

Department of Construction Sciences

Solid Mechanics

ISRN LUTFD2/TFHF-26/5274-SE(1-60)

Simulating X-ray Computed Tomography: Developing a Virtual Representation and Generating Synthetic Training Data for Deep Learning

Master's Dissertation by

Elias Almqvist & Kajsa Fahlgren

Supervisors:

Eskil Andreasson, Tetra Pak

Erik Bergvall, Tetra Pak

Eoin Walsh, Lund University

Examiner:

Stephen Hall, Lund University

Copyright © 2026 by the Division of Solid Mechanics

Printed by Media-Tryck AB, Lund, Sweden

For information, address:

Division of Solid Mechanics, Lund University, Box 118, SE-221 00 Lund, Sweden

Webpage: www.solid.lth.se

Abstract

This thesis investigates simulating X-ray computed tomography with two primary objectives: developing a virtual representation of an X-ray computed tomography scanner and generating labelled synthetic training data for deep learning-based fibre segmentation. A simulation pipeline to generate X-ray projections based on the Beer-Lambert law was developed using the open-source framework gVirtualXray.

The first part explores developing a virtual representation of a laboratory-based X-ray computed tomography scanner at Tetra Pak. This was achieved by optimising simulation parameters such as filtration and detector response in gVirtualXray to match experimentally acquired projections. This method proved successful in improving the accuracy of the virtual representation, which produced simulated projections with a root mean square error of 0.0123 when compared to the experimental projections. Despite promising results, this method is limited in its generalisability and a more advanced optimisation would be required to take multiple X-ray computed tomography settings into account.

The second objective explores generating labelled synthetic X-ray computed tomography datasets for training convolutional neural networks to segment fibres. For this purpose, a pipeline was designed to generate multiple diverse datasets of synthetic fibre scans by varying simulation parameters. A U-NET model architecture was trained exclusively on this synthetic data and then evaluated using an experimental dataset. The models showed good performance where the best model managed to generate an F1-score of 0.956, successfully demonstrating the use of synthetic data for training convolutional neural networks in segmentation tasks. However, the work is limited to binary 2D image segmentation, and future advancements should focus on multi-class and 3D volumetric fibre segmentation.

Acknowledgements

We want to give a huge thanks to our supervisors for all the incredible support during this project. Eskil Andreasson, who gave us the opportunity to be part of this exciting project and who continuously kept up our enthusiasm and connected us with many helpful contacts in both industry and academia. Eoin Walsh, who through his weekly support made us feel confident in our progress and who consistently gave thoughtful feedback on our work. Erik Bergvall, who has always been ready to answer questions and help us solve any problems we might encounter.

We also want to thank Fredrik Linell, Piotr Zakrzewski, Prianka Das, Alexander Mirholm and Ivan Forman for always welcoming us to the X-ray lab at Tetra Pak. We received great support when performing experiments there, and they did not hesitate to give up some of their time to help us out.

Lastly, we want to thank the Virtual Modelling team at Tetra Pak for welcoming us into the group during the course of this project.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Objectives	2
1.3	Thesis Outline	2
2	Theoretical Background	4
2.1	X-ray Computed Tomography	4
2.1.1	X-ray's Interaction with Matter	5
2.1.2	Generation and Detection of X-rays	7
2.1.3	Acquisition and Reconstruction	10
2.2	Neural Networks	13
3	Simulation Method	16
3.1	Background and Motivation	16
3.2	Implementation	17
3.2.1	Configuration of Virtual Imaging System	17
3.2.2	Acquisition of Projections	20
3.2.3	Reconstruction	20
4	Virtual Representation of an XCT Scanner	21
4.1	Experimental XCT Scanner	21
4.2	Development of Virtual Representation	22
4.2.1	Configuration to Match Experimental Scanner	23
4.2.2	Optimisation of Parameters	24
4.2.3	Evaluation Experiment	25
4.3	Metrics for Evaluation	27
4.4	Results of Virtual Representation	28
4.4.1	Pre-Optimisation	28
4.4.2	Optimisation	30
4.4.3	Post-Optimisation	31
4.5	Discussion	35
4.5.1	Virtual Representation Performance	35
4.5.2	Limitations and Future Work	35

5	Segmentation of Fibres	37
5.1	Experimental Fibre Dataset	37
5.2	Generation of Synthetic Training Data	40
5.2.1	Simulation of X-ray Computed Tomography	40
5.2.2	Ground Truth Generation	43
5.2.3	Modelling of Phase Effects	44
5.2.4	Composition of Training Dataset	44
5.3	Segmentation Model	45
5.3.1	Model Architecture	45
5.3.2	Evaluation Metrics	46
5.4	Results of Segmentation	47
5.4.1	Evaluation of Segmentation Models	47
5.4.2	Comparison to Threshold Segmentation	51
5.5	Discussion	52
5.5.1	Segmentation Models	52
5.5.2	Limitations	53
5.5.3	Future Work	54
6	Conclusions	56
	Bibliography	58
	Appendix	i
A	Optimisation	i

Chapter 1

Introduction

1.1 Motivation

X-ray computed tomography (XCT) is a non-destructive imaging method for visualising the internal structure of objects in three dimensions (3D) with a typical spatial resolution on the low micron scale, making it useful for studying materials' internal microstructure [1]. These capabilities have driven its widespread adoption across both academic research and industrial applications. Tetra Pak, the industrial partner for this thesis project, has seven in-house XCT scanners used for research and development, product evaluation and component inspection. Tetra Pak is a food processing and packaging solutions company with an interest in developing more sustainable packaging materials, using XCT for studying their paper-based materials.

Computer simulations and numerical modelling are used in various scientific fields for replicating and complementing physical experiments. The Virtual Modelling team at Tetra Pak, within which this project is conducted, are interested in performing virtual experiments for improving the understanding of materials and processes. In the context of XCT, a virtual representation of both a scanner and scanned object offers several advantages. Because experimental scans require significant time and resources, such a simulation allows settings to be tested and optimised prior to an experimental scan, while doing so in a timely manner. Simulation of XCT is widely used in research and has been applied to the development of medical simulators [2], educational tools for teaching XCT [3][4], prediction of image quality and artefacts in materials science [5], and the generation of large datasets of realistic images for training machine learning algorithms [6].

Experimental XCT produces enormous amount of data.¹ Due to the complex, densely packed cellulose fibre networks with irregular geometry found in paper-based materials, such as those used at Tetra Pak, it remains challenging to fully analyse the resulting large

¹The Human Organ Atlas Hub at the European Radiation Synchrotron Facility (ESRF) produces datasets that can reach hundreds of gigabytes [7]

datasets. Segmenting and tracking individual cellulose fibres through an XCT volume would enable quantitative understanding of the microstructure of paper-based materials, and enable studies on how it is related to mechanical properties and adsorption of fluids on a macro-scale. Manual segmentation is however far too slow to be practical at the required scale where many thousands of fibres are present in one XCT volume. This makes automated approaches essential, with one possible solution being deep learning segmentation. However, training such deep learning models requires large amounts of data with known ground-truth labels for the objects of interest, which are difficult to obtain. One possible application of a virtual XCT scanner, paired with a virtual representation of cellulose fibre networks, would be to generate synthetic training data with an automatically known ground truth for the fibres.

As well as working closely with Tetra Pak to develop a virtual XCT scanner, this project also contributes to the AI-TOMO project. This is an ongoing collaboration between Lund University, MAX IV, the Research Institutes of Sweden (RISE), Tetra Pak, and Billerud, which aims to develop deep learning methods for the segmentation of paper-based materials. [8]

1.2 Objectives

This thesis explores developing a virtual representation of an XCT scanner at Tetra Pak and subsequent deep learning applications by implementing the open-source framework gVirtualXray (gVXR) [5], which simulates X-ray projections based on the Beer-Lambert law of attenuation. Two main objectives are defined:

1. **Develop a virtual representation of an XCT scanner:** The simulation framework is applied to generate a virtual representation of an experimental XCT scanner, with the purpose of simulating data similar to the experimental setup. The simulated data is compared to experimentally acquired data for evaluation.
2. **Generate labelled synthetic training data for deep learning segmentation of cotton fibres:** The simulation pipeline is extended to simulate labelled XCT datasets of a simple cotton fibre network, which are used to train a convolutional neural network (CNN). The model is evaluated on its ability to segment an experimentally acquired cotton fibre dataset.

1.3 Thesis Outline

The outline of the thesis is as follows:

Chapter 2 provides the theoretical background, covering the principles of XCT and the fundamentals of CNNs. The general simulation method is then introduced in Chapter

3, motivating the choice of implementing gVXR and detailing how the framework is implemented to simulate XCT. This method is then applied in Chapter 4 and Chapter 5, each structured as an individual experimental section with its own method, results, and discussion, corresponding to the two objectives defined in the previous section. Chapter 4 concerns the development of a virtual XCT scanner, while Chapter 5 focuses on the generation of labelled synthetic training data for deep learning segmentation of cotton fibres. Finally, Chapter 6 summarises the conclusions of the work.

Chapter 2

Theoretical Background

This chapter presents the theoretical background of the thesis, providing the necessary foundation for the methods and results. In the first section, XCT is explained, with a focus on laboratory-based XCT scanners utilising a cone-beam geometry. It begins with a general overview of X-ray physics, covering how X-rays interact with matter to obtain image contrast, as well as how they are generated and detected. This is followed by a description of the imaging procedure, explaining the acquisition and reconstruction of the 3D image. The second section covers the foundations of deep learning-based segmentation by introducing neural networks and their underlying principles.

2.1 X-ray Computed Tomography

XCT was initially developed for medical imaging in the 1970s, and is widely used in diagnostics [1]. The principle was later adapted for materials science and industrial applications where its use has expanded rapidly over the last two decades [9]. Synchrotron facilities, such as MAX IV located in Lund, Sweden, are the high end of XCT imaging, achieving both high spatial and temporal resolution. However, this thesis focuses on laboratory-based XCT scanners, which are often more accessible in industry and academia.

X-rays are high-energy photons that can penetrate matter, and XCT utilises this property to acquire several two-dimensional (2D) radiographs of the object from different angular positions, a process called a CT-scan. These projections are processed by a reconstruction algorithm to produce cross-sectional slices of the object that constitute the 3D image of the object's internal structure [10]. The general setup for laboratory-based XCT scanners is visualised in Figure 2.1, consisting of an X-ray source, a sample mounted on a rotating stage, and an X-ray detector recording the projections. The source is an X-ray tube generating a polychromatic energy spectrum in the shape of a cone, which is the most common geometry because it utilises most of the emission from the X-ray tube [10].




Figure 2.1: Schematic figure an XCT setup utilizing a cone beam geometry [10].

2.1.1 X-ray's Interaction with Matter

X-rays are high-energy electromagnetic waves with wavelengths in the region of 10^{-10} m which, due to their high energy, are able to penetrate objects that would otherwise seem opaque at visible wavelengths [11]. As X-rays penetrate an object, complex mechanisms cause the intensity of the X-rays to be reduced exponentially. This intensity reduction is referred to as attenuation and is primarily caused by mechanisms such as the photoelectric effect and scattering [1]. The specific refractive index of an object is dependent on the effects of attenuation but also the phase shift the X-rays undergo as they pass through the object [10]. These relations are used in X-ray scans to obtain attenuation contrast and phase contrast of the scanned sample. Of these, conventional laboratory-based XCT scanners primarily operate in attenuation contrast mode [1].

Attenuation Contrast

The amount an X-ray attenuates within a given object depends on the material dependant attenuation coefficient, μ , which is given by

$$\mu = \left(\frac{\rho_m N_A}{M} \right) \sigma_a, \quad (2.1.1)$$

where ρ_m is the mass density, M is the molar mass, N_A is Avogadro's number, and σ_a is the absorption cross-section per atom [12]. The value of σ_a is experimentally determined and can be found in material science databases. The relation of σ_a to an element's atomic number, Z , and the photon energy of the X-ray, $\mathcal{E}$, can be approximated to

$$\sigma_a \propto Z^4 / \mathcal{E}^3 \quad (2.1.2)$$

for the photoelectric effect [1][12]. In general, the attenuation coefficient is not constant throughout an object, but varies with position, x , reflecting spatial variations in material composition and density. Given an incident X-ray intensity, I_0 , the transmitted intensity, I , after propagation through an object along of thickness, L , can be expressed as

$$I(L) = I_0 e^{-\int_0^L \mu(x) dx}. \quad (2.1.3)$$

This expression is known as the Beer-Lambert law of attenuation [1]. The principle is visualised in Figure 2.2.




