



# AUTOMATING MEDICAL IMAGE ANALYSIS OF SR $\mu$ - CT AND EVALUATION OF REGENERATED BONE HEALTH

MASTER THESIS IN COMPUTATIONAL PHYSICS


Written by *Anna Xiao Tan*

19.06.2022

Supervised by

James Emil Avery

UNIVERSITY OF COPENHAGEN





UNIVERSITY OF  
COPENHAGEN

NAME OF INSTITUTE: The Niels Bohr Institute

NAME OF DEPARTMENT: eScience

AUTHOR(S): Anna Xiao Tan

EMAIL: jmz556@alumni.ku.dk

TITLE AND SUBTITLE: Automating Medical Image Analysis of SR $\mu$ -CT and Evaluation of Regenerated Bone Health  
-

SUPERVISOR(S): James Emil Avery

HANDED IN: 19.06.2022

DEFENDED: 01.07.2022

NAME \_\_\_\_\_

SIGNATURE \_\_\_\_\_

DATE \_\_\_\_\_

## Abstract

In this thesis we aim to to fully automate the segmentation and analysis of materials in high-resolution synchrotron radiation micro-computed tomography tomograms consisting of (I) a titanium implant, (II) bone mineral, (III) blood vessels, (IV) osteocytes and (V) air/resin to develop a fully automatic, quantitative model of healthy bone and a fully automatic regenerated bone health evaluation.

The regenerated bone is separated from the old bone by obtaining the boundary between two different screw threads of a titanium implant implanted in the bone sample. The titanium implant is segmented through threshold-based segmentation restricted to an area of interest bounded by a bounding box. The screw thread of the implant is evaluated through wave analysis, through which the boundary between different screw threads is identified.

A fully automatic quantification of bone mineral volume is given through a bone mineral density test. The bone mineral is segmented through a Gaussian mixture modelling of the histogram of voxel intensities in the tomogram. The total bone volume is found through morphological closing of the segmented bone mineral, and the air/resin is segmented as material outside the bone volume. We find that the regenerated bone consists of 57.7% bone mineral, in comparison with 70.3% bone mineral in the old bone.

We further propose a new metric for a fully automatic evaluation of regenerated bone health through assessment of the distribution of blood vessels and osteocytes, respectively. The blood vessels are segmented through connected component analysis, as the blood vessel network comprises the largest connected component of soft tissue. We design a recursive multi-resolution analysis method for segmentation of the blood vessels to alleviate RAM requirements. As the osteocytes have a characteristic volume, they are extracted through a volumetric segmentation method. A threshold for maximum distance to an osteocyte or blood vessel, obtained through a distance transform, will determine the quantity of healthy regenerated bone. Using this metric, 99.85% of the regenerated bone is classified as healthy, in comparison with 99.69% of the old bone.

The algorithm is tested on one 3D tomogram of size  $(3279 \times 3480 \times 3480)$ . The evaluation method should be run on additional bone sample tomograms from the same experiment to substantiate the developed quantitative model of healthy bone, as we expect an accurate model to quantify the old bone in any of these tomogram as healthy.

## Acknowledgements

To Assoc. Prof. James Emil Avery, my supervisor through this master thesis, I offer my deepest gratitude for his time, knowledge and invaluable guidance. *So long, and thanks for all the fish.*

To Z. Tan, where help is given even at the last of minutes, I offer my gratefulness for his advice, support, and lifelong inspiration.

To M. Algren – not even the world would suffice.

# Contents

1	<i>Introduction</i>	1
1.1	<i>Motivation</i>	2
1.2	<i>Problem Formulation</i>	3
1.3	<i>Solution Overview</i>	3
2	<i>Data and Image Formation</i>	6
2.1	<i>From Histology to Synchrotron Radiation CT</i>	6
2.2	<i>SR<math>\mu</math>CT Image Data</i>	8
2.3	<i>X-Ray Physics of Image Formation</i>	9
2.4	<i>Geometry of Image Formation</i>	14
2.5	<i>The European Synchrotron Radiation Facility and High-Brilliance X-rays</i>	16
2.6	<i>Image Reconstruction</i>	18
2.7	<i>Visualisation of 3D Data</i>	21
3	<i>Multi-resolution Representation</i>	23
3.1	<i>Multi-resolution Analysis</i>	23
3.2	<i>Working with Multi-Resolution Image Representation</i>	25
4	<i>Implant Segmentation</i>	28
4.1	<i>Voxel-based Thresholding</i>	28
4.2	<i>Bounding Box</i>	30
4.3	<i>Principal Axes Analysis</i>	30
4.4	<i>Procrustes analysis</i>	32
4.5	<i>Morphology</i>	33
4.6	<i>Implementation and Results</i>	34

5	<i>Implant Orientation and Separation into Parts</i>	39
5.1	<i>Image Projection Profile</i>	39
5.2	<i>Wave Characteristics</i>	40
5.3	<i>Peak Finding</i>	40
5.4	<i>Implementation and Results</i>	41
6	<i>Bone Segmentation</i>	47
6.1	<i>Gaussian Mixture Model</i>	47
6.2	<i>K-Means Clustering</i>	49
6.3	<i>Implementation and Results</i>	50
7	<i>Intermediate Bone Health Assessment</i>	56
7.1	<i>Implementation and Results</i>	56
8	<i>Blood Vessel Segmentation</i>	62
8.1	<i>Connected Component Analysis</i>	63
8.2	<i>More Multi-resolution Analysis</i>	64
8.3	<i>Implementation and Results</i>	66
9	<i>Osteocyte Segmentation</i>	70
9.1	<i>Implementation and Results</i>	71
10	<i>Improved Bone Health Assessment</i>	76
10.1	<i>Distance Transform</i>	77
10.2	<i>Implementation and Results</i>	77
11	<i>Discussion and Future Work</i>	83
12	<i>Conclusion</i>	86

*Appendix A: A Hierarchy of Image Processing Tasks*

*Appendix B: Iterative Image Reconstruction Methods*

*Appendix C: 2D Visualisation of a 3D Surface*

*Appendix D: Implant Segmentation Solution Flow Chart*

*Appendix E: Illustrating Limitations of Threshold-based Segmentation*

*Appendix F: The Bounding Box and Instance Segmentation*

*Appendix G: Illustrating the Projection Profile*

*Appendix H: Minimisation*


*Appendix I: Toy Example for Multi-resolution Blood Vessel Segmentation*

*Appendix J: Recursive Algorithms and The Fibonacci Sequence*

*Appendix K: Osteocyte Segmentation Through Morphology*

# 1 Introduction

Bone is a rigid organ that protects other organs, stores minerals, creates blood cells, regulates hormones and metabolic processes, enables mobility, provides structure and support and is a living tissue undergoing constant change[1]. With so many critical functions, it is essential to maintain good bone health. Bone is subject to a number of impairments, including fractures and cancers, but also bone diseases that arise gradually with advanced age. These include osteoporosis, where the bone mineral density is decreased, as shown in Fig. 1.1, leading to greater chances of fractures. With a world population that is getting older and older, age-progressive bone deficiency diseases such as osteoporosis in particular have become a rising issue, affecting more than an estimated 200 million people worldwide [2]. Thus, the need for understanding bone microbiology and the causes, mechanisms, and effective treatment of bone disease, such as discovering effective methods for stimulating bone regeneration and the evaluation of these methods, is more present than ever[3].



**Figure 1.1:** Biopsies from a patient with normal bone health (left) and a patient with osteoporosis (right). The healthy bone is denser and strongly interconnected, while the osteoporotic bone is sparse and demonstrates a weaker structure and thus weaker bone strength [4].

In [5] and [6], five experiments with the goal of studying bone regeneration methods were conducted on seven goats, with four experiments testing methods for bone regeneration and one experiment for control. Synchrotron radiation micro-computed tomography (SR $\mu$ -CT) at the at the European Synchrotron Radiation Facility (ESRF) was used to produce 35 ultra-precise computed tomography (CT) tomograms<sup>1</sup> of the 35 bone samples, each of size  $(3279 \times 3480 \times 3480)$ .

The goal of this thesis is (I) to automate the segmentation of materials in the tomograms to (II) develop an automatic, quantitative

<sup>1</sup> A 3D volume image produced through tomography.

model of healthy bone through analysis of the old bone, and (III) to automate the evaluation of the quality of the bone regeneration through analysis of the regenerated bone and quantitative comparison of this with the old bone.

<sup>2</sup> "OsteoMorph: Understanding Bone Micro-physics through Computational Morphology of SR $\mu$ -CT" (Villum Experiment 00041017).

<sup>3</sup> <http://www.maxibone.eu>

<sup>4</sup> A pair of bones that form the part of the face surrounding the nose and hold the tooth roots.

<sup>5</sup> Lacking of teeth.

This thesis is part of the research project OsteoMorph<sup>2</sup>, which uses a combination of image analysis methods and computational physics to derive mathematical models of the microbiology and microphysics of bone based on SR $\mu$ -CT images from [5] and [6]. This work, in collaboration with surgeons at Region Syd Hospitals and University of Southern Denmark, will also provide the computational analysis component for the European MAXIBONE project<sup>3</sup>, a programme funded by the European Commission for developing personalised maxillary bone<sup>4</sup> regeneration due to the concern that tens of millions of European citizens are edentulous<sup>5</sup>, but have insufficient bone for placement of dental implants. MAXIBONE involves a trial of 150 human patients across 12 countries, whose bone biopsies will likewise be scanned at the ESRF. In this thesis, the automatic segmentation of materials, quantification of healthy bone and evaluation of bone regeneration will be developed on data obtained in [5] and [6] on biopsies from goats, to later be applied for bone regeneration in human patients through projects such as the MAXIBONE.

This chapter presents the motivation, problem formulation, and solution overview of this thesis.

## 1.1 Motivation

The evident and prevalent evaluation metric for assessing the quality of regenerated bone is the bone mineral density test, a test for the volume of bone mineral within the total bone volume. This evaluation is simple, fast, and can be accomplished with segmentation of the bone mineral alone. However, there is more to bone health than the quantity of bone mineral, as healthy bone contains other tissue types, including blood vessels, which supply nutrients to the bone, and osteocytes<sup>6</sup>, which govern the remodelling of bone and direct the transport of the nutrients throughout the bone [3]. Bone that is comprised of only bone mineral is thus not healthy, living tissue, which must contain blood vessels and osteocytes distributed throughout. It is thus essential to identify and assess these materials in evaluations relating to the bone.

<sup>6</sup> Ellipsoidal bone cells located inside the bone mineral[3].

<sup>7</sup> "Volume elements", the 3D equivalent of the 2D *pixels*, "picture elements".

We must thus identify the material class to which each voxel<sup>7</sup> in the tomogram belongs in order to automatically assess the quality of regenerated bone based on the distribution of materials throughout the bone volume.

Previously, a semi-automatic labelling method involving manual sparse labelling in the form of seeding has been used, where selected voxels or sections are roughly labelled manually as seeds, from which a region-growing algorithm determines the remaining labels [5]. However, this method involves time-consuming and cumbersome manual intervention. In this thesis, we aim to transfer the



entire labour burden to the computer, fully automating the labelling of materials, the quantitative modelling of healthy bone, and the quantitative evaluation of the health of regenerated bone.

## 1.2 Problem Formulation

In the manual process of labelling, the human brain automatically translates the visual information about the voxel in question from the eyes into the cortical information that this voxel belongs to a certain material class. Transferring this task to the computer involves the formulation of mathematical rules from which the computer is able to automatically separate the voxels into their material classes. This image processing classification process is known as segmentation, where each voxel in a 3D image is classified as one of  $N$  classes. For these tomograms, the segmentation is based on the material type, where the material types present are (1) titanium implant, (2) bone mineral, (3) blood vessel, (4) osteocyte, and (5) air/resin. The air and resin will be treated together, as they are not part of the bone sample, but rather materials introduced in the imaging phase.

This segmentation will allow us to make conclusions based on analysis of the segmented image, principally the quantitative modelling of healthy bone and the evaluation of the health of the regenerated bone. For this, we will propose a new automatic bone health evaluation method that is similarly simple and fast to the bone mineral density test, but considers the presence of blood vessels and osteocytes distributed throughout the bone mineral in order to provide a more comprehensive and accurate measure for the quality of regenerated bone. We will use this automatic bone health evaluation method to create an automatic quantitative model for "healthy bone" using the old bone, against which the regenerated bone can be evaluated. This proposed bone health evaluation method, along with the automatic bone health quantification and the material segmentation methods, will be outlined in Section 1.3.

Based upon the introduction above, we formulate the following problem statement.

*Using image processing methods, is it possible to create a fully automatic quantitative model of healthy bone and evaluation of regenerated bone health through automatic segmentation of bone samples imaged through high-resolution synchrotron radiation micro-computed tomography into material classes?*

In the section below, we will provide details on the proposed solution for answering the problem statement.

## 1.3 Solution Overview

This section presents an overview of the solution composed to answer the problem statement. In order to obtain an automatic quantification of healthy bone and evaluation of regenerated bone health, we fully automate the segmentation of voxels into one of five mate-

rial classes present in the tomograms, (I) titanium implant, (II) bone mineral, (III) blood, (IV) osteocyte, or (V) air/resin, using a flight of image processing methods, as will be outlined below. We then propose and implement a method for obtaining an automatic quantitative model of healthy bone and for an automatic evaluation of the regenerated bone against this model based on the segmented tomogram.

In Appendix A, we include a suggested hierarchy of image processing steps that outline a typical path from image formation to image comprehension<sup>8</sup>.

<sup>8</sup>Image comprehension follows the recognition of the contents of an image and involves making conclusions based on the segmented image, e.g., evaluation of bone health.

The tomograms produced are of ultra-high resolution, which is necessary for distinguishing small structures, but unnecessary for larger structures, where the high resolution only carries a superfluous RAM and computational burden. We will use a multi-resolution representation of the tomogram data to analyse large structures at a lower computational cost, while smaller structures can be analysed the original resolution level, as detailed in Chapter 3.

For segmentation of the titanium implant in the tomogram, we employ voxel-based thresholding to roughly segment the implant, as the implant is comprised of voxels of much higher intensity than the rest of the tomogram. However, this approach is sensitive to noise, as surrounding voxels may appear as bright as the implant and thus be misclassified as implant voxels. Thus, we draw an oriented bounding box which limits the area of interest, and the implant can then be segmented through voxel-based thresholding restricted to this area. The methods, implementation and results of the implant segmentation will be given in Chapter 4.

The titanium implant is comprised of a head, a section with a micro screw thread and a section with a macro screw thread. In Chapter 5, we separate the implant into its regions of interest, as these distinguish the old bone from the regenerated bone. A wave characteristic analysis will be used to do this. The outline of the implant is obtained through an image projection, and the orientation and individual components of the implant can be determined through the difference in wavelength and amplitude of the two screw threads.

For segmentation of the bone mineral, a Gaussian mixture model will be used to model the histogram with probability density distributions, where K-means clustering is used to provide a minimisation algorithm with starting guesses for distribution parameters. The method, implementation and the results of the bone mineral segmentation will be given in Chapter 6.

In Chapter 7, we give an intermediate bone regeneration assessment using the segmented bone mineral and a bone mineral density test. We obtain the volume of bone mineral by computing the area under the curve of the projection profile of the segmented bone mineral, and take the ratio of this against the total volume of bone for the bone mineral density test. This gives us a preliminary evaluation of bone regeneration through assessment of bone mineral quantity. For evaluation of the quality of the regenerated bone, we must obtain a

measure of the distribution soft tissue throughout the bone volume.

The soft tissue voxels, hereunder blood and osteocyte voxels, are near identical in voxel intensity value and cannot be separated through voxel-based thresholding methods. Instead, the blood vessel network segmentation will involve connected component analysis, as the blood vessel network comprises the largest connected component of soft tissue. The blood vessel network includes very small blood vessels, which require the full resolution of the tomograms in order to be discerned, and it is not feasible to perform connected component analysis on the full high-resolution tomograms, as this would require enormous amounts of RAM. To alleviate the RAM requirements, we design a recursive multi-resolution analysis, as detailed in Chapter 8.

The osteocytes have characteristic volume, by which they will be extracted using a volume based segmentation method in Chapter 9.

Finally, we require a metric for the automatic evaluation of the health of the regenerated bone. For this, a distance transform will be used to obtain the distance from each bone voxel to the closest osteocyte voxel, as well as the distance from each bone voxel to the closest blood voxel, and a threshold for maximum distance to osteocyte and blood, respectively, will determine the quantity of healthy regenerated bone. The method, implementation and results of the improved bone health quantification will be given in Chapter 10.

The fully automatic segmentation and evaluation methods will be tested on one 3D tomogram of size  $(3279 \times 3480 \times 3480)$ , but the resulting code is developed to run on any of the 35 tomograms produced in [5] and [6] with no additional intervention.

All the source code in the project was coded in Python and is available at: <https://github.com/annaxtan/master-thesis.git>.

# Part I: Background

## 2 *Data and Image Formation*

The data analysed in this thesis are 3D high-resolution SR $\mu$ CT tomograms. Computed tomography (CT) is an imaging technique that involves the reconstruction of a three-dimensional (3D) volume image from two-dimensional (2D) projections from several angles around the sample. In this chapter, a brief history motivating the use of SR $\mu$ CT for imaging is given. The tomogram data analysed in this thesis is described. An introduction to the formation of tomograms, hereunder X-ray physics, X-ray image formation geometry, the ESRF, and the reconstruction of 2D projections into a 3D volume image will be given as well. Finally, we will briefly relay the methods used in this thesis for visualisation of the 3D data. This chapter serves to provide the background for the thesis.


### 2.1 *From Histology to Synchrotron Radiation CT*

Traditionally, bones have been examined through histology, in which a sample, in the form of a 2D slice of bone, is sawed off from a biopsy using a diamond blade, and the sample is then examined under a microscope. To record the presence and distribution of tissue types in the sample, a transparent film with a grid may then be placed over the sample, on which a technician can assign a label to each grid element, which describes the material class to which it belongs. This time-consuming collection and evaluation method carries a substantial manual labour burden. It also invariably requires the removal as well as the destruction of the bone in order to obtain the slices for histological study.

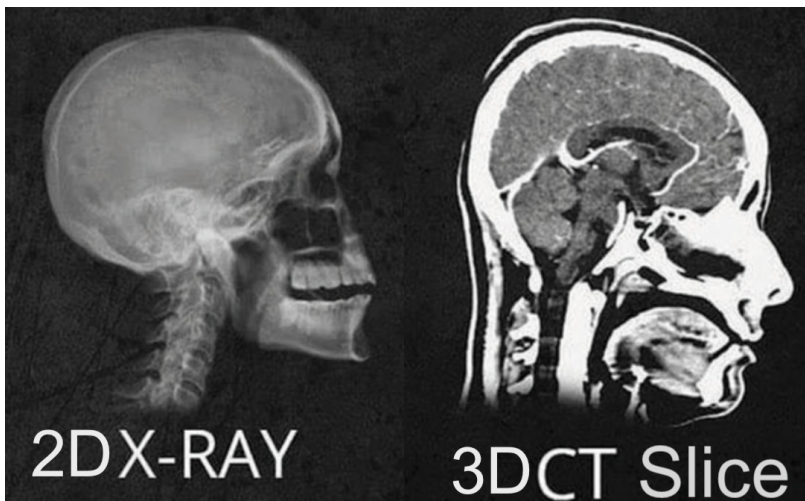
An advancement in imaging technology that circumvented the removal and destruction of an internal structure to enable the study hereof was the 2D X-ray image, where a 2D projection image through a volume is obtained from a single angle using electromagnetic radiation. In Fig. 2.2, one such 2D X-ray projection image is shown. The production of X-rays will be detailed further in Section 2.3. These 2D X-ray images work quite well for examining larger, denser structures, useful for applications such as the identification of broken bones, but are entirely unable to capture finer structures or soft tissue, as they are only able to represent the shadow image from one projection through the whole object volume.

In 3D CT, a 3D volume image is reconstructed through 2D projection images obtained from many angles. Once the 3D volume image

**Figure 2.1:** Two X-ray images obtained through projection from two angles [7].



is reconstructed, a slice through the 3D volume image can be visualised as a 2D slice image. CT image reconstruction will be detailed further in Section 2.6. In Fig. 2.2, we compare a single projection X-ray image with a slice of a 3D CT volume image. We see that CT imaging is able to capture soft tissue and far finer details than the 2D X-ray image.



**Figure 2.2:** 2D X-ray image (left) and 2D slice of 3D CT image (right) of a skull, adapted from [8]. Note that the two are not images of the same patient and are included for a basic illustration of the difference between the two imaging modalities only.

The primary difference lies in that a 2D X-ray image is the result of passing one X-ray beam from one angle through the sample, obtaining a single projection, while 3D CT involves the reconstruction of several X-ray projections from several angles into a 3D volume image. In other words, through CT, we are able to simulate histology without the removal or destruction of the sample. The sample is represented by the reconstructed 3D volume image, and the virtual sample can be examined in slices<sup>1</sup>.

Another significant advantage of the 3D CT is that the samples themselves are three-dimensional and are far better represented by

<sup>1</sup> This is the etymology of the word "tomography", which is derived from the Ancient Greek *tomos*, which means *to section* or *to slice*.

a 3D volume image. In our case, this property of the CT is critical, as there are structures of interest, such as the blood vessel network, where its three-dimensional connectivity is essential for the segmentation and subsequent analysis.

Some of the structures of interest, such as the osteocytes, are very small and require a high resolution in order to be discerned. The resolution of a CT image is given by the voxel length, which is the physical distance represented by the voxel and thus a limit for the smallest representable structure. The resolution can also inversely be cited through the number of voxels used to represent the image. The CT images produced by hospital X-ray equipment have a voxel length of 0.5–0.625mm [9]. A volume of, e.g.,  $(10 \times 10 \times 10) = 1000\text{mm}^3$  is thus represented by  $\sim \frac{1000}{0.5^3} = (20 \times 20 \times 20) = 8000$  voxels.

This resolution is fine for imaging larger structures, such as bone fractures or tumors, but some of the critical structures we are interested in require a much higher resolution to discern. The tomograms analysed in this thesis were produced through synchrotron radiation CT, which is able to achieve far finer resolution than hospital X-ray setups, and have a voxel length of  $1.875\mu\text{m}$ . The example volume from above of  $1000\text{mm}^3$  would in this case be represented by  $\frac{1000}{0.00185^3} \approx 1.579 \cdot 10^{11}$  voxels. Synchrotron radiation will be detailed further in Section 2.5.

## 2.2 SR $\mu$ CT Image Data

The SR $\mu$ CT data analysed in this thesis was obtained in [5] and [6], where five experiments investigating different methods for bone regeneration were conducted.

Bone is regenerated through the work of three bone cell types, osteoclasts, osteoblasts, and osteocytes. Osteoclasts resorb old bone and osteoblasts follow in their path, replacing the old bone with regenerated bone. Osteoblasts may mature into osteocytes, the most abundant bone cell type<sup>2</sup>, which are responsible for monitoring and coordinating bone health maintenance. The osteocytes respond to bone tissue in need of remodelling by communicating this need to the osteoblasts and osteoclasts. The osteocytes are also responsible for directing the transport of nutrients throughout the bone tissue through the lacunae-canalliculi network, the space where the osteocytes reside, and are thus integral to maintaining bone health [3]. When a bone defect arises that the bone cells fail to regenerate spontaneously, surgical intervention in the form of a bone graft may be necessary[10]. Bone grafting is a technique for stimulating the natural bone regeneration process. The bone grafting material is placed in the bone defect, stimulating and accelerating the natural bone regeneration processes to the defect.


For the data analysed in this thesis, the five bone grafts were made of (I) synthetic, resorbable  $\beta$ -tricalcium phosphate ( $\beta$ -TCP), (II) synthetic, partially resorbable  $\beta$ -TCP, (III) autologous bone chips<sup>3</sup> covered by a titanium membrane, (IV) autologous bone chips with no

<sup>2</sup> Comprising 90 – 95% of bone cells, with an estimated total population of 42 billion within a fully developed human skeleton[3].

<sup>3</sup> The bone that is removed to create the defect is ground up into small chips and placed in the defect of the same goat.

membrane and (V) an empty defect covered by a titanium membrane for control [6].

For each experiment, a defect was artificially created in the mandible<sup>4</sup> of a goat. as shown in Fig. 2.3, a titanium dental implant measuring 8mm in length and 3.5mm in diameter was installed for the purpose of distinguishing the original bone from the regenerated bone. The lower 5.5mm of the implant is comprised of macro threads, and the upper 2.5mm is comprised of of micro threads. The lower 5.5mm of the implant was implanted in the mandible and the upper 2.5mm stood free, so that the regenerated bone would grow around the 2.5mm micro threads, distinguishing it from the original bone.



**Figure 2.3:** Goat mandible defect and titanium dental implant (left) and illustration of implant (right) [5]. The lower half of the implant, with a larger screw thread, is implanted in the mandible, with the upper half of the implant, with a smaller screw thread, standing free to be encompassed by regenerated bone, distinguishable from the original bone through the disparity in screw thread.

Following 20 weeks of regeneration, the samples were removed and separated in two parts, where one was used for histological studies and one was fixated in 10% formaldehyde, dehydrated in graded alcohol solutions, embedded in an acrylic cylinder and taken to the ID19 beamline<sup>5</sup> at The European Synchrotron Radiation Facility (ESRF) for scanning. The results of this scanning are the 35 SR $\mu$ -CT tomograms.

In the following sections, a summary of the X-ray physics involved in CT image formation, the optical geometry of CT image formation, the ESRF, and the reconstruction of a number of 2D X-ray projections into a 3D volume image will be given.

### 2.3 X-Ray Physics of Image Formation

In the following sub-sections, an introduction to the formation of X-rays through electron radiation, and the various interactions of X-ray radiation with matter which determine the attenuated intensity values recorded by the sensor, will be given. Note that the X-ray generation process described in the sub-section below pertains specifically to X-rays generated in hospitals<sup>6</sup>, as X-rays generated at the ESRF will be given treatment in its own section in Section 2.5, but the physics involved in both processes, as well as the subsequent interactions of the X-rays with matter in the sample, are identical.

<sup>5</sup> ESRF experimental station.


<sup>6</sup> In a synchrotron, X-rays are produced by decelerating the electrons via magnetic fields, rather than through interaction with atoms in an anode. The method of production and the energy of the incoming electrons are the major differences between hospital setup X-rays and synchrotron radiation X-rays, as will be detailed in Section 2.5.



### 2.3.1 Electron Radiation

X-rays produced by hospital equipment are generated through the interaction of high-energy electrons from an electron beam with atoms in a rotating anode, as shown in Fig. 2.4.

**Figure 2.4:** An illustration of the generation of X-rays in a typical hospital setup. Electrons interact with target atoms in the anode, resulting in the emission of X-ray photons. [11].




When an electron interacts with the electron cloud surrounding a nucleus, the negative charged field of the electron will interact with the field of the electron cloud. The positive charge of the nucleus will attract the electron via the Coulomb force<sup>7</sup> and the electron will be deflected from its path.

Due to energy conservation, this change in momentum causes the electron to lose energy, and this energy is emitted in the form of X-ray photons, as shown in Fig. 2.5, in an effect also known as Bremsstrahlung<sup>8</sup>.

<sup>7</sup> The force of attraction or repulsion between particles as a result of their electric charge.

<sup>8</sup> From German, *bremesen* "to brake" and *Strahlung* "radiation".

**Figure 2.5:** An illustration of the interaction which causes the emission of X-ray photons, adapted from [11]. The incident electron interacts with the electron cloud around the nucleus of the target atoms in the anode. The electron is deflected and decelerated, causing the electron to lose energy emitted in the form of X-ray photons.



As photons have no mass or charge, they interact differently with tissue than charged particles, as they do not form electromagnetic

fields that can interact with tissue. Instead, the intensity of the incoming X-ray beam is reduced exponentially as it passes through the tissue. This is due to the fact that there is a probability for the occurrence of interactions with material in the tissue, given by Lambert-Beer's law,

$$I(d) = I_0 e^{-\mu d} \quad [12], \quad (2.1)$$

where  $I(d)$  is the intensity at depth  $d$  of the tissue,  $I_0$  is the intensity at the starting point, and  $\mu$  is a linear attenuation coefficient that is dependent on the energy of the incoming photons and the material being irradiated. The linear attenuation coefficient depends on the ways the X-rays have interacted with the matter in the sample and is what is solved for during the image reconstruction.


In Eq. 2.1, we assume that the sample is homogeneous, i.e. that it consists of a single material. When the sample is inhomogeneous, as is the case for the the 3D tomogram analysed in this thesis, the Lambert-Beer law is instead given by

$$I(d) = I_0 e^{-\int_0^d \mu(x) dx}, \quad (2.2)$$

The projections recorded by the sensor represents the X-ray attenuation<sup>9</sup>. The X-ray beam interacts with the matter in the sample, which causes the X-ray beam to lose energy with penetration depth. The projections are recordings of the X-rays that exit the object. The higher the measured value is, the lower the X-ray attenuation, as more of the X-rays have passed through the matter without being absorbed or deflected. The probability of interactions increases with density of the tissue, and higher density tissues will cause higher degrees of attenuation, causing denser materials in the reconstructed tomogram to be represented brighter. The five basic radiographic <sup>10</sup> densities are air, fat, water (soft tissue), bone, and metal, where the denser the material, the brighter it will be represented on the tomogram, as shown in Fig. 2.6.

<sup>9</sup> The reduction of the intensity of an X-ray beam as it traverses matter.

<sup>10</sup> i.e. pertaining to medical imaging through X-ray radiation.



**Figure 2.6:** The five basic radiographic densities [13]. Higher density materials will cause higher degrees of attenuation, leading to high density materials appearing brighter.

X-rays interact with matter in three primary ways, photo-absorption, Compton scattering, and pair formation, through which the radiation can deposit energy in the tissue, resulting in reductions in the X-ray beam intensity. These interactions, and the probability of their occurrence, will be summarised below.


### 2.3.2 Photo-absorption

In photo-absorption, the photons in the X-ray beam interact with the atom in a process where an electron is released from the atom and the photon is completely absorbed, as shown in Fig. 2.7. Part of the energy of the photon is used to break the binding energy of the electron, which receives the residual energy as kinetic energy. As energy is required to break the binding energy of the electron, photo-absorption occurs only if the energy of the X-ray beam exceeds the binding energy of the electron.

For an electron from an outer shell to fill the place of the released electron, a photon with an energy corresponding to the difference in the energy levels of the shells between which the electron jumps is required. This is called characteristic radiation. If the characteristic radiation is absorbed internally in the atom, so that it is excited<sup>11</sup>, an Auger electron will be emitted. The electron or electrons emitted are able to interact in further interaction throughout the matter in the same way as electron radiation.

<sup>11</sup> A quantum mechanical process where an electron temporarily occupies an energy state greater than its ground state.

**Figure 2.7:** An illustration of an electron interacting with matter through photo-absorption, adapted from [12]. A photon is absorbed and the energy from the photon, and, if the energy of the photon exceeds the binding energy of the electron, an electron is released. Part of the energy of the photon is used to break the binding energy of the electron, which receives the residual energy as kinetic energy.



The probability of the photoelectric effect occurring is directly proportional with the number of electrons in the absorbing material, and therefore the density of the material and its atomic number, and inversely proportional with the photon energy [14],

$$P_{PE} \propto \frac{(Z_{abs})^5}{E_\gamma^{7/2}}, \quad (2.3)$$

where  $Z_{abs}$  is the atomic number of the absorbing atom and  $E_\gamma$  is the energy of the X-ray beam.


