



HAL
open science

Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts

Franck Vidal, Iwan Mitchell, J.-M. Letang

► To cite this version:

Franck Vidal, Iwan Mitchell, J.-M. Letang. Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts. *Precision Engineering*, 2022, 74, pp.110-125. 10.1016/j.precisioneng.2021.10.014 . hal-03409039

HAL Id: hal-03409039

<https://hal.science/hal-03409039v1>

Submitted on 3 Dec 2021


HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Graphical Abstract

Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts

Franck P Vidal, Iwan T Mitchell, Jean M Létang



Highlights

Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts

Franck P Vidal, Iwan T Mitchell, Jean M Létang

- Fully automatic creation of CAD models by image registration of X-ray projections.
- Automatic, accurate and stable geometric analysis of the material scanned by synchrotron microtomography.
- Simulation of the imaging chain, incl. beam hardening, impulse response of the detector, phase contrast, and photon noise.
- Generation of simulated CT images, including their defects leading to realistic artefacts.
- Fast X-ray simulations on GPU into an objective function to optimise.

Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts

Franck P Vidal^{a,*}, Iwan T Mitchell^a and Jean M Létang^b

^a*School of Computer Science and Electronic Engineering, Bangor University, Bangor, UK*

^b*Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1206, F-69373, LYON, France*

ARTICLE INFO

Keywords:

Computed tomography
X-rays
Numerical simulation
Optimisation
Computer aided analysis
High performance computing
Evolutionary computation

ABSTRACT

The presence of strong imaging artefacts in microtomographic X-ray data makes the CAD modelling process difficult to carry out. As an alternative to traditional image segmentation techniques, we propose to register the CAD models by deploying a realistic X-ray simulation on GPU in an optimisation framework. A user study was also conducted to compare the measurements made manually by a cohort of volunteers and those produced with our framework. Our implementation relies on open source software only. We numerically modelled the real experiment, taking into account geometrical properties as well as beam hardening, impulse response of the detector, phase contrast, and photon noise. Parameters of the overall model are then optimised so that X-ray projections of the registered CAD models match the projections from an actual experiment. It appeared that manual measurements can be variable and subject to bias whereas our framework produced more reliable results. The features seen in the real CT image, including artefacts, were accurately replicated in the CT image reconstructed from the simulated data after registration: (i) linear attenuation coefficients are comparable for all the materials, (ii) geometrical properties are accurately recovered, and (iii) simulated images reproduce observed experimental artefacts. We showed that the choice of objective function is crucial to produce high fidelity results. We also demonstrated how to automatically produce CAD models as an optimisation problem, producing a high cross-correlation between the experimental CT slice and the simulated CT slice. These results pave the way towards the use of fast realistic simulation for accurate CAD modelling in tomographic X-ray data.

1. Introduction


The use of X-ray computed tomography (CT) in precision engineering is becoming commonplace to assess the accuracy of a manufacturing process. The ISO 10360-11 standard has recently been issued to define metrological characteristics and methods for coordinate measuring systems making use of X-ray CT [14]. In particular, dimensional metrology makes use of the CT technique for accurate inspection and geometric analysis, as shown in a review paper [34]. X-ray CT is becoming an industrially relevant technology for additively manufactured materials which require the measurement of complex internal features [27]. In medical imaging, 3D image quality is usually degraded by metal implants or surgical devices in interventional guidance, and the identification of the known computer-aided design (CAD) component in the reconstructed tomographic volume bears a significant promise [25]. One of the main challenges lies in extracting the 3D surfaces from CT data, which are prone to artefacts such as beam-hardening, scatter radiation and partial volume effect [37]. The forward model in analytic X-ray CT reconstruction is usually based on the Radon transform, which amounts to assuming a monochromatic Beer-Lambert attenuation law and therefore translates into discrepancies in the reconstructed CT data. Sophisticated image processing

techniques are therefore necessary to accurately compute the surface [12], and many steps are usually required to extract CAD models from 3D data. When the input data is 3D point clouds of object surfaces, up to 11 phases are needed [18], including segmentation, region classification and surface fitting, and the CAD surfaces to be modelled are often engineering-based like planes and quadrics. A reference to those different stages to determine surfaces from voxel data is also found in the ISO 10360-11 standard [14] as “surfaces” are the mesurand for solving dimensional measurements made with CT-based coordinate measuring systems.

Depending on the final application, it may be necessary to have highly realistic simulation models, including surface models (eg statistical roughness distribution [2]) and a complex forward model. For example this is the case in model-based iterative CT reconstruction where phase contrast, noise and Compton scattering are taken in to account [11]. If traditional segmentation techniques are impractical to deploy, it may be possible to perform manual measurements when the sample is made of simple quadrics. However, manual segmentation by subject experts or produced by algorithms may exhibit various level of variability [17].

In a previous study, we investigated the sources of artefacts in synchrotron microtomography via virtual X-ray imaging [28]. These artefacts made it difficult and time consuming to extract CAD models by segmentation using popular intensity based methods such as threshold techniques and region growing, even though the sample was made of geometrically simple structures (parallel cylinders

*Corresponding author

 f.vidal@bangor.ac.uk (F.P. Vidal); jean.letang@insa-lyon.fr

(J.M. Létang)

ORCID(s): 0000-0002-2768-4524 (F.P. Vidal); 0000-0002-1456-6027 (I.T. Mitchell); 0000-0003-2583-782X (J.M. Létang)

in a box-shaped matrix). It was challenging to accurately locate the boundary between two different materials of the sample. That was a time when computers were still relatively slow. We aim to address these two issues by the registration of highly realistic simulated X-ray images from CAD models onto experimental CT data. Our assumption is that recent advances in computational power makes it possible to deploy high performance X-ray simulation on graphics processing unit (GPU) in the objective function of a global optimisation algorithm.

To simulate realistic X-ray projections from CAD models, Monte Carlo simulation methods are known to produce the most accurate X-ray images [1]. In such methods, X-ray photons are emitted from the source toward the detector. Each photon is tracked independently. At each step of the simulation, interactions of photons as they traverse matter could occur depending on stochastic physics models, i.e. laws of chance that are interaction specific (Photoelectric effect, Compton scattering, and other less dominant interactions). This high fidelity model is, however, very slow. A week may be needed to produce an image without too much noise [33]. Parallelism and variance reduction techniques may be deployed to speed up Monte Carlo simulations, however real-time performance is still not achievable, particularly on desktop computers. Analytic methods based on ray-tracing provide a fast alternative as they rely on a simple forward model [7]. They are particularly suited for GPU implementations [30] but sometimes lack realism to reproduce observed CT artefacts.

We have developed and validated an open source library (gVirtualXRay) to perform the attenuation calculations on GPU using OpenGL [29, 32, 33]. gVirtualXRay is written in C++ and GLSL. Wrappers to other popular languages such as Python, R, Java, are provided. We used Python in this study. Although this choice reduces the speed of the framework, it allows us to prototype it as a proof-of-concept. The beam polychromaticity, the detector response, the phase fringes and the noise have been added to the forward model to simulate realistic projections. Both the geometrical and physical properties of the sample are considered as unknown parameters to tune, which is performed as an optimisation problem. Our code is open source and we produced a Jupyter notebook as a step-by-step tutorial. We show that realistic X-ray simulation is crucial to get an accurate estimation of the CAD model and that manual segmentation of CT data may be prone to bias. We also demonstrate that the choice of objective function is important to produce high fidelity results. The data is analysed using interactive information visualisation (InfoVis). Our visualisations are also available online.

The paper is organised as follows. Section 2 provides some context about artefacts in computed tomography. Section 3 describes the realistic forward model and our registration pipeline. We quantitatively and visually assess our results in Section 4, and we also compare the performance of the framework on nodes with GPU on a supercomputer and a desktop computer. Parameters

extracted manually and automatically are also compared. The paper ends with a discussion and conclusion in Sections 5 and 6 respectively.

2. Context: Artefacts in CT

In X-ray tomographic imaging, the reconstructed volume is ideally a 3D map of linear attenuation coefficients, which corresponds to a probability of interaction of the X-ray photon per unit length. This assumes that the spatial impulse response of the detector is a Dirac, the X-ray source is monochromatic and no secondary radiations resulting from X-ray interactions are detected. However, artefacts may occur in reconstructed volumes [28, 37]. Indeed, the increasing use of laboratory microfocus X-ray sources makes it possible to obtain images with phase contrast without having to apply for beamtime at a synchrotron radiation facility centre. These can aid boundary location but they distort the image contrast locally [6, 37].




