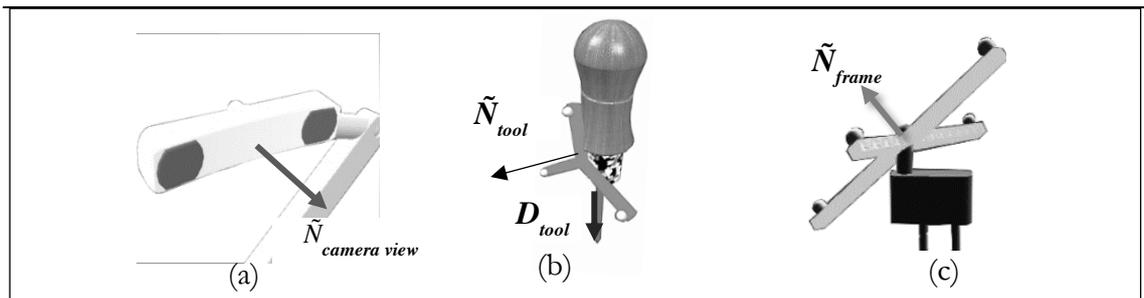


Figure 11 Direction of interactable vectors. (a) Digital model of the infrared camera, $\tilde{N}_{camera\ view}$ represents the unit vector normal to the frontal face of the camera. (b) 3D model attached to the cylindrical tracked object, \tilde{N}_{tool} represents the normal of the mount surface and D_{tool} a normalized vector in direction of the cylinder's axis. (c) Frame of reference fixed on the pelvis. \tilde{N}_{frame} represents the unit vector normal to the surface on which the 3D models of the reflective spheres are supported.

Another essential component in INS procedures is the awareness of the position of the infrared camera (IR). This camera cast rays, which are reflected at the surface of reflective spheres attached to the tracked tool. The volume covered by these rays forms a truncated pyramid, and it is known as the *tracking volume*. To make objects always trackable during surgery, all the spheres on the mounts of the trackable instruments have to be reachable by the infrared rays and be inside the camera's tracking volume at all times. Figure 11 shows a simplified representation of the tracking system. Each mount is assigned a vector normal to its surface (\tilde{N}_{tool} and \tilde{N}_{frame}).

Similarly, the direction of the camera view is represented by the vector $\tilde{N}_{camera\ view}$. The angles between these vectors will determine whether an object is recognisable by the INS. In Figure 11(b), the vector aligned with the symmetry axis of the tracked tool D_{tool} is required after the calibration stage in order to determine the relevant angles during the surgery. As a standalone device, the hardware specifications of the HoloLens limit the rendering quality and mesh size that is able to be used in the simulation. This means that when big meshes have to be rendered simultaneously, the framerate of the simulation drops. An unstable framerate causes the hologram to be unstable, producing jitter as high-frequency shaking of the holograms. Consequently, the number of polygons used in the 3D models was minimised to a total value of around 84000.

3.1.3 Setting up of surgical scene and the navigation system.

First, a printed version of the image target was attached to the surface of a standard surgical bed. During the first stage of the simulation, surgeons were required to adjust the bed to their preferred height, just as they would do in the surgical theatre. When they were satisfied with the height of the table, they sent the system a command to fix the scene in the chosen position using a spatial anchor. Figure 12(a) shows a surgical scene after being anchored to real-world coordinates.

During real surgery, the orientation of the IR camera controls the tracking volume covered by the INS system's field of view. Users must ensure that their workspace and all instruments and targets are inside the modelled IR camera's field of view. In this simulator, surgeons are able to use the HoloLens gesture recognition to modify the orientation of the camera. The tracking volume could also be displayed or hidden at the user's preference during the camera set up phase.

After setting up the IR camera, the acquisition of the points or anatomical landmarks on the pelvis and femur bone surface are the most critical steps during the calibration of INS. These steps allow the system to define the anatomical planes in which the angles of resurfacing and acetabular cup implant insertion will be projected as holograms.

During the point registration stage, an important skill to train is learning to hold the surgical instruments inside the infrared camera tracking volume and not occlude its field of view. Not understanding the cause of instrument tracking loss is the main reason for increased surgical times when using INS. Given that only properly orientated instruments would be able to acquire anatomical landmarks in reality, the relative orientation between the normal of the tracking mounts \tilde{N}_{tool} and the camera view vector \tilde{N}_{camera_view} is evaluated during



Figure 12 (a) Mixed reality capture, surgical scene after setting up and anchoring the holographic content in space. The transparent- yellow parallelepiped gives the user a visual feedback about the location and orientation of the system tracking volume. (b) Mixed reality capture of the simulator.

the simulation at each frame. Therefore, after the user places the probe on an anatomical landmark, this will be acquired during the simulation when the angle of incidence $\theta_{incidence}$ between the tools normal \tilde{N}_{tool} and the IR camera normal \tilde{N}_{camera_view} is lower or equal to 45° . (Figure 13(b)). After a successful point acquisition, the simulation displays a loading screen and plays audio feedback to let users know that they are pointing to the right position.

3.1.4 Tool's orientation.

The tools orientation is measured in two anatomic planes, namely the coronal¹⁴ and sagittal¹⁵ plane (Figure 13(a)). The desired orientation for cup and bone preparation is planned preoperatively and can vary according to the gender, body mass index and other anatomical variations of the patient.

During the insertion of the acetabular component, the two angles that define the proper orientation of the implant, according to anatomical landmarks, are the inclination and antroversion angles. The inclination is defined as the angle between the instrument axis (D_{tool}) and the longitudinal axis when projected on the coronal plane; antroversion is the angle between D_{tool} and the coronal plane projected on the sagittal plane (Bhaskar *et al.*, 2017). $N_{coronal}$ was defined as the vector normal to a plane that contains the first three acquired points on the pelvic bone; namely at the pelvic right and left ASIS¹⁶ and the pubic tubercle.

The angles of antroversion and inclination $\theta_{antroversion}$ and $\theta_{inclination}$ are defined by the equations (1) and (2). Here $D_{tool_projected}$ represents the projection of the vector D_{tool} on the coronal plane, while the longitudinal axis of the patient's L_{axis} is a vector laying in the coronal plane and running straight through the top of the head from down between the feet.

$$\theta_{antroversion} = 90 - ((\cos^{-1}(\overrightarrow{N_{coronal}} \cdot \overrightarrow{D_{tool}})) \frac{360}{2\pi}). \quad (1)$$

$$\theta_{inclination} = ((\cos^{-1}(\overrightarrow{L_{axis}} \cdot \overrightarrow{D_{tool_projected}})) \frac{360}{2\pi}). \quad (2)$$

¹⁴ Plane that divides the body into front and back sections.

¹⁵ Plane that divides the body into left and right sections.

¹⁶ Anterior superior iliac spine

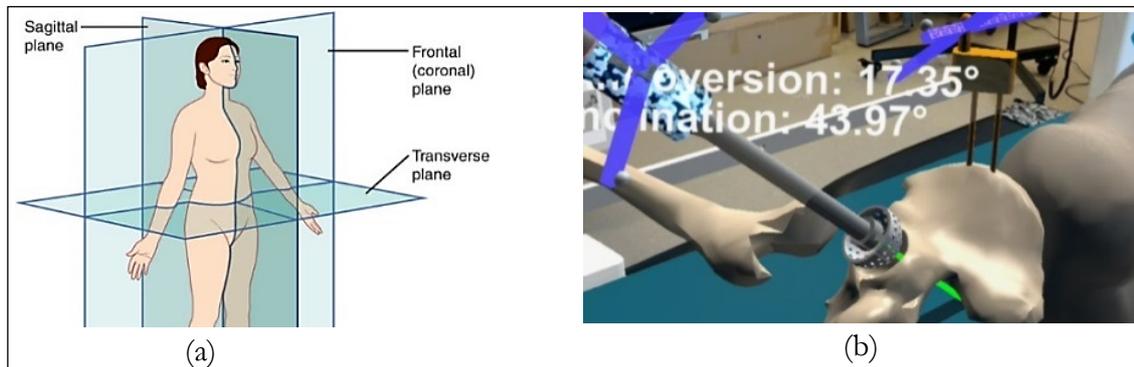


Figure 13 (a) Planes of the body. This resource was used under the image licence notice: Connexions CC BY 4.0 (<http://creativecommons.org/licenses/by/4.0>), via Wikimedia Commons. (b) Mixed reality capture, reamer orientation feedback. A ray is casted from the reamers tip in the direction of the tool, this ray changes its colour according to how close the user is to perform the procedure in an orientation close to the safe zone defined in the literature.

Both previously calculated angles are displayed at the navigation system screen and at the same time attached to the headset display so that the users can see them and use them to guide their movements during the reaming operation (Figure 13b). Additionally, a line is rendered from the tip of the reaming tool in the direction D_{tool} . The colour of this ray changes according to how close the trainee surgeon is to the "safe reaming zone"; defined by the literature as an inclination angle of $40^\circ (\pm 10^\circ)$ and an anteversion angle of $15^\circ (\pm 10^\circ)$ (Chang *et al.*, 2017; Rojas *et al.*, 2018).

3.2 Visual interface, second version

To maintain the learner's motivation, an anatomically accurate and realistic 3D representation of the relevant tissues is required to provide realism and immersion involved in a THR procedure simulation is paramount.

The development of the required 3D models, textures and detail maps was done in collaboration with an undergraduate student research assistant. The student was instructed on the relevant anatomy terminology and guided into modelling the details on the muscles, bones and cartilages exposed during the posterior approach and relevant for the procedure. Other body structures were simplified in order to cover the volume spaces visible through the incision.



Figure 15 Screenshots of the app running on the device

All digital 3D geometrical models were generated with a low polygonal count without compromising appearance; this was necessary because the high polygonal count can decrease the headset rendering framerate and generate discomfort for the user. Additionally, details were included during the texture generation of the acetabular surfaces to generate two versions of them. One texture corresponds to the acetabulum in a preoperative state with clear signs of osteoarthritis, and a second one to the bleeding bone after reaming is completed. This is similar to the rendering approach used by FundamentalVR in their THR simulator (2018).

Additionally, rigging controls and animations were created for all the exposed muscles

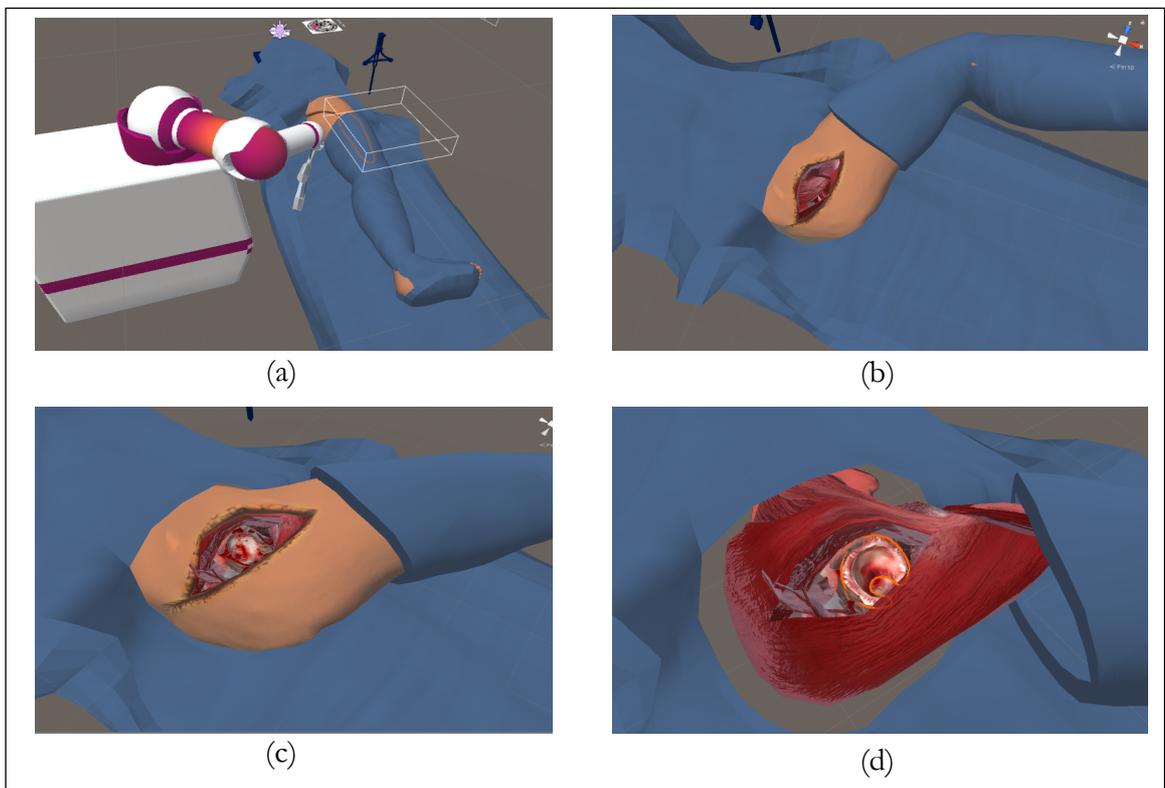
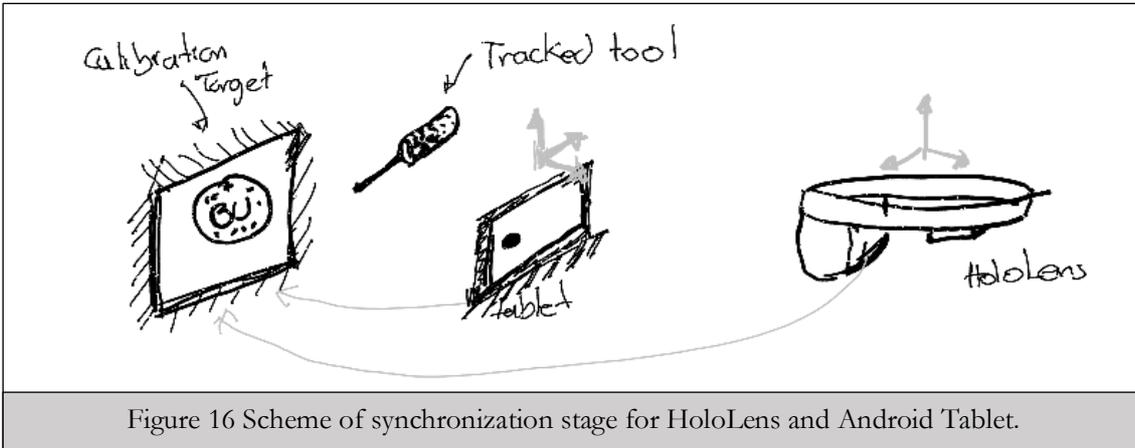


Figure 14 3D model and time captures of the animation of a posterior approach



and cartilages. These animations included the steps to follow to perform a posterior approach to the hip joint, such as muscle incision, traction, hip dislocation, and femoral head removal. So far, this rig is not interactive, and the animation was imported as controlled blend-shapes into the Unity3D game engine to provide an introductory scenario to the simulation of the reaming stage. Figure 14 illustrates some screen captures of the model inside the game engine. Figure 14(a) shows the initial state of the simulation where all models remain in a neutral position until the user decides to trigger the animation. The sequence shows the incision being opened, the skin and the gluteal muscles are retracted to maximise the exposure area, the femoral head is removed, hip rotators and capsule are divided and folded the incision outwards, and the acetabulum is exposed (Figure 14c-d). The orthopaedic surgeons visiting ORI mentored the modelling and the animation stages.

3.3 A communication protocol between Android devices and HoloLens

A common issue observed while performing the validation stage of the previous early prototype was the learning curve involved in getting used to the optical tracking system and keeping the marker always inside the RGB camera field of view. Furthermore, by relying on the AR headset to track the tool's manipulation, we are decreasing the computational capabilities available for holographic rendering.

A tracking alternative was developed to release the computational load on the headset by moving the tool's tracking responsibility from the HoloLens to an external device. This external camera is located in a strategic position to be aware of the tool's movement at all time and share its coordinates through a wireless Transmission Control Protocol (TCP) to the HoloLens (which are then able to transform them to their coordinate frame). This interface allows the development team to control the space area to be tracked by locating the tablet in a convenient fixed position. If a tablet is used as an external camera, the TCP protocol does not require additional hardware to set up the wireless network. The tablet can be both server and client at the same time. A holographic UI interface was created to enter the server address and port from the HoloLens (See Figure 17a).

The device selected was an Android tablet that ran an application based on the Vuforia API. The database containing the image targets was shared between both devices. At the beginning of the experience, both HoloLens and the tablet track the position of the same image target. The measured coordinates from each device's frame of reference determine the transformation matrix between them. The tablet/external camera must be placed on a fixed position at all times. This is because tracked image targets are always relative to the location of the camera in hand-held devices. On the other hand, the global frame of reference in the HoloLens is initialised at the position of the headset when the application is launched. After

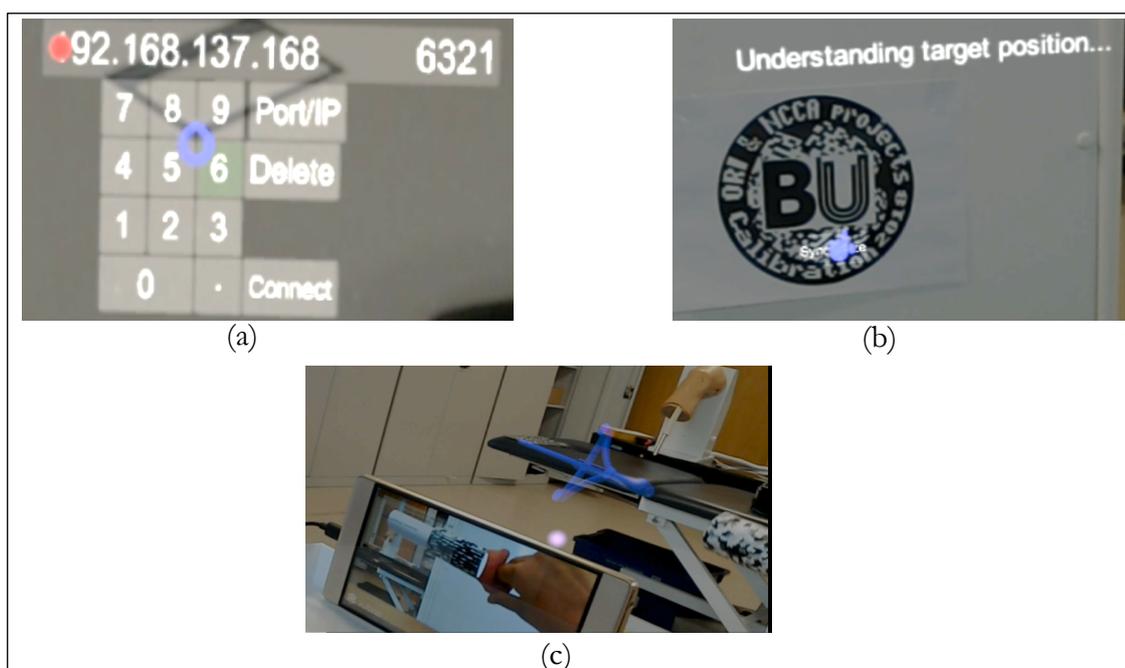


Figure 17 Screenshots taken from the HoloLens first person view. (a) User interface to connect HoloLens to the mobile server. (b) Calibration stage using calibration marker. (c) Hologram projected based on received and corrected coordinates in the HoloLens.

initialisation, this frame of reference remains anchored even after the user movements. After the two devices obtain the position and average orientation of the calibration target, the HoloLens receives the spatial information via Wireless and processes it to calculate the transformation matrix as the difference between both measurements (Figure 16). Once the relative transformation matrix between the two devices is found, the external camera tablet is able to track and send the positions and orientations in real-time to the HoloLens.

3.4 User study.

We conducted a face validation study for the AR simulator for training on INS calibration. Face validation consists of acquiring a subjective opinion about the perceived similarity of the concepts acquired in simulation-based training with those acquired by experience in the real world. For this type of assessment to be valid, both experts and surgeons in training must be included in the group of participants (Schijven and Jakimowicz, 2005).

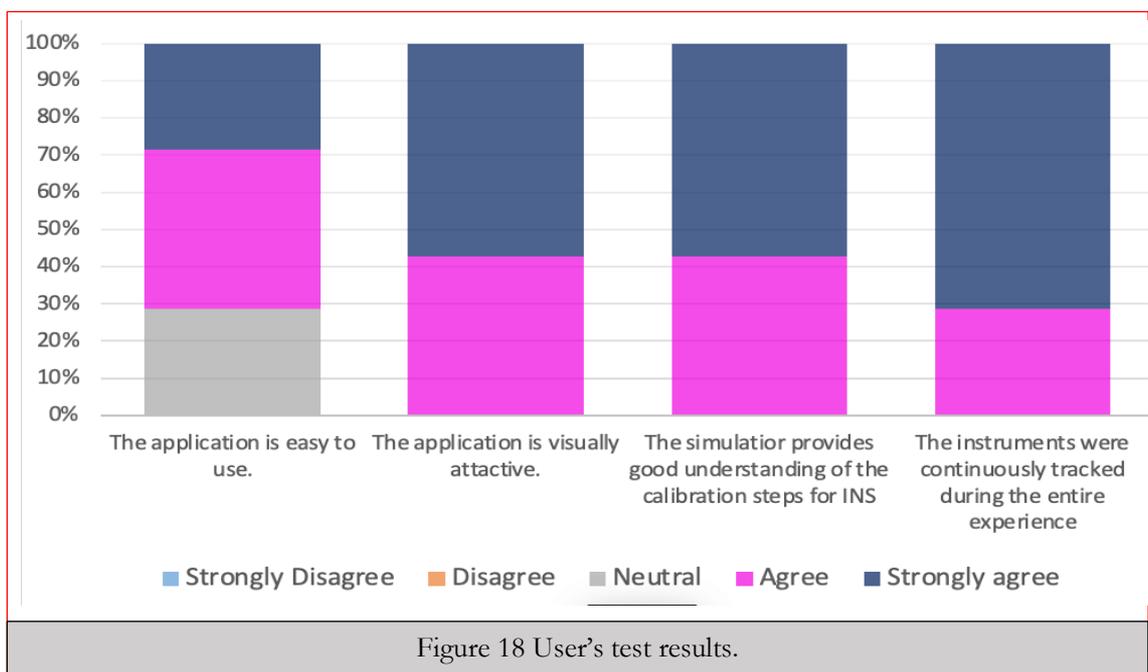
A group of 2 hip surgeons with previous knowledge of INS and 5 orthopaedic surgeon trainees were invited to try the simulator. Before interacting with the simulator, the participants were asked to watch an introductory video that would guide them through the expected steps of the simulation and familiarise them with the headset. To start, users were asked to place a paper sheet with the printed image target on top of a surgical bed fitted to their preferred working height. Once the image target was recognized, it is used as a reference anchor to place the remaining group of holograms forming the surgical scene (Figure 12(a)). The users were able to move this first image target in space until they were comfortable with the position of the surgical scene and fixed its location. Next, the users were required to adjust the orientation of the infrared camera to ensure that its field of view covered the acetabular cavity of the patient. To facilitate this step, the camera's field of view was displayed as a yellow volume (Figure 12(a)). After they were happy with the surgical setup, users had to acquire a total of 27 points displayed one after another and distributed between pelvic anatomical landmarks (Figure 12(b)). This amount of points enabled learning through repetition since the simulator would only allow the user to progress in the experience as a new point was acquired, visual and audio feedback were provided as well. Through this success signal, users associate the absence of registration feedback to a wrong manipulation of the tracked instrument and modify their handling.

3.4.1 Results and discussion

All participants were able to complete all the simulation stages successfully without requiring additional help to the guidance provided inside the app. All participants expressed their approval to the selected headset, and the results from the survey are illustrated in Figure 18. These results validate the medical community's acceptance of mixed reality headsets and their usability as training alternatives for open orthopaedic surgical procedures. While observing the participants interacting, some difficulties with the usability of the headset were present. The Hololens1 does not count with eye-tracking and relies only on the orientation of the users' head to estimate their gaze. This interaction's input is perceived as unnatural for the users. Furthermore, it was a common mistake for the user to place their hands outside the scope of the depth sensors, which prevents the headset from recognizing the hand gestures.

The ease of use for our prototype is related to the estimated future acceptance of similar AR simulations for open surgery simulation. The lack of controllers paired with the headsets and the time required to get used to specific hand gesture can slow down the simulator's learning curve with unrequired cognitive load. These later limitations can be solved by running our simulator on the newest version of the Hololens 2. This new MR display counts with a broader and more natural spectrum of recognized hand gestures, and it is able to track the iris of the user to know the exact point at which the users are looking.

In addition, all of the participants agreed that the simulator allows an understanding of the hand-eye coordination skills needed while calibrating an INS. This shows that the



platform was deemed valuable by the participants, especially experienced surgeons and that the steps simulated correspond to the ones expected to be performed during real INS calibration. When it comes to the visual content, it is evident that surgeons in training value MR experiences that allow them to appreciate the human anatomy without any involved risk to the patient. All participants considered that the visual content was attractive, with 4 of them strongly agreeing with this affirmation. In this question, the simulator's attractiveness referred to its visual appearance and its anatomical accuracy.

Finally, we evaluated the usability of using cylindrical markers with API Vuforia to simulate small hand-held instruments. In 5 out of the 7 cases, the cylindrical marker was tracked with no loss. This result validates the use of Vuforia as long as the size of the marker is adequate for the instrument size and it is visible at all times by the headset's camera.

Overall, these results show that the simulator represents a valuable practice alternative for both surgeons with previous knowledge about intraoperative navigation systems. We have used gaming strategies to create a novel proposal of a functional MR based simulator to help trainee surgeons to understand the principles of imageless navigation systems without high-budget investments. In our AR simulator, the app ends at the beginning of the reaming stage, and no feedback was displayed about the progression of removed bone on this prototype. *A strategy to model the amount of bone removed will be discussed in future chapters.*

3.5 Summary.

In real procedures, the infrared camera (IR) cast rays which are reflected at the surface of some passive spheres attached in a specific position on the surface of a unique mount. To make objects always visible to the camera, this mount must face to a similar direction as the camera view and remain inside the IR camera's tracking volume. This chapter presented an AR simulator that helps surgeons acquire the skills to calibrate an infrared-based navigation system before using them in a THR surgery. The simulator is hosted on a Microsoft HoloLens headset and uses the PTC Vuforia computer vision API to support the tracking of the surgical instruments. This first version of the experience helped us get feedback from orthopaedic surgeons' opinion about both comfort and functionality of the selected hardware and the interaction available in AR devices. A demonstration video from a 1st view camera can be found under the following address:

http://v.youku.com/v_show/id_XMzQwMzMyNDI4MA==.html

Additionally, we developed an alternative communication protocol between an AR headset and an external tracking device. Such protocol is useful for headset devices whose computing capacity must be reserve for graphical computations or enable usage of more advanced and accurate position tracking systems.

Unfortunately, the cost of acquiring a robotic assisting arm can be so high that it makes them unaffordable for many healthcare centres around the world. Therefore we must also provide training opportunities to ensure the correct positioning of the acetabular component. In the following chapters, we will describe the methodology followed to generate a numerical simulation of the reaming procedure in traditional THR surgeries and train a machine learning simulator to predict these results at faster framerates.

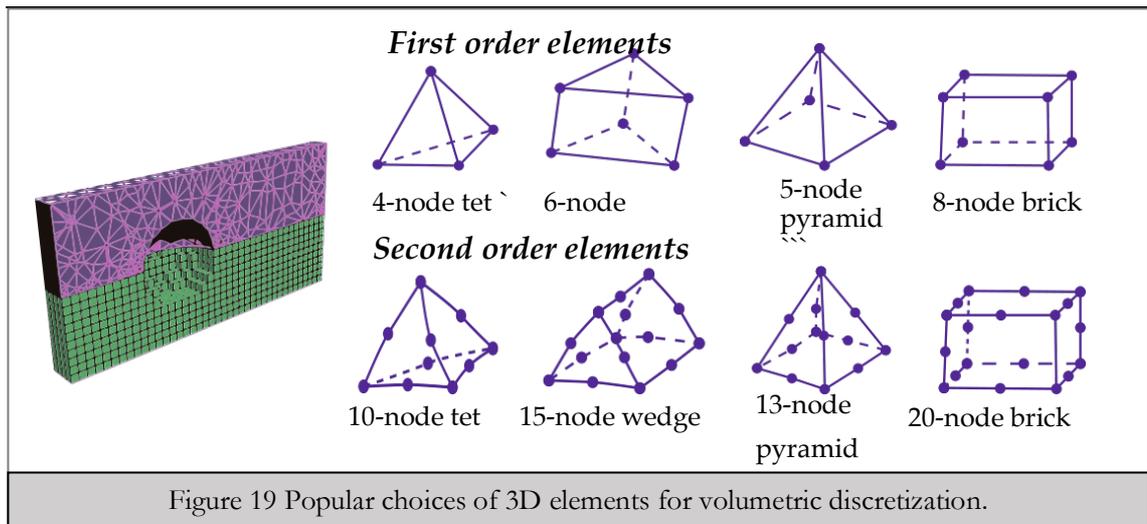
Chapter 4

The Finite Element Method (FEM) for physically-based simulation.

This chapter examines the fundamentals of finite element simulation and its application to analyse resurfacing operations on bone tissue. Additionally, it includes a review of FEM methods used in computer graphics for interactive simulation for surgical procedures and explains its current drawbacks. Finally, we present a literature review of the constitutive material methods used to model cutting and resurfacing operations on bone.

Surgical operations that usually benefit from being represented in a virtual simulation are usually non-reversible ones such as stretching, shearing, tearing, cutting. As part of the computer-simulated experience, 3D meshes can be deformed using geometrical and physically-based methods. Geometrical approaches are computationally efficient as they are applied to surface meshes that can usually be simplified in the form of a sparse adjacency matrix representation. As geometrical deformations only consider spatial relationships between the nodes and faces of the meshes involved, they are less accurate than physically-based. Volumetric meshes are preferred in surgical simulation to shallow ones since they can represent the internal evolution of deformable solids resulting from cutting and shearing operations.

Volumetric meshes are also better at modelling mechanical properties of a deformable object in the continuum domain as well as the effect of the forces exerted on it by the users during the interaction. By including continuum mechanics behaviours, we ensure the



interactive realism of the simulation. This feature is, therefore, vital for designing a successful orthopaedic simulator focused on hand-eye coordination skills. The finite element method is one of the most popular numerical methods to resolve physically-based deformations on complex geometries with non-linear behaviour on volumetric meshes both offline and at interactive framerates.

4.1 The Finite Element Method (FEM).

The Finite Element method is a general discretization of problems from the continuum mechanics domain to analogous mathematical statements (Zienkiewicz and Taylor, 2005). In continuum mechanics, the analysis over a volume is made through integration over a collection of infinitesimal elements conforming the solid. Modelling the mechanical behaviour of complex solid geometries in the infinitesimal domain is computationally intractable even for modern hardware architectures. A FEM analysis overcomes this limitation by subdividing the domain Ω into a finite number of parts (elements) that model the solid's behaviour in a local domain. Consequently, the mechanical behaviour of the solid is governed by each element's shape function and the *constitutive equations*¹ of the material (section 4.1.2.2).

4.1.1 The finite element volumetric mesh.

The solid object is discretised to a volumetric mesh built with one or more type of finite elements. The generated mesh must ensure compatibility in all the volumetric domain Ω . This means that all internal nodes must be connected without discontinuities to ensure that nodal properties are continuous between adjacent elements. The properties estimated on the nodes of the mesh are subsequently interpolated into the element's integration points.

The number of integration points of the elements is usually higher than the number of vertices forming the element, except for the 4-node tetrahedral element, which only has one integration point at its centroid. The location of these integration points can be obtained from the reference documentation of the FEM solver used.

Some of the most common elements available in commercial FEM solvers for 3D solid discretization are illustrated in Figure 19. Elements with triangular faces are usually the most popular choice when discretizing a complex solid. This is because of their easy adaptability to the topology of curvy surfaces. Elements of first-order possess a linear displacement behaviour on their nodes and a strain rate that can be assumed to be constant. Therefore, a fine mesh is needed to obtain an accurate solution in areas where a large deformation is expected.