### 2.3.3 Compton Scattering

Compton scattering is a process in which a photon from the incident X-ray beam collides with a free electron<sup>12</sup> in the tissue, where some of the photon's energy is used to break the bond between the electron

<sup>12</sup> An electron where the binding energy of the electron is much lower than the energy of the X-ray.

and the nucleus and transfer kinetic energy to the electron, as in photo-absorption. From here, the released electron will be emitted at an angle  $\theta$ , as shown in Fig. 2.8. It will interact in the same way as electron radiation.

In addition, a new photon will be emitted, whose path will be deflected by an angle  $\phi$  from the path of the incoming photon, which will continue at a lower energy than that of the incoming photon. The scattered photons may make the final image appear less sharp, as they are able to enter into new interactions with the matter, but as they have a lower energy, they have a larger probability of being absorbed en route to the sensor, as the probability of interaction is inversely proportional with the photon energy.



**Figure 2.8:** An illustration of an electron interacting with matter through Compton scattering, adapted from [12]. An incident photon collides with an electron whose binding energy is much lower than the energy of the photon, and an electron and a new photon with a lower energy than that of the incident photon is emitted.

The probability of the occurrence of Compton scattering is dependent on the energy of the incoming photon, the electron density of the absorbing tissue and in particular on its atomic number, but since Compton scattering involves free electrons it is independent of the total binding energy of the electrons. The probability is given by

$$P_{\text{Compton scattering}} \propto \frac{Z_{abs}}{E_{\gamma}} \quad [14], \quad (2.4)$$


where  $Z_{abs}$  is the atomic number of the absorbing material and  $E_{\gamma}$  is the X-ray energy.

#### 2.3.4 Pair formation

Pair formation is the last significant way in which X-ray radiation interacts with tissue. During pair formation, X-ray radiation is transformed via interaction with the irradiated material into an electron and a positron, i.e. the electron's antiparticle, which behaves as a positively charged electron, as shown in Fig. 2.9.

The probability of interaction via pair formation is described by

$$P_{\text{Pair production}} \propto (\log E_{\gamma})(Z_{abs})^2 \quad [14], \quad (2.5)$$



**Figure 2.9:** An illustration of an electron interacting with matter through pair production, adapted from [12]. The incident photon is absorbed and an electron and a positron are emitted. Pair-formation only occurs at very high X-ray energies (1.02MeV), and is thus not observed at the X-ray energy at which the bone samples analysed in this thesis were imaged.

where  $Z_{abs}$  is the atomic number of the absorbing material and  $E_{\gamma}$  is the energy of the incident photon.

Note that as an electron has a rest mass of  $0.511MeV/c^2$  of energy, pair formation occurs only for X-ray radiation with an energy of at least 1.02 MeV, since the formation of the electron-positron pair requires energy at least equivalent to two rest masses. The bone samples analysed in this thesis were imaged with a X-ray energy of 67.4KeV [6], an energy at which pair formation is not observed.

From Eqs. 2.3, 2.4 and 2.5, we see that the probability of interaction for all interactions increase with increasing the electron density, or simply density, of the material. This results in the radiographic densities shown in Fig. 2.6, where the denser the material, the higher X-ray attenuation.

## 2.4 Geometry of Image Formation


In Fig. 2.10, we show an example of X-ray imaging geometry. The X-rays are sent through the sample, which is placed a distance of  $R_1$  from the X-ray source. The X-rays that exit the sample are recorded by a detector placed a distance of  $R_2$  from the sample. These distances determine the amount the image is magnified with respect to the sample. Some structures of interest in the bone samples are micron sized, and in order to view them, the magnified image must be greater than 100 times the size of the sample[15]. The magnification is given by the ratio between the distance between the detector to the X-ray source and the distance between the sample and the X-ray source,

$$\text{Magnification} = \frac{R_1 + R_2}{R_1} \quad (2.6)$$


Fig. 2.10 showed an ideal point source, i.e. a source of negligible size. The size of the X-ray source is known as the focal spot, and is determined by the area of contact with the anode. The size of the focal spot is a decisive factor in the sharpness<sup>13</sup> of images produced by hospital X-ray setups. The X-rays travel in a straight line from their source. When these focal spot from which the rays are emitted is of a finite size, i.e. not a point source, the X-rays originate from

<sup>13</sup>Sharpness of an image generally refers to the contrast in images, especially near the borders of objects. Sharpness and resolution, whilst sometimes used interchangeably in layman terms, are two distinct features, both of which are sought optimised for the best quality images.

**Figure 2.10:** An example of X-ray imaging geometry [16]. The X-rays are sent through the sample, which is placed a distance of  $R_1$  from the X-ray source, and the X-rays that exit the sample are recorded by a detector placed a distance of  $R_2$  from the sample. These distances determine the amount the image is magnified with respect to the sample.



different points and therefore their projections of a particular feature is not received by exactly the same sensor in the detector, causing a blurring effect around the edges of objects [17]. Ideally, the focal spot is as small as possible, as this will produce the sharpest image, as shown in Fig. 2.11. However, a smaller focal spot runs the risk of melting the anode, as the incoming electrons are concentrated on a small area. A balance must thus be achieved in that the focal spot must be small enough to create a sharp image, but large enough to relieve the thermal stress on the anode.



**Figure 2.11:** The difference in sharpness with a focal spot of finite size (left) and an idealised focal spot of negligible size (right)[18]. The smaller the focal spot, the sharper the image.

## 2.5 The European Synchrotron Radiation Facility and High-Brilliance X-rays

<sup>14</sup> <https://www.esrf.fr/>

The European Synchrotron Radiation Facility<sup>14</sup> (ESRF), founded in 1988 and located in Grenoble, France, is the X-ray facility that produces the highest brilliance X-rays in the world, and the facility at which the bone samples investigated in this thesis were scanned.

<sup>15</sup> X-ray brightness.

The amount of detail that can be discerned from an X-ray image is highly governed by the brilliance<sup>15</sup>,  $B$ , a measure of the total flux of photons in a given six-dimensional phase space per unit bandwidth, of the X-ray source [19],

$$B = I \cdot (4\pi^2 \sigma_x \sigma_y \sigma_{x'} \sigma_{y'} \frac{dw}{w})^{-1}, \quad (2.7)$$

<sup>16</sup> For a set of  $n$  values  $\{x_1, x_1, \dots, x_n\}$ , the root mean square value is given by  $\sigma_x = \sqrt{\frac{1}{n}(x_1^2 + x_2^2 + \dots + x_n^2)}$ .

where  $I$  is the intensity of the beam, given by number of photons per second,  $\sigma_x$  and  $\sigma_y$  are the root mean square values<sup>16</sup> of the width of the beam for the two axes perpendicular to the direction of travel of the beam,  $\sigma_{x'}$  and  $\sigma_{y'}$  are the root mean square values of the beam solid angles<sup>17</sup> for the two axes perpendicular to the direction of travel of the beam, and  $\frac{dw}{w}$  is the bandwidth<sup>18</sup>. Thus, the brilliance is maximised when the beam intensity  $I$  is maximised and the  $(4\pi^2 \sigma_x \sigma_y \sigma_{x'} \sigma_{y'} \frac{dw}{w})$  term in Eq. 2.7 is minimised.

<sup>17</sup> The fraction of source particles that enter the detector aperture[20].

<sup>18</sup> The variance in beam frequency from the central frequency.

The standard X-rays used in hospitals have a brilliance in the ballpark around  $10^6$  photons  $s^{-1}$   $mm^{-2}$   $mrad^{-2}$   $0.1\%BW^{-1}$ <sup>19</sup>. This resolution is excellent for the imaging needs of hospitals, where X-ray imaging is used for e.g. identification of bone fractures, tumors, or pneumonia in the lungs. For specific research purposes, such as the study of microscopic structures in bone, even finer resolution is required.

<sup>19</sup> We have here used the industry standard units for brilliance, in favour of the SI units of  $s^{-1}$   $m^{-2}$   $rad^{-2}$   $100\%BW^{-1}$  [21].

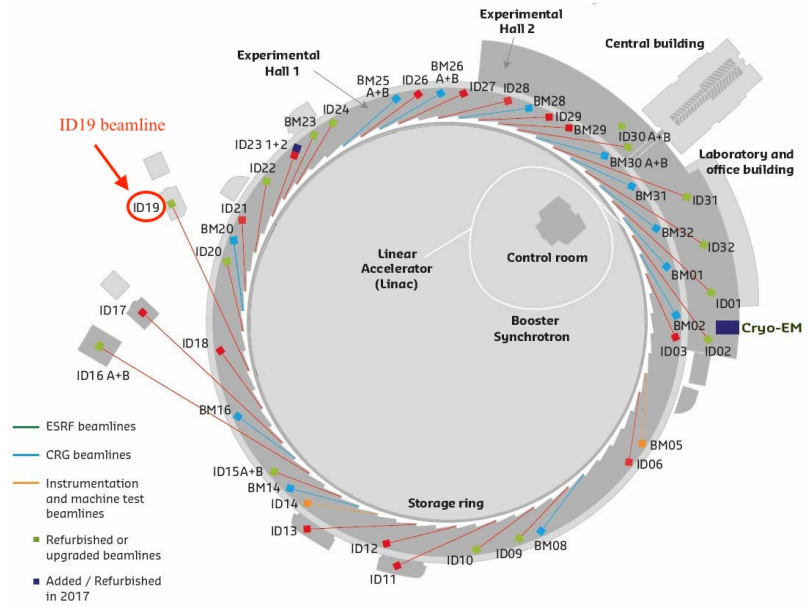
The European Synchrotron Radiation Facility answers this high resolution requirement, producing X-rays 10 trillion times brighter than the X-rays used in hospitals, capable of imaging structures in matter to the atomic level. In Fig. 2.12, we include a diagram of the ESRF.

To produce synchrotron radiation X-rays at the ESRF, the linear accelerator (Linac) emits electrons through an electron gun, which are packed closely together and accelerated by electric fields until they are travelling near the speed of light. The electrons are then sent into the booster synchrotron, a ring with a circumference of 300m, round which the electrons travel a few several thousand times, gaining energy for each lap, until they reach an energy of up to 6 billion electron-volts (6GeV)<sup>20</sup>. As mentioned in Section 2.3, the emission of X-ray photons becomes more likely as the energy of the electron beam increases, and the 6GeV electron beam of the ESRF boasts enormous energy. The resulting X-ray beam will proportionately have a very high intensity.

<sup>20</sup> Compared to 30 – 150 eV in normal hospitals.

The high-energy electrons are subsequently released into the storage ring, a ring with a circumference of 844 metres. Once in the storage ring, the electrons will be deflected from their path by magnets, causing the electrons to lose energy in the form of electromagnetic


**Figure 2.12:** Experimental hall with beamline location markers at the European Synchrotron Radiation Facility [22]. The ID19 beamline used for obtaining the tomograms for this thesis is highlighted in red.



radiation due to deceleration, including X-ray radiation, which are sent through the sample for imaging.

Due to this method of generating X-ray radiation, the synchrotron is designed with a massive advantage over hospital X-ray setups when it comes to the geometry of X-ray imaging. In hospital X-ray setups, the incident electrons produce X-rays via interactions with an anode. This places a limit on how small the focal spot can be. As the X-rays from synchrotron radiation are produced when the high-speed electrons are decelerated by magnetic fields, rather than their contact with an anode, the focal spot in synchrotron radiation can be made extremely small. The ID19 beamline used for obtaining the tomograms for this thesis can achieve a focal spot below 100nm.

In Fig. 2.13, we show the difference in resolution-capability between an X-ray scan made at a hospital and that obtained at the ESRF.



**Figure 2.13:** Difference in resolution-capability between an X-ray scan made at a hospital (left) and that obtained at the ESRF (right) [22].



While the ESRF produces CT images of extraordinary resolution, they are not immune to noise. An effect related to the physics of the X-ray beam is beam hardening, which causes the intensity of the image to vary with penetration depth. This is because the lower energy photons are preferentially attenuated, leaving only higher energy photons to continue interactions deeper in the sample. A few artefacts<sup>21</sup> may arise from beam hardening, including streaking, where streaks of dark bands are formed between denser objects, and cupping, where the center of the tomogram appears darker than the edge [24]. Voxel bleeding is an effect of concern during the image reconstruction process when imaging high density materials, such as the titanium implant, where the signal from the bright implant contaminates the signal from the surrounding voxels, causing these voxels to appear brighter than those further away from the implant [25]. Ring artifacts may arise due to miscalibrations or failures of detector elements in the X-ray sensor. The effect appears as a ring superimposed on the tomogram, as the defective detector element gives an erroneous reading at each angular position [26].

<sup>21</sup> Any feature that appears in an image which was not present in the sample [23].

## 2.6 Image Reconstruction

The CT scanner has obtained a number of 2D projections and now requires a method of reconstructing these into a 3D volume image. A way to obtain such a method is to formally define a mathematical model that simulates the way a CT scan obtains 2D scans from a 3D object, and subsequently derive its inversion formula. For an object function,  $f(x, y)$ , the Radon transform,  $\mathbf{R}\{f(x, y)\}$ , computes the projections of an object at different angles, i.e. gives the sum of the voxel intensity values at each angle through a line integral,  $\mathbf{R}\{f(x, y)\} = \int_L f(x, y) dl$ , as shown in Fig. 2.14<sup>22</sup>, just as a detector records the projections of an object at different angles as it moves around the sample.

<sup>22</sup> Note that the physical sample is continuous, hence why the Radon transform theoretically is an integral, but the reconstruction is discrete as we only have a finite number of projections. In practice, the Radon transform is thus a sum, as it is used to transform images with finite pixel representations, rather than a physical sample, into sinograms for reconstruction algorithm testing.

For better suited maths, as we are concerned with angles, we introduce a change of coordinate system

$$x \cos \theta + y \sin \theta = p \quad \text{and} \quad -x \sin \theta + y \cos \theta = q \quad (2.8)$$


where  $p$  and  $q$  are the new axes and  $\theta$  is the angular displacement of  $p$  with regards to the axis  $x$  [11].

We can now define the two-dimensional continuous Radon transform as

$$\begin{aligned} P(p, \theta) &= \mathbf{R}\{f(x, y)\} \\ &= \int_L f(x, y) dl \\ &= \int_{-\infty}^{\infty} f(x, y) \delta(p \cos \theta - q \sin \theta, p \sin \theta + q \cos \theta) dq \end{aligned} \quad (2.9)$$

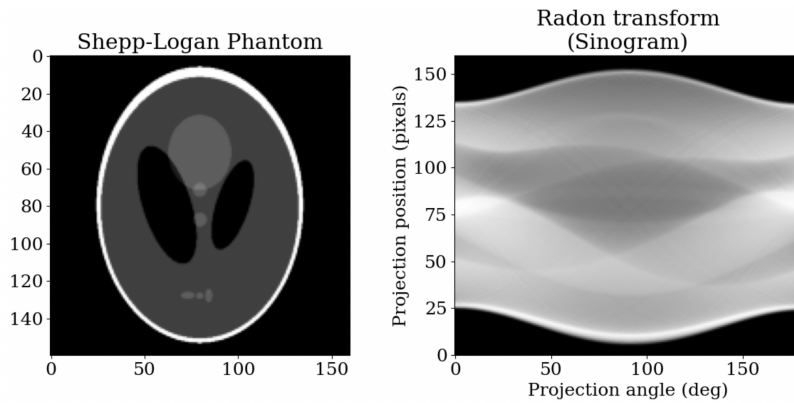
The Radon transform of an image is also called a sinogram of the image. For a 2D image, the sinogram is a 2D mapping of original image, where each column of the sinogram corresponds to a 1D projection, i.e. integral, from a particular angle. Due to symmetry, we

**Figure 2.14:** Illustration of the Radon transform [11].



only need an angular range of  $[0, 180[$  degrees or  $[\pi, 2\pi[$  radians, as  $P(p, \theta) = P(-p, \theta + \pi)$ .

In Fig. 2.15, we show the Shepp-Logan Phantom image, a test image created by [27] for the specific purpose of testing image reconstruction algorithms. We also show the Radon transformed sinogram of the Shepp-Logan Phantom image, from which a reconstruction algorithm will attempt to reconstruct the original image.

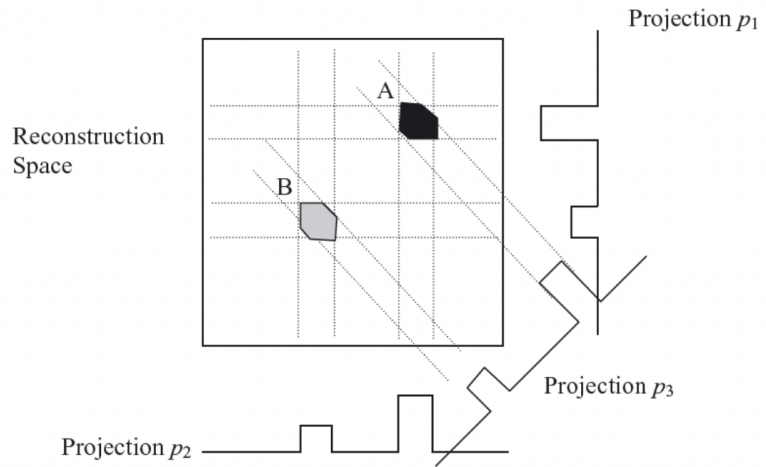


**Figure 2.15:** The original Shepp-Logan Phantom image [28] (left) and a sinogram of the Shepp-Logan Phantom image, constructed using the radon transformation from the Python library scikit-image (right).

The most common method for image reconstruction is back-projection, a Radon transform inversion formula. Through the Radon transform, a sum of pixel intensities,  $I$ , is obtained from a set of angles. Back-projecting "smears" this sum onto the corresponding column in the reconstruction space, i.e. each pixel in that column in the reconstructed image receives the value  $\frac{1}{\Sigma I}$ , as shown in Fig. 2.16. Through

this, an impression of the Radon transformed object will eventually appear.

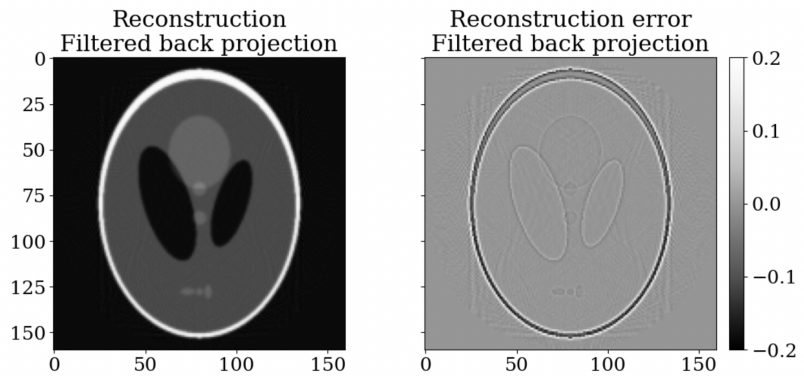
**Figure 2.16:** An illustration of back-projection [11].



<sup>23</sup> Various types of filters may be used in the image reconstruction process to suppress noise in the reconstruction.

In Fig. 2.17, on the left, the back-projection from the Python library `scikit-image` has been used with a ramp filter<sup>23</sup> to reconstruct the original Shepp-Logan Phantom image from the sinogram in Fig. 2.15. On the right, we show the reconstruction error.

**Figure 2.17:** The reconstructed Shepp-Logan Phantom image (left) and the reconstruction error (right).



We see that the algorithm performs quite well. However, in the example above, the back-projection was not applied as simply as back-projecting the projection from each angle in the sinogram once, but rather through an iterative reconstruction algorithm. In Appendix B, we include two examples which illustrate and motivate the rising application of iterative reconstruction. In summation, we can think of the image reconstruction problem as being multiple equations required to solve for multiple unknowns. The tomograms produced for this thesis are of size  $(3279 \times 3480 \times 3480)$ , yielding over a trillion unknowns for a single tomogram. Further, due to noise, arising from anything between movement of the sample to beam hardening, the generated equations are not consistent, adding to the difficulty of


resolving the linear attenuation values. To bypass the need to solve these systems of unknowns exactly, which is not always possible due to noise, iterative reconstruction can be used.

Computationally, image reconstruction is an extremely complex problem, and many algorithms have been created in the effort to solve it. The specific imaging method used for obtaining the tomograms analysed in this thesis is known as Paganin Phase-Contrast CT, which excels in imaging of materials that have naturally low inter-material contrast with other materials in the sample. It thus lends itself well for reconstructing the bone sample tomograms, as blood and osteocytes both belong to the radiographic density of soft tissue, causing them to appear similar in voxel intensity in the tomogram. Further treatment of this imaging method can be found in [29], but is beyond the scope of this thesis.

## 2.7 Visualisation of 3D Data

In this thesis, we will primarily treat the images as discrete 3D arrays,  $I(n, m, l)$ , where the indices  $(n, m, l)$  define the coordinates of a voxel in the array with respect to the image coordinate frame, a frame of reference with origin  $(n, m, l) = (0, 0, 0)$  in the top left corner of the image.


As mentioned, the data is 3D, while the medium through which it is being presented in this thesis is not. Throughout most of the thesis, we will simply visualise a 2D slice of the 3D data. In Fig. 2.18, we show an example with slices  $n = 1740$ ,  $m = 1740$  and  $l = 1740$ , respectively, of the tomogram. In the visualisation, it can be seen that the sample was cut in half before being embedded in an acrylic cylinder, surrounded by air.



**Figure 2.18:** Visualisation of slices  $n = 1740$ ,  $m = 1740$  and  $l = 1740$ , respectively, of the tomogram.

We may also use a 3D rendering software when this is appropriate. For this purpose, we will use `vedo`<sup>24</sup>, a Python library specifically developed for the visualisation of three-dimensional scientific data. `vedo` allows us to view the data in its 3D form and to examine it from every angle. This is useful for analyses where the connectivity of the data is important, for example for viewing the blood vessel network. In Fig. 2.19, we visualise the implant segmented through the analysis in Chapter 4 and the implant surface using the Python library `vedo`. Both of these representations are 3D interactive volume images that we can rotate or scale in real time.

<sup>24</sup> <https://vedo.embl.es/>



**Figure 2.19:** 3D visualisation of segmented implant (left) and segmented implant surface (right).

In Appendix C, we consider visualisation of a 3D surface through projection onto a 2D plane, which can be used for visualising the distributions of materials in contact with the implant surface.

## Part II: Analysis

## 3 *Multi-resolution Representation*

In this chapter, we will describe the use of a multi-resolution representation of the data in this thesis, and the method through which this representation is obtained. The data is of an extremely high resolution, which translates to a large RAM and computational requirement for processing. The multi-resolution representation of the data will be used to alleviate these requirements when possible.

### 3.1 *Multi-resolution Analysis*

Bone is a hierarchical organ, with important structures at a wide range of length scales, from the microscopical to the macroscopical[3]. The successful segmentation of some of the materials in the tomogram is highly dependent on the resolution of the tomogram. As stated in Chapter 2, the resolution of an image is given by the number of pixels or voxels in the array used to represent the image. Changing the resolution will have a number of consequences. When we decrease the resolution, we lose information, as the resolution determines the smallest structure that we are able to represent. Any details smaller than the voxel will not be reproducible. This effect, where structures become indistinguishable, is known as aliasing.

For the finer structures, such as the capillaries<sup>1</sup> in the blood vessel network, or the osteocytes, we in fact require the full resolution of the tomograms in order to avoid aliasing. For the larger, more coarsely defined materials, such as the implant and the bone mineral, it is unnecessary to work with data of such high resolution, as these features will be distinguishable even with fewer voxels to represent them.

Luckily, we are not required to analyse all structures at the same resolution level. We can produce a multi-resolution representation of the tomograms through down-scaling in order to analyse large structures at a lower computational cost, while smaller structures can be analysed the original resolution level. For this purpose, we introduce the image pyramid.

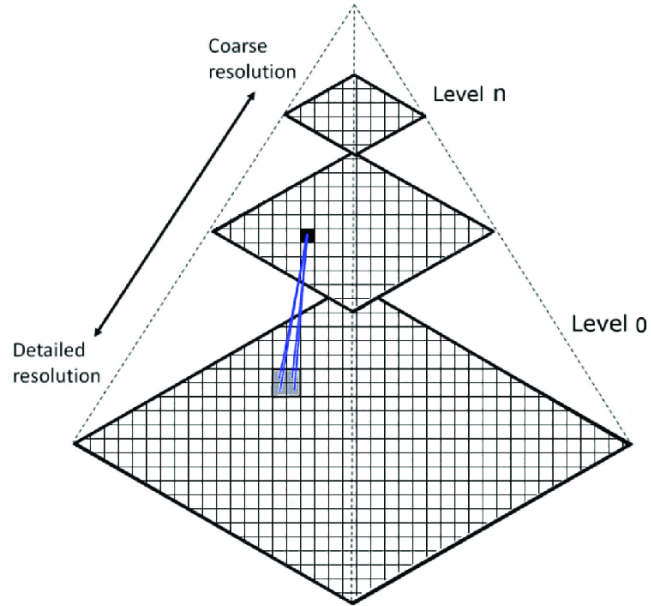
#### 3.1.1 *The Image Pyramid*

The image pyramid is a visual representation of the same imaged physical space represented in levels of decreasing resolution, as shown in Fig. 3.1. The lower resolution representations represent the same physical space, only with fewer voxels. The base image of the pyra-

<sup>1</sup> The smallest and most abundant of the blood vessels[30].

mid, on level 0, is the full resolution image of size  $(N \times N)$ . The image on the next level of the pyramid, level 1, is the 2 times down-scaled image of size  $(N/2 \times N/2)$  and so on until the apex image of size  $(1 \times 1)$  on level  $n$ . In Fig. 3.2, we show an example of a multi-resolution representation of a 1D image of shape  $(1 \times 16)$ .

**Figure 3.1:** An illustration showing an image pyramid. Level 0 represents the  $(N \times N)$  full resolution image, level 1 the  $(N/2 \times N/2)$  down-scaled image, and so forth, until the pinnacle level  $n$  of size  $(1 \times 1)$ . [31]



**Figure 3.2:** A multi-resolution representation of a 1D image of shape  $(1 \times 16)$ . The lower resolution representations represent the same physical space, only with fewer voxels.

Full Resolution, (1, 16)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
2x Down-scaled (1, 8)	1		2		3		4		5		6		7		8	
4x Down-scaled, (1, 4)	1				2				3				4			
8x Down-scaled (1, 2)	1								2							
16x Down-scaled (1, 1)	1															

### 3.1.2 Down-Scaling

As mentioned, the lower resolution representations for the multi-resolution representation are achieved through down-scaling. Down-scaling of an image involves converting an image to the same image with a lower resolution. The factor by which the lower resolution image is smaller, i.e. represented by fewer pixels or voxels, than the higher resolution image is given by  $s^d$ , where  $s$  is the down-scaling factor and  $d$  is the dimension of the image. For a 3D image of shape  $(N \times M \times L)$  that we wish to down-scale by a down-scaling factor  $s = 2$ , the down-scaled image is  $2^3 = 8$  times smaller at size  $(N/2 \times M/2 \times L/2)$ . Each image volume that was previously represented by  $(2 \times 2 \times 2)$  voxels must now be represented by just 1, as the number of voxels in each axis is reduced to half.

There are many ways to down-scale an image, dependent on the chosen method of condensing the information of the original number of voxels to the number of voxels available for the new resolu-

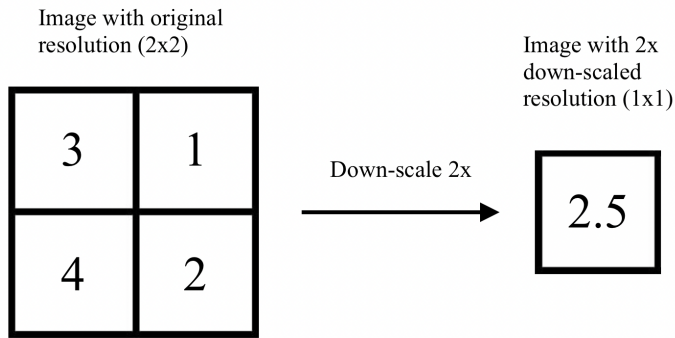


tion. The most commonly used method, and the one used for down-scaling in this thesis, is the averaging method.

In Fig. 3.3, we show an example of an image,  $B$ , with an original resolution of  $(2 \times 2)$ , where the values in the pixels represent the intensity value of a grey-scale image. If we down-scale this image to half the resolution using the averaging method, the  $2x$  down-scaled image,  $A$ , is of size  $(1 \times 1)$  and has an intensity value given by the average of the four pixels in the original resolution image,  $B$ ,

$$A_{1,1} = \frac{B_{1,1} + B_{1,2} + B_{2,1} + B_{2,2}}{2^2} = \frac{3 + 1 + 4 + 2}{4} = 2.5. \quad (3.1)$$

where  $s$  is the down-scaling factor and  $d$  is the number of dimensions.




**Figure 3.3:** Illustration of down-scaling an image by a factor of 2. The down-scaled image is of size  $(1 \times 1)$  and has an intensity value given by the average of the four pixels in the original resolution image

### 3.2 Working with Multi-Resolution Image Representation

The tomograms are constituent of five material classes of interest: implant, bone mineral, blood, osteocyte, and air/resin.

The implant measures  $(8 \times 3.5)$ mm, but some of the implant features, namely the micro screw threads, are far smaller. For this reason, we choose the  $16x$  down-scaled tomograms for segmentation of the implant, where the voxel length is  $1.875 \cdot 16 = 30\mu\text{m}$ , as the representation at this resolution level sufficiently retain all the features of the implant.

The method for segmentation of the bone mineral, described in Chapter 6, merely requires sufficient voxels to create a faithful histogram of voxel intensities in the tomogram. In Fig. 3.4, we show the histograms of voxel intensities in the full resolution tomogram, the  $16x$  down-scaled and the  $32x$  down-scaled tomograms, respectively. We see that the histogram of the  $16x$  down-scaled tomogram appears distributed with similar characteristics as the full resolution histogram, while the histogram of the  $32x$  down-scaled tomogram is too rough due to insufficient statistics. We will thus use the  $16x$  down-scaled tomogram for segmentation of the bone mineral.



**Figure 3.4:** The histograms of voxel intensities in the full resolution, 16x down-scaled and 32x down-scaled tomograms. The 16x down-scaled histogram appears similarly distributed with the full resolution histogram, while the 32x down-scaled histogram is rougher due to insufficient statistics.

The blood vessel network within the bone is itself a material with important structures at varying length scales, varying from a diameter of 3 – 45 $\mu\text{m}$  in nonhuman samples [32][33]. The 16x down-scaled image with a voxel length 30 $\mu\text{m}$  is sufficient for segmentation of the largest blood vessels, while the full-resolution image with a voxel length 1.875 $\mu\text{m}$  is required for segmentation of the smallest blood vessels. Thus, to segment the full blood vessel network, we will use multiple levels of representation in tandem. This method is described in detail in Chapter 8.

The osteocytes will be segmented through their characteristic volume. This volume varies from 31.25 – 2000  $\mu\text{m}^3$  [3]. To segment the smallest osteocytes, we will thus require the full resolution representation of the tomogram, with a voxel volume of  $(1.875 \times 1.875 \times 1.875) \approx 6.592\mu\text{m}^3$ , as the next finest resolution has a voxel volume of  $(3.75 \times 3.75 \times 3.75) \approx 52.73\mu\text{m}^3$ , neglecting the smallest osteocytes.

