



Faculty of Science

Ensuring Flawlessness in Additive Manufacturing: Advances in X-ray Inspection Techniques for Efficient Defect Detection

Thesis submitted in fulfilment of the requirements for the degree of

DOCTOR OF SCIENCE: PHYSICS

at the University of Antwerp by

Domenico Iuso

Antwerpen, 2024

Supervisors
prof. dr. Jan Sijbers
prof. dr. Jan De Beenhouwer

Jury**Chairman**

prof. dr. Gunther Steenackers, University of Antwerp, Belgium

Supervisors

prof. dr. Jan Sijbers, University of Antwerp, Belgium

prof. dr. Jan De Beenhouwer, University of Antwerp, Belgium

Members

prof. dr. Geert Van der Snickt, University of Antwerp, Belgium

dr. Jeroen Soete, Katholieke Universiteit Leuven, Belgium

prof. dr. ing. Gabriel Herl, Technische Hochschule Deggendorf, Germany

The research presented in this thesis was partly supported by the VLAIO grants HBC.2019.2808 and HBC.2022.0094.

Contact

Domenico Iuso

University of Antwerp

Faculty of Physics, imec-Vision Lab

Universiteitsplein 1, 2610 Antwerpen, Belgium

Domenico.Iuso@uantwerpen.be

Copyright © 2024 by Domenico Iuso

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, broadcasting or by any other information storage and retrieval system without written permission from the copyright owner.

Dutch title:

**Zorgen voor Perfectie in Additive
Manufacturing: Innovaties in
Röntgeninspectietechnieken voor
Efficiënte Defectdetectie**

Acknowledgements

June 2020, as I sat on a chair with my chin resting on my right wrist, doing my best impression of Rodin's thinker, my mind was a chaotic mix of PhD daydreams and apocalyptic scenarios. The new social and work environment ahead seemed as uncertain as a bad sci-fi plot, and let's not forget our ramping virulent acquaintance that was determined in turning humanity into a history lesson. Despite these challenges, we moved forward, swapped out old acquaintances for new ones, and bit by bit, my uncertainty was replaced by friendly faces and intriguing research topics.

These research topics, though fascinating, were often a cocktail of bitterness, sourness, and the occasional sip of joy. Over my four-year journey, I frequently found myself questioning my life choices, but thankfully, I had a crew of amazing people who turned this rocky road into an enjoyable adventure with their jokes, support, and willingness to listen to my endless rants. To this end, I thank for their presence and their merits, in no particular order: Soumick, for your moral or intellectual insights, supporting chats and impeccable movie suggestions; Camille, for your support, kindness, and sharing dreams of evasion; Federica, for being a reliable good friend; Pavel & Julian, for sharing their jokes and their office with me; Hana, for the support, fun Bosnian cultural facts, fatty recipes, and chats that lightened my days (but not my belly); Michele, who never refused a cup of coffee with complains and adequately reciprocated all my pranks (and complains); Fatima and Neam, for our supporting chats and tastes of Lebanon; Fardokht, Kavè, MoMo, Andrea and Ilenia, for our many diverse day & night activities, always full of laughter (and perhaps too many beer bottles); Femke & Caroline, for making fun of me just when I needed it the most; Joaquim, for sharing many "it must work at first try" sessions, dealing with ferocious dogs and dog-bones; Gabriele, for our Machiavellian discussions about technology, everyday and future life, or which are the true Italian borders; Rupali, for all the water under the bridge and still sticking around; Negar and Banafshe, for the passionate conversations and the warm introduction to Persian culture. And to countless others who, in their unique ways, sweetened my days and lightened my load. A special thank you to my family for their unwavering support in every way possible. Finally, to my supervisors, whose expertise was invaluable in shaping

the publications that form the basis of this thesis. Without all of you, I might have spent my time sitting in that chair, chin on my wrist, pondering my next move, instead of celebrating the completion of this thesis/journey.

Contents

1	Introduction	15
2	Foundation	17
2.1	The Additive Manufacturing Process	17
2.2	X-ray imaging and Computed Tomography	21
2.3	Deep Learning for semantic segmentation	24
3	Multi-mesh registration from few X-ray projections	27
3.1	Related works	28
3.2	Methods	28
3.3	Experiments	30
3.3.1	Registration with limited scanning system information . . .	31
3.3.2	Pose estimation performance by reducing the number of projections	31
3.3.3	Pose estimation performance by reducing the angular range	31
3.4	Results and discussions	32
3.4.1	Registration with limited scanning system information . . .	32
3.4.2	Pose estimation performance by reducing the number of projections	32
3.4.3	Pose estimation performance by reducing the angular range	40
3.5	Conclusion	44
4	PACS: Projection-driven with Adaptive CADs X-ray Scatter Compensation	45
4.1	Related Works	46
4.2	Materials	47
4.3	Methods	48
4.3.1	Multi-object pose and shape estimation	48
4.3.2	X-ray interaction simulation	49
4.3.3	X-ray scatter compensation	54
4.4	Experiments	54
4.4.1	Spatial frequencies of X-ray scatter on the detector	55

4.4.2	Multi-object pose and shape estimation	55
4.4.3	X-ray scatter compensation	55
4.5	Results and discussions	56
4.5.1	Spatial frequencies of X-ray scatter on the detector	56
4.5.2	Multi-object pose and shape estimation	58
4.5.3	X-ray scatter compensation	60
4.6	Conclusions	70
5	Voxel-wise pore segmentation with 3D patch-based neural models	71
5.1	Materials	73
5.2	Methods	74
5.2.1	Dataset annotation	74
5.2.2	Deep learning models	77
5.2.3	Training	79
5.2.4	Focal Tversky Loss function	79
5.2.5	Post-processing	80
5.3	Experiments	81
5.3.1	Dataset	81
5.3.2	Training	82
5.3.3	Parameter search for the FTL function	82
5.3.4	Cross-validation of performance of the cro:DLDeep Learning (DL) models	82
5.3.5	Cross-validation of performance of post-processed unsupervised models	83
5.3.6	Cross-validation of performance of supervised models re-trained with unsupervised models	83
5.3.7	Model complexity	83
5.3.8	Cross-validation of performance of the best performing model in extreme visual scenarios	83
5.4	Results and discussions	84
5.4.1	Parameter search for the FTL function	84
5.4.2	Cross-validation of performance of the DL models	85
5.4.3	Cross-validation of performance of post-processed unsupervised models	88
5.4.4	Cross-validation of supervised models trained with labels generated by an unsupervised model	92
5.4.5	Model complexity	92

5.4.6	Cross-validation of performance of the best performing model in extreme visual scenarios	94
5.5	Conclusions	95
6	General conclusion & prospects	99
A	Appendices & Supplementary material	101
A.1	Classifier graphs for the voxel-wise segmentation task	101
A.2	Cross-validation graphs for the FTL parameter search per each fold	102
A.3	Random Variable Module: Sigmoid Activation Function	105
A.4	Theoretical white-field image for flat detectors	107
A.5	Supplementary material for Chapter 3	109
B	Scientific contributions	111
C	Abbreviations	115

Summary

Additive Manufacturing (AM), or 3D printing, has revolutionized industrial production by enabling the creation of complex parts and products through the successive layering of materials. This technique offers significant advantages in prototyping and customization, streamlining the transition from digital models to physical objects. However, challenges such as anisotropic properties, surface finish issues, and internal defects necessitate robust quality control measures.

This thesis investigates advanced non-destructive testing methods, focusing on X-ray inspection techniques, to enhance the reliability and quality of AM components. Key contributions include:

- Automatic Simultaneous Multi-Mesh Registration and X-ray system spectral estimation: Introduction of an automated technique for aligning X-ray CT scans with their corresponding CAD models, facilitating precise defect localization and comparison, which can incorporate estimation of the poly-chromatic behaviour of the scanning system.
- Compensation for X-ray Scattering: The development of a novel software method to mitigate X-ray scattering effects, thereby improving image clarity and defect detection accuracy, suitable for radiographic and X-ray Computed Tomography (X-CT) setups.
- 3D Deep Learning Models: Application of state-of-the-art deep learning methods tailored for 3D defect detection in X-CT images, utilising volumetric data in an efficient, 3D patch-wise approach to identify internal flaws.

These approaches, addressing critical issues such as X-ray scattering and beam hardening, aim to jointly improve the image quality and ease the 2D/3D comparison of the manufactured object and its digital model. By cross-validating these techniques on real-world data, through a series of experiments, this research makes a step forward towards ensuring the structural integrity and defectiveness of AM samples. The findings contribute to the broader adoption of X-rays inspection setups in AM industries where safety and reliability of samples are critical, such as aerospace, medical devices, and automotive sectors, and where reduction of operational costs is usually desired.

Samenvatting

Additive Manufacturing (AM), of 3D-printen, heeft de industriële productie gerevolutioneerd door de creatie van complexe onderdelen en producten mogelijk te maken via opeenvolgende lagen materiaal. Deze techniek biedt aanzienlijke voordelen in prototyping en maatwerk, waardoor de overgang van digitale modellen naar fysieke objecten wordt vereenvoudigd. Echter, uitdagingen zoals anisotrope eigenschappen, problemen met de oppervlakteafwerking en interne defecten vereisen robuuste kwaliteitscontrolemaatregelen.

Deze thesis onderzoekt geavanceerde niet-destructieve testmethoden, met de focus op röntgeninspectietechnieken, om de betrouwbaarheid en kwaliteit van AM-componenten te verbeteren. Belangrijke bijdragen omvatten:

- Automatische Gelijktijdige Multi-Mesh Registratie en Spectrale Schatting van het Röntgensysteem: Introductie van een geautomatiseerde techniek voor het uitlijnen van röntgen-CT-scans met hun overeenkomstige CAD-modellen, wat precieze defectlokalisatie en vergelijking mogelijk maakt, waarbij de schatting van het polychromatische gedrag van het scansysteem kan worden opgenomen.
- Compensatie voor Röntgenverstrooiing: De ontwikkeling van een nieuwe softwaremethode om röntgenverstrooiingseffecten te verminderen, waardoor de beeldhelderheid en nauwkeurigheid van defectdetectie verbeteren, geschikt voor radiografische en röntgen Computed Tomography (X-CT) opstellingen.
- 3D Deep Learning Modellen: Toepassing van state-of-the-art deep learning methoden, afgestemd op 3D-defectdetectie in X-CT-beelden, waarbij volumetrische gegevens efficiënt worden gebruikt in een 3D patch-wise benadering om interne fouten te identificeren.

Deze benaderingen, die kritieke kwesties zoals röntgenverstrooiing en verharding van de bundel aanpakken, zijn bedoeld om gezamenlijk de beeldkwaliteit te verbeteren en de 2D/3D-vergelijking van het vervaardigde object en zijn digitaal model te vergemakkelijken. Door deze technieken te cross-valideren op real-world

data, via een reeks experimenten, levert dit onderzoek een stap vooruit in het waarborgen van de structurele integriteit en de foutloosheid van AM-monsters. De bevindingen dragen bij aan de bredere acceptatie van röntgeninspectieopstellingen in AM-industrieën waar veiligheid en betrouwbaarheid van monsters cruciaal zijn, zoals lucht- en ruimtevaart, medische apparaten en de automobielsector, en waar verlaging van operationele kosten meestal gewenst is.

Chapter 1

Introduction

cro:AMAdditive Manufacturing (AM), commonly known as 3D printing, is a new approach to industrial production that enables the creation of parts and products through the successive layering of materials [Gibson2015, Frazier2014]. This technique contrasts with traditional subtractive manufacturing, where material is removed to achieve the desired shape [Bourell2009]. In AM, a digital 3D model designed using cro:CADCComputer-Aided Design (CAD) software is directly translated into a physical object by adding material layer by layer, making the process highly efficient and flexible [Wohlers2011].

Within the framework of Industry 4.0, AM is standing out due to its integration of smart technologies and digital processes in manufacturing [Huang2015]. The ability of AM to quickly prototype and create complex geometries without needing extensive tooling or assembly lines demonstrates the versatility and customization at the heart of Industry 4.0. The straightforward translation of CAD designs into physical items streamlines development, greatly reduces lead times, and promotes innovation and rapid response to market demands.

However, AM has its weaknesses. The layer-by-layer construction can result in anisotropic properties, meaning the material properties may vary in different directions [DebRoy2018]. Surface finish and dimensional accuracy can also be inferior compared to traditional methods, potentially requiring additional post-processing [Herzog2016]. The variability in material properties and the presence of internal defects are critical concerns that necessitate thorough inspection and quality control [Thijs2010], especially in industries such as aerospace, medical devices, and automotive, where the unexpected failure of a part is highly undesirable.

To ensure the reliability and integrity of AM parts, cro:NDTNon-destructive testing (NDT) techniques are commonly employed. These methods allow for the examination of materials and structures without causing damage. Key NDT techniques used in the context of AM include visual inspection, ultrasonic testing, thermography, and X-ray inspection [Leach2020].

X-ray inspection, through *cro:DXR* Digital X-ray radiography (DXR) and *cro:X-CT* X-ray Computed Tomography (X-CT), is widely used in AM due to its ability to reveal internal structures and defects with high resolution. DXR produces a two-dimensional image, capturing variations in material density and thickness. X-ray CT, on the other hand, provides three-dimensional images by compiling multiple radiographic images taken from different angles. This technique allows for detailed analysis of internal features and detection of flaws such as cracks, voids, and inclusions [Seifi2016].

Nevertheless, X-ray imaging is not without its challenges. Two common issues that can hinder defect detection are X-ray scattering and beam hardening [Thompson2020]. X-ray scattering occurs when X-rays deviate from their original path as they pass through a material, leading to image blurring and reduced contrast [Maire2013]. Beam hardening happens when lower energy X-rays are absorbed more than higher energy ones, causing artifacts in the images that can obscure or mimic defects [Smith2017]. Moreover, the lack of detailed information about the X-ray source spectrum and the sensitivity of detectors further complicates the application of poly-chromatic approaches to address beam hardening.

To enhance defect detection in AM parts, comparing X-ray images with their corresponding CAD models can be highly effective [Pavlov2017]. This can be done by directly comparing the surface mesh obtained from X-ray CT scans with the original CAD mesh [Tang2020]. Alternatively, virtual projections of the CAD model can be created and compared with the actual X-ray images [Cloetens2018]. These methods help in identifying deviations and locating defects with greater accuracy [Masood2014].

Within this context, the thesis you are reading delves into the exploration of novel ways for defect detection through automatic multi-mesh registration and scanning system spectral behaviour estimation (Chapter 3), enhancement of X-ray images through X-ray scattering compensation (Chapter 4), and employing the latest deep learning methods tailored efficient 3D segmentation through a patch-wise approach (Chapter 5). By identifying internal flaws that could compromise performance or safety, these techniques help maintain high-quality standards and enhance overall confidence in AM-produced components, while lowering technical and operational requirements.

This thesis explores X-CT image processing and inspection techniques for AM samples. This foundational chapter introduces the AM process, discusses common AM defects, outlines X-CT imaging principles, and delves into the application of deep learning for semantic image segmentation.

2.1 The Additive Manufacturing Process

AM constructs objects by sequentially adding material layer by layer, contrasting sharply with traditional subtractive manufacturing techniques that remove material through cutting, drilling, or milling. The AM process begins with a digital 3D model created using Computer-Aided Design (CAD) software. This model is then sliced into thin cross-sectional layers by the printer's software, and the printer constructs the object by adding material layer upon layer based on these slices [Gibson2010].

Various methods of additive manufacturing are employed to achieve the layering process. Fused Deposition Modeling (FDM) uses a heated thermoplastic filament extruded through a nozzle to build each layer [Crump1992]. SLA (Stereolithography) utilizes a vat of liquid resin that is cured layer by layer using a UV laser [Hull1986]. Selective Laser Sintering (SLS) employs a laser to sinter powdered material, binding it together to form a solid structure [Deckard1990]. Direct Metal Laser Sintering (DMLS) and SLM (Selective Laser Melting) use high-power lasers to fuse metal powders [Kruth1998]. Since SLA and SLM are utilized for the AM samples in this thesis, these processes are explained in further detail.

SLA is one of the earliest and most widely used 3D printing technologies. It employs a vat of liquid photopolymer resin, which is cured layer by layer using a UV laser [Hull1986], as outlined in Fig. 2.1. The process begins with the printer platform being submerged slightly below the surface of the liquid resin.




Figure 2.1: Schematic of the SLA process.

The UV laser traces the first layer of the design onto the surface, polymerizing and solidifying the resin in the desired pattern. After each layer is completed, the platform lowers slightly, and the next layer is traced and cured on top of the previous one. This process is repeated until the entire object is formed, with typical layer thicknesses ranging from 25 to 100 micrometers, allowing for high precision and fine detail [Malek1995].

The advantages of SLA include high precision and smooth surface finishes, making it ideal for applications requiring detailed features and fine tolerances. However, the materials used in SLA, typically photopolymers, can be brittle and may require post-processing steps such as washing to remove excess resin, additional UV curing to ensure full polymerization, and support removal to finalize the part [Gibson2010]. Common defects in SLA include layer delamination, resin pooling, and incomplete curing [Malek1995]. Layer delamination occurs when cured layers fail to bond properly, often due to insufficient laser power or rapid movement of the build platform. Resin pooling happens when excess resin accumulates in certain areas, causing inconsistencies in layer thickness. Incomplete curing can result in soft or uncured spots within the printed object, affecting its mechanical properties and durability. These defects can be identified through visual inspection, mechanical testing, or advanced imaging techniques like X-ray CT and mitigated by optimizing printing parameters and thorough




Figure 2.2: Schematic of the SLM process.

post-processing [Williams2002].

SLM is a highly advanced additive manufacturing technique that uses a high-power laser to fully melt and fuse metallic powders to create solid three-dimensional objects [Kruth1998]. Unlike SLA, which uses liquid resins, SLM works with a variety of metal powders, including stainless steel, titanium, aluminum, and cobalt-chrome alloys. The process begins by spreading a thin layer of metal powder across the build platform, as outlined in Fig. 2.2. The laser then selectively melts the powder according to the CAD data, fusing it to form a solid layer. After each layer is completed, the build platform lowers, and a new layer of powder is spread on top, repeating the process until the entire object is formed [Frazier2014].

SLM is particularly valued for its ability to produce parts with complex geometries and high mechanical properties [Yap2015]. It is extensively used in industries such as aerospace, automotive, and medical, where strong, lightweight, and customized components are essential. Despite its advantages, SLM can be expensive due to the cost of metal powders and the high energy consumption of the lasers. Additionally, the process can induce residual stresses in the parts, which may require heat treatment post-processing [Gibson2010].

Among undesired inner porosity arising from the SLM technique, cro:LoFLack-of-fusion (LoF) and cro:KHKeyhole (KH) pores are known to have the highest share in the total porosity. LoF pores arise when there is insufficient melting and bonding between adjacent powder layers or tracks [Yap2015], as shown in Fig. 2.3.a. This defect typically results from inadequate laser energy, excessive scanning speed, or incorrect layer thickness. When the energy input is too low,

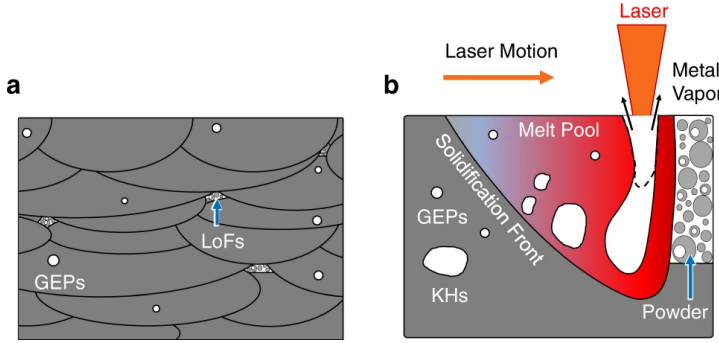


Figure 2.3: Formation of (a) LoFs and (b) KH pores. From Ref. [poudel2022feature], used under Creative Commons CC-BY license.

the powder particles do not fully melt, leading to incomplete fusion between layers. This creates irregular voids within the material, compromising structural integrity and potentially leading to premature failure under load [King2014]. LoF pores are often irregular in shape and size and can be distributed throughout the part, making them challenging to detect and characterize [King2014]. KH pores form during the SLM process when excessive laser energy creates deep, narrow cavities in the melt pool [Yap2015], as shown in Fig. 2.3.b. This occurs when the laser power is too high or the scanning speed is too slow, leading to a keyhole-mode melting. In this mode, the laser penetrates deeply into the powder bed, creating a vapor depression that can collapse, trapping gas and forming elongated, keyhole-shaped pores [King2014]. These pores act as stress concentrators, significantly reducing the material's mechanical properties, such as tensile strength and fatigue resistance [King2014].

The materials used in additive manufacturing are diverse, including polymers, metals, ceramics, and composites. Common polymers such as PLA and ABS are favored for their ease of use and versatility. Engineering-grade plastics like nylon and polycarbonate serve more demanding applications [Gibson2010]. Metals like titanium, aluminum, and stainless steel are utilized in industries such as aerospace and medical for their strength and durability. Ceramics are printed for applications requiring high heat resistance or electrical insulation [Gibson2010]. Composite materials, which combine polymers with fibers like carbon or glass, provide enhanced mechanical properties for high-performance applications [Gibson2010].

A significant challenge in additive manufacturing is the occurrence of defects. These defects act as stress concentrators, reducing mechanical properties such as tensile strength and fatigue resistance [King2014]. Accurate detection and

characterization of these defects are crucial for ensuring the reliability of AM components [Yap2015]. NDT methods are essential across various industries for evaluating material properties without causing damage. These methods ensure product integrity and reliability, detect defects, and prevent failures throughout a product's lifecycle. Surface methods like visual inspection or dye penetrant testing are limited to surface-breaking flaws and cannot provide information about subsurface conditions [Shull2002]. Internal NDT methods, such as X-ray testing, are crucial for identifying internal defects that are not visible from the surface.