On the other hand, second-order elements model the displacement in their nodes through a quadratic model. By having a higher number of nodes per element, meshes built with second-order elements require more memory for a given mesh size than first-order memory. However, the accuracy obtained using a mesh with second-order elements is higher than one evaluated on a first-order mesh with the same number of elements. In general, second-order elements are computationally expensive to use and are ideal for pure mechanical bending analysis, while first-order elements perform better when used for contact analysis.

4.1.2 Primary and secondary unknowns in a FEM analysis for a dynamic system under mechanical loads applied.

4.1.2.1 Nodal Displacements and strains.

The primary unknown to be solved under any mechanical FEM analysis is the vector of nodal displacements \mathbf{u} . The values of strain and stresses field are obtained through relationships given by the material properties and these displacements. Both of these fields measure the response of the object to the applied load. Analytically, the displacement \mathbf{u} of each node is found by measuring the difference between the post-deformation position of the vertex \mathbf{x}_i and the initial position of the same vertex \mathbf{x}_0 (3). The methodology used by the FEM solver to obtain the vector of nodal displacements given an applied force will be defined in detail in Section 4.1.3.

$$\mathbf{u} = \mathbf{x}_i - \mathbf{x}_0 \quad (3)$$

The strain field is a concrete measurement of the deformation on each finite element. It is computed by comparing the length of the linking segments of the mesh with their original length. The differential representations of the strains measured on each coordinate axis are defined by equation (4). Added to the changes in their length, the effect of internal forces can distort the faces of the finite elements. These distortions are measured in terms of the angular deformations and are described by the equations in (5) (Spencer, 2004).

$$\varepsilon_{xx} = \frac{\partial u}{\partial x} \quad \varepsilon_{yy} = \frac{\partial v}{\partial y} \quad \varepsilon_{zz} = \frac{\partial w}{\partial z} \quad (4)$$

$$\gamma_{xy} = \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} \quad \gamma_{xz} = \frac{\partial u}{\partial z} + \frac{\partial w}{\partial x} \quad \gamma_{yz} = \frac{\partial v}{\partial z} + \frac{\partial w}{\partial y} \quad (5)$$

The success of the numerical simulation of a mechanical problem depends on finding the unique strain field corresponding to a given deformation vector. Therefore, we must incorporate compatibility equations to constrain the solution to a unique strains field given a specific displacement. The system of compatibility equations is built with 6 differential conditions; the first three are obtained by derivating over the strains defined in equations (4) and (5); and the remaining equations (9) – (11) are found by derivating twice the normal stresses against their orthogonal axes (Zienkiewicz and Taylor, 2005).

$$\frac{\partial^2 \gamma_{xy}}{\partial x \partial y} = \frac{\partial^2 \varepsilon_x}{\partial y^2} + \frac{\partial^2 \varepsilon_y}{\partial x^2} \quad (6)$$

$$\frac{\partial^2 \gamma_{xz}}{\partial x \partial z} = \frac{\partial^2 \varepsilon_x}{\partial z^2} + \frac{\partial^2 \varepsilon_z}{\partial x^2} \quad (7)$$

$$\frac{\partial^2 \gamma_{zy}}{\partial z \partial y} = \frac{\partial^2 \varepsilon_y}{\partial z^2} + \frac{\partial^2 \varepsilon_z}{\partial y^2} \quad (8)$$

$$\frac{\partial^2 \varepsilon_{xx}}{\partial z \partial y} = \frac{\partial}{\partial x} \left(-\frac{\partial \gamma_{yz}}{\partial x} + \frac{\partial \varepsilon_{xx}}{\partial y} + \frac{\partial \gamma_{xy}}{\partial z} \right) \quad (9)$$

$$\frac{\partial^2 \varepsilon_{yy}}{\partial x \partial z} = \frac{\partial}{\partial y} \left(\frac{\partial \gamma_{yz}}{\partial x} - \frac{\partial \varepsilon_{xx}}{\partial y} + \frac{\partial \gamma_{xy}}{\partial z} \right) \quad (10)$$

$$\frac{\partial^2 \varepsilon_{zz}}{\partial x \partial y} = \frac{\partial}{\partial z} \left(\frac{\partial \gamma_{yz}}{\partial x} + \frac{\partial \varepsilon_{xx}}{\partial y} - \frac{\partial \gamma_{xy}}{\partial z} \right) \quad (11)$$

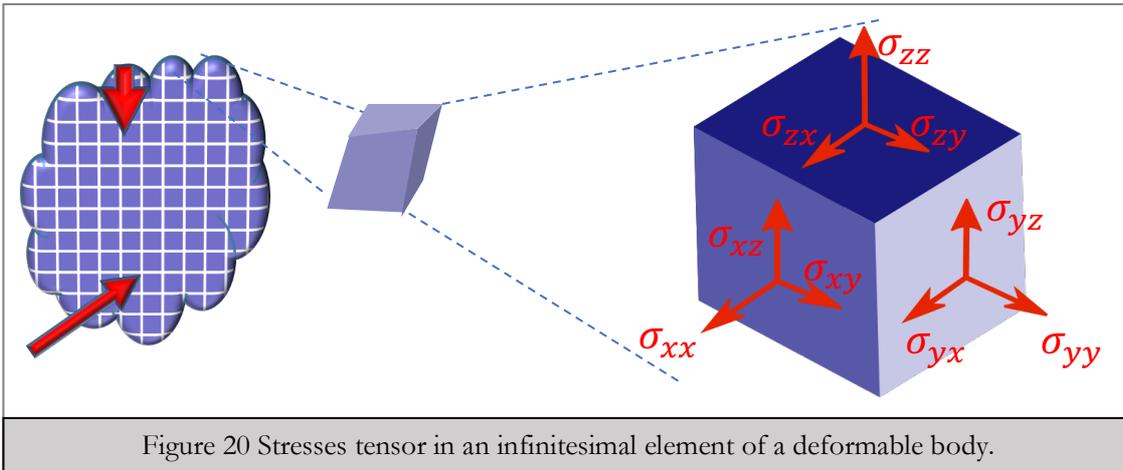


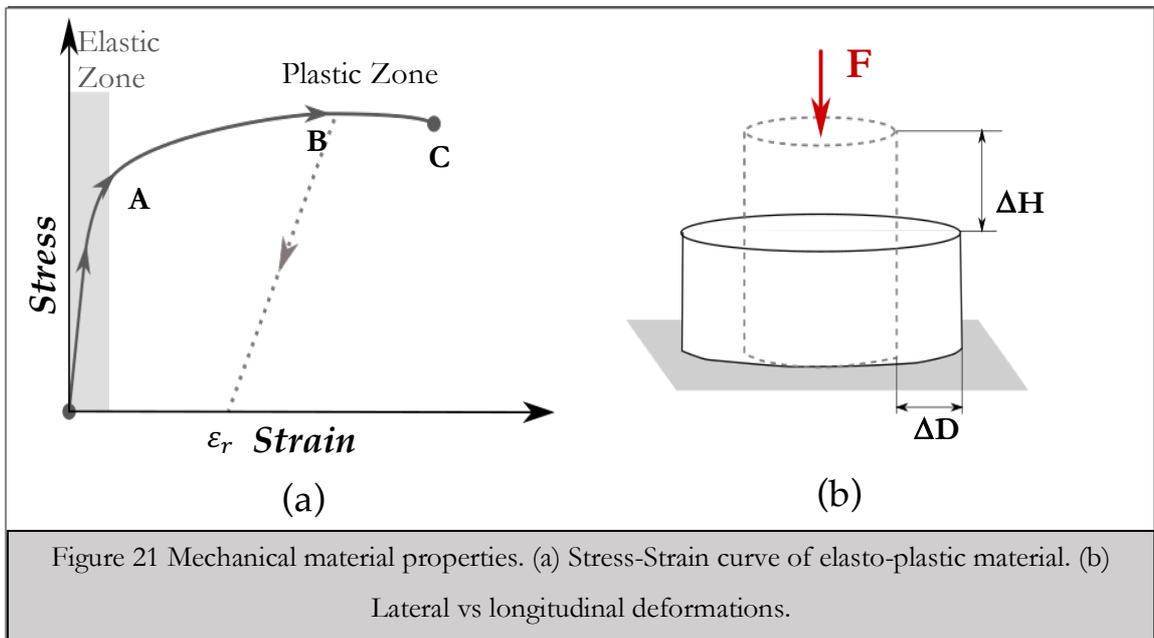
Figure 20 Stresses tensor in an infinitesimal element of a deformable body.

4.1.2.2 Stresses

Once the value of the strain on each segment or edge of the mesh is found, the FEM analysis must find their corresponding stresses. To understand the nomenclature of the vector of stresses, it is better to represent them on an infinitesimal cube with its faces aligned to the orthogonal coordinate axes (Figure 20). Similar to the nomenclature of the strains, the naming convention for each stress obeys the following rules. The first subindex corresponds to the axis perpendicular to the plane or face where the stress is acting. The second subindex determines the direction of the stress in the euclidean coordinate system. The equilibrium conditions hold in the infinitesimal cube when the domain Ω is constrained, and there is no absolute movement of the solid. Therefore, $\sigma_{yz} = \sigma_{zy}$; $\sigma_{xz} = \sigma_{zx}$ and $\sigma_{xy} = \sigma_{yx}$.

Even though we have used an infinitesimal cube in this section to represent the orientation of the stresses inside a solid, this does not mean that the stresses' naming convention is only valid for finite elements of the brick type. The cube in Figure 20 is infinitesimally small, and all the stresses can be evaluated on any node of a finite element chosen. The science of solid mechanics distinguishes the stresses present in a material into two types. On the one hand, the normal stresses σ_{ii} , which are responsible of any volumetric changes by producing strains ϵ_{ii} ; on the other hand, the shear stresses $\sigma_{i,j}$; $i \neq j$, which are responsible for distortion, shear strain and mechanical failure through element separation (Spencer, 2004).

Before we explain how to obtain the stresses in a domain Ω given a field of strains, we must introduce the concept of the *mechanical deformation zone*. Depending on the amount of stress applied and the material behaviour, a solid can be inside its elastic or plastic



deformation zone. This is, the strains of ductile materials with predominantly elastic behaviour only move inside an elastic zone before reaching the fracture point (point C in Figure 21(a)). On the opposite, plastic materials can have a very short elastic zone followed by a main long plastic deformation zone. Figure 21(a) shows the stress-strain curve for an elastic-plastic material. When the solid is initially stressed, its strains and stresses behave following a linear model. All strains return to the initial value of 0 if all external loads are removed in the elastic region. The boundary of the elastic region (point A in Figure 21(a)) is defined as the value of the material's yield stress (σ_0). After this point, even if the external loads are discharged before fracture (point B in Figure 21(a)), a residual strain ϵ_r will remain permanently in the material (Spencer, 2004).

The material constitutive equations describe the behaviour of the material's strain-stress curve for all points in a continuum solid. The Hooke's Law (Hooke, 1978) is used as the constitutive equation for the elastic zone of any material by using a material-specific property called the Young modulus $E = \frac{\sigma}{\epsilon}$ as the elasticity coefficient. Additionally, the relationship between the normal and shear stresses is set by the modulus of shear elasticity G defined in (12) (Polyanin and Chernoutsan, 2010). In this equation, ν is another scalar elastic property of the material called the Poisson's module. The Poissons module is computed as the relationship between the longitudinal and the lateral strain after an external load is exerted on a deformable body (Figure 21(b)).

$$G = \frac{E}{2(1 + \nu)} \quad (12)$$

$$\nu = -\frac{\epsilon_{lateral}}{\epsilon_{axial}} \quad (13)$$

By combining the equation (12), (4) and (5) with Hooke's law, we can assemble a system of equations to describe the stresses in terms of the strains inside a linear elastic zone (15). In this system, the diagonal and symmetric matrix multiplies the vector of strains is known as the *constitutive stiffness matrix* \mathbf{D} (Zienkiewicz and Taylor, 2005).

$$K = \frac{E}{(1 + \nu)(1 - 2\nu)} \quad (14)$$

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{xy} \\ \sigma_{xz} \\ \sigma_{yz} \end{bmatrix} = K \begin{bmatrix} (1 - \nu) & \nu & \nu & 0 & 0 & 0 \\ \nu & (1 - \nu) & \nu & 0 & 0 & 0 \\ \nu & \nu & (1 - \nu) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5(1 - 2\nu) & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.5(1 - 2\nu) & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.5(1 - 2\nu) \end{bmatrix} \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \gamma_{xy} \\ \gamma_{xz} \\ \gamma_{yz} \end{bmatrix} \quad (15)$$

Modelling the plastic region is more complicated as the plastic behaviour of every material is different. The relationships between strain and stress in the plastic region are found after obtaining the strain-stress curve of the material experimentally. In Section 4.3.2 we will describe the plastic behaviour of the bone, which is the deformed material of interest in this thesis.

To solve the system of differential equations previously introduced, the FEM follows the principle of "virtual work". This model establishes that a body is in equilibrium if and only if the internal and external virtual works are equal (18). This constraint creates a direct relationship between the effect of the internal stresses inside the mesh with the effect of the external loading conditions at the external elements of the domain Ω . On the one hand, the external virtual work (16) is defined as the deformation energy made by all the external loading conditions such as weight \mathbf{b} and concentrated nodal forces \mathbf{f} . On the other hand, the internal virtual work corresponds to the energy produced by the internal forces and their triggered nodal displacement $\delta\mathbf{u}$ (17) (Katsikadelis, 2016).

$$\delta W^{ext} = \int_{\Omega} \delta\mathbf{u}^T \mathbf{b} d\Omega + \sum_{n=1}^r \delta\mathbf{u}(x_n)^T \mathbf{f}(x_n) \quad (16)$$

$$\delta W^{int} = \int_{\Omega} (\tilde{\nabla} \delta \mathbf{u})^T \boldsymbol{\sigma} d\Omega \quad (17)$$

$$\delta W^{ext} = \delta W^{int} \quad (18)$$

During the solution of the system, the solver uses the boundary conditions of known displacements. In a static analysis, these boundary conditions correspond to encastre conditions¹ applied to some nodes of the mesh.

4.1.3 FEM dynamic analysis.

A dynamic analysis of a physical problem can be performed in FEM using a Standard or an Explicit solver. Both of these methods require establishing the mass matrix of the solid domain to estimate its inertia. However, these two methods differ in their approach to discretising the time dimension and solving differential equations previously introduced. The standard method holds a dynamic equilibrium condition (19) based on the instant acceleration and velocities of the node's displacements (Zienkiewicz, OC and Parekh, 1970). In (19), the variables ρ and μ correspond to the density of the material and damping coefficient at the contact poits, respectively.

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{C}\dot{\mathbf{x}} + \mathbf{K}\mathbf{x} + \mathbf{f} = \mathbf{0} \quad (19)$$

$$M_{ij}^{(e)} = \int_{\Omega} N_i \rho N_j d\Omega \quad C_{ij}^{(e)} = \int_{\Omega} N_i \mu N_j d\Omega \quad (20)$$

Alternatively, explicit methods have proven to be more computationally efficient in solving dynamic analysis where a fracture in the material is expected (Guo and Lv, 2010). The explicit FEM (E-FEM) uses the explicit central time integration rule to satisfy the dynamic equilibrium condition on each time increment t (Smith, 2016). At the beginning of each time increment, the value of the nodal acceleration must be obtained first. To do so, a dynamic equilibrium expression is manipulated into the form presented in equation (21). In (21) \mathbf{M} corresponds to the Mass Matrix, \mathbf{P} to the nodal forces applied on the external nodes of the mesh and \mathbf{I} the internal forces deforming the internal topology of the mesh. Since the matrix \mathbf{M} is almost always diagonal or sparse, obtaining the value of nodal accelerations $\ddot{\mathbf{u}}_t$ is computationally inexpensive. Once the acceleration of each vertex of the volumetric FEM mesh is found, the velocities (22) and displacements (23) can be computed by advancing

them “explicitly” through time. In other words, the displacements at the end of the time increment are obtained exclusively from position, acceleration and velocities and the beginning of it (Smith, 2016).

$$\ddot{\mathbf{u}}_t = \mathbf{M}^{-1}(\mathbf{P}_t - \mathbf{I}_t) \quad (21)$$

$$\dot{\mathbf{u}}_{t+\frac{\Delta t}{2}} = \dot{\mathbf{u}}_{t-\frac{\Delta t}{2}} + \frac{(\Delta t_{t+\Delta t} + \Delta t_t)}{2} \ddot{\mathbf{u}}_t \quad (22)$$

$$\mathbf{u}_{t+\Delta t} = \mathbf{u}_t + \Delta t_{t+\Delta t} \dot{\mathbf{u}}_{t+\frac{\Delta t}{2}} \quad (23)$$

Once the displacements on the mesh's nodes are computed, the solver proceeds to calculate the strain increments and their corresponding stresses. Finally, the solver assembles the matrix of internal forces \mathbf{I} based on the stresses field in the internal elements of the mesh.

The estimated time increment used by the ABAQUS CAE solver is described in equation (24) (Smith, 2016). Where L_{min} is the size of the smallest element from the mesh, and c_d is a function of the material's Young modulus and Poisson's ratio. The smaller the size of the elements in the mesh, the more computationally expensive the system is, and the bigger is the amount of memory required to store their corresponding stresses, strain and damage status on each time increment of the dynamic analysis.

$$\Delta t = \frac{L_{min}}{c_d(E, \nu)} \quad (24)$$

4.2 Realtime FEM simulation.

As previously mentioned, the simulation of non-reversible deformations can benefit greatly from computer generated training alternatives. This usage of VR/AR is ideal for applications such as surgical operations where the topology of the tissue is usually expected to change. Some offline approaches rely on element deletion, element duplication¹⁷ or adaptative remeshing to represent these discontinuities. Nevertheless, these methods can deteriorate the quality of the mesh and increase the probability of achieving bad quality

¹⁷ Followed by separate mesh refining.

elements¹⁸ on each remeshing iteration. Furthermore, adaptative remeshing around the discontinuities increases the size of the tetrahedral mesh exponentially, making it potentially untractable for further processing. In this section, we explore some techniques found in the literature to improve the performance of finite element simulations to make them suitable for virtual reality applications.

4.2.1 *Precomputed FEM*

These methods involve precomputations of complete or partial results from FEM simulations to ease the calculation required at runtime. These type of algorithms exchange the fine mesh material properties with a new equivalent database of materials that would produce the same variation on the elastic energy on a coarse mesh as the high-resolution one. Chen et al.(2015a) built two data sets for modelling the constitutive equations of hyperplastic materials. One represents the non-linear mechanical properties of a refined mesh, and another set containing the properties of a material that would generate similar results on a coarser mesh. These allowed them to compute the deformations on a mesh with simpler embedded material and find the fine displacements based on the mapping between the 2 sets. The drawback of this approach is that the simplified material model can not be re-used for different geometries and must be regenerated on different simulation instances. A different approach was developed by Banihani et al.(2013), who employed orthogonal decomposition to reduce the number of variables composing the model significantly. By doing so, the complexity of the model was reduced, and its robustness increased. This technique acted as a regularization function, making the differential equations more computationally stable, which adds a big benefit for solving energy-based models. However, the orthogonal decomposition can only be computed and be loaded into runtime memory for a finite number of cases. Therefore, the simulation error increases rapidly as the user deviates from precomputed deformations (Malukhin and Ehmann, 2018).

4.2.2 *Mass-Spring approaches*

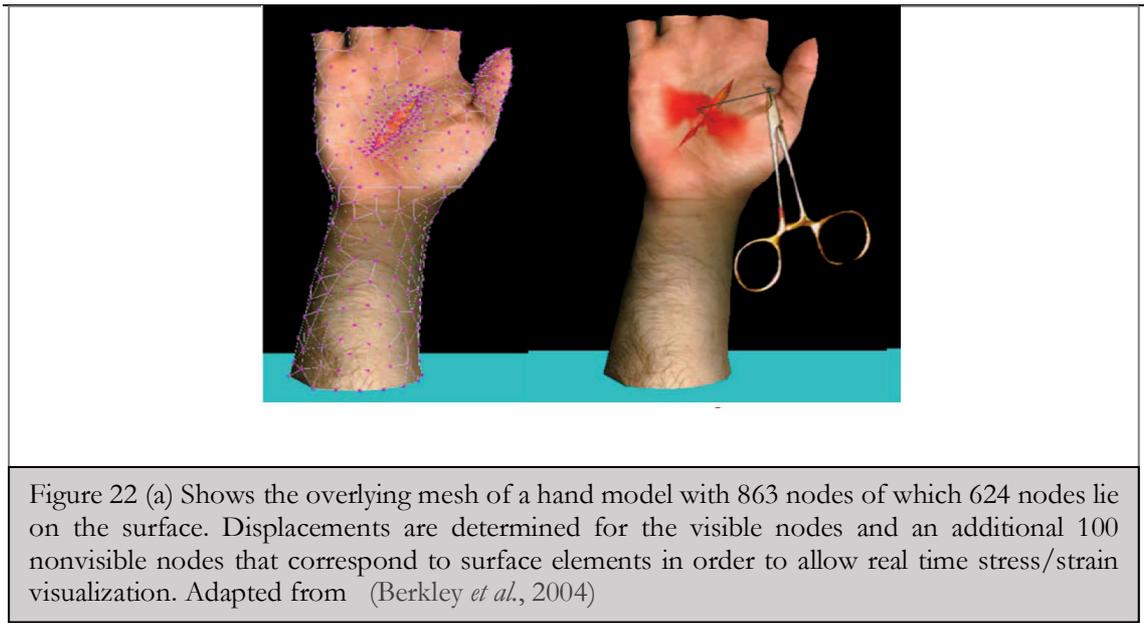
These models are popular for simulating stretching deformation on thin elastic tissue and small cuts (Wu *et al.*, 2015). Their simplicity reduces the computational resources required and ensures a stable and fast framerate during the interaction. The main idea of this model is that elements on the volumetric mesh are represented as a system of masses that

¹⁸ Bad quality elements are elements with one of its edges significantly shorter than the rest of them forming a very sharp internal angle

are connected to each other with a set of springs. Nevertheless, the accuracy of mass-spring models is much lower than a finite element simulation when applied to thicker geometries. This is because the mass-spring model is unable to consider the effect of mechanical non-linearities such as stiffness, uneven mass distribution and damping coefficients. Also, the nature of the model allows the springs to stretch or contract almost in ways that the total model is not preserved, which decreases realism (Wu *et al.*, 2015). Some authors have attempted to use the fast convergence of mass-spring models with targeted modifications to improve their accuracy. A more complex model was later developed by Pan *et al.* (2015) to simulate soft tissue dissection using a multi-layered mass-spring system for training on preparation for resection of rectal cancer. This approach is only appropriate for modelling deformation on thin surfaces. The authors included haptic interaction with fat and bowel membrane tissues. The content was arranged into a system of 4 layers connected according to their adjacency. Each layer was modelled with different mechanical properties allowing the model to consider non-linearities and propagate their effect to all tissue types. However, this model is only valid for thin elastic tissue deformation that can be discretized in a finite number of layers. Therefore, it is not possible to implement a similar approach to simulate resurfacing on volumes of bone.

4.2.3 *Linear FEM*

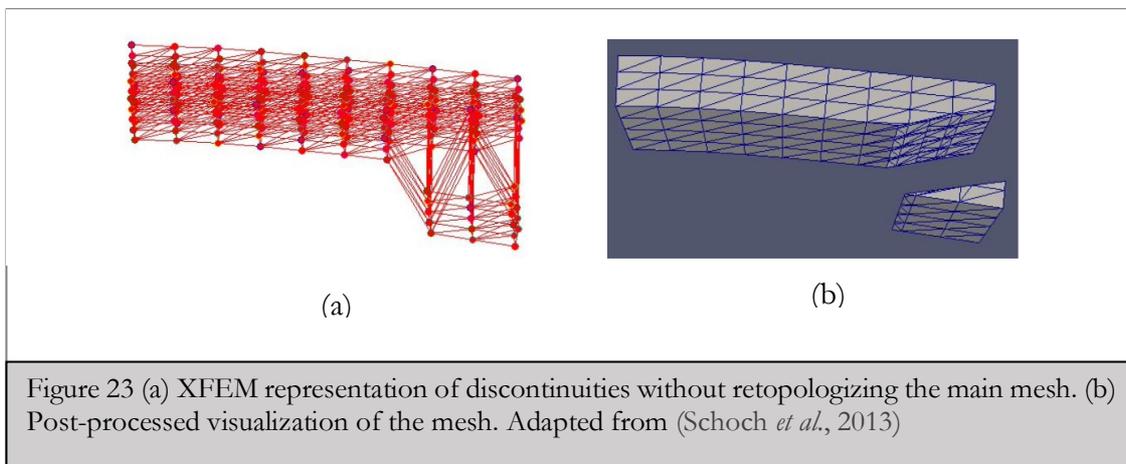
As the name suggests, linear FEM simulations implement a linear stiffness matrix to find the values of nodal deformations according to the method explained in Section 4.1.3. This ensures their computational efficiency and stability on moderate-size meshes. In an Augmented Reality simulator developed for training on suturing (Berkley *et al.*, 2004), linear FEM simulations were used to represent the binding effect of the suture in a hand-cut (Figure 22). The FEM simulation was aligned in AR using ArUco markers. In an attempt to include non-linearities while still keeping the fast performance, Cakir *et al.* (2009) used the Stiffness-Wrapping method to increase the realism of their simulator for cutting on soft-



tissue. Throughout this method, the deformed positions of each node are rotated back to an undeformed frame of reference. This allows the system to calculate a hyper-elastic force using an undeformed stiffness matrix and rotating the result back to the deformed frame of reference. Although, a clear drawback from this method is that extra memory must be reserved for storing the definition of the rotational frames of reference, which are created and updated on each rendered frame. Xie *et al.* (2020) included a Kalman filtering identification process to decrease the size of the deformation system given local measurements of displacement. Their method was applied for soft tissue deformation and was able to achieve realtime performance while maintaining values of nodal deformation similar to traditional FEM. Similar to spring-mass models, the main drawback from linear FEM is that their accuracy is only maintained on small deformation magnitudes (up to 10% strain). This means that objects can increase in volume unnaturally under conditions that trigger large rotational deformations (Malukhin and Ehmann, 2018).

4.2.4 Corotational FEM

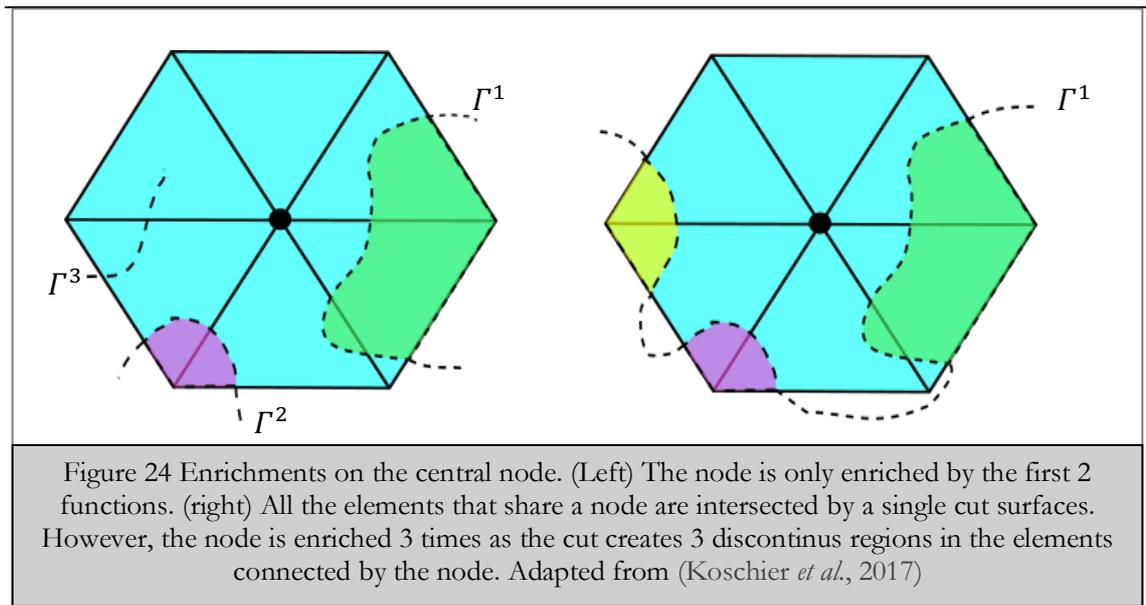
Corotational FEM has been popular for simulation of liver tissue manipulation, given its hyper-elastic nature (Schoch *et al.*, 2013; Bui *et al.*, 2019). This method improves the accuracy of a linear model when the strain values are expected to be higher than 10% while maintaining low complexity in the differential equations that describe it. First, a global stiffness matrix is assigned to the system and used to compute a global estimate of the body's displacement and rotation. Additionally, a local frame of reference is assigned to each node to find their internal vectors of "pure-strain". Therefore, the rotational factor of the



deformation is isolated before the internal forces have to be computed (Bui *et al.*, 2019). Both frames of reference are mapped through linear relationships, which ensures that the interactive framerates can be maintained while the accuracy is increased compared with a single linear system. Most recently, corotational FEM has been used for real-time simulation of needle insertion into brain and liver tissue (Bui *et al.*, 2019). In this simulator, the elements of the mesh were labelled into 3 categories. Namely, internal, splitted and external. Different local matrices with different mechanical properties were assigned to each one of these categories. Although the deformations computed by the model had acceptable accuracy for their application, the existence of a global stiffness matrix added to a progressively growing number of local ones required a significant amount of memory. Therefore, it is strongly recommended to mind the meshes' sizes when using corotational FEM (Bui *et al.*, 2019). Another major drawback of corotational FEM is the high computational cost of inverting the global stiffness matrix while handling the high number of degrees of freedom coming from the local system. This is because the global and local stiffness matrices are not additive.

4.2.5 XFEM

The Extended Finite Element Method is an alternative approach for modelling discontinuities in a volumetric without changing its topology (Figure 23). Instead of having to remesh the solid geometry once an element is separated, the stiffness matrix of the elements is enriched with additional DOFs that can represent the fracture or separation. Since the object remains as a continuum domain, the XFEM model remains numerically stable for large deformation values. However, there is a computational cost added by the presence of the enrichment functions and the additional degrees of freedom. Quesada *et al.* (2016) used the XFEM method to obtain the internal forces on the cornea and liver tissue



given the interaction with a scalpel at frequencies compatible with haptic devices¹⁹. The authors enhanced the mechanical properties of each tetrahedral elements using the cracking node method (Song and Belytschko, 2009). This method parametrizes a cut by using nodally-centred cracks approximations that extend up to the boundary of each tetrahedral element. Therefore, the true geometry of a curved crack is approximated by a collection of straight crack segments passing through the nodes. This reduces the accuracy of the crack propagation but improves the real-time response.

Koschier *et al.* (2017) combined XFEM methods with an implicit dynamic integration to simulate fracture of deformable objects using complex cutting surfaces. One of the highlights of their approach is the ability to simulate finely structured cuts, even when the mesh is made from a coarse tetrahedral discretization. The cutting surfaces are modelled as explicit triangular meshes, and only the nodes of fully intersected elements are enriched (Figure 24). Furthermore, different cutting surfaces must be independent and not intersect, as overlapping cuts are not supported. It is worth noting that the method developed by Koschier *et al.* (2017) was used to generate off-line animations instead of interactive simulations. This means that the main limitation of their method is that it is unable to simulate progressive cuts advancing within a single tetrahedron.

The corotational model can also be enriched with XFEM functions. In order to simulate the surgical cutting of soft tissue, Schoch *et al.* (2013) combined corotational

¹⁹ 500-1000 Hz

methods with implicit time integration techniques and XFEM enrichment to achieve a stable and simplified representation of arbitrary cuts. Furthermore, the corotational fraction of the model computed the complex portion of the deformation before the time integration step and achieved an accurate representation of large deformations in soft tissue. The results presented in Schoch *et al* (2013) were obtained by manipulating meshes with less than 1000 tetrahedral elements. As mentioned in previous sections, the computational requirements from both corotational and XFEM models increase primarily as the number of elements increase. Therefore, the performance of this combined method is expected to be sensitive to the level of refinement of the mesh, and it is likely to slow down when a higher level of refinement is desired.