For most purposes, and certainly for ours, it is not practical to approximate the full resolution image at level  $n$  of the image pyramid. Instead, we truncate our image pyramid at level 4, where the down-scaled image is of size  $(204 \times 217 \times 217)$ , as representations at lower resolutions will not be useful for our analyses, as argued above. Our image pyramid thus has five levels, named as levels 0 through 4, as image pyramids traditionally are zero-indexed.

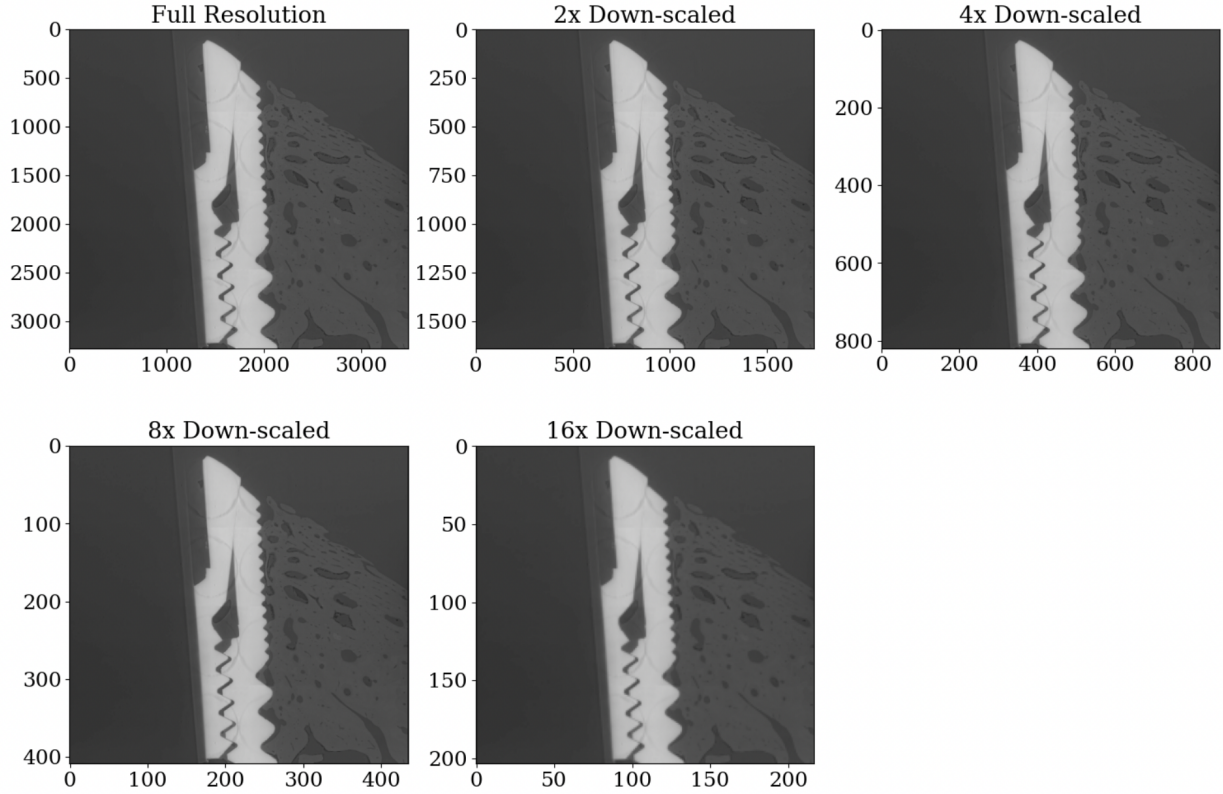
As we are working with 3D data, each down-scaling by a factor 2 of a higher resolution image,  $B$ , results in a lower resolution image,  $A$ , that is  $2^3 = 8$  times smaller,

$$A_{n,m,l} = \frac{B_{n',m',l'}}{8} \quad (3.2)$$

where  $B_{n',m',l'} = (B_{n,m,l} + B_{n+1,m,l} + B_{n,m+1,l} + B_{n,m,l+1} + B_{n+1,m+1,l} + B_{n+1,m,l+1} + B_{n,m+1,l+1} + B_{n+1,m+1,l+1})$ .

The image pyramid thus consists of representations ranging from the full resolution image of size  $(3279 \times 3480 \times 3480)$  to the 16x down-scaled image at size  $(3279/2^4 \times 3480/2^4 \times 3480/2^4) = (204 \times 217 \times 217)^2$ , which is represented by 4134 times fewer voxels. In Fig.3.5, we visualise the tomogram for reach resolution representation in the image pyramid. In Table 3.1, we include a table of resolution, shape, total number of voxels and memory requirement for each level of representation on the image pyramid.

<sup>2</sup> The edges of the tomogram is not vital to the analysis, so these are trimmed by up to 15 voxels in each dimension to ensure an exact down-scaling.



**Figure 3.5:** Full resolution tomogram (top left), 2x down-scaled tomogram (top centre), 4x down-scaled tomogram (top right), 8x down-scaled tomogram (bottom left), 16x down-scaled tomogram (bottom left). Note that the down-scaled tomograms represent the same imaged physical space, only with fewer voxels.

**Table 3.1:** Table of resolution, shape, total number of voxels and memory requirement for each level of representation on the image pyramid.

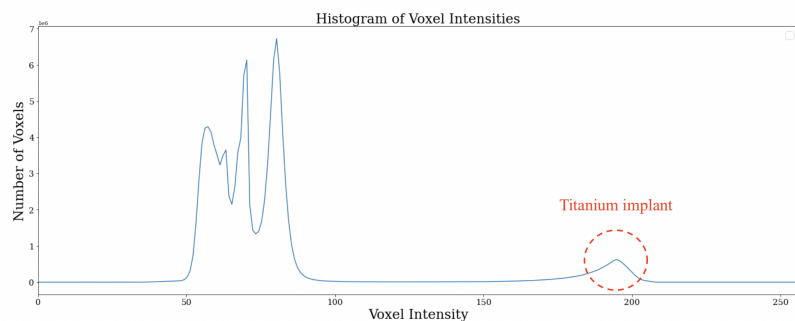
Lvl	Resolution	Shape	Voxels	Size
0	Full	$(3279 \times 3480 \times 3480)$	$3.971 \cdot 10^{10}$	40GB
1	2x Down-scaled	$(1639 \times 1740 \times 1740)$	$4.962 \cdot 10^9$	5GB
2	4x Down-scaled	$(819 \times 870 \times 870)$	$6.199 \cdot 10^8$	625MB
3	8x Down-scaled	$(409 \times 435 \times 435)$	$7.739 \cdot 10^7$	78MB
4	16x Down-scaled	$(204 \times 217 \times 217)$	$9.606 \cdot 10^6$	9.8MB

## 4 Implant Segmentation

The first step in image processing processes often involves feature extraction, where the information in the original data is transformed into numerical features that are more easily processed than the original data. An example is using the histogram of an image rather than the image for further processing.

In Fig. 4.1, we show the histogram of voxel intensities in the tomogram. The  $x$ -axis represents the voxel intensities and the  $y$ -axis the number of voxels in the tomogram with that intensity. The voxels comprising the implant can be seen on the far right of the histogram, as highlighted by the red circle, separate from the rest, as their intensities are far higher than those of the rest of the tomogram. This would lead it to be simply segmented through voxel-based thresholding. Thus, we will begin the segmentation of the tomogram into material classes with segmentation of the implant. In Appendix D, we include a flow chart of the implant segmentation solution.

**Figure 4.1:** Histogram of voxel intensities in the tomogram with titanium implant voxels highlighted by a circle in red..



### 4.1 Voxel-based Thresholding

The simplest method of image segmentation is through voxel-based thresholding. For this, feature extraction is performed by obtaining a histogram of voxel intensities in the image, which is then used to determine the voxel value threshold(s) at which to split the image into partitions.

Thresholding is an example of a voxel-wise operation, namely a group of image processing operations applied to each voxel individually, where the position of the voxel, and therefore the order in which the operation is applied to each voxel, are inconsequential.


When the threshold,  $\tau$ , has been chosen, the image  $I_{\text{in}}(x, y)$  is binarised so that all voxels with values below the threshold are given the value 0, and all voxels with values above the threshold are given the value 1, i.e.,

$$I_{\text{out}}(x, y, z) = \begin{cases} 0, & \text{if } I_{\text{in}}(x, y, z) \leq \tau. \\ 1, & \text{if } I_{\text{in}}(x, y, z) > \tau. \end{cases} \quad (4.1)$$

Thresholding is simple and fast, applying one principle to the whole image at once, but has limited usage, as the voxels of interest must be very similar to each other and very different from the background voxels in order to yield a meaningful result.

Reasons global thresholding becomes problematic include noise, uneven illumination, objects with similar pixel intensity values with the background or several objects with overlapping pixel intensity values. In Appendix E, we include three examples that illustrate these limitations of threshold-based segmentation.

For the segmentation of the implant in the tomogram, the implant appears trivial to segment upon initial inspection, as it is comprised of voxels with values far higher than the rest of the tomogram. In Fig. 4.2, where we show a proposed threshold for segmentation of the implant based on visual inspection of the histogram.



**Figure 4.2:** Histogram of voxel intensities in the tomogram with proposed threshold for voxel-based segmentation of the implant.

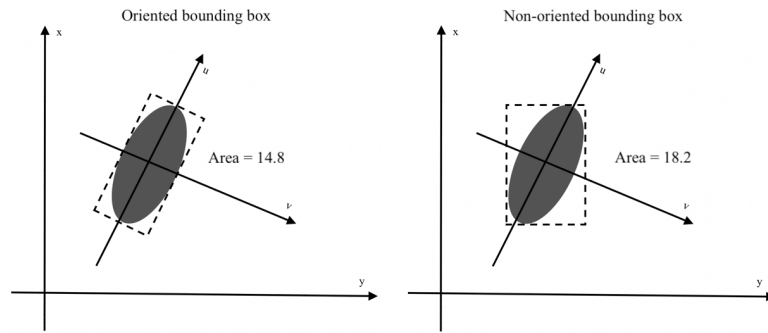
However, bleeding from the implant leading to brightening of voxel values close to the implant, along with beam hardening causing the voxel range to shift with tomogram depth can cause other voxels to have a high enough value to be incorrectly segmented as implant [6].

One way to combat this misclassification is to consider the spatial information of the voxels rather than merely their numerical value. This is the purpose of region identification, which we can accomplish through the drawing of a bounding box for restricting the region of interest for implant segmentation to the area bounded by the box. In Appendix F, we include an introduction to the bounding box in the context of its most frequent application, instance segmentation.

### 4.2 Bounding Box

We are interested in constructing a minimum bounding box, which is defined as the minimum volume box that encloses every point of interest within it. By nature, minimum bounding boxes for a single object are also oriented bounding boxes, as there is no non-oriented box which can contain every point of interest in a smaller area than an oriented box, as shown in Fig. 4.3. Oriented bounding boxes are defined by having faces parallel with the object coordinate frame axes, known as the principal axes, of the object they are enclosing, as opposed to with faces parallel with the image coordinate frame axes. The object coordinate frame and principal axes will be defined in Section 4.3.

**Figure 4.3:** Example of an oriented object with its oriented minimum bounding box (left) and the same object with a non-oriented bounding box. The area of the oriented bounding box will always be smaller, unless the object is entirely symmetric, in which case the two bounding boxes will be identical. The discrepancy in area will increase as the asymmetry increases. The areas are presented unitlessly, as this is irrelevant for comparison.



To draw an oriented bounding box that we can use for restricting the region of interest for implant segmentation to the area bounded by the box we wish to find the vectors that define the principal axes of the implant in order to orient the bounding box.

### 4.3 Principal Axes Analysis

The object coordinate frame is a frame of reference which gives the coordinates  $(u, v, w)$  of a voxel in relation to the object and has origin  $(u, v, w) = (0, 0, 0)$  in the center of mass of the object. The three axes of object coordinate frame, known as the principal axes, are given by the three vectors  $\bar{u}, \bar{v}, \bar{w}$  in the direction of maximum variance from the center of mass of the object. In Fig. 4.3, we illustrated the object coordinate frame in 2D, with principal axes  $u, v$  in the direction of vectors  $\bar{u}, \bar{v}$ .

For a 3D object, the vectors that define the principal axes are given by the three eigenvectors associated with the three eigenvalues of the moment of inertia matrix of the implant data, as these eigenvectors give the directions of maximum variance of the data.

The moment of inertia is a measure of the rotational inertia of an

object, i.e. its resistance to change in angular velocity about a given axis of rotation[34]. The rotational inertia of an object about a given axis of rotation is highly dependent on how the mass of the object is distributed about the axis, as mass further away from the axis of rotation will need to move at a higher velocity than mass nearer to the axis. Changing the rotational velocity of an object about a longer axis thus requires more force than changing the rotational velocity of an object about a shorter axis[34].

Therefore, the moment of inertia is a measure of variance of the mass distribution around a given axis, and the eigenvectors of the moment of inertia matrix will give the directions of maximum variance of the data. The extent of the implant in the directions of the principal axes can further be used for determining the dimensions of the bounding box.

The first step in obtaining the principal axes of the implant is the construction of a moment of inertia matrix,  $\Sigma$ , given by

$$\Sigma = \begin{bmatrix} \mu_{200} & \mu_{110} & \mu_{101} \\ \mu_{110} & \mu_{020} & \mu_{011} \\ \mu_{101} & \mu_{011} & \mu_{002} \end{bmatrix} \quad [35] \quad (4.2)$$

where the elements of the moments of inertia matrix for an image  $I(n, m, l)$  are given by

$$\mu_{pqr} = \sum_n \sum_m \sum_l (n - n_c)^p (m - m_c)^q (l - l_c)^r I(n, m, l) [36] \quad (4.3)$$

where  $n_c, m_c$  and  $l_c$  are the indices of the center of mass of the object in  $I$ , in our case the implant.

We can now find the eigenvalues of the moment of inertia matrix through the determinantal equation,

$$\det \begin{bmatrix} \mu_{200} - \lambda & \mu_{110} & \mu_{101} \\ \mu_{110} & \mu_{020} - \lambda & \mu_{011} \\ \mu_{101} & \mu_{011} & \mu_{002} - \lambda \end{bmatrix} = 0 \quad (4.4)$$

which we can solve for its three roots to obtain the three eigenvalues,  $\lambda_1, \lambda_2, \lambda_3$ . To obtain their corresponding eigenvectors,  $\bar{e}_1, \bar{e}_2, \bar{e}_3$ , we can find these by solving

$$\begin{bmatrix} \mu_{200} & \mu_{110} & \mu_{101} \\ \mu_{110} & \mu_{020} & \mu_{011} \\ \mu_{101} & \mu_{011} & \mu_{002} \end{bmatrix} \begin{bmatrix} e_{i,1} \\ e_{i,2} \\ e_{i,3} \end{bmatrix} = \lambda_i \begin{bmatrix} e_{i,1} \\ e_{i,2} \\ e_{i,3} \end{bmatrix} \quad (4.5)$$

for  $i \in [1, 2, 3]$ .

We now wish to apply the eigenvectors in the transformation of the box from the image coordinate frame,  $(n, m, l)$ , to the object coordinate frame,  $(u, v, w)$ , in order to make the bounding box oriented. This firstly involves the translation of the bounding box to make it centered in the origin of the object coordinate frame, and subsequently rotation using the eigenvectors as a rotation matrix to make align the box with the principal axes of the implant. This transformation between reference frames for image analysis is affectionately known as Procrustes analysis.

#### 4.4 Procrustes analysis

<sup>1</sup>Named after the mythical Greek figure, Procrustes, a bandit who stretched or cut off the limbs of his victims to forcibly fit them to an iron bed. This is analogous with what we are attempting to do with the bounding box, albeit a bit less grisly.

<sup>2</sup>Note that scaling is not used in this thesis work.

To transform bounding box from  $(n, m, l)$  to  $(u, v, w)$  coordinates, we will use the Procrustes analysis<sup>1</sup>, which displaces, rotates and scales<sup>2</sup> an object to align the object with another without altering the shape of the object.

In Fig. 4.4, we give an example of the Procrustes analysis between two objects, defined by the coordinates of their points,  $A = [(x_{a_1}, y_{a_1}), (x_{a_2}, y_{a_2}), \dots, (x_{a_n}, y_{a_n})]$  and  $B = [(x_{b_1}, y_{b_1}), (x_{b_2}, y_{b_2}), \dots, (x_{b_n}, y_{b_n})]$ , where we wish to align  $A$  with  $B$ . The first step is the translation of  $A$  to center it in the center of  $B$ . For this, we calculate the centers of both  $A$  and  $B$ ,

$$c_{A,x} = \frac{x_{a_1} + x_{a_2} + \dots + x_{a_n}}{n} \quad (4.6)$$

and

$$c_{A,y} = \frac{y_{a_1} + y_{a_2} + \dots + y_{a_n}}{n} \quad (4.7)$$

and likewise for  $B$ . We subtract the center of  $A$  from the center of  $B$  to obtain the the offset,  $t = [t_x, t_y]$

$$\begin{bmatrix} t_x \\ t_y \end{bmatrix} = \begin{bmatrix} c_{B,x} \\ c_{B,y} \end{bmatrix} - \begin{bmatrix} c_{A,x} \\ c_{A,y} \end{bmatrix} \quad (4.8)$$

which we add to the coordinates of  $A$  to displace it, so for each coordinate of  $A$ ,  $[x_{a_i}, y_{a_i}]$ , for  $i = [1, n]$ ,

$$\begin{bmatrix} x'_{a_i} \\ y'_{a_i} \end{bmatrix} = \begin{bmatrix} x_{a_i} \\ y_{a_i} \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \quad (4.9)$$

To scale  $A$ , we obtain the sizes of both  $A$  and  $B$ ,

$$s_{A,x} = \sqrt{\frac{||x_{a_1}|| + ||x_{a_2}|| + \dots + ||x_{a_n}||}{n}} \quad (4.10)$$

and

$$s_{A,y} = \sqrt{\frac{||y_{a_1}|| + ||y_{a_2}|| + \dots + ||y_{a_n}||}{n}} \quad (4.11)$$

And likewise for  $B$ . We obtain the ratio between the two to obtain the scaling factor,

$$s_x = \frac{s_{B,x}}{s_{A,x}} \quad (4.12)$$

and

$$s_y = \frac{s_{B,y}}{s_{A,y}} \quad (4.13)$$


and scale  $A$  with this factor for each coordinate of  $A$ ,

$$\begin{bmatrix} x'_{a_i} \\ y'_{a_i} \end{bmatrix} = \begin{bmatrix} x_{a_i} \\ y_{a_i} \end{bmatrix} \begin{bmatrix} s_x \\ s_y \end{bmatrix} \quad (4.14)$$

Finally, we rotate  $A$ , which, as mentioned above, we do by finding the eigenvectors,  $\vec{u} = [u_1, u_2]$  and  $\vec{v} = [v_1, v_2]$ , of the moment of inertia matrix of  $B$ , and using these as the rotation matrix for a coordinate system change. For every  $A$  coordinate,

$$[x'_{a_i}, y'_{a_i}] = \begin{bmatrix} u_1 & u_2 \\ v_1 & v_2 \end{bmatrix} \begin{bmatrix} x_{a_i} \\ y_{a_i} \end{bmatrix} \quad (4.15)$$





**Figure 4.4:** An illustration of the Procrustes analysis.  $A$  is translated, scaled and rotated with respect to  $B$ .

The Procrustes analysis is complete.

Before we carry on with the technical details of the solution, we bring an interlude introduction to morphology, which we will use frequently in this thesis, presently for pre-processing of the implant before the drawing of the bounding box.

#### 4.5 Morphology

The implant has a hollow space in the center, which we would like to close, as it may lead to a skewed distribution of the implant mass for computation of the principal axes. For this purpose, we will use morphological processing. Morphological processing lends itself primarily to the analysis and manipulation of binary images. Binary images have no intensity, colour, contrast, gradients, texture or depth, but have form, shape, structure, size, location and orientation, which are the features that morphology concerns itself with.

Morphological operators are neighbourhood operators, i.e., the operation applied to each voxel depends on the neighbouring voxels. A neighbourhood is defined by a structural element chosen by the user, which moves over the input image data array and applies the morphological operation at each step.

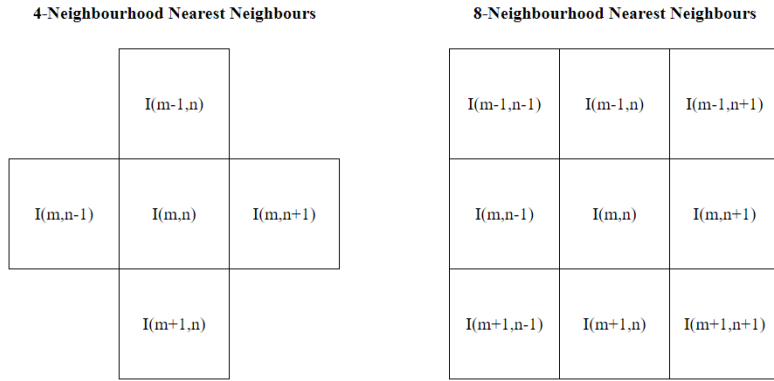
The user-defined structural element is usually a small<sup>3</sup> binary image that can have any shape and reach, but there are only two ways of defining nearest neighbours, as seen in Fig. 4.5. For 2D images, a 4-neighbourhood of an  $I(m, n)$  pixel are the pixels  $I(m, n \pm 1)$  and  $I(m \pm 1, n)$ . An 8-neighbourhood extends the definition to include the diagonal neighbours,  $I(m \pm 1, n \pm 1)$ .

For 3D images, a 4-neighbourhood of an  $I(m, n, l)$  pixel are the pixels  $I(m, n \pm 1, l)$ ,  $I(m \pm 1, n, l)$ , and  $I(m, n, l \pm 1)$ . The 3D 8-neighbourhood further includes the diagonal neighbours,  $I(m \pm 1, n \pm 1, l)$  and  $I(m \pm 1, n \pm 1, l \pm 1)$ .

The simplest of morphological operators, and the basis for all other operators, are erosion and dilation. If we have an input image,  $A$ , and a structuring element  $B$ , and we obtain the output image  $C$  through morphological processing, we write the expression for morphological erosion as

$$A \ominus B = \{z | (B)_z \subseteq A \text{ and } A \subseteq I\} \cup \{A^c | A^c \subseteq I\} = C \quad (4.16)$$

<sup>3</sup> i.e. suitably smaller than the input image



**Figure 4.5:** 4-Neighbourhood Nearest Neighbours (left), and 8-Neighbourhood Nearest Neighbours (right).

and dilation as

$$A \oplus B = \{z | [(B)_z \cap A] \subseteq A\} = C \quad (4.17)$$

<sup>4</sup> An operation  $\circ$  is commutative only if  $x \circ y = y \circ x$  for all  $x$  and  $y$ .

<sup>5</sup> An operation  $\circ$  is distributive only if  $x \circ (y \cup z) = (x \circ y) \cup (x \circ z)$  for all  $x, y$ , and  $z$ .

<sup>6</sup> An operation  $\circ$  is associative only if  $(x \circ y) \circ z = y \circ (x \circ z)$  for all  $x, y$  and  $z$ .

Erosion and dilation are both commutative<sup>4</sup> and distributive<sup>5</sup>, but while dilation is associative<sup>6</sup>, erosion is not. Rather,

$$(A \ominus B) \ominus C = A \ominus (B \oplus C). \quad (4.18)$$

However, erosion and dilation do not commute with each other, a fact from which two simple composite operators can be made: opening and closing. Opening of an image performs erosion on the original image and then dilation on the eroded image, removing small, isolated objects and separating bridged objects, and is written as

$$A \odot B = (A \ominus B) \oplus B. \quad (4.19)$$

Closing of an image performs dilation on the original image and then erosion on the dilated image, filling small holes inside objects and strengthening the bridge between bridged objects, and is written as

$$A \bullet B = (A \oplus B) \ominus B. \quad (4.20)$$

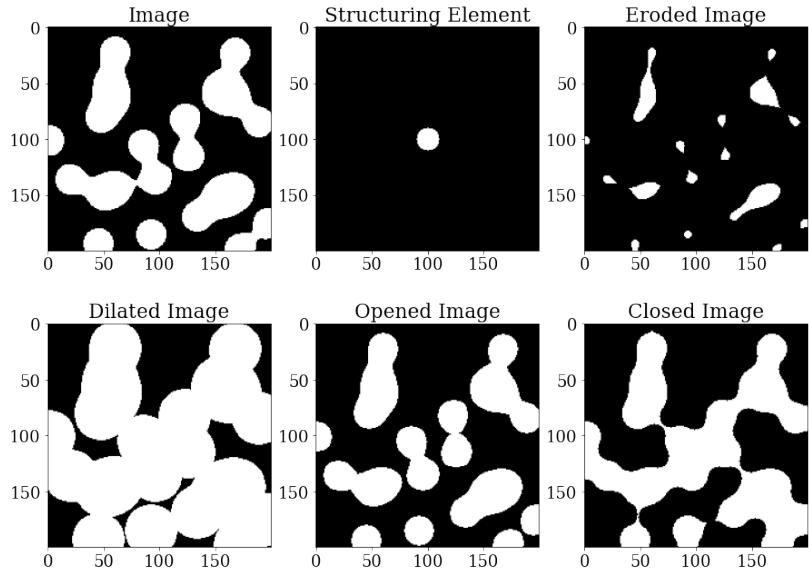
We will use morphological closing to close the hollow space in the center of the titanium implant to obtain an even distribution of implant mass for computing the implant principal axes.

All four basic morphological operations are illustrated in Fig. 4.6


#### 4.6 Implementation and Results

In order to find the central image moments of the implant, we must obtain the implant through segmentation first. Although we have established that a voxel-intensity based thresholding is insufficient to properly segment the implant, we can use it as a pre-processing step in order to obtain a crude segmentation of the implant that we can use for computing the principal axes, as the few misclassified voxels due to noise in the tomogram contribute negligently to the total segmented mass and will thus not impact the principal axes analysis. We

**Figure 4.6:** Original image of size  $(200 \times 200)$  created with `skimage.data.binary_blobs` (top left), disk structuring element of radius (10) (top centre), morphologically eroded image (top right), morphologically dilated image (bottom left), morphologically opened image (top centre), morphologically closed image (bottom right).



also binarise the image as to not give bias to brighter implant voxels, as these might falsely reflect a higher mass, and perform binary morphological closing to eliminate any gaps in the implant, as these gaps may likewise give undue bias to some facets of the implant. The visual output of the pre-processing is shown in Fig. 4.7.



**Figure 4.7:** Tomogram pre-processing: Tomogram before pre-processing (top left), threshold-segmented implant (top right), binarised threshold-segmented implant (bottom left) and morphologically closed binarised threshold-segmented implant (bottom right).

We find the principal axes by constructing a moment of inertia ma-

trix and finding the eigenvalues,  $\lambda = [\lambda_1, \lambda_2, \lambda_3]$ , and eigenvectors,  $\bar{u} = [u_1, u_2, u_3]$ ,  $\bar{v} = [v_1, v_2, v_3]$ ,  $\bar{w} = [w_1, w_2, w_3]$ , as described in Section 4.3.

We wish to draw a box where the lengths are given by the lengths of the implant, i.e. given by the distance between the outermost implant voxels in the  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions. However, we cannot simply take the outermost threshold-segmented implant voxels in the  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions as the lengths, as this is highly susceptible to noise, such as a rogue misclassified voxel far out in the tomogram. Thus, we propose another method for obtaining the lengths of the implant, which involves including only the voxels which fall within the area that contains 99.9% of the threshold-segmented implant voxels.

To this end, we obtain the cumulative sum of the number of threshold-segmented implant voxels in  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions and divide the cumulative sum element-wise with the total number of threshold-segmented implant voxels in the tomogram to obtain the cumulative sum as a percentage. We then take the lengths of the implant to be the difference between two thresholds, determined by requiring  $> 0.0005\%$  for the first threshold and  $> 0.9995\%$ , for a total  $> 99.9\%$  of the voxels included within the thresholds for the  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions, respectively. We draw the lengths of the bounding box accordingly.

To obtain number of threshold-segmented implant voxels in  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions, we first obtain the  $(u, v, w)$ -coordinates by transforming the  $(n, m, l)$ -coordinates of the implant by dotting the  $(n, m, l)$ -coordinates with the eigenvectors found above,

$$[u, v, w] = \begin{bmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{bmatrix} \begin{bmatrix} n \\ m \\ l \end{bmatrix} \quad (4.21)$$


We separate the  $(u, v, w)$  coordinate array into three arrays that hold the  $u$ ,  $v$  and  $w$ -coordinates separately. The array of  $u$ -coordinates contains the  $u$ -coordinates of several voxels that have identical  $u$ -coordinates but varying  $(v, w)$ , and likewise for the two remaining axes. Therefore, we obtain a histogram of  $u$ ,  $v$ , and  $w$  coordinates, respectively, which give the distribution of voxels along the implant in  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions. We can then take the cumulative sum of these and find the thresholds for each direction, as described above.

In Figs. 4.8, 4.9, and 4.10, we show, for  $\bar{u}$ ,  $\bar{v}$  and  $\bar{w}$  directions, respectively, the closed, crudely segmented implant, the projection profile, and the normalised cumulative sum of voxels.


We have now drawn a bounding box with the desired dimensions, but with respect to the image coordinate frame,  $(n, m, l)$ . We now wish to apply Procrustes analysis to transform the box from the image coordinate frame  $(n, m, l)$  to the object coordinate frame,  $(u, v, w)$  in order to make the bounding box oriented.

First, we translate the box. The object coordinate frame has its origin in the mass center of the implant, so we calculate the mass center of the bounding box,  $c_{\text{box}}$ , and subtract this from implant center,

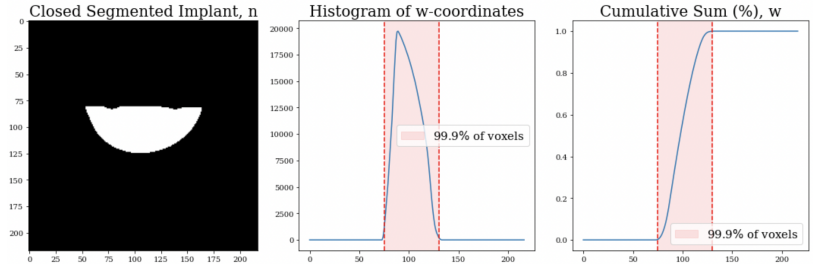
**Figure 4.8:** The closed, crudely segmented implant sliced in the  $l$  direction (left), the projection profile for voxels in  $\bar{u}$  direction (centre), and the normalised cumulative sum of voxels (right).



**Figure 4.9:** The closed, crudely segmented implant sliced in the  $m$  direction (left), the projection profile for voxels in  $\bar{v}$  direction (centre), and the normalised cumulative sum of voxels (right).



**Figure 4.10:** The closed, crudely segmented implant sliced in the  $n$  direction (left), the projection profile for voxels in  $\bar{w}$  direction (centre), and the normalised cumulative sum of voxels (right).



$c_{\text{implant}}$ , to find the offset,  $t = [t_n, t_m, t_l]$ ,

$$[t_n, t_m, t_l] = \begin{bmatrix} c_{\text{implant},n} \\ c_{\text{implant},m} \\ c_{\text{implant},l} \end{bmatrix} - \begin{bmatrix} c_{\text{box},n} \\ c_{\text{box},m} \\ c_{\text{box},l} \end{bmatrix} \quad (4.22)$$

which we add to the  $(n, m, l)$ -coordinates of mask to displace the bounding box,

$$[n', m', l'] = \begin{bmatrix} n \\ m \\ l \end{bmatrix} + \begin{bmatrix} t_n \\ t_m \\ t_l \end{bmatrix}, \quad (4.23)$$

obtaining new coordinates for the bounding box, still in  $(n, m, l)$ -coordinates but centered in the origin of the object coordinate frame, i.e. the center of the implant.