2.2 X-ray imaging and Computed Tomography

X-rays imaging is a NDT method essential for evaluating material properties without causing damage. X-rays penetrate materials to varying degrees, making them particularly useful for identifying internal defects such as cracks, voids, and inclusions. The generation of X-rays typically involves X-ray tubes, where high-energy electrons are accelerated and directed toward a metal target, usually tungsten. When these high-speed electrons collide with the target material, their sudden deceleration causes the emission of X-rays. This process produces characteristic X-rays, specific to the target material, and Bremsstrahlung (braking radiation), a broad spectrum of X-rays generated by the deceleration of electrons.

Once generated, X-rays pass through the object being inspected and are absorbed by a detector. In industrial environment, photon-integrating detectors are the most common choice for their contained cost. These detectors measure the intensity of the transmitted X-rays and convert this information into an image. Areas where X-rays are absorbed more strongly appear darker on the resulting radiograph, indicating denser or thicker regions, while areas where X-rays pass through more easily appear lighter, revealing internal structures. As an example, a bevel gear is shown in Fig. 2.4, where it is noticeable to peripheral part of the gear are less attenuating the X-rays than the central, bulkier part.

In Fig.2.4, careful observation reveals non-uniform grey values in rectangular areas where the imaged objects are not absorbing X-rays. These irregular patterns are caused by variations in the gain and electrical bias of the detector's read-out circuitry. To correct these inconsistencies, a log-correction must be applied, utilizing both white-field and dark-field images from the X-ray scanning system, as illustrated in Fig. 2.5.

The white-field image is captured without any object in the X-ray beam path, representing the detector's response to uniform X-ray exposure. Conversely, the dark-field image is obtained with the X-ray source turned off, capturing the detector's inherent noise and biases. By applying a log-correction using these




Figure 2.4: Few radiographs of a bevel gear.




Figure 2.5: An example of (a) white-field image and (b) dark-field image.




Figure 2.6: Main interaction mechanisms of X-rays with matter, which are associated with image formation. From Ref. [seibert2005x] © SNMML.

reference images, we can normalize the radiographs, thereby reducing the impact of detector non-uniformities and enhancing image quality [Hsieh2003].

X-ray imaging quality can be significantly affected by various factors related to the interaction mechanisms between X-rays and matter, such as scattering and beam hardening. These phenomena are rooted in the fundamental interactions of X-rays with the atoms in the material they penetrate, primarily involving the photoelectric effect, Rayleigh scattering, and Compton scattering, as schematically shown in Fig. 2.6.

Scattering occurs through two primary mechanisms: Rayleigh scattering and Compton scattering. Rayleigh scattering, also known as coherent scattering, involves the deflection of X-rays without a change in their energy. This type of scattering is more prevalent at lower X-ray energies and contributes to a reduction in image contrast and detail due to the scattered X-rays deviating from their original paths [Seibert2005]. Compton scattering, or incoherent scattering, involves the interaction of X-rays with loosely bound electrons in the material. This interaction results in the X-rays being deflected and losing energy in the process [Hsieh2003]. Compton scattering is more significant at higher X-ray energies and, as for Rayleigh scattering, contributes to image degradation.

Beam hardening is primarily caused by the photoelectric effect, which is highly

dependent on the energy of the X-rays and the atomic number of the material. The photoelectric effect occurs when X-ray photons are absorbed by atoms, causing the ejection of inner-shell electrons [Hsieh2003]. This effect is more pronounced at lower X-ray energies and with materials of higher atomic numbers. As the X-ray beam passes through the material, the lower-energy X-rays are absorbed more readily than the higher-energy ones, resulting in a progressive increase in the average energy (or "hardening") of the X-ray beam. This preferential absorption of lower-energy photons leads to artifacts in the X-ray images, such as false density variations or streaks, which can compromise the diagnostic quality of the images.

X-CT builds on traditional X-ray methods by constructing a three-dimensional representation of an object from multiple two-dimensional X-ray images taken at different angles. The X-CT imaging process involves rotating the X-ray source and detectors around the object, capturing numerous radiographs, and using computational algorithms to reconstruct a detailed 3D image of its internal structure, as in Fig. 2.7. X-CT can also be done in a discrete way, which incorporate prior knowledge about the limited number of materials within the object to enhance the reconstruction process [DART]. By recognising that many objects are composed of a few distinct materials, discrete tomography methods can improve the overall quality of the reconstructed images, as in Fig. 2.8. Both traditional and discrete X-CT offers unparalleled insights into the internal features of objects. X-CT is particularly advantageous for its high-resolution imaging capabilities, which allow for the detection of small, subsurface defects that might be missed by other NDT methods.

2.3 Deep Learning for semantic segmentation

DL has emerged as a powerful tool for image processing, particularly in the domain of semantic segmentation, which involves partitioning images into meaningful regions. Segmentation is crucial in X-CT imaging for isolating specific structures or areas of interest within the reconstructed volume, such as detecting defective zones in AM prints. DL models can be trained for segmentation using various training paradigms, broadly categorized into supervised and unsupervised methods, as illustrated in Fig. 2.9.

Supervised training relies on labeled datasets to train models, where regions of interest are already identified and annotated. During this process, the model learns to recognize similar patterns in new images, providing accurate segmentation based on the examples it has seen during training.

Unsupervised training does not rely on labeled data. Neural models in this ap-




Figure 2.7: Three orthogonal cross-sections of a bevel gear after X-CT reconstruction of the radiographs. The planes of the cross-sections are indicated by the colored lines in each image: (a) referenced in green, (b) in blue, and (c) in red. This reconstruction accounts for beam-hardening and X-ray scatter.




Figure 2.8: Cross-sections of a discrete X-CT reconstruction of the bevel gear shown in Fig. 2.7. Compared to Fig. 2.7, it is noticeable the limited colormap with just two colour values, black (air) and white (gear material).

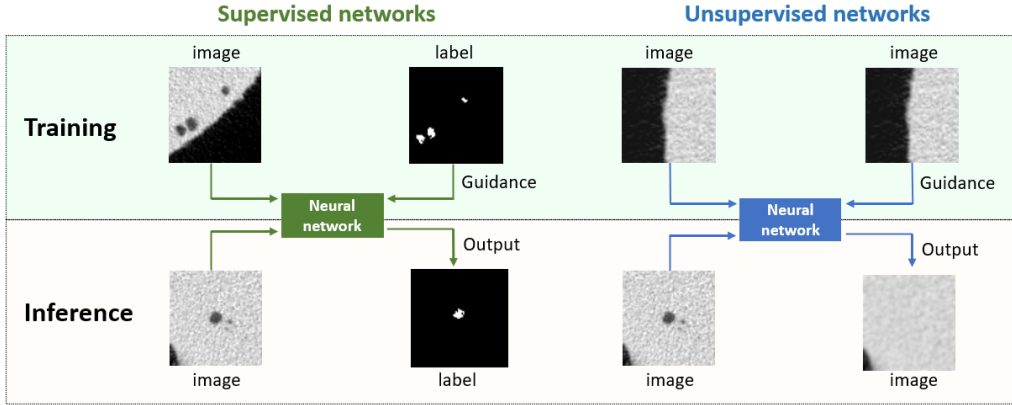


Figure 2.9: Two training paradigms for DL methods in defect segmentation. Supervised networks use labeled data for direct inference of defects in X-CT images. Unsupervised networks aim to generate a defect-free version of the input image, with anomalies identified by differences between the input and output images.

proach are designed to reproduce the input image as the output, learning the statistical characteristics of defect-free images. During inference, the networks cannot reproduce defects, and the difference between input and output images highlights the defective zones. This method identifies patterns and structures without prior knowledge, useful for datasets lacking labeled examples [bozorgtabar2020salad, kim2022virtual]. Autoencoders, variational autoencoders (VAEs), and generative adversarial networks (GANs) are commonly used for unsupervised anomaly detection, trained to reconstruct normal data with anomalies detected through high reconstruction errors [baur2020autoencoders].

The quality of segmented images depends on the initial X-CT reconstruction, scan resolution, and DL model robustness. Higher resolution scans provide more detailed segmentation, while advanced DL models trained on diverse datasets enhance performance by capturing finer details and reducing errors. Despite the computational intensity and the need for large datasets, volumetric information within the X-CT can be fruitfully used through 3D patch-based approaches, as recently proposed, which drastically reduce the memory footprint of DL models [perez2021torchio]. Compared to traditional machine learning methods, DL approaches offer significant improvements in the accuracy and efficiency of image segmentation and defect detection, making them invaluable nowadays in the analysis of complex X-CT data.

Multi-mesh registration from few X-ray projections

One effective strategy for defect detection in AM is the comparison of measured X-ray radiographs with simulated radiographs derived from digital CAD models. This method necessitates accurate simulation of X-ray projections that reflect the multi-chromatic nature of the scanning system and the scanned objects. Key to this approach is the precise alignment or registration of the surface mesh, contained within the CAD model, with the X-ray radiographs, a process complicated by the need for detailed knowledge of the object's properties and the spectral information of the X-ray scanning system.

The challenge lies in the availability of these details. Typically, end-users lack comprehensive information about the object's X-ray attenuation coefficients of the scanned objects and the spectral behavior of the scanning system. Addressing these gaps often involves complex estimations or the use of supplementary systems, posing significant hurdles for practical implementation.

This chapter introduces an innovative methodology utilizing the CAD-ASTRA toolbox [paramonov2024cad], which enables efficient X-ray mesh projection and facilitates multi-object pose estimation with a minimal number of radiographs. CAD-ASTRA's capability to simulate projections from complex geometries makes it particularly advantageous in industrial settings, where spatial constraints on X-ray source and detector placement are common.

Leveraging CAD-ASTRA, which differential implementation has been developed specifically for this task, combined with PyTorch [paszke2019pytorch] for optimisation, this chapter focuses on registering multiple object meshes using sparse X-ray projections. By employing a differentiable projector within CAD-ASTRA, the process integrates seamlessly with gradient-based optimisers, allowing for simultaneous determination of the 3D position, orientation, and spectral characteristics the scanning system. This approach not only enhances the accu-

racy of defect localization but also streamlines the inspection process, aligning with the operational needs of modern manufacturing environments.

Through this methodology, the author aim to address the inherent challenges in mesh registration for X-ray radiograph analysis, providing a robust basis for an efficient framework for defect detection in AM components.

3.1 Related works

The pose estimation of objects within industrial settings traditionally relies on a comprehensive set of projections, through X-CT reconstructions. This technique allows for the inference of pose through the registration of CAD models with point clouds [**kim2021robust**] or extracted meshes [**sukowski2022automated**, **iuso2021cad**]. However, the feasibility and desirability of 3D X-CT images may be limited in certain scenarios. In recent developments, efforts have been directed towards achieving pose estimation based on a low number of X-ray projections. A notable industrial approach utilizing deep learning was proposed by Presenti et al. [**presenti2023fast**], demonstrating pose estimation efficacy with as few as one projection. While such approaches exhibit promising performance in controlled environments, challenges arise from the specialized training procedures and the inherent black-box characteristics of many deep learning methods, impeding widespread adoption. Another recent contribution explored pose estimation from X-ray projections by employing a CAD model and matching 2D-3D image features through mesh projections [**tanmatching**]. The proposed approach shares similarities with this method, with the additional benefit of overcoming challenges associated with complex geometries and overlapping meshes, where image features might be hardly discernible. In this study, a multi-mesh registration from X-ray projections is presented, using a mesh projector implemented as a differential program, elucidating the capabilities of the proposed method.

3.2 Methods

CAD-ASTRA utilizes watertight triangular surface meshes¹ to represent homogeneous volumes as enclosed entities. These meshes incorporate information about the source and detector positions and orientations to simulate X-ray acquisition, considering object attenuation. The initial guess for mesh position and orientation, represented by \mathbf{p} , serves as a starting point, with source and detector positions assumed from scan metadata.

¹A watertight surface mesh is a closed surface mesh free from self-intersections and overlaps.

Each mesh $\mathbf{m} \in \mathbb{R}^{N \times 3}$ comprising N vertices undergoes transformation via a roto-translation operator $\mathbf{A}_{\boldsymbol{\theta}_k} : \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}^{N \times 3}$. Here, the pose vector $\boldsymbol{\theta} = [\alpha, \beta, \gamma, t_x, t_y, t_z]$ defines the 3D mesh's position and orientation, in terms of Euler Z-Y-Z angles (α, β, γ , respectively) and translation along the x, y , and z axis (t_x, t_y, t_z), respectively.

The projection operator Proj_d maps mesh vertices to W projection images, each image consisting of M pixels:

$$\text{Proj}_d : \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}^{M \times W} \quad , \quad (3.1)$$

with d representing metadata linked to the meshes.

Estimation of the pose \mathbf{p}_k , for the k -th of the K meshes composing the scene, is achieved through linear programming, aiming to minimize the sum of squared differences in the projection space between measured projections $\mathbf{P} \in \mathbb{R}^{M \times W}$ (adjusted for dark and bright fields) and projections simulated by Proj_d :

$$\tilde{\boldsymbol{\theta}}_1, \dots, \tilde{\boldsymbol{\theta}}_K = \arg \min_{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K} \left\| \mathbf{P} - \sum_{k=1}^K \text{Proj}_d[\mathbf{A}_{\boldsymbol{\theta}_k}[\mathbf{m}_k]] \right\|_2^2 \quad . \quad (3.2)$$

Any objective function, as the one denoted in 3.2, can be conceptually decomposed as a concatenation of an error function g on the simulated projection function Proj_d , depending on a roto-translation through the operator \mathbf{A} by a pose vector $\boldsymbol{\theta}_k$. This leads to a compact formulation as $f_d(x) = g(\text{Proj}_d(\mathbf{A}_{\boldsymbol{\theta}_k}(x)))$. Optimising the linear programming problem 3.2 through analytical gradient methods requires knowledge of all the Jacobians involved in the Jacobian of the composed function f_d . The analytical gradient is then defined as:

$$\nabla f_d(\mathbf{m}) = \text{Jac}_T^T \text{Jac}_{\text{Proj}}^T \nabla g(\text{Proj}_d(\mathbf{m})) \quad , \quad (3.3)$$

where direct access to the Jacobian-vector product of Jac_{Proj} is provided by CAD-ASTRA, whilst the other Jacobian-vector products are computed through autograd patterns [bradbury2018jax, paszke2019pytorch]. Optimising 3.2 with analytical gradient-based methods, offers computational efficiency, particularly in scenarios with detectors with high spatial-resolving capability.

As real X-ray projections in industrial scenarios are commonly employing a poly-chromatic X-ray source, accurate polychromatic forward model is needed to reduce the likelihood of undesired local minima in 3.2. Therefore, the Proj_d operator is substituted by the poly-chromatic operator PolyProj_d , where:

$$\text{PolyProj}_d = \sum_{e=1}^E s_e D_e \delta_e \exp\left(-\sum_{k=1}^K \mu_{e,k} l_k\right) \quad . \quad (3.4)$$

In the above formula, the X-ray attenuation of photons is based on the energy-dependent intensity of photons s_e emitted by the X-ray source, the detector response D_e and the energy δ_e contained in the energy bin e . The product $\Omega_e = s_e D_e \delta_e$ defines the weight of each energy bin e , contributing to the total spectral behavior $\mathbf{\Omega} = [\Omega_1, \dots, \Omega_E]$, with E denoting the number of energy bins. For each mesh, the attenuation is measured by its spectral linear attenuation $\mu_{e,k}$ and the path-length l_k crossed by a geometrical ray pointing at a detector pixel.

If such a model is implemented in frameworks that leverage automatic differentiation patterns, such as PyTorch, it is possible to optimise even the new linear programming problem using 3.4 without expensive numerical approximations. As the poly-chromatic characteristics of the X-ray source and detector, as well as the spectral linear attenuation of each scanned material, may be unknown, these parameters can be jointly optimized with a proper cost function. For experiments with no prior knowledge on the exact scanning system spectral characteristics, the linear programming problem in 3.2 is enriched with regularization on the first derivative of the system spectral behavior, under the assumption of smooth spectral behavior:

$$\tilde{\theta}_1, \dots, \tilde{\theta}_K = \arg \min_{\theta_1, \dots, \theta_K, \mathbf{\Omega}} \left\| \mathbf{P} - \sum_{k=1}^K \text{PolyProj}_d [\mathbf{A}_{\theta_k} [\mathbf{m}_k]] \right\|_2^2 + \|\nabla_e \mathbf{\Omega}\|_2^2 \quad . \quad (3.5)$$

Pose Refinement through Re-Iteration: To mitigate cases where local minima occur due to symmetry in the object with respects to the vertical axis, an additional step of re-iteration may be introduced (with results presented separately). After the initial registration, the algorithm systematically rotates the objects around their symmetry axis and re-executes the registration procedure. This process helps in overcoming challenges posed by symmetry, enhancing the robustness of the pose estimation, especially in scenarios with highly symmetric objects.

3.3 Experiments

In this section, the experiments conducted on three distinct scanned objects are presented using one or more supporting scanning elements. The scanned objects include an aluminum step-wedge, and two samples printed through AM, which are a cro:SS316LStainless steel 316L (SS316L) cantilever and a SS316L cylinder-like object. The supporting elements, constructed from cro:PA12Polyamide (PA12), consist of cylinders with a height of 1 cm and diameters ranging from 3 to 5

cm, as well as a hollow shaft with a 0.5 cm diameter. The FleXCT scanning system [**FlexCT**] was utilised, with different source-to-object and source-to-detector distances, kVp and pre-filtering settings for each scan.

The objective of these experiments was to showcase two different applications of the registration technique, one involving limited information about the scanning system and the other utilizing limited projection information.

3.3.1 Registration with limited scanning system information

In this experiment, limited scanning system information refers to uncertainty regarding the spectral behavior of the scanning system and the objects' poses. The registration process addresses a challenging scenario by iteratively estimating the spectral behavior and poses through the solution of the linear programming problem in 3.5. The initial pose of the objects and supporting elements is set as a shifted and rotated configuration from a vertically aligned state. The registration is performed using 100 projections acquired in a circular trajectory around the object.

3.3.2 Pose estimation performance by reducing the number of projections

In this experiment, the linear programming problem from 3.2 based solely on projection error is employed. The initial pose of the scene's objects is realistically estimated with the assumption that the objects and supporting elements are vertically aligned. The registration is conducted by reducing the number of projections from 100 to 10 (100, 50, 10), all acquired in a circular trajectory around the isocentre.

3.3.3 Pose estimation performance by reducing the angular range

Similar to the preceding experiment, the linear programming problem presented in 3.2 is employed, assuming that the objects are vertically aligned. In this experiment, only two projections are utilized, chosen from a complete circular scan around the isocenter. The angle between these two projections is systematically decreased (90 degrees, 50 degrees, 10 degrees) to assess its influence on pose estimation stability.

To ensure the accuracy and repeatability of this experiment, the mesh registration is iterated five times, initiating from different projections for each scene and for each angular case. This repetition aids in investigating the consistency and reliability of the obtained results.

3.4 Results and discussions

In this section, the outcomes of the experiments described in section 3.3 are shown through 3D rendering of the scene of the initial and final objects spatial configuration (through Mayavi libraries [[d5725237](#)]) and figures of the residual errors (i.e. difference between projections and simulated projections). For the latter, the projection-wise cross-root mean squared error (RMSE) is computed and shown to ease the comparisons. Animations showing the evolution of the residual errors throughout the optimisation steps are available online as supporting media (<https://osf.io/da6p3/>).

3.4.1 Registration with limited scanning system information

A scaled 3D rendering, showing both the initial and final poses of the objects from Experiment 3.3.1, is presented in Figs. 3.1, 3.2, and 3.3. These figures include the X-ray source and detector, forming a digital twin of the actual scanning setup.

To assess the accuracy of the registration results, attention is directed to the residual images (Fig. 3.4). These images display the residuals for one of the 100 projections utilized in this experiment. During the registration procedure, the position of each mesh in space is adjusted to minimize the residual. Discrepancies that persist between the real and simulated projections may stem from object deformations, especially noticeable in the case of CAD printed samples, uncertainties in the actual chemical composition of samples, and additional physical effects not simulated, such as X-ray scattering.

3.4.2 Pose estimation performance by reducing the number of projections

The results of the registrations are depicted in Figs. 3.5, 3.6, and 3.7, showcasing initial residual errors and final errors (a-c) for three scanning scenarios with 20, 10, and 5 projections. Timings for these registrations are presented in Table 3.1.

For the cantilever, reducing the number of projections does not seem to limit the quality of the registration. However, for the cylinder scene, a mismatch in the identification of teeth in the bottom part of the cylinder indicates convergence to a minimum different from the one identified in the previous experiment (Fig. 3.4). To address this, a few more iterations of the registration algorithm, incorporating a starting rotational offset against the vertical axis, successfully mitigate issues arising from the high symmetry of the cylinder. This refinement leads to more accurate results, as depicted in Fig. 3.6 (d-f).




Figure 3.1: Graphical rendering of the X-ray setup of one projection, showing the initial (left) and final (right) pose of the cantilever and its supporting element.




Figure 3.2: Graphical rendering of the X-ray setup of one projection, showing the initial (left) and final (right) pose of the cylinder and its supporting element.




Figure 3.3: Graphical rendering of the X-ray setup of one projection, showing the initial (left) and final (right) pose of the stepwedge and its supporting elements.




Figure 3.4: An X-ray projection residual showing the initial pose (top row) of the cantilever (left), the cylinder (center) and stepwedge (right) with their supporting element, and the final pose (bottom row) of the objects after the registration procedure.