Chitalu *et al.* (2020) combined XFEM with quasi-static linear elastic fracture mechanics and a high-resolution crack propagation scheme mesh fracture to simulate crack propagation on brittle materials. Their approach requires a predefined cutting surface to initiate the crack as one of its inputs. Realtime performance is achieved by reducing the number of mesh edges, and boundary faces required to describe the fracture by 20% compared with the state of the art methods (Chitalu *et al.*, 2020). However, the approximations performed by the crack propagation scheme reduced the accuracy of the simulation by underestimating the stress intensity factors. Therefore, the authors recommend using their simulation method only on applications that do not require numerical fidelity.

4.3 Bone reaming – tissue deformation

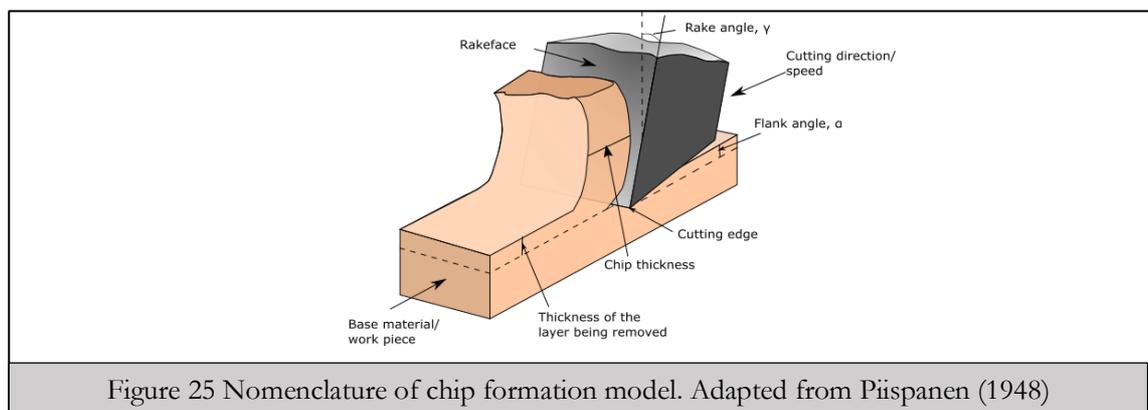
This section will describe the current approaches to FEM simulation around the phenomena of bone cutting. Firstly, the perspective of the manufacturing engineering science of the fracture and plastic yielding involved in material removal will be described. Afterwards, the chapter presents an introduction to the current mathematical models used to describe the behaviour of bone tissue as a material for FEM simulation.

4.3.1 Finite element method and chip formation in machining operations of bone.

According to solid mechanics science, after a deformable body is under a specific loading condition, each element of a volumetric grid suffers a change in its stresses/strain tensor. The FEM solves a system of partial differential equations according to the applied loads and the material behaviour encoded in its stiffness matrix. By solving the differential system, the strains and stresses tensors are updated while holding the condition of forces on

equilibrium for each element (Zienkiewicz and Taylor, 2005).

In order to create a physically-based model of removed material during the reaming stage, we must take a look at the effect of each tool's *tooth* on the surface of an object. Once each cutting edge has penetrated the material at a defined cutting depth, a chip is formed under plasticity in front of the tool's tip and forced until fracture. The deformation occurs as the cutting edges move relative to the base material and compress the material in front of it. Figure 25 illustrates the parameters involved in orthogonal cutting, which is one of the most commonly researched material removal processes. In Figure 25, the rake and flank angle are geometrical design parameters of the tool, usually optimised to maximise the removal rate while minimising the reacting friction. Oblique cutting, such as acetabular reaming, differs from orthogonal cutting (Figure 25) as the relative movement of the base material block is oblique to the tool's cutting edges. Guo and Lv (2010) and Llanos et al. (2009) presented examples for successful FEM models for oblique cutting common engineering materials. They established failure criteria that included both maximum values for shear and normal stresses, and the effect of the friction given the normal stress values between the tool's forces and the forming chip.



4.3.2 FEM models for bone cutting operations

Modelling the amount of bone removed depends on several factors such as the force with which the surgeon supports the reamer on the acetabulum, the motor's torque, the tool's rotational speed, its orientation/position relative to acetabulum shape, and the area of surfaces in contact. These previous variables are not static in practice and can vary

significantly even while performing the same surgery. The bone tissue also has unique properties that differentiate its fracture and chip formation from other popular researched materials (e.g. steel). At a macroscopic level, bones are divided into two main tissue layers (See Figure 26) that differ from each other in terms of density, strength and hardness. These are the cortical bone, made of hard and compact tissue that takes a high percentage of the biomechanical loads (Marco *et al.*, 2015), and the cancellous bone, which is made of thin lighter tissue creating a net-like matrix (Liao and Axinte, 2016), and reducing the bone's weight. Its structure is not isotropic at a microscopic level, meaning that its stiffness and values of plastic yielding stresses are not homogeneous over all the material orientation.

Regarding the chip formation, Mitsuishi *et al.* (2004) described how during machining operations on the cortical bone, the initial crack formation starts at the tool's cutting edge and follows its direction in continuous chip formation until separation. On the contrary, during cancellous bone cutting, the crack begins below the cutting edge, and it fractures in blocks underneath its moving path, leaving a wake of cutting which for some applications can make the depth of cutting of cancellous bone hard to control. In a review published by Marco *et al.* (2015), the authors report that most of the research done on bone modelling focus on operations in orthogonal cortical bone and therefore, more research is required around the cutting operations on the cancellous tissue.

Tu *et al.* (2013) proposed a simple elastic-plastic model for drilling on both cortical and cancellous bone. The model was validated with data acquired from experiments performed on Sawbones^{®20}, and therefore the mechanical properties for the numerical model

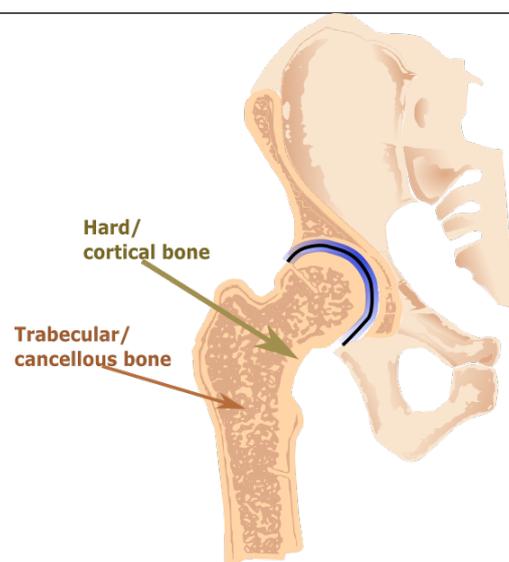


Figure 26 Bone tissue layer in the hip joint.

were provided by the Sawbones manufacturer. The constitutive equations of the material followed the characteristic of an elastic-plastic model. The computed fracture properties of the bone were determined using the *shear damage* model predefined in the FEM solver ABAQUS CAE®. Later, Wang et al. (2016) implemented the same simplifications on an orthogonal cutting 2D FEM simulation with ultrasonic vibrations applied to the workpiece in cancellous bone only. These two articles were the only relevant identified mathematical model for cancellous bone FEM simulation. A validation study using experimental data in Wang et al. (2016) confirmed that the strain predicted matches the measured data for orthogonal cutting cases. Therefore, it was concluded that the elastic-plastic isotropic model with shear damage fracture properties is a good representation of the mechanical behaviour of cancellous bone.

In orthopaedics, machining of cortical bone is usually more challenging than trabecular/cancellous bone, as higher cutting forces are required, and there is a higher risk of heat generation leading to thermal-induced necrosis. Most of the different models available in the literature that describe the chip formation of cortical bone focus on predicting cutting forces to evaluate the tool's design and tool's efficiency and estimate the heat generation produced by the friction during the bone machining operations. This review focuses on mathematical models developed for drilling and orthogonal cutting operations, as these are the closest ones to acetabular reaming in terms of material's induced failure and chip formation.

The most popular model found in the literature to describe the behaviour of bone-cutting was the Johnson-Cook (JC) model (Hage and Hamade, 2013; Baro and Deoghare, 2018). This model describes the plastic behaviour of the material after it has passed its elastic zone. It is usually selected due to its simplicity, fast convergence and its easy implementation inside several engineering FEM software. The JC model, described by (26), considers the effect of strain/work hardening and thermal softening $\Theta(T)$ in the flow of equivalent plastic stress given the values of the average plastic equivalent strain (PEEQ) $\bar{\epsilon}^P$ (25) on each finite element. To compute the value of average stress $\bar{\sigma}$ given a plastic strain $\bar{\epsilon}^P$, the model in Equation (26) takes the average plastic strain $\bar{\epsilon}^P$ and the plastic deformation rate $\dot{\epsilon}^P$ as inputs. Following the numerical update in the values of of the stresses, a new condition for fracture is established by finding the strain to fracture $\bar{\epsilon}_f^P$. This fracture condition is a function of the mean of the three principal stresses $\bar{\sigma}$ and the Von Mises (VM) stress σ_M . Failure occurs

on an element when $\sum \frac{\Delta \varepsilon}{\bar{\varepsilon}_f^{pl}} = 1$ (28), where $\Delta \varepsilon$ is the increment on the equivalent plastic strain for each integration step and $\bar{\varepsilon}_f^p$ is the ultimate strain that triggers fracture.

$$\bar{\varepsilon}^p = \frac{\sqrt{2}}{3} \left[(\varepsilon_{xx}^p - \varepsilon_{yy}^p)^2 + (\varepsilon_{yy}^p - \varepsilon_{zz}^p)^2 + (\varepsilon_{zz}^p - \varepsilon_{xx}^p)^2 + 6(\gamma_{xy}^p)^2 + 6(\gamma_{xz}^p)^2 + 6(\gamma_{yz}^p)^2 \right]^{1/2} \quad (25)$$

$$\bar{\sigma}(\bar{\varepsilon}^p, \dot{\varepsilon}^p, T) = [A + B(\bar{\varepsilon}^p)^n] \left[1 + C \ln \left(\frac{\dot{\varepsilon}^p}{\dot{\varepsilon}_0} \right) \right] [1 - (\Theta(T))^m] \quad (26)$$

$$\bar{\varepsilon}_f^{pl} = \left[d_1 + d_2 \exp \left(d_3 \frac{\bar{\sigma}}{\sigma_M} \right) \right] \left[1 + d_4 \ln \left(\frac{\bar{\varepsilon}_f^{pl}}{\dot{\varepsilon}_0} \right) \right] (1 + d_5 \Theta(T)) \quad (27)$$

$$\sum \frac{\Delta \varepsilon}{\bar{\varepsilon}_f^{pl}} = 1 \quad (28)$$

An evident limitation of the JC hardening law is that this model was proposed originally for ductile chip formation, while bones have a brittle nature. The values of the constants A , B , C and n were found through experimentation to fit the model to the mechanical behaviour of bone under machining operations with results validated with experimental data (Table 4). By using these values, equation (26) describes a strain *hardening* model of the chip formation and fragmentation for a heterogeneous material equivalent to the bone. The values of the obtained cutting force were validated with an experiment carried on a fresh bovine femur. This bone was selected for experimental validation because of its similarities with the human's bone during fracture (Vashishth *et al.*, 2000). Alam *et al.* (2009) also proposed to omit the thermal softening term $\Theta(T)$, which describes a relationship between the working and the melting temperature of the material. The reason being that the temperatures reached in theatre for bone machining operations usually stay below 45°C in order to avoid thermal-induced necrosis²¹.

²¹ The death of the cells conforming a tissue due to high temperatures.

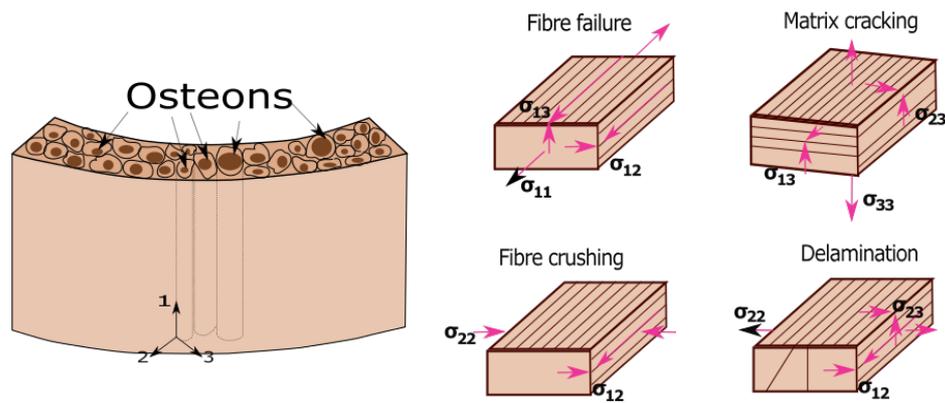


Figure 27 Scheme of the cortical bone structure. Figure adapted from Santiuste et al. (2014) and Hou et al. (2000)

A [Pa]	B[Pa]	n	m	C	ϵ_0
50e6	101e6	0.08	0.02	0.03	0.001

Table 4 Johnson-Cook Plastic Hardening properties for cortical bone.

Table adapted from (Baro and Deoghare, 2018).

Given that in an anisotropic material, the strengths in the orthogonal and longitudinal direction of the osteons can reach a relationship of 3:2 in value. An anisotropic approach would be more aligned with the exact composition of the cortical bone, as mentioned in Section 4.3.1. A successful thermomechanical model of the anisotropic behaviour of cortical bone was proposed by Santiuste et al. (2014). The authors assumed an analogy between bone and a fibre reinforced material, given the bone’s microstructure. Figure 27 illustrates how the cortical bone is composed of a collagen matrix filled with cylindric shaped osteons.

To evaluate failure Santiuste et al. (2014) considered the Hues damage model for composite materials (Hou *et al.*, 2000). This damage model proposes that different types of failure can occur under different kinds of stresses conditions (Figure 27).

Assuming an elongated deformable-body, in Equations (29)-(32), the 1st stresses axis is orientated in the directions of the fibres, the 2nd direction in a transversal from the composite matrix and the 3rd is the through-thickness direction. The material-specific parameters T_i and C_i are the tensile and compressive strength values in the direction i . S_{ij} are the shear strength of the material in the plane ij . Moreover $S_{fi,i}$ are the corresponding materials shear strength.

$$failure_{fiber}: \left(\frac{\sigma_{11}}{T_1} \right)^2 + \left(\frac{\sigma_{12}^2 + \sigma_{13}^2}{S_{f1}^2} \right) \geq 1 \quad (29)$$

$$failure_{m_{cracking}}: \left(\frac{\sigma_{22}}{T_2} \right)^2 + \left(\frac{\sigma_{12}}{S_{12}} \right)^2 + \left(\frac{\sigma_{23}}{S_{f23}} \right)^2 \geq 1 \quad (30)$$

$$failure_{m_{crushing}}: \frac{1}{4} \left(\frac{-\sigma_{22}}{S_{12}} \right)^2 + \frac{Y_c^2 \sigma_{22}}{4S_{12}^2 C_2} - \frac{\sigma_{22}}{C_2} + \left(\frac{\sigma_{12}}{S_{12}} \right)^2 \geq 1 \quad (31)$$

$$failure_{delamination}: \left(\frac{\sigma_{33}}{T_3} \right)^2 + \left(\frac{\sigma_{23}}{S_{f23}} \right)^2 + \left(\frac{\sigma_{31}}{S_{31}} \right)^2 \geq 1 \quad (32)$$

Hou's failure criteria are evaluated in parallel; an element is considered to fail when any of them reaches a value of 1. The mechanical behaviour of this anisotropic model was validated with the reported results from experiments on bovine femur.

4.4 Summary

This chapter explained how FEM is used as an accurate alternative to models of the mechanical behaviour of complex solid geometries in the infinitesimal domain, which are commonly computationally expensive. We also introduced how the type of element chosen to discretize the solid domain can influence the simulation's results and how their shape function can provide alternatives for interpolation of the nodal properties into the entire element. Finally, we presented constitutive equations models found in the literature that describe isotropic or anisotropic relationships between strains and stresses on the cortical bone. Solving a dynamic FEM analysis that includes all of the non-linearities of the plastic region of the bone when the debris is being separated from the main body requires long computational times per time increment.

Chapter 5

XFEM-based reaming simulation.

This chapter describes the methodology followed to define and solve the FEM simulation of acetabular reaming. Reaming on bone tissue is a procedure of destructive and non-reversible nature. Consequently, it is impossible to use experimental data as target and training examples to train an ML estimator. XFEM simulations were selected to construct a training dataset for machine learning as they are capable of providing the strain/stress solution for identical instances of a solid under different loading conditions. Such invariance in the mechanical properties of the deformed material would be impossible to obtain in an experimental setup. This includes the generation of the 3D geometries, the definition of the material's constitutive equations, and the initial boundary constraints for the required dynamic explicit FEM simulation.

We selected XFEM simulation as the source model to compute the expected physically-based deformation given its accuracy and fast convergence to represent cuts, discontinuities and large strain rates. Linear FEM methods are a popular selection for interactive AR/VR simulation. However, they require adaptative remeshing to represent the formation of new discontinuities in the solid. Depending on the size of the mesh, this can exponentially increase the computational load of the system. Furthermore, linear FEM can often present low accuracy when high strain rates are expected. XFEM has proven to be more efficient and accurate than Standard dynamic time discretisation methods to model the material removal process.

The stages followed to complete a FEM analysis are commonly known as pre-

processing, simulation and post-processing. Pre-processing is usually the most arduous step. Here, the models of the physical problem, the material constitutive equations, and the mechanical boundary conditions must be defined. Afterwards, the simulation step is started in the background and is performed entirely by the FEM solver. During the simulation step, the system of differential equations introduced in section 4.1 is assembled to find the displacements in the vertices conforming to the mesh. Finally, the post-processing step consists of the visualisation and analysis of the results according to requirements specific to the problem.

To train a machine learning model on the physical relationships happening during bone resurfacing, it is necessary to create a big dataset of training examples. In order to generate these different training scenarios, the orientation of the reamer head was rotated around the θ and γ angles (Figure 30(c)), which are analogous to the angles of antroversion and inclination $\theta_{antroversion}$ and $\theta_{inclination}$ defined in section 3.1.4.

Our FEM simulation was defined and solved using the software Simulia ABAQUS® CAE. This software is a popular choice by several authors when modelling resurfacing operations on both engineering materials (Liu *et al.*, 2013), and bone (Lughmani *et al.*, 2013; Tai *et al.*, 2013; Tu *et al.*, 2013; Santiuste *et al.*, 2014).

5.1 Mesh generation

Our FEM analysis studies the effect of the interaction between 2 objects, which has been discretised in finite elements. These are the head of an acetabular reamer and a block of cortical bone. Both volumetric geometries were generated in Autodesk Inventor®, a Computer-Aided Design (CAD) software. Later, the object's discretisation and tetrahedrons

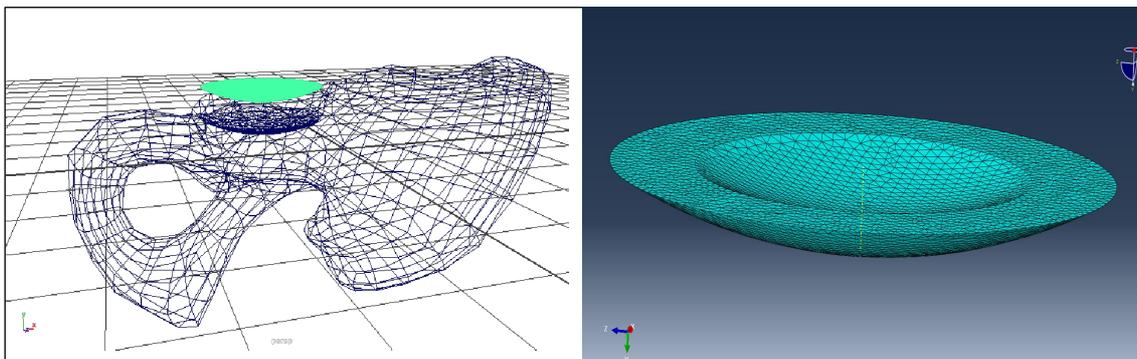


Figure 28 Volume extracted from the geometrical mesh to create the tetrahedral refined mesh. In the picture the offset is intentional to show the fit between the surfaces.

generation was carried out inside the meshing module inside ABAQUS. Since the discretised domain representing the bone will only consider the constitutive equations of cortical bone, we assume that the layer of acetabular cartilage has no significant influence on the mechanical behaviour of the bone during reaming. We also assume that the osteophytes²², if any, have already been manually removed by the surgeon.

5.1.1 Acetabulum

The computational resources required to generate a FEM volumetric mesh with a high number of elements (refined) for the complete pelvic bone would require long computation times and enormous memory resources. Similarly to Kim et al. (2017), we extracted a volume from a bigger object, acetabular cavity from the pelvic bone, and used it to generate the refined FEM tetrahedral mesh similar to minimise memory consumption (Figure 28). This provides a fine level of refinement in the volume of interest while maintaining computation times for each FEM tractable.

The technique used to generate the volumetric mesh of the bone was the free meshing algorithm for tetrahedral elements, which is flexible enough for complex hemispherical surfaces. As an element type, we selected a linear tetrahedron with 4 nodes and one integration point. Because of its linear nature, this type of element performs faster and better when used for contact analysis than second-order elements. This meshing approach was sufficient for the simulation of machining bone as the material behaviour will be assumed as isotropic, and there are no requirements in the orientation of the elements during mesh generation. The level of refinement for the 3D model of the bone is another critical parameter to optimize while generating the tetrahedral mesh. On the one hand, a coarse mesh will produce inaccurate results, requiring unrealistically high cutting forces to trigger fracture. On the other hand, a mesh too fine will yield smaller stable time increments (See section 4.1.3) during processing, requiring higher computational times. The relevant parameters for the generation of the tetrahedral mesh, such as minimum element global size, are listed in Table 5.

In contrast with orthogonal cutting procedures, the high kinetic energy produced by the reamer's rotation is likely to distort the elements of the mesh during a dynamic simulation. Therefore, an additional *distortion control* parameter was enabled on all elements

²² Bony lumps (bone spurs) that grow on the bones around the joints after the cartilage wears off.

of the mesh to ensure convergence during the simulation’s solution. ABAQUS provides an analysis tool for the topology of a generated mesh to allow the user to identify distorted elements or elements with a big ratio between their shorter and their longer edges. **The generated mesh had 20406 elements** with 0% reported errors or warnings triggered by a poor element geometry.

Meshing parameter	Value
Approximate global size	0.45 [mm]
Maximum deviation from the original geometry	10%
Minimum size control	0.045 [mm]

Table 5 Bone model meshing parameters

5.1.2 Reamer

Meshing parameter	Value
Approximate global size	9 [mm]
Maximum deviation from geometry	10%
Minimum size control	0.9 [mm]
Refined tooth edges	0.5[mm]

Table 6 Meshing properties for reamer 3D model.

Surgical reamers used for THR (Figure 29) have a hemispherical shape and cutting edges distributed along its external face. The back of the reamer is hollow to provide a space to collect the pieces of removed bone during surgery. The dynamics of the removed particles mixed with blood is a problem of fluid dynamics where the internal faces of the reamer act as boundary surfaces. This debris behaviour is highly complex and irrelevant for our system since we are interested only in the changes occurring on the surface of the pelvic bone.



Figure 29 Reamer sizes, frontal and back view of the cutting edges. Pictures provided by the Bournemouth University Orthopaedic Research Institute.

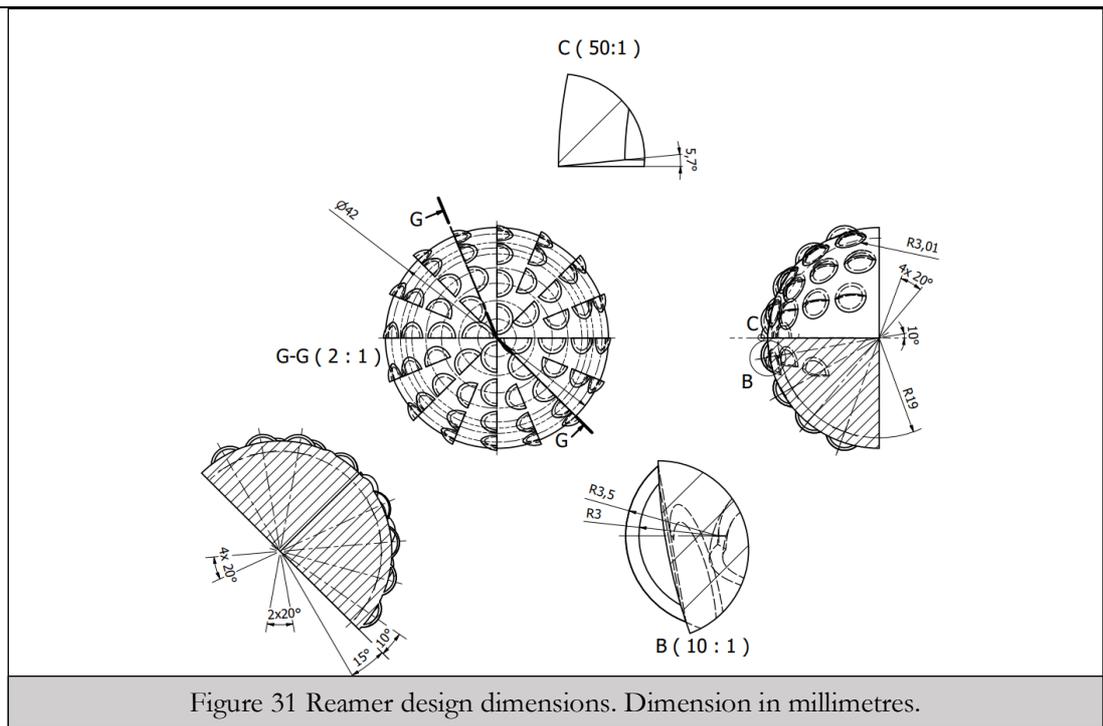


Figure 31 Reamer design dimensions. Dimension in millimetres.

During THR, surgeons start with a small size of the reamer and gradually increase the acetabulum diameter until a patient-specific size. The surgical tool was represented by a simplified solid model of a surgical reamer (Figure 30(b)) with cutting edges to produce element removal when interacting with a volumetric mesh. Depending on their size, the surgical reamer can have between 30 to 60 teeth and a rake angle²³ close to 45°. In real tools, higher rake angles increase the material removal rate but also intensifies the generated heat due to the friction between the debris and the cutting faces. The reamer model used in the FEM simulation has a total of 56 teeth. Additionally, since the thermal effects will be neglected in this numerical analysis, the rake angle was optimised to a value of 85% to maximise the material removal rate. The teeth were modelled with their corresponding rake

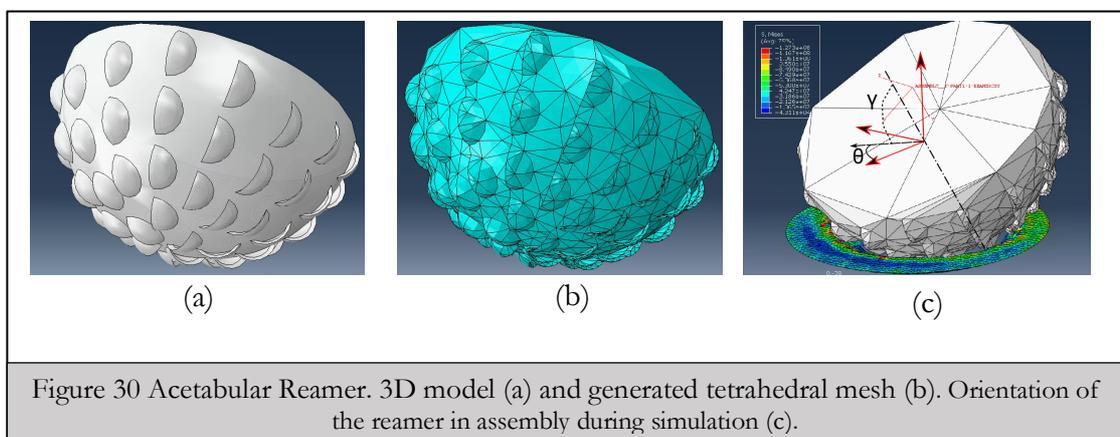


Figure 30 Acetabular Reamer. 3D model (a) and generated tetrahedral mesh (b). Orientation of the reamer in assembly during simulation (c).

²³ The angle between the faces connected by cutting edge.

faces and cutting edges, as shown in the detail view **C** in Figure 31. Since we require results for simulations under different orientation conditions, only one reamer size of 45 mm diameter will be included in the training data. The finite element chosen for the reamer was also a linear tetrahedron. The meshing properties applied to the reamer are listed in Table 6. Unlike the mesh of the acetabulum, not all the elements discretising the reamer have the same size. We performed a refinement in the cutting edges to maintain their spherical shape and ensure accuracy during contact analysis and the collision detected between both meshes.

The reamer's strain and stresses were ignored as we are only interested in its kinematics²⁴. After mesh generation, 4936 elements were generated for the reamer mesh, and 0 errors were triggered during the mesh verification step.

5.2 Material model

5.2.1 Acetabulum

Property	Value
Density [kg/m ³]	1700
Young's modulus [Pa]	17e9
Poisson's Ratio	0.35

Table 7 Mechanical properties of the bone used to model its elastic behaviour.

During total hip replacement, surgeons ream the acetabulum surface while remaining inside the boundaries of the cortical layer, avoiding and exposure to cancellous bone (Beverland *et al.*, 2016) (Figure 26). Therefore, the bone's material constitutive equation used as part of this FEM problem will only consider the behaviour of cortical bone. As described in Section 4.3.2, the nature of the bone tissue is anisotropic. However, it is difficult to determine a predominant osteon orientation (Figure 27) during hip resurfacing simulation because of the biomechanical adaptation that makes the bone increase its osteons density in the direction of the mechanical load (Main *et al.*, 2014). Inside the acetabulum, the osteon orientation is expected to follow a radial direction according to the hemispherical shape of this cavity. **Baro and Deoghare (2018) proved that when the expected coefficient of friction is lower than 0.5,**

²⁴ Kinematic objects are also known as *rigid bodies*.

the isotropic Johnson Cook model yields similar cutting force values to the ones obtained by anisotropic models. Therefore, we will use the JC model as its fast convergence times would allow us to generate a bigger training dataset.

Given the high amount of training examples required, we selected the Johnson-Cook model to yield accurate enough deformation results with relative fast convergence times. The Johnson-Cook constitutive equation, described by equation (26), considers the effect of strain/work hardening and thermal softening $\Theta(T)$ in the flow of tensile stress given a certain PEEQ deformation (See section 4.3.2). The deformation and failure laws applied to the volumetric domain representing the bone are the adaptation of the JC model to machining operations on cortical bone by Childs and Arola (2011) and Baro and Deoghare (2018). The constants used by the authors are listed in Table 4, Table 7 and Table 8. It is essential to mention that all factors related to thermal softening were set to a value of 0 as heat generation will be neglected for our simulation. This assumption can be made as during real surgery, the natural lubrication of the bleeding tissue and the intermittent rinsing with cold water performed by the surgeon avoids a significant rise of the bone temperature and prevents thermal-induced necrosis²⁵. The mechanical properties listed in Table 7 describe the behaviour of the material on the elastic region of its strain-stress curve. Given that the elastic zone of any material is present before the plastic one (see Section 4.1.2.2), the elastic behaviour must be set first and then complemented with the JC damage evolution model (Table 4).