Figure 1: CT slice reconstructed from Figure 2b. Dark streak artefacts are due to beam hardening and the response of the detector, black and white fringes are due to phase contrast [28]. Other artefacts are seen around the tungsten cores.

Figure 1 shows a typical example of a reconstructed slice from a fibre composite where strong CT artefacts are present. Two kinds of artefacts are clearly visible:

1. Black and white fringes at the silicon carbide edges: these phase artefacts are similar to some image enhancement techniques like unsharp masking and can be modelled by an additive term involving the Laplacian of the phase shift [6, 20, 22];
2. Dark streaks with bright borders, located along alignments of high density tungsten cores: these artefacts are due to beam hardening and the impulse response of the detector [28]. As voxel values in these streaks are darker than what they should be, it means that the linear attenuation coefficients are underestimated in the artefact areas. Table 5 compares

the attenuation coefficients in this slice with the corresponding theoretical values [4, 5]. The tungsten coefficients are underestimated by a factor of 2 (162. vs 342. cm^{-1}).

These artefacts make it impractical to extract CAD models by segmentation using popular intensity based methods such as threshold techniques and region growing. It is also extremely challenging to accurately locate the boundary between a fibre and its core, and between the fibres and the matrix. A quantitative analysis is also difficult as the attenuation coefficients are not consistent with the corresponding theoretical values.

An alternative approach to segment these difficult images would be the deployment of deep neural networks (DNNs) that, again, advances in computational power enabled in recent years. DNN-based methods are robust in segmenting or denoising against artefacts but their downside is the need for training data [19, 26]. Such data is often labelled manually, which is labour intensive. Our framework eliminates this need by the use of simulations.

Alternatively, fast analytical simulations can also be used to generate training data in a controlled manner. For example,

Haiderbhai *et al* [8] proposed a method based on a generative adversarial network (GAN), a machine learning approach that can be used to create synthetic images. Images simulated using gVirtualXRay, the same X-ray generator as in our registration framework, are used to create a large database to train the GAN.

3. Materials & Methods

The research presented here is available for download [32]:

- (i) some of the sinograms acquired at ESRF are provided,
- (ii) the code of the X-ray simulation package is open source [29],
- (iii) our registration code is provided on GitHub [31], and as supplementary material with this paper,
- (iv) the code of our visualisations is provided in the same locations and the corresponding HTML pages can be consulted on gVirtualXRay's website [29], and
- (v) a tutorial is provided as a Jupyter Notebook as supplementary material. The steps in the sub-sections below refer to the step numbers in the notebook.

We first present the issues to be addressed, namely (i) how to enhance the Beer-Lambert law to simulate realistic X-ray projections including phase contrast, Poisson noise, impulse response and beam hardening, and (ii) how to select the optimisation algorithm and the objective function. The toy test case and the CAD optimisation framework are then presented.

3.1. Forward modelling

In the literature, a projection is often modelled using the polychromatic version of the Beer-Lambert law:

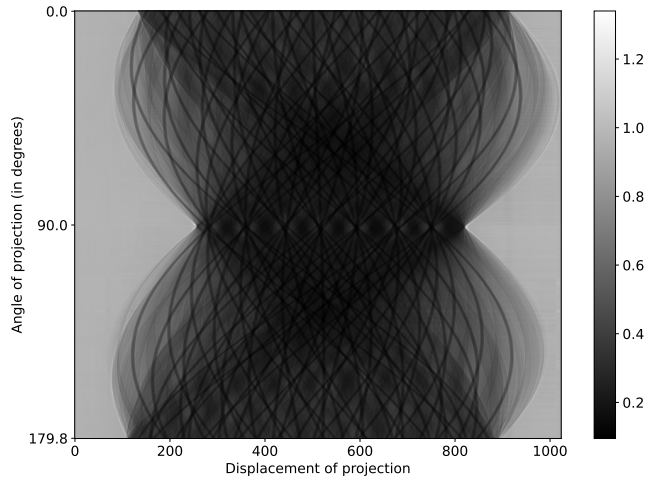
$$I(x, y) = \sum_i R_i N_i \exp\left(-\sum_j \mu_j(E_i) d_j(x, y)\right) \quad (1)$$

with $I(x, y)$ the value of the raw X-ray projection at pixel location (x, y) , and with the sample and with the X-ray beam turned on; i the i -th energy channel in the beam spectrum; E_i the energy in eV; R_i and N_i the detector response and the number of photons at that energy respectively; j the j -th material being scanned, $\mu_j(E_i)$ its linear attenuation coefficient at energy E_i , and $d_j(x, y)$ path length in cm^{-1} of the ray crossing the j -th material from the X-ray source to pixel (x, y) .


Projections are then corrected to account for variations in beam homogeneity and in the pixel-to-pixel sensitivity of the detector [15]. This is the projection with flat-field correction (**Proj**):

$$\text{Proj} = \frac{I - D}{F - D} \quad (2)$$

where F (full fields) and D (dark fields) are projection images without sample and acquired with and without the X-ray beam turned on respectively. Note that with an ideal detector ($R_i = E_i$), pixels of D are null, and pixels of F are equal to $\sum_i E_i N_i$.



(a) Projections for a given detector row after flat-field correction, i.e. **Proj** in Equation 2.



(b) Sinogram for a given detector row, i.e. **Sino** in Equation 4.

Figure 2: Experimental data acquired at ESRF.

Figure 2a corresponds to **Proj**. We can see that when the primary spectrum is not monochromatic the measurement in Equation 1 is the sum of several attenuation laws. We could however compute the effective monochromatic attenuation that would give the same measurement:

$$I(x, y) = I_0(x, y) \exp\left(-\sum_j \mu_j(E_{\text{eff}}) d_j(x, y)\right) \quad (3)$$

with $I_0(x, y) = \sum_i R_i N_i$, and where E_{eff} corresponds to the monochromatic energy that would give the same attenuation than the one measured. We are now able to linearise the transmission tomography data, namely **Proj** in Equation 2, and we get the sinogram:

$$\mathbf{Sino} = -\ln(\mathbf{Proj}) \quad (4)$$

Now we got a sinogram (Figure 2b), we can reconstruct the CT slice. As we used a synchrotron, where the Source-to-Object Distance (SOD) is equal to 140m, we can assume a parallel source has been used. In this case the filtered back-projection (FBP) algorithm can be used instead of the FDK (Feldkamp, Daci, Kress) method.

Figure 1 shows the CT slice reconstructed from Figure 2b. When the simulations are computed without phase contrast, gVirtualXRy calculates Eq. 1 for every angle of the tomography acquisition (Line 9 of Step 44).

When phase contrast is taken into account, gVirtualXRy is used to generate the path length buffers (d_j) for each material j of the sample (Line 24 of Step 49). The contribution to phase contrast for each material is accumulated in a new image (Line 30 to 35 of Step 49):

$$\mathbf{Phase} = \sum_j k_j d_j * \mathbf{Laplacian}_{\sigma(j)} \quad (5)$$

where k_j is a weighting factor corresponding to the contribution of the j -th material to phase contrast, $\sigma(j)$ a parameter controlling the width of the Laplacian operator, and $*$ denotes the convolution of d_j (the path-length image of the j -th material) with a Laplacian kernel implemented as follows:

$$\mathbf{Laplacian}_{\sigma}(x, y) = \left(\frac{x^2 + y^2}{\sigma^4} - \frac{1}{\sigma^2}\right) \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \quad (6)$$

The final phase contrast model includes two real numbers for each material j , k_j and $\sigma(j)$.

A pixel-wise subtraction ‘adds’ the phase contrast to the raw projections (Line 55 of Step 49):

$$I_p = I - \mathbf{Phase} \quad (7)$$

Once the projections are generated, the response of the camera (i.e. the line spread-function (LSF) in our case since the sample is translation invariant) is applied row by row by convolution (Lines 58 to 60 of Step 49):

$$I_{\text{LSF}} = I_p * \mathbf{LSF} \quad (8)$$

The LSF model from our previous investigation [28] was refined and is parameterised as follows

$$\mathbf{LSF} = \left(\frac{a}{c} \left(1 + \frac{x^2}{c^2}\right)^{-1} + \frac{2b}{\sqrt{\pi}d} \exp\left(-\frac{x^2}{d^2}\right)\right) \frac{1}{\pi a + 2b} \quad (9)$$

and the flat-field correction can be performed on I_{LSF} to produce **Proj**_{LSF} using Eq. 2 (Line 63 of Step 49). Poisson noise is added to the projections (Lines 67 to 72 of Step 49):

$$\mathbf{NoiseMap} = \text{Poisson}[(\mathbf{Proj}_{\text{LSF}} + \text{bias}) \times \text{gain}] \quad (10)$$

and

$$\mathbf{Proj}_n = \mathbf{Proj}_{\text{LSF}} + \text{intensity} \times \mathbf{NoiseMap} \quad (11)$$

Proj_n is eventually linearised using Eq. 4 to produce a sinogram suitable for CT reconstruction algorithms (Line 79 of Step 49).