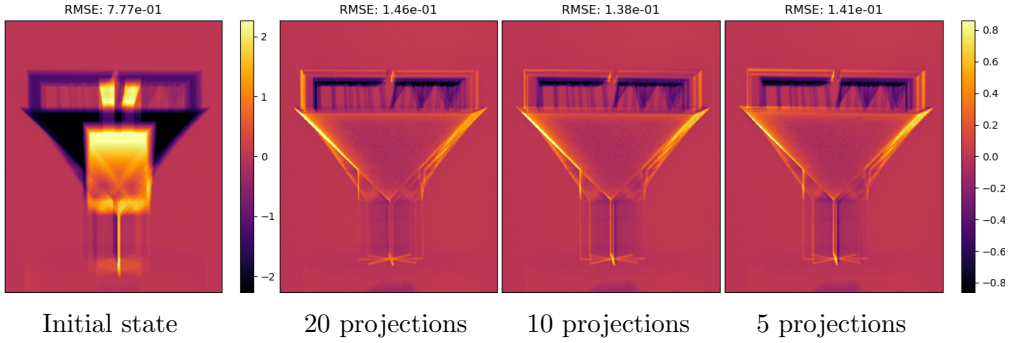


Figure 3.5: A X-ray projection residual showing the initial pose of the cantilever and its supporting element aligned along the vertical axis, and the final pose of the objects after registration using 20 projections, 10 and 5.

Scene	Cantilever			Cylinder			Stepwedge		
# of projections	20	10	5	20	10	5	20	10	5
Time (s)	326	185	113	263	157	138	312	225	216

Table 3.1: Pose estimation timings for each scene, varying the number of projections.

In contrast, the registration of the stepwedge exhibits stable behavior, except for the most challenging case with only 5 projections. In this instance, the supporting straw-like object demonstrates a different pose convergence. This behavior is attributed to the near transparency of the thin plastic straw to the majority of X-ray photons produced by a 230 kVp X-ray source. The attenuation values are comparable to flat-field fluctuations observed during the scans, contributing to the pose convergence variation.

To inspect the stability of the estimated pose as a function of the number of projections, a further analysis is conducted by repeating this last scenario of the stepwedge 5 times, starting from different sets of projections. The results are summarised in Table 3.2, which shows the angle of rotation against the estimated rotational axis in relation to the number of projections. The results consistently indicate that lowering the number of projections increases the likelihood of ending up in undesired minima for the estimated pose parameters.




Figure 3.6: A X-ray projection residual showing the initial pose of the cylinder and its supporting element aligned along the vertical axis, and the final pose of the objects after registration using 20 projections (a), 10 (b) and 5 (c). Systematic rotation of the objects' poses around their vertical axes allows for mitigating the risk of local minima, resulting in more favorable registration outcomes with 20 projections (d), 10 (e), and 5 (f) compared to the configurations in (a-c).




Figure 3.7: A X-ray projection residual showing the initial pose of the stepwedge and its supporting elements aligned along the vertical axis, and the final pose of the objects after registration using 20 projections, 10 and 5.

	20 projections	10 projections	5 projections
Rot. angle (deg)	179.96 ± 0.01	179.94 ± 0.72	177.13 ± 1.28

Table 3.2: Pose estimation results for the scene with the stepwedge, presented as the rotation angle of the stepwedge relative to an estimated rotation axis. The results are provided as the average and standard error across 5 repetitions, starting from different sets of projections.

Scene	Cantilever	Cylinder	Stepwedge
Time (s)	258.0 ± 4.6	382.7 ± 3.2	445.3 ± 6.5

Table 3.3: Pose estimation timings for each scene, by using 2 projections. Results are given as average and standard error of the mean, across all the 5 repetitions.

3.4.3 Pose estimation performance by reducing the angular range

Similar to the preceding experiments, the results, presented as projection residuals, are depicted in Figs. 3.8, 3.9, and 3.10. Timings for these registrations are presented in Tab. 3.3. Given the significant limitation of projective information in this scenario, results are showcased for both the conventional application of the algorithm and a re-iteration of the registration procedure.

In the scanning scenarios involving the cylinder and stepwedge, the results exhibit consistent behavior with the findings of previous experiments. However, in the case of the cantilever, the more challenging registration scenario results in an unrealistic positioning of the main object, noticeable for the case with a angular distance of 50 deg. The difficulty arises from a more impervious solution space, making it easier to fall into local minima, as demonstrated in the more challenging case with 10° in Fig. 3.8.c. Again, re-iteration of the registration algorithm leads to more accurate pose estimation, as graphically shown in the bottom row of Fig. 3.8.

The accuracy and stability of the registration procedure, including the re-iteration procedure, for all scenes in this experiment are extensively reported in the supplementary material (pages 109 and 110), in terms of the average and standard error of the estimated pose parameters for each object. In this analysis, the registration runs 5 times with different pairs of angles. The results indicate that the estimated position is relatively stable throughout the repetitions, and the error of the pose parameters is relatively low. The maximum deviations are recorded for the stepwedge, as its distance from the source (500 mm) is significantly higher than in the other two cases (cantilever 86.68 mm, cylinder 43.33 mm). The rotational angle shows a standard error of 1.76° across the repetitions, while its translation is determined with an error of 2.41 mm. Exceptions are observed for supporting elements, as they appear in their CAD model as perfectly symmetrical around their vertical axis. As also the surface mesh of the cantilever is perfectly symmetric to one of its intersecting planes, the rotation-related parameters have higher error due to the ambiguity arising from its symmetry.

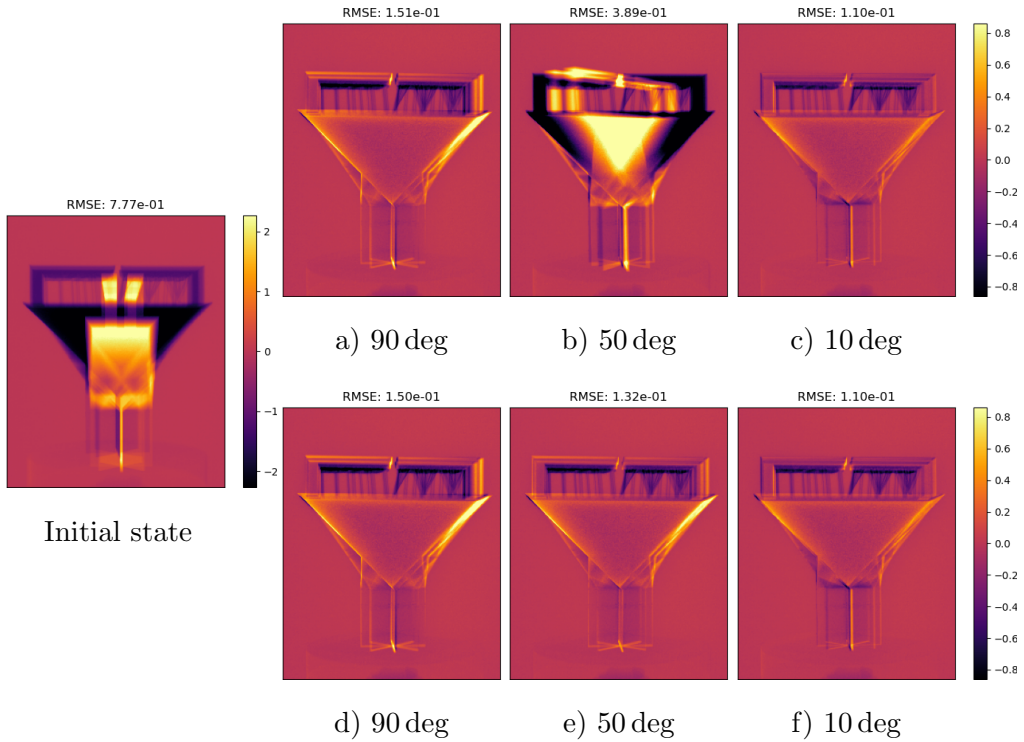


Figure 3.8: A X-ray projection residual showing the initial pose of the cantilever and its supporting element aligned along the vertical axis (left side). The final poses of the objects are presented after executing the registration procedure using 2 projections with angular ranges of 90 deg (a), 50 deg (b), and 10 deg (c). Systematic rotation of the objects' poses around their vertical axes allows for mitigating the risk of local minima, resulting in more favorable registration outcomes with different angular ranges 90 deg (d), 50 deg (e), and 10 deg (f) compared to the configurations in (a-c).




Figure 3.9: A X-ray projection residual showing the initial pose of the cylinder and its supporting element aligned along the vertical axis (left side). The final poses of the objects are presented after re-executing the registration procedure using 2 projections with angular ranges of 90 deg (a), 50 deg (b), and 10 deg (c). Systematic rotation of the objects' poses around their vertical axes allows for mitigating the risk of local minima, which resulted in different outcomes for the angular ranges 90 deg (d), 50 deg (e) and 10 deg (f).

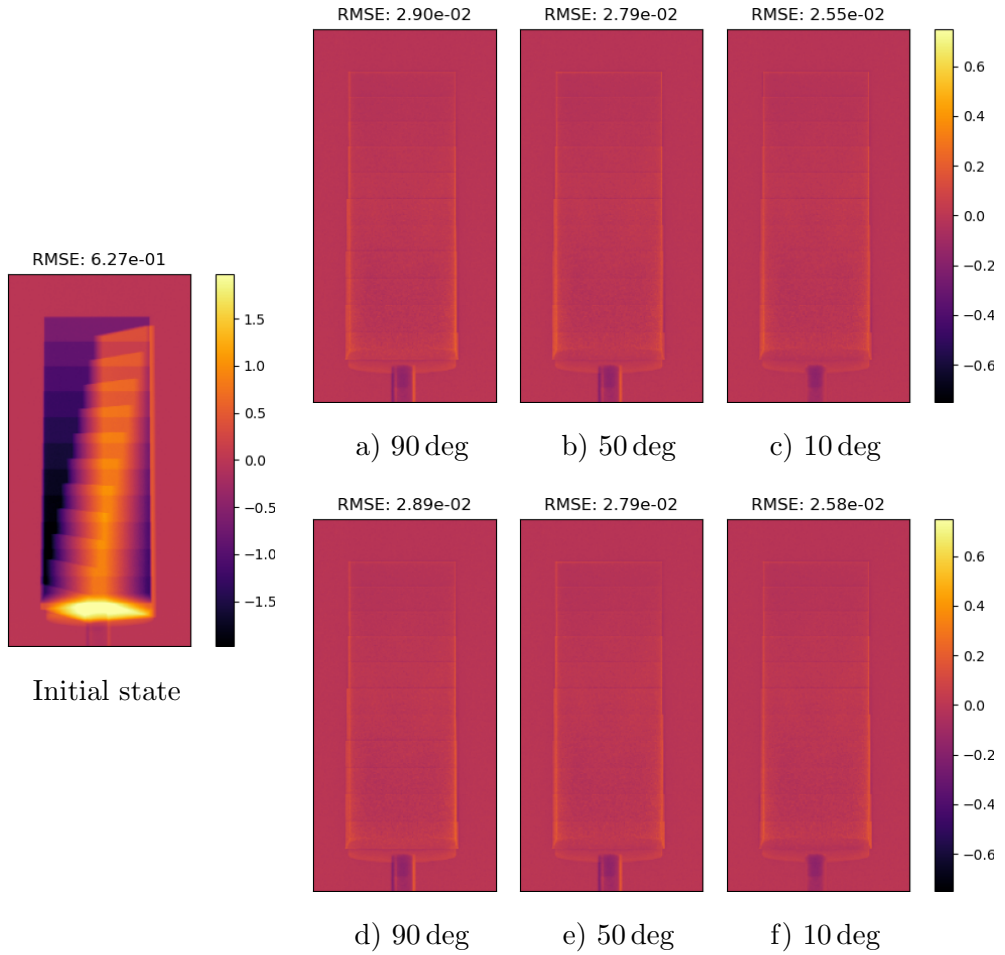


Figure 3.10: A X-ray projection residual showing the initial pose of the stepwedge and its supporting elements aligned along the vertical axis (a). The final poses of the objects are presented after re-executing the registration procedure using 2 projections with angular ranges of 90 deg (b), 50 deg (c), and 10 deg (d). Systematic rotation of the objects' poses around their vertical axes allows for mitigating the risk of local minima, resulting in more favorable registration outcomes with different angular ranges 90 deg (e), 50 deg (f), and 10 deg (g) compared to the configurations in (b-d).

3.5 Conclusion

In conclusion, this novel approach to 3D mesh registration in few-view industrial X-ray imaging, utilizing an X-ray mesh projector with compatibility for the ASTRA toolbox and auto-differentiation libraries like PyTorch, emerges as a resource-efficient alternative. The method, utilising a X-ray mesh projector implemented as differential program specifically for this task, has demonstrated efficacy in achieving 3D multi-mesh registration in multiple X-ray scanning scenarios.

The experiments revealed the robustness of this approach in simultaneous multiple object registration, even under poly-chromatic conditions with limited knowledge about the scanning system’s spectral characteristics or sparsity in projection domain. When using 2 projections, the highest error on the rotational angle was up to 1.76° and 2.41 mm on the translation, for the case of the stepwedge with source-to-object distance of 500 mm. However, challenges surfaced, particularly when objects had a high degree of symmetry or in case projective model inaccuracies were present. Addressing these challenges is crucial for enhancing the applicability and accuracy of the proposed methodology.

Despite identified challenges, this method highlights resource efficiency, eliminating the need for resource-intensive X-CT reconstruction allowing registration even in a fixed multi-head X-ray radiography scanning system. This study marks a significant advancement, showcasing the practicality and efficiency of the proposed methodology. As a future prospect, the method’s adaptability opens possibility to proceed in mesh deformation estimation from X-ray projections. This potential extends the utility of the approach presented here, making it a valuable candidate for enhancing industrial inspection workflows.

PACS: Projection-driven with Adaptive CADs X-ray Scatter Compensation

Ensuring high-quality imaging is of paramount interest during the evaluation of AM components. In X-ray imaging, however, image quality is often compromised by artifacts such as X-ray scattering and beam hardening. These artifacts can obscure defects or mimic their presence, leading to inaccurate assessments.

Among these issues, X-ray scatter presents a significant challenge. As X-rays pass through an object, they can deviate from their original path. This scattering effect is influenced by the material properties and geometry of the scanned object, complicating the accurate detection and characterization of defects. Scatter artifacts cause image blurring, reduced contrast, and streaking artifacts, particularly with high-absorbing objects. Effective scatter compensation methods are essential to enhance image clarity and reliability.

To address this challenge, this chapter introduces cro:PACSProjection-driven Adaptive CADs X-ray Scatter compensation (PACS), which leverages CAD models to simulate and correct for scatter effects, thus enhancing defect detection capabilities without the need for prior CT reconstructions or extensive deep learning models. PACS eliminates the need for deep learning training and geometry-dependent scatter library calculations. Additionally, to expedite simulations, acceleration strategies are described and utilised for downsizing the simulation problem.

The effectiveness of PACS in improving the reliability and accuracy of X-ray inspections for AM components is demonstrated across various scenarios, including poly-chromatic radiography-based analysis, discrete poly-chromatic X-CT reconstruction (based on DART [six2019poly]), and conventional poly-chromatic X-CT reconstructions. Its performance is compared against established scatter compensation methods, highlighting PACS's advantages in practical applications.

4.1 Related Works

The quest for mitigating the effects of X-ray scattering has led to the development of various hardware- and software-based solutions. Hardware approaches include adjustments in detector positioning [**sorenson1985scatter**], the use of bow-tie filters [**liu2014dynamic**], anti-scatter grids [**vogtmeier2008two**], and beam-stop/beam-hole arrangements [**schorner2011comparison**]. While these methods can effectively reduce scattered radiation, they often require specific modifications to the scanning setup, limiting their applicability and flexibility.

In contrast, software-based methods provide versatile alternatives for scatter compensation. Scatter-deconvolution techniques, for example, model the X-ray source as an array of pencil beams, each contributing to the scattered radiation [**rew:firstPaperScatterKernel**]. These methods estimate scatter-kernel parameters based on the known composition of the object, allowing for the subsequent deconvolution of scattered radiation in the projection domain. Various enhancements to this approach have been proposed, including asymmetry considerations [**rew:scatter:asymmetrical-kernels**], attenuation-dependent kernels [**rew:scatter:adaptive-kernels**], accurate scatter models [**bhatia2016separable**], and energy-dependent kernels [**rew:scatter:energy-kernels**].

Deep learning methods have also shown promise in scatter compensation, particularly through the use of U-Net models [**rew:scatter:DL-unet1**, **rew:scatter:DL-unet2**, **rew:scatter:DL-unet3**, **rew:scatter:DL-unet4**]. Despite their potential, these approaches often lack explainability and trustworthiness. Recent advancements aim to address these issues by incorporating more transparent models. For instance, Roser et al. utilize splines to model and suppress low-frequency scatter components, with parameters controlled by a neural network [**roser2020deep**], while Iskender et al. propose a neural model that simulates X-ray propagation and scattering across perpendicular planes [**rew:DL:physicsbased**].

Heuristic techniques leveraging prior knowledge of specific X-CT applications, such as object symmetries [**rew:scatter:heuristic-breast**] or approximating scatter as blurred projections [**rew:scatter:heuristic**], offer targeted solutions for specific artifact types. Analytical methods capable of modeling first-order scattered radiation in homogeneous or composite objects [**rew:scatter:analytical**], as well as simulation-based approaches using Monte Carlo or deterministic methods for voxelized volumes [**maslowski2018acuros**, **rew:scatter:sim-few-proj2**] or surface meshes [**xia2019scatter**, **iuso2021cad**], provide further avenues for scatter compensation.

All of the X-ray scatter compensation methods mentioned above require a training procedure, such as those employing scatter-kernels or deep learning,

impose constraints on the object’s shape, or assume the availability of a prior X-CT reconstruction. The requirement for a computationally expensive and data-demanding prior X-CT reconstruction can be circumvented by employing a mesh projector (e.g., [paramonov2024cad]). The mesh projector efficiently estimates the position and orientation of AM parts during scanning using only a few X-ray projections, eliminating the need for a full X-CT reconstruction. This approach enables the application of simulation-based scatter compensation techniques in scenarios where obtaining an X-CT scan is not feasible, such as when only a few projections are available. Moreover, the originally-designed surface meshes can be deformed and adapted to the real printed sample, allowing for accurate X-ray scatter compensation for both the printed part and any other objects present in the scene, such as physical supports. Finally, the use of the mesh projector avoids the necessity of aligning meshes with different digital origins, which often result in differing sampling patterns between the CAD model and the surface mesh extracted from the X-CT scan [lorensen1998marching].

4.2 Materials

Scans were acquired with the FlexCT X-ray scanning device, as detailed in [FlexCT]. The device was equipped with a 2880×2880 detector with pixels of size $150 \mu\text{m}$. The scanned object was an AM component, specifically a bevel gear, which was fabricated using a Form 3+ Low Force Stereolithography (LFS)TM 3D printer by Formlabs, utilizing Draft Resin (Formlabs). During the scan, the AM part was supported by 3D-printed components constructed from Polyamide 12. The source-to-detector distance was 800 mm and the source-to-object distance was 320 mm. A tube voltage of 150 kV_p was applied and the beam was filtered with 1.5 mm aluminum.

The complete scan comprised 3 000 projections, equiangularly distributed in the range $[0^\circ, 360^\circ]$. However, in the industrial quality inspection of AM samples, such a high number of projection images may not always be available. To accommodate this, three datasets were created, each corresponding to a distinct inspection scenario. The first dataset was designed for projection-based inspection and consisted of ten equiangularly sampled projections, and therefore labelled ‘10P’. The second dataset included 100 equiangularly sampled projections and was intended for X-CT inspection using discrete reconstruction of X-CT images, and therefore labelled ‘100P’. Lastly, the third dataset, designed for X-CT-based inspection through conventional CT reconstruction, encompassed the entire set of 3 000 projection images, and is thus referred to as ‘3000P’. A discrete reconstruction applied to 100P was obtained through the poly-chromatic version of

DART [six2019poly] (poly-DART), while conventional X-CT image reconstruction was achieved using poly-chromatic projective error minimisation based on Barzilai-Borwein optimisation [barzilai1988two], referred to as poly-BB. Both poly-DART and poly-BB were built on a distributed version of the ASTRA toolbox [palenstijn2016distributed].

4.3 Methods

The proposed scatter compensation pipeline relies on three essential operations: 1) estimating the object's pose and shape from few X-ray projections (Section 4.3.1), 2) predicting scatter for each individual X-ray projection (Section 4.3.2), and 3) compensating for scattered radiation in the actual X-ray projections (Section 4.3.3).

4.3.1 Multi-object pose and shape estimation

Accurate scatter compensation in X-CT scans requires estimation of the pose of the scanned objects and possible deviations from their nominal shape (e.g., due to the printing process). To achieve this, the nominal surface model of each object from its CAD model is employed along with a mesh projector (CAD-ASTRA [paramonov2024cad]) to iteratively estimate the object's position and orientation. This rigid pose estimation then serves as the starting point for iterative estimation of the (possibly deformed) object shape.