We can forego the scaling step, as we already drew the box in the desired size. We finish the Procrustes analysis by transforming the  $(n, m, l)$ -coordinates of the bounding box to  $(u, v, w)$ -coordinates by dotting the  $(n, m, l)$ -coordinates with the eigenvectors found above,

which serve as the rotation matrix,

$$[u, v, w] = \begin{bmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{bmatrix} \begin{bmatrix} n \\ m \\ l \end{bmatrix} \quad (4.24)$$


The resulting bounding box is shown in Fig. 4.11. We have padded the tomogram with zeros, as the bounding box may extend beyond the original image. Finally, we use the bounding box mask to the tomogram to restrict the implant area to the area bounded by the mask.

**Figure 4.11:** Bounding box mask shown from the three tomogram orientations.



We use voxel intensity based thresholding and multiply the results with the bounding box mask to segment the implant from the rest of the tomogram. We unpad the tomogram to obtain the resulting segmented implant shown in Fig. 4.12.

**Figure 4.12:** The segmented implant shown from three orientations.



## 5 *Implant Orientation and Separation into Parts*

In order to quantify the regenerated bone, we must be able to distinguish the regenerated bone from the old bone. As shown in Fig 2.3 in Chapter 2, the regenerated bone is grown around the part of the implant with the micro threads. Thus, we can find the threshold between old and regenerated bone by finding the boundary between the macro and micro threads.

To determine the threshold between the old and regenerated bone we want to obtain the screw thread outline, which we can view as a set of consecutive waves. We can then separate the screw threads by their amplitudes and wavelengths, as the macro threads will have larger amplitudes and larger wavelengths when compared with the micro threads.

In Chapter 4, Section 4.6, we obtained histograms of voxel coordinates along each principal axis. These histograms are also known as projection profiles, and we can obtain the screw thread outline as the projection profile of the segmented implant onto the plane that runs vertically through it.

### 5.1 *Image Projection Profile*

The projection profile of the  $\bar{u}$  principal axis exhibits troughs and peaks reflective of the screw thread outline. We will thus use the projection profile of the implant onto this axis as the screw thread outline in order to evaluate this as a sequence of waves, where we can use wave analysis to differentiate between the micro and macro threads.

In Chapter 4, Section 4.6, we found the projection profiles along each principal axis by obtaining histograms of voxel coordinates, but as the projection profile for a given axis of an image gives the sum of pixels along each increment of the axis, we may also obtain the projection profile by "summing away" the other axes,

$$proj(u) = \sum_v \sum_w I(u, v, w). \quad (5.1)$$

Producing the projection profile through summation will often produce a smoother result, as the deliberation over optimal binning of a histogram is eliminated. For the analysis in Chapter 4, smooth-

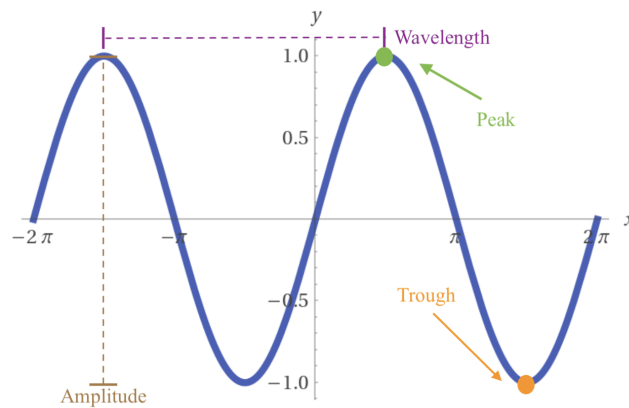
ness of the projection profile was not of concern, as only the position of the extremities was required to determine the lengths of the bounding box. However, for the analyses in this chapter, we will use the projection profile of the tomogram as the screw thread outline, for which smoothness becomes essential, as any noise will greatly influence the later analysis.

In Appendix G, we include an example of the projection profile of an image with text.

### 5.2 Wave Characteristics

To evaluate the screw thread as waves, we first have to find a way to define them as so. A wave is a propagating disturbance that deviates from equilibrium and is defined by two main characteristics, its wavelength,  $\lambda$ , and amplitude,  $A$ . The wavelength is defined as the horizontal distance between two subsequent peaks, or two subsequent troughs, and the amplitude as the vertical distance between a peak and a trough, as shown in Fig. 5.1.

**Figure 5.1:** An example of a propagating wave,  $f(x) = \sin(x)$ , with the wavelength, amplitude, a trough and a peak illustrated.



In order to define the wavelengths and amplitudes of the micro and macro threads, respectively, we must find the peaks and troughs.

### 5.3 Peak Finding

We begin the wave definition by finding the peaks and troughs of the projection profile. The projection profile is one dimensional, for which the construction of a peak finding algorithm is trivial. For an array  $[a, b, c, d, e \dots N]$ , we define  $b$  as a peak if and only if  $b \geq a$  and  $b \geq c$ . For the edges, the only conditions are  $a \geq b$  and  $N \geq (N - 1)$ , respectively. To find the troughs, the opposite conditions must simply be true, i.e. for an array  $[a, b, c, d, e \dots N]$ , we define  $b$  as a trough if and only if  $b \leq a$  and  $b \leq c$ , and, for the edges, if and only if  $a \leq b$  and  $N \leq (N - 1)$ .

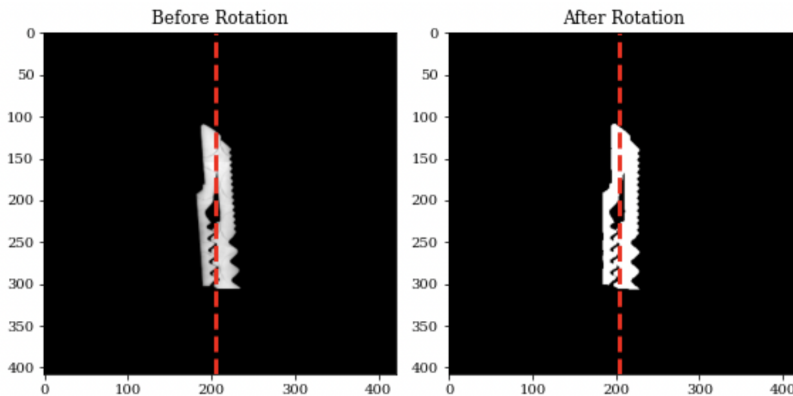


## 5.4 Implementation and Results

We recall that the implant is not aligned with the image coordinate frame. To simplify the computations, we can transform the implant from the object coordinate frame to the image coordinate frame. As we used the eigenvectors of the implant to rotate the bounding box to transform the bounding box from image coordinate frame  $(n, m, l)$  to the object coordinate frame,  $(u, v, w)$ , in Eq.4.24, we can do the reverse transform to rotate the implant by transforming the implant from the object coordinate frame to the image coordinate frame. As the inverse of a rotation matrix is its transpose,

$$[n, m, l] = \begin{bmatrix} u_1 & u_2 & u_3 \\ v_1 & v_2 & v_3 \\ w_1 & w_2 & w_3 \end{bmatrix}^T \begin{bmatrix} u \\ v \\ w \end{bmatrix} \quad (5.2)$$

The result is an implant that is upright with respect to the image coordinate frame, as shown in Fig. 5.2, where we have once again padded the image as the rotation may extend the implant beyond the original image.




**Figure 5.2:** Segmented implant before rotation (left), and segmented implant after rotation (right).

As we are working with 3D data, we must sum away two axes to obtain the desired projection profile. The screw thread outline is obtained as the projection profile of the segmented implant onto the plane that runs vertically through it, i.e. the plane aligned with the  $\bar{u}$  axis,

$$proj(u) = \sum_v \sum_w I(u, v, w). \quad (5.3)$$

The result is shown in Fig. 5.3.

In Section 5.3, we discussed the straight-forward nature of 1D peak-finding. However, this implementation is rather frail in the face of noise. The principal solution to this frailty is often to include a requirement for the minimum amplitude of the peak, i.e.  $b$  is a peak if and only if  $(b - a) \geq x$  and  $(b - c) \geq x$ , where  $x$  is the chosen minimum amplitude. However, this is not a feasible fix for our implementation, as we are looking for peaks at greatly varying amplitudes. We may attempt to limit the noise in our data instead.




**Figure 5.3:** Morphologically closed segmented implant (left), and implant screw thread projection profile (right).

An examination of our data shows that there is some noise due to asymmetry of the implant. Particularly the increasing of implant girth half-way through the region of the micro threads is of note.

Thus, for the sake of robustness, we take the gradient of the screw thread outline, which retains the wave characteristics of the projection profile but reduces noise, as shown in Fig. 5.4. We see that this gives a much cleaner representation of the various screw threads. We append the first and last non-zero values in the array as troughs, as they necessarily are troughs.

**Figure 5.4:** Screw thread projection profile gradient (left), and peaks and troughs in the screw thread projection profile gradient (right).



We define each wave as consisting of two troughs and one peak, or, alternatively, as two peaks and one trough, as shown in Fig. 5.5. We define the waves in this dual manner with the intention of taking the mean of their results as the final wave definitions to further boost the robustness. We thus create two arrays,  $l_{\text{valleys}}$  and  $l_{\text{peaks}}$ , which contain the indices of troughs and peaks, respectively. If there is no peak between two troughs, we automatically add a peak at the highest value between two troughs. Similarly, we automatically add a trough at the lowest value between two peaks if none exists.

When we have ensured that there is one peak between any two troughs, and one trough between any two peaks, we calculate the wavelength as the distance between subsequent troughs as well as the distance between subsequent peaks, resulting in two arrays,  $\lambda_{\text{valleys}}$  and  $\lambda_{\text{peaks}}$ , where

$$\lambda_{\text{valleys},i} = l_{\text{valleys},i+1} - l_{\text{valleys},i} \quad (5.4)$$

and

$$\lambda_{\text{peaks},i} = l_{\text{peaks},i+1} - l_{\text{peaks},i} \quad (5.5)$$

As shown in Fig. 5.6, we take the element-wise mean of the two to

obtain an array with the mean wavelength,  $\lambda_{\text{mean}}$ , where

$$\lambda_{\text{mean},i} = \frac{1}{2}(l_{\text{vals},i} + l_{\text{peaks},i}) \quad (5.6)$$


We obtain the value of the projection profile gradient at the index between subsequent troughs and calculate the difference between this value and the value of the projection profile gradient at the index of the peak between the troughs,

$$A_{\text{vals},i} = \nabla \left[ \frac{1}{2}(l_{\text{vals},i+1} - l_{\text{vals},i}) \right] \quad (5.7)$$


And similarly,

$$A_{\text{peaks},i} = \nabla \left[ \frac{1}{2}(l_{\text{peaks},i+1} - l_{\text{peaks},i}) \right] \quad (5.8)$$


We take the element-wise mean of the two as the amplitudes, as shown in Fig. 5.7.



**Figure 5.5:** Waves found through definition of two troughs and one peak (left), and waves found through definition of two peaks and one trough (right).



**Figure 5.6:** Wavelengths found through definition of two troughs and one peak (left), and through definition of two peaks and one trough (centre), and the mean of the two (right).



**Figure 5.7:** Amplitudes found through definition of two troughs and one peak (left), and through definition of two peaks and one trough (centre), and the mean of the two (right).

Having obtained the characteristics that define the waves, we require a similarity measure in order to distinguish groups of similar waves. The amplitudes and wavelengths of each thread of the two screw threads will not be exactly the same. Rather, the amplitudes and wavelengths of each oscillation of the micro threads will be similarly small, while the amplitudes and wavelengths of each oscillation

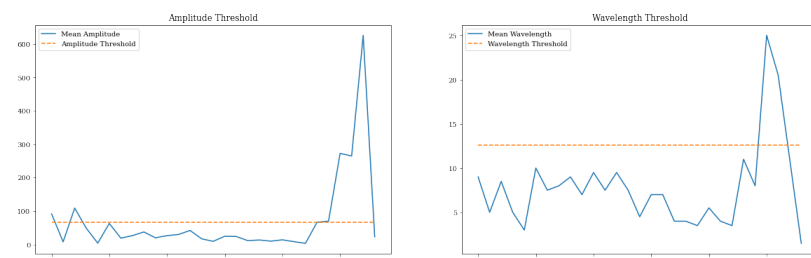
of the macro threads will be similarly large. In order to separate the screw threads by their amplitudes and wavelengths, we must find a definition for similarity. For this, we will use the central limit theorem. The central limit theorem states that for large data samples, the sample mean will be approximately Gaussian distributed, regardless of the distribution from which we are sampling. In practice, this means that if we perform an experiment enough times, e.g. measuring the length of a table, the value measured will eventually become Gaussian distributed. This is because uncertainties lead to the values of the measurements to deviate slightly, even though the length of the table has not changed.

If we take the amplitudes of, e.g., the micro threads as a measurement experiment of the same amplitude several times with small deviations due to uncertainties, we can view the amplitudes as a Gaussian distributed variable. For a Gaussian distribution, 98.7% of values fall within the interval of  $\pm 2.5\sigma$ , so can choose our measurement of amplitude similarity to be values that fall  $\pm 2.5\sigma$  above or below the mean, as 98.7% of values of measurements of the same physical quantity should fall within this interval. We can do the same with the wavelengths.

We will look for similarly small amplitudes to define the micro threads. As we are only considering small amplitudes, we calculate the mean of the amplitudes and only consider the amplitudes below the mean. Amongst these small amplitudes, we are only interested in the ones within our definition of similarity, so we take the mean of these small amplitudes and set a threshold as  $2.5\sigma$  above and below this mean to separate the similarly small amplitudes that comprise the amplitudes of the micro threads, as shown in Fig. 5.8.


For the wavelengths, we can only say something meaningful in relation to the larger wavelengths, as the wavelengths of the head and smaller screw thread of the implant are fairly indistinguishable. Thus, we only endeavour to separate the larger wavelengths. For this, we extract all the wavelengths above the mean, take the mean of these wavelengths and separate the similarly large wavelengths that comprise the wavelengths of the macro threads with a threshold  $2.5\sigma$  above and below this mean, as shown in Fig. 5.8.

**Figure 5.8:** Mean amplitude and amplitude threshold (left), and mean wavelength and wavelength threshold (right).



From the amplitude we are able to segment the head and macro threads from the micro threads, and from the wavelength we are able to segment the head from the macro thread using the thresholding


criteria set above, as shown in Fig. 5.9.




**Figure 5.9:** Small screw thresholds found through amplitude thresholding shown on projection profile gradient (left), and shown on segmented implant (right).

From Fig. 5.9, we see that the separation of the head of the implant and the micro threads is poor. This is because the end of the head of the implant has an amplitude sufficiently close to the amplitudes of the micro threads when the first threshold was found using all the amplitudes, including those of the macro threads. However, now that we have obtained the point of separation between the micro and macro threads, we can disregard the section of the implant with the macro threads to separately analyse the remainder of the implant.

We perform the same analysis as above with only the head and micro threads, as shown in Fig. 5.10, and obtain a new threshold, as seen in Fig. 5.11. When the macro threads are removed, the head of the implant is far easier distinguishable from the micro threads, as, with the very large amplitudes removed, the amplitudes of the head now no longer fall within the similarity measure.



**Figure 5.10:** Gradient of projection profile excluding the large screw thread (left), amplitude thresholding (center) and new implant head threshold (right).




**Figure 5.11:** Small screw thresholds found through amplitude thresholding shown with a implant head segmentation fix (left), and shown on segmented implant (right).

We can now also use the gradient of the projection profile from Fig. 5.4 in conjunction with the threshold between the macro and micro threads to determine the orientation of the image. We see the end

of the macro threads is indicated by the largest peak in the projection profile gradient, so we extract the index of the largest deviation from the mean of the gradient as the end of the macro threads, and if this index is larger than the threshold between the macro and micro threads, the implant is upside down and we can rotate the image  $\pi$  radians to obtain an upright implant.

Finally, we have a segmented, upright, vertically oriented implant, separated into its three constituents, as shown in Fig. 5.12. From this we also know the region of original bone, surrounding the macro threads, and the region of regenerated bone, surrounding the micro threads.



**Figure 5.12:** Implant head (left), small screw thread (center), and large screw thread (right).

## 6 Bone Segmentation

Returning to the histogram of voxel intensities from Fig. 4.2, we see that the rest of the tomogram will be an even greater challenge to segment than the implant, as there is significant overlap in the voxel values comprising the various remaining materials in the tomogram, rendering segmentation through simple thresholding unserviceable. However, this fact far from excludes thresholding from contributing, or the histogram from being valuable information in the segmentation of these materials.

The limitation of thresholding-based segmentation is that it draws a hard limit between groups of materials, even when there is grave uncertainty about the correct class. We wish to find a way to assign a probability that each voxel belongs to a specific class, rather than a binary classification. We can use the histogram to this end.

The histogram is a representation of the frequency distribution of voxel intensities in the tomogram, giving the count of the number of voxels with each specific intensity or over a each specific intensity range. If we normalise the histogram, it becomes the probability density function of the image.

The histogram probability density function is the superposition of the probability density functions of voxel intensities of individual object classes in the tomogram. If we can model these individual object class probability density functions, we can use these to assign a probability that each voxel belongs to a specific class, as desired. We will thus attempt to model the probability density functions using Gaussian functions. This method of describing a histogram with a mixture of Gaussian distributions is aptly named the Gaussian Mixture Model.

### 6.1 Gaussian Mixture Model

The Gaussian Mixture Model is defined as the weighted sum of multiple Gaussian components that represent a density of a particular random variable[37], given by [38]

$$G(\mathbf{x} | \boldsymbol{\mu}_i, \sigma_i) = \sum_{i=1}^M \alpha_i g(\mathbf{x} | \boldsymbol{\mu}_i, \sigma_i), \quad (6.1)$$

where  $x$  is the  $D$ -dimensional data vector,  $\alpha_i$  are the scalar Gaussian Mixture Model weights and  $g(\mathbf{x} | \boldsymbol{\mu}_i, \sigma_i)$  are the Gaussian den-

sity components, [38]

$$g(\mathbf{x} | \boldsymbol{\mu}_i, \boldsymbol{\sigma}_i) = \frac{1}{(2\pi)^{D/2} |\boldsymbol{\sigma}_i|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^T \boldsymbol{\sigma}_i^{-1} (\mathbf{x} - \boldsymbol{\mu}_i)\right), \quad (6.2)$$

where  $\boldsymbol{\mu}_i$  is the  $(1 \times D)$  mean vector, giving the center, and  $\boldsymbol{\sigma}_i$  the  $(D \times D)$  covariance matrix, giving the width.

As the histogram consists of a 1D array of points, for our case, the data vector is the 1D image intensity value array,  $I$ , and the mean vector and covariance matrix simply scalar values.

Thus, we describe the individual object class probability density functions as Gaussian distributions,

$$g_m(I) = \alpha_m e^{-\sigma_m |I - \mu_m|^2}, \quad (6.3)$$

where  $m$  is a given material class,  $\alpha_m$  is the height,  $\sigma$  is the width,  $\mu$  is the center, and the histogram  $H(I)$  as the sum of these distributions, along with a rest component  $g_{rest}(I)$ , which contains everything our model does not describe, so that

$$H(I) = \sum_m g_m(I) + g_{rest}(I). \quad (6.4)$$

We can estimate the probability that a given voxel belongs to a given material class as

$$p(m|I) = \frac{g_m(I)}{H(I)} \quad (6.5)$$

and

$$p(\text{rest}|I) = \frac{g_{rest}(I)}{H(I)} \quad (6.6)$$

which must satisfy the constraint,

$$\sum_m p(m|I) + p(\text{rest}|I) = 1. \quad (6.7)$$

As seen above, a Gaussian is characterised by its center and its standard deviation. The challenge is now to determine the optimal values for these parameters so that the Gaussian Mixture Model best describes the histogram. To find such best parameter values, we define an objective function that mathematically describes the target output so that the optimal parameters are the ones that minimise or maximise the objective function.

For our purpose, we define our objective function as

$$f(\theta) = \|H(I) - \sum_m \alpha_m g_m(I)\|^2, \quad (6.8)$$

<sup>1</sup>  $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_m]$ ,  $\boldsymbol{\mu} = [\mu_1, \dots, \mu_m]$ ,  $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_m]$ .

where  $\theta = [\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{\sigma}]^1$  contains the parameters for optimisation,  $H(I)$  is the histogram and  $\sum_m \alpha_m g_m(I)$  is and what is described by our Gaussian Mixture Model. As the optimal model is the one where this difference is zero,  $f(\theta) = 0$ , we aim to minimise the objective function. In Appendix H, we include further discourse on minimisation.




It is inefficient to blindly choose starting guesses for the parameters to feed to the minimising algorithm, as choosing well qualified guesses can significantly shorten the minimisation algorithm convergence time. To obtain good starting guesses, we will use a segmentation method known as K-means clustering.

## 6.2 K-Means Clustering

K-means clustering is an unsupervised classification algorithm that divides input data into  $K$  clusters, where  $K \in \mathbb{N}$  is chosen by the user.

In the simple 1D example in Fig. 6.1 we have an artificial data set of 20 values and choose the number of clusters  $K = 3$ . The clusters are defined by their respective centroids, the data point with the mean position of all the points in a given cluster. Three random initial centroids are initially selected by the algorithm. The algorithm assigns each data point to the same cluster as the nearest centroid. The measure used to define distance between two points in 1D is simply the difference between their numerical values. The centroids are then updated to be the mean data point of each cluster, and the data points are re-assigned to the cluster containing the nearest updated centroid. This recurs until the centroids do not change. The algorithm has converged.



**Figure 6.1:** An illustration of K-means clustering in 1D. Centroids are initiated, data points are assigned to clusters, centroids are updated, data points are re-assigned to clusters, centroids are updated, until the algorithm converges.

We can use this method to cluster the voxel intensity values that

constitute the histogram, from which thresholds can be obtained as the border between clusters. We can then choose the distance between subsequent thresholds as starting guesses for the variances, the mean value of each cluster as starting guesses for the centres, and the maximum value of the histogram for a cluster as starting guesses for the heights of the Gaussian distributions.

The methods used in this chapter for segmentation of the bone mineral, hereunder the Gaussian Mixture model and the K-means clustering for obtaining starting guesses, were first developed and implemented for 2D analysis in [39]. In this thesis, these methods are adapted for fully automatic segmentation of the bone mineral in 3D.


### 6.3 Implementation and Results

A K-Means Clustering algorithm is dependent on three variables, the data vector of points to cluster, the number of clusters, and the dimension of the data. The histogram contains information only regarding the intensity of the voxels, not their position in space, so the dimension is 1. We know there are at least four radiographic densities represented in the tomogram, excluding the implant: bone, soft tissue (blood, osteocytes), fat (resin), and air. We thus wish to cluster the points into 4 clusters, one for each of the known materials.

The time complexity of the K-Means Clustering algorithm is  $\mathcal{O}(n^{dk+1})$  [40], where  $d$  is the dimension,  $k$  is the number of clusters and  $n$  is the number of points to be clustered. The K-Means Clustering algorithm used in this thesis is from the Python library Scikit-learn and relies on a linear time approximation algorithm which reduces the complexity to  $\mathcal{O}(n)$ . The computation time is thus entirely dependent on the number of points to be clustered.

The histogram in Fig. 4.2 is comprised of  $9.606 \cdot 10^6$  points, one for each voxel in the 16x down-scaled tomogram, which is redundant, as we are not interested in the absolute value of the amplitude of the histogram, only the distribution of voxel intensities. Thus, we sample 5000 points that are distributed as the histogram to use for the K-Means Clustering algorithm. In Fig. 6.2, we check that the points are correctly distributed.


**Figure 6.2:** Histogram of tomogram voxel values (left), and histogram of 5000 values distributed as the original histogram (right).



The final intervals of voxel groups will thus be identical to those obtained with all  $9.606 \cdot 10^6$  points, but the algorithm will converge

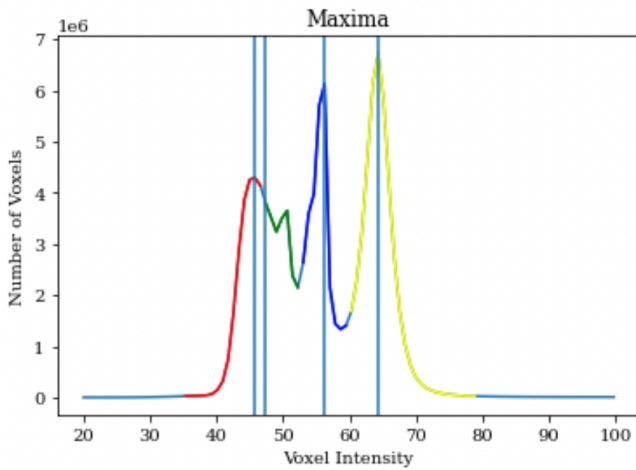
$\frac{9.606 \cdot 10^6}{5000} = 1921$  times faster.

We use the K-Means Clustering algorithm from `scikit-learn` with  $k = 4$ . The output of the algorithm are the beginning and end intensities of each cluster, which we use to construct the cluster regions shown in Fig. 6.3.



**Figure 6.3:** Histogram of voxel intensities with cluster regions found through K-Means Clustering.

As starting guesses, we take the widths of each cluster to be the widths of the distributions, the mean of the intensities in each cluster to be the means of the distributions, and the maxima in each cluster to be the heights of the distributions, as shown in Fig. 6.4. This crude peak-finding method is not perfect, but it provides a sufficient starting guess.



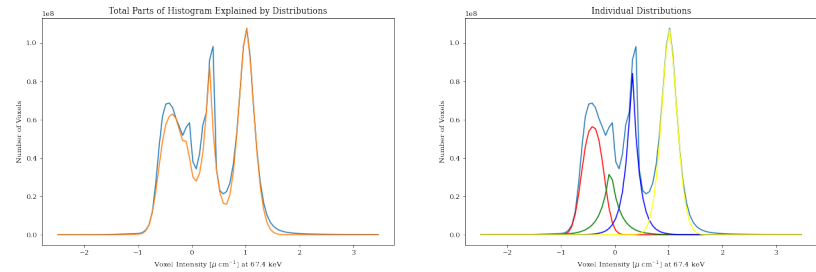
**Figure 6.4:** Histogram of voxel intensities with maxima of each region.

We feed these starting guesses to the minimising algorithm, subject to three inequality constraints,  $\alpha \geq 0$ ,  $\sigma \geq 0$ , and  $\mu \geq 0$ , i.e. the constraints that the parameters cannot be negative, as the variance and amplitude of a Gaussian distribution cannot be negative. The mean of a Gaussian distribution can be negative, but, as our histogram data is comprised of intensity values, the means cannot be negative either. Further, the fitted function must never exceed the histogram in value, as the sum of probabilities may never exceed 1, following the constraint given in Eq. 6.7. We include an overshoot

penalty to the optimisation algorithm to ensure this.


In Fig. 6.5, we show the best fit distributions returned by the minimising algorithm, as well as the portion of the histogram additively explained by them.

**Figure 6.5:** Total parts of histogram explained by distribution (left) and individual distributions (right)



We find the probabilities that any voxel belongs to any material class using Eq. 6.4 and obtain the probability distributions shown in Fig. 6.6. From left to right, the distributions represent the probability that a voxel belongs to air, resin, soft tissue (blood or osteocytes), and bone, respectively, as the higher density materials are represented by higher intensity values, following the radiographic densities.

**Figure 6.6:** Probability distributions. The x-axis represents voxel intensity values, and the y-axis represents the probability of belonging to a given material class. For any voxel intensity, the probability of the voxels with the intensity value belonging to a material class is given by the value of the probability distribution of the material class.




For illustration, in Fig. 6.7 we show a slice of the tomogram, with the implant removed, where each voxel is coloured according to the material it is most likely to comprise.


We can see from the probability distributions visualised in Fig. 6.6 there is a high proportion of the range of voxel intensities for the bone mineral probability that demonstrates a high probability of a voxel in this intensity range being a bone voxel, i.e. many of the voxels in the bone mineral probability array are higher than 90% certain that the voxel in question is a bone mineral voxel, while the probabilities for the other materials are waning in certainty. We will therefore not use the probabilities for segmentation of the osteocytes or blood vessels.

However, we obtain a bone mineral mask by binarizing the bone probability array to obtain an index map of voxels that comprise

**Figure 6.7:** Tomogram colored using probability distributions, where each voxel is coloured according to the material it is most likely to comprise.



bone with a certainty higher than 95%. The bone mineral mask is shown in Fig. 6.8.



**Figure 6.8:** Segmented bone mineral mask overlaid on tomogram slice.


We see that there are some misclassified bone mineral voxels in areas that must contain no bone, such as in the acrylic cylinder on the other side of the implant. We can improve on the segmentation using information we obtained during the implant segmentation. Having segmented the implant, we can begin by using the inverse of the segmented implant as a mask to remove any voxels within the implant area that have been misclassified as bone mineral, as shown in Fig. 6.9.



**Figure 6.9:** Improving on the bone mineral segmentation by limiting the bone mineral segmentation to outside the implant. The segmented implant mask for slice  $n = 60$  (left), the original segmented bone mineral mask (centre), and the improved segmented bone mineral mask (right).


Further, the bone grows around the implant in a half moon. We thus know that for each  $n$ -slice, down the implant, there should be no bone voxels before the first implant voxel in the  $w$ -direction. We obtain coordinates for each voxel belonging to the implant and impose the condition that there should be no bone voxels before the first implant voxel in the  $w$ -direction. In Fig. 6.10, we show an example for slice  $n = 60$ , where we see that it removes some of the misclassified voxels before the onset of the implant, although the half-moon shape of the implant preserves the misclassified voxels inside the crescent shape.

**Figure 6.10:** Improving on the bone mineral segmentation by limiting the bone mineral segmentation to after the first implant voxel. The segmented implant mask for slice  $m = 108$  (left), the once improved segmented bone mineral mask (centre), and the twice improved segmented bone mineral mask (right).



Further, we know that there should be no bone growth above the smaller screw thread, i.e. the bone mineral mask should be limited to below the threshold between the implant head and the smaller screw thread. Luckily, we have already found this threshold as well, so we can imply impose this condition to remove the misclassified voxels in this area from the segmentation. The final, improved bone mineral segmentation is shown in Fig. 6.11.


**Figure 6.11:** Improving on the bone mineral segmentation by limiting the bone mineral segmentation to under the threshold between the implant head and the smaller screw thread. The twice improved segmented bone mineral mask for slice  $l = 108$  and the threshold (left), and the thrice improved segmented bone mineral mask (right).



Having obtained a good separation at the top of the implant between the bone mineral and the imperfection in Fig. 6.10, we may now hope to remove this. To remove this imperfection, we notice that it is relatively small and separate from the bone mineral. This lends the bone mineral to be segmented from the imperfection through largest connected components analysis<sup>2</sup>. We perform morphological closing on the bone mineral mask in order for the bone mineral to be recognised as one connected component. We then extract the largest


<sup>2</sup> Connected component analysis implementation details will be given in Chapter 8, Section 8.1.

connected component, obtained simply as the connected component associated with the most frequently occurring connected component label. We take this as the closed bone mineral mask and multiply this with the thrice improved bone mineral mask to obtain the final improved bone mineral mask, as shown in Fig. 6.12.



**Figure 6.12:** Closed thrice improved bone mask (left), largest connected component of closed thrice improved bone mask (centre), and final improved bone mask (right).