3.2. Optimisation algorithm

Image registration can be defined as the process of overlaying two (or more) images [38]. These images may be from different equipment, time, etc. In medicine they may even be from different modalities or patients. In image registration, a *moving object* is transformed so that its image matches a *target image* [21]. The parameters of the transformation are controlled and iteratively tuned by an optimisation algorithm (see Figure 3). The optimisation algorithm will either minimise a discrepancy metrics or maximise a similarity metrics. The choice of metrics that is optimised will impact the quality of the registration [23, 24]. However, there is no image comparison method that fits all problems. For complex transformations or complex image comparison metrics, sophisticated optimisation algorithms must be used [38] as the choice of optimisation algorithm, stopping criteria, and objective function can significantly affect the result of the registration [36].

The two images that are being registered must be compared in order to minimise their discrepancies or maximise their similarities. This comparison is performed within the objective function of the optimisation algorithm. It is very likely that the derivatives of the various objective functions to optimise are unavailable, and cannot be obtained. It is, indeed, the case here. As a consequence, a derivative-free optimisation algorithm is required. As we make no assumption on the shape of the search space, we also need an optimisation algorithm that is robust and can escape local minima. A stochastic and population-based algorithm based on artificial evolution is favoured for these reasons. Artificial evolution is a class of optimisation algorithms inspired by nature, here Darwin’s principle of natural selection. The solutions of the optimisation problem are represented by *individuals*. The encoding of individuals is problem-dependent. In our case an individual is a sequence




Figure 3: Image registration as an optimisation algorithm.

of real numbers. Individuals are grouped in a *population*.

The fittest individuals of the population have a higher probability of breeding. This is determined by the values of the objective function.

During an optimisation loop, called a *generation*, new individuals are created (offspring) from the current population (*parents*) by repetitively applying genetic operators such as *selection*, *recombination* and *mutation*.

The selection makes use of the objective values to randomly select parents. The genes of parents are ‘mixed’ to produce a child, this is the recombination (also called cross-over). There are various ways to implement the recombination. The child’s gene are randomly altered during the mutation. Again, there are various ways to implement the mutation. Successive generations are created until a stopping criterion is met.

We use one of today’s most popular global optimisation algorithm: covariance matrix adaptation evolution strategy (CMA-ES) [10]. It is an EA designed for difficult non-linear non-convex optimisation problems in continuous domain and it is considered as state-of-the-art in evolutionary computation [9]. CMA-ES does not require a tedious parameter tuning for its application as finding suitable internal parameters is part of the algorithm design, which also makes it attractive. Only an initial solution and an initial standard deviation must be set by the user. The default population size is relatively small for fast convergence. However, if the algorithm becomes stuck in a local minimum, restart with with increasing population size improves the global search performance [3] (note that we purposely did not use this possibility to assess the

framework). Stopping criteria may be tweaked by the user.

3.3. Objective Functions

Objective functions used in image registration aim to quantify how dissimilar or how similar two images are. There are many ways to compare images (Step 50). Mean absolute error (MAE) and root mean squared error (RMSE) are two of the most common metrics to quantify differences (see Equations 12 and 13 respectively). There is no clear consensus of which one is the best. They both provide an average model prediction error. MAE is considered more robust to outliers as the squared values in RMSE give a higher weight to large errors. It does not mean that MAE is better. For example, RMSE should be preferred when large errors must be prevented.

$$\text{MAE}(\mathbf{m}, \hat{\mathbf{m}}) = \frac{1}{w h} \sum_q^h \sum_p^W |m_{p,q} - \hat{m}_{p,q}| \quad (12)$$

$$\text{RMSE}(\mathbf{m}, \hat{\mathbf{m}}) = \sqrt{\frac{1}{w h} \sum_q^h \sum_p^W (m_{p,q} - \hat{m}_{p,q})^2} \quad (13)$$

where \mathbf{m} is the reference image (i.e. the experimental one), $\hat{\mathbf{m}}$ is its approximation (i.e. the simulated one), w is the number of pixels in a projection (here 1024 since a single row is considered), and h is the number of angles in the sinogram (here 900). When $\text{MAE}(\mathbf{m}, \hat{\mathbf{m}})$ or $\text{RMSE}(\mathbf{m}, \hat{\mathbf{m}})$ are close to 0, the approximated image $\hat{\mathbf{m}}$ is similar to the reference image \mathbf{m} .

In computer vision, the Zero Mean Normalized Cross-Correlation (ZNCC) is a popular way to compare images as it embeds the zero-mean, unit-variance normalisation (Equation 17), which tackles changes in brightness and contrast between the two images.

$$\text{ZNCC}(\mathbf{m}, \hat{\mathbf{m}}) = \frac{1}{w h} \sum_q^h \sum_p^W \frac{m_{p,q} - \bar{m}}{\sigma_m} \frac{\hat{m}_{p,q} - \bar{\hat{m}}}{\sigma_{\hat{m}}} \quad (14)$$

The ZNCC provides values between -1 and 1 and is often presented as a percentage. The cross-correlation between \mathbf{m} and $\hat{\mathbf{m}}$ is 1 when the two images are 100% correlated; it is 0% when they are non-correlated; and -100% when they are anticorrelated (e.g. when one image is the negative of the other image). As our optimisation algorithm minimises the objective function, we use:

$$\frac{1 - \text{ZNCC}(\mathbf{m}, \hat{\mathbf{m}})}{2} \quad (15)$$

The structural similarity index measure (SSIM) is also popular in computer vision as it considers structural information. It is typically used to assess the perceived quality of digital images in television and cinema, as well as other kinds of digital images and videos. It is the structural similarity index measure (SSIM) [35]. It considers the image degradation as ‘perceived change in structural information’, e.g. due to the use of lossy compression algorithms to encode images. Traditional techniques such as MAE and RMSE consider every pixel independently from its neighbours. In SSIM, however, pixels have strong inter-dependencies, especially when they are spatially close. When $\text{SSIM}(\mathbf{m}, \hat{\mathbf{m}})$ is close to 1, $\hat{\mathbf{m}}$ is similar to \mathbf{m} . As our optimisation algorithm minimises the objective function, we use the structural dissimilarity (DSSIM):

$$\text{DSSIM}(\mathbf{m}, \hat{\mathbf{m}}) = \frac{1 - \text{SSIM}(\mathbf{m}, \hat{\mathbf{m}})}{2} \quad (16)$$

To register simulated data onto the real experiment’s data, we will evaluate these objective functions computed either on the projections after flat-field correction (**Proj** in Equation 2) or the sinogram (**Sino** in Equation 4), and either with or without zero-mean, unit variance normalisation (also known as standardisation or Z-score Normalisation in machine learning):

$$\mathbf{m}_o = \frac{\mathbf{m} - \bar{m}}{\sigma_m} \quad (17)$$

where \mathbf{m}_o is the image after normalisation of Image \mathbf{m} , \bar{m} is the average pixel value of Image \mathbf{m} , and σ_m its standard deviation. After normalisation, the average pixel value is null and the standard deviation of pixel values is equal to one.

Whilst this strategy (using four different error metrics) increases the runtime, it allows us to assess the effect of the log function used to convert the projections into a sinogram, and if normalisation is needed or not. One of the

difficulty is the relative wide range of attenuation properties amongst the material used in the sample. In total, 14 different objective functions will be evaluated.

3.4. Toy test case

In this paper, we use a typical reconstructed slice of an Ti/SiC fibre composite as an example (see Figure 4 for a photograph of the sample and an optical microscopy of the slice). The matrix is a mixture composed of titanium




Figure 4: Scanned object: SiC fibres with a tungsten core in a Ti90/Al6V4 matrix. From the microscopic image, the tungsten core diameter is 15µm and the fibre diameter 105µm.

(90%), aluminium (6%) and vanadium (4%). Fibres are silicon carbide and their cores tungsten. The data has been acquired at ESRF and a multilayer monochromator used to make the beam spectrum almost monochromatic. We selected an energy of 33 keV. As the source is so far away from the sample and the detector (see experimental set-up in Figure 5), a parallel geometry can be considered during the tomography reconstruction using the standard filtered back-projection (FBP).