During real surgery, the bone mixed with the natural bleeding is determined by a fluid and particles behaviour. Modelling the behaviour of the debris once it is separated on a numerical simulation increases the complexity and time required to perform the FEM analysis. This is to say, once removed from the main object; the debris does not affect the deformation of the bone. Therefore, tetrahedral elements removed by the cutting edges were immediately excluded from calculations. This was necessary to avoid an exponential growth in the computation times or even failure in the solver due to the excessive deformations product of floating debris.

d1	d2	d3	d4	d5	ϵ_f [m]
-0.09	0.08	0.006	0	0.01	5e-4

Table 8 Johnson-Cook Damage values

²⁵ The death of the cells conforming a tissue due to high temperatures.

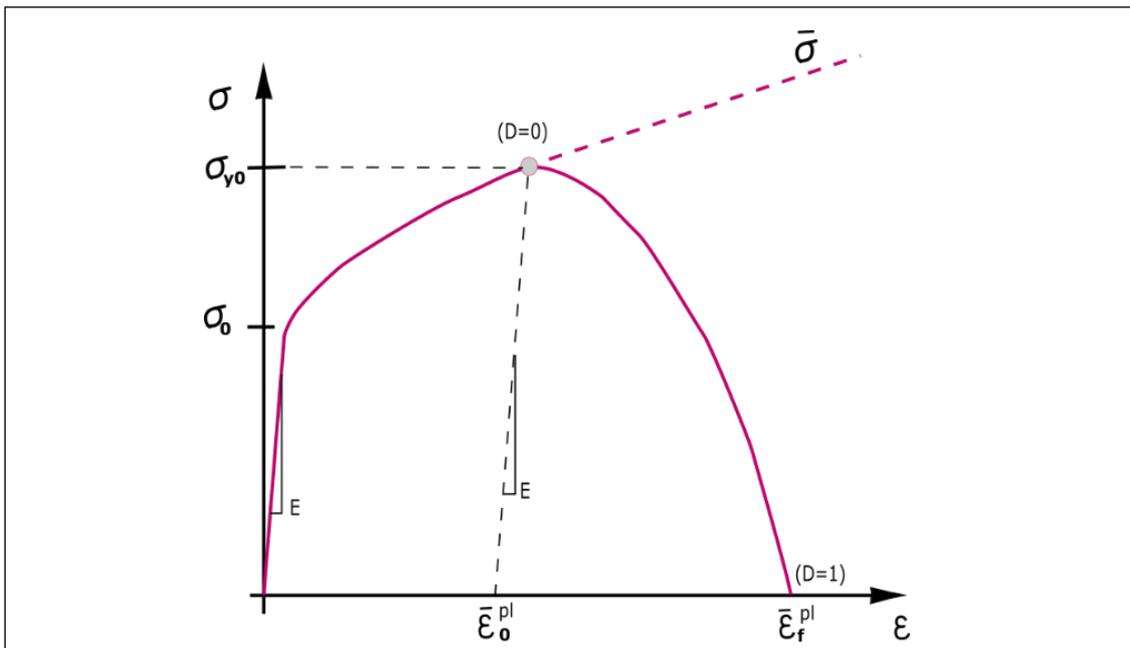


Figure 32 Stress-strain curve with progressive damage degradation. Adapted from [online] <https://abaqus-docs.mit.edu/2017/English/SIMACAEMATRefMap/simamat-c-damageevolductile.htm>

During the damage evolution of a fractured element (Figure 32), the constitutive equation that governs the stress-strain relation of the material follows the Johnson-Cook elastic-plastic hardening law (See section 4.3.2). The initial deformation follows a linear elastic behaviour until the yielding stress σ_0 . Later, once the element Von Mises stress reaches the plastic yielding stress σ_{y0} , the damage D starts being described by (28) with a value equals to 0. The equivalent plastic strains ϵ increases as the element detach itself from the main material block, and therefore, the stress in the element decreases and final fracture occurs when $D=1$ at the strain to fracture ϵ_f^{pl} . In Figure 32, the dashed line represents the element’s plastic behaviour in the absence of yielding plastic stress and when no damage is initiated. The values listed in Table 8 are the constants used to describe the evolution of damage on each element after fracture is initiated (Figure 32). These values were selected to yield a small value of deformation during fracture as it is typical for brittle materials. For our simulation, the final strain at complete fracture was constrained to 0.1mm to avoid excessive ductile deformation in the elements before being removed from the main tetrahedral mesh.

5.2.2 Reamer

Property	Value
Density [kg/m ³]	7800

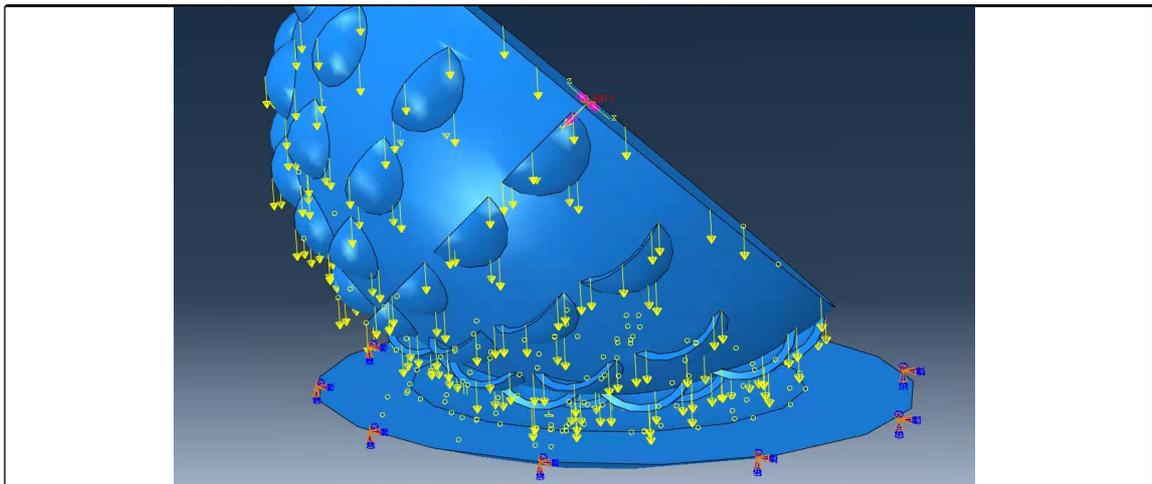


Figure 33 Loading and boundary conditions applied to the assembly

Young Modulus [Pa], E	200e9
Poissons ratio, ν	0.3

Table 9 Mechanical properties for AISI 420

In the industry of medical technologies, the material used to manufacture the surgical reamer is chemically optimised to ensure hardness on the cutting teeth without compromising elasticity and cutting performance. Different surgical manufacturers differ in their chosen reamer geometry and used material. Surgical tool's makers develop specific metal alloys for their products to ensure their competence in the market. Hence, the mechanical characterisation of these materials is challenging to find in the public domain. For our FEM analysis, the material assigned was a Stainless Steel AISI 420 (high-carbon stainless steel with a minimum chromium content of 12%), a generic choice for surgical tools due to its hardness and ease to be sterilised. The values used to describe the material's linear elastic behaviour are listed in Table 9 and were taken from the standards of the American Society for Testing and Materials (ASTM) (ASTM International, 2004).

5.3 Boundary conditions

External loading conditions such as applied forces on the mesh's vertices and kinematic constraints that limit the displacement of some external areas of the mesh must be set during the pre-processing step of the FEM analysis to ensure the validity of the simulation. This section will describe the definition of each one of these conditions for our FEM simulation.

5.3.1 Contact interaction

By default, the ABAQUS FEM solver evaluates collision and contact only between

exterior elements of the FEM meshes. It was necessary to create a set of interactable elements that include both interior and exterior elements to ensure that the reamer will affect all elements conforming to the mesh. In a similar way to Childs and Arola (2011) and Tu et al. (2013), the energy loss during reaming due to the friction between the bone and the cutting tool was modelled through the Columbus friction law using a coefficient of 0.3. Furthermore, the bouncing effect between the reamer's edges and bone was minimised by enabling the effect of gravity and including a damping coefficient between both solids of value 0.2. The effect of gravity (represented as yellow arrows in Figure 33) was included to avoid undesired bouncing of the tool due to the normal force reactions in the reamer teeth while reaming. This resistance to bouncing is usually achieved in theatre via a manual control of the reamer by the surgeon.

5.3.2 Loads.

The magnitude of the force used to push the reamer head into the acetabular cavity during our FEM simulation was obtained from measurements performed on a PrimusRS rehabilitation system from BTE technologies (See Figure 34). The PrimusRS is a device designed for multijoint, musculoskeletal biomechanical measurement. It provides an objective evaluation of the interactive rehabilitation experience by allowing force application in multiple directions. Three fellow surgeons from the Orthopaedic Research institute were asked to replicate the average pressure they would typically exert while performing acetabular reaming. A handle similar to the one used in reaming devices was provided to make sure that the experience was perceived as realistic and align the measurement protocol with the clinical practice. After this experiment, the average value obtained was 230 N.



Figure 34 Protocol used to measure exerted axial force during reaming.

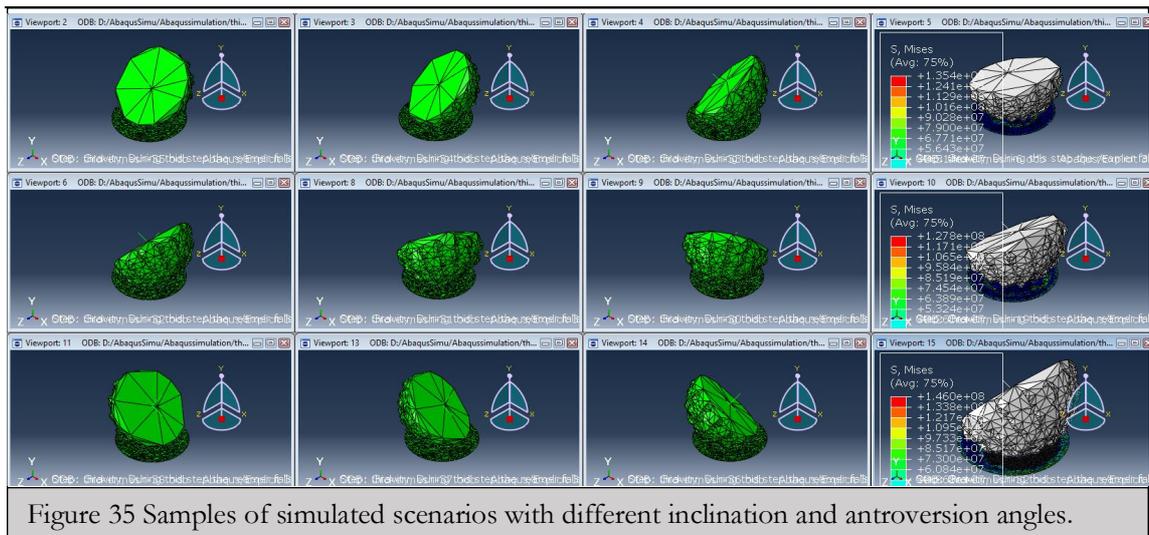


Figure 35 Samples of simulated scenarios with different inclination and antroversion angles.

The applied force in the FEM simulation can be appreciated in the ABAQUS pre-processing view (Figure 33) as a magenta arrow that is aligned with the reamer's rotation axis. Similarly to Tunotti et al. (2017a), the main force was applied a concentrated nodal force to the 3D object representing the tool. To do so, an additional reference coordinate system was attached to the reamer located over its axis of rotation (Figure 33). In other words, the axial force was applied on a node located at the centre of this new reference system and aligned with its axis of rotation during all the simulated orientations of the reamer. In real practice, the force magnitude and direction is a variable that surgeons would adapt, given the visual feedback of bone reaming progress. In our FEM database, the direction of the force varied between the values of 0 and 45 on both the inclination and antroversion angles (See Section 3.1.4).

5.3.3 Displacement boundary conditions

An encastre²⁶ was applied to all the mesh vertices located on the lower face of the bone to simulate a supporting surface capable of providing a reaction force to the reaming axial load. The displacement on the other vertices composing the bone mesh was allowed to ensure convergence of the Johnson-Cook fracture model.

On the reamer, all boundary conditions were applied on its attached reference coordinate system to ensure that they will remain independent of the changes in the reamer's orientation. The displacement of the reamer was enabled only along its axis of rotation. This allowed us to simulate the penetration of the reamer into the bone as the reaming process

²⁶ Boundary condition where no displacement or rotation is allowed.

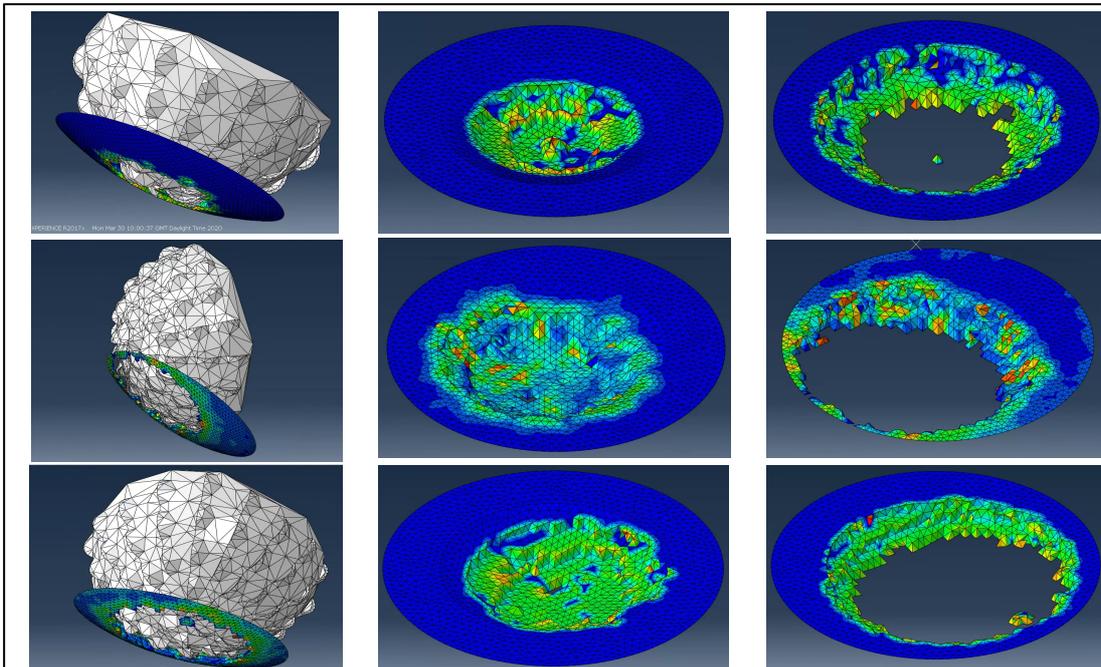


Figure 36 Examples of some simulation under difference force directions, From left to right: Assembly view, results in bone tetrahedral mesh on frame 24 and 94.

progresses. The rotational speed was set to 136 rad/s along its axis of symmetry which is about four times faster than rotational speeds obtained by traditional surgical power systems. Its value was increased to maximise the amount of removed material per frame.

5.3.4 Simulated training examples.

In real practice, the angles of antroversion and inclination (θ and γ) have different range values given the anatomical spatial constraints. The antroversion angle θ could vary from 0° to 180° ; the inclination angle γ can only reach values up to 45° degrees. Therefore, the simulated training cases were built around four different values of inclination, namely $[0,15,30,45]$, while the antroversion angle varied between 0° and 180° every 15° degrees for each inclination value. Figure 35 illustrates the orientation of the reamer, given different combinations of θ and γ . In the right column of Figure 35, the reamer at inclinations angles of 15° , 30° and 45° while the other three most left columns show a variation of the antroversion angle while maintaining the inclination angle static.

5.4 Results

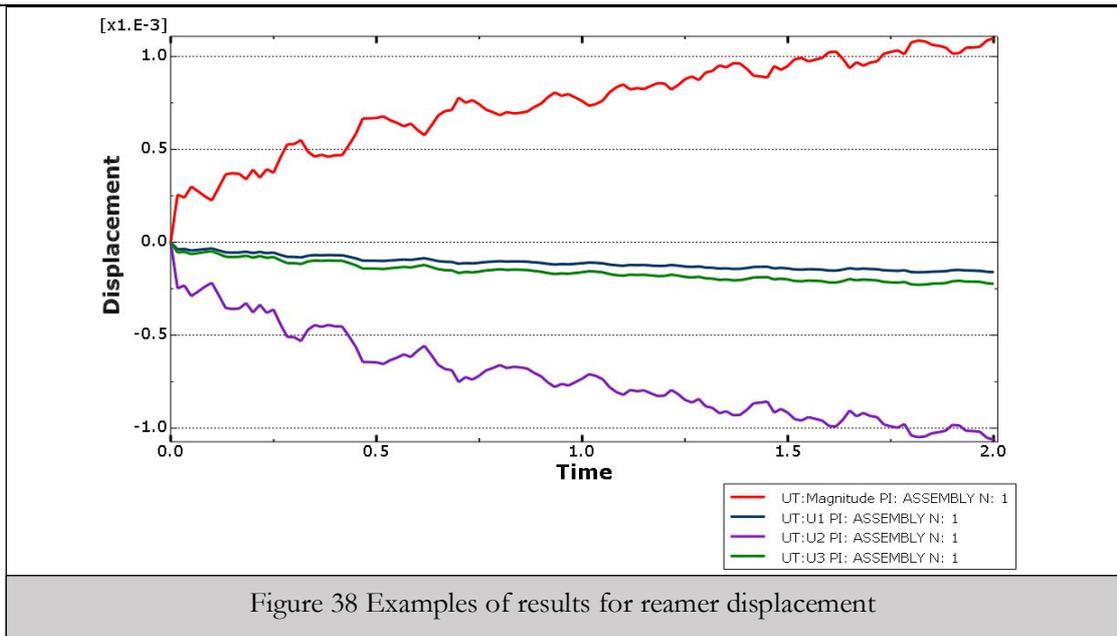


Figure 38 Examples of results for reamer displacement

Figure 36 shows the post-processing view of results from the combinations of different orientations of the tool under the same loading conditions. The tool was hidden during post-processing to visualise the resurfaced area. However, its position can be inferred from the effect of the reaming operation on the volumetric mesh. The elements conforming to the mesh are shaded according to their values of stresses. The colour heatmap used in Figure 37 and Figure 39 assign the colour red to higher values of stress/strain close to the material's failure threshold, while the lowest values are rendered in blue.

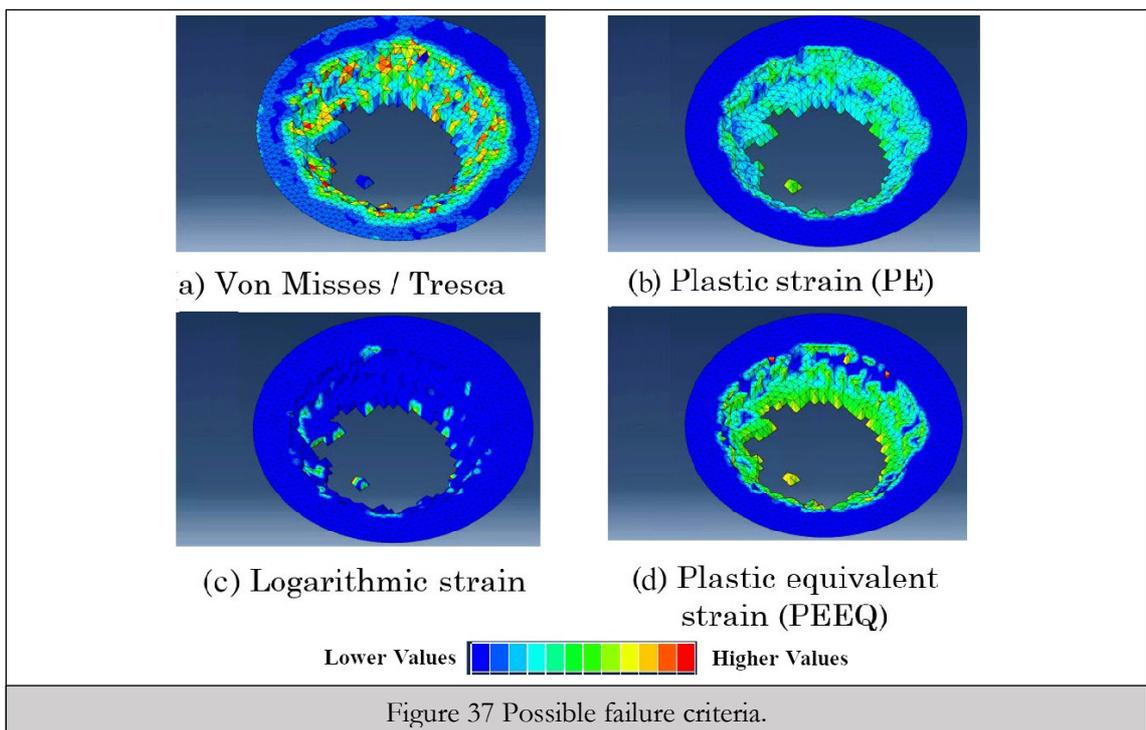


Figure 37 Possible failure criteria.

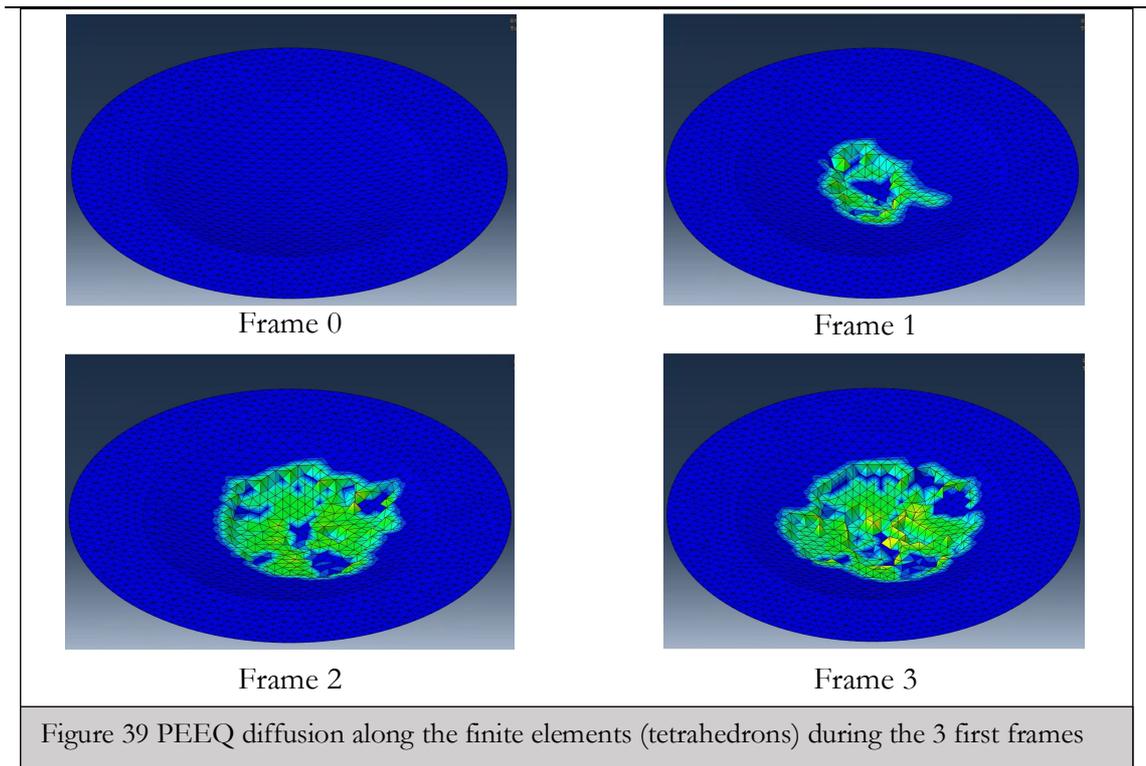
For each unique combination of θ and γ angles describing the orientation of the tool, we obtained 60 reports/frames per simulated second. All tetrahedrons were assigned a binary value at each frame corresponding to their STATUS as active or removed. Elements identified as removed are not rendered in the Abaqus CAE post-processing tool. A total of 120 frames per inclination permutation were generated after an average of 3.5 hours per analysis. The FEM simulation ran on a machine with an Intel processor model Xeon E5-1650, 16GB of Memory RAM, and CPU parallelisation.

The dynamic behaviour of this rigid body has an important meaning in our analysis. The exerted force on the reamer, the reaction forces acting on each one of the edges, and the rate of removed bone produce an effect on the displacement of the reamer while penetrating the deformable mesh. During post-processing, it is possible to evaluate the tool's displacement path product of the penetration of the tool in the machined surface. Figure 38 shows the magnitude of the displacement in each one of the coordinate axis and the absolute magnitude of the displacement vector during the simulated time range.

5.4.1 Discussion

The binary label STATUS calculated per element on each frame is not enough for an ML model to learn the diffusion of accumulated energy transmitted from the tool to the bone responsible for the element's fracture. As shown in Figure 37, different field value stress/strain distributions were explored to identify the best parameter for the training features. The theory of continuum mechanics uses the Von-Mises ($\bar{\sigma}$) and Tresca equivalent stresses as common criteria to evaluate failure/fracture on deformable bodies (Beer, 2011). However, as it can be appreciated in Figure 37, the magnitudes of these two failure criteria are not distributed uniformly over the tetrahedrons that are in close contact with the reamer's edges and are expected to be removed in the subsequent frames. The Logarithmic and Plastic Strain fields, shown in Figure 37(b) and Figure 37(c), succeed in segmenting the tetrahedral mesh elements in contact with the tool and therefore accumulating deformation energy. However, the variation range of the values for these fields is too narrow and hinders their use as a parameter to describe the energy diffusion on each element prior to failure. The Plastic equivalent Strain (PEEQ) segments the mesh most appropriately so that elements accumulating stresses are easy to detect with suitable variation range of its values over the whole mesh.

Ideally, we would include several of the failure criteria available as outputs from the



FEM solution (Figure 37) as features to train our ML learned simulator and describe the energy accumulation inside the mesh. As described in section 6.1.2, a machine learning predictor will require to predict all of these properties on each rollout and taken as input for the following iteration. Thus, increasing the estimated array per element would increase the complexity of the problem and optimise the learning for more than one target variable might accelerate the drift error over high-step rollouts. Consequently, the PEEQ value and the Von Mises criteria are selected as the best feature candidates that best describes the diffusion of deformation energy inside the volumetric mesh through time.

Figure 39 shows a post-processing view where the values of the PEEQ view of the first three simulated frames during the case where the reamer orientation was determined by the angles $\theta = 0$ and $\gamma = 0$. As in previous figures, the reamer was hidden to facilitate the visualisation of the PEEQ values in the mesh surface. Similar to previous figures, the colour heatmap assigns the colour red to higher values of stress/strain close to the material's failure threshold, while the lowest values are rendered in blue. In this figure, it can be appreciated the evolution and diffusion of energy in the time dimension. Figure 39 also shows a highlighted "selection" of the elements across the mesh in contact with the reamer's cutting edges and are candidates to be removed in the following frames.

5.5 Summary

This chapter presented the materials' constitutive equation, geometries, and boundary conditions to define our FEM simulation on the reaming procedure of a piece of cortical bone. For this, we have created digital representations of both the head of a standard surgical reamer and a piece of the acetabulum. Next, we have discretised these volumetric domains using linear tetrahedral elements with an average edge size of 0.45mm. We included in our FEM analysis the effect of nonlinearities such as friction and damping between the surgical reamer and the bone. Finally, we have shown how the PEEQ and Von Mises equivalent stress provide a reliable segmentation of the elements in the volumetric mesh with high accumulated contact energy and high probabilities of being removed.

Even after using an isotropic material's constitutive equation on the bone, the nodal deformations on each frame required an average of 1.75 minutes to be resolved. Considering that each frame corresponds to a time increment of 0.01s, these solver durations are not suitable to be used for interactive training of the reaming procedure. In the next chapter, we will use the data from our FEM results to train a simulator capable of predicting the amount of removed material and identifying the pieces of removed bone at faster rates.

Chapter 6

Machine learning and volumetric deformation.

This chapter describes the state of art ML techniques to speed up calculations of physically-based mesh deformations and provide real-time interaction. Normally, it is impossible to represent FEM meshes in the Euclidean domain because of its elements' variable size and orientation. **This is, the number of neighbours connected to each element forming the volumetric mesh depends on the element's location in the mesh.** Consequently, a graph is the most appropriate data structure to keep the geometric topology and element-to-element relationships of a tetrahedral FEM mesh. Therefore, the graph-structured data and learning concepts in the graph domain are introduced as a tool to imitate the plastic deformation behaviour in continuum mechanics.

6.1 Data-driven / Machine learning (ML) methods for mesh deformation

Modelling complex physically-based deformations, such as the ones required to represent bone resurfacing, can be computationally expensive. Therefore, they are challenging to perform at the interactive framerates required in simulation-based training. However, the current development of machine learning algorithms has allowed some authors to estimate deformations results based on examples calculated offline.

6.1.1 Algorithms to speed up FEM calculations

Several authors have attempted to use ML to reduce the computational complexity of FEM simulations. One approach found in the literature models a linear FEM stiffness matrix via machine learning to fit empirical deformation data (Bickel *et al.*, 2009; Mosbech *et al.*,

2009). Theoretical mechanical characterisation of the material is avoided, and the FEM model is downgraded in complexity to the linear domain. It is important to note that the implementation of this method was applied exclusively to elastic deformations, and the built estimators were object-specific. One shared limitation on Bickel et al. (2009) and Mosbech (Mosbech *et al.*, 2009) is that the ML model had to be fitted again every time the shape or material of the target object changes. Mosbech et al. (2009) modelled the vertex displacement deformations resulting from a nodal force using compactly supported radial basis functions. The authors achieved an error of 1.41 ± 0.87 mm on their predictions for their mechanical model. In a similar work published by Bickel et al. (2009), the authors obtained deformation estimations with a minimum mean error reported of 0.8mm when using a biharmonic RBF kernel. They did so by including a force sensor and a high-resolution stereo vision system during the deformation acquisition stage. In the work of Bickel et al. (2009), the properties fitting process was aimed to parametrise the object's Young module based on its local coordinates and the applied stress conditions. At runtime, the model was treated as elastic static FEM simulation with a material's Young modulus (See section 4.1.2.2) that varied along its surface according to the applied force input.

Kim et al. (2017) implemented a dynamic data-driven volumetric human body model to approximate the variable value of several mechanical properties used in FEM (i.e Elastic module, Poisson's module and damping coefficient. See section 4.1.2.2). The volumetric model was trained using 4D captured data of a human subject performing different types of movements. This dataset contained the tracked displacements of the fat and skin tissues by the action of gravity and the dynamics associated with each activity (running, walking, jumping). Kim et al. (2017) simplified their analysis by assuming fat and skin as the same material with the same mechanical properties. Furthermore, in order to reduce the number of mesh elements and allow the simulation to run at interactive framerates, the FEM calculations were only carried out in the external volumetric layer of the 3D model.

Simplifications in the mesh refinement increase the online FEM speed but could reduce the accuracy of the results. Seiler et al. (2012), Si et al. (2017), and Chen et al. (2015a) created databases with two versions of physically-based deformations; a high and low resolution one. These databases were mapped using linear relationships between them. The numerical simulations on a coarse mesh (low resolution) ran in real-time. Next, its results were enriched thanks to interpolation with the precomputed results of the high-resolution

database, maintaining accuracy. The authors validated this approach under interactions not considered in the training database. The interpolation to unexpected input conditions was possible by assigning weights or relevance relationships for each available training set relative to the coarse displacement calculated in real-time. Thus, allowing the model to find the most suitable candidate to enrich the simulation at each request.