Figure 2.2: Schematic image demonstrating the principle of Beer-Lambert law of attenuation, as presented in Equation 2.1.3.

However, Equation 2.1.3 assumes a monochromatic beam, i.e., all photons have the same energy, which is not true for laboratory-based XCT scanners. Since the attenuation coefficient, μ , depends on energy *via* σ_a (Equation 2.1.1 and 2.1.2), the attenuation through an object varies with a dependency on the incident polychromatic energy spectrum. For a polychromatic X-ray beam with a maximum photon energy of $\mathcal{E}_{max}$, the Beer-Lambert law is instead given by

$$I(L) = \int_0^{\mathcal{E}_{max}} I_0(\mathcal{E}) e^{-\int_0^L \mu(\mathcal{E}, x) dx} d\mathcal{E}. \quad (2.1.4)$$

Phase Contrast

As an X-ray enters an object it undergoes a phase shift that depends on the spatial frequency of the material [13]. The difference in phase of the X-rays emerging from the object to those that have propagated only through air give rise to a detectable phase contrast. For materials with low atomic number in which X-rays attenuate weakly, such as tissue or fibres, the phase shift will be much greater than the attenuation [13]. The resulting phase contrast image will therefore be much better at distinguishing these materials compared to an attenuation contrast image of the same sample [10].

2.1.2 Generation and Detection of X-rays

X-ray Tube and Generated Spectrum

The generation of X-rays occurs within the X-ray tube, which consists of a filament (cathode) and a metal target (anode) connected across a high voltage supply called the tube voltage, as seen in Figure 2.3. The filament, which is usually a metal wire, is heated from an applied current to emit electrons through thermionic emission. The tube voltage gives rise to a potential difference between the filament and the target that will cause the emitted electrons to accelerate towards the target with a kinetic energy equal to the applied tube voltage multiplied by the electron charge [14]. The flow of electrons created is referred to as the tube current. As the target is hit with electrons, X-rays will be generated in one of two ways [15]. The atomic nucleus of the target may cause the incoming electrons to decelerate, and the lost kinetic energy is emitted as X-rays in a phenomenon called 'Bremsstrahlung'. Alternatively, the incoming electron may collide with a target atom's orbital electron and eject it from the orbital. This will result in an electron from a higher energy orbital taking its place and the excess energy will be emitted as a characteristic X-ray with an energy specific to the target's material [15]. The emitted X-ray spectrum is polychromatic, meaning it consists of multiple wavelengths and therefore a distribution of energy levels. The spectrum contains a sharp peak from the characteristic X-rays [10], while the energy width of the spectrum depends on the Bremsstrahlung X-rays and ranges from zero up to the maximum energy of the accelerated electrons, determined by the applied tube voltage [15].

The distribution and amplitude of the generated polychromatic X-ray energy spectrum are mainly affected by the following factors:

- Tube voltage
- Tube current-time product
- Anode angle and material
- Filtration

The tube voltage affects both distribution and amplitude, where increasing the voltage shifts the spectrum to a higher mean energy while increasing the overall intensity [1]. The tube current controls the rate of emitted photons and the exposure time defines the duration of emission. They both scale the amplitude of the spectrum without affecting the distribution of X-ray energies, so their product, the current-time product, determines the total number of emitted X-rays. The anode angle and material both affect the intensity of Bremsstrahlung radiation. Materials with a higher atomic number produce more emissions, which is why tungsten is commonly used as a target material. Increasing the anode angle yields a similar intensity improvement [1].




Figure 2.3: Schematic image of the X-ray tube generating X-rays [15].

To remove unwanted X-ray energies emitted from the polychromatic spectrum, a metal sheet filter can be applied between the X-ray source and the sample. The most common use of filters is to absorb the low energy X-rays that would otherwise mostly be scattered or absorbed by the sample before hitting the detector. Applying an X-ray filter will increase the quality of the radiograph [16] and reduce artefacts in the reconstructed volume [15]. However, this will also decrease the intensity of the desired X-ray energies and increase the required exposure time [16]. Additionally, the X-ray tube housing provides an inherent filter to the X-ray output. The inherent filter of the X-ray tube is conventionally expressed in aluminium thickness equivalent [16] [17].

X-ray Detector

X-ray detectors are designed to capture the spatial distribution of transmitted X-ray intensities. Most XCT systems employ indirect detection schemes, in which X-rays are converted to visible light by a scintillator, and then subsequently detected and converted into digital signals by a photodetector pixel array [10]. The general principle is visualised in Figure 2.4.

When an X-ray is absorbed or scattered in a scintillator, it deposits energy locally in the material by creating electronic excitations. The medium de-excites by emitting photons in a lower energy range around the visible light spectrum. The efficiency of this conversion process depends on the scintillator material and its thickness. A thicker layer absorbs a larger fraction of the incident X-rays, increasing detector efficiency. At the same time, the generated photons spread laterally as they cross the scintillator, leading to optical blur. Therefore, a thin layer results in improved spatial resolution [14]. Additionally, the probability of absorption in the scintillator is strongly dependent on X-ray energies, resulting in an energy-dependent detector response.

The visible light emitted by the scintillator is detected by a photodetector array, commonly a charge-coupled device (CCD) or complementary metal-oxide-semiconductor (CMOS)

photodetector. Most detectors operate as energy integrating detectors, in which each pixel integrates the incident light over the exposure time and converts it to a proportional charge or voltage [18].




Figure 2.4: Schematic image demonstrating the operation of an indirect X-ray detector.

Noise and Calibration

The integrated signal recorded by the detector includes dark current accumulated independently of incident photons. This together with variations in sensitivity across pixels, introduce a fixed noise pattern in the detector. Additionally, the illumination provided by X-ray sources is far from uniform. To compensate for these effects, a calibration step is often required [10]. Flat-field correction uses a dark-field image I_{dark} , recorded with no X-ray exposure to capture the dark current offset, and a flat-field image I_{flat} , recorded with a uniform X-ray beam and no sample present. A corrected projection image I_{corr} is then obtained as

$$I_{\text{corr}} = \frac{I_{\text{raw}} - I_{\text{dark}}}{I_{\text{flat}} - I_{\text{dark}}}, \quad (2.1.5)$$

where I_{raw} is the uncorrected measurement. A systematic detector miscalibration across projections will appear as ring artefacts in the reconstructed volume.

However, in the resulting image there is always noise present due to the randomness of light, so called photon shot noise. The number of photons detected per pixel fluctuates statistically between exposures, which follow a Poisson distribution. If a pixel detects on average N photons, the signal-to-noise ratio (SNR) is:

$$\text{SNR} = \sqrt{N}, \quad (2.1.6)$$

which shows that increasing the number of detected photons reduces the noise level.

Resolution and Geometric Blur

Due to the X-ray beam being emitted in the form of a cone, a geometric magnification occurs, where the sample appears magnified on the detector. The geometric magnification is given by

$$\text{magnification} = \frac{SDD}{SSD}, \quad (2.1.7)$$

where SDD is the source-detector distance and SSD is the source-sample distance [19]. Hence moving the sample closer to the source will make it appear bigger. The observable area, or the field of view (FOV), is also affected by the geometric magnification as FOV is given by

$$\text{FOV} = \frac{\text{detector size}}{\text{magnification}}. \quad (2.1.8)$$

Due to the size limitations of the detector, the FOV is limited as magnification increases [20]. In a digital representation of a volume, the basic unit of each element is called a voxel [10]. The resulting voxel size, affecting the resolution of a CT-scan is given by

$$\text{Voxel size} = \frac{\text{pixel size}}{\text{magnification}}, \quad (2.1.9)$$

where the pixel size is the dimension of the individual pixels in the detector.

The spot where the electron beam is focused on the anode in the X-ray tube, called the focal spot, has a finite size. This will result in X-rays being emitted from multiple ideal point sources, causing geometric blurring in the recorded radiograph, given by

$$\text{geometric blur} = f(\text{magnification} - 1), \quad (2.1.10)$$

where f is the focal spot size. As the blur scales with f , it directly affects the resolution of the XCT scanner. Conventional scanners have a spot size larger than 0.1 mm, microfocus scanners a spot size of a few microns, while nanofocus scanners require a spot size lower than 1 μm [1].

2.1.3 Acquisition and Reconstruction

Acquisition of Projections

To reconstruct an XCT volume, a set of projections first has to be acquired. The sample is rotated while X-ray projections are recorded from different angles over a full 360° rotation at uniform angular intervals [14]. Since each projection encodes the linear attenuation coefficient along a given ray path, combining projections from multiple angles allows a three-dimensional map of $\mu(x, y, z)$ to be constructed. Most reconstruction algorithms operate on the projection value p , defined as the negative logarithm of the normalized

transmitted intensity [1]. Under the idealised conditions of the monochromatic version of the Beer-Lambert law (Equation 2.1.3), this represents the integrated attenuation along the path L from source to pixel,

$$p = -\ln \frac{I}{I_0} = \int_L \mu(x, y, z) d\ell. \quad (2.1.11)$$

Each projection row has an attenuation profile, $p(v, \theta)$, dependent on the vertical detector location, v , and angle, θ . For a simple egg-like object with a "yolk", examples of the attenuation profile for a fixed vertical position v at $\theta = 0$ and $\theta = 90^\circ$ are shown in Figure 2.5.




Figure 2.5: X-ray attenuation profiles of a simple egg-like object at two different projection angles [21].

The theoretical number of projections, N_p , needed to accurately reconstruct an image slice, derived from the Nyquist theorem, is

$$N_p = N_x \frac{\pi}{2} \quad (2.1.12)$$

where N_x is the number of horizontal detector pixels in each projection [10]. The effect of varying the number of projections can be visualised in Figure 2.6 for the simple egg-like object, where the Nyquist limit is 120 projections. While a few projections provide very little information about the scanned object, increasing the number of projections will gradually increase the accuracy of the reconstructed object [21].




Figure 2.6: Resulting vertical slice for different total of back projections at equiangular rotations [21].

Reconstruction Algorithms

A variety of algorithms have been developed for reconstructing XCT projection data into volumetric images, including algebraic reconstruction techniques (ART), iterative reconstruction methods, convolution back-projection approaches, and more recently, deep learning-based reconstruction methods [21]. Among these, filtered back projection (FBP) remains the most widely used [21]. However, the conventional FBP algorithm is limited to XCT systems utilizing parallel-beam geometry. In laboratory-based XCT systems, which often utilize cone-beam geometries, the Feldkamp–Davis–Kress (FDK) algorithm is instead utilized [22]. The FDK algorithm extends the principles of FBP to cone-beam geometries.

In FDK, initial distance-based weighting is applied to weight each detector element at coordinate (u, v) of the projection value $p(u, v, \theta)$ [23][24]. The weighted projection is calculated as

$$p_w(u, v, \theta) = \frac{SDD}{\sqrt{SDD^2 + u^2 + v^2}} p(u, v, \theta). \quad (2.1.13)$$

The weighted projection compensates for intensity fall-off due the divergence of the cone beam and the associated variation in ray path length [23]. The weighting is followed by a spatial filtering process in order to reduce blurring in the final reconstruction [14]. Here $p_w(u, v, \theta)$ undergoes a row wise 1D Fourier transformation $\mathcal{F}$ for each row v in every projection at angle θ [21]. A frequency filter $H(\omega)$ is applied to the Fourier transformed weighted projections before applying an inverse Fourier transform to yield

$$\tilde{p}(u, v, \theta) = \mathcal{F}_z^{-1} (H(\omega) \cdot \mathcal{F}_z(p_w(u, v, \theta))), \quad (2.1.14)$$

where $\tilde{p}(u, v, \theta)$ are the filtered projections [25]. The unfiltered and filtered of the egg-like object can be observed in Figure 2.7. Some of the most common filters are Ramp, Ram-Lak, Shepp-Logan and Hann [26].




Figure 2.7: Comparison of using no filter and using a ramp-filter for the reconstructed egg-like object using 120 projection [21].