Let K denote the number of objects, where each object $k = 1, \dots, K$ is represented by its surface mesh denoted as $\mathbf{m}_k = (\mathbf{V}_k, \mathbf{N}_k)$. Here, $\mathbf{V}_k = \{\mathbf{v}_{ik}\} \in \mathbb{R}^{V \times 3}$ represents the V vertices of the mesh, and $\mathbf{N}_k = \{\mathbf{n}_{jk}\} \in \mathbb{R}^{N \times 3}$ the normals of its N triangular faces. During pose estimation, the mesh vertices \mathbf{V}_k undergo a roto-translation described by a pose vector $\boldsymbol{\theta}_k \in \mathbb{R}^6$, consisting of three parameters for translation and three for orientation. The iterative refinement begins with an initial position and orientation assumption, where all objects are initially aligned along the vertical axis passing through the center of rotation of the scan. Subsequently, the estimation process minimises the discrepancy between the measured projections $\mathbf{P} \in \mathbb{R}^M$ and the simulated projections, where M is the product of the number of projections and their number of pixels, leading to:

$$\tilde{\boldsymbol{\theta}}_1, \dots, \tilde{\boldsymbol{\theta}}_K = \arg \min_{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K} \left\| \mathbf{P} - \sum_{k=1}^K \text{PolyProj}_d[\mathbf{A}_{\boldsymbol{\theta}_k}[\mathbf{V}_k]] \right\|_2^2. \quad (4.1)$$

Here, the $\text{PolyProj}_d : \mathbb{R}^{V \times 3} \rightarrow \mathbb{R}^M$ operator simulates the poly-chromatic projec-

tions through the metadata d representing the vertices connectivity and material optical properties, while $\mathbf{A}_{\boldsymbol{\theta}_k} : \mathbb{R}^{V \times 3} \rightarrow \mathbb{R}^{V \times 3}$ represents a roto-translation operator that depends on the pose parameters $\boldsymbol{\theta}_k$.

Subsequently, registered meshes are employed for estimating the deformed shape of the AM samples. This deformation is described as per-vertex shifts from the nominal registered surface mesh of the CAD model. Similar to pose estimation, this involves iterative refinement of a linear programming problem. Besides a data consistency term, regularisation terms are introduced penalise significant changes in mesh topology. This linear programming problem is formulated as:

$$\begin{aligned} \tilde{\mathbf{V}}_1, \dots, \tilde{\mathbf{V}}_K = \arg \min_{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K} & \left\| \mathbf{P} - \sum_{k=1}^K \text{PolyProj}_d \left[\mathbf{A}_{\tilde{\boldsymbol{\theta}}_k} [\mathbf{V}_k] \right] \right\|_2^2 \\ & + \sum_{k=1}^K \left(\alpha \|\mathbf{V}_k^* - \mathbf{V}_k\|_2^2 + \beta \|\mathbf{N}_k^* - \mathbf{N}_k\|_2^2 \right) \quad . \end{aligned} \quad (4.2)$$

Here, for each k -th mesh, the distance between the vertices \mathbf{V}_k and original vertices \mathbf{V}_k^* is measured, where the difference $\mathbf{V}_k^* - \mathbf{V}_k$ symbolically represents the difference $\mathbf{v}_{ik}^* - \mathbf{v}_{ik}$ between each of the i vertices of the k -th mesh. An analogous explanation applies to $\mathbf{N}_k^* - \mathbf{N}_k$ and $\mathbf{n}_{jk}^* - \mathbf{n}_{jk}$. These regularisation terms are governed by the parameters α and β during minimisation of the projection error of (4.2). Both the linear programming problems (4.1) and (4.2) are solved using the Adam optimiser.

4.3.2 X-ray interaction simulation

From the available simulation software [guillemaud2003sindbad, badal2009acceleration, sarrut2021advanced, kawrakow2000accurate, ay2005development], as well as specialised X-ray scatter simulation [maslowski2018acuros, elshemey2009monte], the Monte Carlo (MC) photon transport simulator GATE [sarrut2021advanced] (version 9.2, built on Geant4 version 11.1.2) was chosen due to its widespread adoption and extensive validation in literature. Information about the scanning system, such as scanning geometry and X-ray source operating parameters, was obtained from the experimental scan data and used for the GATE simulations. Additionally, the adapted surface meshes of all objects in the scan (as detailed in Section 4.3.1), were incorporated into the simulation. The simulation was configured to produce two outputs: the X-ray radiation detected by the detector, $\hat{\mathbf{P}}$, and the portion of $\hat{\mathbf{P}}$ that underwent at least one either Compton or Rayleigh scattering event within the object or its supporting elements. This portion is

referred to as $\hat{\mathbf{S}}$. Since the FleXCT detector, along with the majority of industrial X-CT systems, is energy-integrating, the simulated detector is designed to compute the cumulative energy contribution of each incident photon. Each contribution is multiplied by the detector's sensitivity corresponding to the photon's energy.

In addition to $\hat{\mathbf{P}}$ and $\hat{\mathbf{S}}$, the white-field image $\hat{\mathbf{W}}$, which represents the unattenuated X-ray radiation, is obtained through an identical simulation without attenuating objects. To minimise the statistical MC fluctuations in the radiation detected by each pixel of the simulated white-field image, the simulated white-field image is fitted with an analytical function for point-wise radiative emission over a finite square detector element. The derivation of this analytical function for flat detectors is provided in A.4.

Accelerating MC simulations

X-ray physical interaction simulations conducted with GATE are known for their high accuracy, but also for slow execution times. To address this issue, the proposed approach focuses on two main strategies: parallelisation of simulations and leveraging the physical behavior of X-ray scattering. For the parallelisation of simulations, an MPI-based manager that is tailored for the GATE simulations has been developed. This manager is responsible for partitioning simulation tasks across CPU-clusters, running parallel instances across computing nodes, and merging the results. The MPI-manager is publicly accessible on GitHub [[mpiforgate](#)], it is platform-independent in terms of workload management systems (e.g., PBS, SLURM), uses the ITK [[mccormick2014itk](#)] and MPI [[dalcin2021mpi4py](#)] libraries, and is fully implemented in Python.

The knowledge on the scattering pattern of X-rays is harnessed to reduce computational complexity by optimizing the number of simulated photons per projection and the number of projections themselves. Previous research by Colijn et al. [[colijn2004accelerated](#)] offered foundational insights into accelerating simulations by exploiting the spatial smoothness of scatter signals. However, self-absorption of scattered radiation by the object potentially leads to higher scatter frequencies for complex object shapes. Building upon this foundation, the simulation software is utilised for determining how many photons will give a reliable measure of scattered radiation, to identify the minimal number of photons needed to simulate each projection. This is done by monitoring the power of the spatial scatter signal (see Fig. 4.1) for the first projection, for increasing number of simulated photons. As the photon count at each pixel increases along with the signal-to-noise ratio of the scatter signal, the bandwidth Ω_m of the scatter signal




Figure 4.1: Every minuscule region of the scanned object (brown) scatters X-ray radiation with a specific density per steradian. This figure illustrates the contribution of a scattering region of the object on the detector reading throughout three subsequent X-ray projections. Spatial frequencies of the scattered radiation are observed across the detector for a specific projection, while temporal frequencies are observed across subsequent projections for a specific detector pixel.

(representing the highest frequency encompassing 95% of the total power within $[0, \Omega_m]$) gradually converges to a stable value (as determined experimentally in Section 4.4.1). The number of photons determined in this way was used for all subsequent projections. As a consistent part of the scatter signal is contained in Ω_m , the rest of the spectrum can be cleaned from the MC noise with a low-pass Butterworth filter, known for its maximally flat response in the passband Ω_m . By utilising the filter as a zero-phase filter, the filter order is chosen as the lowest order necessary to achieve a maximum power suppression of 0.025 dB in the Ω_m -band, along with a minimum -30 dB power suppression of frequencies beyond the stopband Ω_s , which is 10% higher than Ω_m . These numeric values that select the filter order were found to be sufficient for the denoising task.

By exploiting the temporal smoothness of the scattered radiation on the detector pixels (i.e., smoothness between subsequent projections, see Fig. 4.1), the computational burden of simulating the entire X-ray projection stack is reduced to the simulation of just a few projections. These few scatter estimates, through Fourier interpolation, provide the X-ray scatter estimates cor-




Figure 4.2: The X-ray scatter footprint of a minuscule region of the scanned object exhibits low spatial frequencies in conventional X-CT acquisition setups. This makes two scattering regions indiscernible if they are close enough, despite the linear absorption associated with these regions being resolved by the tomographic reconstruction.

responding to all other projections. This is a substantial difference from previous works [maslowski2018acuros, xia2019scatter, bootsma2013spatial], where scatter estimates were computed using linear interpolation, and the selection of key projections was done manually. In principle, nearby scattering regions of the scanned object at a distance d may not be clearly discernible from their scatter footprint (Fig. 4.2). Consequently, during an X-CT reconstruction that ignores the scattered radiation in the forward model (e.g. FDK, SART, SIRT), the scatter-affected projections drive the reconstructed volume to be explanatory of the (mainly low) spatial frequencies of the scatter signal, as well as the scatter signal variation among projections. By sub-sampling these projections, the variation of the scatter signal among these projections can be constrained to have the same impact on the reconstruction as for the spatial scatter signal. In case of circular cone-beam computed tomography (CBCT), the constraint is imposed through known relations [izen2012sampling], based on the data completeness theorem for the central cross-section [wu2023cone], which establish a connection between spatial sampling and angular sampling for reconstructing a given object. Let $g(x, y) \in \mathbb{R}$ indicate the linear attenuation of said object, contained in a circle of radius r , such that $g(x, y) = 0; \forall (x, y) \in \mathbb{R}^2, |\sqrt{x^2 + y^2}| \geq r$. Assuming that the spectral band of g can be reasonably Ω_g -limited (as is commonly done for X-CT reconstructions [izen2012sampling, zhao2015generalized]), it follows that a condition on the minimal number of samplings along each view is necessary to recover the Ω_g -limited g [izen2012sampling]:

$$\Delta y \leq \frac{\pi}{\Omega_g} \quad . \quad (4.3)$$

Here, Δy represents the scaled spacing between pixels on the flat detector, measured relative to the object center ($\Delta y = \Delta s \cdot \frac{R}{D}$, where R is the source-to-object distance and D the source-to-detector distance), where Δs denotes the spacing between pixels. Moreover, to ensure that g can be fully recovered, the angular spacing $\Delta\beta$ among views of the circular trajectory must also satisfy the inequality ([izen2012sampling]):

$$\Delta\beta \leq \frac{\pi}{r\Omega_g} \left(\frac{R+r}{R} \right) \left(1 - \frac{r^2}{R^2} \right) \quad . \quad (4.4)$$

Given that the spatial bandwidth Ω_m of the scattered radiation is already established, the object bandwidth Ω_g that described the spatial scatter is linearly determined and consequently constrains the angular sampling, as indicated by (4.4). Finally, after determining the subset of X-ray projections to simulate using this angular sampling, resampling in Fourier space provides the missing estimates for scatter.

Variance reduction techniques in MC simulations of X-ray photons are known to improve the efficiency of the simulations. These methods strategically allocate computational resources to areas where interactions are most important, reducing statistical noise. As these techniques may introduce a bias in scatter estimates, which is adversary to the scope of the study, they are not used here.

4.3.3 X-ray scatter compensation

After obtaining the estimate of scattered radiation, the next step is to compensate for the scatter in the projections. It is important to note that the measured and simulated projections depend on the number of photons physically generated and simulated, respectively. Additionally, the presence of dark currents affecting the measured projections must be properly handled during X-ray scatter compensation — a factor that was overlooked in previous studies [rew:scatter:sim-all-proj, maslowski2018acuros, rew:scatter:sim-few-proj2, xia2019scatter]. To address these factors, the following corrections are applied.

First, the projections \mathbf{P} are corrected using the white-field \mathbf{W} and dark-field \mathbf{D} images to account for gain and dark currents. This results in

$$\mathbf{P}_{corr} = \frac{\mathbf{P} - \mathbf{D}}{\mathbf{W} - \mathbf{D}} \quad . \quad (4.5)$$

Similarly, the same correction is applied to the simulated projections, which do not have dark currents, leading to $\hat{\mathbf{P}}_{corr} = \hat{\mathbf{P}}/\hat{\mathbf{W}}$ and $\hat{\mathbf{S}}_{corr} = \hat{\mathbf{S}}/\hat{\mathbf{W}}$. In the absence of scatter, the dynamic range of \mathbf{P}_{corr} and $\hat{\mathbf{P}}_{corr}$ would be within to the interval $[0, 1]$. However, the presence of scattered radiation causes some pixel values to exceed 1. Since a portion of the absolute value of \mathbf{P}_{corr} and $\hat{\mathbf{P}}_{corr}$ is attributed to scattered radiation, this scattered radiation is compensated in the measured projections with $\mathbf{P}_{sc} = \mathbf{P}_{corr} - \hat{\mathbf{S}}_{corr}$. Finally, the dynamic range of the measured projections can be restored by reversing the white/dark-field correction using $\mathbf{P} = \mathbf{P}_{sc}(\mathbf{W} - \mathbf{D}) + \mathbf{D}$ and are ready for reconstruction.

4.4 Experiments

The three experiments presented in this section are conducted sequentially, with the information gathered from experiments in Section 4.4.1 and 4.4.2 serving as a foundation for the last one (Section 4.4.3). After determining the optimal number of photons required for a reliable estimation and establishing the bandwidth of the X-ray scattering signal to downsize the simulation task (Section 4.4.1), the next step involves multi-mesh registration and adaptation based on a minimal

set of X-ray projections (dataset 10P), as detailed in Section 4.4.2. This step enables the subsequent simulation and compensation of X-ray scattering from the scanned X-ray projections 4.4.3.

4.4.1 Spatial frequencies of X-ray scatter on the detector

In this experiment, the scattered radiation of the scanned sample is measured through simulations that were run with different number of photons (10^9 , 10^{10} , ..., 10^{13}), with the information provided in Section 4.2 and 4.3.2. The aim of the experiment is two-fold: to find the number of photons to be used for a reliable estimation of the X-ray scattered radiation for the X-ray scan, and to find the bandwidth of the X-ray scattered signal. The latter is used to determine the angular downsampling of the X-ray scatter simulations. The analysis is conducted in the frequency domain, where the bandwidth Ω_m of the X-ray scattered radiation is defined as the band in which 95% of the power is contained.

4.4.2 Multi-object pose and shape estimation

In this experiment, the aim is to infer the pose of the objects (the gear and its supporting elements) from the scanned X-ray projections and to adapt any surface mesh to the corresponding object, to compensate for deformation arising from a non-ideal 3D printing process. The operation is conducted in two subsequent phases. First, the pose of the object is estimated through minimisation of the linear programming problem (4.1) while the poly-chromatic behaviour of the source and detector are continuously estimated jointly with the object pose. Then, while the spectral behaviour and the pose remain fixed, the linear programming problem described with (4.2) is optimised by adjusting the vertices' positions in order to account for the objects' deformations. The optimisation of both the linear programming problems is performed with the lowest number of projections (dataset 10P).

4.4.3 X-ray scatter compensation

In this experiment, the surface meshes of all the components of the X-ray scene, along with their chemical composition information, are used for the compensation of the X-ray scattered radiation. The full compensation procedure is described in section 4.3.3. This process was systematically applied to all datasets, and the outcomes were analysed according to the inspection scenario — whether it is projection-based inspection, X-CT inspection using discrete reconstruction, or conventional X-CT inspection. Projection-based inspection, suitable for rapid

assessments where tomographic information is unnecessary, relies on the sparse set of projections from the 10P dataset. In contrast, X-CT inspection with discrete reconstruction offers enhanced analysis and tomographic visualization of internal features and defects compared to projection-based methods, leveraging the denser 100 equiangularly sampled projections of the 100P dataset. Conventional X-CT inspection, providing the highest level of tomographic detail and resolution for precise reconstruction, is employed when thorough examination and accurate measurement are imperative, despite the increased computational and time demands. This method utilizes the extensive data from the 3000P dataset. Specifically, for projection-based and conventional X-CT inspection inspections, the results are supplemented with an assessment of residual errors using the mesh-projector (cfr. Section 4.4.2). In each case, the obtained results were benchmarked against the performance of a recent flavour of scatter kernel superposition method [bhatia2016separable] for scatter compensation.

4.5 Results and discussions

4.5.1 Spatial frequencies of X-ray scatter on the detector

Fig. 4.3 shows the results of simulation runs with increasing number of photons, and all replicating the same geometrical setup of a scanned X-ray projection. More in-depth insights can be obtained from Fig. 4.4, which shows the power spectrum of the detector readout. As expected, the white noise related to the MC-statistics diminishes its power by 10 dB for every order of magnitude of the number of simulated photons. As a consequence, the 95% of the power of the detected signal falls between a smaller bandwidth, as the number of photons increase, which is in accordance with the literature affirming that scattered radiation is mostly composed of low-frequencies. The 95% power bandwidth is reported in Table 4.1 and shows that after a certain number of simulated photons (10^{11}), the bandwidth of the signal does not vary anymore. This is the lowest number of photons that is simulated during the simulation of each projection during the scatter compensation phase. The related bandwidth has been used to set up the low-pass filter, which led to Fig. 4.5a. For this bandwidth, the number of projections to be simulated during the scatter compensation phase, as selected according to Eq. (4.4), is 30, which resulted in almost 11 h of computation time with 2048 cores (AMD Epyc 7452).




Figure 4.3: Rayleigh and Compton scatter for different numbers of simulated photons, divided by the average value.




Figure 4.4: One-sided power spectrum, extracted from the vertical distribution of frequencies of the 2-D Fourier space, for the scattered radiation shown in Fig. 4.3. The white-noise, related to the MC-statistics, lowers its power with higher number of photons, while the low-frequency scatter components become noticeable.

	10 ⁹ ph.	10 ¹⁰ ph.	10 ¹¹ ph.	10 ¹² ph.	10 ¹³ ph.
Ω_m [rad/px]	3.138	9.942e-03	6.628e-03	6.628e-03	6.628e-03

Table 4.1: The 2D bandwidth of the scattered radiation, which contains 95% of the total power, is reported here for different numbers of simulated photons.




Figure 4.5: (a) The X-ray projection after application of the Butterworth filter, using 10¹¹ photons, and (b) the relative difference between the filter and related unfiltered image, showing that only the MC-white noise has been removed by the filter.

4.5.2 Multi-object pose and shape estimation

The optimisation of the linear programming problem (4.1) refined the initial pose of the objects within the scanned scene, which is noticeable in Fig. 4.6 from residual images (i.e. difference between simulated projections and scanned projections) and measures of the Root Mean Squared residuals (RMSres). Starting from the initial object pose configuration that resulted in high projection errors (Fig. 4.6a), the optimization led to improved spatial positioning of the objects (Fig. 4.6b), highlighting a significant reduction in projection-wise error. However, some regions in the post-registration image still exhibit high errors, which can be attributed to actual 3D printing deformations.

By optimising the linear programming problem (4.2), the surface meshes of the CAD model were adjusted to match the true shape of the 3D printed objects (Fig. 4.6c). A visual representation of this deformation is provided in Fig. 4.7, where the displacement of each vertex is measured as Euclidean distance, influencing the heatmap colour representation in the figure.

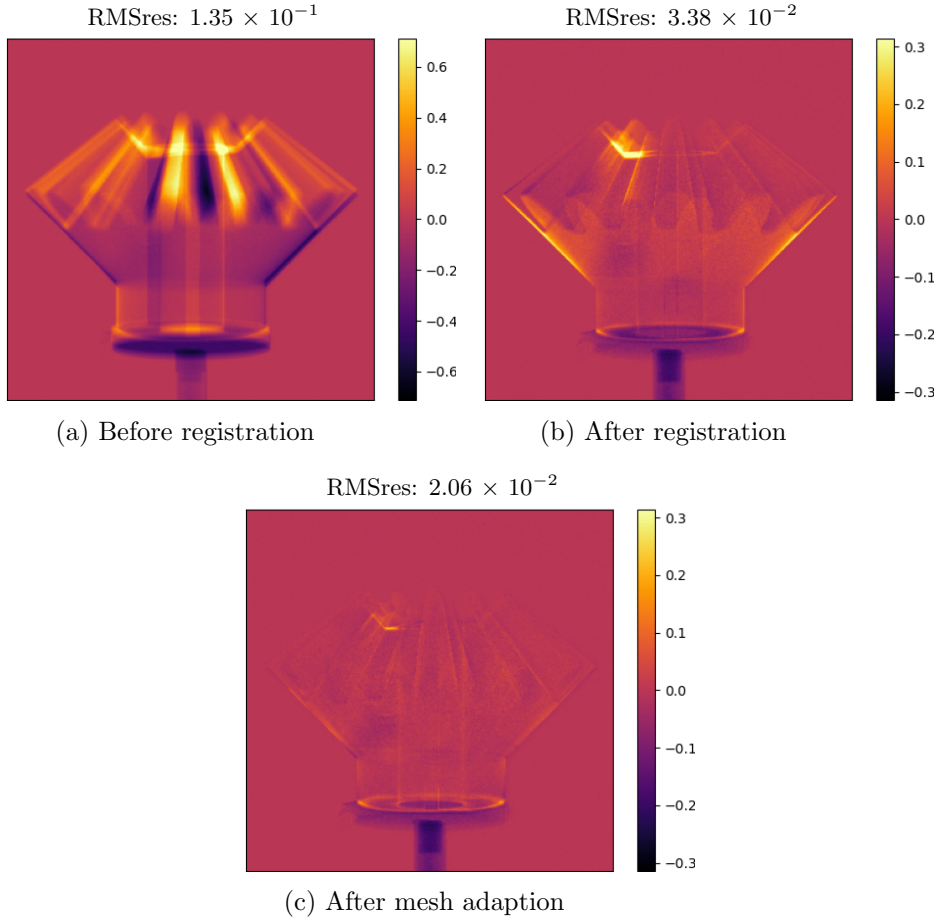


Figure 4.6: Residual images of the bevel gear and its supporting elements for three different instants of the multi-mesh registration/adaption procedure. The figure shows how an initial pose guess of the various elements can be adapted to the actual scan, while accounting for deformations, in an automatic way.




Figure 4.7: The adaption of the surface mesh of the CAD model to the actual object is considered to arise from a faulty 3D printing process, A heatmap of the Euclidean displacement of each surface mesh vertex provides a clearer visualisation of the deformed shape compared to the projection-wise representation of Fig. 4.6b and 4.6c.