Seiler et al. (2012) added skin meshes to both the fine and coarse versions of the 3D FEM volumetric object but used them exclusively for rendering purposes. These geometrical meshes were linked to their corresponding models (coarse and fine) by linear interpolation so that deformations could be extrapolated quickly. After calculating the deformation on the low-resolution geometrical model, the corresponding high-resolution values were rendered. This high-resolution skin mesh update was performed by blending detail around the contact point where the nodal force was applied in like a stamp. Since the contact point is considered as the origin for this stamp-like detail enrichment, this method is constrained to just one collision point, and the detail enrichment is limited to only affect the surface of the geometrical mesh. Alternatively, Si et al. (2017) were able to predict deformation on all the vertices of a volumetric mesh with the behaviour of a composite material. Their model included coupling forces between different adjacent objects by including the effect of inner forces in a runtime simulation. The mapping relationships between both materials were linear regression optimised by the moving least squares method (MLS). This parametrisation allowed the algorithm to insert new vertices in the coarse mesh and interpolate the corresponding deformation values.

Chen et al. (2015a) proposed another alternative for reducing the number of elements involved in a FEM mesh while maintaining accuracy. Their framework works independently of the object geometry and material composition. Unlike previously mentioned authors, Chen et al. (2015a) mapped the relationship between the mechanical properties of two different sets of materials on a coarse and fine mesh instead of the deformations on their geometries. This algorithm exchanges the fine mesh material properties with a new equivalent database of materials that would produce the same change in the potential energy in the coarse mesh as the high-resolution one. Chen et al. (2015a) built the databases with hyperplastic materials as they are more suitable for fitting the mechanical properties in the strain-energy space. Chen et al. (2015a) were the only authors found to include anisotropic materials (See section 4.3.2) in the coarse metamaterial database. The authors also reported a performance increase from 51 to 489 times, maintaining relative errors under 0.003%

compared with the high-resolution FEM simulation.

6.1.2 Methods to approximate deformations without running FEM simulations in real-time.

This section includes identified research methods that predicted deformation values based on an initial geometry and force input without performing physic calculations at runtime. Instead, they are able to estimate their results based on Neural Networks (MLP) trained with FEM or empirical data. All of the identified methods trained a separate MLP for each component (x,y,z) of the deformation or stress to be approximated. The first reported attempt to learn FEM simulation results using artificial Neural Networks was made by Bhise and Pratihar (2006). This work addressed the problem in a 2D domain for cylindrical deformation analysis. Only recently, Tawbe and Cretu (2017) and Tonutti et al. (2017b) have made attempts in the 3D domain to estimate elastic physically-based deformation.

Bhise and Pratihar (2006) proposed an algorithm able to estimate the behaviour of a simple 2D finite element simulation while using different mesh element types. In their experiments, an elastic cylinder was compressed between two plates. Once a significant number of FEM cases were simulated offline, the data was used to train through backpropagation a Genetic Neural Network and an MLP. The article reported a difference between the behaviour of the learned models for tensive and compressive loading conditions. The genetic algorithm yielded lower Mean Squared Errors (MSE) when predicting tensile stresses (36.03% vs 79.08%) and deformation (3.94% vs 3.71%) than under compression. The MLP model yielded lower MSE when estimating the effects of compressive deformation (1.0% vs 1.06%) and stresses (15.35% vs 19.69%). This method has a reduced application in interactive animation development as in computer graphics, the type of the mesh element is maintained during the whole simulation

Recently, Tawbe and Cretu (2017) and Tonutti et al.(2017b) used the initial object geometry and the parameters describing the applied force to compute changes on the surface of a volumetric object interactively to train Neural Networks and Support Vector Machines (SVM). The main difference between these two contributions is the type of data used as ground truth during training.

Tawbe and Cretu (2017) used a robotic arm to measure the input force magnitude and orientation applied to an elastic object while a Kinect RGB-D sensor was tracking the changes on the surface of this object. A synchronisation stage was carried out between the

refresh rates of the Kinect and the pressure sensor to ensure the correspondence of data. The motivation behind the approach of Tawbe and Cretu (2017) was to accurately approximate the deformation based on an empirical characterisation of the object's mechanical properties bypassing any numerical model. This method has the advantage of being equally valid for linear, anisotropic, hyperplastic or composed materials. The authors reported maximum errors of 23.6% on their approximations when the model was validated around the entire surface and 9.2% when only the deformed area was analysed. Nevertheless, its accuracy in terms of the deformation magnitude is constrained by the RGB-D sensor resolution, and the quality of the geometrical mesh reconstructed from a Kinect point cloud.

Tonutti et al. (2017b) used offline FEM simulations of a heterogeneous material as training data. In their FEM model, the forces were applied on external mesh nodes while the deformation analysis was carried on internal ones. After building their database with FEM results, the authors performed a data exploration analysis of the stress-strain results versus the dynamic parameters used in the simulation. This explorative data analysis aimed to identify the variables best suitable to be used as features to train their neural network. Therefore, they identified the parameters that had a strong influence on the final deformation of the nodes of the inner mesh. As selected features, the authors considered the distance between the applied nodal force and each vertex of the mesh, the F_x , F_y and F_z components of the applied force and the angle between the applied force and a vector describing its relative position to each vertex. Tonutti et al. (2017b) trained artificial neuro-networks (ANN) and support vector regression (SVR) models and compared their performance. The SVM regressor yielded lower mean displacement errors of 0.191 ± 0.201 mm compared to 0.456 ± 0.384 mm for the ANN with one hidden layer. With errors below 0.5 mm and computational times below 10ms, this is so far the most accurate and fastest estimator of physically-based deformation found in this literature review. Compared to Tawbe and Cretu (2017), using the FEM data to train the ML estimator achieved more accurate results as the object topology remained unchanged for examples and predicted cases.

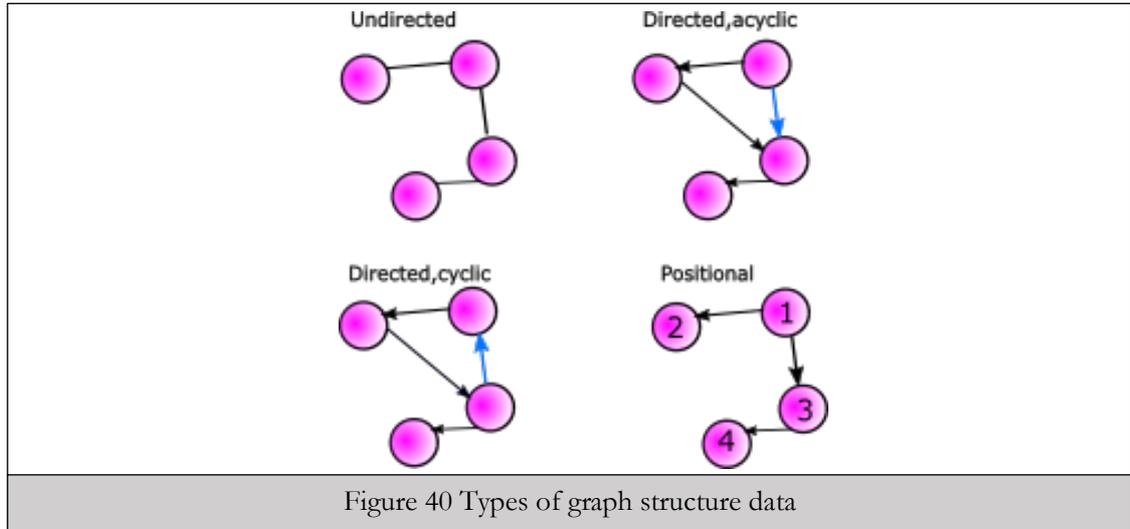
Recently, the “Graph Network-based Simulator” (Sanchez-Gonzalez *et al.*, 2020) was proposed. An interactive simulator of particle physics that discretises the mechanical state of deformable objects and fluids as a particles-based mesh-free simulation. By doing so, the material's deformation is computed through kinetic energy broadcasting between all particles of the system.

6.2 Machine learning on graph-structured data.

In the previous section, we discussed how standard multilayer perceptrons (MLP) had been trained with finite element simulation data to speed up calculations of elastic deformations over the vertices of a volumetric mesh. These techniques treat each vertex of the mesh as an independent training example ignoring its relationships with its neighbours. In numerical simulations, the dynamic calculation of stress accumulation depends on the communication of stresses between mesh elements. This phenomenon is an important feature that would be ignored by flattening the input elementwise, as previously explored by Tonutti et al. (2017b). Popular networks to keep relational learning between data structures are Convolutional Neural Networks (CNN) for grid-shaped data and Recursive Neuro Networks (RNN) for sequential data. However, unlike FEM meshes, relationships between entities in regular grids or 1D sequences are constant over all the instances. In a volumetric FEM tetrahedral mesh, the number of neighbours connected to each element forming the volumetric mesh depends on the location of the element in the mesh. In cases where the mechanical behaviour of composite materials is modelled, the data can require different types of representations of the relationships between the elements. We believe that learning on graph-structured data provides a suitable alternative to keep the geometric topology and element-to-element relationships of a tetrahedral FEM mesh. This chapter explains the evolution of learning in the graph domain, emphasising its uses as estimators in physics engines. In some cases, the notation used by different authors has been adapted from their original publication to keep consistency and facilitate the comparisons between relevant contributions found in the literature.

6.2.1 *Graph nomenclature*

A graph \mathbf{G} is a data structure that possesses a total of N nodes \mathbf{n}_i , each one with attributes stored in the features vector \mathbf{n}_i . The nodes are usually connected through edges \mathbf{e}_{ij} (except sets of nodes with no connection between them, like point clouds). This relationship between nodes can also contain weights or information stored in the vector \mathbf{e}_{ij} . In some cases, external effects that influence all nodes and/or edges can also be included as part of the graph, and its attributes are stored as global/universal features (\mathbf{u}). Following Battaglia (2018) notation, we will refer to the *nodes*, *edges* and *global* objects as *entities* of the graph.



Graph data structures can also be classified according to their internal organisation (Figure 40). When the relationships between two nodes $(\mathbf{n}_i, \mathbf{n}_j)$ remain immutable if it is evaluated in the inverse order $(\mathbf{n}_j, \mathbf{n}_i)$ the graph is referred to as *undirected*. However, in some applications, the order of how the message is passed between nodes requires making a distinction between the edge's sender and receiver. These types of graphs are known as *directed*. Similarly, if a directed graph contains a closed circuit, we say that the graph is *cyclic*. An example of a non-cyclic graph is an RNN whose hidden states are comparable to directed acyclic graphs. Finally, if a unique logical index for each node is required to determine the relative position between nodes, the graph is considered positional.

The graph's connectivity can be described by its adjacency matrix \mathbf{A} and degree matrix \mathbf{D} . The adjacency matrix of size $N \times N$ is a sparse matrix where $A_{ij} \neq 0$, only if \mathbf{e}_{ij} exists. When \mathbf{A} is not a binary matrix, it includes weights for the connection between nodes (edges). The degree matrix \mathbf{D} is a diagonal matrix with values $D_{ii} = \sum_{j=1}^N A_{ij}$

6.2.2 Machine learning on graphs for dynamics and physically-based interactions.

The Graph Neural Network (GNN) is a model proposed for supervised learning on any type of graph-structured data. Gori et al. (2005) and Scarselli et al. (2009b) suggested that a learnable function τ would be able to map the graph \mathbf{G} , and its nodes \mathbf{n}_i into an output feature vector $\mathbf{o} \in \mathbb{R}^m$. Instead of flattening the data and using each node in the graph as a separate input for an MLP; this type of learning is valid for all the types of graphs described in Figure 40.

The graph neural network (GNN)(33)-(35) assigns and encoded status \mathbf{x}_i (33) to each

node depending on its own feature values \mathbf{n}_i , the features assigned to their connected edges \mathbf{e}_{ij} and the encoded and feature status of its neighbours $\mathbf{x}_{ne[i]}$ and $\mathbf{n}_{ne[i]}$. The values \mathbf{x}_i per node are updated recursively (34) until converging to a stable solution. Scarselli et al. (2009b) compared this recursive process to an encoder network. A fast convergence to a unique solution of $\mathbf{X}(t)$ is ensured given the Banach’s fixed point theorem (Khamsi and Kirk, 2001), for any initialisation of $\mathbf{X}(0)$. For an in-depth demonstration of how the learnable update functions ϕ^n and ϕ^o are ensured to be unique contraction maps, we refer the reader to Scarselli et al. (2009b). After a stable status for \mathbf{x}_i has been reached, the status variable and the original node features are taken as inputs to compute the decoded output (35).

$$\mathbf{x}_i = \phi^n(\mathbf{n}_i, \mathbf{e}_i, \mathbf{x}_{ne[i]}, \mathbf{n}_{ne[i]}) ; \quad 1 \leq i \leq N \tag{33}$$

$$\mathbf{X}(t + 1) = \phi^n(\mathbf{X}(t), \mathbf{N}) ; \quad \mathbf{x} \in \mathbf{X} \tag{34}$$

$$\mathbf{o}_i = \phi^o(\mathbf{x}_i, \mathbf{n}_i) ; \quad \mathbf{o} \in \mathbb{R}^m \tag{35}$$

Li et al. (2015) proposed the Gated-Graph neural network (GGNN), a framework for directed graph classification tasks. Instead of having to ensure ϕ^n to be a contraction map so that any initialisation for the encoded node’s states can be used, Li et al. (2015) suggested initialising $\mathbf{x}_i(0)$ with one-hot²⁷ encoded *annotations* \mathbf{s}_i . This is, $\mathbf{x}_i(0) = [s_i, \mathbf{0}] ; 1 \leq i \leq N$. Instead of learning a single function ϕ^n , this model expands the update of the encoded node

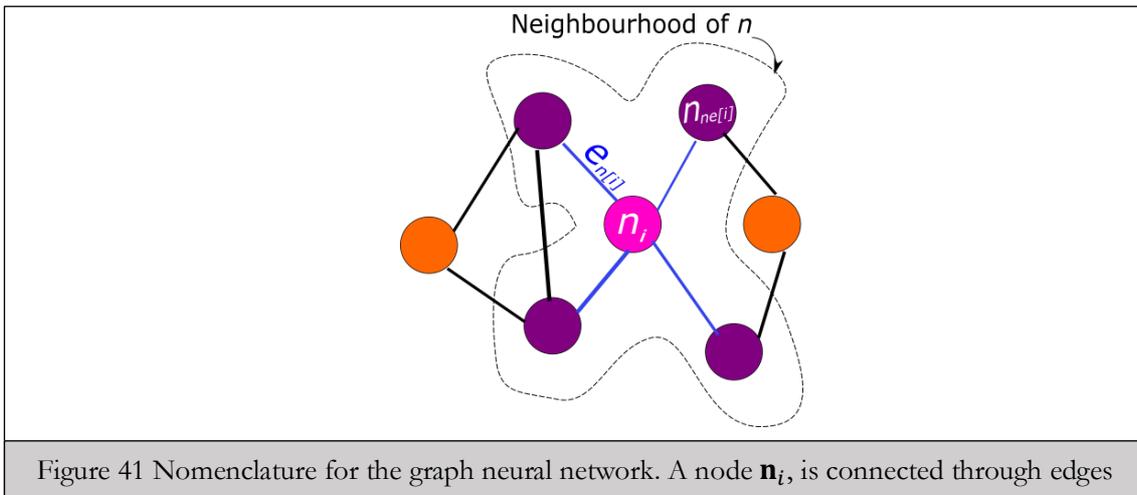


Figure 41 Nomenclature for the graph neural network. A node \mathbf{n}_i , is connected through edges

²⁷ Binary representation of categorical variables where only 1 bits are “hot” or TRUE.

representations into several steps in (36)-(39). The convergence of the encoding network is reached by iterating $\mathbf{X}(t)$ through gated recurrent units (Cho *et al.*, 2014). The nodes to edges information broadcast is performed more efficiently by learning a matrix $\mathbf{C} = \phi^e(\mathbf{e})$, where ϕ^e is a learnable function applied to the edges and shares some shape similarities with the graph’s adjacency matrix \mathbf{A} and holds different values for incoming and outgoing messages.

$$\mathbf{a}_i(t) = \mathbf{C}[\mathbf{X}(t - 1)]^T + b \tag{36}$$

$$\mathbf{z}_i(t) = \text{act}(W^z \mathbf{a}_i(t) + U^z \mathbf{x}_i(t - 1)) \tag{37}$$

$$r_i(t) = \text{act}(W^r \mathbf{a}_i(t) + U^r \mathbf{x}_i(t - 1)) \tag{38}$$

$$\mathbf{x}_i(t) = (1 - \mathbf{z}_i(t)) \odot \mathbf{x}_i(t - 1) + \mathbf{z}_i(t) \odot \left[\tanh \left(W \mathbf{a}_i(t) + U(r_i(t)) \odot \mathbf{x}_i(t - 1) \right) \right] \tag{39}$$

When a graph-level output is desired, instead of connecting a “supernode” to all the nodes in the graph (Scarselli *et al.*, 2009b); the GGNNs computes an output based on aggregation over all nodes in the graph and their encoded \mathbf{s}_i values (40).

$$\tau(\mathbf{G}, \mathbf{n}) = \tanh \left(\sum_{i=1}^N \text{act}(\text{NN}_1(\mathbf{x}_i(t), \mathbf{s}_i)) \odot \tanh(\text{NN}_2(\mathbf{x}_i(t), \mathbf{s}_i)) \right) \tag{40}$$

Battaglia *et al.* (2016) proposed the Interaction Network (IN), another alternative to replace computing the encoded status \mathbf{x}_i per node \mathbf{n} applied to estimate dynamic interaction between *rigid bodies*²⁸. This model initialises the attributes of edges and nodes based on the dynamic conditions of each object (position, speed, acceleration) and triggers the first update of the edges’ attributes in terms of the *object-to-object* exerted forces. This step is followed by an aggregation of the newly updated edges into all connected node status, in their case positions and velocities. Compared with the GNN, the GN’s ϕ^n is broken down into an update of the edges attributes performed by the learnable function ϕ^e (41) and an output function ϕ^o (42) that takes as its inputs the aggregated updated edges and the attributes of each node. The Interaction Networks allow the inclusion of external effects to be concatenated with the edges and node features instead of treating them as a special node in the graph as proposed by Scarselli *et al.* (2009b). In the IN, the relationships are aggregated only at the receiver nodes; $\rho_{\mathbf{v}_i}^{e_{i,j} \rightarrow n_j}(\mathbf{e}_{i,j})$, where n_j receives the effect of all the relationships

²⁸ Objects with neglectable deformation.

from all nodes \mathbf{n}_i . One of the main benefits of using IN is the ability to compute the dynamic status of objects, even thousands of steps in the future after being trained on a frame-in-frame-out basis. Chang et al. (2016) explain how this rollout over long time sequences can be possible given the Markovian nature of physics laws.

Instead of explicitly setting up the relationships between the nodes of a graph, Chang et al. (2016) trained their model to recognise interactions from all possible permutation of nodes $(\mathbf{n}_i, \mathbf{n}_j)$ through self-attention labels. Santoro et al.(2017) combined a similar approach with CNN to detect the presence of objects on a video sequence and learn the relationships between them (43).

$$\mathbf{e}'_{ij} = \phi^e(\mathbf{n}_i, \mathbf{n}_j, \mathbf{e}_{ij}) \quad (41)$$

$$\mathbf{o}_j = \phi^o(\mathbf{n}_j, \rho_{\mathbf{v}_i}^{e_{ij} \rightarrow \mathbf{n}_j}(\mathbf{e}_{ij}), \mathbf{u}) \quad (42)$$

On a survey by Gilmer et al. (2017), models that update the graph's relationships values before updating the nodes are defined as Message Passing Networks (43). Gilmer et al. (2017) also implemented GGNN and IN for comparing the estimated properties on molecular compounds where each chemical element corresponded to a node on the graphs. They found that GGNN yields better estimations when multiplying the matrix C (36), with the original node features of the receiver node. Here, undirected graphs were adapted to fit the directed nature of (43) by duplicating and inverting all original edges. A global entity was also considered, but in this case, the authors modified Scarselli's (2009a) concept of a supernode by assigning special features to its edges.

$$\mathbf{e}'_{ij} = \text{matmul}(\phi^e(\mathbf{e}_{ij}), \mathbf{n}_j) \quad (43)$$

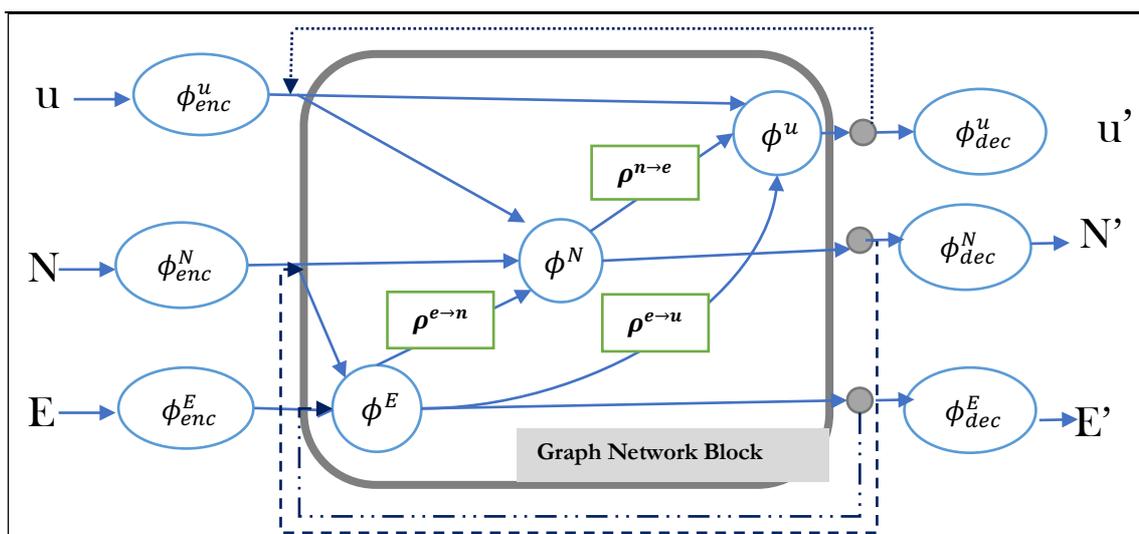


Figure 42 Graph network block. Encode-Process decode in Sanchez-Gonzalez et al. (2018). Figure adapted from Battaglia et al. (2018)

In an attempt to integrate different learning on graph techniques, Battaglia et al. (2018) created the Graph Network (GN²⁹) framework (44)-(46). An open-source framework built based on the Google Tensorflow and DeepMind Sonnet³⁰ libraries. The API is compatible with positional graphs as it stores an identifier of each one of the nodes and edges. As in Gilmer et al. (2017), the analysed graphs are directed, but undirected cases can be included by duplicating and mirroring the directed edges of the graph. A key feature that differentiates Graph Networks (GN) from other techniques that learn in the graph domain is that its outputs update the attributes in all of the graph's entities³¹ while maintaining the connectivity of the input graph. This graph-to-graph nature allows GN blocks to be used to generate complex learning architectures. However, GN also requires maintaining shape invariance between its inputs and outputs as they are expected to have the same number of nodes. Shape invariance applies only to single-pass outputs, and it does not mean that all graphs used during training or validation should have the same structure. GN allows for the update functions ϕ to be any function, not restricted to neural networks. In case that NN are implemented, the trainable parameters are shared across all of the same entities. This allows the generalisation for graphs of different sizes during generalisation and takes advantage of parallelisation in training.

²⁹ Not to be confused with Scarselli's GNN.

³⁰ https://github.com/deepmind/graph_nets.

³¹ Entities of the graph are nodes, edges or globals.

$$e'_{ij} = \phi^e \left(\rho^{n \rightarrow e} (\forall n_i \in U_{e_{ij}}), e_{ij}, u \right) \quad (44)$$

$$n'_i = \phi^n (\rho^{e \rightarrow n} (\forall e_{ij} \in U_{n_i}), n_i, u) \quad (45)$$

$$u' = \phi^u (\rho^{n \rightarrow u} (\forall n \in \mathbf{G}), \rho^{e \rightarrow u} (\forall e \in \mathbf{G}), u) \quad (46)$$

The graph network model, represented as Graph Network Block in Figure 42, updates the attributes stored in the entities e_{ij} , n_i and u following the Equations (44)-(46). Before the network ϕ^e performs the first update call, the features of each edge are concatenated with the attributes stored on its connected nodes and in the global entity. The aggregation step performed by the function $\rho^{n \rightarrow e}$ broadcasts the value of the accumulated feature values from the nodes connected by the edge e_i to the updated attributes of the edge e_{ij}' . Similarly, the network responsible for the update of the nodes' attributes ϕ^n takes as its input the feature vectors of the nodes after being concatenated with the updated attributes of their adjacent edges and the attributes of the global entity. Finally, the global features are updated by ϕ^u , also taking the aggregation of the updated features n_i' and e_{ij}' as an input. In the aggregation functions $\rho^{l \rightarrow m}$ in Equations (44)-(46)), l and m can be any entity of the graph (e_{ij} , n_i and u). The aggregated values are the features in all the instances of the entity l contained in the neighbourhood U of m_i . Sanchez-Gonzalez (2018) suggested linking the GN to a decoder and an encoder neural network responsible of finding the best representations for the values of each entity independently. The authors also suggest carrying out k recursions over the block to allow the network to pass messages over a specific range of neighbours in a similar way to the original GNN. By doing so, each nodes' "level of influence" increases by one connection on each recursion and allows nodes to communicate their attributes to others up to k edges away. For many applications, the number of recursions k is not higher than 3, as the effects of a node on another node decrease with the number of connections between them.

A similar structure to the GN has also been applied for graph generation (Li *et al.*, 2018). This model overcomes the limitations of a static number of nodes per graphs by recurrently evaluating the graph structure and adding edges or nodes as required. It updates edges and nodes features to determine whether to include a new node and define its connectivity. The apparent limitation of this Generative Network is that once the size of the graph reaches a specific limit, updating the attributes inside the node and edges is expensive to compute as the probability of creating a new edge must be evaluated against all existing

nodes.

GN can also be applied to reinforcement learning (Zambaldi *et al.*, 2018), where the graph structure is built on runtime by a self-attention mechanism in charge of inferring relationships between the player and interactable agent in the environment. The status of the environment is passed to the model as an image translated to a GN Block via a CNN. Another implementation of GN was proposed for physics and gravitational law (Cranmer *et al.*, 2019) and particle mechanics (Sanchez-Gonzalez *et al.*, 2019). In the latter one, the rollout performance over thousands of steps in the future was increased by adding a Runge-Kutta integrator and describing the behaviour of the particles through Hamiltonian equations (Hamrick *et al.*, 2018).

Recently, an application of GN was developed for computing deformations in discretised objects (Sanchez-Gonzalez *et al.*, 2020). In our understanding, this is the closest work to the model we propose in Chapter 7.3. The “Graph Network-based Simulator” also implements an encoder-process-decoder model (Sanchez-Gonzalez *et al.*, 2018). Deformations are handled by discretising objects and fluids as a mesh-free/particle-based simulation. The graph update focuses on estimating each particle's linear and angular acceleration to resolve the next system state by an Euler integrator. However, the prospective interactions \mathbf{e}_{ij} between all particles in the system are required to be evaluated after each rollout step of the GNS. This continuous search for potential relationships between the particles of the system can cause the interactive speed of the system to drop as the number of particles increase significantly.

6.3 Summary.

In this chapter, we presented a review of data-driven techniques that merge the accuracy of numerical simulations with the efficiency of data-driven and machine learning techniques to estimate deformations based on physical inputs at faster computation times. The main approaches identified include models that linearize the FEM stiffness matrix and constitutive equations of the materials using machine learning techniques. However, some physical deformations can have a mechanical behaviour so complex that it is impossible to be accelerated using such linearization approaches. In those cases, other authors have proposed methods to approximate deformations based on offline precomputed FEM simulations.

Typically, it is not possible to represent the relationship between the elements on a FEM meshes in the Euclidean domain because of their variable size and orientation. Therefore, a graph is the most appropriate data structure to keep the geometric topology and element-to-element relationships of a tetrahedral FEM mesh. This chapter explains the fundamentals of graph-structured data and learning in the graph domain. We have introduced the concept of machine learning on graphs as they provide a suitable architecture for learning through message passing mechanisms

Except in Chen et al.(2015a), the literature showed that the data-driven methods proposed so far are highly constrained to the geometry and material used during training. Additionally, the cumulative effect of plastic deformation was not found to be considered in any of the reported models. This shows that more research is needed in this area to allow simulations that include a change of the mesh topology. One possible solution to include discontinuities is to implement the embedded finite element method or increase the training data's size to include the analysis of several stages of the defined mesh.

It is also clear from the literature review that approximation errors of mesh deformations below 1mm require a large dataset on which an exploratory data analysis must be performed to identify the relevant training features to be used. Approaches where the training data comes from machine vision-based capture systems are highly realistic because they bypass any simplification performed while building the FEM model. However, they do not apply to this research since acquiring a significant amount of visual data intraoperatively would compromise the safety of the patient undergoing surgery. Also, since reaming is a plastic non-reversible procedure, it would be necessary to find a significant number of samples with comparable mechanical properties, which is extremely hard in practice. Therefore, the best approach to train our ML learned simulator would be through offline simulated data obtained by theoretical biomechanical models, as reported in Tonutti et al. (2017b).

Finally, more research is required to include the effects of more than one input force for neural-network models. As most of the algorithms can be parallelised and multiple force input points and their cumulative effect can be further researched.

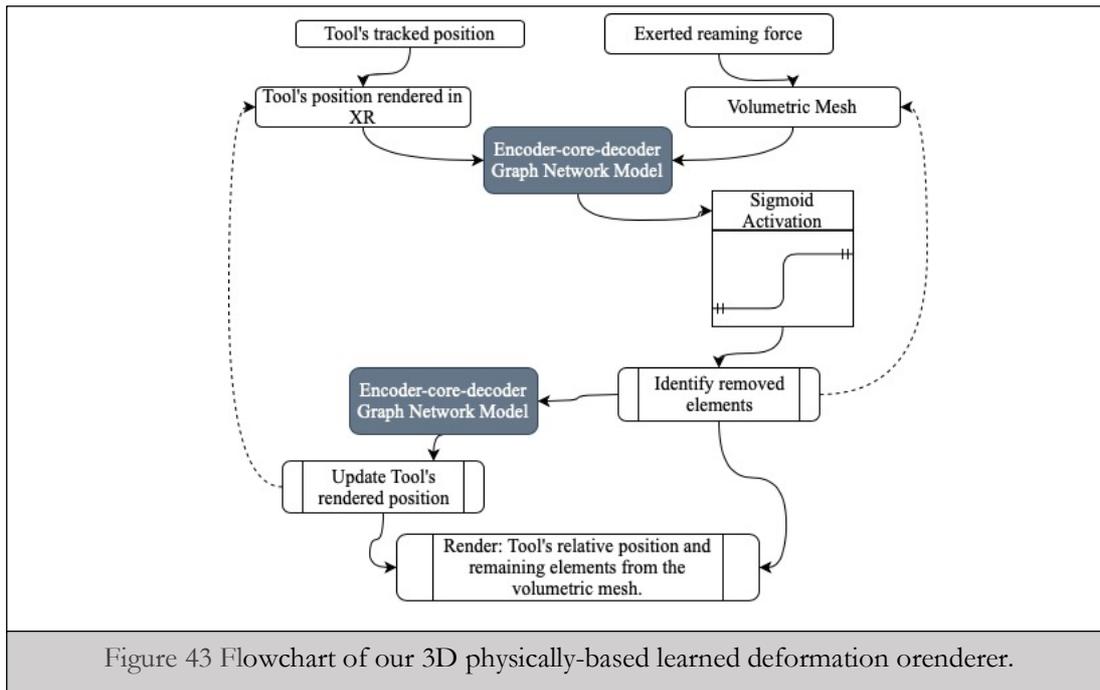
The word node in this document will refer exclusively to the entities in graph-structured data to avoid confusion. In contrast, the corners of a volumetric or geometrical mesh will be referred to as *vertices*.

Chapter 7

Machine Learning for reaming simulation.

This chapter describes the structure of the training data and learning strategies used to build a learned simulator capable of predicting the amount of removed bone during reaming operations. Each graph was formed by 20406 nodes corresponding to the number of tetrahedrons in a volumetric mesh processed via XFEM. In total, 4000 graphs were used for training and 1000 for validation. Finally, we describe the visualization pipeline responsible for applying the linear threshold to the output of the GN model and provide a 3D visualization of the resurfacing results. This final step provides the connection between machine learning into the computer graphics domain to ensure that the developed machine learning-based 3D deformation renderer is compatible with the design process of common VR/AR simulators.