Lastly, the filtered and weighted projection rows, $\tilde{p}(v, \theta)$, are at this stage back projected across image space for each angle θ [22]. All back projections are then overlapped to generate a vertical slice. The final attenuation coefficient distribution, $\mu(x, y, z)$, is obtained by integrating over the full range of projection angles.

$$\mu(x, y, z) = \int \tilde{p}(u(x, y, z, \theta), v(x, y, z, \theta), \theta) d\theta. \quad (2.1.15)$$

Here, each angle, θ , is mapped onto a corresponding detector position, (u, v) , and the filtered projection values corresponding to each voxel position, (x, y, z) , is accumulated.

Artefacts

Artefacts in XCT are systematic deviations between the reconstructed voxel values and the true attenuation coefficients of the scanned sample [27]. They arise from violations of assumptions underlying the reconstruction process. In real systems, the measured projections deviate from the line integral model (Equation 2.1.11) due to the polychromatic X-ray spectrum, the presence of scattered radiation, and detector imperfections. Artefacts can appear in various ways in the reconstructed volume, including streaking, shading, and rings.

2.2 Neural Networks

In supervised machine learning, models are trained to connect input data to a corresponding mask [28]. A widely studied problem in supervised learning is segmentation, where a label is predicted for each individual element in the input, for example a pixel in an image [28]. This section will focus on how these problems are solved using artificial neural networks (ANNs) to segment data into different classes.




Figure 2.8: The basic structure of an ANN [29].

ANNs are computational machine learning models inspired by the structure and function of biological neural networks [30]. An ANN typically consists of three main layers, an input layer, a hidden layer and an output layer, as seen in Figure 2.8. Each layer consists of interconnected processing units which are commonly referred to as neurons. Each neuron in the hidden layer receives multiple input signals, x , where each signal is associated with a weight, w . The weighted sum of each signal is computed and passed into the activation function, ϕ , together with a bias, b [30]. The non-linear output, y , of a neuron can be expressed as

$$y = \phi \left(\sum_{i=1}^n w_i x_i + b \right). \quad (2.2.1)$$

An ANN learns by adjusting its weights and biases, through a process called backpropagation. In this process, the network first performs a forward pass to generate predictions and compute a loss value. The loss value measures how accurate a prediction is compared to the ground-truth, where a low loss value shows that the prediction is accurate. The backpropagation algorithm then calculates the gradients of the loss function with respect to each parameter using the chain rule. These gradients are used in gradient-based optimization methods, such as gradient descent, to iteratively update the parameters in a direction that minimizes the loss. Over time, this enables the network to learn meaningful patterns from the data [30].

In segmentation tasks, ANNs operate by mapping input features to output labels at a finer level, performing classification at the level of individual pixels or regions [31]. The ANN learns to associate these feature representations with corresponding class labels and errors are computed using loss functions to update weights and minimize future prediction errors. This learning process allows the network to capture non-linear relationships in data and improve segmentation accuracy compared to traditional methods. However, since standard ANNs are usually fully connected [30], they do not explicitly preserve spatial structure,

which limits their effectiveness for image segmentation tasks and motivates the use of other machine learning architectures [32], such as convolutional neural networks (CNNs).

Unlike standard fully connected ANNs, each neuron in a CNN is connected to a small region of different inputs which preserves spatial information throughout the network [33]. This makes CNNs highly suitable for processing data with a grid structure, such as images [30]. The most basic hidden layers in a CNN are convolutional layers, activation layers and pooling layers [34]. The convolutional layers consist of filters that are able to detect spatial hierarchies [30]. Convolutional layers are often followed by an activation function to introduce non-linearity. The pooling layers reduce the spatial dimensions, further reducing the number of parameters within that activation [34].

Chapter 3

Simulation Method

In this work, the open source framework gVXR was employed to simulate X-ray projections. The simulation pipeline was implemented in Python. As gVXR requires no commercial licence, the source code is accessible and reproducible. This chapter describes the simulation method, detailing the implementation of the framework to simulate XCT. We start by providing a background to simulation of XCT, motivating the choice for employing gVXR. Subsequently, the implementation is described in detail, including the configuration of the virtual XCT scanner through the definition of the X-ray source, sample, and detector, followed by the acquisition of projections and the reconstruction approach.

3.1 Background and Motivation

There are several approaches to simulating XCT, with many commercial and open-source software are available. Monte Carlo (MC) simulations are regarded as the most accurate method for simulating X-ray projections. They can model the underlying physics by tracking a high number of photons as they are emitted from the X-ray source, interact with the sample through absorption and scattering processes, and reach the detector. Geant4 [35] is widely considered the state-of-the-art MC framework, offering comprehensive physical models for photon transport. However, MC methods are computationally heavy, and a high resolution XCT volume with realistic noise levels can take several days to simulate [36]. Deterministic simulations based on the Beer-Lambert law of attenuation, described in Section 2.1.1, can provide a good compromise between speed and accurate simulation. These models neglect scattering noise, although the latter can be approximated in a post-processing step if required.

After reviewing several deterministic simulation frameworks, the open source framework gVXR was selected to simulate X-ray projections. It is optimized for computational speed through a GPU-based implementation of Beer-lambert law of attenuation. This aspect was valued because of the prospect of quickly generating large labelled synthetic datasets for deep learning. The framework has been validated against both MC simulations performed

with Geant4 and experimental projections, demonstrating its capability to produce realistic results [36]. In a comparison to the frameworks ASTRA and TIGRE, gVXR provides greater flexibility in configuring the source, detector, and sample [37]. For example, gVXR supports modelling of a polychromatic X-ray spectrum, while ASTRA and TIGRE are limited to a monochromatic spectrum.

As gVXR enables modelling of a detector response and objects consisting of several materials, this results in the following polychromatic version of the Beer-Lambert law of attenuation, where $\mathbf{I}(x, y)$ is the integrated energy received by detector pixel (x, y) :

$$\mathbf{I}(x, y) = \sum_i \mathbf{R}(E_i) \mathbf{D}(E_i) e^{-\sum_j \mu_j(E_i) \mathbf{d}_j(x, y)}. \quad (3.1.1)$$

The polychromatic spectrum is discretized into energy bins with energies E_i , where $\mathbf{D}(E_i)$ represents the number of photons emitted in each bin. The detector response is described by $\mathbf{R}(E_i)$, which accounts for the energy-dependant efficiency of the scintillator. Each scanned material, indexed j , is characterized by an energy-dependent linear attenuation coefficient, $\mu_j(E_i)$. The path length through material j from the source to pixel (x, y) is denoted $\mathbf{d}_j(x, y)$ [5].

3.2 Implementation

3.2.1 Configuration of Virtual Imaging System

Configuring the virtual XCT scanner is the first step in the simulation pipeline, and defines how the resulting X-ray projections are simulated. It involves defining the X-ray source, sample, X-ray detector and their relative positions. Figure 3.1 visualizes an example of the simulation environment after configuration, where the sample is a cube.




Figure 3.1: Illustration of the gVXR simulation environment, visualising the point source, cone beam, sample, and detector. The sample is composed of two different materials, which is reflected by the intensity distribution on the detector. [36]

X-ray Source

The X-ray source is configured by setting the X-ray tube voltage, tube current-time product, anode angle, and eventual X-ray filter material and thickness. From these parameters, a polychromatic spectrum is modelled from a tungsten-anode X-ray tube, where the validated spectrum-generation toolkit SpekPy is employed [38]. In the simulation, the resulting photon fluence energy distribution at the specified source-detector distance is expressed in units of photons per pixel and discretized into energy bins. In this work 50 bins are used as a trade-off between spectral accuracy and computational speed.

As discussed in Section 2.1.2, inherent filtration occurs in the X-ray source, where materials in the X-ray tube and its housing absorb low-energy photons. As inherent filtration is conventionally expressed in units of aluminium-equivalent thickness, this phenomenon is modelled by adding an aluminium filter with specified thickness to the source. Figure 3.2 illustrates this effect by displaying two spectra generated with identical settings: one without filtration and one with a 0.5 mm aluminium filter.¹




Figure 3.2: Simulated spectra from a tungsten-anode X-ray tube with 60 kV tube voltage, 1 mAs tube current-time product, 12° anode angle, 100 cm source-detector distance and 1 mm² pixel size. The spectra are shown without filtration and with a 0.5 mm aluminium filter.

The X-ray source in this work is modelled as a cone beam with an ideal point source. This is a simplification of real XCT systems, where a finite focal spot introduces geometrical blur, discussed in Section 2.1.2. Accurately modelling this effect would require simulating multiple point sources, which increases computational time significantly. This is possible to implement in gVXR, but was excluded to reduce computational cost.

¹This example is for demonstration purposes; actual simulations may use different settings depending on the experiment.

Sample

In gVXR, samples are represented using triangular meshes, a technique commonly used in computer graphics for describing three-dimensional structures. The sample can consist of one or several triangular meshes. Once the mesh is added to the simulation, its material properties are defined. This involves defining the elements present in the material, the relative mass fraction of each element, and the material density. Based on this provided information, gVXR internally approximates the linear attenuation coefficient for the material. Xraylib [39], a library connected to multiple X-ray databases is used by gVXR to fetch the absorption cross section for each of the elements. These are directly related to the linear attenuation coefficients, described by Equation 2.1.1.

X-ray Detector

The detector is configured by specifying its pixel size and the number of pixels along both the height and width. It is operated in energy-integrating mode, where each pixel records the total energy deposited by incident X-rays over the exposure time. A scintillator material with a specified thickness can be assigned to model a more realistic energy-dependant detector response, as discussed in Section 2.1.2. Then, the detected signal is reduced relative to the incident photon energy to mimic the actual efficiency of the scintillator. Note that while scintillator thickness also governs detector blur in practice, this effect is not captured by the scintillator model in gVXR. The blur can however be approximated by specifying a line spread function (LSF), which describes how the detector spreads the signal from an ideal line source, and is a one-dimensional approximation of the point spread function (PSF). In this work, the LSF is approximated as a Gaussian:

$$\text{LSF}(x, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad (3.2.1)$$

where x is the pixel coordinate relative to the centre and σ is the standard deviation in units of pixels, controlling the degree of blur. If no line spread function is assigned, the detector is idealized using a Dirac PSF leading to no correlation between neighbouring pixels.

Photon shot noise is modelled using a feature available in gVXR. As this noise is proportional to the number of incident photons per pixel, it is directly influenced by the source parameters used to generate the spectrum and detector settings. The number of incident photons is mainly affected by the tube-current time product, the aluminium filter thickness and pixel size. It is important to note that the projection model (Equation 3.1.1) is derived based on a idealized detector without photon shot noise, and that these effects are added as a post-processing step within gVXR, introducing additional computational cost while enhancing realism.

3.2.2 Acquisition of Projections

Once the sample, source and detector is configured, the imaging procedure is carried out. The number of projections is set, and the virtual sample is rotated incrementally while projections are simulated at each position. Before reconstruction, the projections are converted by taking the negative logarithm of the normalized transmitted intensity (Equation 2.1.11).

Flat-field correction can be enabled during acquisition by defining a number of flat-field images. No dark images are acquired since the virtual detector does not contain any dark current, and would just consist of an array of zeros. Each projection is then divided by the average flat field image. In experimental XCT, this procedure calibrates the detector and reduces ring artefacts in the reconstructed volume as discussed in Section 2.1.2. However, in the virtual system the effect is the opposite: since the flat-field images are simulated with photon shot noise, dividing by them imprints a fixed spatial noise pattern onto every projection, which then simulates ring artefacts.

3.2.3 Reconstruction

In this work, the Core Imaging Library (CIL) [40] is used to perform the reconstruction. It is an open-source library that provides many of the popular reconstruction algorithms, such as FDK for cone beam geometry, which is employed in this work. Many reconstruction filters are available, including Ramp, Ram-Lak, Shepp-Logan and Hann. The resulting volume is saved as a Tagged Image File Format (TIFF), which is a stack containing each reconstructed vertical slice.

Chapter 4

Virtual Representation of an XCT Scanner

A realistic virtual representation of an XCT scanner enables virtual experiments to be performed, producing results comparable to those obtained experimentally. Tetra Pak are interested in performing virtual experiments as a preparation for and compliment to experimental scans. Here we develop a virtual representation of one of Tetra Pak’s experimental XCT scanners, applying the simulation method using gVXR, described in Chapter 3. As some simulation parameters are not known or directly measurable, an optimisation approach is used to estimate them by matching simulated data to experimentally acquired data.

By recreating an experimental XCT scanner virtually, it becomes possible to scan the same sample both experimentally and virtually, enabling quantitative comparison of the generated data down to the pixel level. The results reveal the capabilities and limitations of the simulation approach, forming a foundation for its use in future experiments.