4.5.3 X-ray scatter compensation

In Fig. 4.8, an analysis of the scatter compensation for the projection from the dataset 10P in projection space is presented. Despite the visual similarity between the chosen projection and its scatter-compensated versions, certain details are more apparent in the residual images. These images depict the absolute residual between the dark/white-field corrected X-ray projections and the poly-chromatic projections of the meshes.

Examining the images in Fig. 4.8, the central part of the residual image for a measured projection (Fig. 4.8b) reveals a whitish halo. This halo is less pronounced in the scatter-compensated projection using the Scatter-Kernel Superposition (SKS) method [bhatia2016separable], as noticeable in Fig. 4.8d, and almost imperceptible in the scatter-compensated projections using PACS (Fig. 4.8f). Additionally, all residual images (Fig. 4.8b, 4.8d and 4.8f) display the general shape of the scanned sample, resulting from minor inaccuracies in mesh estimation. Furthermore, two macroscopic defects can be observed. Close-ups and line profiles of these macroscopic defects, provided in Fig. 4.9, indicate that the scatter-affected projection has a non-zero value even outside the defective zone. These values arise from the mismatch between the actual object surface and its surface mesh. To quantitatively evaluate the performance of the scatter compensation methods of the results shown in Figures 4.8 and 4.9, the RMSres is

	Unprocessed	SKS	PACS
RMSres	$(3.778 \pm 0.046) \times 10^{-2}$	$(2.094 \pm 0.046) \times 10^{-2}$	$(1.539 \pm 0.044) \times 10^{-2}$

Table 4.2: Measure of the mean squared residual throughout all projections of dataset 10P. The residual is interpreted as the difference between the projections and the simulated poly-chromatic projections.

	Unprocessed	SKS	PACS
# of pore-voxels	1816	1001	368
Sørensen–Dice coefficient	0.061	0.069	0.074

Table 4.3: Measure of the number of detected pores through the poly-DART reconstruction of the few projections of dataset 100P. The similarity of the detected pores with those of the ground truth is measured with the Sørensen–Dice score.

computed along with its standard deviation and it is reported in Table 4.2. This evaluation, encompassing all projections in dataset 10P, highlights that PACS yields the least amount of residuals than the SKS method.

With the dataset 100P, a discrete X-CT reconstruction is performed. Employing this technique enables the creation of a (voxel-wise) multi-label volume, offering a three-dimensional discrete representation, as depicted in Fig. 4.10. The application of this technique highlights macroscopic defects more distinctly, as observed in the images in Fig. 4.11b, 4.11f and 4.11h. Furthermore, it allows for the detection of pores within the X-CT. However, in practice, the presence of imaging noise and scattered radiation may lead to the identification of spurious pixels unrelated to real pores.

In Fig. 4.11, the 3D rendering visually displays the distribution of spurious pores in the reconstructed volume. Notably, scatter-compensated projections result in a reduced occurrence of these defects. The number of detected pores and their similarity to the voxel-wise manually-labelled ground truth of pores are quantified and presented in Table 4.3. This analysis suggests that scatter compensation significantly contributes to a higher fidelity of pore representation, with PACS outperforming SKS in terms of the Sørensen–Dice coefficient. The reason for higher concentration of detected pores in the top part of the gear is related to X-ray scattering artefacts, and it is explained by using the dataset 3000P, where more information can be extracted from X-ray projections.

Utilising the complete set of projections (dataset 3000P) allows for regular X-CT reconstruction, ensuring faithful results. In Fig. 4.12, a cross-section anal-

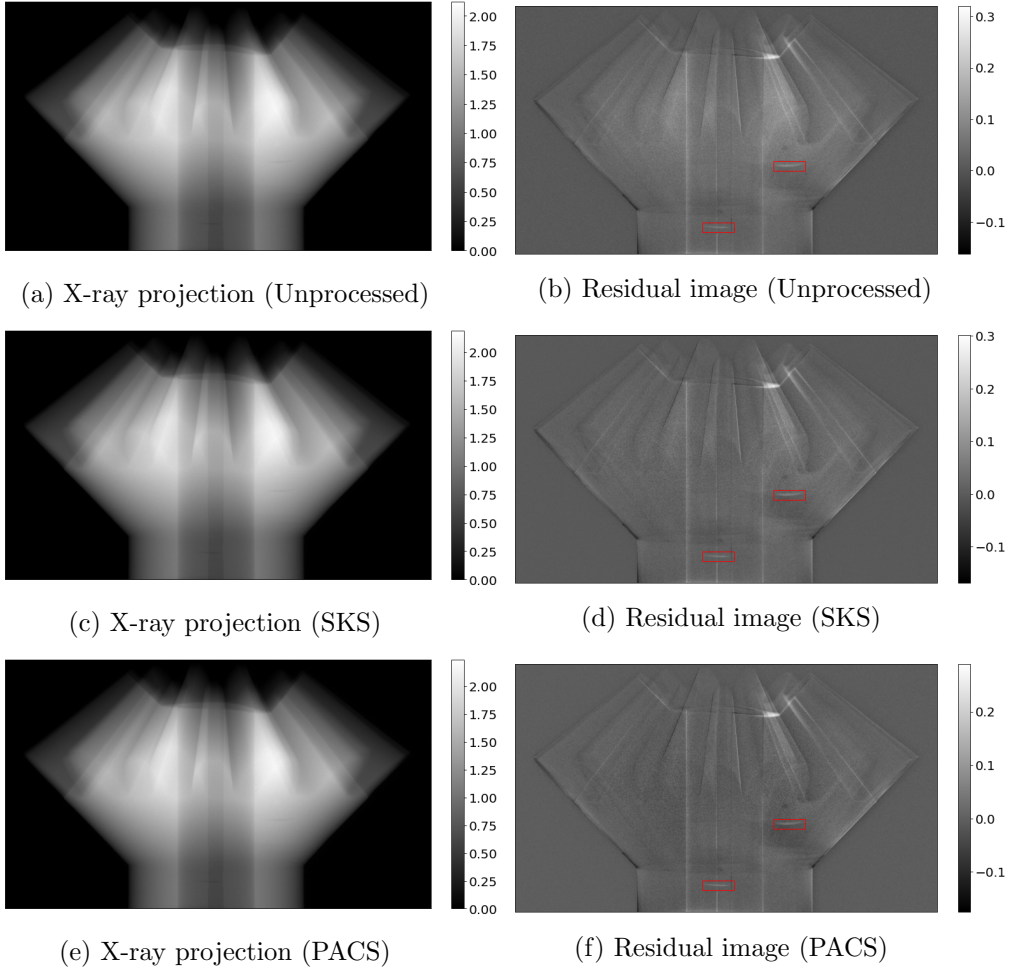


Figure 4.8: X-ray projections are presented on the left for three cases: (a) without scatter correction, (c) with scatter correction using SKS, and (e) with scatter correction using PACS. Accompanying these, the residual images are depicted as (b) uncorrected, (d) corrected with SKS, and (f) corrected with PACS. Two red boxes highlight macroscopic defects for a detailed analysis in Fig. 4.9.

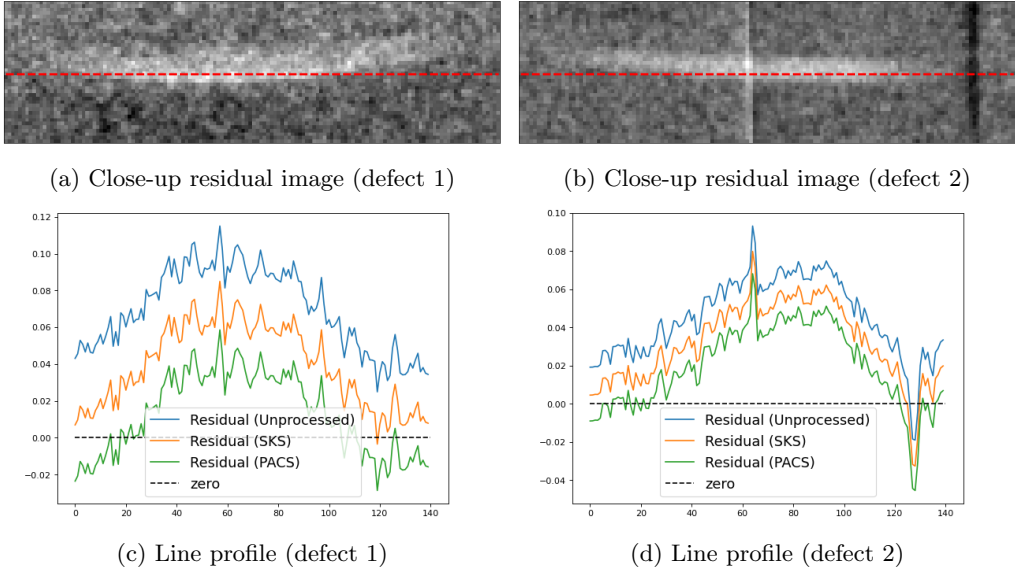


Figure 4.9: Close-ups of two defects highlighted in the residual images of Fig. 4.8. Line profiles through these close-ups reveal the varying behavior of residual gray values across the defects, where it is desirable that high values are associated with the defective area, while zero values are associated with the absence of defects.






Figure 4.10: From the limited data in the 100P dataset, the binary poly-DART reconstruction (a) of the bevel gear holds meaningful information, as for the 3D printing defects (b).



(a) 3D view of pores (unprocessed)




(b) Gear defect (unprocessed)




(c) Region with many detected pores (unprocessed)



(d) 3D view of pores (SKS)




(e) Gear defect (SKS)




(f) Region with many detected pores (SKS)



(g) 3D view of pores (PACS)



(h) Gear defect (PACS)



(i) Region with many detected pores (PACS)

Figure 4.11: Comparison of Poly-DART reconstructions: unprocessed X-ray projections (a-c), scatter-compensated with SKS (d-f), and scatter-compensated projections with PACS (g-i). The 3D visualization (a,d,g) highlights all detected pores within the Poly-DART reconstruction, with a visibly higher number in the top part of the bevel gear. The images (b,e,h) show a cross-section with the correct reconstruction of a defective area, while the (c,f,i) images focus on the variable number of pores detected by each Poly-DART reconstruction. The correlation of the detected pores with the ground truth is shown in Table 4.3.

	Unprocessed	SKS	PACS
Gray value	$(2.654 \pm 0.001) \times 10^{-2}$	$(2.719 \pm 0.001) \times 10^{-2}$	$(2.768 \pm 0.001) \times 10^{-2}$
SNR	87.2	92.6	96.9

Table 4.4: Measure of the average gray value level and SNR throughout the bevel-gear region of the poly-BB reconstruction of dataset 3000P, for unprocessed X-ray projections, SKS scatter-compensated projections and the scatter-compensated projections with PACS.

	Unprocessed	SKS	PACS
Average CNR	4.23 ± 0.37	4.45 ± 0.40	4.70 ± 0.40

Table 4.5: Measure of the average CNR of pores in the poly-BB reconstructions of the uncorrected, SKS scatter-compensated and scatter-compensated projections with PACS, for the dataset 3000P.

ysis of the poly-BB reconstruction is presented in a region previously identified as challenging for poly-DART. The line plot in the figure illustrates that the standard deviation and average gray values of poly-BB are smaller for scatter-compensated images, providing an explanation for the spurious pores detected with the 100P dataset. These values, evaluated across the volumetric gear region, are detailed in Table 4.4. A consequence of the higher signal-to-noise (SNR) value is that the contrast of pores is enhanced. This is evaluated across all pores and reported in Table 4.5, through their contrast-to-noise (CNR) ratio, which confirmed the expectations.

Examining the residual values (Fig. 4.13) offers a deeper insight into the results observed in Figures 4.11 and 4.12. The presence of spurious pores (Fig. 4.11) and the additional noise and non-constant trend of line profiles (Fig. 4.12), unexpected for a homogeneous material, is attributed to X-ray scattering artefacts. These artefacts gradually diminish with the application of scatter-compensation techniques. The evaluation of average absolute residuals, confirming the visual findings, is reported in Table 4.6.

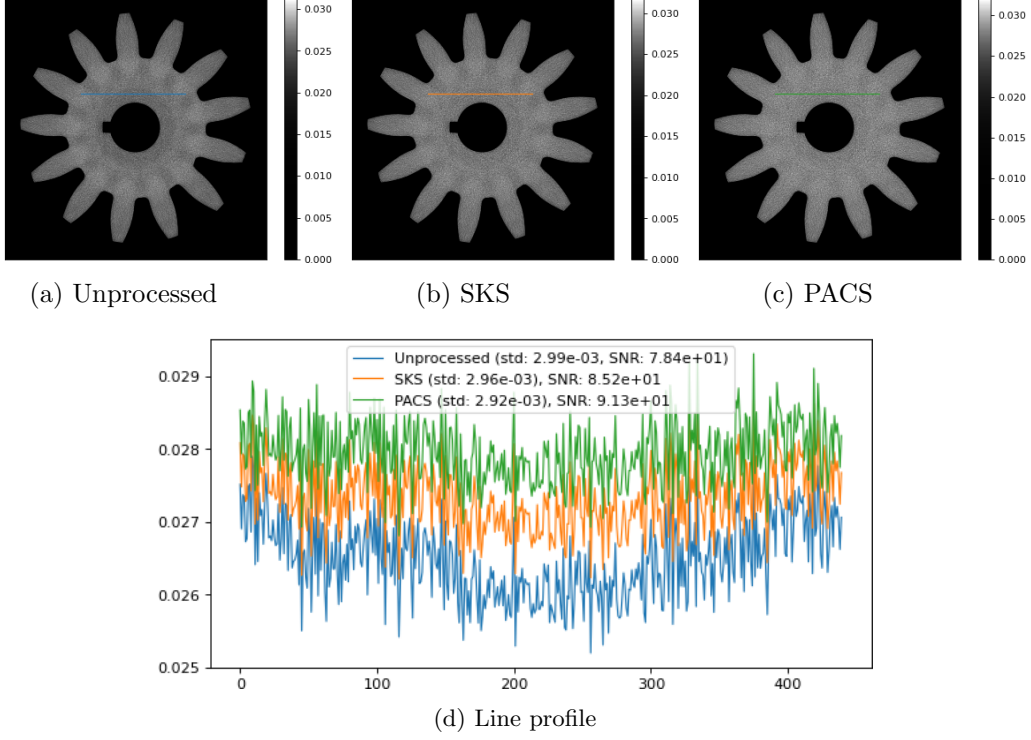


Figure 4.12: A cross-section of the poly-BB reconstruction of (a) unprocessed X-ray projections, (b) SKS scatter-compensated projections and (c) the scatter-compensated projections with PACS of the complete 3000P dataset.

	Unprocessed	SKS	PACS
RMSres	$(3.146 \pm 0.034) \times 10^{-3}$	$(2.909 \pm 0.035) \times 10^{-3}$	$(2.819 \pm 0.035) \times 10^{-3}$

Table 4.6: Measure of the root mean squared residual throughout the dataset 3000P. The residual is interpreted as the difference between the poly-BB reconstruction and the simulated poly-BB reconstruction in the space occupied by the gear.




Figure 4.13: A cross-section of the residual poly-BB reconstruction of unprocessed X-ray projections (a), SKS scatter-compensated projections (b) and the scatter-compensated projections with PACS (c) of the complete 3000P dataset. The Root Mean Squared Error (RMSE) evaluated throughout the whole image is reported on top of each figure.

4.6 Conclusions

A new approach (PACS) was proposed to accurately compensate for X-ray scatter in projection images. This method involves registering and adapting CAD models of AM objects, followed by a Monte Carlo simulation of X-ray scatter in the projections, and then compensating for it in the actual radiographs. Results demonstrated that accurate registration of CAD models is possible using only few projections, rendering simulation-based X-ray scatter compensation approaches a viable option in scenarios where few projections are available or desirable. The adaptability of such approach addresses common constraints encountered in industrial scanning setups, including throughput, cost, and spatial limitations. Replacing the CAD model with another object before scanning, in response to changing industrial needs, incurs no additional overhead compared to alternative methods. Furthermore, this method, inherently coupled with a mesh projector, allows for the analysis of residual differences between real and simulated projections. This feature facilitates the identification of deformities or defects within the scanned object.

Multiple experiments consistently demonstrate superior accuracy in X-ray scattering artefact reduction of PACS compared to one of the latest SKS scatter compensation methods. The absence of training procedures and the ability to easily adapt to any scanned object make PACS the method of choice for scans in the context of AM. As PACS relies on Monte Carlo simulations for X-ray scatter compensation, strategies to reduce the computational load were investigated. However, employing GATE as a simulation backend remains computationally intensive, which can be further reduced using modern GPU solutions specialised in X-ray scatter simulations, as presented in recent works [maslowski2018acuros].

A noteworthy future prospect arising from this work is the development of an X-ray scatter estimation pipeline that would rely solely on the mesh projector. Coupling the mesh projector with an analytical description of the 1st order scatter [yao2009analytical], evaluated across the entire surface mesh, is expected to yield significant improvements in speed. This is particularly true in cases where the inspected object is small, or the detector is positioned sufficiently far from the X-ray source.

Voxel-wise pore segmentation with 3D patch-based neural models

In this chapter, the analysis and segmentation of porosity and defects within AM samples using deep learning techniques are the focus. Detecting anomalies from X-CT data is a challenging task due to factors such as inhomogeneous density, low contrast-to-noise ratio, and beam hardening artifacts, which can lead to incorrect segmentation.

Data-driven DL approaches have shown superior performance over traditional machine learning techniques in handling complex and varied definitions of anomalies [wong2021automatic, bihani2022mudrocknet, kim2022achieving, sarkon2022s]. Anomalies can be detected using supervised or unsupervised methods. While supervised methods require annotated datasets, unsupervised methods are more desirable as they do not require annotated training data. This reduces the technical overhead and mitigates the impact of noisy annotations on model performance. However, unsupervised methods often face challenges such as high recall rates and low precision compared to their supervised counterparts [yang2022visual].

Most studies on voxel-wise segmentation tasks using DL techniques focus on analyzing stacks of 2D images [bouget2019semantic, ar2020segmentation, fend2021reconstruction, wang2022deep, mehta2022federated, wang2022centerne]. For voxel-wise segmentation of pores in AM samples, a 2D approach is sub-optimal since small pores usually span only a few voxels in three dimensions in X-CT images and suffer from a low contrast-to-noise ratio. Additionally, pores often exhibit anisotropy, increasing the risk of being overlooked by 2D pixel-wise segmentation methods [maskery2016quantification]. Recognizing this short-fall, Wong et al. introduced 3D pore detection models, initially using a UNet architecture, which demonstrated promise but did not explore deep supervision, alternative neural models, or training patterns extensively [wong2021automatic].

Deep supervision can yield more reliable results by encouraging hidden layers of the models to comply with the desired output [li2022comprehensive].

However, training supervised models on datasets with reduced porosity may severely affect detection performance due to class imbalance between pore and non-pore voxels [bria2020addressing]. Additionally, supervised models are highly sensitive to training labels. Unsupervised models, particularly those based on cro:VAE Variational Autoencoder (VAE), often produce blurry representations of input images due to learning a low-dimensional representation that may not capture fine details [guo2020variational]. Therefore, the voxel-wise anomaly score from these models can be enhanced with more complex anomaly scores or dedicated post-processing [zimmerer2018context, baur2021autoencoders].

With this chapter, 2D supervised and unsupervised DL models are revisited and extended to 3D for voxel-wise segmentation of pores in X-CT samples of varying alloys. Utilizing a 3D patch-based approach and integrating data augmentation, this segmentation method aims to be independent of the material and shape of AM samples, ensuring spatial consistency by operating within the 3D image domain. Several deeply supervised models are trained, including UNet++ [zhou2018unet++], UNet 3+ [huang2020unet], MSS-UNet [ZHAO2020100357], and ACC-UNet [ibtehaz2023acc], with a traditional UNet [ronneberger2015u] serving as a baseline for comparison. To address class imbalance due to the low amount of defects, models are trained using the cro:FTLFocal Tversky Loss (FTL) function, which penalizes anomalies more effectively [abraham2019novel]. Optimal parameters for the FTL function were determined through parameter search.

A roster of unsupervised models, including VAE [kingma2013auto], ceVAE [zimmerer2018context], gmVAE [dilokthanakul2016deep], vqVAE [van2017neural], and RV-VAE [nicodemou2023rv], were additionally evaluated, comparing older and novel architectures. To reduce misclassifications, the anomaly scores of these models are post-processed due to their inability to represent object surfaces adequately. Finally, the supervised models are retrained using the post-processed output of an unsupervised model instead of potentially noisy annotations, effectively making the training process unsupervised to evaluate the impact on model performance. The best performing model is further tested to assess the decrease in performance when lowering the number of X-ray projections and exposure.

Summarizing, the main contributions of the work presented in this chapter are as follows:

- First cross-validated assessment of multiple 3D DL models for voxel-wise pore segmentation in AM samples, comparing supervised and unsupervised approaches using a patch-based method. The neural networks, initially designed as 2D models, were tailored for the 3D context to harness volumetric



Figure 5.1: Some samples used in this study. From left to right, a stainless steel 316L (CLM), a CoCr-DG1 (PLM), and a TiAl6V4 (CLM) sample.

information effectively.

- A post-processing algorithm is proposed and evaluated to address the issue of blurry image representation in VAE models.
- The impact of using unsupervised model labels instead of heuristic algorithm labels for training the DL models is evaluated.