None of the simulators available for THR identified in Section 2.2.1 provides an accurate representation of resurfacing procedures such as acetabular reaming based on the interactive force inputs exerted by the user. Furthermore, state of the art FEM and XFEM solvers are unable to deliver results at interactive framerates for meshes with over 5k elements (Section 4.2). Therefore, a novel approach is required to speed up the estimation of material removed during acetabular reaming and allow for interactive VR/AR simulation. As mentioned in Section 6.1.2, current machine learning models have only been used to estimate elastic deformations based on FEM data. In the literature review described in 6.2.2, we identified the GN (Battaglia *et al.*, 2018) as a suitable ML architecture to capture the physical relationships between elements of a graph and sustain accuracy over up to a thousand rollouts. Therefore, by using these models to estimate plastic deformations on



bone, we will introduce the current state of the art of data-driven and Machine Learning techniques that can be used to speed up these calculations and take advantage of their accuracy at interactive framerates.

7.1 Learned 3D physically-based simulator.

The second objective proposed for this research in Chapter 1 *train a Machine Learning (ML) model capable of estimating the removed material on a refined area of a volumetric mesh faster than traditional FEM methods?*. To be compatible with the current technological trends in surgical simulation, the ML model should take as inputs the readings from a triaxial force sensor and a 3D position tracker. Figure 43 lists the stages proposed for a deformation model to estimate the change on the volumetric mesh during reaming operations. During the interaction, the value of the accumulated deformation energy per volumetric element is dynamically updated in a similar way as the results obtained from the FEM simulation (Figure 39). This update is performed by a trained encoder-core-decoder GN model that reads the tracked position of the tool and the magnitude/direction of the applied force as its inputs. After the strains and accumulated energy in the volumetric mesh have been updated, the removed elements are easy to identify by applying a linear threshold to their accumulated energy value. Then, the volumetric change in the mesh is sent as an input to a separate model that estimates the relative displacement of the reamer based on the amount of material removed. Once both the mesh and tool's relative position have been computed, they are transformed into inputs

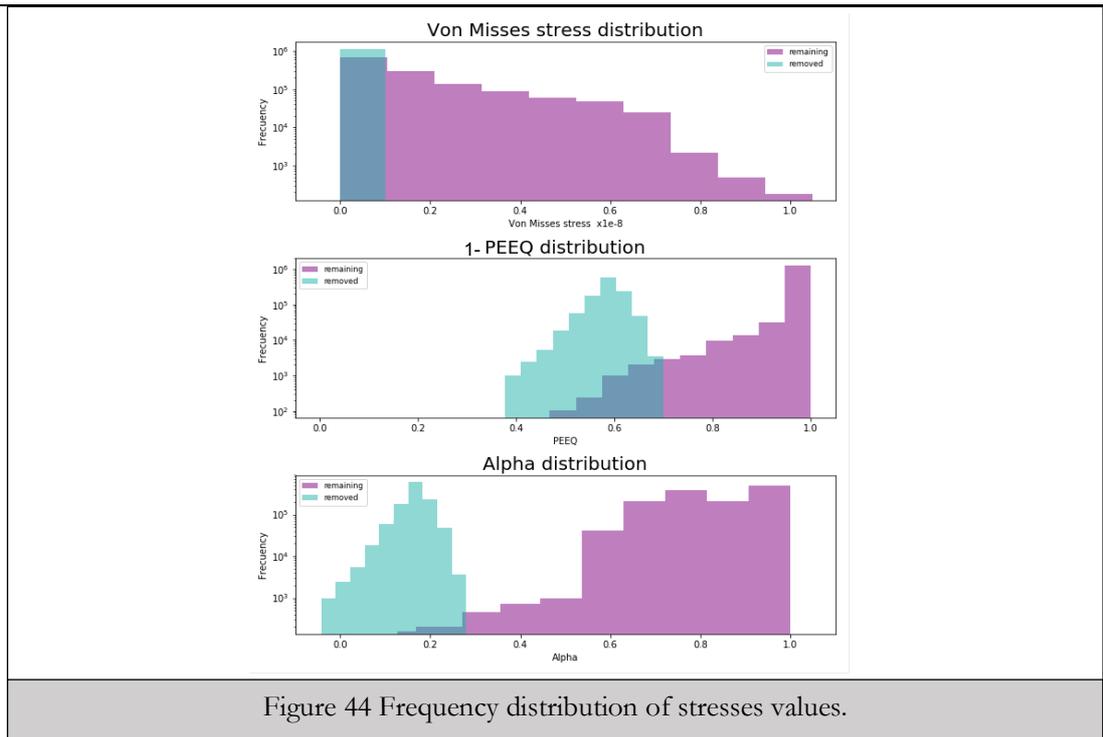


Figure 44 Frequency distribution of stresses values.

for estimating the status of the mesh in the next frame; and the process repeats.

7.2 Feature engineering and data cleaning

7.2.1 Feature selection

At first, identifying the elements removed per frame seems like a classification problem. However, a single-frame classification is not sufficient to ensure interactive energy accumulation during the learned simulation and ensure accuracy over several cumulative rollouts. Therefore, it is essential to build a machine learning model based on a regression step that handles the diffusion of energy inside the material during cumulative deformation. This section describes how results from finite element simulation are transformed into features for a Graph Network model.

Given that the FEM simulation of the acetabular reaming was executed on a dynamic solver, the results were stored in frames describing the dynamic evolution of the system per time. As discussed in Section 5.4, useful parameters to define the element status of each element will be the binary *STATUS* label, the Von Mises ($\bar{\sigma}_M$) equivalent stress and the PEEQ value. To better understand how each of these parameters can help a model identify removed elements from the mesh, Figure 44 illustrates the histograms with the frequency distribution of these metrics at a random frame. The top graph reinforces that even though ($\bar{\sigma}_M$) is a popular parameter used in failure analysis, it is not a perfect feature to separate

both classes as elements stresses are zero after failure. The middle histogram shows the frequency distribution of the PEEQ value. The histogram was inverted and shifted by computing $1 - PEEQ_i$ to produce a feature distribution that would yield lower values for elements that have been removed from the mesh and therefore should not contribute to the energy diffusion process. However, the values of $1 - PEEQ_i$, are still not sufficient to produce an adequate separation boundary between the two classes.

We propose a novel equation that combines the PEEQ and $\bar{\sigma}_M$ into the parameter α (47). This parameter makes it possible to acquire a more clean separation between both classes via a threshold value to assist the classification of removed elements. The labelling parameter α acts as a mapped probability distribution of the element to be still part of the base mesh or to have been removed due to strain accumulation. In other words, removed nodes yield values of α close to zero and elements under no stress condition will have a label of 1 and contribute to the aggregated passed message in the graph. In (47), μ and τ are scalar parameters that depend on the material constitutive equations and damage model chosen during FEM simulation (See section 5.2.1). For the material models used in Chapter 5, their values are 0.38 and 0.42, respectively. After transforming the PEEQ and VM values to α given the equation (47), the resulting frequency distributions of the elements *removed* and/or *remaining* (Figure 44) are easier to separate linearly given a *threshold* value determined during data exploration.

$$\alpha_i = 1 - (PEEQ_i + (\tau * \max(0, 1 - (\frac{\mu \sigma_M}{10^7}))) \quad (47)$$

Other geometrical features included during the training stage were the static position of each tetrahedron's centroid and the position of the tool measured at the centre of the reference system attached to the centre of the reamer's head as mentioned in section 5.3. The role of each one of these features in the graphs will be explained in the next section.

7.3 Graph network model

7.3.1 *Graph structure*

To fit into a graph architecture, each element from the FEM mesh was translated into a node of a graph. Undirected edges were created between tetrahedral mesh elements with adjacent triangular faces. Finally, the loading conditions are added as the global entity (Figure

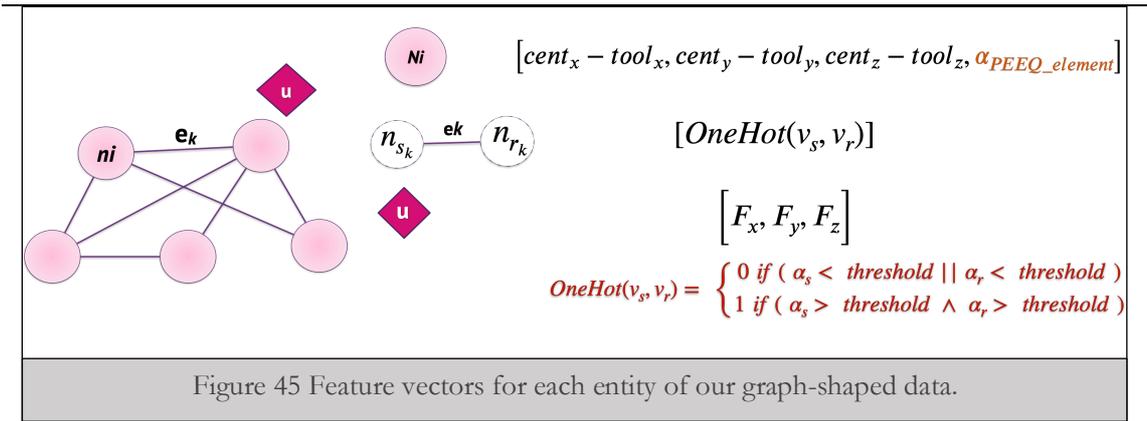


Figure 45 Feature vectors for each entity of our graph-shaped data.

45).

- n_i**: The features stored in the nodes hold the relative position of each tetrahedral element to the tool’s reference point and a parameter α per node. The labelling parameter α is computed to map the PEEQ and $\bar{\sigma}_M$ values to a probability distribution of the likelihood of each element to remain in the base mesh due to strain accumulation. In other words, the energy diffusion on the mesh weakens α at each element and therefore weakens its contribution to the graph's structure and probability to remain active. Therefore, removed nodes yield values of α close to zero and elements under no stress condition will have a label of 1. Adding this extra label to the attributes of the nodes creates internal attention labels that will encourage the model to aggregate features of active nodes over removed ones. The training data generated in the previous chapter has a single material model that is constant over the entire mesh domain. Hence, the properties of the material are not included as attributes of the graphs. However, the model is easy to adapt to compound models by adding the material properties corresponding to each group of elements in the mesh.
- e_{ij}**: The edges’ features vectors include a binary one-hot encoded representation f_{oneHot} of the status of both connected nodes (n_i, n_j) . This is, if for at least one of its nodes $\alpha < threshold$ then the edge will be labelled as “broken” in a one-hot encoded way. This status feature represents “present or broken” connections due to the progression of the machining operation.
- u**: This entity holds the three principal components x,y,z of the tool’s force measured at its centre of reference.

7.3.2 Graph network (GN) application

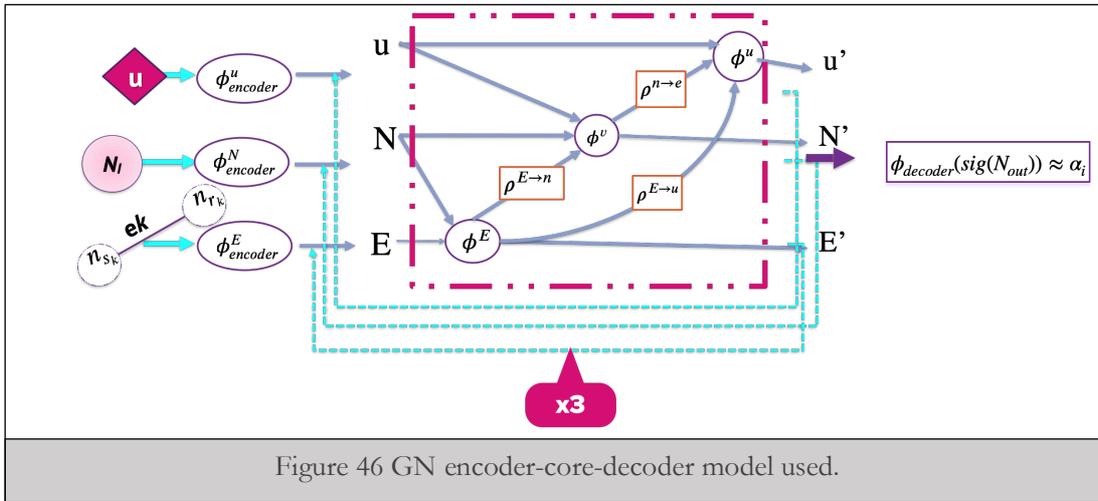


Figure 46 GN encoder-core-decoder model used.

The GN framework is based on early work referred to in the literature as the graph neural network model (GNN) (Gori *et al.*, 2005; Scarselli *et al.*, 2009b); a model for supervised learning on any graph-structured data. GNN proposed that a function τ that would be able to map the graph G , and its nodes n_i into an output feature vector $\mathbf{o} \in R^m$. Instead of flattening the data and setting it as input for an MLP, this type of learning conserve relational inductive biases for different types of graphs (Scarselli *et al.*, 2009a). An essential step of the GNN is the embedded aggregation of the features of the nodes to their connected neighbours. This node-to-node message passing step is performed recursively. By doing so, each nodes' "level of influence" increases by one connection on each iteration. Several authors have proposed mechanisms to enhance the message-passing between nodes (Zhou *et al.*, 2018). The GN framework was selected as the deformation estimator as it has effectively proven to be successful in learning simulations in mechanics over extensive rollouts while being trained on a single-step basis (Sanchez-Gonzalez *et al.*, 2018, 2019, 2020). Furthermore, its graph-in graph-out data flow facilitates the sequential predictions desirable for an interactive estimation of the surface changes during a reaming procedure.

The network structure used follows the encoder-core-decoder model shown in Figure 46. After the results from the FEM simulation are reshaped into the desired graph format (Figure 45), the encoding and decoding neural networks are trained to find the best representations for the GN block (Sanchez-Gonzalez *et al.*, 2018, 2019, 2020). The recursion in the GN block is performed k times to allow nodes to influence others up to k edges away. At the end of the network, a final linear transformation \mathcal{W} is applied to the updated decoded nodes attributes \mathbf{n}'_i to reshape the network output to the R^1 domain.

For our model shown in Figure 46, the update functions are formed by Multi Layer

Perceptrons and GN aggregators listed below:

- $\Phi_{encoder/decoder}^N$: One hidden layer with 40 neurons.
- ϕ^n : 2 hidden layers with 300 neurons each.
- $\Phi_{encoder/decoder}^E$: 2 hidden layers, 20 neurons each.
- ϕ^e : One hidden layer with 30 neurons.
- $\Phi_{encoder/decoder}^u$: 2 hidden layers, 11 neurons each.
- ϕ^u : One hidden layer with 20 neurons.
- $\rho^{l \rightarrow m}: \sum l ; \forall l \in U_m$

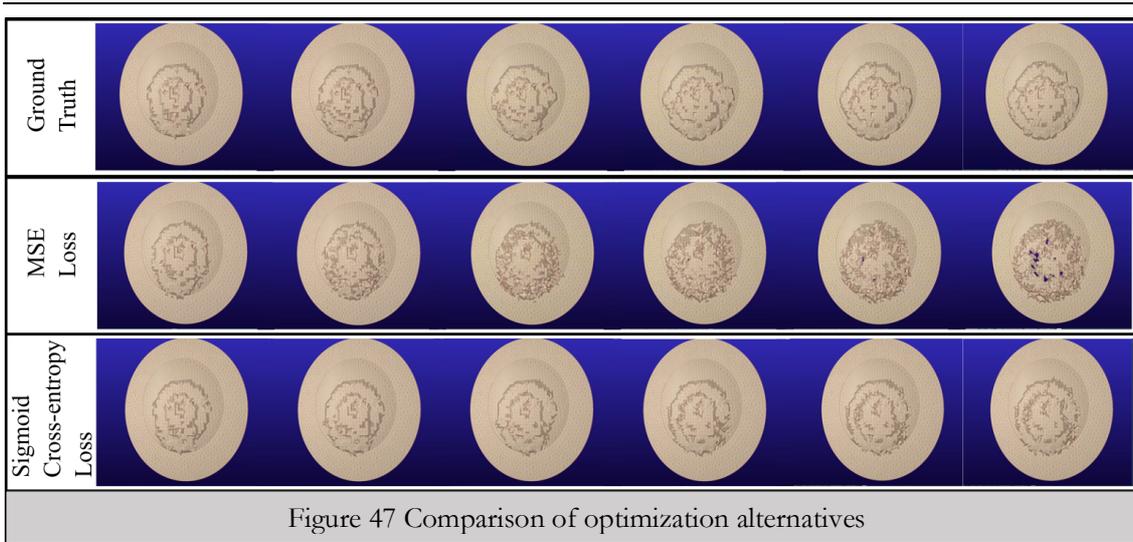
As shown in Figure 46, after the encoder-core-decoder model has generated an output \mathbf{n}'_i , this tensor is mapped to a probability distribution α'_i at $t+1$ computed as described in (48). The optimization of the α'_i as target variable is performed by minimizing the value of the sigmoid-cross entropy cost function $sig(N')$ as described in (49). Throughout the development of the volumetric learned simulator, the Mean Squared Error was explored as an option of the target cost function to minimize. However, Figure 47 exemplifies that optimizing the sigmoid cross-entropy cost produces a smoother output surface that is less influenced to the drift product of the accumulated error.

$$\alpha'_i = \frac{1}{1 + e^{-W \mathbf{n}'_i}} \quad (48)$$

7.3.3 Learning

$$err_w = - \sum_{i=0}^N (m_{in}^t \odot \alpha_i^{t+1}) \odot \ln(m_{in}^t \odot \alpha'_i) \quad (49)$$

Given that the GN was proven effective in maintaining accuracy on rollout given a single-step training by its authors (Sanchez-Gonzalez *et al.*, 2018, 2019, 2020), the training of the Graph Network was performed on a single-frame basis. During training, graphs were built on a frame t will be considered as inputs, and the updated stress-strain state will be referred to as $t+1$. In (49), since only nodes where $\alpha_i^t > threshold$ are considered to be active in the input graph, the cost function err_w is only evaluated on these nodes. To do so, the binary filter m_{in} is applied to both the predicted α'_i and target α_i^{t+1} values based on the status of each node at the input t . Therefore, the model neglects the updates on the energy of nodes that are considered inactive or removed in the input graph.

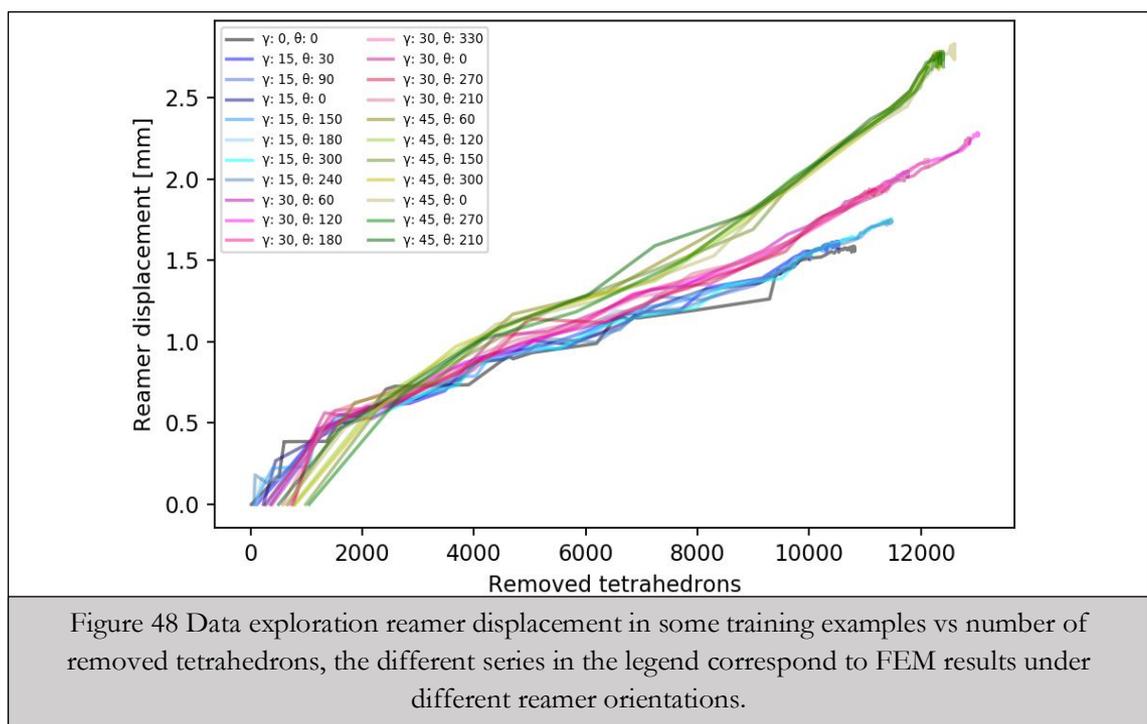


Once the accumulated energy per element α_i has been updated, elements that were considered *removed* from the base material at the input frame t , will keep their previous α_i value instead of continuing to update their attention label. The tool's displacement is updated given the removed tetrahedrons (See next section 7.3.4), and it is used to update the relative position between each node and the reamer. Finally, the edges and global features are reconstructed as described in section 7.3.1. Once the new graph is created, the updated α_i per element works as a long-term memory value inversely proportional to the element's accumulated stress values.

7.3.4 Regression for reamer displacement.

After a new group of nodes are *removed*³² from the graph on each frame, we must estimate the displacement of the tool to determine the penetration of its teeth into the machined material. However, in a system such as the one proposed in Figure 55, the tracked tool will not change its real position due to the physical constraints added by the passive haptic model. Therefore, after the first frame of interaction, only the tracked orientation is valuable, and the tool's displacement must be overlaid digitally in the rendered VR/AR environment. Estimating the displacement of the reamer is essential for valuable simulation-based learning experiences as it would not be possible to compute the relative position required in the graph architecture described in section 7.3.1 without an estimation of the updated tool's position.

³² The nodes of the graph are never removed from the graph but their presence is ignored as their α attribute yield 0. This behaviour corresponds to a removed tetrahedron in the FEM mesh.



Ideally, the Graph Neural Network is able to compute the displacement of the tool as an update of the global attributes in the graph. Nevertheless, optimizing the learning for two independent cost functions increases the complexity of the problem. Thus, the GN was trained to focus on the volume change of the mesh so that no trade out is necessary when optimizing for the alpha values and the tool's displacement. An independent Machine Learning model Φ was trained to estimate the displacement on the material removal tool during the interaction.

During the design stage of the 3D deformation renderer, several machine learning models Φ were trained with the same input and compared in their performances. The selected models included Random Forest, Support Vector Machine regressors, Ridge and Linear regression. A Grid Search was applied, varying the hyper-parameter for each model available in scikit-learn. The k-neighbours algorithm was not included in the comparison since such regressor would not be able to extrapolate over cumulative deformations not similar to the included in the training data. This project proposes a Random Forest regressor that takes as inputs position of the reamer in 3D space, the inclination and antroversion angles (θ and γ) and the number of new tetrahedrons classified as *removed* elements since the previous frame. The output of this regressor is an estimation of the displacement in the direction given by the angles θ and γ . This value is used later as input for an integrator to update the tool's position.

During FEM simulation, the direction of the force determines the unique direction along which the reamer is expected to move. Therefore, only the magnitude of such tool's penetration in the bone per frame is required. Figure 48 shows the correlation between the displacement magnitude at the frame $t+1$ against the amount of removed bone since frame t . In Figure 48, different colours correspond to FEM results under different reamer orientations.

$$\Delta position = \Phi(position, \Delta mesh, \theta, \gamma) \quad (50)$$

7.3.5 Results

The trained encode-core-decode learned simulator requires 579.11ms to build the input graph per each frame and to render the selected tetrahedral elements (Algorithm 1, 3-7 and 14). Once the data is shaped as an adequate graph, an average of 127ms are required to compute an output α_i' values (Algorithm 1, 8-12). This is 1639 times faster than the time required to generate the results per frame after the FEM simulation.

During the interactive simulation, the volumetric renderer reads the α_i' values of each node, corresponding to the accumulated reaming energy on each element. Then, it renders only the tetrahedrons with α_i' values above the threshold established to build the mask m_{in} during training. The computed output on frame $t+1$ is used as the input to compute the next frame $t+2$. In Figure 51, the GN's prediction is computed interactively while the Ground Truth is read from a database containing the results of a FEM simulation. Figure 51(b) illustrates the resurfacing results by using the mesh in Figure 51(a) as input, and the tool's orientation of $\theta:15^\circ, \gamma:45^\circ$. Figure 51(d) shows the resulting surface after 60-step rollouts.

Algorithm 1 GN for machined surface prediction.

```

1: While Machining do
2:   Input Graph connectivity, tool's position, tetrahedral centroid position,
    $\alpha'_i{}^{t-1}$ , tool's axial force.
3:   Build  $\mathbf{G}_{in}$ :
4:      $rel_i \leftarrow (tetrahedron_{pos} - tool_{pos})$ 
5:      $\mathbf{n}_i \leftarrow (rel_i, \alpha_i)$ 
6:      $\mathbf{e}_{ij} \leftarrow (f_{oneHot}(n_i, n_j))$ 
7:      $\mathbf{u} \leftarrow (\mathbf{force}_{axial})$ 
8:      $\mathbf{G}_{enc} \leftarrow \Phi_{encoder}(\mathbf{G}_{in})$ 
9:      $\mathbf{G}' \leftarrow \mathbf{G}_{enc}$ 
10:    for k steps do
11:       $\mathbf{G}' \leftarrow GN(\mathbf{G}_{enc}, \mathbf{G}')$ 
12:       $\alpha'_i \leftarrow sigmoid(W\Phi_{decoder}(\mathbf{G}'))$ 
13:      if  $\alpha'_i > threshold$  then
14:        Render  $tetrahedron_i$ 

```

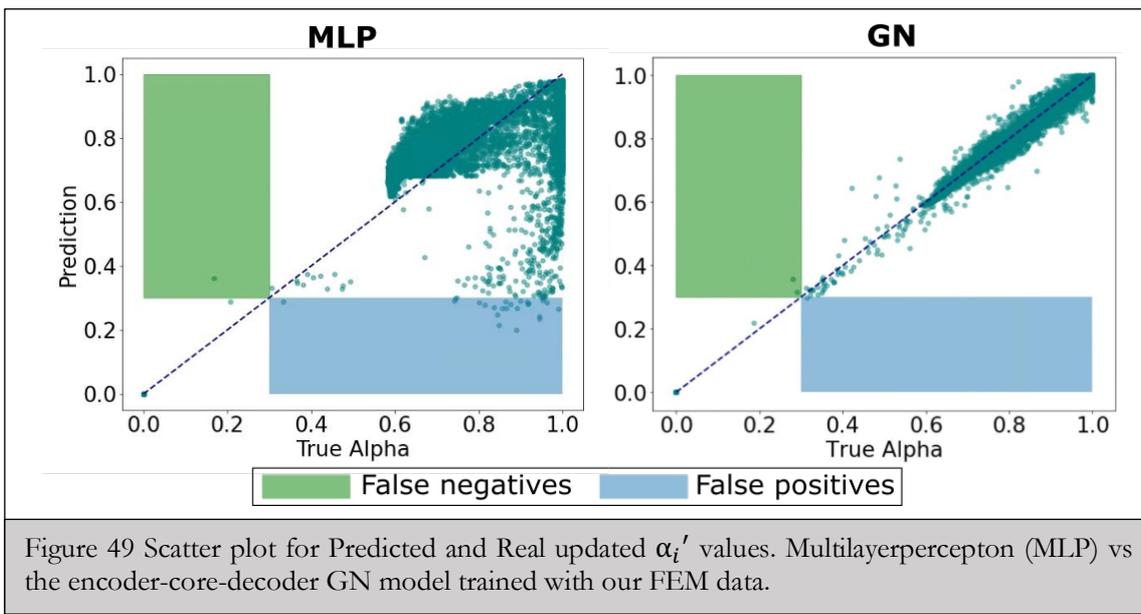
The tool's geometry was hidden to illustrate the changes on the material's surface. In Figure 51 (b,c,d) the tool is displaced and rotated thanks to its axial load, but it was hidden to display the change of status on elements of the mesh. The remaining unconnected tetrahedrons at the bottom of the ground truth mesh are a result of the encastre boundary condition set at the bottom vertices of the deformable mesh during the FEM pre-processing stage.

Given the size of the mesh, there is clearly a class imbalance between the elements classified as removed and remaining/active. Therefore, accuracy is not a sufficient metric for determining the performance of the trained GN model. Precision (51) and Recall (52) are more adequate performance metrics for such imbalance.

$$Precision = \frac{truePositives}{truePositives + falsepositives} \quad (51)$$

$$Recall = \frac{truePositives}{truePositive + falsenegatives} \quad (52)$$

$$F1score = 2 \frac{precision * recall}{precision + recall} \quad (53)$$



Given an unseen orientation of the reamer, the trained GN model updates the value of accumulated stress on each of the 20406 tetrahedral elements that constitute our mesh during the prediction task. Once this value is updated, the tetrahedrons to be removed from the mesh are identified using a threshold condition. After using each single-frame output as input for the following prediction repeatedly for up to 60 iterations, our model can maintain an accuracy of up to 90.8% in identifying the status of each element given their value of accumulated stress. The results for the mean accuracy and precision metrics on unseen data are listed in Table 10. These values are computed as an average for the listed angles γ over values of θ distributed between 0° and 360° . These performance measurements are computed relative to the tetrahedrons classified as removed from an initial mesh state t at the end of the rendering pipeline. It is worth noting that the definition of precision and recall makes them unstable when only a small number of elements (*true positives*) is removed. In other words, in cases where both the ground truth and the learned simulator removed 0 elements, these metrics yield values of 0 even though a good performance was observed. .

The Graph Network shares information between the nodes of the graph (tetrahedral finite elements) during the learning stage according to the relationships described by the graph's edges. To illustrate the advantage of including the geometrical relationships between entities, we compared the predicted α_i' from a trained GN block and a Multi-Layer Perceptron (MLP) in Figure 49. The Multilayer perceptron (MLP) model used had the same architecture as the network ϕ^n described in section 7.3.1 as the one responsible exclusively from the update of the node's attributes. Both GN and MLP models were trained for a total

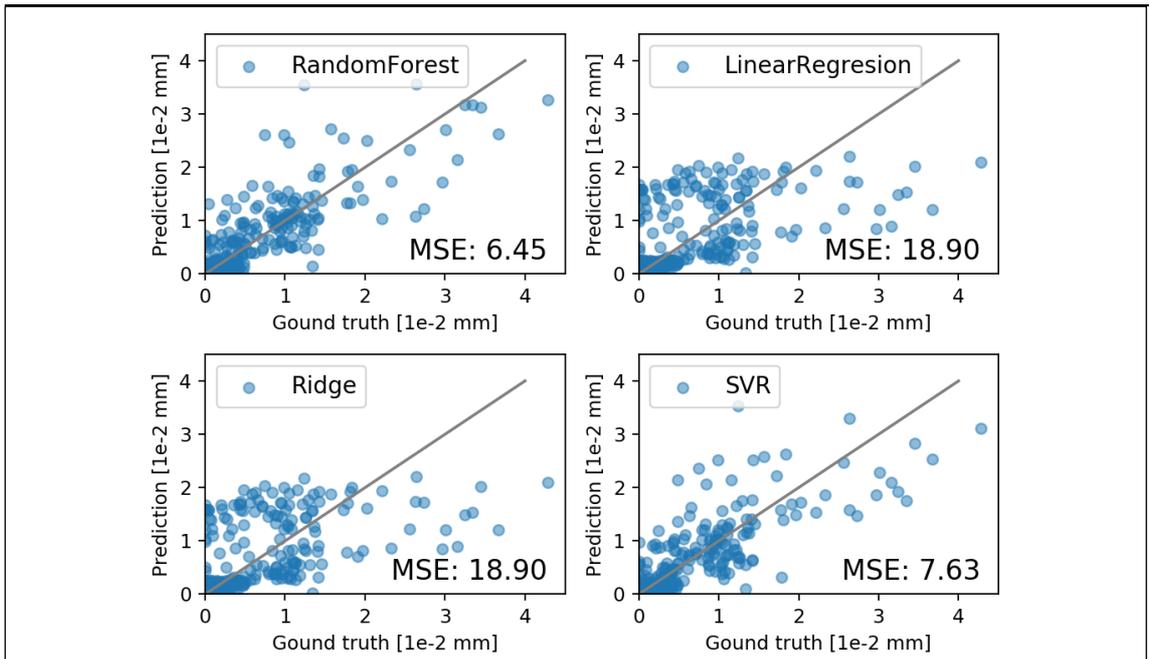


Figure 50 Comparison of regression algorithms for prediction of reamer displacement. The estimations were evaluated from the test data.

of 350 epochs. The GN encode-process-decode model holds a better correlation between the ground truth and the predicted α_s . It is also highlighted that for the training data used and labelling the elements as removed when their α_i' is lower than 0.3, we obtained the lowest quantity of false negatives and/or false positives when using GN based model.