4.1 Experimental XCT Scanner

The experimental XCT scanner of this work is a Phoenix Nanotom M (Waygate Technologies) [41] at Tetra Pak’s Lund site, visualised in Figure 4.1. It features a nanofocus X-ray tube, designed for high resolution imaging. The operating principles of laboratory-based XCT scanners are described in the theoretical background (see Section 2.1), and more detailed information on this particular scanner can be found in the manual provided by the manufacturer [41]. The XCT volume reconstruction is performed using the integrated Phoenix datos|x [41], which operates with limited transparency, leaving many processing steps such as calibration, artefact reduction, and details about the reconstruction algorithm unknown.



Figure 4.1: Phoenix Nanotom M (Waygate Technologies) [41].

4.2 Development of Virtual Representation

The general methodology we employ for designing a virtual representation of the Nanotom XCT scanner is summarised in Figure 4.2. First, a real-world sample is scanned to acquire experimental data. From the reconstructed volume, a triangular mesh of the sample is extracted through segmentation, where voxels exceeding a threshold are classified as sample rather than background. This mesh serves as the geometric input to the simulation, where the scan is recreated virtually using identical scanning parameters to those used in the experimental scan. Finally, the simulated and experimental data are compared quantitatively. This experiment primarily focused on comparing projections and not reconstructed volumes.




Figure 4.2: Workflow of the method used to develop a virtual representation of an XCT scanner.

This section first covers how the virtual scanner is configured to match the experimental scanner, explaining both known and unknown simulation parameters. An optimisation method is then described for determining the unknown parameters, using a sample with known composition and density. Finally, the evaluation experiment is presented, where a different sample is used to assess the generalisation of the optimised simulation.

4.2.1 Configuration to Match Experimental Scanner

To recreate an experimental scan virtually, the simulation must be configured to match the experimental scanner as closely as possible. Table 4.1 provides an overview of the adjustable parameters when configuring the virtual scanner, categorised into X-ray source, sample, X-ray detector, and geometry. We describe these parameters in detail in Section 3.2.1. In this experiment detector blur is not modelled. Some parameters are directly known from the experimental setup, and these are marked in the known column. These include the tube voltage, tube current- time product, pixel array and geometry between source, detector and sample which are all set when configuring the experimental scanner. The pixel size of the detector is a known property, which is found in the manual [41]. Sample composition and density are not marked, as these depend on the specific sample being scanned and must be determined or estimated separately.

Category	Parameter	Known
Source	Tube voltage (kV)	Yes
	Tube current-time product (mAs)	Yes
	Anode angle ($^{\circ}$)	No
	Al filter thickness (mm)	No
Detector	Pixel array (width, height)	Yes
	Pixel size (μm)	Yes
	Scintillator material	No
	Scintillator thickness (μm)	No
Sample	Composition	–
	Density (g/cm^3)	–
Geometry	Source-sample distance (cm)	Yes
	Source-detector distance (cm)	Yes

Table 4.1: Adjustable parameters in virtual XCT scanner configuration, and whether they are directly known from the experimental setup.

The anode angle, scintillator material, and scintillator thickness are internal hardware properties not reported by the scanner manufacturer. The aluminium filter thickness is an approximation of the inherent filtering of the source, and is not an actual property, and therefore needs to be estimated. These are unknown properties to the XCT scanner and cannot be measured directly, and must instead be determined through an optimisation procedure, described in the following section.

4.2.2 Optimisation of Parameters

In order to set up an accurate virtual representation of the experimental XCT scanner, an optimisation method was designed. The aim of this method was to find the optimal values for the unknown filter thickness, anode angle, scintillator material and scintillator thickness (see Table 4.1). An aluminium crucible (see Figure 4.3a) was chosen as the reference material for this optimisation method. It was chosen because of its simple structure and because it is made purely out of aluminium. With pure aluminium, applying the correct chemical composition and density in gVXR is straightforward to implement, and is therefore assumed to be a known property in the virtual XCT scanner configuration.

We imaged the aluminium crucible using the Nanotom XCT scanner at Tetra Pak. The settings used for this scan are displayed in Table 4.2. To image the aluminium crucible, it was attached to a holder by its bottom half. Due to only wanting to include pure aluminium in the scan, only half of the aluminium crucible was imaged to avoid parts of the holder being visible. After the scan, the 2000 projections obtained were used to reconstruct the XCT volume. This allowed for the extraction of a 3D mesh of the aluminium crucible, which can be seen in Figure 4.3b. As observed at the bottom of the aluminium crucible mesh, the mesh is missing a small part of its volume. This is due to reconstruction artifacts appearing at the edges of the reconstructed volume, making the bottom of the sample appear less dense than it actually is. Because of this, the intensity values near the bottom of the sample appears close to that of air, making the extraction of the entire mesh volume difficult. Further, both experimental and simulated projections will be cropped to not include the bottom in order to make comparisons fairer. Using gVXR, the mesh was imported into the virtual Nanotom XCT scanner and the material properties were defined as pure aluminium with a density of 2.698 g/cm^3 . The virtual XCT parameters were set as the parameters used in the experimental XCT scan (Table 4.2).




Figure 4.3: (a) The aluminium crucible and (b) the mesh of the aluminium crucible.

Setting	Value
Pixel array (width, height)	(2048, 1000)
Source-detector distance (mm)	500
Source-sample distance (mm)	25
Tube voltage (kV)	60
Current (mA)	0.3
Exposure time (ms)	375
Number of projections	2000

Table 4.2: Experimental parameters of the Nanotom XCT scanner for the scan of the aluminium crucible.

The optimisation algorithm used to find the optimal filter thickness, anode angle, scintillator material and scintillator thickness was based on automatically testing all combinations of these parameters within a reasonable range of values (see Table 4.3). All the scintillator materials available in gVXR were tested in the optimisation, these can be found in Appendix A. For each parameter combination, a projection at 0° was generated and compared to the corresponding experimentally acquired projection, and their similarity evaluated by calculating the selected metrics (see Section 4.3). The set of parameters that gave the best result was saved for a combined score that took all metrics into consideration.

Parameter	Min	Max	Spacing
Filter thickness (mm)	0.1	1.1	0.01
Anode angle ($^\circ$)	6	20	1
Scintillator thickness (μm)	100	1100	5

Table 4.3: Range of values and materials tested for the optimisation of the unknown parameters.

4.2.3 Evaluation Experiment

To evaluate the virtual representation of the Nanotom XCT, a LEGO[®] Minifigure (see Figure 4.4a) was chosen as a simple test object. The aim of this experiment was to test the accuracy of the virtual representation using a different object than used for the optimisation. As previously described, an experimental XCT scan was initially performed to acquire both the mesh of the LEGO Minifigure and experimental data for comparison. The parameters of the experimental XCT scan are displayed in Table 4.4. For this purpose a holder was built using LEGO bricks (see Figure 4.4b) to minimize the number of different materials used.




Figure 4.4: (a) The LEGO Minifigure and (b) the experimental setup of the Nanotom XCT scanner.

Setting	Value
Pixels array (width, height)	(1300, 2048)
Source-detector distance (mm)	400
Source-sample distance (mm)	95
Tube voltage (kV)	60
Current (mA)	0.3
Exposure time (ms)	375
Number of projections	2042

Table 4.4: Experimental parameter of the Nanotom XCT scanner for the scan of the LEGO Minifigure.

Following the experimental scan and reconstruction it became apparent that three different materials are present in the LEGO Minifigure. The hands, the body and the white print on head and torso had given rise to different attenuations of the X-rays, meaning they possess different densities. In this experiment the minifigure was therefore exported as three different meshes. These are shown separate and together in Figure 4.5. These were imported into gVXR and a CT-scan was performed using the same number of projections.

The elemental composition and density of the polymers of the LEGO Minifigure were unknown. To find these, Fourier-transform infrared spectroscopy (FTIR) was performed at Tetra Pak's material analysis lab. The FTIR showed that the body and holder is made of acrylonitrile butadiene styrene (ABS) and that the hands are made of Polyoxymethylene (POM). The print on the torso and head were too small and thin to test using the FTIR tool. Therefore, the print was assumed to be ABS and was modelled as part of the body.

ABS is a polymer with the chemical composition $(C_8H_8)_x(C_4H_6)_y(C_3H_3N)_z$ and a density between 0.9–1.53 g/cm³. The different ratios of the monomers (x, y, z) in ABS gives rise to an uncertainty in both exact elemental composition and density. In this work ABS was modelled with an elemental mass fraction of 87.1% carbon, 8.1% hydrogen and 4.8% nitrogen, and a density of 1.04 g/cm³. This mass fraction and density was chosen for ABS was based on the typical values given by standard grade ABS [42]. POM is a polymer with the chemical composition $(CH_2O)_n$ and a density of 1.41 g/cm³. POM was therefore modelled with an elemental mass fraction of 40% carbon, 6.7% hydrogen and 53.3% oxygen.




Figure 4.5: The mesh of (a) the LEGO Minifigure bricks and body, (b) hands and (c) the print and (d) the meshes in (a), (b) and (c) together.

4.3 Metrics for Evaluation

To evaluate the quality of the simulation, a projection of the simulation, $\mathbf{P}_S$, is compared to the corresponding experimental projection, $\mathbf{P}_M$. Their similarity is evaluated based on two metrics: Root Mean Square Error (RMSE) and Histogram Intersection (HI). The RMSE measures the pixel-wise differences between $\mathbf{P}_M$ and $\mathbf{P}_S$, and it is defined as the square root of the mean squared error. RMSE is always non-negative, and lower values indicate that the projections are similar on a pixel level. It is defined as:

$$\text{RMSE} = \sqrt{\frac{1}{rc} \sum_{i=0}^{r-1} \sum_{j=0}^{c-1} (\mathbf{P}_S(i, j) - \mathbf{P}_M(i, j))^2}, \quad (4.3.1)$$

where r is the number of rows and c is the number of columns of the projections. A histogram of a projection is the distribution of intensity values. The HI metric compares

the histograms H_M and H_S , derived from P_M and P_S , respectively. It is defined as:

$$HI = \sum_{i=1}^n \min(h_S(i), h_M(i)) \quad (4.3.2)$$

where h_M and h_S are the normalized histograms with equal binning n . The score is between 0 and 1, where a higher value indicates more overlap between the histograms. HI is especially useful for capturing differences and similarities in global patterns, including the background noise. However, two projections can have identical histograms yet differ locally, and is therefore complemented by RMSE.

Prior to comparison, both projections are divided by the background intensity to obtain the normalised transmitted intensity. For $\mathbf{P}_S$ this is straightforward, as the background intensity is known directly from the simulation. For $\mathbf{P}_M$ this is estimated by averaging pixel values in rows where no sample is present. Theoretically, flat-field correction (Equation 2.1.5) should already yield normalised projections. However, the experimental projections obtained from Phoenix datos|x are not output in this range, and additional normalisation was therefore required.

4.4 Results of Virtual Representation

4.4.1 Pre-Optimisation

A first virtual projection was acquired before the optimisation, using no filter applied to the source and no scintillator applied to the detector, while anode angle was set as the default (12°). This virtual projection along with the experimental projection can be observed in Figure 4.6, along with each projection's respective histogram plot. It is clear from both the projection comparison and the histogram comparison that the intensity values in the simulated projection are much lower than those in the experimental projection. The mean intensity in the experimental projection is 0.909 and the mean intensity in the simulated projection is 0.819, yielding a HI value of 0.541 and a RMSE of 0.135. The lower intensity values of the simulated projection compared to the experimental projection is also highlighted in the intensity line plots in Figure 4.7. The intensity line plot compares the intensities of the experimental projection (green) and simulated projection (orange) for a given detector row. At every point of the LEGO Minifigure the simulated intensity (orange) is displaced at a much lower intensity value compared to the experimental intensity (green).




Figure 4.6: Experimental and simulated projections of the LEGO Minifigure acquired at $\theta = 0^\circ$, along with corresponding histogram distributions.




Figure 4.7: Intensity line plots for three different detector rows, comparing the experimental (green) and the simulated (orange) intensities, for a projection taken at $\theta = 0^\circ$.

4.4.2 Optimisation

To find the unknown parameters of the Nanotom XCT scanner (see Table 4.1), the optimisation method described in Section 4.2.2 was performed, and the results of this method will be presented in this section. The parameters that resulted in the best combined HI and RMSE score during the optimisation can be found in Table 4.5. The combined score weighted HI and RMSE equally and was calculated as $HI/2 + (1 - RMSE)/2$.