The Ti90/Al6/V4 matrix is a parallelepiped. It has a centre, a size and orientation. These are unknown and will be tuned by the registration pipeline. The tungsten cores are cylinders. Each core has a centre and radius. The centres can easily be retrieved from the image (see below). The radius is, however, unknown. The SiC fibres are cylinders.

Each fibre has a centre and radius. This radius is also unknown. SiC fibres and tungsten cores share the same centres. We subtract the fibres from the matrix using the difference operator in constructive solid geometry (CSG). We subtract the cores from the fibres using the difference operator in CSG. The cores have a smaller radius than the fibres.




Figure 5: Experimental set up at ESRF’s ID19 - Microtomography beamline.

Table 1
Summary of the parameters that need to be optimised.

	Parameter	Notation
1	Position of the matrix along the primary axis (in μm)	u
2	Position of the matrix along the secondary axis (in μm)	v
3	Size of the matrix along the primary axis (in μm)	w
4	Size of the matrix along the secondary axis (in μm)	h
5	Rotation angle of the matrix (in degrees)	θ
6	Radius of the cores (in μm)	r_W
7	Radius of the fibres (in μm)	r_{SiC}
8	Percentage of 33 keV photons in the beam spectrum	$k_{33\text{keV}}$
9	Percentage of 66 keV photons in the beam spectrum	$k_{66\text{keV}}$
10	Percentage of 99 keV photons in the beam spectrum	$k_{99\text{keV}}$
11	Bias controlling the Poisson noise	bias
12	Gain controlling the Poisson noise	gain
13	Intensity of the Poisson noise	intensity
14	Intensity of the phase contrast for the tungsten core	k_W
15	Spread of the phase contrast for the tungsten core	σ_W
16	Intensity of the phase contrast for the SiC fibres	k_{SiC}
17	Spread of the phase contrast for the SiC fibres	σ_{SiC}
18	Intensity of the phase contrast for the Ti90Al6V4 matrix	$k_{\text{Ti90Al6V4}}$
19	Spread of the phase contrast for the Ti90Al6V4 matrix	$\sigma_{\text{Ti90Al6V4}}$
20-23	Parameters of the LSF (see Equation 9)	a, b, c, d

3.5. Optimisation Framework

In our context, the target are the projections (either **Proj** or **Sino**) provided by the experiment at ESRF. The moving image is created by simulation using the CAD models and gVirtualXRy. The simulation model is initialised with all the known parameters. The unknown parameters controlling the simulation model (see Table 1) are repetitively tuned by a global optimisation algorithm until a stopping criterion is met. The optimisation algorithm will minimise (or maximise) a numerical value, the *objective function*. The comparison between the target and moving images measures how different (or similar) the two images are. It is performed within the objective function.

To improve the performance of our registration pipeline, we use a divide and conquer strategy where we split this complex optimisation problem into smaller problems that are easier to solve. In total there are 23 parameters to estimate (see Table 1). Our registration pipeline is divided in 8 successive steps:

1. Registration of a cube to approximate the Ti90Al6V4 matrix: 5 real numbers: u and v for the centre of the matrix, w and h for its width and height, and θ a rotation angle (Step 52).
2. Registration of the radii of the tungsten cores and SiC fibres: 2 real numbers: r_{SiC} the radius of the SiC fibres, and a scaling factor so that $r_W = \text{scaling} \times r_{\text{SiC}}$ with $\text{scaling} < 1$ (Step 66).
3. Recentring the tungsten cores and SiC fibres (Steps 72 & 73).
4. Refining the registration of the radii of the tungsten cores and SiC fibres: 2 real numbers: the radius of the SiC fibres, and the scaling factor (Step 75).
5. Optimisation of the weights of the different energies in the beam spectrum: 3 real numbers $k_{33\text{keV}}$, $k_{66\text{keV}}$,

and $k_{99\text{keV}}$ so that $k_{33\text{keV}} + k_{66\text{keV}} + k_{99\text{keV}} = 1$ (Step 82).

6. Adding the phase contrast: 8 real numbers: 2 per material (k_W , k_{SiC} & $k_{\text{Ti90Al6V4}}$, and σ_W , σ_{SiC} & $\sigma_{\text{Ti90Al6V4}}$), and one for the radii of the tungsten cores r_W and SiC fibres r_{SiC} (Step 92).
7. Optimising the LSF parameters and reoptimising the phase contrast: 7 real numbers: 4 for the LSF (a, b, c, d), and 1 for the phase contrast of each material (k_W , k_{SiC} & $k_{\text{Ti90Al6V4}}$) (Step 100).
8. Optimisation of the noise level: 3 real numbers (bias, gain and intensity) (Step 111).

The registration framework will vary these 23 parameters using an optimisation algorithm to improve the value of an objective function. This objective function is based on the comparison of the experimental projections acquired at ESRF with the projections simulated using the parameters being optimised. To compute these simulated projection, we need to:

1. Set the fibre and core geometries and material properties (Step 45)
2. Set the matrix geometry and material properties (Step 46)
3. Simulate the raw projections for each angle:
 - Without phase contrast (Line 9 of Step 49), or
 - With phase contrast (Lines 14-55 of Step 49)
4. Apply the LSF (Lines 58-60 of Step 49)
5. Apply the flat-field correction (Line 63 of Step 49)
6. Add Poisson noise (Lines 69-76 of Step 49)
7. Apply the minus log normalisation to compute the sinogram (Step 8 and Lines 79 of Step 49)

Table 2

Main simulation parameters. These parameters are fixed, i.e. they are constant during the registration process.

Parameter	Value
Source shape:	cone beam
source-object distance (SOD):	145m
source-detector distance (SDD):	145m + 80mm
Pixel spacing:	1.9 μ m
Pixel resolution:	1024 pixels per line
Number of projections:	900
Angular span:	180°
Beam spectrum:	primary energy: 33keV 1 st harmonic: 66keV 2 nd harmonic: 99keV

3.6. Initialisation of the Simulation Model with Known Parameters

During the initialisation phase, the input data, here the X-ray projections, is loaded (Step 8). The corresponding sinogram is computed (Steps 12 & 13) and it is used to create the reference CT slice (Step 16). If the objective function used in the registration makes use of the zero-mean unit-variance normalisation, it can be applied on the reference images during the initialisation step so that it is performed only once (Steps 19 & 20). Note that the reference CT slice is not actually used during the registration. It is used afterwards to assess the accuracy of the registration.

For the X-ray simulation, we first create an OpenGL context (Step 21). Geometric parameters corresponding to the experiment conducted at the ESRF are set (Steps 23 & 24). They are summarised in Figure 5 and Table 2. A beam spectrum of 97% of primary energy at 33keV, 2% of 1st harmonic at 66keV and 1% 2nd harmonic at 99keV is used as an initial guess (Step 25). The respective weights will be optimised at some point during the registration, once geometric properties of the matrix and fibres are known (Step 82).

The material properties (chemical composition and density) of the matrix, tungsten cores and silicon carbide fibres are then set (Step 27).

The convolution kernel corresponding to the line spread function measured at ESRF is then loaded as an initial guess (Step 28) [28]. Its parameters will be tuned at some point during the registration (Step 100).

The final step of the initialisation is the localisation of the centre of all the fibres. As the fibres and the cores correspond to circles in the CT images, the obvious technique to try is the Hough Circle Transform (HCT) [13]. It is a feature extraction technique used in image analysis that can output a list of circles (centres and radii).

However, as Figure 6 shows, 13 fibres were missed and many centres were misplaced (Steps 32-34). Controlling the meta-parameters of the algorithm can be difficult, particularly in a fully-automatic registration framework, and we eventually discarded this method. As the tungsten cores have a much higher intensity than any other region of




Figure 6: Circle detection using the Hough Circle Transform on the reference CT slice.

the image, binary thresholding can be used as an alternative. We rely on the popular Otsu's method that creates a histogram and uses a heuristic to determine a threshold value (Step 35). It highlights the pixels of the




Figure 7: Overlay of the binary image showing the tungsten cores in green and the reference CT slice.

tungsten cores (see Figure 7). Each distinct tungsten core is assigned a unique label, i.e. a unique pixel intensity (see Figure 8) (Step 41). For each labelled tungsten core, we extract the centroid (Steps 43 & 44). We now have a list of the centres of all the fibres that can be used as the parameter of a function to create the cylinders




Figure 8: Overlay of the labelled image showing the tungsten cores in different colours and the reference CT slice.

corresponding to the cores and the fibres (Step 44).