γ	$t+1$			$t+30$			$t+60$		
	Acc	Rcll	F1	Acc	Rcll	F1	Acc	Rcll	F1
0	75.3	79.4	77.3	76.2	98.8	86.1	90.8	95.3	92.9
15	69.3	59.4	63.97	92.1	93.5	92.8	92.2	93.5	92.8
30	69.3	58.1	60.6	85.2	85.1	85.1	85.5	84.7	85.1
45	76.7	71.1	73.8	86.1	83.3	84.7	85.2	87.1	86.1

Table 10 Accuracy (Acc) and Recall (Rcll) scores for selected removed elements in mesh after 1,30 and 60-step rollouts on the test dataset. Values in percentages.

Figure 50 shows the correlation scatter plot for the ground truth and predicted values of reamers displacement magnitude on different estimators evaluated on unseen data. As

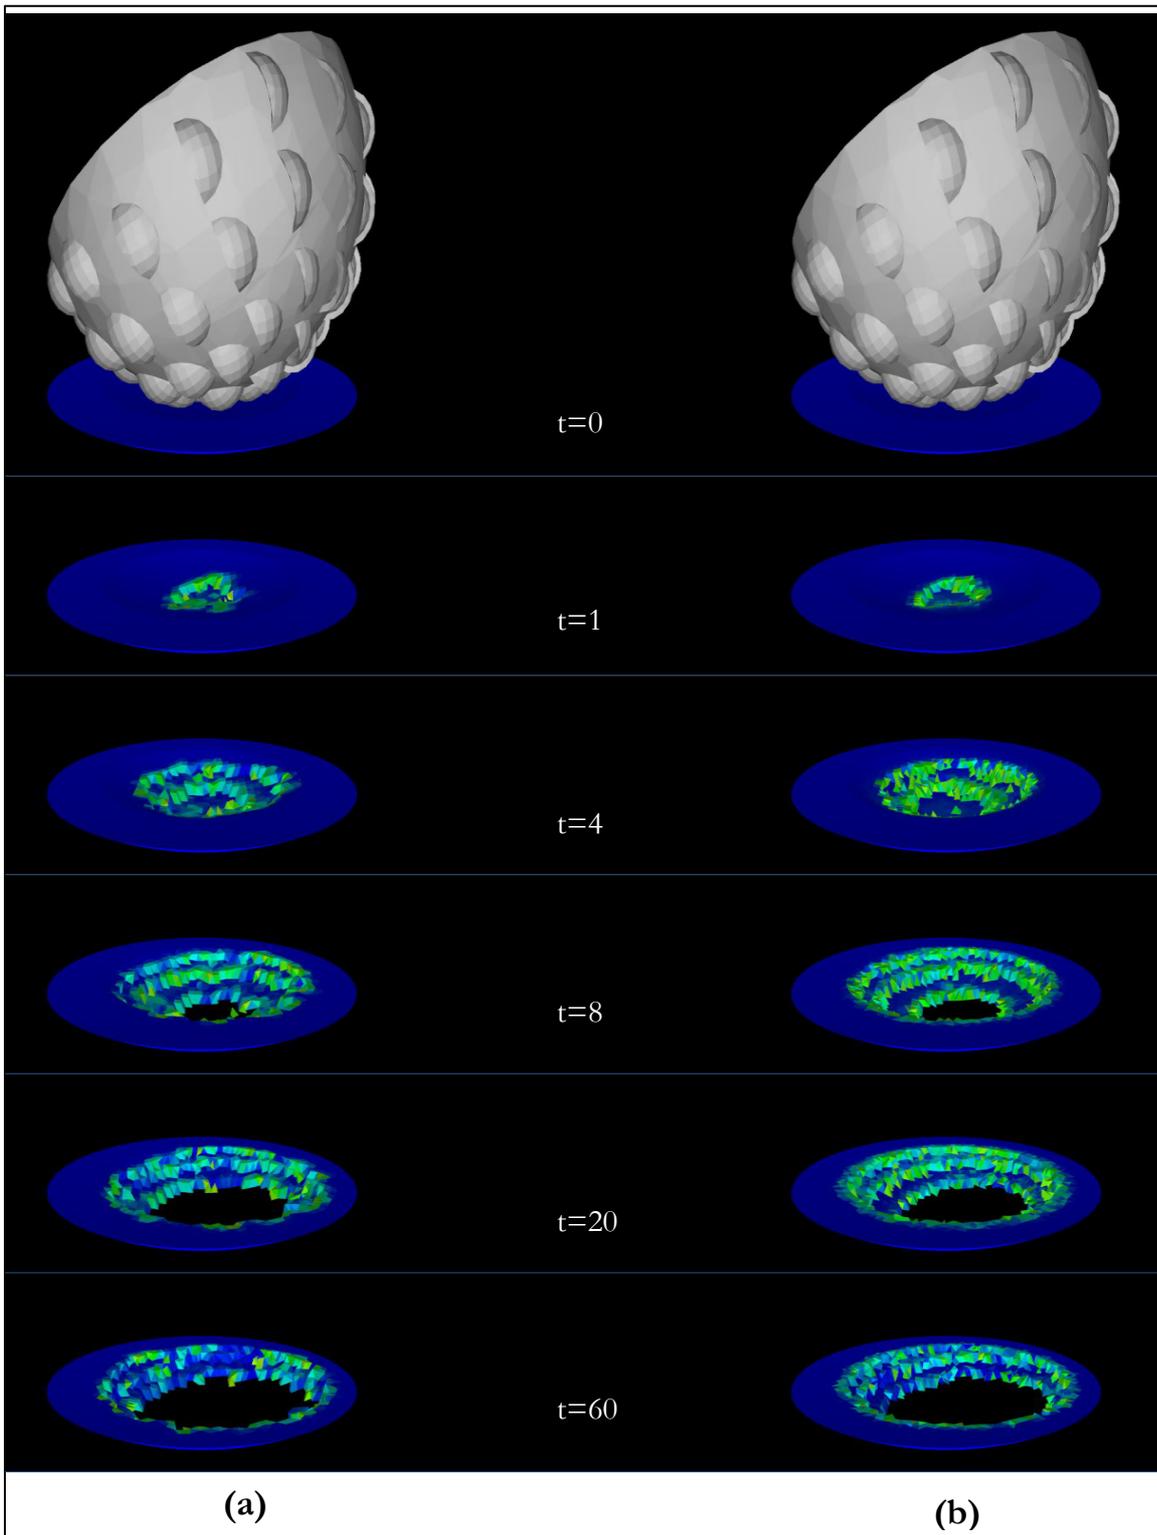


Figure 51 Results of resurfacing on curved surface. (a) Heatmap of ground Truth values on each tetrahedron (b) Heatmap representation of predicted values of accumulated stress on each tetrahedron on frame t .

mentioned in Section 7.3.4, a Random Forest model with 100 trees was selected as a regressor to estimate the displacement of the tool due to the removal of material as its MSE was the

smallest one from all considered regression models.

7.4 Modern GL – openGL for python

During the simulation, the nodes' values of the graphs representing the volumetric mesh are dynamically updated on each output frame. As described in Chapter 1, VR/AR systems are able to track the tool's position and connected force exerted by the user to estimate an interactive resurfacing update on the machined surface. At the initial input frame t , the values of the node's labels α_i are updated. It is worth noting that elements considered to have been *removed* from the base material already will keep their previous α_i value. The tool's path is also updated given the removed tetrahedrons, and it is used to compute each node relative distance feature.

Finally, the edges and global features are further constructed, as described in section 7.3.1. Once the new graph is created, the updated α_i work as a long-term memory value of the accumulation effect of the remaining energy. Since the total number of nodes and edges in the graph remains immutable during the entire simulation, the selection of tetrahedral elements is made inside the simulator's visualization pipeline. Since the total number of nodes and edges inside the graph remain immutable during the entire interaction, the selection of tetrahedral elements is done inside the simulator's visualization pipeline. The rendering pipeline is based on the one provided for volumetric meshes inside the modernGL library (Forselv 2020) as it was easy to connect to the results of the GN API. During the interaction, the geometry shader reads the updated α_i' values per element and constructs the tetrahedrons

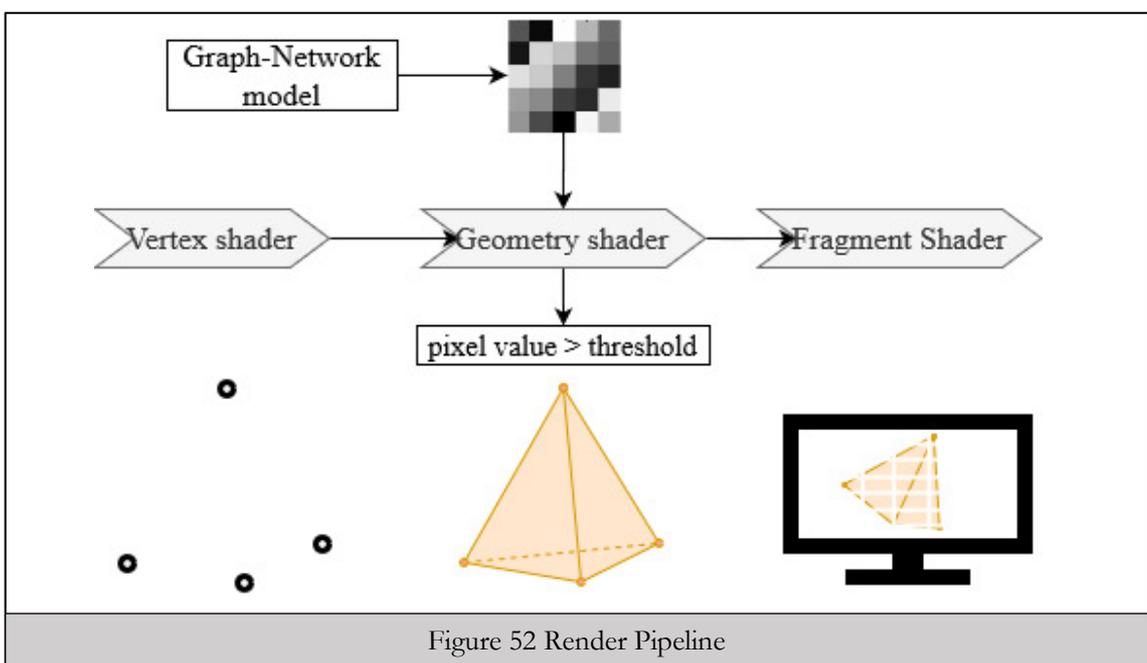
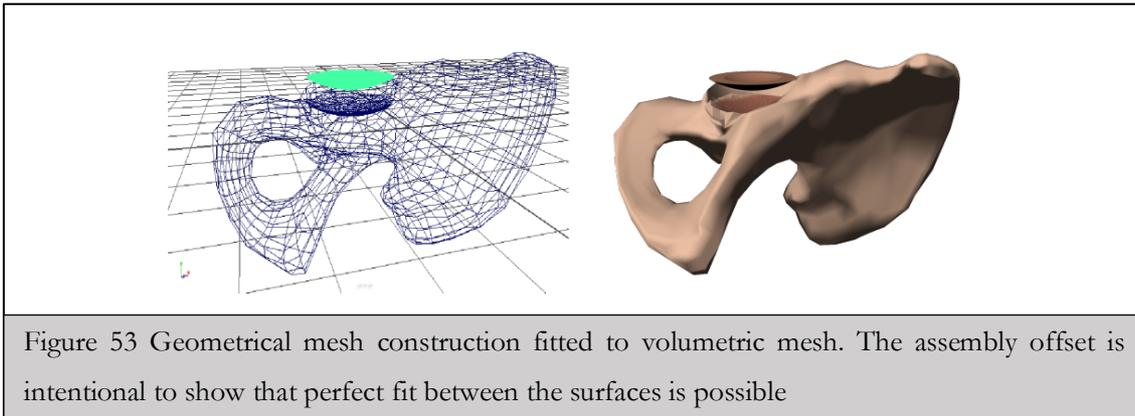


Figure 52 Render Pipeline



where the threshold condition is satisfied. The updated α_i' values feed the construction of the new edges and nodes features for the graph used as input for the next frame.

Once the updated α_i' s are estimated for each frame, they are written into a 2D greyscale texture of height H , and width W : where $H \times W$ is equal to the total amount of nodes in the output graph and/or the total amount of tetrahedral elements in the volumetric mesh Figure 52. The primitive identifier available in OpenGL allows the geometry shader to match each pixel of the 2D texture with its corresponding tetrahedron. Therefore, each pixel of the texture corresponds to the α_i' value of each tetrahedron. The vertex coordinates are passed from the vertex shader along with the texture to the geometry shader, which reads the grayscale value in the texture and only renders the triangular primitives for tetrahedrons with α_i' above the defined *threshold*. Once the triangular faces and face's normals are built, these values are passed to the fragment shader responsible for the final image rasterization on the display device.

As mentioned in Section 5.1.1, given the high GPU memory requirements for training GN-based models with fine volumetric meshes, it would be inefficient to simulate areas of the pelvic bone that the user will not be interacting with during simulation-based training. Furthermore, the FEM output for a bigger mesh would be exponentially larger as the total size of the meshed object would increase. Alternatively, we propose to re-align the volumetric mesh extracted from its geometrical parent mesh, as demonstrated in Figure 53. After achieving a visually pleasing integration of the volumetric model with a geometrical mesh, the rendering pipeline can be linked with other 3D geometrical meshes and fed into a game engine platform to help achieve visual immersion and provide the interaction desired for the training environment (Figure 54).

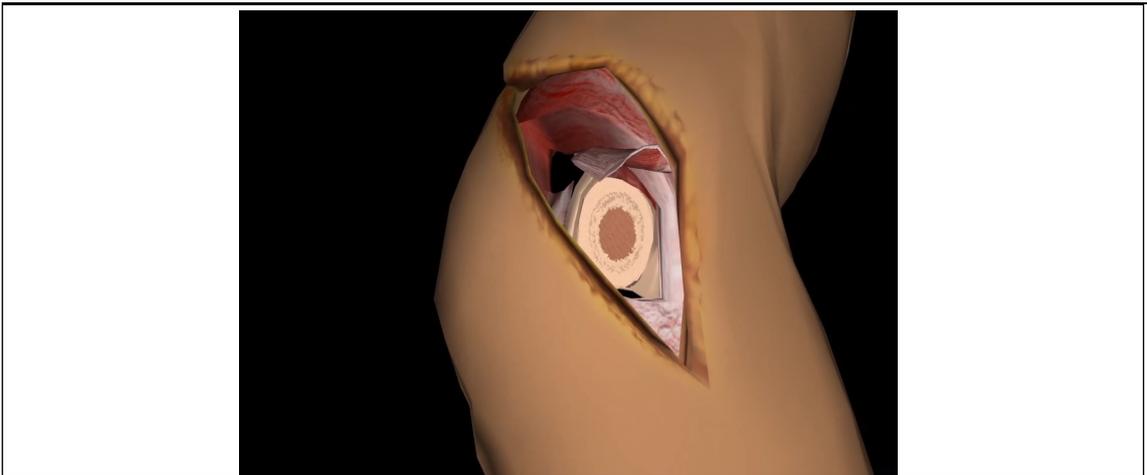


Figure 54 Acetabular cavity exposed in hip model AR headsets. For Augmented reality applications the background is set to black corresponding to the absence of projected light on top of the environment.

7.5 Summary

This chapter explains the training methodology for a GN-based model capable of estimating the physical results obtained from our FEM simulation. The proposed model computes an accurate forecast of cumulative plastic deformations based on the force applied on the reamer and a given volumetric mesh.

Our GN-based model not only identifies the elements to be removed on each frame but also estimates the accumulated stress evolution in all the elements conforming a volumetric mesh. Using data generated from the finite element method (FEM) results, the GN can learn to estimate the effects of nonlinearities such as friction or damping without extra processing time. Given that the GN requires its input and output graph to have the same number of nodes, we included self-attention labels to allow the model to distinguish between removed and remaining active elements.

Our GN model can also be connected to a visualization pipeline similar to the one included in several game engines used to develop VR/AR simulators.

Chapter 8

Conclusions and future work

One of the common limitations of open surgery simulators is the use of commercial haptic units to provide feedback. These devices shared a restriction to be implemented to replicate orthopaedic surgeries. This is, during joint replacement, the forces involved can reach values up to 2kN during implant fitting through impact and 300N during bone resurfacing stages (Mathieu *et al.*, 2013). PHANTOM devices, developed by 3D systems, are the most popular implemented models. However, the technical specifications of their most capable devices report a limited tool's range of motion relative small to provide the movements often required during surgery, a small handler's size and maximum exerted force below 40N (3D Systems, 2017). These technological constraints make the human-machine interaction unnatural and reduce the value of the transferred skills.

As a consequence of this technological gap, current VR simulators lack training arthroplasty skills with accurate haptic representation so that the skills can be translated into the theatre. Enhanced algorithms and hardware designs are required to build simulators easy to be accredited and incorporate in the training curriculum with certainty on their impact on skills acquisition. For that reason, this project proposed a solution with an alternative way of interaction, using passive haptics that would provide a more accurate force input into the deformable models.

In this thesis, we have used gaming strategies to create a novel proposal for an AR simulator to help trainee surgeons understand imageless navigation systems' principles with lower financial investments than assigning one of these devices for training purposes. This simulator helped surgeons to understand how to set up an INS inside their surgical theatre

and manipulate their tools without interfering with the tracking performed by the infrared cameras on both the patient and the robot. All participants who try the simulator successfully finished all the training steps required and expressed their comfort using the HoloLens as a headset. In addition, all of them agreed that the simulator allows an understanding of the skills needed in INS. Most participants believed that the system was easy to use and it had attractive visual content. Additionally, we developed an alternative communication protocol between an AR headset and an external tracking device. Such protocol is useful for headset devices with limited computing capacity that must be reserved for graphical computations or to enable the usage of advanced accurate position tracking systems.

The Johnson-Cook (JC) model integrated into the ABAQUS FEM solver was chosen as a numerical representation of the effect of the reamer procedure over bone tissue. This model has been validated in the literature as an accurate representation of the constitutive equation of the bone during cutting operations. JC is easy to implement in commercially available FEM simulators, and its required computational times allowed to create a dense database of simulated cases with relatively low pre-processing requirements. The fracture criteria are dependant on the evolution of the PEEQ stress. Therefore, this parameter provides a valuable metric to build features vectors for a machine learning model to describe the evolution of energy diffusion due to the interaction with the cutting force applied by the tool's cutting edges.

This thesis presents an ML model capable of learning to simulate the removed material on a refined area of a volumetric mesh faster than traditional FEM methods. This efficient interactive model can estimate the evolution of surfaces undergoing machining operations with results comparable to accurate numerical simulations. The trained ML model does not involve any bounding box simplification for collision detection or complex models for geometry contact evaluation to compute its output. Furthermore, by using outputs from a FEM analysis as training examples, it was possible to include in the estimations the effect of non-linearities usually neglected in an interactive simulation, such as friction between the tool and machined piece.

The proposed parameter α , which describes the accumulation of reaming energy, and the relative position of each tetrahedral element to the reamer, are adequate features to include on the features stored as attributes of the nodes of a graph. Likewise, the binary encoded status of the edges works as an attention mechanism to identify broken connections between the removed element and therefore adds a barrier to the energy propagation through

such edges. The trained GN-based reaming model returned the updated α'_i requiring in average 706.11ms to compute the output at each frame. Compared with the results produced by finite element simulation, this is more than 1639 times faster. Since only 127ms are taken from the GN model, parallelization techniques using multiple GPUs can speed up constructing the input graphs and improve framerates. This time efficiency on estimating FEM results is much faster than the one reported by Chen et al. (2015b) where results were obtained 489 times faster than the numerical FEM simulation.

The GN encode-core-decoder learned simulator trained in this thesis differs from the GNS physics engine proposed by Sanchez-Gonzalez (2020), as in our analysis, the relative movement between particles is not enough to describe machining operations. Our graph architecture is similar to the one used in the GNS. However, instead of computing the relative movement between particles, the 3D reaming deformation renderer focuses on the stress propagation that triggers the fracture of the tetrahedral primitives from a volumetric mesh. Moreover, instead of allowing the graph to interact only with physical boundaries in the simulation, we also included the influence of an external tool as a global entity affecting all elements of the graph to allow user's interaction. Similarly, our learned simulator computes the deformation energy per volumetric element of a mesh at each time step instead of requiring an additional integration step. Finally, unlike a particle-based system, the connections between the nodes of our graph never change since there is no relative movement between the tetrahedrons conforming the mesh. Therefore, we maintain the graph structure static and add dynamic attention labels to determine whether a graph node still has influence over its neighbours or if it has been *removed* from the graph.

The learned simulator developed throughout this research also predicts the global position of the centre of the reamer and its orientation according to the current bone volume distribution and force applied on a specific frame. Compared to MLP, the GN-based model was able to have a higher correlation between the predicted values and the ground truth. However, it is worth noting that learning on the graph domain requires significantly more memory than treating the stress-load status on each element individually. This is because of the parameters that determine the graph topology and the extra information stored in the edges.

An interesting behaviour is illustrated in Figure 51 where some elements at the bottom of the ground truth mesh are clearly unconnected and they appeared to be not yet removed

from the volumetric mesh. This occurred due to the boundary condition applied to the elements at the bottom face of the mesh, which hinders their displacement in all directions. Since the trained GN simulator is agnostic of those boundary conditions, it learned to remove loose elements from the mesh without remaining debris at the bottom. Additionally, even though the encode-core-decode model is completely agnostic to the tool's geometry, the results reflect the spherical footprint of the reamer. Moreover, on a visual comparison between the surface resulting from both the FEM simulation and the GN estimator, the GN's generalisation produces a smoother output surface than the FEM ground truth.

8.1 Future work

- In the simulator presented in Chapter 3, the limited number of available orthopaedic surgeons in the Bournemouth area constrained the assessment of our INS AR simulator to a face validation³³ test. Furthermore, conducting further test required human participants (surgeons) who were classified as front-line NHS workers during the COVID-19 global pandemic. However, to increase the confidence in the simulator to be used in formal medical training, it is important to assess its characteristics and technical specification given according to the validation strategies from Schijven & Jakimowicz (2005). Construct validity is an evaluation of whether the simulator can identify a novice from an expert user showing a contrast between the levels of expertise (Schijven and Jakimowicz, 2005). To this aim, it is crucial to count on a larger number of experts than the ones available in our area. We also suggest performing a content validity study. Content validity refers to the relevance and correctness of the abilities to be trained and their value in the real-life situations. Content validity studies usually identify undesired skills transferred from simulation to the real surgery. To perform a content evaluation, it is necessary to assess trainees' performance using INSs after training with the AR simulator against a control group. For THR procedures, a surgeon's performance is assessed given the stability of the implant system after a 5-year follow-up as complications can take a long time to develop (Schnurr *et al.*, 2009). Unfortunately, this challenge was impossible to overcome, given that

³³ An evaluation of the similarity between the skills that can be acquired using the simulator with the ones in the real world.

our research was constrained to 3 years.

- Our trained model for implementation in a simulator can also be enhanced in different facets to achieve better generalization. Many different adaptations, training data, and experiments have been left for the future due to lack of time. Firstly, the FEM results used as training data were generated covering a portion of the acetabular cavity with a mesh of only 20406 elements big. By refining the mesh, the network could be trained on more refined behaviour producing smoother results. Additionally, the dimensions of the volume extracted from the acetabulum can also be enlarged to include a more comprehensive portion of the bone into the interactive volume of our physically-based deformer. Furthermore, the training data was generated by solving the dynamic system for the first 2 seconds after contact is started. This duration was enough for a thin mesh to be perforated entirely by the reamer. However, if the volume of volumetric mesh is increased, a bigger tetrahedral mesh can be simulated during longer time ranges, increasing the model's generalisation for longer interaction times.
- We believe that enhancing the accuracy of the material's constitutive equations used to define the FEM is very likely to increase the accuracy of the FEM results. However, a more complex model is also likely to increase exponentially the processing times that the FEM solver requires to find the system of dynamic equations described in Section 4.1.3. To make the constitutive equations more realistic, we suggest including the anisotropy of the bone. As mentioned in Section 4.3.2, the direction of anisotropy for the acetabulum, i.e the orientations of the osteons and the strength of the bone of the principal directions, which have to be characterized experimentally to define an anisotropic model successfully. These topological and mechanic characterization of the pelvic bone can be considered a whole research project on its own, and due to our time constraints, they were not included in our project. The osteons orientation is typically found through microscopic examination of a sample from the acetabulum. The bone sample should be pre-treated with a pigment to enhance the fibre-like orientation of the osteons forming the bone. Once the direction of anisotropy has been determined, the mechanical failure test must be conducted to identify the strength of the

the user's movements to be constrained by a haptic arm. Instead, the movements of the reamer and its exerted force are monitored in real-time by adequate sensors such as a triaxial force sensor and a 3D position tracker. To provide unlimited attempts inside the simulation, the physical reamer dummy should not induce any sign of wear on the surface of the passive haptics models. Therefore, the simulator must take advantage of the spatial alignment to display the progress of the reaming operation on the surface of the digital pelvis.

References

- 3D Systems, I. (2017) 'Specifications for Touch X haptic device'. Available at: https://uk.3dsystems.com/haptics-devices/touch-x?_ga=2.233009757.253924424.1571306723-986983799.1571306723.
- 3Dsystems (2017) 'Arthro-mentor'. Available at: <https://symbionix.com/simulators/arthromentor/>.
- Aggarwal, R., Ward, J., Balasundaram, I., Sains, P., Athanasiou, T. and Darzi, A. (2007) 'Proving the effectiveness of virtual reality simulation for training in laparoscopic surgery', *Annals of Surgery*. LWW, 246(5), pp. 771–779. doi: 10.1097/SLA.0b013e3180f61b09.
- Aguilera-Canon, M. C., Wainwright, T., Yang, X. and Nait-Charif, H. (2019) 'Mixed Reality-Based Simulator for Training on Imageless Navigation Skills in Total Hip Replacement Procedures', in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, pp. 30–34. doi: 10.1007/978-3-030-23712-7_5.
- Alam, K., Mitrofanov, A. V. and Silberschmidt, V. V. (2009) 'Finite element analysis of forces of plane cutting of cortical bone', *Computational Materials Science*. Elsevier, 46(3), pp. 738–743. doi: 10.1016/j.commatsci.2009.04.035.
- Alexander, J. W., E. Kamaric, and P. C. N. (1999) 'The accuracy of acetabular reaming in total hip replacement', *Acc*, 1200, p. 800.
- Alimisis, D. and Zoulias, E. (2013) 'Aligning technology with learning theories: A simulator-based training curriculum in surgical robotics', *Interactive Technology and Smart Education*. Emerald Group Publishing Limited, 10(3), pp. 211–229. doi: 10.1108/ITSE-05-2013-0011.
- Antonis, J., Immins, T., Bahadori, S., Wainwright, T. W., Gallagher, K. and Middleton, R. G. (2019) 'Validation of the anterior cruciate ligament (ACL) module of the virtamed virtual reality arthroscopy trainer', *Surgical Technology International*. Bournemouth University, Fern Barrow, Poole, Dorset, BH12 5BB, UK, 35, pp. 1–8.

- Apple (2018) *ARKit: Apple Developer Documentation*, Apple Developer.
- ASTM International (2004) *A276-04 - Standard Specification for Stainless Steel Bars and Shapes*. West Conshohocken, PA. doi: 10.1520/A0276-06.
- Baby, B., Srivastav, V., Singh, R., Suri, A. and Banerjee, S. (2016) 'Idfwruv Dqg Wkhlu Dssoldwlrq Lq Vxujldo Vnloov', in *2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom)*, pp. 2564–2569.
- Banihani, S., Rabczuk, T. and Almomani, T. (2013) 'POD for real-time simulation of hyperelastic soft biological tissue using the point collocation method of finite spheres', *Mathematical Problems in Engineering*, 2013. doi: 10.1155/2013/386501.
- Baro, D. K. and Deoghare, A. B. (2018) 'Investigation of cutting force and temperature variation in orthogonal cutting of cortical bone', *Materials Today: Proceedings*. Elsevier Ltd, 5(5), pp. 12595–12608. doi: 10.1016/j.matpr.2018.02.242.
- Battaglia, P., Pascanu, R., Lai, M., Rezende, D. and Kavukcuoglu, K. (2016) 'Interaction networks for learning about objects, relations and physics', *Advances in Neural Information Processing Systems*, pp. 4509–4517.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., et al. (2018) 'Relational inductive biases, deep learning, and graph networks', pp. 1–40. Available at: <http://arxiv.org/abs/1806.01261>.
- Beer, F. P. (2011) *Statics and Mechanics of Materials, Connect, learn, succeed*. doi: 10.1080/03043799608913030.
- Berkley, J., Turkiyyah, G., Berg, D., Ganter, M. and Weghorst, S. (2004) 'Real-Time Finite Element Modeling for Surgery Simulation: An Application to Virtual Suturing', *IEEE Transactions on Visualization and Computer Graphics*, 10(3), pp. 314–325. doi: 10.1109/TVCG.2004.1272730.
- Beverland, D. E., O'Neill, C. K. J., Rutherford, M., Molloy, D. and Hill, J. C. (2016) 'Placement of the acetabular component', *Bone and Joint Journal*. British Editorial Society of Bone and Joint Surgery, 98B(1), pp. 37–43. doi: 10.1302/0301620X.98B1.36343.

Bhaskar, D., Rajpura, A. and Board, T. (2017) ‘Current concepts in acetabular positioning in total hip arthroplasty’, *Indian Journal of Orthopaedics*. Wolters Kluwer--Medknow Publications, 51(4), pp. 386–396. doi: 10.4103/ortho.IJOrtho_144_17.

Bhise, O. P. and Pratihari, D. K. (2006) ‘Neural network-based expert system to predict the results of finite element analysis’, in *Advances in Soft Computing*. Springer, pp. 231–240. doi: 10.1007/978-3-540-36266-1_22.

Bickel, B., Bächer, M., Otaduy, M. A., Matusik, W., Pfister, H. and Gross, M. (2009) ‘Capture and modeling of non-linear heterogeneous soft tissue’, in *ACM Transactions on Graphics*, p. 89. doi: 10.1145/1531326.1531395.

Bloom, B. S. (1965) *Taxonomy of Educational Objectives: Cognitive domain (Vol. 1)*, New York: McKay. New York: McKay. Available at: [http://nancybroz.com/nancybroz/Literacy_I_files/Bloom Intro.doc](http://nancybroz.com/nancybroz/Literacy_I_files/Bloom%20Intro.doc).

Blyth, P., Stott, N. S. and Anderson, I. A. (2008) ‘Virtual reality assessment of technical skill using the Bonedoc DHS simulator’, *Injury*, 39(10), pp. 1127–1133. doi: 10.1016/j.injury.2008.02.006.