Parameter	Optimal setting
Filter thickness (mm)	0.26
Anode angle ($^\circ$)	12
Scintillator thickness (μm)	525
Scintillator material	CdWO ₄

Table 4.5: The resulting optimal settings for the unknown parameters.

With these settings applied, a virtual XCT scan of the aluminium crucible was performed. A projection and histogram comparison of the simulated projection (orange) and experimental projection (green) are shown in Figure 4.8 and yields a HI value of 0.930 and a RMSE of 0.0118. The average HI value and RMSE for ten different projections taken at even intervals between 0° and 360° were calculated to be 0.923 and 0.0110, respectively. Figure 4.9 shows intensity line plots for three different projections. As can be observed in all three line plots, the intensities at this line in every projection overlap very closely.




Figure 4.8: Experimental and simulated projections of the Al crucible with a histogram comparison, after the optimisation was performed.




Figure 4.9: Intensity line plots for the same row taken at (a) 0° , (b) 72° and (c) 288° .

4.4.3 Post-Optimisation

Applying the parameters found in the optimisation step to the virtual XCT scanner and performing an acquisition generated a new set of simulated projections of the LEGO Minifigure. A resulting simulated projection of the minifigure compared to the experimental projection is shown in Figure 4.10, along with their respective histogram plots. Both projections are acquired at the same angle, $\theta = 0^\circ$. The mean intensity value of the simulated projection is now 0.911, yielding a HI value of 0.873 and a RMSE of 0.0112. The




Figure 4.10: Experimental and simulated projections of the LEGO Minifigure, along with corresponding histogram distributions.

average HI value and RMSE for all 2042 projections acquired is 0.824 and 0.0123 respectively. The line intensity plot for the same detector rows as in Figure 4.7, but now for the optimised projection, is shown in Figure 4.11. Comparing Figure 4.6 to Figure 4.10, as well as Figure 4.7 to Figure 4.11, it is clear that optimising the unknown parameters had positive effect on the accuracy of the virtual representation of the Nanotom XCT scanner.

The most noticeable variation in the histograms in Figure 4.10 can be observed in the peaks at the intensity value approximately at 0.73 (marked by the red vertical line) and at 0.75 (marked by the blue vertical line). Here the small peak of the simulated projection's histogram curve (orange) is shifted to a lower intensity value relative to that of the peak of the experimental projection's histogram curve (green). The intensity line plot seen in Figure 4.11c shows a similar offset in intensity values at the position off the hands. The same values as for the peaks in the histograms are here marked by a red and blue horizontal line. This indicates that the observed offset in the histogram peak can be connected to the incorrect attenuation in the hands of the LEGO Minifigure, where the simulated projections show higher attenuation than the experimental projections. The greatest deviation




Figure 4.11: Intensity line plots for three different detector row positions, comparing the experimental (green) and the simulated (orange) intensities, for a projection taken at $\theta = 0^\circ$.

in intensity can however be seen in the hair of the LEGO Minifigure. This is clear in both the intensity line plot in Figure 4.11a and in the absolute errors seen in Figure 4.12. Figure 4.11a indicates that the attenuation of the hair region in the simulated projection (orange) is underestimated relative to the experimental projection (green). It is also observable in Figure 4.12 that the intensity errors are greater at some angles than others. This would explain why the average HI value is lower and the RMSE is higher compared to the HI value and RMSE of the single projection at $\theta = 0^\circ$.

Due to computational limitations, full-scale reconstruction of the simulated projections was not feasible using CIL. Reconstruction was therefore performed on a reduced dataset, using 100 central detector rows, and the resulting central slice was used for evaluation. An attempt was made to use a common reconstruction method for both simulated and experimental projections, but this was not achieved within the scope of the project. The central slice from both datasets are displayed in Figure 4.13, where the experimental volume was reconstructed using the integrated software of the Nanotom XCT scanner, Phoenix datos|x, and the simulated volume using CIL with the Shepp-Logan filter. No quantitative comparison is performed, as the reconstructed intensity values are arbitrary and dependent on the reconstruction pipeline. However, the slices appear visually similar. The hands appear brighter due to their higher attenuation coefficient (red box), and reconstruction artefacts are visible as dark streaks parallel to the body. The metallic print on the body appear bright in the experimental reconstruction (blue box), but as it was assigned the same density as the body in the simulation, it does not appear bright in the simulated volume.




Figure 4.12: The absolute error of simulated projections of the LEGO Minifigure at (a) 0°, (b) 72°, (c) 144° and (d) 216°. The error is defined as the absolute difference between the experimental and simulated projection intensity values.




Figure 4.13: A reconstructed slice from the (a) experimental projections using `datos|x` software, and (b) simulated projections using CIL. The vertical slice is located at the centre of the reconstructed volume, and displays the inertial structure of the LEGO Minifigure body and arms. The red box in (a) highlights a hand, while the blue box marks the print.

4.5 Discussion

4.5.1 Virtual Representation Performance

This experiment demonstrated a procedure for obtaining a virtual representation of an experimental XCT scanner, using an optimisation approach to determine unknown parameters of the simulation. The resulting simulated projections show strong similarity to experimentally acquired projections, where the average HI value was 0.824 and the RMSE was 0.0123 for the entire acquisition set, suggesting that the simulation captures the essential behaviour of the scanner.

The main deviations of the simulated projections compared to the experimental projections are found in the attenuation of the hands and the hair of the LEGO Minifigure. The attenuation error in the hands is most likely caused by inaccurate modelling of the material properties of POM. This is unexpected since there was a greater uncertainty in the material properties of ABS (see Section 4.2.3), which the body was modelled after. However, as seen in Figure 4.11c, the body is the most accurate in terms of intensity as the intensities overlap very closely. The observation that the hair was more attenuating in the experimental projections than in the simulated projections could be caused by several factors. Because of the hair being a brighter colour than the rest of the minifigure, there is a possibility that the ABS of the hair is mixed with a pigmentation that contains slightly heavier particles which would result in more attenuation in the experimental projection. This effect would not be captured in the simulation as everything except the hands was modelled as pure ABS.

The errors that arise in intensity of the simulated projections largely seem to be connected to uncertainties in material modelling, rather than inaccurate simulations of X-ray characteristics. The presence of multiple materials within the sample increases the complexity of accurately reproducing X-ray attenuation behaviour, an issue that becomes more prominent when modelling polymers. Even though material modelling can explain the main deviations between simulated and experimental projections there are some effects that might be attributed to other causes. As seen in the absolute error plot in Figure 4.12, the edges of the LEGO Minifigure often contain an error, which could be explained by imperfections in the extracted mesh. Another possible reason is geometric and optical blurring effects, which are not captured in the simulation, which would make edges appear unnaturally sharp in the simulation. As the simulation in gVXR is attenuation based, the lack of simulating scattering and phase effects could also contribute to errors.

4.5.2 Limitations and Future Work

The main limitation of this experiment is that the optimisation and evaluation were performed using a single set of scanning parameters, with tube voltage, tube current, and exposure time kept constant throughout. Therefore, no conclusions can be drawn about

how well the virtual representation generalises to other parameter combinations. This is particularly relevant for the scintillator response, which is energy dependent. The optimal scintillator, $525\text{ }\mu\text{m CdWO}_4$, may be an accurate match at the evaluated energy range, but could differ from the experimental detector response at higher tube voltages. Future work should therefore repeat the optimisation across a range of scanning parameters and to evaluate and improve the generalisability of the virtual representation. Applying the optimisation procedure to another experimental XCT scanner would further assess how well the developed method generalises beyond the Nanotom.

Another limitation of the optimisation method is that it is based solely on one projection. This choice was made to speed up the process, but makes the results sensitive to that specific angle of the aluminium crucible. The LEGO Minifigure demonstrated that metric values vary considerably across angles, as reflected in the difference between the single-projection HI of 0.873 at $\theta = 0^\circ$ and the all-angle HI of 0.824. Incorporating multiple angles into the optimisation would reduce this sensitivity.

Finally, the evaluation would benefit from testing on additional samples with known material properties, which would provide a more decisive assessment of how accurately the virtual representation recreates the Nanotom XCT scanner. In addition to comparing projections, extending the evaluation to reconstructed volumes would provide a more complete assessment. A final evaluation of the virtual representation would also require investigating its performance for various resolutions. As magnification increases, the geometric blur should have greater effects (Equation 2.1.10), and blurring might be necessary to incorporate to the virtual representation.

Chapter 5

Segmentation of Fibres

XCT can be used to study the internal material structure of fibre-based materials. These materials often consist of complex matrices of closely packed fibres, such as paperboard. Tracking individual cellulose fibres through such volumes would enable obtaining essential material properties such as fibre length and orientation. Simple segmentation approaches such as thresholding do not produce reliable results for this purpose, making deep learning-based segmentation a more promising alternative. However, training such deep learning models requires large amounts of labelled data, which is difficult to obtain. One solution is to generate synthetic datasets from simulated XCT volumes, which are already labelled by construction.

Here we investigate whether the simulation method described in Chapter 3 and 4 can generate synthetic fibre XCT data, used to train a CNN segmentation model that generalises to experimentally acquired fibre data. The experiment is designed as a preliminary validation using a dataset of a simple fibre structure consisting of a single cotton thread. Additionally, the segmentation is a binary segmentation task, classifying voxels as either fibre or air, rather than tracking individual fibres through the volume. This experiment, therefore, serves as a first step in the broader goal of segmenting individual fibres in a complex volume.

5.1 Experimental Fibre Dataset

A cotton thread was selected as the subject of the experimental fibre dataset. Compared to a paperboard fibre network, where cellulose fibres are randomly oriented and densely packed, it is a simpler and more manageable case, as the fibres are aligned in a single direction. However, the individual fibres still share the most important structural features.

The initial idea was to obtain the fibre dataset using the Nanotom Phoenix M Scanner, described in Section 4.1. The thread sample is visualized in Figure 5.1a, mounted in its holder and wrapped in tape to avoid movement during the scan. According to the manufacturer,

the scanner is capable of a minimum voxel resolution of $0.3\ \mu\text{m}$ [41]. However, in practice, achieving low voxel sizes proved difficult, and after several attempts the best resolution obtained was $1.5\ \mu\text{m}$. At this resolution it is possible to distinguish the individual fibres, but it is insufficient to resolve the fibre structure in detail, as seen in Figure 5.1c and 5.1d.




Figure 5.1: (a) The thread sample. (b) The experimental setup. (c) A reconstructed slice, with a white square marking the close-up of fibres shown in (d).

Therefore, the fibre dataset used in this work was instead provided by the 4D Imaging Lab at LTH, which was acquired using a Zeiss XRM520 at the Division of Solid Mechanics, Lund University. The cotton thread sample was scanned at a voxel resolution of $0.72\ \mu\text{m}$, providing a higher image quality and better resolved fibre structure. A representative slice of the reconstructed volume is shown in Figure 5.2, with a close-up on one fibre. Individual fibres are clearly distinguishable from the background, and in some cases the hollow lumen can be observed as a dark interior cavity. A characteristic feature of the dataset is the phase contrast fringe appearing at the fibre boundaries. Constructive interference of X-rays

produces a bright edge, while destructive interference results in a dark edge surrounding each fibre. Although laboratory-based XCT scanners often focus on attenuation contrast, phase effects are always present to some degree and become particularly noticeable when imaging low attenuation materials such as cotton fibres, as discussed in Section 2.1.1. Streak artefacts are also present, likely arising from scattering.




Figure 5.2: Representative preprocessed slice from the dataset. A close-up of a single fibre clearly illustrates phase effects, visible as a dark outer edge and a bright inner edge.

To obtain a test set for evaluation of segmentation models, the experimental fibre dataset was preprocessed and manually segmented. The raw volume, with a shape of $996 \times 996 \times 996$, contains a circular masked region at the edges, probably introduced to suppress artefacts in those areas. To focus the dataset on the fibres, the masked region was removed through a crop in the transverse plane to 636×636 pixels. The top and bottom 1 % of intensity values throughout the volume were then clipped to remove extreme outliers caused by reconstruction artefacts, producing a more even intensity range. To reduce the amount of manual segmentation required, every 50th slice was retained, resulting in a final dataset of 20 slices and a total test set of size $20 \times 636 \times 636$. These slices were manually segmented using the software tool 3D Slicer [43], where fibres were parted from background to generate a ground-truth mask. The hollow lumen was not included in the mask.