3.7. Validation Methodology

To validate our approach we must demonstrate (i) that the choice of objective function is key to success, (ii) results are consistent for a given objective function, (iii) that results can be obtained in a reasonable amount of time with a typical desktop computer, and (iv) the benefit(s) of our framework over manual measurements.

Due to the stochastic nature of the optimisation algorithm, each experiment is repeated 25 times. This number is arbitrary. The standard error of the arithmetic mean is σ/\sqrt{n} . The variability between the samples was sufficiently low for a value of $n = 25$ to be acceptable. This methodology is commonly used in evolutionary computing to provide statistically meaningful results. As we compare 14 objective functions (**Proj** vs **Sino**, with or without zero-mean unit-variance normalisation, using MAE, RMSE, ZNCC or DSSIM), that is 350 registrations in total. These test registrations will be performed on GPU nodes of a supercomputer such as Supercomputing Wales (SCW, the supercomputer accessible by Welsh universities, <https://www.supercomputing.wales/>). For the best performance, we use a single GPU node per registration. Otherwise hardware resources might be shared with other users. A statistical analysis and interactive visualisations will be used to explore the large amount of data that will be generated. The aims are (i) to identify which objective function(s) is significantly better than others and (ii) to assess, again, variability of the results.

Conducting 350 registrations on a desktop computer may not be feasible. However, we will perform 25 registrations with the objective functions that are identified as the ‘best’ ones. This way we will compare the relative computational

performance between a GPU node on the supercomputer and a regular desktop computer, which is just the office PC of one of the co-authors.

A user study will be conducted. Suitable volunteers will be recruited online. To contribute to the user study, volunteers must be familiar with ImageJ or another image analysis software that is equivalent. The reference image is provided in full resolution and with single-precision floating point accuracy. We prepared an online form with instructions and a questionnaire (see supplementary materials). The volunteers are asked:

1. To make sure they record the time when they started the experiment.
2. To find the structure that corresponds to the Ti90Al6V4 matrix in the image and measure or estimate its:
 - (a) Width,
 - (b) Height,
 - (c) Centre of this structure,
 - (d) Rotation angle.

In addition they are asked to evaluate the difficulty of these tasks using a Likert scale.

3. Find in the image the structures that corresponds to three pre-selected fibres. In the form, they must report for each of the three fibres:
 - (a) Its centre,
 - (b) The diameter of its inner white circle,
 - (c) The diameter of its outer dark circle.

In addition they are asked to evaluate the difficulty of these tasks using a Likert scale.

4. Record the time when you stopped the experiment.

The aim of this study is to manually extract geometrical parameters and conduct a statistical analysis to assess the variability of the responses. Ten respondents participated to the study. In practice, if no registration was available the median value for each geometrical parameter could be used in the simulation to compare the corresponding slice with the ground truth. Median values are used instead of average values to avoid using parameters skewed by outliers.

3.8. Pseudo-3D Extension and Robustness to Changes of Acquisition Parameters

We conducted another experiment with different acquisition parameters: incident beam 80keV, distance sample-detector 270mm, distance source-detector 92m, pixel spacing $1.22\mu\text{m}$, angular span 360° over 3500 projections, detector size 1424×1024 pixels. The sample is larger than the detector. To accommodate for this, the centre of rotation of the sample is off compared to the central column of the detector. To avoid sample truncation in the projection domain for the registration optimisation, it is necessary to reformat the projections using redundant information so that they span over 180° , and the centre of rotation and the central column of the detector becomes aligned. Figure 9 shows that the new dataset is corrupted by similar streak and phase contrast artefacts.




Figure 9: ROI of the first slice reconstructed from experimental data with different acquisition parameters (left) and of the last slice (right). Both ROIs have been extracted at the same pixel locations. The differences are due to the fact that the fibres are slightly tilted.

As 2D projections are now available, it makes it possible to extend and alter our proposed method from 2D to pseudo-3D (see Algorithm 1). Indeed, running the pipeline presented in Section 3.5 for every slice is not necessary and it would be far to computationally intensive. Only a single full registration is needed to estimate parameters that will be constant amongst every slice, such as the radii, spectrum, phase contrast, LSF and noise. Note that it is not necessary to optimise the LSF and noise to generate CAD models. We propose in this section to register the radii only once, at the same time as the phase contrast. There are two reasons for that: i) To limit the number of optimisation steps as much more data will be processed in the 3D case, and ii) An estimate is available from Figure 4b. For subsequent slices, only the geometrical parameters of the matrix must be optimised. Parameters of the previous slice can be used as an initial guess. The centre of fibres must be estimated for every slice in case they are slightly tilted. Comparing the core centres in the ROI of the first slice with those of the last slice shows that the fibres are indeed tilted (see Figure 9). Note that the rotation parameters do not need to be modelled and optimised as they can be estimated using traditional image analysis techniques in the projection domain, such as the Hough transform. Our measurements on the projection data show that fibres are parallel and tilted by 0.8° .

4. Results

In this section, we aim to demonstrate the benefit of our framework by performing a series of validation studies. To identify which objective functions provide the best results, we will compare the different objective functions in terms of accuracy, stability, and relative runtime. We will also compare corresponding runtimes on a GPU node on the supercomputer and a regular desktop computer. We will compare the parameters provided by the registration and parameters extracted manually by a cohort of volunteers. A registration corresponds to optimising the 23 unknowns of the model, which was broken down into 8 successive

Algorithm 1 Pseudo-3D extension

```

1: Initialise known parameters
2: Initialise X-ray simulation library
3: for each slice (i) do
4:   if i is 0 then
5:     Initialise matrix with initial guess    ▷ can be
       random or use rough manual measurements
6:   else
7:     Use parameters of slice i-1 as initial guess
8:   Register the matrix
9:   Extract fibre centres
10:  Simulate the corresponding sinogram
11:  Reconstruct the corresponding CT slice
12:  Recentre the fibres
13:  if i is 0 then
14:    Optimise the beam spectrum
15:    Optimise the phase contrast and radii
16:    Optimise the LSF and phase contrast ▷ optional
17:    Optimise the Poisson noise        ▷ optional
18:  Save corresponding CAD models
19:  Simulate the corresponding sinogram  ▷ optional
20:  Reconstruct the corresponding CT slice ▷ optional

```

steps, amongst which 7 are optimisation problems solved with CMA-ES. Note that CMA-ES has been used as a black-box optimisation algorithm, its hyperparameters have not been tuned to achieve the best performance. Only the initial guess and standard deviations have been set. The termination tolerance in x-changes (tol_x) and on the function value (tol_{fun}) have been increased from 10^{-11} to 10^{-2} or 10^{-4} (except for the noise where they are 10^{-8}) to speed up the computations. Again, we made no particular effort to tweak them. The value of other parameters, including population size and other stopping criteria, are set to their default values.

On average each registration takes 1 hour and 35 minutes, which is achievable on a desktop computer or laptop.

However, the cumulative runtime for 350 registrations is 23 days (*i.e.* $350 \times 1:35$), which is only achievable thanks to the use of GPUs on a supercomputer.

4.1. Quantitative Evaluation of the Objective Functions

Figure 10 presents the distributions of ZNCC between the experimental CT slice and the final simulated CT slice for the 25 runs of each objective function using boxplots. Our assumption is that the objective function providing the highest correlation and smallest spread is the best method, *i.e.* RMSE on **Sino** with normalisation. We computed the p -values between the distribution of ZNCC for the 25 runs of this objective function and the distributions for the other objective functions. The aim is to assess if RMSE on **Sino** with normalisation is significantly better ($p < .001$) than all the other objective functions (it is the case) or if there are other objective functions that perform as well (this is not the case). As a consequence we can use RMSE on **Sino**




Figure 10: Boxplot of the ZNCC between the experimental CT slice and the final simulated CT slice for the 25 runs of each objective function.

with normalisation as our baseline as it is significantly better than all the other objective functions. However, there are a few outliers that need to be investigated. We can identify these outliers using a parallel coordinate plot (see Figure 11), a technique popularised by Alfred Inselberg [16]. Such plots enable us to visualise and analyse high-dimensional datasets.

When highlighting the 25 runs of RMSE on **Sino** with normalisation in Figure 12, we can easily identify the outliers, the lines with the lowest ZNCC (see http://gvirtualxray.sourceforge.net/parallel_coordinates-ZNCC.php for an interactive version). It is clear that the registration of the radii of the tungsten cores and SiC fibres did not perform as well as for the other runs. Considering that the runtime is similar for all the runs, it is likely that the optimisation algorithm fell into a strong local minimum. Techniques used in evolutionary computing, such as *restart*, can be deployed to address this problem. Figure 11 also shows that the runtime of the ZNCC-based optimisations is significantly lower than for others. In case speed is favoured over accuracy, it is a good contender, although care must be given to the data used in the fitness function. Comparing the highlighted lines in Figures 13 and 14 clearly shows that ZNCC on **Sino** is a valid choice ($92.20\% \pm 1.02\%$) whereas ZNCC on **Proj** ($90.68\% \pm 1.49\%$) must be discarded.