Bruyns, C. D., Senger, S., Menon, A., Montgomery, K., Wildermuth, S. and Boyle, R. (2002) ‘A survey of interactive mesh-cutting techniques and a new method for implementing generalized interactive mesh cutting using virtual tools’, *Journal of Visualization and Computer Animation*. Wiley Online Library, 13(1), pp. 21–42. doi: 10.1002/vis.275.

Bui, H. P., Tomar, S. and Bordas, S. P. A. (2019) ‘Corotational cut finite element method for real-time surgical simulation: Application to needle insertion simulation’, *Computer Methods in Applied Mechanics and Engineering*, 345, pp. 183–211. doi: 10.1016/j.cma.2018.10.023.

Cakir, Omer, Yazici, R. and Cakir, Oguzhan (2009) ‘Real-time cutting simulation based on stiffness-warped FEM’, *2009 24th International Symposium on Computer and Information Sciences, ISCIS 2009*, pp. 721–724. doi: 10.1109/ISCIS.2009.5291912.

Chang, J.-D., Kim, I.-S., Bhardwaj, A. M. and Badami, R. N. (2017) ‘The Evolution of Computer-Assisted Total Hip Arthroplasty and Relevant Applications’, *Hip & Pelvis*, 29(1), p. 1. doi: 10.5371/hp.2017.29.1.1.

Chang, M. B., Ullman, T., Torralba, A. and Tenenbaum, J. B. (2016) ‘A compositional object-based approach to learning physical dynamics’, *5th International Conference on Learning*

Representations, ICLR 2017 - Conference Track Proceedings.

Chen, D., Levin, D. I. W., Sueda, S. and Matusik, W. (2015a) 'Data-driven finite elements for geometry and material design', *ACM Transactions on Graphics*, 34(4). doi: 10.1145/2766889.

Chen, D., Levin, D. I. W., Sueda, S. and Matusik, W. (2015b) 'Data-driven finite elements for geometry and material design', *ACM Transactions on Graphics*, 34(4). doi: 10.1145/2766889.

Childs, T. H. C. and Arola, D. (2011) 'Machining of cortical bone: Simulations of chip formation mechanics using metal machining models', *Machining Science and Technology*. Taylor & Francis, 15(2), pp. 206–230. doi: 10.1080/10910344.2011.580699.

Chitalu, F. M., Miao, Q., Subr, K. and Komura, T. (2020) 'Displacement-Correlated XFEM for Simulating Brittle Fracture', *Computer Graphics Forum*, 39(2), pp. 569–583. doi: 10.1111/cgf.13953.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. and Bengio, Y. (2014) 'Learning phrase representations using RNN encoder-decoder for statistical machine translation', *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pp. 1724–1734. doi: 10.3115/v1/d14-1179.

Clapper, T. C. (2010) 'Beyond Knowles: What Those Conducting Simulation Need to Know About Adult Learning Theory', *Clinical Simulation in Nursing*, 6(1), pp. e7-- e14. doi: 10.1016/j.ecns.2009.07.003.

Cranmer, M. D., Xu, R., Battaglia, P. and Ho, S. (2019) 'Learning Symbolic Physics with Graph Networks', (NeurIPS). Available at: <http://arxiv.org/abs/1909.05862>.

Cui, N., Kharel, P. and Gruev, V. (2017) 'Augmented reality with Microsoft HoloLens holograms for near infrared fluorescence based image guided surgery', in *Molecular-Guided Surgery: Molecules, Devices, and Applications III*, p. 100490I. doi: 10.1117/12.2251625.

Damewood, A. M. (2016) 'Current Trends in Higher Education Technology: Simulation', *TechTrends*. Springer, 60(3), pp. 268–271. doi: 10.1007/s11528-016-0048-1.

Ericsson, K. A. (2012) 'The Influence of Experience and Deliberate Practice on the Development of Superior Expert Performance', *The Cambridge Handbook of Expertise and Expert Performance*, 38, pp. 683–704. doi: 10.1017/cbo9780511816796.038.

European Parliament (2003) 'Directive 2003/88/EC of the European Parliament and of the

Council of 4 November 2003 concerning certain aspects of the organisation of working time', *Official Journal of the European Union*, 299, pp. 9--19. Available at: <https://www.bma.org.uk/pay-and-contracts/working-hours/european-working-time-directive-ewtd/doctors-and-the-european-working-time-directive>.

Fotouhi, J. (2018) 'Plan in 2-D, execute in 3-D: an augmented reality solution for cup placement in total hip arthroplasty', *Journal of Medical Imaging*. International Society for Optics and Photonics, 5(02), p. 1. doi: 10.1117/1.jmi.5.2.021205.

FundamentalVR (2018) 'Fundamental Surgery - Total hip replacement'. Available at: <https://www.fundamentalsurgery.com/total-hip-arthroplasty-posterior/#hipvideo>.

Gardner, A. K., Scott, D. J., Pedowitz, R. A., Sweet, R. M., Feins, R. H., Deutsch, E. S. and Sachdeva, A. K. (2015) 'Best practices across surgical specialties relating to simulation-based training', *Surgery (United States)*. Elsevier, 158(5), pp. 1395–1402. doi: 10.1016/j.surg.2015.03.041.

Gillies, D. F. and Bourmpos, M. (2003) 'Modelling soft tissue deformations using principal component analysis and haptic devices', *Proceedings of the LASTED International Conference on Biomechanics*, pp. 297–302.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. and Dahl, G. E. (2017) 'Neural message passing for quantum chemistry', *34th International Conference on Machine Learning, ICML 2017*, 3, pp. 2053–2070.

Google (2018) 'ARCore Overview | ARCore | Google Developers', *Developers.Google.Com*. doi: 10.1021/jo020560b.

Gori, M., Monfardini, G. and Scarselli, F. (2005) 'A new model for learning in graph domains', *Proceedings of the International Joint Conference on Neural Networks*. IEEE, 2, pp. 729–734. doi: 10.1109/IJCNN.2005.1555942.

Guo, J. and Lv, M. (2010) 'Explicit finite element simulation of oblique cutting process', in *Key Engineering Materials*, pp. 297–300. doi: 10.4028/www.scientific.net/KEM.431-432.297.

Hage, I. S. and Hamade, R. F. (2013) 'Micro-FEM orthogonal cutting model for bone using microscope images enhanced via artificial intelligence', *Procedia CIRP*. Elsevier, 8, pp. 385–390. doi: 10.1016/j.procir.2013.06.121.

Hamrick, J. B., Allen, K. R., Bapst, V., Zhu, T., McKee, K. R., Tenenbaum, J. B. and Battaglia,

- P. W. (2018) 'Relational inductive bias for physical construction in humans and machines', 1. Available at: <http://arxiv.org/abs/1806.01203>.
- Healthcare, C. A. E. (2017) 'VimedixAR'. Available at: <https://caehealthcare.com/hololens>.
- Hetaimish, B. M. (2016) 'Sawbones laboratory in orthopedic surgical training', *Saudi Medical Journal*. Saudi Medical Journal, 37(4), pp. 348–353. doi: 10.15537/smj.2016.4.13575.
- Hololens, M. (2017) 'Hololens Gesture'. Available at: <https://developer.microsoft.com/en-us/windows/mixed-reality/gestures>.
- Hooke, R. (1978) 'De Potentia Restitutiva, or of Spring', *Explaining the Power of Springing Bodies*. Printed for John Martyn Printer to the Royal Society, London.
- Hou, J. ., Petrinic, N., Ruiz, C. and Hallett, S. R. . (2000) 'Prediction of impact damage in composite plate', *Composites Science and Technology*. Elsevier, 60(2), pp. 273–281. doi: 10.1051/mattech/2019006.
- Huang, H. M., Rauch, U. and Liaw, S. S. (2010) 'Investigating learners' attitudes toward virtual reality learning environments: Based on a constructivist approach', *Computers and Education*. Elsevier, 55(3), pp. 1171–1182. doi: 10.1016/j.compedu.2010.05.014.
- Jang, D., Kim, K. and Jung, J. (2000) 'Voxel-based virtual multi-axis machining', *The International Journal of Advanced Manufacturing Technology*. Citeseer, 16(10), pp. 709–713.
- Katsikadelis, J. T. (2016) *The Boundary Element Method for Engineers and Scientists: Theory and Applications: Second Edition*, *The Boundary Element Method for Engineers and Scientists: Theory and Applications: Second Edition*. doi: 10.1016/C2014-0-01297-X.
- Khamsi, M. A. and Kirk, W. A. (2001) *An Introduction to Metric Spaces and Fixed Point Theory*, *An Introduction to Metric Spaces and Fixed Point Theory*. John Wiley & Sons. doi: 10.1002/9781118033074.
- Kim, M., Pons-Moll, G., Pujades, S., Bang, S., Kim, J., Black, M. J. and Lee, S. H. (2017) 'Data-driven physics for human soft tissue animation', *ACM Transactions on Graphics*, 36(4). doi: 10.1145/3072959.3073685.
- Knowles, M. S. (1984) 'The adult learner: a neglected species Houston: Gulf Pub', *Kobasa, SC (1982). The Hardy Personality: Toward a Social Psychology of Stress and Health*. In GS Sanders & J. Suls (Eds.), *Social Psychology of Health and Illness*, pp. 3–32.

-
- Kohn, L. T., Corrigan, J. M., Donaldson, M. S. and others (2000) 'To err is human: building a safer health system. A report of the Committee on Quality of Health Care in America, Institute of Medicine'. Washington, DC: National Academy Press.
- Kolb, D. (1984) 'Experiential education: Experience as the source of learning and development', *Englewood Cliffs, NJ*.
- Koschier, D., Bender, J. and Thuerey, N. (2017) 'Robust extended finite elements for complex cutting of deformables', *ACM Transactions on Graphics*, 36(4). doi: 10.1145/3072959.3073666.
- Kotsis, S. V and Chung, K. C. (2013) 'Application of the "see one, do one, teach one" concept in surgical training.', *Plastic and reconstructive surgery*, 131(5), pp. 1194–1201. doi: 10.1097/PRS.0b013e318287a0b3.
- Krathwohl, D. R. and Anderson, L. W. (2009) *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.
- Lainema, T. (2008) 'Perspective making: Constructivism as a meaning-making structure for simulation gaming', *Simulation & Gaming*. SAGE Publications Sage CA: Los Angeles, CA, 40(1), pp. 48–67. doi: 10.1177/1046878107308074.
- Li, L., Yu, F., Shi, D., Shi, J., Tian, Z., Yang, J., Wang, X. and Jiang, Q. (2017) 'Application of virtual reality technology in clinical medicine', *Am J Transl Res*, 9(9), pp. 3867–3880.
- Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R., Brockschmidt, M. and Tarlow, D. (2015) 'Gated graph sequence neural networks', *arXiv preprint arXiv:1511.05493*, (1), pp. 1–20.
- Li, Y., Wu, J., Tedrake, R., Tenenbaum, J. B. and Torralba, A. (2018) 'Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids', *arXiv preprint arXiv:1810.01566*, pp. 1–15.
- Liao, Z. and Axinte, D. A. (2016) 'On monitoring chip formation, penetration depth and cutting malfunctions in bone micro-drilling via acoustic emission', *Journal of Materials Processing Technology*. Elsevier, 229, pp. 82–93.
- Liu, X. Bin, Wang, X. Bin, Li, C. N. and Deng, S. P. (2013) 'Finite element simulation of the orthogonal cutting based on abaqus', in *Advanced Materials Research*, pp. 1410–1413.
- Llanos, I., Villar, J. A., Urresti, I. and Arrazola, P. J. (2009) 'Finite element modeling of oblique machining using an arbitrary Lagrangian--Eulerian formulation', *Machining science and*

technology. Taylor & Francis, 13(3), pp. 385–406.

Lughmani, W. A., Bouazza-Marouf, K. and Ashcroft, I. (2013) ‘Finite element modeling and experimentation of bone drilling forces’, in *Journal of Physics: Conference Series*, p. 12034.

Mabrey, J. D., Reinig, K. D. and Cannon, W. D. (2010) ‘Virtual reality in orthopaedics: is it a reality?’, *Clinical Orthopaedics and Related Research* \textregistered. Springer, 468(10), pp. 2586–2591. doi: 10.1007/s11999-010-1426-1.

Magic Leap Inc. (2018) *Magic Leap One Creator Edition - Device, Florida*.

Main, R. P., Lynch, M. E. and Van Der Meulen, M. C. H. (2014) ‘Load-induced changes in bone stiffness and cancellous and cortical bone mass following tibial compression diminish with age in female mice’, *Journal of Experimental Biology*, 217(10), pp. 1775–1783. doi: 10.1242/jeb.085522.

Malukhin, K. and Ehmann, K. (2018) ‘Mathematical Modeling and Virtual Reality Simulation of Surgical Tool Interactions With Soft Tissue: A Review and Prospective’, *Journal of Engineering and Science in Medical Diagnostics and Therapy*, 1(2). doi: 10.1115/1.4039417.

Marco, M., Rodríguez-Millán, M., Santiuste, C., Giner, E. and Miguélez, M. H. (2015) ‘A review on recent advances in numerical modelling of bone cutting’, *Journal of the mechanical behavior of biomedical materials*. Elsevier, 44, pp. 179–201.

Mason, W. T. M. and Strike, P. W. (2003) ‘See one, do one, teach one—is this still how it works? A comparison of the medical and nursing professions in the teaching of practical procedures’, *Medical teacher*. Taylor & Francis, 25(6), pp. 664–666.

Mathieu, V., Michel, A., Lachaniette, C.-H. F., Poignard, A., Hernigou, P., Allain, J. and Haiat, G. (2013) ‘Variation of the impact duration during the in vitro insertion of acetabular cup implants’, *Medical engineering & physics*. Elsevier, 35(11), pp. 1558–1563. doi: 10.1016/j.medengphy.2013.04.005.

McAlinden, M. G. and Dougherty, P. J. (2014) ‘Orthopaedic education in the United Kingdom’, *Clinical Orthopaedics and Related Research*, 472(6), pp. 1697–1702. doi: 10.1007/s11999-014-3573-2.

McNutt, A. J. (2017) *Model Preparation and User Interface Aspects for Microsoft HoloLens Medical Tutorial Applications*. Wright State University.

Merriam, S. B., Caffarella, R. S. and Baumgartner, L. M. (2012) *Learning in adulthood: A comprehensive guide*. John Wiley & Sons.

Microsoft-OpenSource (2018) 'MixedRealityToolkit-Unity'. Available at: <https://github.com/Microsoft/MixedRealityToolkit-Unity>.

Microsoft (2017) 'Hololens Health & Safety'. Available at: <https://www.microsoft.com/en-us/hololens/legal/health-and-safety-information>.

Mitchell, L. (2017) *Safer surgery: analysing behaviour in the operating theatre*. CRC Press.

Mitsuishi, M., Warisawa, S. and Sugita, N. (2004) 'Determination of the machining characteristics of a biomaterial using a machine tool designed for total knee arthroplasty', *CIRP Annals-Manufacturing Technology*. Elsevier, 53(1), pp. 107–112.

Morgan, M., Aydin, A., Salih, A., Robati, S. and Ahmed, K. (2017) 'Current Status of Simulation-based Training Tools in Orthopedic Surgery: A Systematic Review', *Journal of Surgical Education*. Elsevier. doi: 10.1016/j.jsurg.2017.01.005.

Mosbech, T. H., Ersbøll, B. K. and Christensen, L. B. (2009) 'An efficient data-driven tissue deformation model', in *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops 2009*, pp. 1771–1777. doi: 10.1109/ICCVW.2009.5457497.

National_Joint_Register (2018) 'Procedure details by type of provider for hip procedures for 2017'. Available at: http://www.njrreports.org.uk/hips-all-procedures-activity/H01v2NJR?reportid=C6F582E2-140D-4D22-8C4E-2C354EDB1B41&defaults=DC__Reporting_Period__Date_Range=%22MAX%22,JYS__Filter__Calendar_Year__From__To=%22max-max%22,H__Filter__Joint=%22Hip%22.

National_Joint_Registry (2018) 'National Joint Registry | 15th Annual Report 3'. Available at: http://www.njrreports.org.uk/hips-all-procedures-activity/H03v2NJR?reportid=FFCEA144-54BC-486D-81A6-C6A58DDCA079&defaults=DC__Reporting_Period__Date_Range=%22MAX%22,JYS__Filter__Calendar_Year__From__To=%22min-max%22,R__Filter__Country=%22All%22,H__Filter__.

NDI Medical, N. (2017) 'Medical Polaris Spectra and Vicra'. Available at: <https://www.ndigital.com/medical/products/polaris-family/>.

Newman, J. M. (2014) 'Robotic-assisted total hip arthroplasty: a critical appraisal'.

- Newman, S., Gulati, V., Bahadori, S., Wainwright, T. and Middleton, R. (2018) 'Content and Face Validity Assessment of the Sim-K Haptic-Feedback Enhanced Total Knee Replacement Virtual Reality Simulator.', *The Internet Journal of Orthopedic Surgery*, 27(1). doi: 10.5580/IJOS.53658.
- Ng, E.-G. and Aspinwall, D. K. (2002) 'Modelling of hard part machining', *Journal of materials processing technology*. Elsevier, 127(2), pp. 222–229.
- Niu, Q. and Leu, M. C. (2007) 'Modeling and rendering for a virtual bone surgery system.', *Studies in health technology and informatics*, 125, pp. 352–354.
- Orthopaedics, T. (2016) 'TOTAL HIP REPLACEMENT - UNCEMENTED', [*onlinevideo*]. Available at: <https://www.youtube.com/watch?v=sMBx3C8hhJg&t=217s>.
- OsSimTech (2018) 'Sim ortho'. Available at: <https://ossimtech.com/en-us/Simulators>.
- OssoVR (2016) 'OssoVR Virtual Surgery. Real Results.' Available at: <http://ossovr.com/>.
- Pan, J. J., Chang, J., Yang, X., Liang, H., Zhang, J. J., Qureshi, T., Howell, R. and Hickish, T. (2015) 'Virtual reality training and assessment in laparoscopic rectum surgery', *International Journal of Medical Robotics and Computer Assisted Surgery*, 11(2), pp. 194–209. doi: 10.1002/rcs.1582.
- Piispänen, V. (1948) 'Theory of formation of metal chips', *Journal of Applied Physics*. AIP, 19(10), pp. 876–881.
- Polyanin, A. and Chernoutsan, A. (2010) *Elements of Strength of Materials, A Concise Handbook of Mathematics, Physics, and Engineering Sciences*. doi: 10.1201/b10276-29.
- Quesada, C., González, D., Alfaro, I., Cueto, E. and Chinesta, F. (2016) 'Computational vademecums for real-time simulation of surgical cutting in haptic environments', *International Journal for Numerical Methods in Engineering*. doi: 10.1002/nme.5252.
- Rainger, P. (2019) 'Validation of Fundamental Surgery, a Haptic VR education platform, in conjunction with the British Orthopaedic Training Association'. Available at: https://www.fundamentalsurgery.com/wp-content/uploads/2019/03/FundamentalVR_ValidationFeb2019_BOTA.pdf.
- Redmond, J. M., Gupta, A., Hammarstedt, J. E., Petrakos, A., Stake, C. E. and Domb, B. G. (2016) 'Accuracy of Component Placement in Robotic-Assisted Total Hip Arthroplasty',

Orthopedics. SLACK Incorporated, 39(3), pp. 193–199.

Rohrich, R. (2006) ‘“See One, Do One, Teach One”: An Old Adage with a New Twist’, *Plastic and Reconstructive Surgery*, 118(1), pp. 257–258. doi: 10.1097/01.prs.0000233177.97881.85.

Rojas, J., Bautista, M., Bonilla, G., Amado, O., Huerfano, E., Monsalvo, D., Llinás, A. and Navas, J. (2018) ‘A retrospective study on the relationship between altered native acetabular angle and vertical implant malpositioning’, *International Orthopaedics*, 42(4), pp. 769–775. doi: 10.1007/s00264-017-3584-1.

Salas, E., Rosen, M. A., Held, J. D. and Weissmuller, J. J. (2009) ‘Performance measurement in simulation-based training: A review and best practices’, *Simulation & Gaming*. SAGE Publications Sage CA: Los Angeles, CA, 40(3), pp. 328–376. doi: 10.1177/1046878108326734.

Sanchez-Gonzalez, A., Heess, N., Springenberg, J. T., Merel, J., Riedmiller, M., Hadsell, R. and Battaglia, P. (2018) ‘Graph networks as learnable physics engines for inference and control’, *35th International Conference on Machine Learning, ICML 2018*, 10, pp. 7097–7117.

Sanchez-Gonzalez, A., Bapst, V., Cranmer, K. and Battaglia, P. (2019) ‘Hamiltonian Graph Networks with ODE Integrators’. Available at: <http://arxiv.org/abs/1909.12790>.

Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J. and Battaglia, P. W. (2020) ‘Learning to Simulate Complex Physics with Graph Networks’. Available at: <http://arxiv.org/abs/2002.09405>.

Santiuste, C., Rodríguez-Millán, M., Giner, E. and Miguélez, H. (2014) ‘The influence of anisotropy in numerical modeling of orthogonal cutting of cortical bone’, *Composite Structures*. Elsevier, 116, pp. 423–431.

Santoro, A., Raposo, D., Barrett, D. G. T., Malinowski, M., Pascanu, R., Battaglia, P. and Lillicrap, T. (2017) ‘A simple neural network module for relational reasoning’, *Advances in Neural Information Processing Systems*, 2017-Decem, pp. 4968–4977.

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M. and Monfardini, G. (2009a) ‘Computational capabilities of graph neural networks’, *IEEE Transactions on Neural Networks*, 20(1), pp. 81–102. doi: 10.1109/TNN.2008.2005141.

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M. and Monfardini, G. (2009b) ‘The graph

neural network model', *IEEE Transactions on Neural Networks*, 20(1), pp. 61–80. doi: 10.1109/TNN.2008.2005605.

Schermann, T., Marsolek, J., Schmidt, C. and Fleischer, J. (2006) 'Aspects of the simulation of a cutting process with ABAQUS/Explicit including the interaction between the cutting process and the dynamic behavior of the machine tool'.

Schijven, M. P. and Jakimowicz, J. J. (2005) 'Validation of virtual reality simulators: Key to the successful integration of a novel teaching technology into minimal access surgery', *Minimally Invasive Therapy & Allied Technologies*. Taylor & Francis, 14(4–5), pp. 244–246.

Schnurr, C., Michael, J. W. P., Eysel, P. and König, D. P. (2009) 'Imageless navigation of hip resurfacing arthroplasty increases the implant accuracy.', *International orthopaedics*, 33(2), pp. 365–372.

Schoch, N., Suwelack, S., Speidel, S., Heuveline, V. and No, P. (2013) 'Simulation of Surgical Cutting of Soft Tissue using the X-FEM', *Preprint Series of the Engineering Mathematics and Computing Lab*, 0(04). doi: 10.11588/emclpp.2013.04.11825.

Seiler, M., Spillmann, J. and Harders, M. (2012) 'Enriching coarse interactive elastic objects with high-resolution data-driven deformations', in *Computer Animation 2012 - ACM SIGGRAPH / Eurographics Symposium Proceedings, SCA 2012*, pp. 9–17.

Seymour, N. E., Gallagher, A. G., Roman, S. A., O'Brien, M. K., Bansal, V. K., Andersen, D. K. and Satava, R. M. (2002) 'Virtual reality training improves operating room performance: results of a randomized, double-blinded study', *Annals of surgery*. Lippincott, Williams, and Wilkins, 236(4), p. 458. doi: 10.1097/00000658-200210000-00008.

Si, W., Liao, X., Wang, Q. and Heng, P. A. (2017) 'Personalized heterogeneous deformable model for fast volumetric registration', *BioMedical Engineering Online*, 16(1). doi: 10.1186/s12938-017-0321-3.

Sifakis, E., Der, K. G. and Fedkiw, R. (2007) 'Arbitrary cutting of deformable tetrahedralized objects', in *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pp. 73–80.

Silvennoinen, M., Helfenstein, S., Ruoranen, M. and Saariluoma, P. (2012) 'Learning basic surgical skills through simulator training', *Instructional Science*. Springer, 40(5), pp. 769–783.

- Silvennoinen, M. (2017) 'Displaying Augmented Reality content on HoloLens environment'. Metropolia Ammattikorkeakoulu.
- Simonetti Ibanez, A. and Paredes Figueras, J. (2013) *Vuforia v1. 5 SDK: Analysis and evaluation of capabilities*. Universitat Politècnica de Catalunya.
- Smith, M. (2016) *ABAQUS/Standard User's Manual, Version 2016*. United States: Dassault Systèmes Simulia Corp.
- Snijders, T., Gaalen, S. M. and Gast, A. (2017) 'Precision and accuracy of imageless navigation versus freehand implantation of total hip arthroplasty: A systematic review and meta-analysis', *The International Journal of Medical Robotics and Computer Assisted Surgery*. Wiley Online Library.
- Song, J. H. and Belytschko, T. (2009) 'Cracking node method for dynamic fracture with finite elements', *International Journal for Numerical Methods in Engineering*. doi: 10.1002/nme.2415.
- Spencer, A. J. M. (2004) *Continuum mechanics*. Courier Corporation.
- Stirling, E. R., Lewis, T. L. and Ferran, N. A. (2014) 'Surgical skills simulation in trauma and orthopaedic training', *Journal of Orthopaedic Surgery and Research*, 9(1), p. 126. doi: 10.1186/s13018-014-0126-z.
- Swemac (2017) 'TraumaVision - medical orthopedic simulator'. Available at: <http://www.swemac.com/simulators/traumavision>.
- Syed, S., Mirza, A. H. and Ali, A. (2009) 'A brief comparison of orthopaedic training in English-speaking countries', *The Annals of The Royal College of Surgeons of England*. The Royal College of Surgeons of England, 91(3), pp. 226–231.
- Tai, B. L., Zhang, L., Wang, A., Sullivan, S. and Shih, A. J. (2013) 'Neurosurgical bone grinding temperature monitoring', *Procedia Cirp*. Elsevier, 5, pp. 226–230.
- Tawbe, B. and Cretu, A. M. (2017) 'Acquisition and neural network prediction of 3D deformable object shape using a kinect and a force-torque sensor', *Sensors (Switzerland)*, 17(5). doi: 10.3390/s17051083.
- Thorey, F., Klages, P., Lerch, M., Flörkemeier, T., Windhagen, H. and von Lewinski, G. (2009) 'Cup positioning in primary total hip arthroplasty using an imageless navigation device: is there a learning curve?', *Orthopedics*, 32(10 Suppl), pp. 14–17. doi: 10.3928/01477447-20090915-52.

- Tilly, B. (2016) 'Mako Robotic-Arm Assisted System: A Clinical and Economic Analysis for Health Plans and Providers.'
- Tonutti, M., Gras, G. and Yang, G. Z. (2017a) 'A machine learning approach for real-time modelling of tissue deformation in image-guided neurosurgery', *Artificial Intelligence in Medicine*. Elsevier B.V., 80, pp. 39–47. doi: 10.1016/j.artmed.2017.07.004.
- Tonutti, M., Gras, G. and Yang, G. Z. (2017b) 'A machine learning approach for real-time modelling of tissue deformation in image-guided neurosurgery', *Artificial Intelligence in Medicine*. Elsevier, 80, pp. 39–47. doi: 10.1016/j.artmed.2017.07.004.
- TouchSurgery (2017) 'TouchSurgery'. Available at: <https://www.touchsurgery.com/>.
- Tu, Y.-K., Chen, L.-W., Ciou, J.-S., Hsiao, C.-K., Chen, Y.-C. and others (2013) 'Finite element simulations of bone temperature rise during bone drilling based on a bone analog', *Journal of Medical and Biological Engineering*. 中華民國生物醫學工程學會, 33(3), pp. 269–274.
- Turini, G., Ganovelli, F. and Montani, C. (2006) 'Simulating drilling on tetrahedral meshes', in *Proceedings of Eurographics Conference - Short Papers*, pp. 127–131. Available at: <http://vcg.isti.cnr.it/Publications/2006/TGM06/drillingtetra.pdf>.
- Unity Technology (2018) 'Unity 3D', *Unity Technology*.
- Vashishth, D., Tanner, K. E. and Bonfield, W. (2000) 'Contribution, development and morphology of microcracking in cortical bone during crack propagation', *Journal of Biomechanics*. Elsevier, 33(9), pp. 1169–1174.
- Vaughan, N., Dubey, V. N., Wainwright, T. W. and Middleton, R. G. (2016) 'A review of virtual reality based training simulators for orthopaedic surgery', *Medical engineering & physics*. Elsevier, 38(2), pp. 59–71. doi: 10.1016/j.medengphy.2015.11.021.
- VirtaMed AG (2014) *Virtamed Medical education with virtual reality simulators*. Available at: <https://www.virtamed.com/en/> (Accessed: 26 July 2020).
- Vozenilek, J., Huff, J. S., Reznick, M. and Gordon, J. A. (2004) 'See one, do one, teach one: advanced technology in medical education', *Academic Emergency Medicine*. Wiley Online Library, 11(11), pp. 1149–1154. doi: 10.1197/j.aem.2004.08.003.
- VR, F. (2017) 'Multiperson Medical Training'. Available at: <https://www.fundamentalvr.com/services/hololens-studio/>.

- Vuforia, P. T. C. (2018) 'Vuforia Developer Library'. Available at: <https://library.vuforia.com/articles/Training/Image-Target-Guide.html>.
- Wang, D., Roy, A. and Silberschmidt, V. V (2016) 'Hybrid cutting of bio-tissues'. \copyright The Authors. Published by Elsevier Ltd.
- Wang, E. E. (2011) 'Simulation and adult learning.', *Disease a month: DM*, 57(11), pp. 664–678. doi: 10.1016/j.disamonth.2011.08.017.
- Wolf, B. R., Lu, X., Li, Y., Callaghan, J. J. and Cram, P. (2012) 'Adverse outcomes in hip arthroplasty: Long-term trends', *Journal of Bone and Joint Surgery - Series A*, 94(14), p. e103(1). doi: 10.2106/JBJS.K.00011.
- Wolford, M. L., Palso, K. and Bercovitz, A. (2015) 'Hospitalization for total hip replacement among inpatients aged 45 and over: United States, 2000-2010', *NCHS data brief*, pp. 1–8. Available at: <https://www.cdc.gov/nchs/products/databriefs/db186.htm>.
- Wu, J., Yu, G., Wang, D., Zhang, Y. and Wang, C. C. L. (2009) 'Voxel-based interactive haptic simulation of dental drilling', in *Proceedings of the ASME Design Engineering Technical Conference*, pp. 39–48. doi: 10.1115/DETC2009-86681.
- Wu, J., Westermann, R. and Dick, C. (2015) 'A Survey of Physically Based Simulation of Cuts in Deformable Bodies', *Computer Graphics Forum*, 34(6), pp. 161–187. doi: 10.1111/cgf.12528.
- Xie, H., Song, J., Zhong, Y. and Gu, C. (2020) 'Kalman Filter Finite Element Method for Real-Time Soft Tissue Modeling', *IEEE Access*. IEEE, 8, pp. 53471–53483. doi: 10.1109/ACCESS.2020.2981400.
- Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E. and others (2018) 'Deep reinforcement learning with relational inductive biases'.
- Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C. and Sun, M. (2018) 'Graph Neural Networks: A Review of Methods and Applications', pp. 1–22. Available at: <http://arxiv.org/abs/1812.08434>.
- Zienkiewicz, OC and Parekh, C. (1970) 'Transient field problems: two-dimensional and three-dimensional analysis by isoparametric finite elements', *International Journal for Numerical Methods in Engineering*, 2, pp. 61–71. doi: 10.1002/nme.1620020107.

Zienkiewicz, O. C. and Taylor, R. L. (2005) *The finite element method for solid and structural mechanics*. Elsevier.

Zigmont, J. J., Kappus, L. J. and Sudikoff, S. N. (2011) ‘Theoretical foundations of learning through simulation.’, *Seminars in perinatology*, 35(2), pp. 47–51. doi: 10.1053/j.semperi.2011.01.002.

Zinar, R. and Schmalzried, T. P. (2015) ‘Why hip implants fail: Patient, surgeon, or device?’, in *Seminars in Arthroplasty*, pp. 118–120.