5.1 Materials

Various DL models for voxel-wise segmentation of pores were trained using 3D X-ray CT images of AM samples. To this end, AM samples were manufactured through the selective laser melting process, in a continuous (CLM, [meiners1998shaped]) or pulsed laser melting (PLM, [abe2001manufacturing]) strategy. Five cylindrical samples of different materials were 3D printed (as shown in Fig. 5.1): one with TiAl6V4, two with CoCr-DG1 alloy and two with SS316L. Printing the test objects in multiple materials allowed to assess the effectiveness of voxel-wise pore segmentation across different materials. In the CAD model used for the 3D printing, the cylinders had an height of 20 mm and a diameter of 5 mm. In addition to the cylindrical samples, a SS316L cube with an edge length of 9 mm was also printed. The cube was specifically printed to provide an object with different shape and poorer X-CT image quality, which is useful for evaluating porosity in a challenging visual environment and to ensure that DL models are not learning information regarding the shape of the object. These samples were essential for this study as their X-CTs provided the digital dataset with which the neural networks could be trained to classify the porosity. Porosity was intentionally induced in all samples using controlled laser parameters, as described in [booth2022encoding].

Next, 3D images of the AM samples were generated by scanning them with a micro-CT X-ray system [**FlexCT**] and reconstructed with the FDK algorithm [**feldkamp1984practical**] with a $10\text{ }\mu\text{m}$ resolution. The imaging settings, such as filament power, peak kV of the anode, exposure time, source filter, etc., were selected for each cylindrical sample to ensure comparable image quality. However, the geometrical distances and the number of projections were kept constant for all cylindrical samples, with a source-to-detector distance (SDD) of 650.0 mm, a source-to-object distance (SOD) of 43.33 mm, and 4283 projections. The cubic sample was scanned with different SDD (950.0 mm) and SOD (63.33 mm) and had a lower number of X-ray projections (2878) than the other scans. The X-CT of the cubic sample was also affected by severe cone-beam artefacts and poor beam-hardening compensation. The cubic sample was particularly challenging due to its different geometry and visual environment (as noticeable in Fig. 5.2), making it useful for evaluating porosity.

5.2 Methods

Several DL models were trained to segment porosity from X-CT scans of AM samples at the voxel level. Voxel-wise annotations, necessary for both training and performance evaluation, were provided using the method described in section 5.2.1. These models employed either supervised or unsupervised approaches, detailed in section 5.2.2.

For the training of supervised models, the class imbalance of labels was addressed using the FTL function, which will be discussed in section 5.2.4. The class imbalance arose from the low amount of pores (positive instance of labels) within the training dataset. After training, and only for the unsupervised models, the anomaly score is post-processed, as unsupervised models are known to produce blurry representations of the input. The post-processing procedure is explained in section 5.2.5.

5.2.1 Dataset annotation

To assign a label to each voxel of the X-CTs comprising the datasets, which indicates whether it is a pore or not, a 3D processing algorithm was applied. The high-level pseudo-code in Algorithm 1 outlines the pore identification process.

The algorithm for extracting pores from volumetric X-CT images begins by creating a binary mask to distinguish low-value voxels using Otsu thresholding ($Otsu_{thr}$). Subsequently, a background mask is obtained through FloodFill starting from a corner of the X-CT image, isolating low-attenuating values. The




Figure 5.2: A slice of the X-CT of a cylindrical sample (left) and of the cubic sample (right), with equal colour-map and scale. While all the X-CT of cylindrical samples share similar image quality, the cube has stronger artefacts (which are particularly visible at the extremities of the cube) and consequently less contrast. The histograms (a) and (b) refer to the cylindrical sample and of the cubic volume, respectively. The two peaks in each histogram are related to the background (lower) and foreground (higher) colours. The quality of each sample is defined by its distance between the peaks and the broadness of the bells, which are influenced by artefacts and noise.

Algorithm 1 Pore Extraction from X-CT Images

```

1: Input
2:    $CT$       Volumetric X-CT image
3: Output
4:    $pore_{mask}$  Volumetric binary mask representing pores
5:
6: Get low-value voxels through  $Otsu_{thr}$  of  $CT$ 
7: Get the background mask by FloodFilling the low-values from a corner of  $CT$ 
8: Get the watertight object mask from binary inversion of the background mask
9: Get low-value voxels by  $Otsu_{thr}$  of  $CT$  inside the object mask
10:  $pore_{list} \leftarrow$  Collect connected low-value voxels inside the object
11:
12: for all pores in  $pore_{list}$  do
13:   if size of  $pore <$  minimal size then
14:     remove  $pore$  from  $pore_{list}$ 
15:   end if
16: end for
17:  $pore_{mask} \leftarrow$  Convert  $pore_{list}$  to a volumetric mask

```

binary inversion of this background mask yields the watertight object mask, effectively separating the image into air and the watertight object. To identify low-value voxels corresponding to pores, a second Otsu thresholding operation is applied within the object. To address the potential misclassification of pores due to imaging noise, pores-voxels are screened based on shape criteria. Initially, pores in a 6-connected 3D neighbourhood are identified and listed. The boundary box of each pore is then examined, and the pore is excluded from the list if its boundary box is smaller than 2 in at least one dimension. This shape-based filtering is implemented to improve the reliability of the pore identification process [DUPLESSIS20181102, kim2017investigation]. The filtered $pore_{list}$ is then converted into a volumetric binary mask ($pore_{mask}$), providing a voxel-wise representation of pore locations. It's important to note that any residual misclassification arising from partial-volume effects and imaging artefacts contributes to the overall noise of the labels.

Accurately and reliably labelling the X-CT scan of the cubic sample was a challenging task due to its poor image quality, as discussed in section 5.1. Given the limitations of automated voxel-wise annotation, manual labelling was the only viable option to achieve the desired level of accuracy and dependability in the labels.

5.2.2 Deep learning models

The study used two types of models: VAE-based models (VAE [kingma2013auto], ceVAE [zimmerer2018context], gmVAE [dilokthanakul2016deep], vqVAE [van2017n] and RV-VAE [nicodemou2023rv]) and UNet-based models (UNet [ronneberger2015u], MSS-UNet [ZHAO2020100357], UNet++ [zhou2018unet++], UNet 3+ [huang2020u] and ACC-UNet [ibtehaz2023acc]). The VAE-based models were trained in an unsupervised manner using unlabelled data, while the UNet-based models were trained in a supervised manner. Starting from their original 2D implementation, these networks were extended to accept 3D inputs of size 64^3 by substituting all 2D layers with their 3D counterparts.

Supervised models

UNet is a popular encoder-decoder architecture that has shown promising results in many semantic voxel-wise segmentation tasks. MSS-UNet, UNet++, and UNet 3+ are extensions of the original UNet architecture. MSS-UNet incorporates multi-scale guidance in the decoding process during training, enabling it to capture more fine-grained details and to have a more coherent processing of information in the decoding stage. UNet++ includes a nested and dense skip-connection structure to capture more multi-scale features, while UNet 3+ uses a more powerful encoder with multi-resolution inputs. To ensure consistency, UNet and MSS-UNet were built using the same encoding/decoding building blocks as for UNet++ and UNet 3+ [huang2020unet]. This approach made it easier to compare the results of different architectures and understand how they impact the final outcome in voxel-wise segmentation tasks. Vision Transformers have recently addressed complexity challenges, making them a viable and competitive solution for visual tasks, where a notable work is [liu2021swin]. Building on these advancements, the core concepts of Transformers have been integrated into ResNet models, surpassing the performance of Swin Transformers. Another notable development involves incorporating essential Transformer ideas into a convolution-based neural model called ACC-UNet. This model has shown promise in segmentation tasks, motivating its use in the current study. MSS-UNet, UNet++, UNet 3+ and ACC-UNet are deeply supervised during this study, which means they are trained with a loss function calculated on multiple inner layers to supervise the learning process effectively. In contrast, the original UNet architecture is not deeply supervised.

Unsupervised models

The VAE-based models were trained in an unsupervised manner to learn a compressed and disentangled representation of the input data. During training, the VAE models learned to reconstruct images from the compressed representations that resemble the input images as closely as possible. The reconstruction error, which quantifies the discrepancy between the input and output of the unsupervised models, was adopted as the anomaly score. Since the introduction of the VAE model in 2014 by Kingma and Welling, it has been used in a variety of studies for voxel-wise anomaly detection (e.g. [chen2019unsupervised, lin2020anomaly, chatterjee2021unsupervised]). The ceVAE model has similar architecture as VAE but a more complex definition of the loss. During training, ceVAE uses "masked" input data where certain patches within the image are fixed to a specific value. The model uses an ad-hoc loss function to infer the missing or distorted voxels within the masked zone, which helps the network to capture the context of the image. This peculiarity of the model may have a positive impact on the score, since it can prevent the network to learn to represent the pores within the training dataset. On the contrary, the gmVAE and vqVAE models are more complex than the VAE architecture, enabling them to catch features of the input 3D images that could not be interpreted by the coarser architecture of VAE. The gmVAE model assumes that each input data point's latent representation is generated by one of several possible Gaussian distributions, each with a different mean and variance, and identifies which distribution in the mixture is most likely to have generated the latent representation of each input data point during training. The vqVAE model is based on the idea of vector quantisation, where the continuous latent space is discretised into a set of discrete codes. The model comprises an encoder network that maps the input images to a discrete code book, followed by a decoder network that maps the discrete codes to the reconstructed input images. The vqVAE model was adapted to 3D inputs without additional alterations, except for an extra encoding/decoding stage that processes larger input patches of 64^3 instead of the default 32^3 . The RV-VAE model eliminates stochastic sampling, directly incorporating latent space information into decoder layers as continuous random variables. Applying the inherent mathematical prior during decoding leads to a more precise representation, making it appealing for segmentation tasks. As a final sigmoid activation function was used for all the neural models, the related RV-model for this function is provided in the appendix A.3.

5.2.3 Training

The deep learning framework was based on the Pytorch [paszke2019pytorch] 2.0.1, Pytorch-lightning [falcon2020pytorchlightning] 2.0.2 and the CUDA [cuda] 11.6 libraries and it is publicly available (https://github.com/snipdome/nm_3D-anomaly-detection). The 3D patch extraction, aggregation and data augmentation were based on the TorchIO libraries [perez2021torchio] version 0.18.84. A unique main seed propagated throughout the libraries ensures that all the extraction from random distributions were reproducible. Each of the models was trained with the Adam optimiser (learning rate of 0.0001) and halted through early stopping when the loss value did not decrease by more than 0.0001 for 40 consecutive epochs.

5.2.4 Focal Tversky Loss function

In this pore segmentation task, the number of voxels belonging to the foreground class (pores) is much smaller than the number of voxels belonging to the background class, in the training dataset. This class imbalance results in a bias towards the background class during training, which leads to poor voxel-wise segmentation performance. In order to address the problem of class imbalance in semantic segmentation tasks, the FTL was proposed as a modification to the Tversky Loss [abraham2019novel], and is defined as follows:

$$\text{FTL} = \left(1 - \frac{\text{TP}}{\text{TP} + \alpha \text{FN} + \beta \text{FP}}\right)^{\gamma} \quad (5.1)$$

The FTL depends on the number of true negatives (TN), false negatives (FN), and false positives (FP), where FN and FP are weighted by α and β , respectively. By adjusting the values of these parameters, the FTL can be fine-tuned to emphasise either precision or recall. In addition, the FTL also includes a parameter γ , which controls the degree to which the FTL prioritises correcting misclassifications by adjusting the weight given to the Tversky Loss function. If $\gamma = 1$, the FTL reduces to the standard Tversky loss and, if is also true that $\alpha = \beta = 0.5$, to the Dice-Sørensen loss. If $\gamma > 1$, the FTL function will assign a higher weight to the correction of misclassifications. This means that the loss function will be more sensitive to false negatives and false positives, and the model will prioritise the correction of misclassifications over the correct classification of the majority class. As a result, the model will be better at identifying instances of the minority class but may struggle to accurately classify instances of the majority class. The degree to which the model's sensitivity to misclassifications increases will depend on the value of γ . In case of deep supervision, the FTL is calculated at each

supervised stage and averaged with geometric progression weights (1, 1/2, 1/4, etc.).

5.2.5 Post-processing

During the prediction or testing procedure, each of the models inferred patches belonging to the X-CT scan and then aggregated them back together to obtain an output volume with the same size as that of the input.

Only for the unsupervised models, the output was post-processed to amend the scarce quality that these models have in representing the fine details of the samples, as the surface. The surface of each of the samples has unique characteristics, due to different printing processes and polishing procedures, which can never be properly represented with an Autoencoder (AE). While AEs are designed to learn a concise representation of the input, their ability to faithfully reproduce high-fidelity images depends on factors such as the training dataset's size and diversity, the complexity of input data, and the model's architecture and hyperparameters. As this is a beneficial feature that makes the AEs potentially unable to reproduce anomalies that may be present in the training dataset, it comes with the cost of inaccuracies near the surface of the samples. To counteract this, a compensation mechanism that suppresses the anomaly score near the sample surface is introduced. The computation of the new voxel-wise anomaly score, denoted as A_{pores} , involves subtracting the spatially blurred derivative D of the inferred volume \hat{V} from the original anomaly score A . As previously mentioned, the neural models struggle to faithfully represent the surface of samples, leading to pronounced derivatives of the inferred volume along the border. The elements of D are determined by the sum of the absolute voxel-wise derivatives in the x , y , and z directions of the predicted volume \hat{V} . These derivatives are represented as $d_{ijk} = \|\partial_x \hat{v}_{ijk}\| + \|\partial_y \hat{v}_{ijk}\| + \|\partial_z \hat{v}_{ijk}\|$, where \hat{v}_{ijk} corresponds to the (i, j, k) voxel in \hat{V} .

The formulation for A_{pores} is expressed as:

$$A_{pores} = \max(0, A - \lambda^* G_{\sigma^*}(D)) \quad (5.2)$$

Values for the standard deviation σ^* of the Gaussian smearing kernel G and the scaling factor λ^* are determined through an on-the-fly optimisation process outlined in Formula 5.3. This optimisation process aims to minimise the disparity between the anomaly score and the Gaussian-blurred absolute sum of derivatives, utilising the mean of the L1-norm as a metric. Both λ and σ are considered to be positive parameters in this context.

The optimisation problem is formally stated as:

$$\lambda^*, \sigma^* = \arg \min_{\lambda, \sigma} \|A - \lambda G_{\sigma}(D)\|_1 \quad (5.3)$$

The results of the experiment 5.3.5 show the benefits of applying the proposed technique.

5.3 Experiments

The X-CT images were organized into training, validation, and testing sets, as explained in section 5.3.1. All models were trained using a common training framework, detailed in section 5.3.2. For the evaluations presented in this section, the labelled X-CT volumes were compared with the output of the DL models, after the output 3D patches were aggregated.

More specifically, the patch-extraction pipeline extracted overlapping patches from the input volume, each with half of their length overlapping with neighboring patches. These patches were segmented by the neural networks and then combined by computing an average value among the overlapping patches. This approach ensured a comprehensive evaluation of the model’s performance on the X-CT volumes.

5.3.1 Dataset

The X-CT images of several AM samples composed the digital dataset for training, validation, and testing of the DL models. In a 5-fold manner, the X-CT images of the cylindrical samples were organised into 4 samples for the train-set and 1 sample for the validation-set. Noise, image artefacts, and misclassified voxel-wise labels (commonly referred to as ‘noisy labels’) can negatively affect training and lead to inaccurate predictions. To mitigate the influence of noisy labels during training and to expand the training sets, data augmentation was employed [song2022learning]. The data augmentation created novel spatial configurations by flipping of patches in random directions and elastic distortion while teaching the networks to be resilient against noise, specific attenuation of samples, and artefacts such as cone-beam and beam-hardening. After data augmentation was applied at every training epoch to each of the cylindrical samples, which have around 800x800x2000 voxels, 3D patches of 64x64x64 voxels were extracted and supplied to the neural networks.

5.3.2 Training

The deep learning framework was based on the Pytorch [paszke2019pytorch] 2.0.1, Pytorch-lightning [falcon2020pytorchlightning] 2.0.2 and the CUDA [cuda] 11.6 libraries and it is publicly available (https://github.com/snigdome/nm_3D-anomaly-detection). The 3D patch extraction, aggregation and data augmentation were based on the TorchIO libraries [perez2021torchio] version 0.18.84. A unique main seed propagated throughout the libraries ensures that all the extraction from random distributions were reproducible. Each of the models was trained with the Adam optimiser (learning rate of 0.0001) and halted through early stopping when the loss value did not decrease by more than 0.0001 for 40 consecutive epochs.

5.3.3 Parameter search for the FTL function

As different values of the α, β, γ parameters sensibly affect the performance of models trained with the FTL function [iuso2022evaluation], the optimal values were identified with grid search approach. A 5-fold cross-validation strategy evaluated the performance of the model with different parameter combinations, while the γ parameter was kept at 0.5 (as in [iuso2022evaluation]). The grid search space spanned the parameter-space uniformly from 0.1 to 0.9 for each of the variables, for a total of 4 steps. For each combination of α and β , the model was trained in a 5-fold cross-validation, resulting in a total of 16 different combinations of α and β and a total of 80 model trainings. In addition to the α and β parameters, another grid search identified the optimal γ parameter in the FTL. A higher value of γ puts more emphasis on minimising false positives and false negatives, which can be useful in tasks where the cost of misclassification is high. So, even though the author of the FTL had suggested a value of 4/3 for the γ parameter [abraham2019novel], the optimal γ parameter turned out to vary for the current application of this work. The γ grid search had a total of 8 steps ranging from 1/3 to 2, for a total of 40 trainings.

5.3.4 Cross-validation of performance of the DL models

All the supervised and unsupervised models have been trained in a 5-fold cross-validation, for a total of 50 trainings. In the case of supervised models, they were trained with the optimal parameters found during the experiment 5.3.3. After training, the performance has been evaluated, for each fold, on both the validation-set and the challenging test-set.

5.3.5 Cross-validation of performance of post-processed unsupervised models

For this experiment, the unsupervised models are compared in cross-validation before and after the application of a post-processing algorithm presented in section 5.2.5. Since the post-processing happens after the aggregation of all the patches composing a X-CT volume, it is possible to compare the models before and after post-processing, without the need to re-train the models. Also in this case, the performance has been evaluated, for each fold, on both the validation-set and the challenging test-set.

5.3.6 Cross-validation of performance of supervised models re-trained with unsupervised models

In this experiment, the anomaly score of the (best performing) unsupervised model of experiment 5.3.5 was used as label for the training of supervised models, for each fold. Training in such a way would make the overall pipeline unsupervised, which, apart from being a favourable feature for the user, it would theoretically allow the UNet-family to reproduce the task of the unsupervised model (and its post-processing algorithm). A total of 25 trainings has been performed.

5.3.7 Model complexity

For this experiment, all the neural models have been compared with regards to their memory footprint and computational cost. The networks were fed with a one-element batch with size $1 \times 64 \times 64 \times 64$ and analysed during their complete forward and backward operation.

5.3.8 Cross-validation of performance of the best performing model in extreme visual scenarios

In this final experiment, the best performing model in the previous experiments has been tested when the image quality of the challenging test-set has been worsened by lowering X-ray exposure and number of projections. This test is designed to show how the performance decreases in extreme visual scenarios. The number of X-ray projections of the challenging test-set was reduced to 50% and 33.3%. The simulation of lower exposure of X-ray projections is achieved by adding Poisson distributed noise. The exposure was lowered to 75%, 50% and 25% of the original values, which corresponded in an increase in the imaging noise over the X-ray projections.

Dice-Sørensen score					
α	0.9	0.74 ± 0.06	0.74 ± 0.04	0.75 ± 0.07	0.75 ± 0.07
	0.63	0.78 ± 0.04	0.76 ± 0.06	0.76 ± 0.06	0.73 ± 0.07
	0.37	0.75 ± 0.05	0.76 ± 0.07	0.74 ± 0.07	0.70 ± 0.05
	0.1	0.76 ± 0.06	0.73 ± 0.07	0.72 ± 0.07	0.71 ± 0.06
		0.1	0.37	0.63	0.9
		β			

Table 5.1: Average Dice-Sørensen score and standard error of the models evaluated across the related validation dataset, depending on the α/β parameters of the FTL.

5.4 Results and discussions

Section 5.4.1 presents the cross-validation results for selecting the optimal parameters of the FTL function. These parameters were used to train all the supervised models employed in the voxel-wise segmentation task cross-validation, whose results are shown in sections 5.4.2 and 5.4.4. Section 5.4.4 compares the supervised models trained with the FTL function using heuristic labels and labels generated by the post-processed output of the best performing unsupervised model. The best performing unsupervised model was established based on the performance results presented in Section 5.4.3.

5.4.1 Parameter search for the FTL function

The initial parameter search for α and β has been conducted on all the folds of the cross-validation, and the average results are shown in Table 5.1. As apparent from the results, the optimal values for the α and β parameters are 0.633 and 0.1, respectively. Subsequently, with these optimal parameters, the optimal γ parameter has been searched for each fold, and summary results are shown in Fig. 5.2. In this case, there is good agreement among folds that $\gamma = 1$ ensures the best performance. For the sake of completeness, the fold-wise results have been included in the appendix for both parameter searches (A.2).

Dice-Sørensen score								
	0.76 \pm 0.04	0.78 \pm 0.04	0.79 \pm 0.03	0.82 \pm 0.04	0.72 \pm 0.03	0.70 \pm 0.06	0.64 \pm 0.06	0.59 \pm 0.04
γ	0.33	0.5	0.67	1.0	1.33	1.5	1.67	2.0

Table 5.2: Average Dice-Sørensen score and standard error of the models evaluated across the related validation dataset, depending on the γ parameters of the FTL.