To further assess if the RMSE on **Sino** with normalisation is the best objective function, we consider the distribution of ZNCC values of the second best objective function (namely DSSIM on **Proj** with normalisation). We aim to assess their relative performance in light of the runtime. For this purpose, a scatter plot is generated for these two objective functions (see Figure 15). The runtime is a value that should be as low as possible, and the ZNCC as high as possible: Points in the bottom right corner of Figure 15 highlight the objective function that should therefore be favoured. From Figure 15, it is clear that the spread in term

of ZNCC is smaller for RMSE on **Sino** with normalisation, than DSSIM on **Proj** with normalisation. It is also clear that ZNCC on **Sino** is a good solution if speed is preferred over accuracy.

4.2. Runtime Comparison between a GPU Node on a Supercomputer and a Desktop PC

We run our framework on two different systems:

Supercomputing Wales, and a desktop computer. Their specifications are given in Table 3. As we can repetitively perform the registration on a supercomputer, we make no assumption on the best choice of image metrics, data and normalisation: Many stochastic registrations will be performed to gather statistically meaningful results.

Hardware and software resources are shared on a supercomputer. Depending on availability, up to 5 GPU nodes could be used at the same time. Each node is equipped with a NVIDIA Tesla V100 GPU. Each node performs a single registration. In addition, we compare the run times with a desktop computer for three objective functions: RMSE on **Sino** with normalisation, ZNCC on **Sino**, and DSSIM on **Proj** with normalisation.

Table 4 summarises the run times for each optimisation for each objective function. For each corresponding image metrics, it always took longer on the sinograms than on projections. However, minimising the same metrics on sinogram does not necessarily leads to better results than on projections (see Figure 10). Some of the optimisations were rather quick, e.g. three to five minutes to estimate the parameters of the matrix. The first optimisation of the radii always took more time than the second one (after recentring). This indicates that initialising the meta-parameters of CMA-ES with a good initial guess has the potential to decrease the overall run time by a large factor. However, for convenience we let the computer do all the work and minimise the amount of human inputs, even if it increases the run time. Phase contrast is by far the longest to tune, by a large factor.

Table 4 also shows that the desktop system performs faster than a GPU node on the supercomputer. Using the supercomputer, it took 1 hour and 47 minutes on average per registration for RMSE on **Sino** with normalisation. The desktop computer was 1.40× faster (it only took 1 hour and 17 minutes on average). The same pattern is observed for ZNCC on **Sino**, the desktop computer was 1.16× faster (34 minutes vs 29 minutes); for DSSIM on **Proj** with normalisation, the desktop computer was 1.41× faster (1 hour and 39 minutes vs 1 hour and 10 minutes). The reasons for this difference are hard to pin point as the systems (both hardware and software) are different. NVIDIA Tesla V100 on SCW are ten months older than the NVIDIA GeForce RTX 2080 Ti on the desktop computer. The Tesla GPUs use the Volta architecture whereas the GeForce RTX ones use the Turing architecture.

4.3. Visual Comparison of the Results

As RMSE on **Sino** with normalisation performed significantly better than any other objective function, we




Figure 11: Parallel coordinate plot showing the successive ZNCC values between the experimental CT slice and the CT slices reconstructed using simulated sinograms for the 25 runs of each objective function. Each line corresponds to an individual run. The first three axes identify the objective function. The different colours indicate the image metrics used in the objective function. Other axes indicate the ZNCC values after the consecutive stages of the whole registration pipeline.




Figure 12: The data is filtered to highlight the 25 runs of RMSE on **Sino** with normalisation. We can identify the outliers where the optimisation algorithm fell in a local minimum at an early stage of the registration.




Figure 13: The data is filtered to highlight the 25 runs of ZNCC on **Sino**.




Figure 14: The data is filtered to highlight the 25 runs of ZNCC on **Proj**.




Figure 15: Scatter plot of the two best objective functions for 25 runs each. Points in the lower right corner correspond to runs that produced the most accurate results in the smallest amount of time.

Table 3

Systems used to perform our tests.

	GPU node at SCW	PC
RAM	395 GB	32 GB
CPU	Intel® Xeon® Gold 6248 CPU @ 2.50GHz × 40	Intel® Core™ i9-9900K CPU @ 3.60GHz × 16
GPU chip	NVIDIA Tesla V100	NVIDIA GeForce RTX 2080 Ti
GPU Architecture	Volta	Turing
GPU Release date	December 7, 2017	September 27, 2018
OS	Red Hat Enterprise Linux Server 7.8	openSUSE Leap 15.2
OpenGL	4.5.0	4.5.0
Drivers	460.27.04	460.67

only consider it in this section. We extract the result of the worst, median and best runs in terms of ZNCC. We compare these three images with the real CT slice using the checkerboard method, which alternates tiles from the first and the second images (see Figures 16 and 17). It allows us to visually assess if the fibre, core, streak artefacts and the phase contrast are perfectly aligned on both the real CT and the simulated CT. It is the case for both the median and best runs, but not the worse one, which failed to estimate the radii. The ZNCC of the median and best runs are within 0.15% of each other.

It is clear on the corresponding profiles shown in Figure 18 that the features seen in the real CT image are replicated in the median and best runs: (i) The pixel values in cm^{-1} are comparable, (ii) the core radius is similar, (iii) the fibre radius is similar, (iv) the dip at the centre of the core is similar, (v) phase contrast around the fibre is similar, (vi) phase contrast around the core is present although a bit attenuated. In practice, it is hard to distinguish the median run from the best one, which demonstrate the stability of our solution.

To further validate the accuracy of our overall numerical model, we compare in Table 5 the input linear attenuation coefficients μ from the literature [5], from the CT slice of the experiment at ESRF and the median run using the RMSE on **Sino** with normalisation. It is clear that the reconstructed values from our model are relatively close to those extracted from the CT slice of the experiment at ESRF for all the material (162 vs 154 cm^{-1} for the tungsten, 6 vs 5 cm^{-1} for the silicon carbide, and 13 vs 11 cm^{-1} titanium-aluminium-vanadium alloy).

In supplementary material, we provide a Jupyter Notebook. When the notebook was executed on the desktop computer, successive results obtained during the artificial evolution were recorded and we created a video for each optimisation of the registration pipeline. They are provided as supplementary material.

4.4. Parameter Sensitivity

We now assess the response of our framework in terms of parameter sensitivity. In particular, we must ascertain that the radii and linear attenuation coefficients are consistent

Table 4

Run times in MM:SS (or HH:MM:SS when appropriate) over 25 runs for each objective function. 14 objective functions were evaluated with a NVIDIA Tesla V100 GPU on Supercomputing Wales, and 3 objective functions were evaluated with a NVIDIA GeForce RTX 2080 Ti GPU on a desktop computer.