Figure 5.3: Processing pipeline and segmentation for the experimental fibre dataset, displayed for one: (a) raw slice, (b) preprocessed slice, including cropping and removal of outliers, (c) manually segmented slice.

5.2 Generation of Synthetic Training Data

The initial idea was to generate the synthetic data using the virtual representation of the Nanotom XCT scanner, described in Chapter 4, and to base the scanning parameters on the experimental scanner configuration used to acquire the experimental fibre dataset. However, the experimental dataset was ultimately acquired using a different scanner for which no virtual representation was developed, so the simulation parameters could not be based on a known scanner configuration. Instead, simulation parameters were randomly sampled in a plausible range, with the purpose of generating a diverse dataset acquired under different imaging conditions. The idea is that a trained model can learn to generalise and segment the experimental dataset, even though none of the synthetic datasets directly represent the experimental scanner, a strategy known as domain randomisation.

5.2.1 Simulation of X-ray Computed Tomography

As input to the simulation model, 10 virtual cotton fibres were used, each represented as a triangular mesh as shown in Figure 5.4. The meshes were segmented from the experimental fibre data described in the previous section, ensuring that the fibre geometries are representative of actual cotton fibres. Experimental cotton fibres consist of approximately 90% cellulose, but are in this work modelled as pure cellulose ($\text{C}_6\text{H}_{10}\text{O}_5$), with elemental mass fractions 44.4% carbon, 6.2% hydrogen, and 49.4% oxygen, and a density of 1.5 g/cm^3 . However, to vary the linear attenuation coefficient across datasets, the density was treated as a free parameter and sampled within a plausible range. The sample orientation was varied by applying two rotations: pitch and yaw. Although the same 10 fibres and relative geometry is preserved across all generated XCT volumes, this ensures fibres appear at different positions and orientations within the reconstructed volume.



Figure 5.4: Triangular meshes of cotton fibres used as input to the simulation model.

The pixel array was kept constant at 996×996 to ensure that the generated XCT volumes and the raw experimental dataset had the same dimensions. Later, the volumes were cropped in the transversal plane, exactly like the experimental dataset, to 636×636 pixels. The geometry between source, detector, and sample as well as pixel size were configured to obtain approximately the same resolution of $0.72 \mu\text{m}$ in the experimental dataset. Otherwise, parameters were varied to generate diverse datasets, with varying contrast, noise levels and spatial blur. The ranges of parameters is summarised in Table 5.1. The volumes were reconstructed using a variety of filters. Cross sectional slices with the same index from three volumes generated with different combinations of parameters are visualised in Figure 5.5.




Figure 5.5: The same index cross sectional slice for three XCT volumes generated with parameters within ranges of Table 5.1. The effect of yaw and pitch rotation causes the same 10 fibres to appear in different locations. Slice (a) is generated without flat field correction, while (b) and (c) are generated with flat field correction, which produces the circular noise pattern. Note that these slices have been cropped to 636×636 pixels.

The flat field correction feature was used, where specifying a number of flat field images greater than zero simulates ring artefacts in the reconstructed volume. Among the XCT

Varied Parameter	Range	Unit
Tube voltage	40 – 80	kV
Tube current time-product	2 – 8	mAs
Aluminium filter thickness	0.2 – 1	mm
Source-sample distance (SSD)	8 – 13	mm
Sample density	1 – 1.52	g/cm ³
Yaw (sample offset around z-axis)	–60 – 60	°
Pitch (sample offset around y-axis)	–20 – 20	°
Line spread std. deviation (σ)	0.01 – 0.05	–
Number of flat field images	0 – 50	–
Number of projections	1200 – 1601	–
Fixed Parameter	Value	Unit
Anode angle	12	°
Source-detector distance (SDD)	20.42	mm
Detector pixel size	1.35	μm
Pixel array (width, height)	(996, 996)	
Scintillator material	CdWO ₄	–
Scintillator thickness	525	μm
Fibre composition	C ₆ H ₁₀ O ₅	–
Photon shot noise	Enabled	–

Table 5.1: Parameters used for XCT data generation. Varied parameters were sampled within the stated ranges; fixed parameters were held constant across all volumes.

slices visualised in Figure 5.5, this was applied in the generation of (b) and (c). Although the experimental fibre dataset does not exhibit visible ring artefacts, the flat field correction is included to introduce additional variability, making the trained segmentation models more robust to artefacts.

The Line Spread Function (LSF) was applied when generating the XCT volumes. In experimental datasets, photon shot noise exhibits spatial correlation between neighbouring voxels, caused by optical blur in the detector system, as discussed in Section 2.1.2. Without the LSF, we found that the reconstructed slices exhibited an unnatural pixelated background noise inconsistent with this behaviour. Applying a Gaussian LSF, as described in Section 3.2.1, introduces spatially correlated noise more representative of experimental data, as illustrated in Figure 5.6. Therefore the standard deviation (σ), which determines the amount of blur, was treated as a free parameter.




Figure 5.6: Two slices from datasets generated with identical parameters except for (a) no LSF, (b) Gaussian LSF with $\sigma = 0.02$. The slices are cropped to display 100×100 pixels.

5.2.2 Ground Truth Generation

The ground truth mask is a binary image marking the pixels where a fibre is present. There are several ways to obtain this mask when the geometric input to the simulation model is known. In this work we designed a simple method where the mask of a synthetic XCT volume was obtained through a noiseless scan using identical parameters but with photon shot noise disabled, no LSF and a monochromatic energy spectrum. The monochromatic energy spectrum, consisting of 1600 incident photons with the assigned tube voltage energy, is used to avoid beam-hardening artefacts. Since a polychromatic energy spectrum, noise modelling, and LSF convolution all increase the computational cost of generating a projection, the noise-free scan was significantly faster to simulate than the corresponding synthetic XCT volume. The resulting noiseless volume was then segmented using Otsu thresholding, producing a binary mask. Since the fibre and background intensities are well separated in the noiseless volume this segmentation method is reliable.

5.2.3 Modelling of Phase Effects

As gVXR is solely attenuation-based, phase contrast effects are not captured in the simulation. In an attempt of matching experimental fibre dataset characteristics, a method for modelling the phase contrast fringe was developed. This is a post-processing step applied after the synthetic XCT volume and corresponding ground truth mask has been generated. For each slice, the binary mask is morphologically dilated and eroded to define two edge regions: an outside edge (3 voxels wide) and an inside edge (2 voxels wide). The dark outside edge is assigned an intensity lower than the local air value, defined as the slice minimum. The bright inner edge is assigned an intensity greater than the local fibre intensity, defined as the mean intensity within the masked region. The intensities are controlled by empirically determined scaling factors of 0.3 and 4.3 for the dark and bright edges respectively, ensuring the edges remain darker than the surrounding air and brighter than the fibre interior.




Figure 5.7: Slice from synthetic XCT volume with (a) no phase effects, (b) added phase effects in a post-processing step. Both slices are cropped to display 350×350 pixels.

5.2.4 Composition of Training Dataset

A synthetic training dataset consists of reconstructed slices and their corresponding ground truth mask, combined from multiple XCT volumes generated with different parameter combinations. Since neighbouring slices within a XCT volume are nearly identical, only every 10th slice is retained to avoid over fitting and avoid redundancy. Phase effect modelling is applied at the volume level, meaning all slices within a given volume either have phase effects modelled or not. When splitting into training and validation sets, all slices from a given volume are assigned to the same split, ensuring that the specific scanning parameters of any volume appear in only one set. To match the experimental fibre dataset dimensions, which serves as the test set in the segmentation model, all volumes are cropped in the transversal plane from 996×996 to 636×636 pixels.

5.3 Segmentation Model

For this work a simple U-NET architecture was chosen for the CNN model since this architecture is commonly used for image segmentation purposes. Here each pixel is either classified as either foreground (fibre) or background (air). The input for the training and validation of this model is the synthetic datasets generated described in the previous section, together with their respective binary masks. Each slice in the training and validation data is normalised between 0 and 1 as part of the training.

5.3.1 Model Architecture




Figure 5.8: Schematic visualisation of the U-NET architecture used in this chapter.

The model follows the common encoder-decoder architecture for U-NETs with a symmetric structure and skip connections and is displayed in Figure 5.8. The encoder progressively extracts hierarchical features using convolutional blocks combined with downsampling. Each convolutional block in this model consists of two convolutional layers that are activated using a rectified linear unit (ReLU) function. At the deepest level, the bottleneck performs a convolution on a highly compressed representation of the input using a large number of filters. This stage serves as a transition between the encoder and decoder. The decoder reconstructs the spatial resolution of the feature maps through upsampling operations. Each upsampling is concatenated with corresponding feature map in the encoder via skip connections. The skip connections help maintain the finer spatial details that are otherwise lost during the downsampling. The networks then outputs a segmentation map through a 1x1 convolutional layer with a sigmoid activation function. The sigmoid activation function is given by

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (5.3.1)$$

and produces a value in the range $[0,1]$ for each pixel. The output prediction for each image will be a probability map where each value is the probability of that pixel belonging to the true class (fibre). This prediction is turned binary by applying a threshold to decide the upper limit at which probability a pixel will be considered a fibre. The optimal threshold is found by calculating the precision (Equation 5.3.2) and recall (Equation 5.3.3) for a large range of thresholds. Precision and recall are then used to calculate the F1-score (Equation 5.3.4) and the threshold at which the F1-score is at its maximum is chosen as the threshold for the entire predicted probability map.

5.3.2 Evaluation Metrics

The predicted classification of images can be evaluated by comparing the positive and negative class distribution of the test data. These are the true positive, false positive, false negative and true negative values for the binary classes which are explained below. These quantifiers are important to assess whether there is a bias towards either of the classes [44]. The number of correctly classified positive samples is denoted by a true positive (TP). The number of pixels in an XCT slice containing a fibre is for example correctly predicted as pixels containing a fibre. The number of incorrectly classified positive samples is denoted by a false positive (FP). The pixels in an XCT slice not containing a fibre is for example falsely predicted as pixels containing a fibre. The number of incorrectly classified negative samples is denoted by a false negative (FN). The pixels in an XCT slice containing a fibre is for example falsely predicted as pixels not containing a fibre.

The mentioned quantifiers are used to calculate different metrics such as precision, recall and F1-score. Precision measures the ratio between the correctly predicted pixels belonging to the true class and the total number of pixels predicted as true, and is expressed as

$$\text{Precision} = \frac{TP}{TP + FP}. \quad (5.3.2)$$

Recall measures the ratio of the correctly predicted pixels belonging to the true class and all the pixels belonging to the true class, and is expressed as

$$\text{Recall} = \frac{TP}{TP + FN}. \quad (5.3.3)$$

Both precision and recall are scores ranging between 0 and 1, where 1 represents a perfect precision or recall. The harmonic mean of precision and recall is known as the F1-score and it penalizes extreme values of either of the two [44]. The F1-score ranges between 0 to 1 and is expressed as

$$\text{F1} = 2 \cdot \frac{\text{Recall} \cdot \text{Precision}}{\text{Recall} + \text{Precision}} \quad (5.3.4)$$

where 1 represents both a perfect precision and recall score and 0 represents a zero for either or both precision and recall.

5.4 Results of Segmentation

5.4.1 Evaluation of Segmentation Models

The U-NET model described in Section 5.3.1 was used to train four different models. Each model was trained on reconstructed slices from volumes generated with scanning parameters in Table 5.1, varied to create variation in the dataset. The resulting size of the synthetic dataset is determined by the number of XCT volumes it is based on, as each volume contributes 99 slices. Generation of one XCT volume and corresponding ground-truth mask took approximately 30 minutes. This means, a dataset composed from 12 volumes can be generated in 6 hours. Each model is presented with their precision, recall and F1-score, as well as the standard deviation for each metric, calculated for the entire test set of 20 slices.

Model 1

This model was trained using a dataset comprised of six XCT volumes where four of them were used for training data and two for validation. None of these volumes had been post-processed with modelled phase effects. Model 1 generated the following metric scores:

- Precision = 0.888 ± 0.0063
- Recall = 0.914 ± 0.0220
- F1 = 0.901 ± 0.0105

The lower precision score means that the model generated more FP (red) than FN (blue). This can be observed in Figure 5.9 where the dark phase effect around the edges of the fibres are often predicted falsely as fibres.