Model	AP	AUC
MSS-UNet \triangle	0.784 \pm 0.050	0.975 \pm 0.013
UNet \triangle	0.815 \pm 0.025	0.982 \pm 0.009
UNet++ \triangle	0.750 \pm 0.026	0.974 \pm 0.009
UNet-3+ \triangle	0.873 \pm 0.036	0.992 \pm 0.003
ACC-UNet \triangle	0.658 \pm 0.078	0.955 \pm 0.014
VAE \diamond	0.711 \pm 0.101	0.999 \pm 0.001
ceVAE \diamond	0.746 \pm 0.094	0.999 \pm 0.001
gmVAE \diamond	0.607 \pm 0.156	0.974 \pm 0.014
vqVAE \diamond	0.602 \pm 0.129	0.990 \pm 0.004
RV-VAE \diamond	0.728 \pm 0.082	0.999 \pm 0.001

Table 5.3: Average ROC-AUC and AP scores (with confidence interval) of the supervised (\triangle) and unsupervised (\diamond) models evaluated on the validation dataset.

5.4.2 Cross-validation of performance of the DL models

The segmentation results of the cross-validation technique were evaluated using two metrics: the cro:AUCArea under the curve (AUC) of the cro:ROCReceiver operating characteristic (ROC) curve and the cro:APAverage precision (AP) of the cro:PRPrecision-recall (PR) curve. While the AUC is a commonly used metric, it can be misleading in the presence of class imbalance [hanczar2010small, saito2015precision]. To address this issue, PR curves were used to evaluate the performance of algorithms, as recommended by [saito2015precision]. Therefore, both PR and ROC curves were used to evaluate the models.

The voxel-wise classification task of the models was evaluated for each fold, whose summary ROC-AUC and AP values are shown in Fig. 5.3 for the validation dataset and the challenging test set. The cross-validated results related to the validation dataset (represented with blue colour in Fig. 5.3) indicate that




Figure 5.3: Point-plots of the average ROC-AUC and AP scores (with confidence interval) of the models evaluated on the validation dataset and on the challenging dataset. The quantitative values are shown in Table 5.3 and Table 5.4.

Model	AP	AUC
MSS-UNet \triangle	0.572 ± 0.019	0.856 ± 0.017
UNet \triangle	0.581 ± 0.008	0.880 ± 0.014
UNet++ \triangle	0.583 ± 0.021	0.848 ± 0.014
UNet-3+ \triangle	0.541 ± 0.010	0.882 ± 0.008
ACC-UNet \triangle	0.418 ± 0.018	0.786 ± 0.016
VAE \diamond	0.615 ± 0.038	0.990 ± 0.002
ceVAE \diamond	0.635 ± 0.021	0.990 ± 0.001
gmVAE \diamond	0.374 ± 0.092	0.838 ± 0.044
vqVAE \diamond	0.313 ± 0.025	0.871 ± 0.009
RV-VAE \diamond	0.634 ± 0.014	0.985 ± 0.001

Table 5.4: Average ROC-AUC and AP (with confidence interval) of the supervised (\triangle) and unsupervised (\diamond) models evaluated on the challenging test-set.

supervised models have been generally better trained to be consistent with labels than the unsupervised methods. The results on the challenging dataset with high artefacts and manually labelled (represented with orange colour) show a clear drop of the score for all the models, as expected for the considerations in 5.2.1. Moreover, it is noticeable that the score of some of the unsupervised models is even higher than that of the supervised ones for the challenging dataset. Although these results may not seem consistent with the validation dataset, it should be noted that in both cases the labels were generated in different ways: either with a heuristic labelling algorithm or via manual annotation. Among the supervised models, there is no significant difference in performance, which suggests that deep supervision and the different architecture of the models is not inducing a significant difference in performance. On the other hand, a noticeable difference in scores is present between ceVAE and gmVAE/vqVAE on the challenging dataset, which is significant for vqVAE with a confidence of 95% (Welch's t-test, p-value $1.98 \cdot 10^{-4}$ (AUC) and $1.05 \cdot 10^{-5}$ AP). The higher degree of complexity of gmVAE and vqVAE is not favourable to the segmentation task by the mean of the anomaly score. These models have been capable of learning how to reproduce defects within the input samples, so the reconstruction error is not as high in the proximity of defects as it is with simpler VAEs. On another note, VAE and ceVAE are most robust with respect to the quality of the input image, since the AP/AUC scores are almost unvaried between the validation and the challenging test-set (AP/AUC differences lower than or approximately equal to a decimal point), when opposed to the other models (AP/AUC differences

Model	AP	AUC
VAE	0.964 ± 0.020	0.998 ± 0.001
ceVAE	0.964 ± 0.021	0.999 ± 0.001
gmVAE	0.516 ± 0.144	0.913 ± 0.033
vqVAE	0.512 ± 0.122	0.948 ± 0.019
RV-VAE	0.951 ± 0.020	0.998 ± 0.001

Table 5.5: Average ROC-AUC and AP (with confidence interval) of the unsupervised models evaluated on the validation dataset, with post-processing of the output. Solely the performance of unsupervised models is shown, since the post-processing of the output is defined for them only.

Model	AP	AUC
VAE	0.824 ± 0.007	0.989 ± 0.002
ceVAE	0.830 ± 0.003	0.989 ± 0.001
gmVAE	0.234 ± 0.089	0.555 ± 0.099
vqVAE	0.138 ± 0.020	0.587 ± 0.028
RV-VAE	0.777 ± 0.004	0.981 ± 0.001

Table 5.6: Average ROC-AUC and AP (with confidence interval) of the unsupervised models evaluated on the challenging test set, with post-processing of the output. Solely the performance of unsupervised models is shown, since the post-processing of the output is defined for them only.

exceeding a decimal point).

5.4.3 Cross-validation of performance of post-processed unsupervised models

By applying post-processing to the output of the VAE models (Fig. 5.4), the considerations of the previous section about supervised models become more evident. When post-processing is applied to the output of the VAE and ceVAE models, which have not learned to visually represent pores, their AP scores increase by almost 2 decimal points on both datasets, while their AUC remains almost unchanged. On the other hand, post-processing adversely affected the performance of gmVAE and vqVAE, which is to be expected since the derivative of the output of these models is non-negligible near the edge of the sample as well as near the pores. This behaviour is noticeable in the ROC and PR classifier curves for the




Figure 5.4: Point-plots of the average ROC-AUC and AP of the models (with confidence interval) evaluated on the validation dataset and on the challenging dataset, with and without post-processing. Solely the performance of unsupervised models is shown, since the post-processing of the output is defined for them only. The values in textual form are shown in Table 5.5 and Table 5.6.

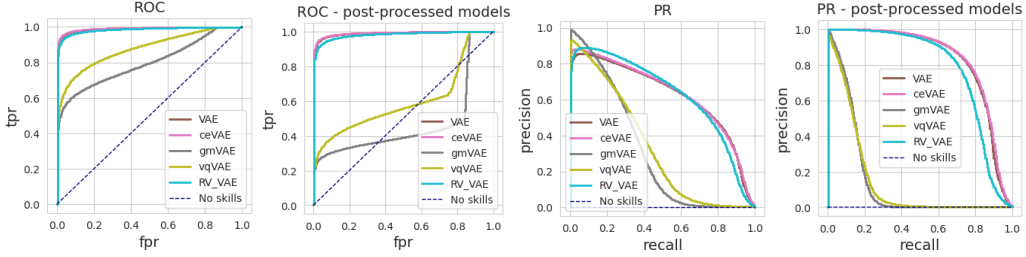


Figure 5.5: Graph of the ROC and PR curves of cross-validated performance for all models. The graphs represent the median trend of the fold-wise performance on the challenging dataset without (left) and with post-processing (right) of the aggregated output.

Model	AP	AUC
MSS-UNet	0.651 ± 0.008	0.889 ± 0.005
UNet	0.639 ± 0.008	0.882 ± 0.004
UNet++	0.751 ± 0.030	0.902 ± 0.015
UNet-3+	0.627 ± 0.006	0.894 ± 0.006
ACC-UNet	0.586 ± 0.008	0.874 ± 0.004

Table 5.7: Average ROC-AUC and AP (with confidence interval) of the supervised models re-trained with the labels generated by ceVAE and evaluated on the challenging dataset.

challenging case as shown in Fig. 5.5 (other ROC and PR graphs are shown in the A.1). The greater complexity of gmVAE/vqVAE models enables them to replicate defects within the samples, leading to a reduction in anomaly scores and compromising performance. This effect intensifies with the application of post-processing, as illustrated in Fig. 5.6, where a validation sample is inferred by both ceVAE and gmVAE with and without post-processing of the anomaly scores. These results highlight that a more complex architecture is not always advantageous, particularly when anomalies exist within the training dataset. Additionally, it can be observed from Fig. 5.3 and Fig. 5.4 that the scores of VAE and ceVAE are still resilient against the poor image quality of the challenging test-set, compared to the drastic drop in performance of the supervised networks.




Figure 5.6: A slice took from a validation dataset (a) and its voxel-wise anomaly score accordingly to ceVAE (b) and gmVAE (d). Post-processing the anomaly scores (c, e) reveals a beneficial impact, particularly for models that unequivocally classify pores as anomalies. The color-scale represents the intensity levels in the anomaly score images.

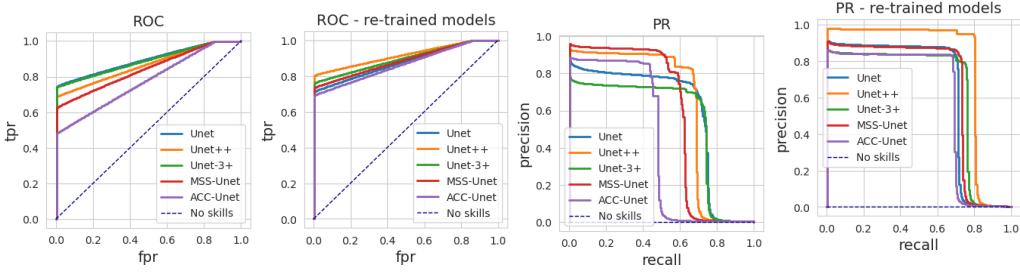


Figure 5.7: Graph of the ROC and PR curves of cross-validated performance for all models. The graphs represent the median trend of the fold-wise performance on the challenging dataset, with Otsu-based labels (left) and post-processed ceVAE-generated labels (right).

5.4.4 Cross-validation of supervised models trained with labels generated by an unsupervised model

By using ceVAE (the best performing model) to generate labels for the samples, the supervised models could be trained from scratch to detect pores. The necessary steps for the production of these labels by ceVAE were the post-processing (with the algorithm described in section 5.2.1) and the suppression of smaller pores. The results are shown in Fig. 5.8 and Fig. 5.7. Higher performance is achieved by using the unsupervised labels, confirmed by both AUC and AP for all the models. These results confirm the observations in section 5.4.2 that the different architectures of the models are not significantly affecting the scores for this voxel-wise segmentation task.

5.4.5 Model complexity

Table 5.8 presents key metrics related to the model complexity of each neural model, including the number of parameters, peak memory usage, and Multiply-Accumulate Operations (MACs). The number of parameters indicates the quantity of floating-point numbers that need to be stored in video memory, reflecting the minimal memory occupancy required to store the model. Conversely, the forward/backward peak memory highlights the memory needed to process an input with a batch size of 1. Lower memory requirements lead to larger permissible batch sizes, consequently reducing training times. The MACs value encapsulates information about the speed of the neural models to process a single 3D patch. In the case of X-CT volumes sized at 800x800x2000, comprised of numerous overlapped patches by half of their patch-length, the forward operation during the inference phase necessitates multiple repetitions to process the entire volume.




Figure 5.8: Point-plots of the average ROC-AUC and AP (with confidence interval) of the supervised models evaluated on the challenging dataset. The graphs highlight the different performance when these models were supervised by the Otsu-based method and with the labels provided by the unsupervised models. The values in textual form are shown in Table 5.7

Model	# Parameters	Forward/Backward Peak Memory	MACs	Total MACs
MSS-UNet	1.328 M	383.740 / 424.840 MB	14.270 G	69.678 T
UNet	1.325 M	353.924 / 390.925 MB	14.124 G	68.967 T
UNet++	1.503 M	933.490 / 1005.831 MB	34.821 G	170.024 T
UNet-3+	1.672 M	1571.642 / 1720.766 MB	84.881 G	414.460 T
ACC-UNet	5.062 M	6897.734 / 7269.893 MB	39.724 G	193.966 T
VAE	29.024 M	44.703 MB / 189.918 MB	3.698 G	18.058 T
ceVAE	140.650 M	33.765 MB / 778.901 MB	8.344 G	40.742 T
gmVAE	383.650 M	774.129 MB / 1842.710 MB	207.48 G	1013.096 T
vqVAE	2.511 M	17.688 MB / 32.701 MB	8.471 G	41.361 T
RV-VAE	29.024 M	223.288 MB / 230.196 MB	0.456 G	2.176 T

Table 5.8: Model complexity metrics for each neural model, including forward/backward peak memory usage and MACs, are specified for batch-size 1. Total MACs represent operations for processing an 800x800x2000 voxel volume, with a 3D patch overlap of half the patch-length.

The cumulative MACs operations, represented as "Total MACs" in the table, quantify the overall computational workload.

It is noteworthy that the memory usage of the UNet-family generally exceeds that of the VAE-family in forward/backward passes, with the exceptions of ceVAE and gmVAE. Specifically, the high memory requirements of ceVAE are visible only during the training procedure, as it is related solely to the backward pass. Nevertheless, ceVAE has shown good performance during the previous experiments (Section 5.4.3 and 5.4.4). Conversely, the huge memory requirement of gmVAE and MACs do not directly translate in outstanding performance for the prior experiments.

5.4.6 Cross-validation of performance of the best performing model in extreme visual scenarios

By reducing the number of X-ray projections of the challenging X-CT scan and reducing the exposure of each X-ray projection, the quality of the reconstructed X-CT scan decreased. The best performing model, which was shown to be the post-processed ceVAE, was applied to these X-CT scans. An exemplary visual representation of the voxel-wise segmentation is shown in Fig. 5.10, related to the post-processed output of the ceVAE model, trained on the 1st fold. In this figure, a small portion of a slice of the cube is shown, in which pores are visible that were induced with off-nominal parameters of the melting laser during the printing. The degradation of the segmentation performance is noticeable due to the increasing number of voxels classified as pores (as shown in Fig. 5.9). Inter-

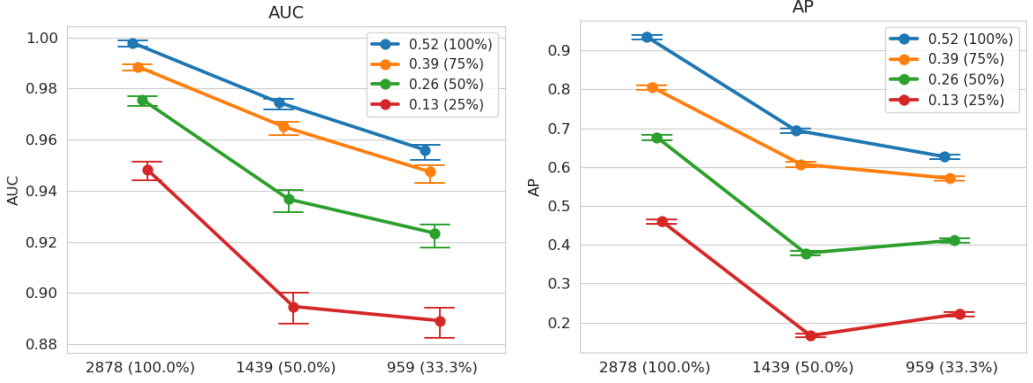


Figure 5.9: Point-plots of the average ROC-AUC and AP (with confidence interval) of the anomaly score of ceVAE evaluated on the challenging test-set when the image quality is lowered by reducing the number of X-ray projections or exposure.

estingly, while reducing the number of X-ray projections from 4283 (the dataset used for training/validation) to 2878 (the original challenging test-set) did not significantly affect the performance (Fig. 5.4), further reductions in the number of projections had a significant impact on the performance scores (Fig. 5.9). Another point to note is the trend exhibited by the AP scores at low exposure levels ranging from 50-25%. Specifically, reducing the number of projections from 50% to 33.3% led to a slight increase in the AP scores. When data is highly noisy and the number of projections is relatively low, adding some more X-ray projections may not always lead to better image quality of the reconstructed X-CT scans. This is because the additional (noisy) projections can also introduce more noise into the reconstructed images. This can be observed from the fact that the trend gradually disappears as the exposure level increases from 25% to 100%.

5.5 Conclusions

This study explores recent Deep Learning techniques for voxel-wise pore segmentation in X-CT images of AM samples. Employing Tversky focal loss, deep supervision, and 3D patch-based training, various 2D neural models (UNet, UNet++, UNet 3+, MSS-UNet, ACC-UNet, VAE, ceVAE, gmVAE, vqVAE, RV-VAE) were adapted to 3D, with both supervised and unsupervised training strategies. Post-processing of unsupervised models and training supervised models with unsuper-




Figure 5.10: A portion of a X-CT slice is shown in each row and column by modifying the number of X-ray projections and exposure of each X-ray projection. Each input slice is shown together with the label mask predicted by ceVAE (trained on the 1st fold). The degradation of the segmentation performance is noticeable from the raising number of voxels that are classified as pores (white colour in the predicted mask).

vised inferred labels are also investigated.

The comprehensive comparison of all neural models reveals that supervised models (UNet-3+, AP 0.873 ± 0.036) outperform unsupervised models (ceVAE, AP 0.746 ± 0.094), a trend not upheld when tested on a challenging X-CT test set. In this scenario, ceVAE (AP 0.635 ± 0.021) outperforms supervised neural models (UNet++, AP 0.583 ± 0.021). The application of additional post-processing, beneficial for VAE and ceVAE (AP 0.830 ± 0.003 on the challenging test set), proves counterproductive for gmVAE and vqVAE due to the more complex architecture of these models. This complexity lead the models to be able to replicate defects within the training samples, thereby impairing the voxel-wise anomaly score. Although using an ideal pore-free training dataset might improve the scores of gmVAE and vqVAE models, it would hinder supervised models' performance due to the absence of pores. Overall, the resulting VAE/ceVAE models exhibit resilience to lower image quality, unlike supervised models.

Training supervised models with labels derived from the best unsupervised model (ceVAE) enhances their performance (UNet++, AP from 0.583 ± 0.021 to 0.751 ± 0.030 on the challenging testset) but does not surpass that of the unsupervised model. The study confirms that unsupervised ceVAE, robustly captures the statistical properties of 3D patches compared to the supervised UNet family. This finding aligns with analogous results in anomaly detection in MRI images [**chatterjee2022strega**], endorsing unsupervised learning as a viable training paradigm for addressing anomaly segmentation in AM samples without the need for labelled data.

Looking ahead, future endeavours may involve developing efficient models capable of detecting pores from X-CT scans at a faster rate, with fewer projections or shorter scan times, in coherence with the future trends foreseen by Khosravani&Reinicke [**khosravani2020use**], which will expand the experiment 5.3.8. This would facilitate the use of X-CT in streamlined evaluations of entire sample batches. Furthermore, while the research presented here primarily focuses on porosity analysis in the AM process, it opens avenues for broader anomaly detection applications, including identifying impurities, microstructural inhomogeneities, or alloying element loss due to vaporisation.

General conclusion & prospects

With the work presented in this thesis, the possible applications of X-ray devices in the AM production pipeline have been extended and new paths for further investigation have been opened.

Firstly, a novel automated registration method for comparing CAD models with their physical realisations through a few X-ray radiographs has been developed. This method eases the automated rejection of defective AM parts by analyzing projective residuals, which measure discrepancies between the actual X-ray radiographs and their ideal, defect-free counterparts. The method relies on the application of a mesh projector, implemented as a differential program specifically for this task. The automated registration method has been proven effective on real radiographs utilizing a poly-chromatic model and performing simultaneous registration of multiple meshes, based solely on scanning geometry and CAD model information, achieved without prior training procedures or spectral information regarding the scanning system. However, challenges such as high object symmetry and projective model inaccuracies must be addressed to enhance the method's applicability and accuracy. As a future prospect, the method can be investigated to include a joint estimation of the shape of the scanned object by as few as two X-ray projections.

The thesis also introduces a novel approach to compensate for undesirable X-ray scattering effects through a simulation-based method, leveraging the exact positioning of scanned objects and their supports, from just few X-ray radiographs. This marks the first instance of using a simulation-based scatter compensation method without prior X-CT scanning. The adaptability of this method, free from training procedures, allows it to be easily integrated into new factory productions with different designs and materials. While the current use of the GATE simulation back-end is computationally intensive, modern GPU solutions for X-ray scatter simulations present a promising future direction for reducing this computational load. Additionally, developing an X-ray scatter estimation pipeline

relying solely on the mesh projector and an analytical description of first-order scatter is expected to yield significantly improve processing speed.

Finally, the thesis explores the use of DL techniques for analyzing X-CT images, specifically for voxel-wise pore segmentation in AM samples using volumetric information. By employing various 2D neural models adapted to 3D and using supervised and unsupervised learning strategies, the research demonstrates the effectiveness of a 3D patch-based approach for porosity analysis. Results indicate that unsupervised models, particularly ceVAE, outperform supervised models in challenging test scenarios. Future research may focus on accurately segmenting pores with fewer X-ray radiographs and extending these methods to broader anomaly detection applications, such as identifying impurities, microstructural inhomogeneities, or alloying element loss due to vaporization.

In summary, this thesis presents a suite of innovative methodologies that enhance the efficiency and applicability of X-ray imaging in industrial AM processes. The advancements in automated multi-mesh registration, X-ray scatter compensation, and deep learning-based porosity analysis represent significant strides forward, setting the stage for continued research and development in these areas.