Objective		System	Matrix	Radii 1	Radii 2	Spectrum	Noise	Phase & Radii	Phase & LSF	Overall	
Proj	w/ norm.	DSSIM	SCW	04:37 ± 01:14	10:47 ± 01:49	06:19 ± 02:22	09:01 ± 01:43	00:01 ± 00:00	50:39 ± 11:48	17:45 ± 02:27	01:39:12 ± 00:13:14
		PC	03:40 ± 01:20	09:49 ± 04:19	04:21 ± 01:24	06:25 ± 00:47	00:01 ± 00:00	32:58 ± 11:41	13:11 ± 01:34	01:10:28 ± 00:11:55	
		MAE	SCW	03:43 ± 01:12	11:54 ± 03:11	06:24 ± 04:53	09:36 ± 01:15	00:01 ± 00:00	43:04 ± 21:31	12:17 ± 03:11	01:27:02 ± 00:21:42
	w/o norm.	RMSE	SCW	03:58 ± 01:41	11:56 ± 03:49	06:38 ± 04:48	08:49 ± 02:15	00:01 ± 00:00	29:41 ± 09:23	13:05 ± 03:02	01:14:11 ± 00:12:47
		ZNCC	SCW	03:46 ± 00:50	10:15 ± 02:47	02:15 ± 02:03	03:32 ± 01:32	00:01 ± 00:00	02:47 ± 03:50	00:15 ± 00:01	00:22:54 ± 00:05:55
		DSSIM	SCW	04:35 ± 00:50	14:44 ± 11:08	09:07 ± 07:14	08:08 ± 00:51	00:01 ± 00:00	47:46 ± 12:29	17:39 ± 02:26	01:42:03 ± 00:18:38
Sino	w/ norm.	MAE	SCW	04:26 ± 02:25	12:33 ± 04:56	05:51 ± 05:10	09:36 ± 01:47	00:01 ± 00:00	25:57 ± 12:02	07:29 ± 04:48	01:05:57 ± 00:12:36
		RMSE	SCW	03:45 ± 00:57	12:22 ± 05:17	04:59 ± 02:33	09:10 ± 01:49	00:01 ± 00:00	19:51 ± 05:33	00:53 ± 00:31	00:51:03 ± 00:09:55
		ZNCC	SCW	04:47 ± 01:06	12:43 ± 01:57	07:40 ± 07:08	08:17 ± 01:34	00:01 ± 00:00	56:00 ± 23:28	18:14 ± 02:17	01:47:45 ± 00:23:04
	w/o norm.	MAE	SCW	04:05 ± 01:14	15:04 ± 10:07	05:53 ± 02:33	30:53 ± 07:39	00:01 ± 00:00	47:42 ± 16:51	15:41 ± 01:58	01:59:23 ± 00:19:26
		RMSE	SCW	03:23 ± 00:46	13:50 ± 05:36	06:56 ± 03:44	24:37 ± 08:04	00:01 ± 00:00	42:20 ± 09:47	16:16 ± 01:21	01:47:26 ± 00:09:02
		ZNCC	PC	02:59 ± 00:56	10:19 ± 02:48	04:28 ± 01:02	19:11 ± 04:55	00:00 ± 00:00	28:17 ± 05:00	11:41 ± 01:04	01:16:59 ± 00:07:10
w/o norm.	MAE	SCW	03:25 ± 00:43	11:38 ± 02:45	04:12 ± 02:15	05:28 ± 01:09	00:01 ± 00:01	07:05 ± 08:47	01:45 ± 01:29	00:33:37 ± 00:10:25	
	RMSE	SCW	03:09 ± 01:10	11:45 ± 06:48	03:16 ± 01:06	04:28 ± 00:51	00:00 ± 00:00	05:31 ± 04:38	00:40 ± 00:33	00:28:54 ± 00:09:13	
	DSSIM	SCW	05:08 ± 01:19	12:31 ± 01:52	05:43 ± 01:13	08:34 ± 01:18	00:01 ± 00:00	06:24 ± 11:47	14:33 ± 05:08	01:52:58 ± 00:13:45	
w/o norm.	MAE	SCW	04:24 ± 01:50	17:49 ± 04:38	08:44 ± 01:54	30:46 ± 08:04	00:01 ± 00:00	21:28 ± 17:33	07:07 ± 03:14	02:50:23 ± 00:18:15	
	RMSE	SCW	04:10 ± 01:32	16:38 ± 04:20	08:32 ± 02:04	25:47 ± 09:59	00:01 ± 00:00	27:49 ± 18:22	29:30 ± 04:35	02:52:30 ± 00:24:02	

Worse run, ZNCC: 90.44%

Median run, ZNCC: 94.04%

Best run, ZNCC: 94.19%




Figure 16: Checkerboard comparison between the real CT slice and the simulated ones for the worse, median and best runs using the RMSE on **Sino** with normalisation. A checkerboard image alternates tiles from the real and the simulated CT slices. Red squares are two examples of tiles from the real CT slices; blues squares from the simulated slices. The colorbar shows the lookup table for all the images in cm^{-1} .

Worse run, ZNCC: 90.44% Median run, ZNCC: 94.04% Best run, ZNCC: 94.19%




Figure 17: Checkerboard comparison between the fibre in the centre of the real CT slice and the corresponding simulated ones for the worse, median and best runs using the RMSE on **Sino** with normalisation. Red squares are two examples of tiles from the real CT slices; blues squares from the simulated slices.

for the runs producing the best results in terms of ZNCC. Again, we use a parallel coordinate plot (see Figure 19 and http://gvirtualxray.sourceforge.net/parallel_coordinates_all_data.php) that we filter to highlight the runs producing the highest possible ZNCC (Figure 20). The radius of the tungsten cores is $8 \mu\text{m}$, and $52\text{-}54 \mu\text{m}$ for the silicon carbide fibres. However, the range of attenuation

Table 5

Linear attenuation coefficients in cm^{-1} from the literature (Theoretical), from the CT slice of the experiment at ESRF (Experimental), and of the median run using the RMSE on **Sino** with normalisation (Simulated).

Structure	Material	Theoretical	Experimental	Simulated
Core	W	341.61	162.34 ± 21.67	153.88 ± 19.97
Fibre	SiC	2.74	5.61 ± 5.73	4.52 ± 3.98
Matrix	Ti90Al6V4	13.13	12.87 ± 3.57	11.34 ± 2.19

coefficient for tungsten is very wide, 110 to 175 cm^{-1} . We further filter the data to highlight the runs providing both the highest possible ZNCC and the most realistic μ_W (162 cm^{-1}) (Figure 21). There is now a clear consensus on the size of the radii (8 and $54 \mu\text{m}$), and the fibre and matrix linear attenuation coefficients (4 , and $12\text{-}13 \text{ cm}^{-1}$ respectively).

After elimination of the four outliers for RMSE on **Sino** with normalisation, the radius of the tungsten cores is $7.91 \pm 0.05 \mu\text{m}$, and $53.71 \pm 0.08 \mu\text{m}$ for the silicon carbide




Figure 18: Profiles extracted from the fibre in the centre of the real CT slice, and the corresponding simulated ones for the worst, median and best runs using the RMSE on **Sino** with normalisation. Grey dotted lines correspond to the average radii after optimisation (see Table 7).

fibres. This is in very strong agreement with the values extracted from the optical microscopic image in Figure 4b.

On this CT slice, the matrix is a rectangle of $(1451.93 \pm 12.32) \times (1062.58 \pm 9.77) \mu\text{m}^2$. It is rotated by 90.68 ± 0.44 degrees. These results show an accurate and stable geometric analysis of the scanned material when using our registration framework with this objective function.

4.5. User study

As demonstrated above, we use RMSE on **Sino** with normalisation as a baseline. We compare the statistics of the parameters (matrix and fibres) that were estimated (i) using our optimisation framework with realistic GPU simulation, and (ii) manually by 13 volunteers using ImageJ or equivalent. Results are shown in Table 6 for the matrix and in Table 7 for the fibres.

The manual estimation of the matrix parameters are within a few percent of the optimisation, although the standard deviations of the matrix sizes are much larger than the optimised ones; this is due to the non regular parallelepiped shape which caused more dispersion in the manual estimations.

The diameter of the tungsten core estimated manually is biased (by +28%) with respect to the nominal value recovered from the microscopic image ($15 \mu\text{m}$ see Figure 4b), whilst the optimised core diameter is not. Possible reasons are that the tungsten cores are affected by strong artefacts, and that they are comparably smaller than the SiC fibres. Both fibre diameter estimations are close to the nominal value ($105 \mu\text{m}$), the optimisation one is a bit better.

Most of the respondents found measuring the properties of the matrix moderate to easy (3.70 ± 1.06 , with 1: very hard, 5: very easy). One participant found it very easy.

They found measuring the radii slightly easier (3.90 ± 0.88). Three participants found it very easy. One participant commented that “manipulating the mouse cursor to the centre of the tungsten core is a bit fiddly”. It

Table 6

Comparison of the size and orientation of the matrix estimated (i) using optimisation and (ii) manually by 13 volunteers.

	Optimisation	Manual
Width (in μm)	1452.8 ± 12.1	1509.5 ± 57.8
Height (in μm)	1061.6 ± 9.9	1081.0 ± 40.3
Rotation (in degrees)	90.7 ± 0.4	91.6 ± 1.5

Table 7

Comparison of the core and fibre diameters estimated (i) using optimisation and (ii) manually by 13 volunteers.

	Optimisation	Manual
Core diameter (in μm)	15.80 ± 0.12	19.44 ± 1.45
Fibre diameter (in μm)	104.96 ± 7.54	104.39 ± 2.03

Table 8

Run times in HH:MM:SS for the registration of the first slice of the 3D dataset.