In Fig. 6.13, we show the final segmented bone mineral mask for slices  $n = 108$ ,  $m = 108$  and  $l = 108$ , respectively.



**Figure 6.13:** The final segmented bone mineral mask for slices  $n = 108$ ,  $m = 108$  and  $l = 108$ , respectively.

With the bone segmented, we are able to assess the regenerated bone by quantifying the amount of regenerated bone. In the following chapter, we will perform such an intermediate quantification of the amount of the regenerated bone.

## 7 *Intermediate Bone Health Assessment*

The bone grows around the implant and upwards from the old bone region, starting at the surface of the artificially created defect. Thus, we would like the quantification of regenerated bone to reflect the amount of regenerated bone as a function of position along the implant. We will obtain both a measure of the amount of regenerated bone and conduct a bone mineral density test.

We can obtain a measure of the quantity of regenerated bone mineral as a function of position along the implant simply by taking the projection profile of the bone mineral mask obtained in Chapter 6. However, the segmented bone mineral mask is not a direct representation of the amount of bone regeneration, only the amount of bone mineral regeneration, as the segmented bone mineral mask neglects the volume of blood and osteocytes, which are part of the total regenerated bone volume. For the quantification of regenerated bone through a bone mineral density test, we require a measurement of both the regenerated bone mineral volume and the total regenerated bone volume in order to find the ratio between the two. We can obtain a measure of the quantity of total regenerated bone by performing morphological closing on the segmented bone mineral mask and taking the projection profile of this.


### 7.1 *Implementation and Results*

We obtain a measure of the amount of regenerated bone mineral by taking the projection profile of a binary mask of the segmented bone mineral. We close the segmented bone mineral mask to obtain the total bone volume mask, as shown in Fig. 7.1, which reflects the amount of regenerated bone, as it encompasses the blood vessels and osteocytes as well. We then obtain a measure of the amount of total regenerated bone by taking the projection profile of the total bone volume mask. The projection profile of bone mineral mask and the projection profile of the bone volume mask are both shown in Fig. 7.2.


It is now quite simple to perform a bone mineral density test. We simply take the area under the curve of projection profile of the bone mineral mask,  $f(i)$ , which is a representation of only the volume of bone mineral, and divide this by the area under the curve of the projection profile of the bone volume mask,  $g(i)$ , which is a representation of the total bone volume. We can obtain these areas by summing



**Figure 7.1:** Visualising a slice  $l = 108$  of bone mineral mask (left) and total bone volume mask, obtained through morphological closing of the bone mineral mask (right).



**Figure 7.2:** Projection profile of bone mineral mask and projection profile of bone volume mask.



away the remaining axis in each projection profile. Multiplying the areas under the curve of the projection profiles with the voxel volume will give the volumes of bone mineral and total bone, respectively. As the bone mineral mask was obtained using the 16x down-scaled tomograms, the voxel volume is  $(30 \times 30 \times 30)\mu\text{m} = 27000\mu\text{m}^3$ . Thus,

$$\begin{aligned} V_{\text{Bone Mineral}} &= \sum_{i=0}^n f(i) \cdot 27000\mu\text{m}^3 \\ &= 2177206 \cdot 27000\mu\text{m}^3 = 5.8785 \cdot 10^{10}\mu\text{m}^3 \end{aligned} \quad (7.1)$$

and

$$\begin{aligned} V_{\text{Bone}} &= \sum_{i=0}^n g(i) \cdot 27000\mu\text{m}^3 \\ &= 3135940 \cdot 27000\mu\text{m}^3 = 8.4670 \cdot 10^{10}\mu\text{m}^3 \end{aligned} \quad (7.2)$$

which yields a bone mineral density of

$$\text{BMD} = \frac{V_{\text{Bone Mineral}}}{V_{\text{Bone}}} = \frac{2177206}{3135940} \approx 0.694 \quad (7.3)$$


Thus, the bone mineral comprises 69.4% of the total bone volume.

However, we have performed this bone density quantification using all of the bone in the tomogram, when we would like to obtain a quantification of the old and regenerated bone separately. We can

then use the quantification of the old bone as a model for healthy bone, and evaluate the bone mineral density of the regenerated bone against the bone mineral density of the old bone to evaluate the bone mineral quantity of the regenerated bone. We can separate the old and regenerated bone by the threshold between the micro and macro threads found in Chapter 5.

In Fig. 7.3, we visualise the separate projection profiles of the bone mineral and bone, respectively.

**Figure 7.3:** Projection profiles of bone mineral masks of old and regenerated bone (left), and projection profiles of total bone masks of old and regenerated bone (right).



We obtain the bone density of the old bone first by calculating the area under the curve of the projection profile of the bone mineral mask for the old bone,

$$\begin{aligned} V_{\text{Old Bone Mineral}} &= \sum_{i=0}^n f_{\text{Old Bone Mineral}}(i) \cdot 27000\mu\text{m}^3 \\ &= 795251 \cdot 27000\mu\text{m}^3 = 2.1472 \cdot 10^{10}\mu\text{m}^3 \end{aligned} \quad (7.4)$$

and the area under the curve of the projection profile of the bone mask for the old bone,

$$\begin{aligned} V_{\text{Old Bone}} &= \sum_{i=0}^n g_{\text{Old Bone}}(i) \cdot 27000\mu\text{m}^3 \\ &= 1131986 \cdot 27000\mu\text{m}^3 = 3.0564 \cdot 10^{10}\mu\text{m}^3 \end{aligned} \quad (7.5)$$

which yields a bone mineral density of

$$\text{BMD}_{\text{Old Bone}} = \frac{V_{\text{Old Bone Mineral}}}{V_{\text{Old Bone}}} = \frac{908357}{985984} \approx 0.703 \quad (7.6)$$

Thus, the bone mineral comprises 70.3% of the total bone volume for the old bone.

We then obtain the bone density of the regenerated bone for comparison. We once again divide the area under the curve of the projection profile of the bone mineral mask for the regenerated bone with the area under the curve of the projection profile of the total bone mask for the regenerated bone,

$$\begin{aligned} V_{\text{regenerated bone Mineral}} &= \sum_{i=0}^n f_{\text{regenerated bone Mineral}}(i) \cdot 27000\mu\text{m}^3 \\ &= 1155364 \cdot 27000\mu\text{m}^3 = 3.1195 \cdot 10^9\mu\text{m}^3 \end{aligned} \quad (7.7)$$


$$\begin{aligned} V_{\text{regenerated bone}} &= \sum_{i=0}^n g_{\text{regenerated Mineral}}(i) \cdot 27000\mu\text{m}^3 \\ &= 2003954 \cdot 27000\mu\text{m}^3 = 5.4107 \cdot 10^{10}\mu\text{m}^3 \end{aligned} \quad (7.8)$$

which yields a bone mineral density of

$$\text{BMD}_{\text{regenerated bone}} = \frac{V_{\text{regenerated bone Mineral}}}{V_{\text{regenerated bone}}} = \frac{1493435}{1829618} \approx 0.577 \quad (7.9)$$

Thus, the bone mineral comprises 57.7% of the total bone volume for the regenerated bone. This is 12.6% less than for the old bone, indicating that the regenerated bone is more porous than the old bone.

In Fig. 7.4, we also obtain a slice by slice bone mineral quantification. We see that the bone of the regenerated bone at the top of the implant is significantly less dense in bone mineral than further down the implant. This could indicate that regenerated bone nearer the old bone is regenerated structurally more similar to the old bone. Looking at Fig.7.1, this may also be contributed to an effect of the bone mineral segmentation, as the segmentation at the edges at the top of the implant is subpar.



**Figure 7.4:** Slice by slice bone mineral density quantification for regenerated bone (left) and old bone (right).

In [41], the normal bone mineral density for men aged 20 – 29 is given as  $(1.134 \pm 0.112)\text{g}/\text{cm}^3$ . To obtain the measurement of bone mineral content for the regenerated bone in units of  $\text{g}/\text{cm}^3$ , we take the volume of bone mineral for the regenerated bone,  $3.1195 \cdot 10^9 \mu\text{m}^3$  and convert this to units of  $\text{cm}^3$  by multiplying by a factor of  $10^{-12}$ , obtaining a bone mineral volume of  $3.1195 \cdot 10^{-2}\text{cm}^3$ .

As bone weighs  $1.9\text{g}/\text{cm}^3$ [42], the weight of the regenerated bone mineral is  $3.1195 \cdot 10^{-2}\text{cm}^3 \cdot 1.9\text{g}/\text{cm}^3 = 5.771 \cdot 10^{-2}\text{g}$ .

To obtain the measure in  $\text{g}/\text{cm}^3$ , we must find the total volume of the bone in units of  $\text{cm}^3$ , which we can find by converting the volume of the total bone to units of  $\text{cm}^3$ ,  $5.4107 \cdot 10^{10} \mu\text{m}^3 \cdot 10^{-12} = 5.4117 \cdot 10^{-2}\text{cm}^3$ .

We can now obtain the measure bone mineral content for the regenerated bone in units of  $\text{g}/\text{cm}^3$  as  $\frac{5.771 \cdot 10^{-2}\text{g}}{5.4117 \cdot 10^{-2}\text{cm}^3} = 1.0666\text{g}/\text{cm}^3$ . We repeat the calculation for the old bone and obtain a measure of bone mineral content of  $1.2997\text{g}/\text{cm}^3$ .

For bone to be characterised as healthy, the  $t$ -score,

$$t = \frac{\text{measured} - \text{mean}}{\text{standard deviation}} \quad (7.10)$$

must satisfy  $t > -1$  [43], where no industry standard upper limit exists [44]<sup>1</sup>. However, we know that bone comprised of 100% bone mineral is not healthy, as living tissue requires, e.g., blood vessels

<sup>1</sup> In [44], the drawbacks related to the lack of an upper limit for the bone mineral density test are discussed.

for the transport of oxygen and nutrients. In [45], an upper limit of 2.5 was suggested. We will use this moderated condition and classify bone with a  $t$ -scores between  $-1$  and  $2.5$  as healthy bone.

For the regenerated bone,

$$t = \frac{1.0666 - 1.134}{0.112} = -0.6 \tag{7.11}$$

and for the old bone,

$$t = \frac{1.2997 - 1.134}{0.112} = 1.5 \tag{7.12}$$

The old and regenerated bone are thus both classified as healthy.

This bone regeneration evaluation should be run on all of the 35 tomograms produced in [5] and [6] to substantiate this quantitative model of healthy bone, as we expect an accurate model to classify the old bone in all 35 tomograms as healthy.

While we have now obtained a quantification of the amount of bone regeneration, this evaluation metric faces a glaring flaw in that we are in fact unable to say much accurately about the health of this regenerated bone, despite traditional evaluation methods, such as the bone mineral density test above, considering the two equivalent. Below, we provide a simple illustration of the limitations of the bone mineral density test.

Say that a checkerboard is a representation of healthy bone, where the white squares represent bone. In this case, a bone mineral density test will characterise all of the cases in Fig. 7.5, where black and white squares each have equity of about half the total area, as healthy bone. By the criteria for healthy bone given in [43], it will even characterise a board comprised entirely of white squares as healthy bone. It is indifferent to the distribution of bone mineral, as well as the presence, or absence, and distribution of critical materials, such as blood vessels and osteocytes.

**Figure 7.5:** Illustration of the shortcomings of the bone mineral density test.



(a) A checkerboard representing healthy bone. (b) An image with equal area covered by black and white squares, respectively, representing unhealthy bone that is inobservable by the bone mineral density test. (c) An image with equal area covered by black and white squares, respectively, representing unhealthy bone that is inobservable by the bone mineral density test.

Of particular note, two of the regeneration methods from [5] and [6] involve the placement of bone chips, i.e. 100% bone mineral, in the defect, which is intended to be resorbed by the osteoclasts and remodelled into living bone by the osteoblasts. However, a test in-

volving merely the amount of bone mineral could mistakenly classify the bone chips as healthy bone, even if no remodelling had taken place and the bone chips remained unresorbed.

Thus, the quantification above gives the amount of bone regeneration but reveals little about the health of the regenerated bone. For that, we must know information about the distribution of other materials throughout the bone, namely blood and osteocytes.

## 8 *Blood Vessel Segmentation*

The blood voxels are almost indistinguishable from osteocytes by voxel intensity value, especially with the noise present in the tomograms. However, the blood vessels have the distinction of belonging to a connected network of blood vessel voxels running throughout the bone. The blood vessels must thus comprise the largest connected component in the soft tissue of the tomogram.


In Chapter 6, we refrained from using the probabilities for segmentation of the soft tissue, as the analysis was not sufficiently confident in identifying the voxel intensities belonging to these tissue classes, i.e. blood and osteocyte. Inspecting Fig. 6.7, we see that the parts of the tomogram outside the bone volume, obtained through morphological closing of the bone mineral mask, is responsible for quite some noise, with the acrylic cylinder that the bone sample was suspended in for imaging possessing voxel intensities similar to those of the blood vessels and osteocytes. It is irrelevant to search for soft tissue in these volumes, as it exclusively exist inside the bone. We can thus use the segmented bone mineral to limit the segmentation efforts to the volume of interest, namely the bone volume.

For this purpose, we would like a binary mask that reflects the volume occupied by bone, including bone mineral, blood vessels, osteocytes, and anything else that exists within this volume. We already obtained this bone volume mask in Chapter 7 by performing morphological closing on the segmented bone mineral mask.

The bone mineral mask was obtained in the 16x down-scaled resolution, so to obtain the closed bone volume mask in the full resolution, we can simply resize or up-scale it, producing a multi-resolution representation of the mask as well. This resizing is the inverse to down-scaling. Rather than condensing the information from  $N$  voxels into 1 voxel, the information of 1 voxel is spread out on  $N$  voxels. As we are dealing with a binary image, the voxel intensity of the  $n$  voxels will not be distributed over the  $N$  voxels, but rather all  $N$  voxel will receive the value of its equivalent voxel in the lower resolution image. In Fig. 8.1, we illustrate this with a 2D example, where resizing the image by a factor of 2 results in the voxel intensity of 1 voxel to be received by  $(2 \times 2) = 4$  voxels in the resized image. In 3D, for resizing the closed bone volume masks, resizing by a factor 2 will result in the voxel intensity of 1 voxel to be received by  $(2 \times 2 \times 2) = 8$  voxels in the resized image.

Multiplying the binary bone volume mask with the tomogram,

**Figure 8.1:** A 2D example where resizing the image by a factor of 2 results in the voxel intensity of 1 voxel to be received by  $(2 \times 2) = 4$  voxels in the resized image.



we are left with only the bone volume of the tomogram, with bone mineral, blood vessels and osteocytes included, but without the surrounding acrylic cylinder surrounding the bone. Using the bone volume mask, we segment the air/resin surrounding the bone sample as the material in the tomogram not within the bone volume mask or implant mask. We separate the soft tissue from the bone mineral in the bone volume through automatic Otsu thresholding, which computes a suitable threshold by minimising the within-class variance between the voxels on both sides of the threshold. We can now perform a connected component analysis on the soft tissue in order to identify the largest connected component herein.

### 8.1 Connected Component Analysis

As the name implies, connected component analysis identifies distinct, disjoint objects in an image, for which we would like the computer to recognise their pixels or voxels as connected components and label each object with a unique ID[46].


Connected component analysis originates from graph theory. The binary mask is viewed as a graph, where each foreground voxel is a node, which is connected to all neighbouring foreground voxels. The connected components are subgraphs, within which all nodes have a path to all other nodes in the subgraph.

The most commonly used connected component algorithms in image analysis identify the connected components by performing a raster scan. Whenever a foreground voxel is encountered, the algorithm checks the prior neighbourhood, and one of the following three steps are performed for the voxel in question,

1. If all the neighbouring voxel are background voxel, the voxel in question is labelled with a new label, i.e. it belongs to a new object.
2. If there is only one foreground voxel in the neighbourhood, the voxel in question is given the same label as the label of this foreground voxel, i.e. they belong to the same object.
3. If there is more than one foreground voxel and they do not have the same label, the labels are merged: one label is chosen and all voxel(s) with the other label(s) are updated to have the chosen label, as they must all belong to the same object.

In Fig. 8.2 we show an example of connected component labelling. On the left, the original image is shown. Four shapes are immediately discernible for human perception. Through connected component analysis, four shapes are likewise discernible for computer perception, as shown on the right, where each shape of connected component is labelled.

**Figure 8.2:** Image with four discernible shapes (left), and connected component labels of image (right).



When attempting to segment the blood vessels, a few problems are encountered, most pressingly the problem of large data. The resolution of a down-scaled tomogram is too poor for the smaller capillaries to be made out, while the tomogram at full resolution is too large to be processed in RAM. Simply cutting the tomogram at full resolution into boxes and performing connected component analysis on these separately is faulty, as the separation of the tomogram may sever the very capillaries in question from the blood vessel network, leading the connected component analysis to misclassify them as separate objects, when they are in fact connected to the blood vessel network.

To circumvent this quandary, we propose a multi-resolution analysis method as a solution. The larger blood vessels can be obtained through the extraction of the largest connected component in a connected component analysis of low resolution data. The finer blood vessels can be extracted separately through finer resolution data and connected to the blood vessel network through multi-resolution analysis.

## 8.2 More Multi-resolution Analysis

As described in Chapter 3, we have been using multi-resolution analysis for the entirety of the thesis. We used down-scaled low resolution samples for the segmentation of the implant and bone, as this was sufficient for the identification of these high contrast or large structures. Smaller, low contrast objects, such as capillaries or osteocytes, require higher resolutions. For the segmentation of the entire blood vessel network, we require a number of various resolutions, and thus we will use multi-resolution analysis more explicitly.



We begin our multi-resolution analysis at the top of the image pyramid, the level  $n = 4$  image of size  $(204 \times 217 \times 217)$ . At this resolution level, the voxel length is  $(1.875 \cdot 2^4)\mu\text{m} = 30\mu\text{m}$ , which is just fine enough to segment the largest blood vessels. At lower resolutions, even the largest blood vessels, at a diameter of  $\sim 45\mu\text{m}$ [33], will be lost to aliasing.

We want the computations for the next layer of the image pyramid, the level 3 image of size  $(3279/2^3 \times 3480/2^3 \times 3480/2^3) = (409 \times 435 \times 435)$ , to cost the same as the level 4 image, which we can accomplish by processing eight  $(204 \times 217 \times 217)$  blocks of the  $(409 \times 435 \times 435)$  image separately. The trick is that we will likewise split the level 4 image into eight  $(102 \times 108 \times 108)$  blocks and compare each of these blocks to their equivalent  $(204 \times 217 \times 217)$  block of the level 3 image, as each down-scaled image voxel corresponds to eight image voxels of the image twice its resolution, i.e. one level higher on the image pyramid.

For each non-zero voxel in the level 4 segmented blood vessel image, we look at its corresponding eight voxels in the level 3 connected component labelling. If any of these eight are non-zero, we label them as blood, regardless of their component size, as they are connected to the blood vessel network through the lower resolution image. Further, any voxels in the level 3 image with the same label as the connected voxel in the given block are likewise labelled as blood, as they are connected to the blood vessel network through the connected voxel. We repeat this process until we reach the full resolution at level 0.

When we refer to splitting an image into blocks, we mean that we load only the part of the image that we are currently processing. We do this as we want out of core computation to avoid having the entire problem in RAM at once but rather continuously read partial problems from and write partial solutions to the disk. This means that when we are on level 4 of the image pyramid, we load the entire image, perform connected component analysis on this and extract the largest component as the blood vessel network for this level. On level 3 of the image pyramid, we load the first of eight blocks. However, in same same fashion as there are eight level 3 image blocks for the level 4 image, there are eight level 2 image blocks for each of the eight level 3 image blocks for a total of 64 level 2 image blocks and so on until we reach the base of the pyramid.

In Appendix I, we include a 2D toy example that was constructed to test and illustrate the multi-resolution analysis method.

In short, we have designed a set of sub-problems where each sub-problem is structurally identical to its parent problem. We have done so as there is a genus of algorithms, the recursive algorithms, that specifically lend themselves to this problem structure. Recursive algorithms are algorithms that arrive at a solution by calling itself until the solution, a stopping condition in the form of the base case, is reached. In Appendix J, we include a quip on the best known example of recursion, the Fibonacci sequence.

The elegance of the recursive algorithm is that it works as proof by induction. This will provide us with confidence in the performance of the algorithms on levels of the image pyramid that may be too large for us to visualise in full in three dimensions.

In the section below, we will detail how we construct and apply a recursive algorithm for segmentation of the blood vessel network.

### 8.3 Implementation and Results

As the sub-problems of a recursive algorithm are structurally identical to the parent problem, we design a solution for going from the base case  $f_0$  to the first sub-problem  $f_1$  and combine it with the recursive pattern. We will therefore first consider the problem of connecting the level 3 image connected component blocks with the level 4 segmented blood vessel image.

The level 4 image is of size  $(204 \times 217 \times 217)$ , while the level 3 image has twice as many voxels on each axis, yielding an image with eight times as many voxels at  $(409 \times 435 \times 435)$ . This means each level 4 voxel  $(n, m, l)$  corresponds to the eight level 3 voxels  $(n, m, l), (n + 1, m, l), (n + 1, m + 1, l), (n + 1, m, l + 1), (n, m + 1, l), (n, m + 1, l + 1), (n, m, l + 1), (n + 1, m + 1, l + 1)$ . Therefore, we need to process eight blocks of the level 3 image separately to proceed with an computational cost identical to the level 4 image. We will describe the blocks by their down-sampling factor and the position of the top, leftmost, front voxel of the block, i.e. the first level 3 block, comprised of the voxels with indices  $[0:204,0:217,0:217]$ , will be referred to as  $8_{(0,0,0)}$ , the next block, comprised of the voxels with indices  $[204:409,0:217,0:217]$ , as  $8_{(204,0,0)}$  and the final block, comprised of the voxels with indices  $[204:,217:,217:]$ , as  $8_{(204,217,217)}$ .

To avoid holding the entire solution in RAM at once we create an empty .hfd5 file to hold the level 3 segmented blood vessel image, to which we can write the partial solutions. We load the first of the eight level 3 image blocks,  $8_{(0,0,0)}$ , which spans the indices  $[0:204,0:217,0:217]$  of the full level 3 tomogram image, along with the corresponding  $16_{(0,0,0)}$  block spanning the indices  $[0:102,0:108,0:108]$  of the full level 4 segmented blood vessel image, which holds information about the same volume of the tomogram, only at half the resolution. We perform connected component analysis on the level 3 image block.

We are only interested in the voxels in the level 3 connected component analysis that connect to blood voxels in the level 4 segmented blood image, so we extract all the indices for which the level 4 segmented blood vessel image is non-zero,  $(\bar{n}, \bar{m}, \bar{l})$ . The corresponding indices for corresponding voxels in the level 3 tomogram can be obtained as  $(\bar{n} + i, \bar{m} + j, \bar{l} + k)$ , where  $i, j$  and  $k$  can be either 0 or 1.

We obtain a unique list of the values occupying positions in the level 3 tomogram by the indexes found as above. We create a zero array of the same size as the level 3 tomogram block and set all voxels

that was occupied by a value in the unique list to be 1, i.e. we label these as blood, and write this to the `.hfd5` file.

We iteratively load each of the eight level 3 image blocks, overriding the previous block each time to relieve the RAM.

Now, for each of the eight level 3 image blocks, there are eight level 2 image blocks, for a total of 64 level 2 image blocks, that require connection to the level 3 segmented blood vessel image in the same manner as the connection of the level 3 image blocks to the level 4 segmented blood vessel image. This is where we turn to the recursive algorithm.

In order to set up a recursive algorithm, a formal relation between the layers that constitute its solution is required. For our case, the relation between an image on level  $n$  and level  $n - 1$  is that the image on level  $n$  has half as many voxels on each axis and thus has eight times fewer voxels. We already set up our initial solution to consider the eight corresponding blocks at a level higher on the pyramid.


We require, as input, information about the position of the algorithm on the recursive algorithm tree, i.e. which layer and which node of the layer is currently being processed. For our case, this information is provided by the block notation we introduced earlier, the down-sampling factor indicating the layer and the position of the top, leftmost, front voxel of the block indicating the node.

In short, we have all the components to construct our recursive algorithm for multi-resolution blood vessel segmentation. We use the recursive algorithm to perform multi-resolution matching to obtain the largest connected component in the full-resolution tomogram. In Fig. 8.3, we visualise a slice of the resulting segmented blood vessel network image at each resolution. We see that as the resolution increases, the segmentation at the very edges of the image becomes poor.

This could be due to the blocks at higher resolutions having the same number of voxels as blocks at lower resolutions, leading the higher resolution blocks to cover less physical space. Thus, a  $8 \times 8 \times 8$  block may have multiple coarser defined blood vessels, which will be connected to the larger blood vessel network if any one of those blood vessel voxels are connected to a blood vessel voxels in its corresponding  $16 \times 16 \times 16$  down-scaled block. However, a full resolution block may contain only a single capillary, which may be connected through vessels that are not in the block, but are not disconnected from the blood vessel network because it in itself is not connected to a blood vessel voxels in its corresponding  $2 \times 2 \times 2$  down-scaled block. Introducing an overlap between blocks could remedy this effect.

In Figs. 8.4, 8.5 and 8.6, respectively, we visualise the  $16 \times 16 \times 16$  down-scaled and  $4 \times 4 \times 4$  down-scaled segmented blood images in `vedo`. We see that many more capillaries have been connected to the blood vessel network through the recursive multi-resolution matching algorithm.


The full-resolution matched blood vessel network image of size  $(3279 \times 3480 \times 3480)$  is too large to visualise, as the maximum volume representable by `vedo` is  $(2048 \times 2048 \times 2048)$ . Instead, we visu-




**Figure 8.3:** Visualisation of a slice of the resulting segmented blood vessel network image at each resolution. Note that the slices are the equivalent of each other in different resolutions, as  $l = 72$  for the 16x down-scaled tomogram is the same slice as  $l = 1160$  for the full resolution tomogram.

alise a  $(1500 \times 1600 \times 200)$  section of the full resolution segmented blood image in vedo.


**Figure 8.4:** Visualisation of 16x down-scaled segmented blood image in vedo.



**Figure 8.5:** 3D Visualisation of 4x down-scaled segmented blood image in vedo.



**Figure 8.6:** A  $(1500 \times 1600 \times 200)$  section of the full resolution segmented blood image visualised in 3D in vedo. The full full-resolution matched blood vessel network image of size  $(3279 \times 3480 \times 3480)$  is too large to visualise, as the maximum volume representable by vedo is  $(2048 \times 2048 \times 2048)$ .




## 9 Osteocyte Segmentation

Last but not least, the small but mighty osteocytes. The osteocytes are also difficult to distinguish from all other soft tissue by their voxel values alone, but they have a characteristic shape and volume that sets them apart.

As shown in Fig. 9.1, the osteocytes are ellipsoidal cells, typically  $7\mu\text{m}$  deep,  $7\mu\text{m}$  wide and  $15\mu\text{m}$  long, connected to other cells in the bone mineral through dendrites, appendages that facilitate communication between cells, and canaliculi, small cellular passageways responsible for distributing nutrition.

**Figure 9.1:** Diagram of an osteocyte [47].



In Appendix K, we discuss the segmentation of the osteocytes by their characteristic shape, through shape-based morphological operations such as the Hit-or-Miss and TopHat operations. We ultimately find that these methods are flawed, as the osteocytes are oriented in all directions within the bone, and the structural elements of morphological operations cannot account for orientation.

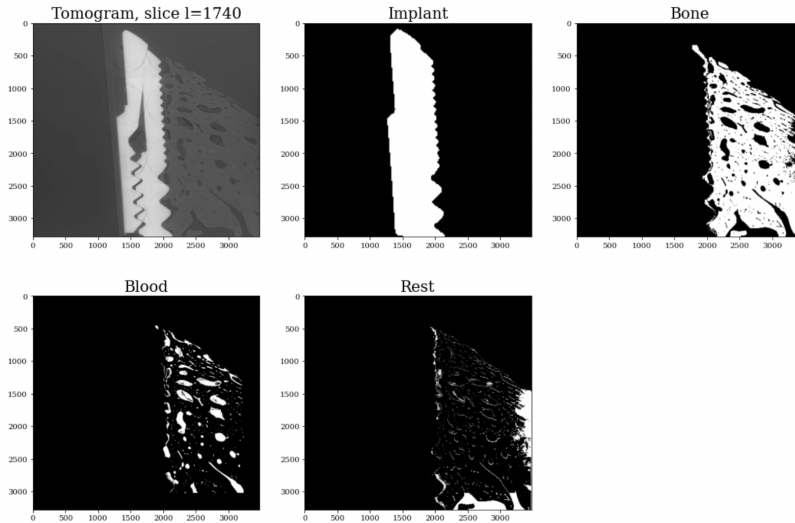
The osteocytes are characterised not only by a certain shape but also by a certain volume, so we may attempt to segment the osteocytes through their volume alone. As mentioned, an osteocyte typically measures about  $(7 \cdot 7 \cdot 15) \mu\text{m}^3 = 735\mu\text{m}^3$  in volume, but they can vary in diameter between  $5$  and  $20 \mu\text{m}$  [3]. We can perform a connected component analysis on the remaining material in the tomogram and filter out any connected components with a volume larger or smaller than the expected volume of an osteocyte. We expect the largest osteocyte to be  $20 \mu\text{m}$  in diameter, with a width and depth of  $10 \mu\text{m}$  to realise its ellipsoidal shape, for a volume of  $(10 \cdot 10 \cdot 20)\mu\text{m}^3 = 2000\mu\text{m}^3$ , and the smallest osteocyte to be

$(2.5 \cdot 2.5 \cdot 5) \mu\text{m}^3 = 31.25 \mu\text{m}^3$ . We filter out any connected components with volumes that fall outside this range.

This volumetric approach to osteocyte segmentation was first conceptualised and naively implemented for a small, 2D sample in [39]. However, the fully automatic implementation for the full resolution 3D tomogram is entirely distinct and unique to this thesis.

### 9.1 Implementation and Results

In Fig. 9.2, we show the segmentation progress so far. We show a slice  $l = 1740$  of the full resolution tomogram, segmented implant, bone mineral, blood vessel network, and the remaining material in the tomogram, where the implant and bone mineral have been resized from 16x down-scaled to full resolution. We will now attempt to segment the osteocytes from the remaining material in the tomogram.



**Figure 9.2:** A slice  $l = 1740$  of the full resolution tomogram (top left), segmented implant mask (top centre), bone mineral mask (top right), blood vessel network mask (bottom left), and the remaining material in the tomogram (bottom right).

For the osteocyte segmentation, we require the full-resolution tomograms as well. However, the osteocytes are not under the condition of being connected to a greater connected component network, so there is no need for multi-resolution analysis in this case. We can simply load one full-resolution block at a time, process this, save the results to a .hdf5 file, and override it with the next block.


For each of the 4096 full-resolution tomogram blocks, which now only contain osteocytes and unresolvable material, we perform a connected component analysis to obtain all connected objects in the block. The voxels that belong to each individual object is now identified by a label specific to that object. For each unique label obtained, we sum the number of voxels occupied by this label. This represents the object volume per voxel volume. The voxel volume is  $((1.875 \cdot s) \times (1.875 \cdot s) \times (1.875 \cdot s)) \mu\text{m} = 6.592s^3 \mu\text{m}^3$ , where  $s$  is the down-scaling factor. The voxel volume for the full-resolution tomogram is thus The voxel volume is given  $(1.875 \times 1.875 \times 1.875) \mu\text{m}^3 =$

$6.592\mu\text{m}^3$ . We multiply the osteocyte voxel count by the voxel volume to obtain the osteocyte volume for each osteocyte. We then filter out any connected components with a volume outside the expected volume range of osteocytes to obtain the segmented osteocytes.