Appendices & Supplementary material

A.1 Classifier graphs for the voxel-wise segmentation task

The ROC and PR graphs of the voxel-wise segmentation results that were not shown in previous sections are reported here. In Fig. A.1, there are the performance graphs of supervised and unsupervised models evaluated on the related validation dataset. The graphs are aligned with the findings discussed in section 5.4.2 and 5.4.4. In Fig. A.2 are shown the performance of the unsupervised models only, since they show the segmentation scores of the post-processed output. The scores were obtained from the fold-wise performance on the related validation dataset, where is noticeable an increase of performance for VAE/ceVAE and a decrease for gmVAE/vqVAE if compared with Fig. A.1 (right), in accordance with the findings in section 5.4.4.

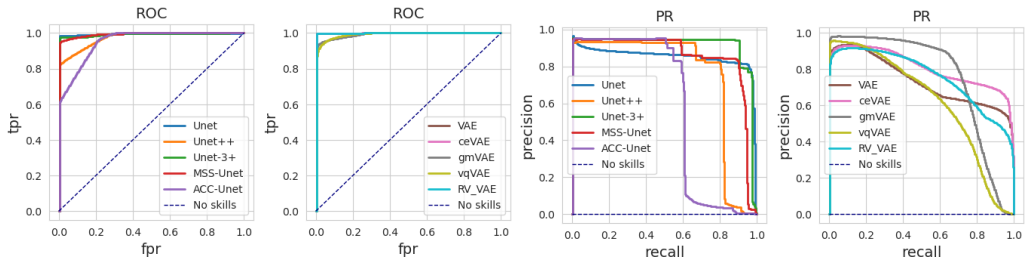


Figure A.1: Graph of the ROC and PR curves of cross-validated performance for all models. The graphs represent the median trend of the fold-wise performance on related validation dataset.

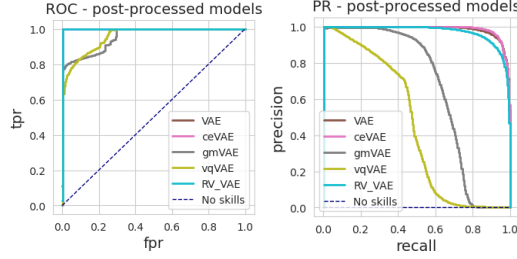


Figure A.2: Graph of the ROC and PR curves of cross-validated performance for the unsupervised models. The graphs represent the median trend of the fold-wise performance on the related validation dataset, when the output of the models is post processed.

A.2 Cross-validation graphs for the FTL parameter search per each fold

For each of the 5 folds of the cross-validation, there is a total of 16 trainings for the α/β parameter, which are presented in Table A.1. For the γ parameter, there is a total of 8 trainings per fold and the values of the Dice-Sørensen are shown in Table A.2.

Dice-Sørensen score - Fold 1					Dice-Sørensen score - Fold 2						
α	0.9	0.82	0.81	0.97	0.97	0.9	0.76	0.63	0.61	0.60	
	0.63	0.86	0.97	0.97	0.96	0.63	0.69	0.61	0.60	0.59	
	0.37	0.96	0.97	0.96	0.83	0.37	0.64	0.60	0.58	0.57	
	0.1	0.97	0.95	0.94	0.84	0.1	0.63	0.54	0.53	0.57	
		0.1	0.37	0.63	0.9			0.1	0.37	0.63	0.9
		β						β			

Dice-Sørensen score - Fold 3					Dice-Sørensen score - Fold 4						
α	0.9	0.94	0.88	0.87	0.88	0.9	0.62	0.71	0.68	0.69	
	0.63	0.92	0.88	0.87	0.85	0.63	0.66	0.72	0.74	0.69	
	0.37	0.70	0.87	0.85	0.83	0.37	0.70	0.75	0.69	0.72	
	0.1	0.87	0.82	0.80	0.79	0.1	0.74	0.78	0.78	0.79	
		0.1	0.37	0.63	0.9			0.1	0.37	0.63	0.9
		β						β			

Dice-Sørensen score - Fold 5					
α	0.9	0.58	0.67	0.61	0.61
	0.63	0.78	0.62	0.63	0.58
	0.37	0.77	0.61	0.59	0.57
	0.1	0.61	0.56	0.55	0.55
		0.1	0.37	0.63	0.9
		β			

Dice-Sørensen score						
γ	2.0	0.50	0.60	0.71	0.52	0.64
	1.67	0.45	0.75	0.82	0.58	0.62
	1.5	0.49	0.75	0.86	0.61	0.80
	1.33	0.67	0.74	0.84	0.64	0.74
	1	0.68	0.81	0.91	0.80	0.87
	0.67	0.72	0.81	0.90	0.70	0.84
	0.5	0.86	0.69	0.92	0.66	0.78
	0.33	0.75	0.66	0.90	0.70	0.80
	1	2	3	4	5	
Fold						

Table A.2: Fold-wise Dice-Sørensen score for the networks evaluated on the related validation dataset, depending on the γ parameter of the FTL.

A.3 Random Variable Module: Sigmoid Activation Function

In this section, we extend the discussion on random variables applied after the encoding layer of Autoencoder-based neural models, as presented in a previous article [nicodemou2023rv]. We maintain the assumptions established in that work, which include the absence of correlations between random variables. Furthermore, we leverage the ability to represent arbitrary probability distributions of real numbers through an expected value and a variance, a condition supported by the validity of the central limit theorem resulting from the summation of unrelated random variables.

Our focus here is to provide a means of obtaining the first two moments (expected value and variance) of a random variable Y resulting from the application of the sigmoid function S to its input random variable X .

Let us begin by defining the sigmoid function:

$$S(x) = \frac{1}{1 + \exp(-x)} \quad , \quad (\text{A.1})$$

alongside its first and second derivatives with respect to x

$$\dot{S}(x) = S(x)(1 - S(x)) \quad \ddot{S}(x) = S(x)(1 - S(x))(1 - 2S(x)) . \quad (\text{A.2})$$

These derivatives will prove useful in deriving the expected value and variance of $Y = S(X)$, where X is considered to be a random variable.

For the calculation of $\mathbb{E}[S(X)]$, we employ a Taylor expansion centred at $X_0 = \mathbb{E}[X]$:

$$\begin{aligned} S(x) &= S(\mathbb{E}[X]) + (X - \mathbb{E}[X])\dot{S}(\mathbb{E}[X]) + \frac{1}{2}(X - \mathbb{E}[X])^2\ddot{S}(\mathbb{E}[X])^2 \\ &\quad + \frac{1}{3!}(X - \mathbb{E}[X])^3\ddot{\ddot{S}}(\mathbb{E}[X])^3 + \dots \quad . \end{aligned} \quad (\text{A.3})$$

From which we extract the expected value as

$$\begin{aligned} \mathbb{E}[S(x)] &= \mathbb{E}[S(\mathbb{E}[X]) + (X - \mathbb{E}[X])\dot{S}(\mathbb{E}[X]) + \frac{1}{2}(X - \mathbb{E}[X])^2\ddot{S}(\mathbb{E}[X])^2 \\ &\quad + \frac{1}{3!}(X - \mathbb{E}[X])^3\ddot{\ddot{S}}(\mathbb{E}[X])^3 + \dots] \quad . \end{aligned} \quad (\text{A.4})$$

Given the assumption that the distribution of the random variable X behaves as a normal distribution, all odd central moments are expected to be null. This leads to a simplified formula for the expected value of $Y = S(X)$

$$\mathbb{E}[Y] = S(\mathbb{E}[X]) + \frac{1}{2}\ddot{S}(\mathbb{E}[X])\text{Var}[X] + M_4 \quad , \quad (\text{A.5})$$

where M_4 collects all the moments after the third and can be neglected under the assumption of smooth distribution. To calculate the expected variance of Y , we can utilise equations A.3 and A.5, so to obtain

$$\begin{aligned}
 \mathbb{V}\text{ar}[Y] &= \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 = \mathbb{E}[S^2(\mathbb{E}[X]) + 2(X - \mathbb{E}[X])S(\mathbb{E}[X])\dot{S}(\mathbb{E}[X]) \\
 &\quad + (X - \mathbb{E}[X])^2(\dot{S}^2(\mathbb{E}[X]) + S(\mathbb{E}[X])S''^2(\mathbb{E}[X])) \\
 &\quad + (X - \mathbb{E}[X])^3(\dot{S}(\mathbb{E}[X])\ddot{S}(\mathbb{E}[X]) + \frac{2}{3!}S(\mathbb{E}[X])\ddot{S}(\mathbb{E}[X])) + R_4] \\
 &\quad - S^2(\mathbb{E}[X]) - \frac{1}{4}\ddot{S}^2(\mathbb{E}[X])\mathbb{V}\text{ar}^2[X] - S(\mathbb{E}[X])\ddot{S}(\mathbb{E}[X])\mathbb{V}\text{ar}[X] - \tilde{M}_4,
 \end{aligned} \tag{A.6}$$

with \tilde{M}_4 being analogous to M_4 in A.5 and R_4 collecting all the central differences above the third exponent. By discarding all moments above the third, a compact approximation for the variance of Y is given by

$$\mathbb{V}\text{ar}[Y] \approx \dot{S}^2(\mathbb{E}[X])\mathbb{V}\text{ar}[X] - \frac{1}{4}\ddot{S}^2(\mathbb{E}[X])\mathbb{V}\text{ar}^2[X]. \tag{A.7}$$

A.4 Theoretical white-field image for flat detectors

This section provides a practical approximation for the amount of radiation crossing a detector pixel in the case of isotropic X-ray emission.

Assuming isotropic X-ray emission with total emission denoted as J^* , the emission per steradian is a constant referred to as $J = J^*/4\pi$. The radiation crossing a detector pixel p is given by the integration of the emission J over the solid angle $d\omega$ related to the detector pixel area A_p :

$$I_p = \int_{A_p} J d\omega \quad , \quad (\text{A.8})$$

where I_p represents the radiation crossing p . The infinitesimal $d\omega$ is linked to the infinitesimal area da through $d\omega = \mathbf{n} \cdot \hat{\mathbf{v}} da / r^2$, with \mathbf{n} being the detector pixel normal and $\hat{\mathbf{r}}$ the versor of the line r connecting the detector pixel center and the X-ray source position. As the versor $\hat{\mathbf{r}} = (x, y, z)$ has components

$$\begin{cases} x = \sin \theta \cos \phi \\ y = \sin \theta \sin \phi \\ z = \cos \theta \end{cases} \quad , \quad (\text{A.9})$$


then it is possible to rewrite Equation A.8 as

$$I_p = J \int_{A_p} \frac{\mathbf{n} \cdot \hat{\mathbf{v}}}{r^2} da = J \int_{A_p} \frac{\sin \theta \cos \phi}{r^2} da \quad , \quad (\text{A.10})$$

and, expressing it in terms of the position of the detector pixel center $P(X_p, Y_p)$, it becomes

$$I_p = J \int_{A_p} \frac{R da}{(R^2 + X_p^2 + Y_p^2)^{\frac{3}{2}}} \quad . \quad (\text{A.11})$$

Under the assumption that the detector distance R from the source is much greater than X_p and Y_p and of the pixel area A , then the denominator is nearly



constant throughout the integration on the pixel area and it can be brought outside the integration. This allows us to conveniently approximate Equation A.11 with

$$I_p \approx \frac{JAR}{(R^2 + X_p^2 + Y_p^2)^{\frac{3}{2}}} \quad . \quad (\text{A.12})$$

A.5 Supplementary material for Chapter 3

Scene	Diff. Angle	Transl. x	Transl. y	Transl. z	Rot. Angle	Rot. Axis x	Rot. Axis y	Rot. Axis z
cantilever	90	-1.09 ± 0.82	-1.30 ± 1.72	2.28 ± 0.07	64.01 ± 36.09	-0.00 ± 0.00	-0.04 ± 0.04	1.00 ± 0.00
cantilever	50	0.21 ± 0.39	0.59 ± 1.85	2.20 ± 0.01	-86.36 ± 45.97	-0.00 ± 0.00	-0.00 ± 0.00	1.00 ± 0.00
cantilever	10	-0.03 ± 0.36	-0.98 ± 1.59	2.28 ± 0.02	95.86 ± 45.96	-0.00 ± 0.00	0.00 ± 0.00	1.00 ± 0.00
cylinder	90	-0.02 ± 0.00	-0.01 ± 0.00	-0.99 ± 0.00	-103.76 ± 0.07	-0.01 ± 0.00	0.01 ± 0.00	1.00 ± 0.00
cylinder	50	-0.02 ± 0.00	-0.01 ± 0.00	-0.99 ± 0.00	-103.98 ± 0.04	-0.01 ± 0.00	0.01 ± 0.00	1.00 ± 0.00
cylinder	10	-0.03 ± 0.02	-0.01 ± 0.02	-0.99 ± 0.00	-104.05 ± 0.18	-0.01 ± 0.00	0.01 ± 0.00	1.00 ± 0.00
stepwedge	90	1.17 ± 1.07	2.78 ± 1.75	-1.22 ± 0.33	177.41 ± 1.41	0.43 ± 0.18	-0.82 ± 0.07	0.02 ± 0.01
stepwedge	50	2.01 ± 2.41	0.72 ± 0.68	-1.01 ± 0.25	178.08 ± 1.70	0.14 ± 0.35	-0.42 ± 0.28	0.02 ± 0.02
stepwedge	10	1.93 ± 1.94	1.27 ± 1.02	-0.99 ± 0.23	178.03 ± 1.76	-0.14 ± 0.35	-0.14 ± 0.35	0.02 ± 0.02

Table A.3: Mean and standard error of pose parameters for the main object across various scenes and projection angles. Units are in mm and deg.

Scene	Diff. Angle	Transl. x	Transl. y	Transl. z	Rot. Angle	Rot. Axis x	Rot. Axis y	Rot. Axis z
cantilever	90	1.14 ± 0.06	-0.51 ± 0.04	-17.99 ± 0.04	8.00 ± 15.19	0.09 ± 0.06	0.20 ± 0.15	0.91 ± 0.07
cantilever	50	1.15 ± 0.07	-0.51 ± 0.05	-18.03 ± 0.01	-1.25 ± 19.70	0.01 ± 0.01	-0.06 ± 0.06	0.99 ± 0.01
cantilever	10	1.06 ± 0.09	-0.41 ± 0.11	-17.89 ± 0.02	10.60 ± 24.04	0.01 ± 0.00	0.05 ± 0.04	1.00 ± 0.00
cylinder	90	1.03 ± 0.43	-0.28 ± 0.40	-12.68 ± 0.02	-90.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	1.00 ± 0.00
cylinder	50	0.61 ± 0.68	-0.11 ± 0.68	-12.49 ± 0.05	-93.57 ± 3.22	0.00 ± 0.00	0.00 ± 0.00	1.00 ± 0.00
cylinder	10	0.48 ± 1.25	0.04 ± 1.28	-12.13 ± 0.16	-82.81 ± 3.95	0.00 ± 0.00	0.00 ± 0.00	1.00 ± 0.00
stepwedge	90	-6.00 ± 2.81	-1.25 ± 2.60	-66.91 ± 0.40	-5.43 ± 4.95	0.00 ± 0.00	0.00 ± 0.00	0.20 ± 0.20
stepwedge	50	-4.10 ± 2.73	1.45 ± 0.27	-66.93 ± 0.62	-4.72 ± 20.89	0.00 ± 0.00	0.00 ± 0.00	0.60 ± 0.24
stepwedge	10	-3.53 ± 1.88	0.61 ± 0.41	-67.07 ± 0.71	-18.15 ± 30.55	0.00 ± 0.00	0.00 ± 0.00	0.60 ± 0.24


Table A.4: Mean and standard error of pose parameters for the upper support across various scenes and projection angles. Units are in mm and deg.

Scene	Diff. Angle	Transl. x	Transl. y	Transl. z	Rot. Angle	Rot. Axis x	Rot. Axis y	Rot. Axis z
stepwedge	90	-0.05 ± 0.23	0.53 ± 0.29	-89.45 ± 0.23	-5.43 ± 4.95	0.00 ± 0.00	0.00 ± 0.00	0.20 ± 0.20
stepwedge	50	-0.53 ± 0.16	0.46 ± 0.48	-89.63 ± 0.23	-4.72 ± 20.89	0.00 ± 0.00	0.00 ± 0.00	0.60 ± 0.24
stepwedge	10	-0.88 ± 0.31	1.51 ± 0.32	-89.18 ± 0.67	-18.15 ± 30.55	0.00 ± 0.00	0.00 ± 0.00	0.60 ± 0.24

Table A.5: Mean and standard error of pose parameters for the lower support across various scenes and projection angles. Units are in mm and deg.


Scientific contributions

Journal articles


1.  **Domenico Iuso**, Pavel Paramonov, Jan De Beenhouwer and Jan Sijbers. "PACS: Projection-driven with Adaptive CADs X-ray Scatter Compensation for Additive Manufacturing inspection." Precision Engineering
2. **Domenico Iuso**, Pavel Paramonov, Jan De Beenhouwer, and Jan Sijbers. "Practical Multi-Mesh Registration for Few-View Poly-Chromatic X-Ray Inspection." Journal of Nondestructive Evaluation 43, no. 2 (2024): 63. <https://doi.org/10.1007/s10921-024-01071-y>
3. **Domenico Iuso**, Soumick Chatterjee, Jan De Beenhouwer, and Jan Sijbers. "Voxel-wise classification for porosity investigation of additive manufactured parts with 3D unsupervised and (deeply) supervised neural networks." arXiv preprint arXiv:2305.07894 (2023). <https://doi.org/10.48550/arXiv.2305.07894>
4. Anh-Tuan Nguyen, Jens Renders, **Domenico Iuso**, Yves Maris, Jeroen Soete, Martine Wevers, Jan Sijbers, and Jan De Beenhouwer. "MIRT: a simultaneous reconstruction and affine motion compensation technique for four dimensional computed tomography (4DCT)." arXiv preprint arXiv:2402.04480 (2024). <https://doi.org/10.48550/arXiv.2402.04480>
5. Pavel Paramonov, Nicholas Francken, Jens Renders, **Domenico Iuso**, Tim Elberfeld, Jan De Beenhouwer and Jan Sijbers. "CAD-ASTRA: a versatile and efficient mesh projector for X-ray tomography with the

- ASTRA-toolbox.” *Optics Express* 32, no. 3 (2024): 3425-3439. <https://doi.org/10.1364/OE.498194>
6. Mathieu Vandecasteele, Rob Heylen, **Domenico Iuso**, Aditi Thanki, Wilfried Philips, Ann Witvrouw, Dries Verhees, and Brian G. Booth. ”Towards material and process agnostic features for the classification of pore types in metal additive manufacturing.” *Materials & Design* 227 (2023): 111757. <https://doi.org/10.1016/j.matdes.2023.111757>
 7. Negar Chabi, **Domenico Iuso**, Oliver Beuing, Bernhard Preim, and Sylvia Saalfeld. ”Self-calibration of C-arm imaging system using interventional instruments during an intracranial biplane angiography.” *International Journal of Computer Assisted Radiology and Surgery* 17, no. 7 (2022): 1355-1366. <https://doi.org/10.1007/s11548-022-02580-9>
 8. Rob Heylen, Aditi Thanki, Dries Verhees, **Domenico Iuso**, Jan De Beenhouwer, Jan Sijbers, Ann Witvrouw, Han Haitjema, and Abdellatif Bey-Temsamani. ”3D total variation denoising in X-CT imaging applied to pore extraction in additively manufactured parts.” *Measurement Science and Technology* 33, no. 4 (2022): 045602. <https://doi.org/10.1088/1361-6501/ac459a>
 9. Ehsan Nazemi, Nathanaël Six, **Domenico Iuso**, Björn De Samber, Jan Sijbers, and Jan De Beenhouwer. ”Monte-Carlo-based estimation of the X-ray energy spectrum for CT artifact reduction.” *Applied Sciences* 11, no. 7 (2021): 3145. <https://doi.org/10.3390/app11073145>

Conference proceedings

1.  Neam Halat, **Domenico Iuso**, Jan Sijbers and Jan De Beenhouwer. ”KBNNet-Based Noise Suppression in Edge Illumination X-Ray Phase Contrast Imaging”, IEEE RTSI 2024
2. Paramonov, Pavel, Joaquim Sanctorum, **Domenico Iuso**, Jan Sijbers, and Jan De Beenhouwer. ”High-resolution tiled X-ray cone-beam CT using the ASTRA toolbox.” In *e-Journal of Nondestructive Testing, 13th Conference on Industrial Computed Tomography (iCT)*, 2023. <https://doi.org/10.58286/29258>

3. **Domenico Iuso**, Soumick Chatterjee, Rob Heylen, Sven Cornelissen, Jan De Beenhouwer and Jan Sijbers. "Evaluation of deeply supervised neural networks for 3D pore segmentation in additive manufacturing." In Developments in X-Ray Tomography XIV, vol. 12242, pp. 345-352. SPIE, 2022. <https://doi.org/10.1117/12.2633318>
4. **Domenico Iuso**, Ehsan Nazemi, Nathanaël Six, Björn De Samber, Jan De Beenhouwer, and Jan Sijbers. "CAD-Based scatter compensation for polychromatic reconstruction of additive manufactured parts." In 2021 IEEE International Conference on Image Processing (ICIP), pp. 2948-2952. IEEE, 2021. <https://doi.org/10.1109/ICIP42928.2021.9506536>



Abbreviations

X-CT X-ray Computed Tomography

AM Additive Manufacturing

CAD Computer-Aided Design

DXR Digital X-ray radiography

SLA Stereolithography

SLM Selective Laser Melting

KH Keyhole

LoF Lack-of-fusion

DL Deep Learning

FTL Focal Tversky Loss

SS316L Stainless steel 316L

PA12 Polyamide

AP Average precision

PR Precision-recall

ROC Receiver operating characteristic

AUC Area under the curve

RMSE Root mean squared error

NDT Non-destructive testing

VAE Variational Autoencoder

FTL Focal Tversky Loss

PACS Projection-driven Adaptive CADs X-ray Scatter compensation