Figure 5.9: Input slice with respective GT/prediction overlap and absolute error for slice 8, as predicted by Model 1.

Model 2

To evaluate the effect of increasing the amount of XCT volumes, Model 2 was trained on an expanded dataset relative to Model 1. The dataset used to train Model two was comprised of twelve volumes, of which eight were allocated for training and four were reserved for validation. Similar to Model 1, this model was also excluded from being post-processed with modelled phase effects. Model 2 generated the following metric scores:

- Precision = 0.884 ± 0.0187
- Recall = 0.918 ± 0.0696
- F1 = 0.901 ± 0.0310

As seen in Figure 5.10, Model 2 also generates FP around the outskirts of the fibres. By comparing the precision and recall, it's clear that FP dominates slightly over FN.




Figure 5.10: Input slice with respective GT/prediction overlap and absolute error for slice 8, as predicted by Model 2.

Model 3

To evaluate whether model performance could be enhanced by incorporating modelled phase effects around the fibres in the synthetic data, a third model was developed. This model was trained using the same dataset as Model 2, but this time all volumes in the dataset were subjected to such post-processing. Model 3 generated the following metric scores:

- Precision = 0.955 ± 0.0190
- Recall = 0.940 ± 0.0785
- F1 = 0.947 ± 0.0517

Model 3 showed a noticeable increase in all of the above scores. As seen in Figure 5.11, this model showed a more accurate prediction of fibre edges than Model 1 and Model 2.




Figure 5.11: Input slice with respective GT/prediction overlap and absolute error for slice 19, as predicted by Model 3.

Model 4

To evaluate the effect of including more variation in the training data, this model was trained using a dataset where only some of the samples had been post-processed to include modelled phase effect. This model was trained using the complete set of twelve volumes, consistent with the approach applied in Models 2 and Model 3. Within this dataset, five of the eight training volumes and two of the four validation volumes were post-processed to include modelled phase effects around the fibres. Model 4 generated the following metric scores:

- Precision = 0.962 ± 0.0087
- Recall = 0.949 ± 0.0104
- F1 = 0.956 ± 0.0075

As seen in Figure 5.12, Model 4 also shows a more accurate prediction of fibre edges compared to Model 1 and Model 2.




Figure 5.12: Input slice with respective GT/prediction overlap and absolute error for slice 19, as predicted by Model 4.

Model Comparison

The resulting precision, recall and F1-score for all the models is summarised in Table 5.2. The difference in scores generated by Model 1 and Model 2 is negligible, with a mean F1 score of approximately 0.901 for both. A substantial increase was observed for Model 3 and Model 4, which achieved mean F1 scores of 0.947 and 0.956 respectively. Model 4 achieved the highest scores for all three metrics, while also exhibiting the lowest standard deviation across slices.

	Model 1		Model 2		Model 3		Model 4	
Number of samples	6		12		12		12	
Post-processed	None		None		All		7/12	
	Value	Std	Value	Std	Value	Std	Value	Std
Precision	0.888	0.0063	0.884	0.0187	0.955	0.0190	0.962	0.0087
Recall	0.914	0.0220	0.918	0.0696	0.940	0.0785	0.949	0.0104
F1-Score	0.9011	0.0105	0.9008	0.0310	0.947	0.0517	0.956	0.0075

Table 5.2: Comparison of precision, recall, and F1-score across the four trained models under different datasets. The lowest and highest value for each metric is highlighted in red and green respectively.

All four models shared a common result in that the lowest individual F1 score was consistently obtained for either the first or last slice of the test dataset, out of the total 20 slices. These slices especially generate low recall, shown in the greater standard deviation of this metric, meaning they struggle to find the fibres and instead classify them as air. The prediction for the last slice is visualised for Model 3 in Figure 5.13, where the high amount of

FN is seen as blue areas in the overlap. Model 4 was the exception which although it also yielded its lowest individual slice score for these particular slices, all performance metrics remained above 0.93. The prediction of slice 20 as predicted by Model 4 can be seen in Figure 5.14.




Figure 5.13: Input slice with respective GT/prediction overlap and absolute error for slice 20, as predicted by Model 3.




Figure 5.14: Input slice with respective GT/prediction overlap and absolute error for slice 20, as predicted by Model 4.

5.4.2 Comparison to Threshold Segmentation

To enable comparison with the trained CNN models, a baseline segmentation was performed using Otsu thresholding. The resulting metric scores are presented below:

- Precision = 0.9940 ± 0.0180

- Recall = 0.7779 ± 0.0143
- F1 Score = 0.8728 ± 0.0129

It is clear that all trained CNN models achieve higher Recall and F1 score. However, the thresholding model achieves the highest precision, which can be explained by the low amount of FP, meaning that almost all predicted fibre voxels are fibre. The model however struggles to predict the interior of the fibre including the lumen, visualised in Figure 5.15.




Figure 5.15: Input slice with respective GT/prediction overlap and absolute error for slice 17, as generated by threshold segmentation.

5.5 Discussion

5.5.1 Segmentation Models

One surprising result was that the CNN segmentation model did not improve when training on a synthetic dataset combined from 12 compared to 6 generated XCT volumes, as no significant difference in segmentation metrics was observed between Model 1 and Model 2. This suggests that these first 6 volumes generalise sufficiently to the experimental data, and that generating additional volumes did not provide further useful information to the model. It can be observed in Figure 5.9 and Figure 5.10 that both these models struggled with accurate edge detection, where the dark phase effect outside the fibres are mistakenly predicted as fibres.

It is clear that adding post-processed phase had a positive effect for the model, demonstrated in an increase in all segmentation metrics for Model 3. As seen in Figure 5.11, the prediction by Model 3 contains noticeable less false positives around the outskirts of the experimental fibres, which is also reflected by the increase in F1-score. However, the results suggest that greater variation in the synthetic training dataset is favourable, demonstrated

by Model 4 achieving the highest segmentation metrics overall, where the training data combined slices with and without post-processed phase edges.

The key difference between Model 3 and Model 4 is reflected in the standard deviation of their scores. Model 3 exhibited greater variable performance across slices, which was attributed to poor predictions on the first and last slices of the test dataset. These boundary slices correspond to the the first and last slice in the raw experimental fibre dataset, which are known to exhibit different properties than the rest of the volume due to difficulties reconstructing accurate intensities at greater angles of the cone beam. As the modelled phase edges are defined relative to the local intensity distribution of each slice, the post-processing homogenises the intensity distribution across slices, reducing the overall intensity variation in the synthetic dataset. The lack of exposure to intensity variations during training could explain the poor predictions of Model 3 on the first and last slice of the test dataset, the latter displayed in Figure 5.13. Model 4, by contrast, retained greater intensity variation through its mix of post-processed and unprocessed volumes, and was therefore more robust to the boundary slices, as seen in Figure 5.14. One suggestion for improving the phase effect modelling in this aspect could be to vary the scaling factors of the phase edges.

The overall results of this experiment demonstrate that a CNN segmentation model trained on solely synthetic computed tomography datasets can segment an experimental dataset to a high degree of accuracy, demonstrated by a mean F1 score of 0.95 for the best model. The experimental fibre dataset of this work is not particularly difficult to segment, and even a simple threshold segmentation model obtains relatively good result, with a mean F1 score of 0.87. The CNN models still outperform thresholding by capturing the full fibre structure, including the fibre interior, which is reflected in the higher recall score of 0.95 compared to 0.78. It was demonstrated that the Otsu threshold model struggles to capture the centre of the fibres, likely because the threshold is biased by the bright phase contrast fringe at the fibre boundary. Lowering the threshold would recover the fibre interior but would risk misclassifying background as fibre.

5.5.2 Limitations

One of the main limitations of the developed model is the is the restricted number of fibre meshes available as input to the simulation. With only 10 meshes available, the geometric variation within the simulated dataset is limited, which could affect the models ability to generalise to fibres with other shapes. In this work, this was addressed by varying the positions and orientations of fibres across volumes, and the results suggest that the model can segment all fibre types present in the experimental fibre dataset. However, it is important to note that the meshes were segmented from the same experimental dataset used for evaluation, introducing a degree of geometric dependency between the training and test data. A stronger evaluation would use fibre meshes sourced independently of the test dataset. Expanding the model to more complex fibre networks would require a larger mesh library.

These meshes were produced through manual segmentation, an approach that does not scale well, and an automated mesh generation would therefore be necessary.

Another limitation concerns the physics of the simulation, as gVXR is solely attenuation-based. By varying scanning parameters, different imaging conditions can be reflected, but effects from scattering and phase contrast cannot be simulated directly. In this work, phase contrast effects were approximated through a post-processing step, which proved effective for segmenting the experimental cotton fibre dataset. However, we are concerned this approach might become more complex when applied to denser fibre networks with overlapping fibres in several directions, such as in paperboard. It should be noted that models trained without phase effects could still segment the fibres, but with less generalisation to the fibre edges, as demonstrated by Model 1 and Model 2. This suggests that attenuation-based simulation alone may be sufficient for segmentation, though whether this holds for more complex fibre structures would need to be further evaluated.

5.5.3 Future Work

The work described in this chapter also shows a lot of promise for future experiments which could not be performed due to time constraints. Three main future work approaches have been identified with the overall goal to of combing them into a single approach.

Volumetric Segmentation

This work focused on the segmentation of fibres in individual slices of a volumetric dataset. A natural extension of this work is to perform segmentation at the volumetric level. This would involve using synthetic XCT volumes as training and validation data, while experimental volumes would be employed for testing.

Multi-class Segmentation of Fibres

Another area of interest is the segmentation of individual fibres. This is a complex problem, largely limited by the availability of sufficiently labelled datasets for training. The results of this work demonstrate that synthetic XCT data can be effectively used for binary segmentation. Extending this approach to incorporate multi-class labelling in the synthetic data could therefore represent a promising step toward enabling accurate segmentation of individual fibres.

Segmentation of Synthetic Paperboard Fibres

The work in this chapter focuses on segmenting fibres from XCT data of a cotton thread. As mentioned, this is a simplified problem to a much more complex one, which is to segment fibres from XCT data of paperboard. The fibres in paperboard are not as unidirectional as the fibres in a cotton thread, which all follow the same direction in a twinned manner. Extending this approach to paperboard would not only require the generation of synthetic

XCT data, but also the creation of realistic synthetic fibre structures to serve as input for the simulations. Recently, a synthetic paperboard dataset developed by Morgan Görtz *et al.* [45] has been made available, consisting entirely of individually generated synthetic fibres. Integrating such datasets with XCT simulation frameworks would enable further investigation into the use of fully synthetic data, both in terms of geometry and imaging, for training deep learning models in fibre segmentation tasks.

Chapter 6

Conclusions

The work in this thesis has focused on two distinct parts with their own objectives: developing a virtual representation of an XCT scanner and generating labelled synthetic training data for deep learning-based fibre segmentation.

In the first part, a virtual representation of the Nanotom XCT scanner located at Tetra Pak was developed using the open-source framework gVXR. The workflow included designing an optimisation method to find the unknown internal parameters of the Nanotom XCT scanner. By using a sample with known composition, different parameters were tried iteratively to match experimentally acquired data as best as possible. The post-optimised virtual representation produces radiographs where the intensity values strongly overlap with the experimental equivalent, producing an average HI value of 0.824 and a RMSE of 0.0123 for the full acquisition while the non-optimised projection produced an HI of 0.541 and a RMSE of 0.135. The offsets in intensity values are thought to be related to uncertainties in correct modelling of the chemical composition and properties of the virtually scanned LEGO Minifigure. Even though promising results are exhibited, the extent of the virtual representation’s accuracy still remains inconclusive due to only performing one evaluation experiment, which covers a specific resolution with specific acquisition settings. The specific acquisition settings and the fact that only one projection angle was used in the optimisation procedure are limitations of this work, which creates uncertainties in the found XCT parameters and their compatibility with other acquisition settings.