Although the segmentation of the osteocytes require the full resolution data, it can be difficult to evaluate the solution in this resolution. To provide confidence in our solution, we will build a solution that works for any resolution, which can then use on lower resolution data to confirm that the solution works, and then the algorithm can be run on the full-resolution data.

In Fig. 9.3, we show the results of using the osteocyte segmentation method on a  $(50 \times 50 \times 50)$  block within the old bone volume of the 8x down-scaled data. At this resolution, the voxel volume is  $((1.875 \cdot 8) \times (1.875 \cdot 8) \times (1.875 \cdot 8))\mu\text{m} = 3375\mu\text{m}^3$ , and thus we should expect to see no osteocytes at this resolution, as even the largest osteocytes will have been lost to aliasing. Performing a connected component analysis on this result, we can count the number of unique labels in the connected component image to obtain the number of osteocytes in the segmented osteocyte image. We do indeed find that no osteocytes have been segmented at this resolution.

**Figure 9.3:** A slice of a  $(50 \times 50 \times 50)$  block of the 8x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). No osteocytes are identified for this block at this resolution.




In Fig. 9.4, we show the results of using the osteocyte segmentation method on the equivalent  $(100 \times 100 \times 100)$  block of the 4x down-scaled data. At this resolution, the voxel volume is  $((1.875 \cdot 4) \times (1.875 \cdot 4) \times (1.875 \cdot 4))\mu\text{m} \approx 422\mu\text{m}^3$ , and we can expect to see some osteocytes at this resolution, even though the smaller osteocytes will have been lost to aliasing. The connected component label tallying reveals that we in fact find 4612 osteocytes for this block at this resolution.

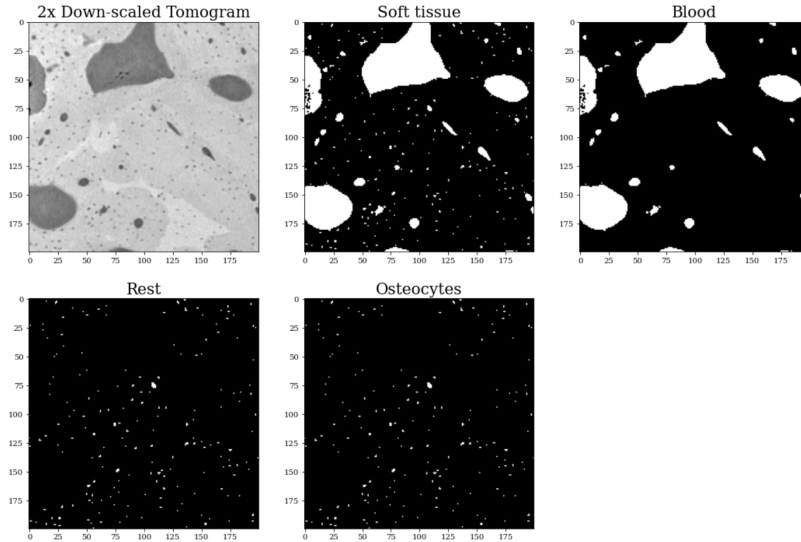
Running the osteocyte finding algorithm on the equivalent  $(200 \times 200 \times 200)$  block of the 2x down-scaled data returns 23,445 osteocytes, as shown in Fig. 9.5. The voxel volume at this resolution is



**Figure 9.4:** A slice of a  $(100 \times 100 \times 100)$  block of the 4x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 4612 osteocytes are identified for this block at this resolution.

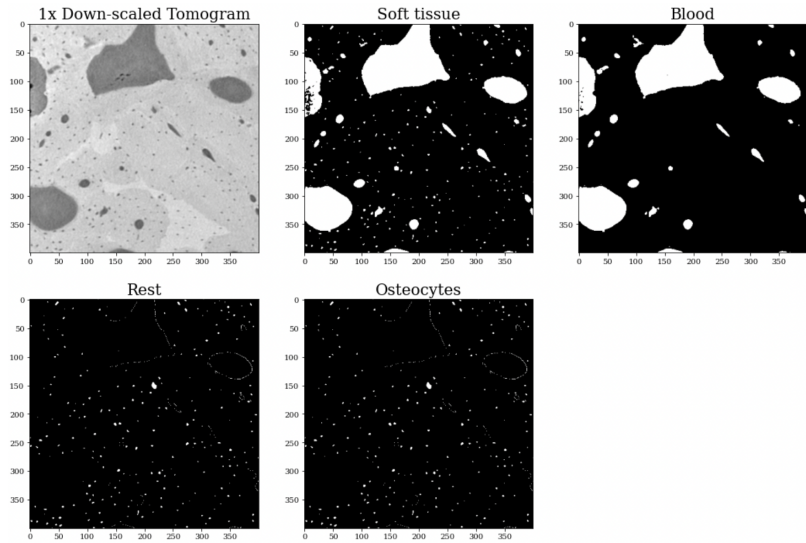


$(1.875 \cdot 2) \times (1.875 \cdot 2) \times (1.875 \cdot 2) \approx 52.73 \mu\text{m}^3$ , and many of the osteocytes should be identifiable at this resolution. Finally, we run the osteocyte finding algorithm on the equivalent  $(400 \times 400 \times 400)$  block of the full-resolution data, as shown in Fig. 9.6. With a voxel volume of  $(1.875 \times 1.875 \times 1.875) \mu\text{m}^3 \approx 6 \mu\text{m}^3$ , all the osteocytes should be identifiable at this resolution. We find 104,295 osteocytes for this block at the full resolution.




**Figure 9.5:** A slice of a  $(200 \times 200 \times 200)$  block of the 2x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 23,445 osteocytes are identified for this block at this resolution.

We now run the osteocyte finding algorithm block-wise on the entirety of the full-resolution data. In Fig. 9.7, we show a slice  $l = 1740$  of the full resolution segmented osteocyte image overlaid on a slice  $l = 1740$  of the full resolution tomogram. As the osteocytes are very small and thus quite difficult to see, we have reduced the opacity of the tomogram to  $\alpha = 0.1$ . The full-resolution data is too large to visualise in *vedo*, so we visualise a  $(1500 \times 1600 \times 200)$  section




**Figure 9.6:** A slice of a ( $400 \times 400 \times 400$ ) block of the full-resolution tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 104,295 osteocytes are identified for this block at this resolution.

of the full resolution segmented osteocyte image in Fig. 9.8.



**Figure 9.7:** A slice  $l = 1740$  of the full resolution segmented osteocyte mask (left) and of the full resolution segmented osteocyte mask overlaid on a slice  $l = 1740$  of the full resolution tomogram (right). The osteocytes are shown in red.

**Figure 9.8:** 3D visualisation of a  $(1500 \times 1600 \times 200)$  section of the full resolution segmented osteocyte image.



## 10 *Improved Bone Health Assessment*

We have reached the final step in our objective to answer the problem statement of whether it is possible to fully automate the segmentation of materials, the quantification of healthy bone, and the evaluation of bone regeneration health in bone samples imaged through high resolution SR $\mu$ -CT.

We have developed a method for automatic segmentation of the implant, bone mineral, air/resin, blood, and osteocytes. We have also obtained a measurement for the amount of regenerated bone by conducting a bone mineral density test from the segmentation of the bone mineral alone, but we are not quite satisfied.

As mentioned, the goal is to evaluate the health of the regenerated bone from the bone regeneration experiments, which cannot be accomplished through the bone mineral density test, as the quantification of the amount of regenerated bone is not equivalent with the quantification of the health of the regenerated bone.

Instead, we propose a method that is approaching the bone mineral density test in simplicity and speed, especially when we have already segmented the materials of interest, but can take the geometric distribution of the blood vessels and osteocytes into consideration. This method involves quantifying the distance from bone mineral to nearest blood vessel, and the distance from bone mineral to nearest osteocyte to ensure the bone has adequate blood flow and nutrient transport and a sensible osteocyte network embedded throughout.

For this, we must be able to describe the position of each bone voxel by its distance to the nearest blood voxel and nearest osteocyte voxel. For this purpose, we remove the segmented material of interest from the tomogram, leaving zero-valued voxels in the tomogram in its place, and use a distance transform to obtain a mapping where each voxel value represents the distance to the nearest voxel of interest. We can then apply a mask of the bone mineral to the distance transform to obtain the distance from each bone voxel to the nearest voxel of interest, blood or osteocyte.


A version of this bone health evaluation metric was first proposed and naively implemented for a small, 2D sample in [39]. However, the fully automatic implementation for the full resolution 3D tomogram is entirely distinct and unique to this thesis.

### 10.1 Distance Transform

A distance transform labels each voxel with the distance to the nearest zero-valued voxel. In Fig. 10.1, we show a simple binary image and the output of subjecting the image to a distance transform. The measure we use to define distance is Euclidean distance, where the distance between two points is given by the length of the shortest straight line that can be drawn between them, given, in three dimensions, for a point  $a$  with coordinates  $(a_x, a_y, a_z)$  and a point  $b$  with coordinates  $(b_x, b_y, b_z)$ , by

$$\text{Euclidean}(a, b) = \sqrt{(b_x - a_x)^2 + (b_y - a_y)^2 + (b_z - a_z)^2}. \quad (10.1)$$

In the example in Fig. 10.1, for the sake of simplicity, the pixel length is 1 and unit-less. When working with the tomograms, the distance transform algorithm will retain the convention of a voxel length as 1, but we will multiply the result with the true voxel width,  $1.875\mu\text{m}$ , to obtain the true distance transform values.



**Figure 10.1:** Binary image (left) and distance transform of binary image (right).

### 10.2 Implementation and Results

We are working with the full resolution data, i.e. data of a size where we are unable to hold all the data in RAM at once, and unable to visualise the full results in 3D. We will therefore once again process the full-resolution data in blocks, and visualise only a part of the final solution in 3D.

To provide confidence in our solution, we will again build a solution that works for any resolution, which can then use on lower resolution data to confirm that the solution works, and subsequently run on the full-resolution data.

We remove the osteocytes from the tomogram by multiplying an inverted mask of the osteocytes with the tomogram, leaving zero-valued voxels in place of the osteocytes. We can then obtain the distance from any voxel in the tomogram to the nearest osteocytes by taking a distance transform of this. We multiply the distance transform with the binary bone mineral mask to obtain the osteocyte distance transform within the bone mineral. We perform the same analysis using an inverted mask of the blood vessels to obtain the blood vessel distance transform within the bone mineral.

To quantify the bone health, we wish to obtain a binary mask of only the healthy bone. For this, we threshold the distance transform within the bone mineral to obtain only the parts of the bone mineral

that have a distance shorter to the maximum healthy distance to an osteocyte and a distance shorter than the maximum healthy distance to a blood vessel. [48] gives the normal distance between osteocyte centres as  $40\mu\text{m}$ . [49] gives the the farthest an osteocyte would be from an Haversian canal<sup>1</sup> as  $172\mu\text{m}$ . We can reasonably take this number as the farthest a bone voxel would be from a Haversian canal, and thus a blood vessel, as well.

<sup>1</sup> The microscopic tubes through which blood vessels travel through bone.


We use these numbers as the thresholds to obtain only the distance transform within the bone volume with values below these thresholds for the osteocyte distance transform and blood vessel distance transform, respectively. For the osteocyte distance transform, we require that the maximum distance to an osteocyte *or* a blood vessel is  $40\mu\text{m}$ . We include the blood vessels in this criteria with the rationale that bone voxels near large blood vessels are healthy bone, but could be far from the nearest osteocyte due to the large blood volume. For the blood vessel distance transform, we require that the maximum distance to a blood vessel is  $172\mu\text{m}$ . For bone to be classified as healthy, it must fulfil both criteria given above. Binarising the result gives us the a mask of the healthy bone mineral. We can now take the sum of projection profile of of the healthy bone mineral mask and divide this with the sum projection profile of of the total bone mineral mask to obtain a percentage measure of the quantity of healthy bone mineral.

We examine the old and regenerated bone separately by using the screw thread threshold to separate the two. Separating the tomograms into the regenerated and old bone regions, respectively, the full resolution tomogram is of shapes  $(2064 \times 3480 \times 3480)$  and  $(1215 \times 3480 \times 3480)$  for the regenerated and old bone regions, respectively. This means the separated 16x down-scaled tomogram is of shape  $(129 \times 217 \times 217)$  and  $(75 \times 217 \times 217)$  for the regenerated and old bone regions, respectively, rather than the original  $(204 \times 217 \times 217)$  for the whole bone sample. If we use the old method of separation into blocks by splitting the data along each axis by  $2^{8/s}$ , where  $s$  is the down-scaling factor, on the regenerated and old bone, respectively, will now yield twice the number of blocks for each resolution. However, we do not need to split the data into quite so small blocks. Thus, instead of splitting the data along each axis by  $2^{8/s}$ , we will continue to split the  $m$  and  $l$  axis by  $2^{8/s}$ , but the  $n$  axis for the regenerated bone region will be split by  $\frac{2}{3}2^{8/s}$  and for the old bone by  $\frac{1}{3}2^{8/s}$ , where we in both cases round up to nearest integer, to achieve blocks of size approximately  $(204 \times 217 \times 217)$  for both the regenerated and old bone regions for each resolution.


At 4x down-scaled resolution, the smallest osteocytes and blood vessels will be have been lost in this resolution due to aliasing, so this will not provide an accurate quantification of the bone health, but it will provide confidence that the solution works in that with fewer blocks, we can examine the solution progress at each step. We will thus test our solution on lower resolution data before running the code on the full resolution data. At 4x down-scaled resolution,

the new method of block splitting described above yields 48 blocks of size  $(172 \times 217 \times 217)$  for the regenerated bone, and 32 blocks of size  $(151 \times 217 \times 217)$  for the old bone, for a total of 80 rather than the 128 that would be produced by the old method of separation into blocks. We run the code for the 4x down-scaled data, as the voxel length at this resolution is  $(1.875 \cdot 4)\mu\text{m} = 7.5\mu\text{m}$ , for a voxel volume of  $(7.5 \cdot 7.5 \cdot 7.5)\mu\text{m}^3 \approx 422\mu\text{m}^3$ , meaning some osteocytes will be discernible at this resolution.

In Fig. 10.2, we show a  $(200 \times 200 \times 200)$  block of the 4x down-scaled tomogram from the old bone volume, and its corresponding bone mineral mask, blood vessel mask, and osteocyte mask. In Fig. 10.3, we show the results of using the bone health quantification on the block. For this block, at the 4x down-scaled resolution, only 53.46% of the bone is classified as healthy. This low percentage is due to the resolution at 4x down-scaled being too poor for proper segmentation of the osteocytes.



**Figure 10.2:** A  $(200 \times 200 \times 200)$  block of the 4x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right). Very few osteocytes are segmented at this resolution, as most are lost to aliasing.




**Figure 10.3:**  $(200 \times 200 \times 200)$  block of the 4x down-scaled tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed, where brighter voxels are closer to osteocytes (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 53.46% of the bone is classified as healthy for this block at this resolution.

Running the code on all 80 4x down-scaled data blocks, we obtain a measure of 41.15% healthy bone for the old bone and 52.77% for the regenerated bone. We see that the regenerated bone is seemingly classified as more healthy than the old bone, where we expect all the old bone to be healthy, but at this resolution we are not able to make any accurate conclusions on the health of the bone. As the resolution increases, more osteocytes and capillaries will be identifiable, and the

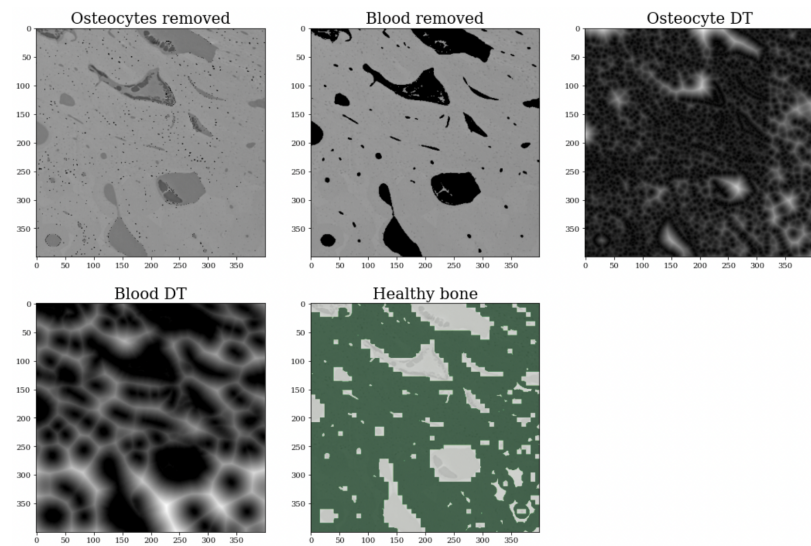
distance to nearest osteocyte and blood vessel will decrease, leading to an increase in the percentage of bone classified as healthy bone.

We run the code for the 2x down-scaled data, first by testing the code on the  $(400 \times 400 \times 400)$  block of the 2x down-scaled data equivalent to the  $(200 \times 200 \times 200)$  block of the 4x down-scaled data in Fig. 10.2. In Fig. 10.4, We show the tomogram, bone mineral mask, blood vessel mask, and osteocyte mask for this block. In Fig. 10.3, we show the results of using the bone health quantification on the block. For this block, 96.27% of the bone mineral is classified as healthy at this resolution. We see that the proportion of bone classified as healthy has seemingly increased. This makes good sense, as many more osteocytes and capillaries are identifiable at this resolution level, and thus the bone to osteocyte and bone to blood vessel distances becomes smaller.

**Figure 10.4:** A  $(400 \times 400 \times 400)$  block of the 2x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right).



**Figure 10.5:**  $(400 \times 400 \times 400)$  block of the 2x down-scaled tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 96.27% of the bone is classified as healthy for this block at this resolution.




Running the code on the entirety of the 2x down-scaled data as well, we obtain a measure of 96.18% healthy bone for the old bone and 97.41% for the regenerated bone. As we saw for the blocks we tested the algorithm on, the proportion of bone classified as healthy has seemingly increased, as more osteocytes and capillaries are identifiable at the higher resolution level, decreasing the bone to osteocyte and bone to blood vessel distances. The voxel length at the 2x down-scaled resolution has a voxel volume of  $(7.5 \cdot 7.5 \cdot 7.5) \mu\text{m}^3 \approx 52.73 \mu\text{m}^3$ . With the typical osteocyte measuring  $(7 \cdot 7 \cdot 15) \mu\text{m}^3 = 735 \mu\text{m}^3$  in

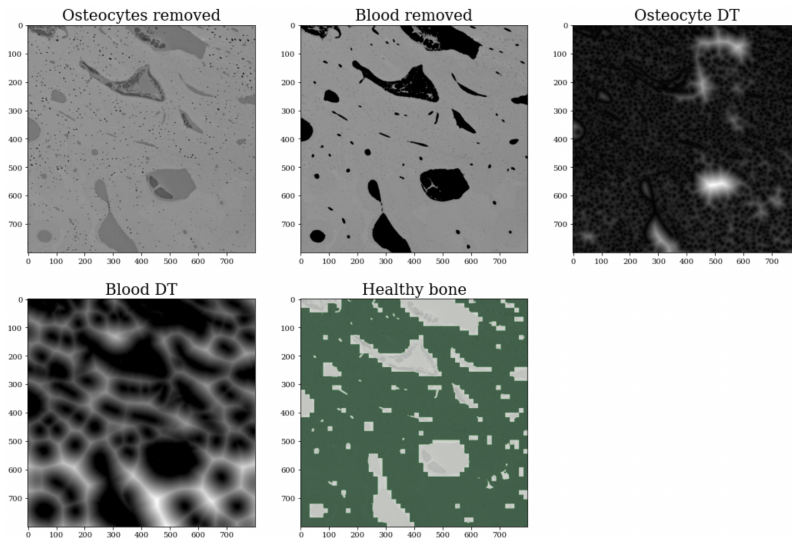


volume, most osteocytes will be discernible at this resolution.

Finally, we run the code for the full-resolution data. First, we test the code on a  $(800 \times 800 \times 800)$  block of the full-resolution data, the tomogram, bone mineral mask, blood vessel mask, and osteocyte mask of which we have shown in Fig.10.6. At this resolution, with a voxel length of  $1.875\mu\text{m}$  and voxel volume of  $6.59\mu\text{m}^3$ , even the smallest capillaries, with diameters of  $\sim 3\mu\text{m}$ , and smallest osteocytes, with volumes of  $\sim 31\mu\text{m}^3$ , will be discernible. Running the bone mineral health quantification algorithm on this block of the full-resolution data returns a healthy bone percentage of 99.06%. This relatively high percentage of healthy bone is consistent with our expectations, as the block is from the old bone volume and we expect the old bone to be healthy bone.



**Figure 10.6:** A  $(800 \times 800 \times 800)$  block of the 4x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right).



**Figure 10.7:**  $(800 \times 800 \times 800)$  block of the full resolution tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 99.06% of the bone is classified as healthy for this block at full resolution.

We run the code on the entirety of the full resolution data, obtaining a measure of 99.69% healthy bone for the old bone and 99.85% for the regenerated bone. We see that the regenerated bone is consistently quantified to have a higher percentage of healthy bone than the old bone. This could be due to the suggested bone health evaluation metric favouring porous bone. In Chapter 7, we suggested that the regenerated bone was more porous than the old bone, as the regenerated bone displayed a 12.6% lower bone mineral density. The bone health evaluation metric proposed in this chapter uses the distance to soft tissue voxels to characterise bone health. As more porous bone

contains more soft tissue, it stands to reason that the regenerated bone is quantified as more healthy by this metric. For future work, a bone health evaluation metric that considers the distribution of bone mineral in tandem with the distribution of soft tissue is sought.

Finally, the evaluation should be run on all of the 35 tomograms produced in [5] and [6] to substantiate this quantitative model of healthy bone, as we expect the model to quantify all the old bone as healthy.

## 11 *Discussion and Future Work*

Image processing at the segmentation and labelling level is a trial and error industry. Throughout the thesis we have made choices for methods to solve the sub-problems of the image segmentation problem. However, alternative solution methods could have yielded similar or better results, and improvements can certainly be made to all results obtained in this thesis. Below, we discuss some possible alternative solution methods and improvements.

For segmentation of the implant, in lieu of the bounding box, we could have used connected component analysis to ensure that high-value voxels further out in the tomogram were not misclassified as implant. The implant is by far the largest connected component in the resulting image from simple threshold-based segmentation, so the implant could have been extracted as the largest connected component similarly to the 16x down-scaled blood voxels.

For segmentation of the bone mineral, the segmentation is affected by uneven lighting due to the X-ray intensity diminishing with penetration depth. Thus, the voxel intensity values in the center of the tomogram will be higher. Another form of uneven lighting arises from light bleeding from the bright implant. A way to address this unevenness is to make Gaussian mixture model distributions region by region as a function of distance from the implant. We can remove the implant from the tomogram and perform a distance transform to obtain a map of distances  $d$  from every voxel in the tomogram to the implant. We can then take equidistant regions  $d = 1, 2, \dots, n$  and create distributions for each region. Both these improvements would require automating the choice of  $k$  for the K-means algorithm, as the number of materials present in a given region is not constant.

During pre-processing for the blood vessel segmentation, we used simple thresholding for separation between the bone mineral and soft tissue in the bone volume, as we found that the resin and air surrounding the bone volume imparted too much noise for the GMM analysis to provide reliable segmentation of the soft tissue. We could have separated the two by multiplying the bone volume mask with the inverse of the segmented bone mineral mask, but we chose separation through thresholding in preparation for future improvements. In future work, the current method for segmenting the bone mineral in Chapter 6 should be used only for obtaining the bone volume mask, and the Gaussian Mixture Model analysis should be repeated for the bone volume alone, eliminating the noise from the surround-

ing resin and air. This will provide a better segmentation of the bone mineral, and a better separation between bone mineral and soft tissue.

In Chapter 8, we also discussed the segmentation of blood vessels at the very edges of the image becoming poor as the resolution increases, which could be remedied through the introduction of an overlap between blocks.

Further, for segmentation of the blood vessels, we used multi-resolution analysis to go from low resolution to high resolution. However, we can further obtain a better segmentation for each resolution level by using multi-resolution analysis to go from high resolution to low resolution. In lower resolutions, many artefacts arise due to aliasing, such as objects blurring together or becoming thicker. We can remove these artifacts by requiring that a voxel in a lower resolution image is only a blood vessel if and only if one of its equivalent 8 higher resolution voxels has been classified as a blood voxel.

For segmentation of the osteocytes, we could look to ensuring the ellipsoidal shape of the osteocytes. We could construct a covariance matrix for each potential osteocyte, as we did in Chapter 4, the three largest eigenvalues of which are proportional to the height, width and depth of the potential osteocyte. We could then impose the condition that for a potential osteocyte to be an osteocyte, the largest eigenvalue must be roughly twice as large as the width and depth, which must be similar in size. However, this is a time-consuming computation, and alternative solutions should be considered.

We saw that the old bone was consistently quantified as less healthy than the regenerated bone. In Chapter 10, we speculated that this could be due to the suggested bone health evaluation metric favouring porous bone. For future work, a bone health evaluation metric that considers the distribution of bone mineral in tandem with the distribution of soft tissue is therefore sought.

The bulk of the run time of the segmentation algorithm is spent on the segmentation of the blood vessels, osteocytes, and the final bone health assessment, where the full resolution data is in play. The full segmentation algorithm takes  $\sim 16$  hours to run for one tomogram, where the final bone health assessment accounts for  $\sim 10$  of the hours. This is not an insignificant run time, and a more effective implementation is sought in the future.

For further future work, improvements are to be made to the segmentation via the methods detailed in the discussion above.

The fully automatic segmentation and evaluation method should be run on all of the 35 tomograms produced in [5] and [6] to evaluate the quality of regenerated bone for each of the 5 regeneration methods. Running the method on all tomograms will also substantiate the quantitative model of healthy bone brought forward in this thesis, as we expect an accurate model to quantify all the old bone as healthy.

Finally, while synchrotron radiation answered the high resolution and image quality requirement for imaging of fine structures in bone, it is not a scalable solution, as synchrotron radiation is ex-

pensive to produce and synchrotron radiation sources are few. The large cost associated with imaging samples through synchrotron radiation and the scarcity of synchrotron radiation facilities ultimately greatly limit the amount of research that can be carried out. To circumvent the need for synchrotron radiation resolution, a future supervised machine learning model can be trained on the high resolution tomograms. When the fully automatic labelling initiated in this thesis is satisfactory, the high resolution labels can be used as ground truth, and low resolution test images can be simulated from the high-resolution images using down-sampling and simulated distortion. With the latter, we are able to simulate specific types of noise, e.g. using Salt-and-Pepper noise to simulate detector sensor component failures or modelling X-ray effects in cheaper equipment, such as beam hardening. The trained model can then be used to classify voxels into their material classes for lower quality tomograms, which could conceivably be produced with hospital X-ray setups, enabling more research, which in turn enables better treatment for patients.

## 12 Conclusion

In this thesis, we fully automated the segmentation of a ( $3279 \times 3480 \times 3480$ ) high resolution SR $\mu$ -CT tomogram into its constituent materials, (I) implant, (II) bone mineral, (III) blood vessels, (IV) osteocytes and (V) air/resin. We also created an automatic quantitative model of healthy bone, and an automatic evaluation of regenerated bone health.

The fully automatic segmentation and evaluation methods were tested on one 3D tomogram of size ( $3279 \times 3480 \times 3480$ ), but the resulting code was developed to run on any of the 35 tomograms produced in [5] and [6].

The segmentation of the implant was done using a combination of voxel-based thresholding, morphological closing, the construction of a bounding box, principal component analysis, and Procrustes analysis. First, the implant was crudely segmented through voxel-based thresholding. To ensure that voxels not belonging to the implant had not been misclassified due to uneven illumination, a bounding box was drawn to identify a region of interest to which the segmentation was restricted. The bounding box was drawn using the three largest eigenvalues and their associated eigenvectors, found through a principal component analysis. Morphological closing was used for pre-processing of the implant to ensure the correct distribution of voxels for the construction of the covariance matrix from which the eigenvalues and eigenvectors were found.

The implant was then separated into its components, a head, micro threads and macro threads. For this, the projection profile of the implant was obtained in order to analyse the screw threads as waves. To minimise noise, the gradient of the projection profile was used in further analysis. A peak finding algorithm was constructed to identify the peaks and troughs of the gradient of the projection profile, from which the wavelengths and amplitudes of the projection profile were found. The micro threads were then identified by their similarly small amplitudes, and the macro threads by their similarly large wavelengths. This similarity was defined on the basis of the central limit theorem. The orientation of the implant was subsequently obtained from the knowledge that the macro threads were to be under the micro threads.

The bone mineral was segmented using a Gaussian mixture model to model the density of voxel intensities in the tomogram as probability density distributions. A K-means clustering algorithm was used

to provide a minimisation algorithm with starting guesses for the Gaussian component distribution parameters. The bone mineral was segmented as voxels with higher than a 95% probability of being a bone voxel.

An intermediate bone health assessment was conducted at this juncture, as a bone mineral density test requires only the bone mineral. For this, the bone mineral volume was evaluated against the total bone volume. The total bone volume was obtained by performing morphological closing on a binarised bone mineral mask. We evaluated the old bone and the regenerated bone separately, with their separation given by the threshold between the micro and the macro threads found during the separation of the implant into its constituents. The sum of the projection profile of the bone mineral mask gave the volume of bone mineral and the sum of the projection profile of the total bone mask gave the total volume of bone. The ratio between the two gave the bone mineral percentage. For the old bone, we found that the bone volume consisted of 70.3% bone mineral. For the regenerated bone, we found that the bone volume consisted of 57.7% bone mineral. The bone mineral content was then obtained in units of  $\text{g}/\text{cm}^3$ . The bone mineral density of the old bone was found to be  $1.2997\text{g}/\text{cm}^3$ , and the bone mineral density of the regenerated bone was found to be  $1.0666\text{g}/\text{cm}^3$ . The resulting  $t$ -scores were 1.5 and  $-0.6$ , respectively, which were within the bounds of  $-1$  to  $2.5$ , and the old and regenerated bone were both classified as healthy.

The blood vessels were segmented using largest connected component analysis and multi-resolution analysis through a recursive algorithm. The full resolution data was too large to be processed at once, and the separation of the data into blocks risked the severance of capillaries, leading the largest connected component analysis to misclassify these. Therefore, a recursive multi-resolution analysis was constructed, in which the largest connected component was extracted from lower resolution data, which could be processed all at once. The higher resolution data was then separated into blocks, a connected component analysis performed for each block, and each voxel compared to its lower resolution equivalent. If the lower resolution equivalent was a blood voxel, the higher resolution voxel and all the voxels with the same label in the higher resolution connected component analysis were labelled as blood as well.

The osteocytes were identified by their characteristic volume. A connected component analysis was performed on the remaining soft tissue in the tomogram, and the volume of each connected component was calculated. If the volume of a connected component was within the accepted range of osteocyte volumes, it was segmented as an osteocyte.

Finally, an improved bone health assessment was given. The distance from any voxel in the bone mineral volume to an osteocyte was found by removing the osteocytes from the tomogram and performing a distance transform. The distance transform was multiplied by the bone mineral mask, as only the distance to an osteocyte within

the bone mineral was of interest. The distance from any voxel in the bone mineral volume to a blood vessel was found through the same method. The threshold for longest accepted distance to an osteocyte for healthy bone was set at  $40\mu\text{m}$ , and the threshold for longest accepted distance to a blood vessel for healthy bone was set at  $172\mu\text{m}$ . Any bone mineral voxels within the accepted thresholds was segmented as healthy bone. The sum of the projection profile of the healthy bone mask gave the volume of healthy bone mineral and the sum of the projection profile of the total bone mineral mask gave the total volume of bone mineral. The ratio between the two gave the bone health percentage. Once again, we evaluate the old bone and the regenerated bone separately. For the old bone, 99.69% of the bone volume was classified as healthy. For the regenerated bone, 99.85% of the bone volume was classified as healthy.

Future work includes improvements to the segmentation made through methods proposed in the discussion, as well as improvements to the efficiency of the algorithm, and the training of a machine learning model for use on lower quality tomograms producible by hospital X-ray setups, enabling more research, which in turn enables better treatment for patients.



# Bibliography

- [1] United States Public Health Service and Surgeon General of the United States. *Bone Health and Osteoporosis: A Report of the Surgeon General*. University Press of the Pacific, 2004. ISBN: 9781410219275. URL: <https://www.ncbi.nlm.nih.gov/books/NBK45504/>.
- [2] Başaran NÇ Sözen T Özişik L. "An overview and management of osteoporosis." In: *European Journal of Rheumatology* 77 (2016), 46–56. DOI: [10.5152/eurjrheum.2016.048](https://doi.org/10.5152/eurjrheum.2016.048). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5335887/>.
- [3] Fiona Linnea Bach-Gansmo. "Bone Ultrastructure With Focus on Osteocyte Lacunae." In: (2015).
- [4] United States. Public Health Service. Office of the Surgeon General. *Bone Health and Osteoporosis: A Report of the Surgeon General*. Free online access: Pubmed. U.S. Department of Health and Human Services, Public Health Service, Office of the Surgeon General. ISBN: 9781588080240.
- [5] Camilla Albeck Neldam et al. "Application of high resolution synchrotron micro-CT radiation in dental implant osseointegration." In: *Journal of Cranio-Maxillofacial Surgery* 43.5 (2015), pp. 682–687. ISSN: 1010-5182. DOI: <https://doi.org/10.1016/j.jcms.2015.03.012>. URL: <https://www.sciencedirect.com/science/article/pii/S1010518215000669>.
- [6] Camilla Albeck Neldam et al. "Synchrotron radiation CT and histology evaluation of bone-to-implant contact." In: *Journal of Cranio-Maxillofacial Surgery* 45.9 (2017), pp. 1448–1457. ISSN: 1010-5182. DOI: <https://doi.org/10.1016/j.jcms.2017.05.019>. URL: <https://www.sciencedirect.com/science/article/pii/S1010518217301853>.
- [7] Liyue Shen et al. "Novel-view X-ray projection synthesis through geometry-integrated deep learning." In: *Medical Image Analysis* 77 (2022), p. 102372. ISSN: 1361-8415. DOI: <https://doi.org/10.1016/j.media.2022.102372>. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522000251>.
- [8] Michael S. Tehrani. *What's the difference between all the different head scans (Xray, CT, MRI, MRA, PET scan) and what they show in the head*. 2019. URL: <https://sdbif.org/index/wp-content/>

- [uploads/2020/02/Differences - Between - Different - Head - Scan - Types .pdf](#) (visited on 05/02/2022).
- [9] E. Lin and A. Alessio. "What are the basic concepts of temporal, contrast, and spatial resolution in cardiac CT?" In: *Journal of Cardiovascular Computed Tomography* 3.6 (2009), 403–408. DOI: [10.1016/j.jcct.2009.07.003](#).
- [10] Rosenbaum AJ. Roberts TT. "Bone grafts, bone substitutes and orthobiologics." In: *Organogenesis* 8.4 (2012), 114–124. DOI: [10.4161/org.23306](#). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3562252/>.
- [11] A.P. Dhawan. *Medical Image Analysis*. IEEE Press Series on Biomedical Engineering. Wiley, 2011. ISBN: 9780470922897.
- [12] F.M. Khan. *The Physics of Radiation Therapy*. Sans tache. Williams & Wilkins, 1984. ISBN: 9780683045017.
- [13] D.E. Thrall. *Textbook of Veterinary Diagnostic Radiology - E-Book*. Elsevier Health Sciences, 2017. ISBN: 9780323482462.
- [14] W.D. Ehmann, V.D. E, and D.E. Vance. *Radiochemistry and Nuclear Methods of Analysis*. Chemical Analysis: A Series of Monographs on Analytical Chemistry and Its Applications. Wiley, 1991. ISBN: 9780471600763.
- [15] Y. Deshko. *Special Relativity: For Inquiring Minds*. Undergraduate lecture notes in physics. Springer International Publishing, 2022, p. 20. ISBN: 9783030911423.
- [16] Sheridan Mayo et al. "Data-Constrained Microstructure Characterization with Multispectrum X-Ray Micro-CT." In: *Microscopy and microanalysis : the official journal of Microscopy Society of America, Microbeam Analysis Society, Microscopical Society of Canada* 18 (June 2012), pp. 524–30. DOI: [10.1017/S1431927612000323](#).
- [17] DDS Stuart C. White and DDS Michael J. Pharoah. *Oral Radiology: Principles and Interpretation*. Elsevier Health Sciences, 2013, pp. 84–85. ISBN: 9780323096331.
- [18] General Electric Company. *High-Flux Target*. 2017. URL: [http://elas.hu/wp-content/uploads/2020/03/geit-31356\\_en\\_high-flux\\_target.pdf](http://elas.hu/wp-content/uploads/2020/03/geit-31356_en_high-flux_target.pdf) (visited on 05/10/2022).
- [19] D. McMorrow and J. Als-Nielsen. *Elements of Modern X-ray Physics*. Wiley, 2011. ISBN: 9781119970156. URL: <https://books.google.dk/books?id=rlqlboWlTRMC>.
- [20] Jeffrey A. Favorite. "The solid angle (geometry factor) for a spherical surface source and an arbitrary detector aperture." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 813 (2016), pp. 29–35. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2015.12.060>. URL: <https://www.sciencedirect.com/science/article/pii/S016890021501640X>.

- [21] D. M. Mills et al. "Report of the Working Group on Synchrotron Radiation Nomenclature – brightness, spectral brightness or brilliance?" In: *Journal of Synchrotron Radiation* 12.3 (2005), p. 385. DOI: [10.1107/S090904950500796X](https://doi.org/10.1107/S090904950500796X). URL: <https://doi.org/10.1107/S090904950500796X>.
- [22] The European Synchrotron Radiation Facility. *The European Synchrotron Radiation Facility*. URL: <https://www.esrf.fr/home.html> (visited on 05/02/2022).
- [23] Joseph Hornak. "The basics of MRI." In: Jan. 2010. Chap. 11.
- [24] Kazuhiro Matsumoto et al. "Virtual Monochromatic Spectral Imaging with Fast Kilovoltage Switching: Improved Image Quality as Compared with That Obtained with Conventional 120-kVp CT." In: *Radiology* 259 (Feb. 2011), pp. 257–62. DOI: [10.1148/radiol.11100978](https://doi.org/10.1148/radiol.11100978).
- [25] Ioannis Tsougos et al. "Differentiation of glioblastoma multiforme from metastatic brain tumor using proton magnetic resonance spectroscopy, diffusion and perfusion metrics at 3 T." In: *Cancer imaging : the official publication of the International Cancer Imaging Society* 12 (Oct. 2012), pp. 423–36. DOI: [10.1102/1470-7330.2012.0038](https://doi.org/10.1102/1470-7330.2012.0038).
- [26] Suheil Artul. "Ring artefact in multidetector CT." In: *BMJ case reports* 2013 (Dec. 2013). DOI: [10.1136/bcr-2013-201379](https://doi.org/10.1136/bcr-2013-201379).
- [27] L. A. Shepp and B. F. Logan. "The Fourier reconstruction of a head section." In: *IEEE Transactions on Nuclear Science* 21.3 (1974), pp. 21–43. DOI: [10.1109/TNS.1974.6499235](https://doi.org/10.1109/TNS.1974.6499235).
- [28] USC Viterbi School of Engineering. *The USC-SIPI Image Database*. URL: <https://sipi.usc.edu/database/> (visited on 05/02/2022).
- [29] David M. Paganin and Daniele Pelliccia. *Tutorials on X-ray Phase Contrast Imaging: Some Fundamentals and Some Conjectures on Future Developments*. 2019. DOI: [10.48550/ARXIV.1902.00364](https://doi.org/10.48550/ARXIV.1902.00364). URL: <https://arxiv.org/abs/1902.00364>.
- [30] National Cancer Institute U.S. National Institutes of Health. *SEER Training Modules, Classification Structure of Blood Vessels*.
- [31] Celestino Ordóñez et al. "Detection of human vital signs in hazardous environments by means of video magnification." In: *PLOS ONE* 13 (Apr. 2018), e0195290. DOI: [10.1371/journal.pone.0195290](https://doi.org/10.1371/journal.pone.0195290).
- [32] Robert A. Freitas Jr. *Nanomedicine, Volume I: Basic Capabilities*. Landes Bioscience, 1999. ISBN: 1570596808. URL: <http://www.nanomedicine.com/NMI/8.2.1.2.htm>.
- [33] Scali S Craig OE Grandi M Sokol RJ. Cattaneo C DiMartino S. "Determining the human origin of fragments of burnt bone: a comparative study of histological, immunological and DNA techniques." In: *Forensic science international* 102(2-3) (June 1999), pp. 181–91. DOI: [10.1016/s0379-0738\(99\)00059-6](https://doi.org/10.1016/s0379-0738(99)00059-6).

- [34] S.J. Ling, J. Sanny, and W. Moebis. *University Physics Volume 1*. University Physics. Samurai Media Limited, 2017. ISBN: 9789888407606.
- [35] Ahmad Raza Khan et al. "3D structure tensor analysis of light microscopy data for validating diffusion MRI." In: *NeuroImage* 111 (2015), pp. 192–203. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2015.01.061>. URL: <https://www.sciencedirect.com/science/article/pii/S1053811915000889>.
- [36] L. Rocha, L. Velho, and P.C.P. Carvalho. "Image moments-based structuring and tracking of objects." In: *Proceedings. XV Brazilian Symposium on Computer Graphics and Image Processing*. 2002, pp. 99–105. DOI: [10.1109/SIBGRA.2002.1167130](https://doi.org/10.1109/SIBGRA.2002.1167130).
- [37] Sina Khanmohammadi and Chun-An Chou. "A Gaussian mixture model based discretization algorithm for associative classification of medical data." In: *Expert Systems with Applications* 58 (2016), pp. 119–129. ISSN: 0957-4174. DOI: <https://doi.org/10.1016/j.eswa.2016.03.046>. URL: <https://www.sciencedirect.com/science/article/pii/S0957417416301440>.
- [38] Douglas Reynolds. "Gaussian Mixture Models." In: *Encyclopedia of Biometrics*. Ed. by Stan Z. Li and Anil Jain. Boston, MA: Springer US, 2009, pp. 659–663. ISBN: 978-0-387-73003-5. DOI: [10.1007/978-0-387-73003-5\\_196](https://doi.org/10.1007/978-0-387-73003-5_196). URL: [https://doi.org/10.1007/978-0-387-73003-5\\_196](https://doi.org/10.1007/978-0-387-73003-5_196).
- [39] Anna Xiao Tan. "Automatic Analysis of High Resolution Synchrotron Radiation  $\mu$ -CTs." B.S. Thesis. University of Copenhagen, 2020.
- [40] Mary Inaba, Naoki Katoh, and Hiroshi Imai. "Applications of Weighted Voronoi Diagrams and Randomization to Variance-Based k-Clustering: (Extended Abstract)." In: *Proceedings of the Tenth Annual Symposium on Computational Geometry*. SCG '94. Stony Brook, New York, USA: Association for Computing Machinery, 1994, 332–339. ISBN: 0897916484. DOI: [10.1145/177424.178042](https://doi.org/10.1145/177424.178042). URL: <https://doi.org/10.1145/177424.178042>.
- [41] .
- [42] Y.H. An and R.A. Draughn. *Mechanical Testing of Bone and the Bone-Implant Interface*. CRC Press, 1999. ISBN: 9781420073560.
- [43] Paul Miller. "Guidelines for the diagnosis of osteoporosis: T-scores vs fractures." In: *Reviews in endocrine metabolic disorders* 7 (July 2006), pp. 75–89. DOI: [10.1007/s11154-006-9006-0](https://doi.org/10.1007/s11154-006-9006-0).
- [44] Celia Gregson et al. "Friend or foe: High bone mineral density on routine bone density scanning, a review of causes and management." In: *Rheumatology (Oxford, England)* 52 (Feb. 2013). DOI: [10.1093/rheumatology/ket007](https://doi.org/10.1093/rheumatology/ket007).

- [45] Michael P. Whyte. "Misinterpretation of Osteodensitometry With High Bone Density: BMD Z + 2.5 Is Not "Normal"." In: *Journal of Clinical Densitometry* 8.1 (2005), pp. 1–6. ISSN: 1094-6950. DOI: <https://doi.org/10.1385/JCD:8:1:001>. URL: <https://www.sciencedirect.com/science/article/pii/S1094695006602989>.
- [46] Lifeng He et al. "The connected-component labeling problem: A review of state-of-the-art algorithms." In: *Pattern Recognition* 70 (2017), pp. 25–43.
- [47] Nobuhito Nango et al. "Osteocyte-directed bone demineralization along canaliculi." In: *Bone* 84 (2016), pp. 279–288. ISSN: 8756-3282. DOI: <https://doi.org/10.1016/j.bone.2015.12.006>. URL: <https://www.sciencedirect.com/science/article/pii/S8756328215004287>.
- [48] Jake P. Taylor-King et al. "Modeling Osteocyte Network Formation: Healthy and Cancerous Environments." In: *Frontiers in Bioengineering and Biotechnology* 8 (2020). ISSN: 2296-4185. DOI: [10.3389/fbioe.2020.00757](https://doi.org/10.3389/fbioe.2020.00757). URL: <https://www.frontiersin.org/article/10.3389/fbioe.2020.00757>.
- [49] Harold M. Frost. "Human Haversian System Measurements." In: *Henry Ford Hospital Medical Bulletin* 9 (1961), pp. 145–147. URL: <https://scholarlycommons.henryford.com/hfhmedjournal/vol9/iss1/23>.
- [50] B. Jähne. *Digital Image Processing*. EngineeringPro collection. Springer Berlin Heidelberg, 2005. ISBN: 9783540275633. DOI: [10.1007/3-540-27563-0](https://doi.org/10.1007/3-540-27563-0).
- [51] Koukalaka's Math. *Physics in Medicine Lectures Week 2 CT Imaging*. 2012. URL: <https://koukalaka.wordpress.com/2012/01/18/physics-in-medicine-lectures-week-2-ct-imaging/> (visited on 05/10/2022).

## List of Figures

- 1.1 Biopsies from a patient with normal bone health (left) and a patient with osteoporosis (right). The healthy bone is denser and strongly interconnected, while the osteoporotic bone is sparse and demonstrates a weaker structure and thus weaker bone strength [4]. 1
- 2.1 Two X-ray images obtained through projection from two angles [7]. 7
- 2.2 2D X-ray image (left) and 2D slice of 3D CT image (right) of a skull, adapted from [8]. Note that the two are not images of the same patient and are included for a basic illustration of the difference between the two imaging modalities only. 7
- 2.3 Goat mandible defect and titanium dental implant (left) and illustration of implant (right) [5]. The lower half of the implant, with a larger screw thread, is implanted in the mandible, with the upper half of the implant, with a smaller screw thread, standing free to be encompassed by regenerated bone, distinguishable from the original bone through the disparity in screw thread. 9
- 2.4 An illustration of the generation of X-rays in a typical hospital setup. Electrons interact with target atoms in the anode, resulting in the emission of X-ray photons. [11]. 10
- 2.5 An illustration of the interaction which causes the emission of X-ray photons, adapted from [11]. The incident electron interacts with the electron cloud around the nucleus of the target atoms in the anode. The electron is deflected and decelerated, causing the electron to lose energy emitted in the form of X-ray photons. 10
- 2.6 The five basic radiographic densities [13]. Higher density materials will cause higher degrees of attenuation, leading to high density materials appearing brighter. 11
- 2.7 An illustration of an electron interacting with matter through photo-absorption, adapted from [12]. A photon is absorbed and the energy from the photon, and, if the energy of the photon exceeds the binding energy of the electron, an electron is released. Part of the energy of the photon is used to break the binding energy of the electron, which receives the residual energy as kinetic energy. 12

- 2.8 An illustration of an electron interacting with matter through Compton scattering, adapted from [12]. An incident photon collides with an electron whose binding energy is much lower than the energy of the photon, and an electron and a new photon with a lower energy than that of the incident photon is emitted. 13
- 2.9 An illustration of an electron interacting with matter through pair production, adapted from [12]. The incident photon is absorbed and an electron and a positron are emitted. Pair-formation only occurs at very high X-ray energies (1.02MeV), and is thus not observed at the X-ray energy at which the bone samples analysed in this thesis were imaged. 14
- 2.10 An example of X-ray imaging geometry [16]. The X-rays are sent through the sample, which is placed a distance of  $R_1$  from the X-ray source, and the X-rays that exit the sample are recorded by a detector placed a distance of  $R_2$  from the sample. These distances determine the amount the image is magnified with respect to the sample. 15
- 2.11 The difference in sharpness with a focal spot of finite size (left) and an idealised focal spot of negligible size (right)[18]. The smaller the focal spot, the sharper the image. 15
- 2.12 Experimental hall with beamline location markers at the European Synchrotron Radiation Facility [22]. The ID19 beamline used for obtaining the tomograms for this thesis is highlighted in red. 17
- 2.13 Difference in resolution-capability between an X-ray scan made at a hospital (left) and that obtained at the ESRF (right) [22]. 17
- 2.14 Illustration of the Radon transform [11]. 19
- 2.15 The original Shepp-Logan Phantom image [28] (left) and a sinogram of the Shepp-Logan Phantom image, constructed using the radon transformation from the Python library `scikit-image` (right). 19
- 2.16 An illustration of back-projection [11]. 20
- 2.17 The reconstructed Shepp-Logan Phantom image (left) and the reconstruction error (right). 20
- 2.18 Visualisation of slices  $n = 1740$ ,  $m = 1740$  and  $l = 1740$ , respectively, of the tomogram. 21
- 2.19 3D visualisation of segmented implant (left) and segmented implant surface (right). 22
- 3.1 An illustration showing an image pyramid. Level 0 represents the  $(N \times N)$  full resolution image, level 1 the  $(N/2 \times N/2)$  down-scaled image, and so forth, until the pinnacle level  $n$  of size  $(1 \times 1)$ . [31] 24
- 3.2 A multi-resolution representation of a 1D image of shape  $(1 \times 16)$ . The lower resolution representations represent the same physical space, only with fewer voxels. 24

- 3.3 Illustration of down-scaling an image by a factor of 2. The down-scaled image is of size  $(1 \times 1)$  and has an intensity value given by the average of the four pixels in the original resolution image. 25
- 3.4 The histograms of voxel intensities in the full resolution, 16x down-scaled and 32x down-scaled tomograms. The 16x down-scaled histogram appears similarly distributed with the full resolution histogram, while the 32x down-scaled histogram is rougher due to insufficient statistics. 26
- 3.5 Full resolution tomogram (top left), 2x down-scaled tomogram (top centre), 4x down-scaled tomogram (top right), 8x down-scaled tomogram (bottom left), 16x down-scaled tomogram (bottom right). Note that the down-scaled tomograms represent the same imaged physical space, only with fewer voxels. 27
  
- 4.1 Histogram of voxel intensities in the tomogram with titanium implant voxels highlighted by a circle in red.. 28
- 4.2 Histogram of voxel intensities in the tomogram with proposed threshold for voxel-based segmentation of the implant. 29
- 4.3 Example of an oriented object with its oriented minimum bounding box (left) and the same object with a non-oriented bounding box. The area of the oriented bounding box will always be smaller, unless the object is entirely symmetric, in which case the two bounding boxes will be identical. The discrepancy in area will increase as the asymmetry increases. The areas are presented unitlessly, as this is irrelevant for comparison. 30
- 4.4 An illustration of the Procrustes analysis. *A* is translated, scaled and rotated with respect to *B*. 33
- 4.5 4-Neighbourhood Nearest Neighbours (left), and 8-Neighbourhood Nearest Neighbours (right). 34
- 4.6 Original image of size  $(200 \times 200)$  created with `skimage.data.binary_blobs` (top left), disk structuring element of radius (10) (top centre), morphologically eroded image (top right), morphologically dilated image (bottom left), morphologically opened image (bottom left), morphologically closed image (bottom right). 35
- 4.7 Tomogram pre-processing: Tomogram before pre-processing (top left), threshold-segmented implant (top right), binarised threshold-segmented implant (bottom left) and morphologically closed binarised threshold-segmented implant (bottom right). 35
- 4.8 The closed, crudely segmented implant sliced in the *l* direction (left), the projection profile for voxels in  $\bar{u}$  direction (centre), and the normalised cumulative sum of voxels (right). 37
- 4.9 The closed, crudely segmented implant sliced in the *m* direction (left), the projection profile for voxels in  $\bar{v}$  direction (centre), and the normalised cumulative sum of voxels (right). 37
- 4.10 The closed, crudely segmented implant sliced in the *n* direction (left), the projection profile for voxels in  $\bar{w}$  direction (centre), and the normalised cumulative sum of voxels (right). 37



- 4.11 Bounding box mask shown from the three tomogram orientations. 38
- 4.12 The segmented implant shown from three orientations. 38
- 5.1 An example of a propagating wave,  $f(x) = \sin(x)$ , with the wavelength, amplitude, a trough and a peak illustrated. 40
- 5.2 Segmented implant before rotation (left), and segmented implant after rotation (right). 41
- 5.3 Morphologically closed segmented implant (left), and implant screw thread projection profile (right). 42
- 5.4 Screw thread projection profile gradient (left), and peaks and troughs in the screw thread projection profile gradient (right). 42
- 5.5 Waves found through definition of two troughs and one peak (left), and waves found through definition of two peaks and one trough (right). 43
- 5.6 Wavelengths found through definition of two troughs and one peak (left), and through definition of two peaks and one trough (centre), and the mean of the two (right). 43
- 5.7 Amplitudes found through definition of two troughs and one peak (left), and through definition of two peaks and one trough (centre), and the mean of the two (right). 43
- 5.8 Mean amplitude and amplitude threshold (left), and mean wavelength and wavelength threshold (right). 44
- 5.9 Small screw thresholds found through amplitude thresholding shown on projection profile gradient (left), and shown on segmented implant (right). 45
- 5.10 Gradient of projection profile excluding the large screw thread (left), amplitude thresholding (center) and new implant head threshold (right). 45
- 5.11 Small screw thresholds found through amplitude thresholding shown on projection profile gradient with a implant head segmentation fix (left), and shown on segmented implant (right). 45
- 5.12 Implant head (left), small screw thread (center), and large screw thread (right). 46
- 6.1 An illustration of K-means clustering in 1D. Centroids are initiated, data points are assigned to clusters, centroids are updated, data points are re-assigned to clusters, centroids are updated, until the algorithm converges. 49
- 6.2 Histogram of tomogram voxel values (left), and histogram of 5000 values distributed as the original histogram (right). 50
- 6.3 Histogram of voxel intensities with cluster regions found through K-Means Clustering. 51
- 6.4 Histogram of voxel intensities with maxima of each region. 51
- 6.5 Total parts of histogram explained by distribution (left) and individual distributions (right) 52

- 6.6 Probability distributions. The x-axis represents voxel intensity values, and the y-axis represents the probability of belonging to a given material class. For any voxel intensity, the probability of the voxels with the intensity value belonging to a material class is given by the value of the probability distribution of the material class. 52
- 6.7 Tomogram colored using probability distributions, where each voxel is coloured according to the material it is most likely to comprise. 53
- 6.8 Segmented bone mineral mask overlaid on tomogram slice. 53
- 6.9 Improving on the bone mineral segmentation by limiting the bone mineral segmentation to outside the implant. The segmented implant mask for slice  $n = 60$  (left), the original segmented bone mineral mask (centre), and the improved segmented bone mineral mask (right). 53
- 6.10 Improving on the bone mineral segmentation by limiting the bone mineral segmentation to after the first implant voxel. The segmented implant mask for slice  $m = 108$  (left), the once improved segmented bone mineral mask (centre), and the twice improved segmented bone mineral mask (right). 54
- 6.11 Improving on the bone mineral segmentation by limiting the bone mineral segmentation to under the threshold between the implant head and the smaller screw thread. The twice improved segmented bone mineral mask for slice  $l = 108$  and the threshold (left), and the thrice improved segmented bone mineral mask (right). 54
- 6.12 Closed thrice improved bone mask (left), largest connected component of closed thrice improved bone mask (centre), and final improved bone mask (right). 55
- 6.13 The final segmented bone mineral mask for slices  $n = 108$ ,  $m = 108$  and  $l = 108$ , respectively. 55
  
- 7.1 Visualising a slice  $l = 108$  of bone mineral mask (left) and total bone volume mask, obtained through morphological closing of the bone mineral mask (right). 57
- 7.2 Projection profile of bone mineral mask and projection profile of bone volume mask. 57
- 7.3 Projection profiles of bone mineral masks of old and regenerated bone (left), and projection profiles of total bone masks of old and regenerated bone (right). 58
- 7.4 Slice by slice bone mineral density quantification for regenerated bone (left) and old bone (right). 59
- 7.5 Illustration of the shortcomings of the bone mineral density test. 60
  
- 8.1 A 2D example where resizing the image by a factor of 2 results in the voxel intensity of 1 voxel to be received by  $(2 \times 2) = 4$  voxels in the resized image. 63
- 8.2 Image with four discernible shapes (left), and connected component labels of image (right). 64

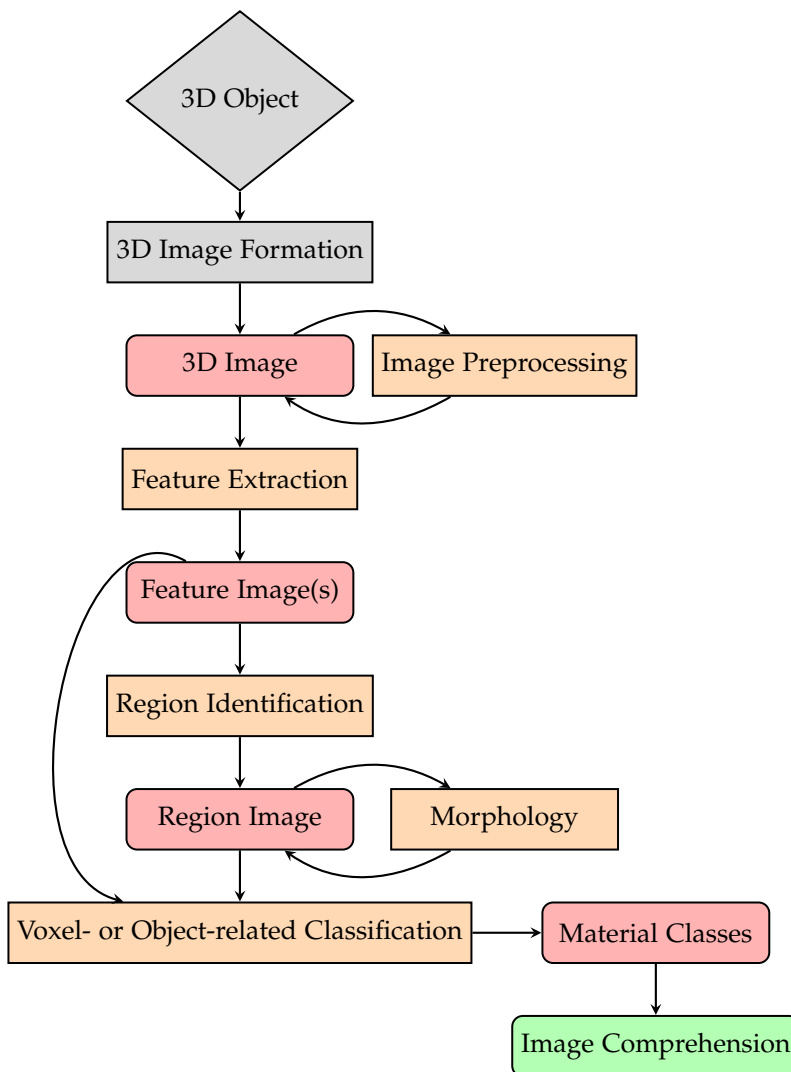
- 8.3 Visualisation of a slice of the resulting segmented blood vessel network image at each resolution. Note that the slices are the equivalent of each other in different resolutions, as  $l = 72$  for the 16x down-scaled tomogram is the same slice as  $l = 1160$  for the full resolution tomogram. 68
- 8.4 Visualisation of 16x down-scaled segmented blood image in vedo. 68
- 8.5 3D Visualisation of 4x down-scaled segmented blood image in vedo. 69
- 8.6 A  $(1500 \times 1600 \times 200)$  section of the full resolution segmented blood image visualised in 3D in vedo. The full full-resolution matched blood vessel network image of size  $(3279 \times 3480 \times 3480)$  is too large to visualise, as the maximum volume representable by vedo is  $(2048 \times 2048 \times 2048)$ . 69
- 9.1 Diagram of an osteocyte [47]. 70
- 9.2 A slice  $l = 1740$  of the full resolution tomogram (top left), segmented implant mask (top centre), bone mineral mask (top right), blood vessel network mask (bottom left), and the remaining material in the tomogram (bottom right). 71
- 9.3 A slice of a  $(50 \times 50 \times 50)$  block of the 8x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). No osteocytes are identified for this block at this resolution. 72
- 9.4 A slice of a  $(100 \times 100 \times 100)$  block of the 4x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 4612 osteocytes are identified for this block at this resolution. 73
- 9.5 A slice of a  $(200 \times 200 \times 200)$  block of the 2x down-scaled tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 23,445 osteocytes are identified for this block at this resolution. 73
- 9.6 A slice of a  $(400 \times 400 \times 400)$  block of the full-resolution tomogram (top left), soft tissue mask (top centre), blood mask (top right), mask of the rest of the material in the tomogram (bottom left) and the mask of identified osteocytes in the block (bottom right). 104,295 osteocytes are identified for this block at this resolution. 74
- 9.7 A slice  $l = 1740$  of the full resolution segmented osteocyte mask (left) and of the full resolution segmented osteocyte mask overlaid on a slice  $l = 1740$  of the full resolution tomogram (right). The osteocytes are shown in red. 74
- 9.8 3D visualisation of a  $(1500 \times 1600 \times 200)$  section of the full resolution segmented osteocyte image. 75

- 10.1 Binary image (left) and distance transform of binary image (right). 79
  - 10.2 A (200 × 200 × 200) block of the 4x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right). Very few osteocytes are segmented at this resolution, as most are lost to aliasing. 79
  - 10.3 (200 × 200 × 200) block of the 4x down-scaled tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed, where brighter voxels are closer to osteocytes (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 53.46% of the bone is is classified as healthy for this block at this resolution. 79
  - 10.4 A (400 × 400 × 400) block of the 2x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right). 80
  - 10.5 (400 × 400 × 400) block of the 2x down-scaled tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 96.27% of the bone is is classified as healthy for this block at this resolution. 80
  - 10.6 A (800 × 800 × 800) block of the 4x down-scaled tomogram (left), bone mineral mask (centre left), blood vessel mask (centre right) and osteocyte mask (right). 81
  - 10.7 (800 × 800 × 800) block of the full resolution tomogram with the osteocytes removed (top left), with the blood vessels removed (top centre), distance transform of tomogram block with osteocytes removed (top right), distance transform of tomogram block with blood vessels removed (bottom left) and results of running the bone health quantification method on the block (bottom right). 99.06% of the bone is is classified as healthy for this block at full resolution. 81
- 1 A hierarchy of image processing tasks. Rectangular blocks represent processes and curved-rectangular blocks represent data or outputs. The blocks in grey represent steps that have been completed before the onset of this thesis. This thesis will traverse the steps from "3D Image" to "Image Comprehension".
  - 2 A square grid representing a 2D image with four pixels.
  - 3 The first three iterations of an iterative reconstruction method [51].
  - 4 Visualisation of slices  $n = 108$ ,  $m = 108$  and  $l = 108$ , respectively, of the segmented implant surface.