In the second part, the simulation model developed in the first part was used to generate labelled synthetic XCT data of cotton thread fibres to train a convolutional neural network for binary segmentation of experimental XCT data. The workflow developed to generate the synthetic XCT data of fibres included varying the simulation parameters to generate diverse datasets, as well as a post-processing step to add a modelled phase effect around the fibres in the simulated XCT slices. A U-NET architecture was used for the CNN models, which trained on synthetic XCT data produced accurate segmentation results. The best performing model produced an F1-score of 0.956, where a perfect and worst score possible is 1 and 0 respectively. Overall, it was clear from the increase in F1-score by Model 3

compared to Model 2 (see Table 5.2) that adding the post-processing step of modelled phase effects to the fibres improved the segmentation models ability to detect edges accurately. The best segmentation model was the one containing the most variation in the training dataset (see Table 5.2). Besides the variation of simulation parameters, whether the modelled phase effect around the fibre edges was added or not also varied in the training data for this model. The limitations of this work include the access to accurate fibre meshes to generate simulations with, as well as the simulation's inability to generate phase contrast in images. This introduces the need for incorporating modelled phase effects as a post-processing step.

Due to time constraints, there still remains details to explore within both parts of this work. To determine the general accuracy of the virtual representation, further evaluation tests where the acquisition settings are varied would be required. This would also most likely introduce the need for a more developed optimisation method that takes multiple acquisition settings, angles and resolutions into account. The segmentation of fibres, although successful, was a rather simple problem of binary segmentation of 2D images as a proof of concept for using synthetic XCT data in deep learning. The future work is to segment individual fibres with multi-class segmentation in 3D volumes, as well as doing this for paperboard fibres, which are more complex and challenging than cotton thread fibres.

Bibliography

- [1] Carmignato S, Dewulf W, Leach R. Industrial X-ray Computed Tomography; 2017.
- [2] Racy M, Barrow A, Tomlinson J, Bello F. Development and Validation of a Virtual Reality Haptic Femoral Nailing Simulator. *Journal of Surgical Education*. 2021;78(3):1013-23. Available from: <https://www.sciencedirect.com/science/article/pii/S1931720420303810>.
- [3] Sujar A, Kelly G, García M, Vidal FP. Interactive teaching environment for diagnostic radiography with real-time X-ray simulation and patient positioning. *International Journal of Computer Assisted Radiology and Surgery*. 2022;17(1):85-95. Available from: <https://doi.org/10.1007/s11548-021-02499-7>.
- [4] Corbi A, Burgos D, Vidal FP, Albiol F, Albiol A. X-ray imaging virtual online laboratory for engineering undergraduates. *European Journal of Physics*. 2020;41(1):014001.
- [5] Vidal FP, Afshari S, Ahmed S, Albiol A, Albiol F, Béchet E, et al. X-ray simulations with gVXR in education, digital twining, experiment planning, and data analysis. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*. 2025;568:165804. Available from: <https://www.sciencedirect.com/science/article/pii/S0168583X25001946>.
- [6] Friemann J, Mikkelsen LP, Oddy C, Fagerström M. Synthetic, automatically labelled training data for machine learning based X-ray CT image segmentation: Application to 3D-textile carbon fibre reinforced composites. *Composites Part B: Engineering*. 2025;305:112656.
- [7] Walsh CL, et al. The Human Organ Atlas. *Science Advances*. 2026.
- [8] RISE Research Institutes of Sweden. AI-Tomo: Accelerated Materials Characterisation by AI & X-ray Tomography; 2024. Accessed: 2026-05-27. Available from: <https://www.ri.se/en/materials-and-durability/composites/project/ai-tomo-accelerated-materials-characterisation-by-aix>.
- [9] Villarraga-Gómez H, Herazo EL, Smith ST. X-ray computed tomography: from medical imaging to dimensional metrology. *Precision Engineering*. 2019;60:544-69.
- [10] Withers PJ, Bauman C, Carmignato S, Cnudde V, Grimaldi D, Hagen CK, et al. X-ray computed tomography. *Nature Reviews Methods Primer*. 2022.
- [11] Fernandez JE. In: Interaction of X-rays with matter. In book: *Microscopical X-Ray Fluorescence Analysis*; 2000. .
- [12] Als-Nielsen J. *Elements of Modern X-Ray Physics*. vol. 55; 2002.
- [13] Wu J. Propagation based X-ray phase-contrast imaging technique for the microstructure analysis of biological soft tissues. *Optik*. 2014;125(3):1062-4. Available from: <https://www.sciencedirect.com/science/article/pii/S0030402613011510>.
- [14] Russo P. *Handbook of X-ray Imaging Physics and Technology*. 1st ed. Boca Raton: CRC Press; 2017.

- [15] Keklikoglou K, Faulwetter S, Chatzinikolaou E, Wils P, Brecko J, Kvaček J, et al. Micro-computed tomography for natural history specimens: a handbook of best practice protocols. *European Journal of Taxonomy*. 2019 04:1-55.
- [16] Academy DI. All About X-Ray Tubes: Construction, Components & Working;. Visited 14 Apr 2026. Available from: https://doctorinsideacademy.com/blog/all-about-x-ray-tube/#X-Ray_Filters.
- [17] The Conference of Radiation Control Program Directors (CRCPD). Beam Quality: Total Filtration and Half-Value Layer; 2001. Visited 14 Apr 2026. Available from: https://crcpd.org/wp-content/uploads/2023/11/10-01_QA_Filtration-2.pdf.
- [18] Badea CT. Chapter 4 - Principles of Micro X-ray Computed Tomography. In: Ross BD, Gambhir SS, editors. *Molecular Imaging* (Second Edition). second edition ed. Academic Press; 2021. p. 47-64.
- [19] Talebinezhad H, Fischer R, Shmatok A, Prorok BC. A straightforward and analytical method for precise X-ray CT magnification correction. *Precision Engineering*. 2023;82:338-49. Available from: <https://www.sciencedirect.com/science/article/pii/S0141635923000788>.
- [20] Zhong G, Li F, Yu H, Liu C, Yang R, Zhou R. Method for expanding the field-of-view of micro-CT imaging using array micro-focus X-ray source. *Opt Express*. 2025 Apr;33(7):16141-60. Available from: <https://opg.optica.org/oe/abstract.cfm?URI=oe-33-7-16141>.
- [21] Takase A. How Does CT Reconstruction Work?; 2023. Visited: 19 Feb 2026. Available from: <https://rigaku.com/products/imaging-ndt/x-ray-ct/learning/blog/how-does-ct-reconstruction-work>.
- [22] Schofield R, King L, Tayal U, Castellano I, Stirrup J, Pontana F, et al. Image reconstruction: Part 1 – understanding filtered back projection, noise and image acquisition. *Journal of Cardiovascular Computed Tomography*. 2020;14(3):219-25. Available from: <https://www.sciencedirect.com/science/article/pii/S1934592519300607>.
- [23] Kim H, Choi JS, Lee Y. Assessment of Feldkamp-Davis-Kress Reconstruction Parameters in Overall Image Quality in Cone Beam Computed Tomography. *Applied Sciences*. 2024 08;14:7058.
- [24] Biguri A, Dos Santos R, Dos Santos M, Blazquez A, Schönlieb CB. TIGRE: Tomographic Iterative GPU-based Reconstruction Toolbox; 2016. Accessed: 2026-05-12. Available from: <https://github.com/CERN/TIGRE>.
- [25] Sun Y. Cone Beam FDK Reconstruction; 2025. Available from: https://syypsyp97.github.io/diffct/fdk_cone_example.html.
- [26] Seung-Wan L, Chang-Lae L, Hyo-Min C, Hye-Suk P, Dae-Hong K, Yu-Na C, et al. Effects of Reconstruction Parameters on Image Noise and Spatial Resolution in Cone-beam Computed Tomography. *Journal of the Korean Physical Society*. 2011 10;59:2825.
- [27] Barrett JF, Keat N. Artifacts in CT: recognition and avoidance. *Radiographics*. 2004;132.
- [28] Virdhi R, Gupta V, Kaur G. Research Paper Based on Artificial Neural Network. *International Journal of Research and Analytical Reviews (IJRAR)*. 2023;10(2):1188-94. Available from: <https://ijrar.org/papers/IJRAR23B1188.pdf>.
- [29] Harsh. Understanding Artificial Neural Network (ANN); 2020. Accessed: 2026-06-18. Available from: <https://www.codespeedy.com/understanding-artificial-neural-network-ann/>.
- [30] Chaudhary S, Sahani A, Akhtar S, Pandey S. A Comprehensive Review of Artificial Neural Networks: Architectures, Learning Algorithms, and Real-World Applications. *Journal of Scientific Innovation and Advanced Research*. 2025 April;1(1). Available from: <https://jsiar.com/2025-April/JSIAR-A-25-04001.pdf>.

- [31] Minaee S, Boykov Y, Porikli F, Plaza A, Kehtarnavaz N, Terzopoulos D. Image Segmentation Using Deep Learning: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2022.
- [32] LeCun Y, Bengio Y, Hinton G. Deep Learning. *Nature*. 2015;521:436-44.
- [33] Unzueta D. Fully Connected Layer vs. Convolutional Layer: Explained; 2025. Accessed: 2026-05-04. <https://builtin.com/machine-learning/fully-connected-layer>.
- [34] O'Shea K, Nash R. An Introduction to Convolutional Neural Networks; 2015. Available from: <https://arxiv.org/abs/1511.08458>.
- [35] Agostinelli S, Allison J, Amako K, Apostolakis J, Araujo H, Arce P, et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. 2003;506(3):250-303.
- [36] Pointon JL, Wen T, Tugwell-Allsup J, Sújar A, Létang JM, Vidal FP. Simulation of X-ray projections on GPU: Benchmarking gVirtualXray with clinically realistic phantoms. *Computer Methods and Programs in Biomedicine*. 2023;234:107500. Available from: <https://www.sciencedirect.com/science/article/pii/S0169260723001669>.
- [37] Huang W, Shao H, Petrò S, Moroni G. The X-ray computed tomography simulation in geometric metrology: a review and case study. *Procedia CIRP*. 2024;129:145-50. 18th CIRP Conference on Computer Aided Tolerancing (CAT2024). Available from: <https://www.sciencedirect.com/science/article/pii/S2212827124011624>.
- [38] Bujila R, Omar A, Poludniowski G. A validation of SpekPy: A software toolkit for modelling X-ray tube spectra. *Physica Medica*. 2020;75.
- [39] Schoonjans T, Brunetti A, Golosio B, Sanchez del Rio M, Solé VA, Ferrero C, et al. The xraylib library for X-ray-matter interactions. Recent developments. *Spectrochimica Acta Part B: Atomic Spectroscopy*. 2011;66(11):776-84. Available from: <https://www.sciencedirect.com/science/article/pii/S0584854711001984>.
- [40] Jørgensen JS, Ametova E, Burca G, Fardell G, Papoutsellis E, Pasca E, et al. Core Imaging Library - Part I: a versatile Python framework for tomographic imaging. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2021 Jul;379(2204):20200192. Available from: <http://dx.doi.org/10.1098/rsta.2020.0192>.
- [41] Waygate Technologies. Phoenix Nanotom M: 180 kV / 20 W X-ray nanoCT System for High-Resolution Analysis and 3D Metrology; 2021. Accessed: 2026-05-06. Product brochure, Baker Hughes Company, document no. BHCS38584. Available from: https://www.bakerhughes.com/sites/bakerhughes/files/2024-10/bhcs38584_phoenix_nanotom_m_brochure_r2_1.pdf.
- [42] Engineering Reference. ABS Plastic Properties, Density & Applications: Complete Engineering Reference; 2026. Accessed: 2026-05-25. Available from: <https://plasticweightcalculator.com/guides/abs-plastic-properties-applications/>.
- [43] Fedorov A, Beichel R, Kalpathy-Cramer J, Finet J, Fillion-Robin JC, Pujol S, et al. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*. 2012;30(9):1323-41. Quantitative Imaging in Cancer.
- [44] Hicks S, Strumke I, Thambawita V, Hammou M, Riegler M, Halvorsen P, et al. On evaluation metrics for medical applications of artificial intelligence. *Scientific Reports*. 2022 04;12.
- [45] Görtz M, Kettl G, Målvik A, Fredlund M, Edelvik F. Iterative method for large-scale Timoshenko beam models assessed on commercial-grade paperboard. *Computational Mechanics*. 2024 05;75:1-13.

Appendix

A Optimisation

Scintillator materials available in gVirtualXray:

- CsI
- CdWO₄
- Y₂O₃
- Gd₂O₃
- (Y, Gd)₂O₃
- Gd₂O₂S
- Gd₂O₂S DRZ-Plus
- NaI
- La₂HfO₇
- Gd₃Ga₅O₁₂
- Y₃Al₅O₁₂