- 5 Projection of 3D implant surface on 2D plane through conversion to cylindrical coordinates with constant radius.
- 6 Flow chart of implant segmentation solution. Rectangular blocks represent processes and curved-rectangular blocks represent data or output.
- 7 Image of a centered square with added placed Salt and Pepper noise (left), histogram (center), and thresholded binary image (right). The pepper noise inside the square is misclassified as background and the salt noise outside the square as part of the square.
- 8 Image of an unevenly illuminated page with words [28] (left), image histogram (center) and thresholded binary image. The darker illuminated portion of the page is misclassified as words.
- 9 Image of two toy cars, one light and one dark [28] (left), image histogram (center), and thresholded binary image (right). The algorithm misclassifies the lighter car as the background.
- 10 Example of semantic segmentation (left) and instance segmentation (right). Semantic segmentation concerns itself only with the identification of which object group to which an object belongs, not the identification of individual objects. Instance segmentation assigns a separate label to each individual object within an object group.
- 11 Image with text (left), and the horizontal projection profile of the image (right).
- 12 A function  $f(x) = x^2 - 2x$  and its derivative function  $f'(x) = 2x - 2$ . The derivative function  $f'(x) = 0$  at the minimum of  $f(x)$ .
- 13 Full resolution toy example image (left), 2x down-scaled toy example image (centre left), 4x down-scaled toy example image (centre right), and 8x down-scaled toy example image (right).
- 14 8x down-scaled toy example image (left), and largest connected component in 8x down-scaled toy example image (right).
- 15 The 8x down-scaled image block, 8x largest connected component blood block, the 4x down-scaled image block, and the resulting matched 4x down-scaled blood block for each of the four blocks.
- 16 8x down-scaled image (left), 4x down-scaled image (centre left), 8x down-scaled largest connected component blood image (centre right), and 4x down-scaled matched blood image (right).
- 17 4x down-scaled image (left), 2x down-scaled image (centre left), 4x down-scaled largest connected component blood image (centre right), and 2x down-scaled matched blood image (right).
- 18 2x down-scaled image (left), full resolution image (centre left), 2x down-scaled largest connected component blood image (centre right), and full resolution matched blood image (right).

- 19 Fibonacci sequence tree for  $n = 5$ .
- 20 Hit-or-miss using the first letter as structural element.
- 21 An illustration of a proposed structural element that is able to identify osteocytes regardless of the orientation of these, but is vulnerable to misclassification of objects too large to be osteocytes as osteocytes.

## Appendix A: A Hierarchy of Image Processing Tasks



**Figure 1:** A hierarchy of image processing tasks. Rectangular blocks represent processes and curved-rectangular blocks represent data or outputs. The blocks in grey represent steps that have been completed before the onset of this thesis. This thesis will traverse the steps from "3D Image" to "Image Comprehension".

In Fig. 1, we show a typical path from image formation to image comprehension, where the image processing steps have been ordered in a suggested hierarchy to provide overview. In practice, the order of image processing steps can be switched and steps may be omitted. Below, we will provide first an overview of the general image pro-

cessing steps, and subsequently the solution methods used in this thesis for segmentation of each of the materials in the tomograms. Finally, we outline the method proposed for obtaining an automatic quantitative model of healthy bone and for an automatic evaluation of the regenerated bone against this model.

As shown in Fig. 1, the road from image formation to image comprehension goes through many steps. The first few steps involve the formation of the 3D image [50]. A 3D image of a 3D object is obtained through a 3D image formation module, such as a CT scanner.

From then, the image processing begins. The 3D image is pre-processed, which refers to a group of image processing manipulations applied to the raw image for the purpose of improving or ensuring the performance of subsequent image processing steps. Examples of image pre-processing include the reduction of noise or the number of data samples.

The next step is feature extraction. In feature extraction, the information in the original data is transformed into numerical features that are more easily processed than the original data. An example is using the histogram of an image rather than the image for further processing.

The feature image may then be subject to region identification, a process to restrict the region of interest, e.g. through a bounding box. However, the region image may be the entire image if the objects of interest are existent throughout the image.

The region image is then manipulated through morphology, a class of image processing methods that concern themselves with the form, shape, structure, size, location and orientation of objects.

Finally, the morphologically processed region image is subject to voxel- or object-related classification, which is concerned with the intensity, colour, contrast, gradient, texture or depth of objects, and returns the labels for each voxel.

The image is now separated into material or object classes. Image comprehension goes beyond the recognition of the contents of an image and involves making conclusions based on the segmented image, e.g., whether, when given the distributions of materials throughout the bone volume, the bone is healthy.

At the onset of this thesis, the 3D images have been obtained by [5] and [6]. The goal of this thesis is to traverse the rest of the road from 3D image to reach the image comprehension step.

To reach the image comprehension step for this thesis, we fully automate the segmentation of voxels into one of five material classes present in the tomograms, (1) titanium implant, (2) bone mineral, (3) blood, (4) osteocyte, and (5) air, using a flight of image processing methods, as will be outlined below. Image comprehension is then possible in the form of automatic quantification of healthy bone and evaluation of the health of the regenerated bone.



## Appendix B: Iterative Image Reconstruction Methods

We can think of the image reconstruction problem as being multiple equations required to solve for multiple unknowns.

$x_{1,1}$	$x_{1,2}$
$x_{2,1}$	$x_{2,2}$

**Figure 2:** A square grid representing a 2D image with four pixels.

Say we have an object that fills four pixels, as shown in Fig. 2, and we pass a single beam of intensity  $I_0$  the width of one pixel through the top row of pixels, for which we obtain the read of the amount of radiation incident on the detector,  $I_1$ .

The discretised form of Lambert-Beer's law for inhomogeneous materials is given by

$$I(d) = I_0 e^{-d \sum_{n=1}^N \mu_n}. \quad (1)$$

We can now write an equation with two unknowns,

$$I_1 = I_0 e^{(-d(x_{1,1} + x_{1,2}))} \quad (2)$$

where  $x_{1,1}$  and  $x_{1,2}$  are the linear attenuation values, and  $d$  is the thickness of the pixels. If we take the natural log of both sides,

$$\begin{aligned} \ln(I_1) &= \ln(I_0 \exp(-d(x_{1,1} + x_{1,2}))) \\ &= \ln(I_0) + \ln(\exp(-d(x_{1,1} + x_{1,2}))) \end{aligned} \quad (3)$$

we can rewrite this as

$$\begin{aligned} \ln(I_1) - \ln(I_0) &= \ln(\exp(-d(x_{1,1}, x_{1,2}))) \\ \ln\left(\frac{I_1}{I_0}\right) &= -d(x_{1,1} + x_{1,2}) \\ -\frac{\ln\left(\frac{I_1}{I_0}\right)}{d} &= x_{1,1} + x_{1,2} \end{aligned} \quad (4)$$

The value of all the variables on the left-hand-side are known, as they are the intensity of the incident beam,  $I_0$ , the value read by the

detector,  $I_1$ , and the thickness of the pixels,  $d$ , which is given by the resolution. Thus, to simplify the expression, we group these terms in a constant,  $C_1 = -\frac{\ln(\frac{I_1}{I_0})}{d}$ . The final expression is then,

$$C_1 = x_{1,1} + x_{1,2}. \quad (5)$$

We pass another beam with the same intensity through the bottom row of pixels and obtain another reading of beam attenuation,  $I_2$ , and, thus, another equation,

$$I_2 = I_0 \exp(-d(x_{2,1} + x_{2,2})) \quad (6)$$

which we rewrite in an identical manner as above to obtain

$$C_2 = x_{2,1} + x_{2,2}. \quad (7)$$

Rotating the incident beam 90 degrees, we can obtain a measurement of the intensity of a beam passing through  $x_{1,1}$  and  $x_{2,1}$ , as well as one for  $x_{1,2}$  and  $x_{2,2}$ , yielding the relations

$$C_3 = x_{1,1} + x_{2,1} \quad (8)$$

and

$$C_4 = x_{1,2} + x_{2,2}. \quad (9)$$

We now have four linear equations with four unknowns, which we can solve analytically quite simply.

The trouble is that the example above represented an image of size  $(2 \times 2)$ , which yielded four unknowns to solve for. In real-world applications, this is rarely the case. The tomograms produced for this thesis are of size  $(3279 \times 3480 \times 3480)$ , yielding over a trillion unknowns for a single tomogram. Further, due to noise, arising from anything between movement of the imaged object to beam hardening, the generated equations are not consistent, adding to the difficulty of resolving the linear attenuation values.

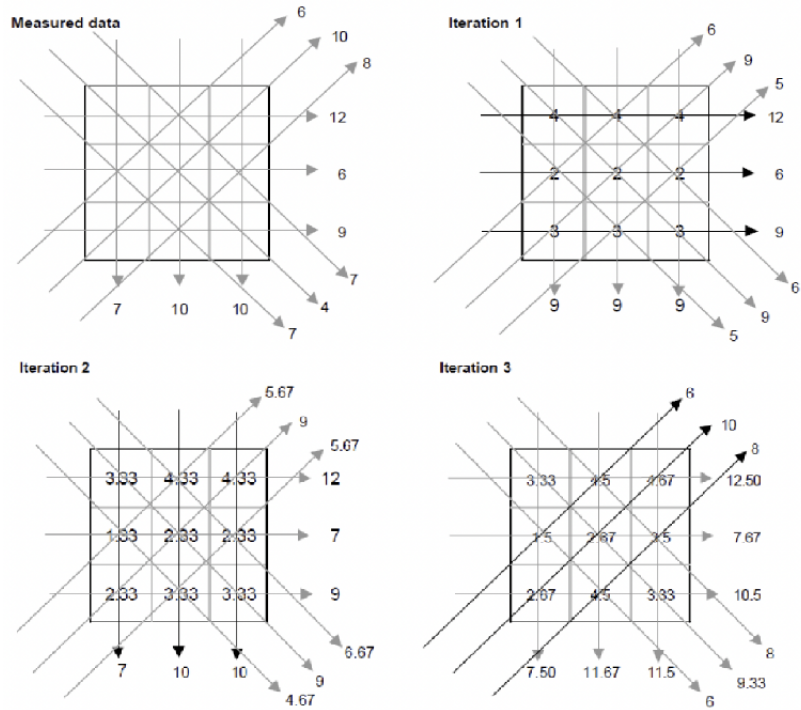
To bypass the need to solve these systems of unknowns exactly, which is not always possible due to noise, iterative reconstruction can be used.

The initial step of iterative reconstruction is making a guess for the true appearance of the imaged object at a given slice, often using a method such as back-projection to conjure the guess.

In Fig. 3, we show an example illustrating the first three iterations of an iterative reconstruction method. On the top left, the values measured by the detector from four different angles are shown. On the top right, an initial guess for the values of the attenuation constants are made using back-projection, where the value measured over each row and column is distributed evenly across that row or column.

The iterative reconstruction method then forward projects the guess, i.e. it computes the projection values for the guessed attenuation constants. It then compares these values with the measured values by obtaining the error as the difference between the measured value and the guessed value.

**Figure 3:** The first three iterations of an iterative reconstruction method [51].



In iteration 2, the algorithm performs a back-projection of this difference, i.e. distributes this evenly across the row or column, and forward projects the guess again. It obtains the difference between the guessed projection values and the measured values, and in iteration 3, it performs a back-projection of this difference, forward projects the guess, and obtains the difference between the guessed projection values and the measured values once again. The algorithm will continue to do this until the ratio is sufficiently small, where "sufficiently small" is user-defined.

## Appendix C: 2D Visualisation of a 3D Surface

For 2D visualisation of a 3D surface, we can project the 3D surface onto a 2D plane. This is accomplished by performing a transformation from Cartesian coordinates  $(x, y, z)$  to cylindrical coordinates  $(r, \theta, z)$ . To achieve a 2D projection, we discard the radial component,  $r$ , as we are unconcerned with the radius when we are visualising a surface. Each point on the projection may thus be made up of several points on the surface that have identical  $(\theta, z)$  coordinates but varying  $r$ .

The conversion from Cartesian coordinates  $(x, y, z)$  to cylindrical coordinates with a constant radius,  $(\theta, z)$  is given by

$$\theta = \tan^{-1} \left( \frac{y}{x} \right) \quad (10)$$


and

$$z = z, \quad (11)$$


where  $z$  simply remains the Cartesian coordinate system  $z$ .

In Fig. 4, we visualise the implant surface via three 2D slices through the 3D volume, as above. In Fig. 5, we show an example of projecting the surface of the implant onto a 2D plane.

**Figure 4:** Visualisation of slices  $n = 108$ ,  $m = 108$  and  $l = 108$ , respectively, of the segmented implant surface.



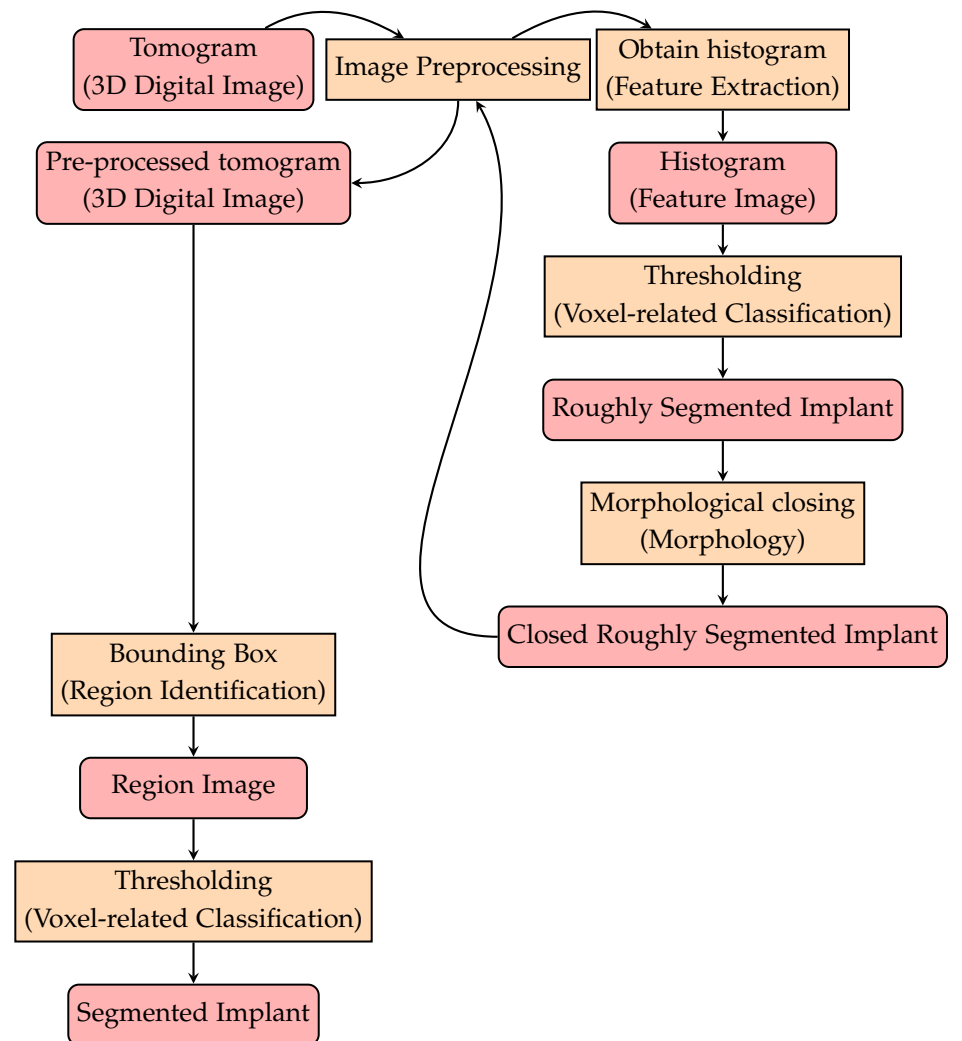
**Figure 5:** Projection of 3D implant surface on 2D plane through conversion to cylindrical coordinates with constant radius.



## Appendix D: Implant Segmentation Solution Flow Chart

In Fig.6, we include a flow chart of the implant segmentation solution.


**Figure 6:** Flow chart of implant segmentation solution. Rectangular blocks represent processes and curved-rectangular blocks represent data or output.



## Appendix E: Illustrating Limitations of Threshold-based Segmentation

Reasons global thresholding becomes problematic include noise, uneven illumination, objects with similar pixel intensity values with the background or several objects with overlapping pixel intensity values, as we will illustrate below.


In Fig. 7, we have used an artificially constructed image of a centered square and added placed Salt and pepper noise <sup>1</sup> to illustrate the effect of noise on global thresholding. This leads the algorithm to classify the pepper noise inside the square as background and the salt noise outside the square as part of the square.



<sup>1</sup> Salt and pepper noise is an artificial noise that simulates sensor element failures in a camera. The effect of the noise is that random pixels are set either to the largest or the smallest value in the pixel range.


**Figure 7:** Image of a centered square with added placed Salt and Pepper noise (left), histogram (center), and thresholded binary image (right). The pepper noise inside the square is misclassified as background and the salt noise outside the square as part of the square.

In Fig. 8, we have used an image of an unevenly illuminated page with words to illustrate the effect of uneven illumination on global thresholding. The human brain is able to automatically correct for the illumination, but the global thresholding classifies the darker illuminated portion of the page as words.



**Figure 8:** Image of an unevenly illuminated page with words [28] (left), image histogram (center) and thresholded binary image. The darker illuminated portion of the page is misclassified as words.

In Fig. 9, we have used an image of two toy cars, one light and one dark, to illustrate the effects of multiple objects of interest with varying intensity values of global thresholding. Say we wish to segment both cars from the background. This is impossible to do with global thresholding, as one car is dark, one bright, and the background falls in the middle of the histogram. The algorithm in this case classifies the lighter car as the background.

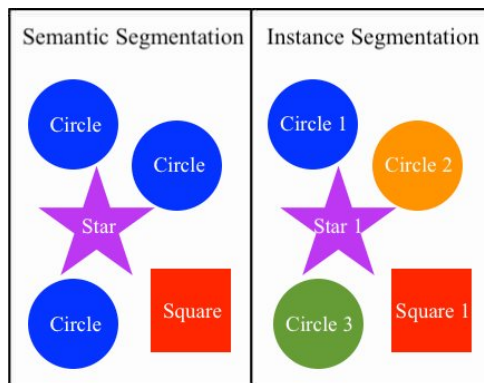


**Figure 9:** Image of two toy cars, one light and one dark [28] (left), image histogram (center), and thresholded binary image (right). The algorithm misclassifies the lighter car as the background.



## Appendix F: The Bounding Box and Instance Segmentation

The Bounding Box is a tool in image segmentation that is most frequently deployed for instance segmentation as opposed to semantic segmentation. In semantic segmentation, we are interested in which object class each voxel in the image belongs to, but we are unconcerned with whether those voxels belong to the same spatial object. Each object class is given one representation, regardless of whether they belong to the same object. In instance segmentation, however, we want to distinguish between object classes, but also between objects in object classes such that each object is given a separate representation. In Fig. 10, we have included an example of these.



**Figure 10:** Example of semantic segmentation (left) and instance segmentation (right). Semantic segmentation concerns itself only with the identification of which object group to which an object belongs, not the identification of individual objects. Instance segmentation assigns a separate label to each individual object within an object group.

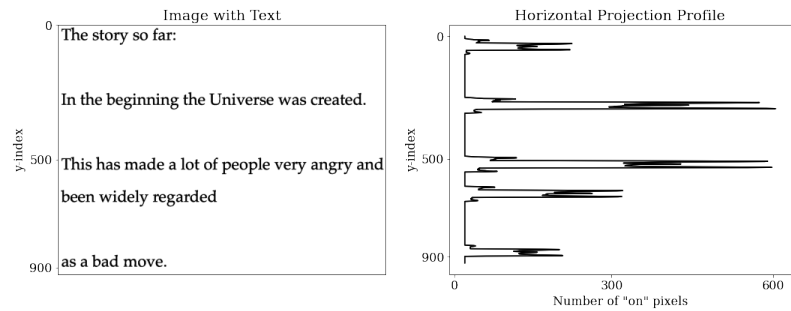
The Bounding Box is used for the purpose of providing separate representation for instance segmentation by bounding the voxels that belong to a specific spatial object. We use this same principle to restrict the region of interest for implant segmentation to the area bounded by the box.

## Appendix G: Illustrating the Projection Profile

In the example in Fig. 11, we show the horizontal projection profile of an image,  $I(n, m)$ , with text. The two-dimensional image has a horizontal and a vertical axis, aligned with the  $n$  and  $m$  axes, respectively. To obtain the horizontal projection profile for the image, we take the sum of pixels for each increment of the vertical axis,

$$proj(m) = \sum_n I(n, m) \quad (12)$$


**Figure 11:** Image with text (left), and the horizontal projection profile of the image (right).



The projection profile is commonly used for text processing, as we have done in the example, as the horizontal projection profile is useful for separation of blocks of text into sentences, as the projection profile goes to zero in the gaps between sentences. The vertical projection profile is able to provide the same analysis for separation of individual words for a single sentence.

## Appendix H: Minimisation

Analytically, the minimisation takes the first derivative of the objective function, sets it equal to zero, and solves for the variables, as the gradient of a minimal point on a function must be zero, as shown in Fig. 12. However, in our case, we do not know the function that describes the histogram,  $H(I)$ , the approximation of which is the entire objective of the Gaussian Mixture Model, and we are unable to differentiate an objective function of which we possess only part of.



**Figure 12:** A function  $f(x) = x^2 - 2x$  and its derivative function  $f'(x) = 2x - 2$ . The derivative function  $f'(x) = 0$  at the minimum of  $f(x)$ .

The information we do possess is the output value of  $H(I)$  for any given point within the constraints. Thus, while we cannot gain a full picture of the objective function, we can easily evaluate it at an arbitrary point. Say we make a guess for the variables,  $\alpha$ ,  $\mu$  and  $\sigma$ , and obtain the resulting Gaussian Mixture Model. We can substitute this into the objective function and obtain the value of this for the guess parameters. We can then carry on choosing new variables for  $\alpha$ ,  $\mu$  and  $\sigma$  until we find the values that minimise the objective function. This is the numerical minimization problem.

Naturally, it is inefficient to blindly choose new parameters for each iteration. The objective function will then randomly increase or decrease and the minimum is unlikely to be found and, if somehow identified, it would be impossible to confirm it as the minimal possible value for the objective function. Instead, various methods may be deployed to determine the best successive guesses, the simplest of which is the gradient descent.


The gradient descent takes starting guesses for the parameters, calculates the value of objective function for these parameters, calculates the negative gradient of the objective function at this point, which gives the direction in which the function becomes smaller. For the next iteration, the algorithm takes a step in this direction, computes the value and negative gradient of the objective function at the

new point, and takes another step in the direction of the negative gradient. It continues to do so until it arrives at a point where the value of the objective function is small enough, defined by a stopping criteria chosen by the user. The size of the steps taken by the algorithm at each iteration is defined by the user, and should be sufficiently small as to not bypass the minimum point, but sufficiently large as to limit the convergence time.

The gradient descent is simple to understand and implement, but ineffective. For the minimisation problem presented in this chapter, we choose the BFGS method, a quasi-Newtonian optimisation method. Rather than merely the first order derivatives computed for the gradient descent, Newtonian optimisation methods compute the inverse Hessian, an inverted matrix of second order derivatives, as well, as the information contained within can significantly decrease convergence time by guiding the step size at each iteration. Further, for quasi-Newtonian optimisation methods, the inverse Hessian approximated rather than computed exactly, as this can be computationally costly. The result is a effective, fast converging optimisation method.

## Appendix I: Toy Example for Multi-resolution Blood Vessel Segmentation


In Fig. 13, we have constructed a 2D toy example to test and illustrate the multi-resolution analysis method. At full resolution, the toy example image has a bright cross in the middle to symbolise a large blood vessel, a dim grid of lines to symbolise the smaller capillaries, two bright lines to symbolise large blood vessel connected to the blood vessel network via the smaller capillaries, and two small, bright crosses to symbolise osteocytes. We down-scale this image by  $2x$ ,  $4x$ , and  $8x$ . We thus have an image pyramid of four levels, each with  $2 \times 2 = 4$  times as many pixels as the image on the level above. In the  $8x$  down-scaled image, the capillaries have disappeared due to aliasing, leaving the two lines of large blood vessels unconnected to the rest of the blood vessel network. Thus, in a largest connected component analysis, these are misclassified as not being blood vessels, as shown in Fig. 14. Further, blurring due to aliasing causes the objects to appear much thicker.



**Figure 13:** Full resolution toy example image (left),  $2x$  down-scaled toy example image (centre left),  $4x$  down-scaled toy example image (centre right), and  $8x$  down-scaled toy example image (right).


We now split the  $4x$  down-scaled image in four blocks so that each of the  $4x$  down-scaled image blocks are of the same size as the full  $8x$  down-scaled image. We likewise split the  $8x$  down-scaled largest connected component blood image in four blocks for multi-scale comparison with the  $4x$  down-scaled image blocks. In Fig. 15, we show the  $8x$  down-scaled image block,  $8x$  largest connected component blood block, the  $4x$  down-scaled image block, and the resulting matched  $4x$  down-scaled blood block for each of the four blocks.

The resulting  $4x$  down-scaled matched blood image is shown in Fig. 16. We see that thickening of the main blood vessel due to aliasing is reduced, but the smaller capillaries and the two larger blood vessels it connects have not yet been identified. In Fig. 17, we show the result of performing the same matching of the  $2x$  down-scaled




**Figure 14:** 8x down-scaled toy example image (left), and largest connected component in 8x down-scaled toy example image (right).

image blocks against the matched 4x down-scaled blood. The 2x down-scaled image has four times as many pixels than the 4x down-scaled image, and thus has 16 times fewer pixels than the 8x down-scaled image, resulting in 16 2x down-scaled image blocks for matching to the 4x down-scaled image. In the 2x down-scaled image the capillaries are visible, enabling this iteration of multi-scale resolution matching to identify these as blood, and thus also connect the two brighter blood vessels to the blood vessel network which were previously disconnected. There is still some blurring due to aliasing, but this is removed in the matching of 64 full-resolution image blocks to the 2x down-scaled matched blood image, as shown in Fig. 18. We see that the multi-resolution analysis method is able to identify the largest connected component in a high-resolution image through multi-resolution analysis.




**Figure 15:** The 8x down-scaled image block, 8x largest connected component blood block, the 4x down-scaled image block, and the resulting matched 4x down-scaled blood block for each of the four blocks.


**Figure 16:** 8x down-scaled image (left), 4x down-scaled image (centre left), 8x down-scaled largest connected component blood image (centre right), and 4x down-scaled matched blood image (right).



**Figure 17:** 4x down-scaled image (left), 2x down-scaled image (centre left), 4x down-scaled largest connected component blood image (centre right), and 2x down-scaled matched blood image (right).



**Figure 18:** 2x down-scaled image (left), full resolution image (centre left), 2x down-scaled largest connected component blood image (centre right), and full resolution matched blood image (right).





## Appendix J: Recursive Algorithms and The Fibonacci Sequence

The best known example of recursion is the Fibonacci sequence, in which a given number in the sequence is defined as the sum of the previous two numbers in the sequence,  $f_n = f_{n-1} + f_{n-2}$ , typically starting with a base case of  $f_0 = 0$  and  $f_1 = 1$ , subject to function or custom. In Fig.19, we show the Fibonacci sequence tree, starting at  $n = 5$ . The recursive algorithm for the Fibonacci sequence, the pseudo code for which is included in Algorithm 1, involves a function,  $\text{Fib}(n)$ , which takes an integer value,  $n$ , and returns  $n$  if the input is  $\leq 1$ , and if else it calls itself as  $\text{Fib}(n - 1) + \text{Fib}(n - 2)$  until the base case,  $f_0$  or  $f_1$ , is reached.




Figure 19: Fibonacci sequence tree for  $n = 5$ .

---

### Algorithm 1: Fibonacci Sequence, $\text{Fib}(n)$

---

**Input** Integer  $n$

**Output** The Fibonacci Sequence from the  $n$ th number

**if**  $n \leq 1$  **then**

**return**  $n$

**end**

**else**

**return**  $\text{Fib}(n - 1) + \text{Fib}(n - 2)$

**end**

---

## Appendix K: Osteocyte Segmentation Through Morphology

We already touched upon the image processing treasure trove that is morphological operations, but many other morphological operations exist for an array of applications, as morphology simply concerns itself with the shape of objects. This is a promising place to start the hunt for the osteocytes, as they have a characteristic shape.

One promising morphological operation for the segmentation of the osteocytes is the Hit-or-Miss operation, which extracts objects from an image that are identical to the structural element by sliding over the image and identifying areas of the image where the structural element is fitted exactly to the foreground. We may be able to utilise the Hit-or-Miss operation using a structural element in the shape of an osteocyte for the segmentation of these.

In Fig. 20, we test this method using an image with six rows of letters, where we wish to extract all Xs from the image using the first X in the top left corner as our structural element.

**Figure 20:** Hit-or-miss using the first letter as structural element.



We see that the method is able to identify some letters, it is far from an impressive result, as digits are not identified unless they fall exactly within the structural element, leaving objects that are smaller or oriented differently than the structuring element automatically out of contention.


While modifications can be made to the structural element to mediate some of the shortcomings of the method, the Hit-or-Miss method remains severely limited by its rigid definition, in that it can only identify objects that are exact matches for the structuring element, whereas the osteocytes can vary in shape and orientation,

which the Hit-or-Miss operation cannot account for.

Another attractive morphological operation for the purpose of segmenting osteocytes is the BottomHat operation, which subtracts the original image from the closing of the image. It is used for extracting objects that are smaller than the structural element, as well as darker than their surroundings, as these are the "holes" that will be filled during the dilation phase of the opening and not be restored during the subsequent erosion. We may thus be able to use the BottomHat operation for extracting the osteocytes using a structural element that is larger than the osteocytes, as the osteocytes are darker than their bone tissue surroundings.

The shortcomings of this operation are similar to those of the Hit-or-Miss operation, namely that the structural element cannot account for orientation. The ellipsoidal shape of the osteocytes makes the BottomHat operation vulnerable, as we cannot create an oriented structural element. As illustrated in Fig. 21, the structural element for the BottomHat operation must be larger than the osteocytes, regardless of the orientation of these, but not so large as to identify objects as osteocytes that are far too large in volume to conceivably be osteocytes, a balance that is near impossible to achieve.

The Top-Hat Problem



**Figure 21:** An illustration of a proposed structural element that is able to identify osteocytes regardless of the orientation of these, but is vulnerable to misclassification of objects too large to be osteocytes as osteocytes.