Optimising	Run time
Matrix geometry (first run)	00:15:41
Matrix geometry (second run)	00:19:47
Beam spectrum	00:37:53
Phase contrast and radii	01:05:42
LSF	N/A
Poisson noise	00:00:33
Total	03:57:22

took 22 ± 12 minutes to complete the tasks. These results (variability, bias, perceived difficulty, and time to complete) demonstrate that extracting the geometrical properties was not trivial, even with simple geometric shapes.

4.6. Change of Acquisition Parameters and Pseudo-3D

New experimental parameters were set to take into account changes of beam spectrum, SOD and SSD. Algorithm 1 was used to generate fully-3D data. A full registration

(including the matrix geometry, the core and fibre geometries, beam spectrum, phase contrast, and noise) was performed for the first slice. To recover the tilt of fibres, it is necessary to perform a partial registration (only the matrix geometry and the extraction of the fibre centres) of another slice, e.g. the last one. We performed one single run only with RMSE on **Sino** with normalisation as it was shown in the previous sections that it was producing good results and it was stable. This run is completed using one

GPU node on the supercomputer. We optimised i) the matrix geometry (with tolerance= 10^{-4} and with restart), ii) the beam spectrum, iii) the phase contrast and radii, and iv) the Poisson noise for the first slice. Restart is used with a smaller tolerance to make sure the guess for the first slice is accurate as it will be used as initial guess for the next slice.

For other slices, only the matrix geometry is optimised (with tolerance= 10^{-2} and without restart). The ZNCC of the first slice is 90.81%. It is 91.37 ± 0.33 for the other slices. These results show that the strategy of optimising




Figure 19: Parallel coordinate plot of the subset of the data collected during the registration runs performed on the supercomputer.




Figure 20: The data is filtered to highlight the runs producing the highest possible ZNCC.




Figure 21: The data is filtered to highlight the methods producing both 1) the highest possible ZNCC and 2) the most realistic μ_W .

Table 9

Comparison of the average ZNCCs of all the slices of the 3D scan for images simulated i) without artefact (i.e. using the monochromatic Beer-Lambert attenuation law only), ii) with beam hardening (i.e. using the polychromatic Beer-Lambert attenuation law), and iii) with beam hardening and phase contrast.

Simulation model	ZNCC in %
Without artefacts	90.80 ± 0.30
With beam hardening	90.85 ± 0.30
With phase contrast	91.37 ± 0.33

the whole model on the first slice and using it for subsequent slices is appropriate as the ZNCCs are comparable. Three slices (#739, 740, & 741) had a ZNCC relatively smaller (88.94, 88.75, & 89.63). A closer inspection of the projection data showed that four successive columns were incorrectly recorded in the image files during the experiment. It led to a ring artefact in these slices. These were not modelled in the simulation, which lead to this lower correlation.

Run times in HH:MM:SS for the first slice are given in Table 8. The run time to register subsequent slices is $00:00:41 \pm 00:00:03$. It took 14:10:25 to register 1024 slices in total. To further improve this performance, it is possible to use a more recent GPU on a desktop computer, or use several GPU nodes at the same time. In this case every GPU node would register a different subset of the whole dataset.

Streaks and phase contrast artefacts in the simulated slices are, again, consistent with the experimental data (see Figure 22). Figure 23 shows a blend comparison between a region of interest (ROI) of the first and last slices of the real experiments, and of the simulated slices after registration.

It shows that the tilt of the fibres is also recovered. Once the geometrical properties extracted, it is possible to generate a perfect 3D CT scan without artefacts using Equation 3. Figure 24 shows the first slice with and without artefacts. Table 9 provides a comparison of the average ZNCCs for images simulated i) without artefact (i.e. using the monochromatic Beer-Lambert attenuation law only), ii) with the beam hardening (i.e. using the polychromatic Beer-Lambert attenuation law), and iii) with beam hardening and phase contrast. It is clear that modelling the monochromatic Beer-Lambert attenuation law fails to generate realistic simulations in this experimental context. Considering the polychromatic case marginally improves the results: The ZNCC only improved by 0.05%. Extending the polychromatic Beer-Lambert attenuation law with the phase contrast to simulate X-ray projections lead to the most accurate CT reconstructions. Estimating the impulse response of the detector has the potential to improve this score further. We provide a video in supplementary material showing the same ROI as above for each slice to compare the experimental data with the monochromatic Beer-Lambert attenuation law and with extended model.

A difficulty with this new experimental dataset was that the




Figure 22: ROI of the first slice reconstructed from simulated data after registration (left) and of the last slice (right).




Figure 23: Blend comparison between a ROI of the first and last slices of the real experiment (left), and of the simulated slices after registration (right).




Figure 24: ROI of the first slice reconstructed from simulated data with (left) and without (right) artefacts.

sample was larger than the detector. Consequently, the centre of rotation was off by a large margin and an angular span of 360° was used. When reformatting the projections over 180° to avoid sample truncation in the projection domain, misalignment of the centre of rotation with the central column of the projections may occur. The consequence of this misalignment was that the phase contrast artefact around a fibre would not be rotationally symmetrical, which affected the optimisation of the phase contrast parameters.

5. Discussion

The extraction of CAD surfaces from X-ray computed tomography is the key aspect in dimensional metrology. We have seen have strong artefacts might occur in the reconstructed CT data, and two major sources of artefacts preclude an accurate surface extraction: (i) dark lines in the direction of the alignments of strongly attenuated regions and (ii) phase patterns at the materials interfaces. Both type of CT artefacts are not restricted to ESRF data, but can occur in volumes acquired with standard industrial CT

instrumentation. With such degraded CT data, standard image processing techniques (segmentation, labelling...) cannot address the surface extraction issue. Manual measurements, whilst they may seem to be relatively easy to perform, can be prone to bias and unreliable. It is therefore essential to have a realistic X-ray simulation tool embedded in an optimisation scheme to extract the CAD surfaces. The method proposed here is suitable when it is possible to generate surface models. One limitation is that all the geometrical properties, e.g. tilting and warping the shapes, must be implemented. Taking into account features such as pores or fibre tows in real materials, even if made of basic shapes, is not straightforward.

Many feature parameters need to be determined by optimisation, but the ones that are not sample specific can be fixed: for example, the X-ray spectrum and the CT geometry can be calibrated off-line. We clearly see from this investigation study that manual tuning and ad-hoc processes might be required, especially for the initialisation of the feature extraction routine.

The optimisation has been carried out for a single row as the sample was translation invariant. The generalisation to 3D reconstructed volumes with more complex quadrics is straightforward and would only require more computation time. This is particularly attractive if only a partial registration is needed for subsequential slices. The choice of the objective function for the optimisation is very important, and from our tests it seems that RMSE on the projection data (after taking the log) is the most accurate. A few outliers were observed using this objective function due to strong local minima. Using black-box optimisation, one might miss the opportunity to tweak the algorithm to maximise the outcome. The result of the evolutionary algorithm will strongly depend on initial meta-parameters such as boundary conditions, probabilities of genetic operators and stopping criteria. Despite of this, it is the most commonly used approach when solving optimisation problems in engineering. To avoid this problem, the algorithm's meta-parameters could be tweaked. For example, we could perform quick manual measurements.

We saw such measurements tend to be plausible and inaccurate at the same time. Another advantage of using manual measurements as an initial guess would be to speed up the convergence of the algorithm. It could also be used to limit the size of the search space. Alternatively, if using manual measurements is impractical, techniques used in evolutionary computing, such as *restart*, would be successful in avoiding the outliers we observed.

This study was implemented as a proof of concept using Python. More can be done to improve its computational efficiency. For example, it is possible to compute the objective values using GPU computing. It would address some of the bottlenecks, such as the use of the \ln function in Equation 4 (which is rather slow on CPU compared to GPU), avoid data transfer from the GPU to the main memory (RAM) (which, again, is rather slow), and the lack of parallelism when computing the RMSE in Equation 13.

This preliminary study can be extended to include this improved forward model within an iterative CT reconstruction algorithm [11]. It has the potential to attenuate the artefacts seen in the reconstructed CT slices, making the data not only more accurate, but also easier to exploit.

However, another limitation is that it is of importance to replicate the same experimental conditions in the simulation. For example, in our pseudo-3D example, the sample is larger than the actual detector. To accommodate for this, the centre of rotation of the sample is off compared to the central column of the detector. The centre of rotation of the projection data may be slightly misaligned when converting from 360° to 180° . This is not the case in the simulated experiment as it directly used projections over 180° . To address this, it would be necessary to model and optimise the centre of rotation of the sample in the simulation and convert simulated projections from 360° to 180° , in the same way real projections were transformed before the tomography reconstruction using standard algorithms.

6. Conclusion

Our aim was to show that a realistic projection model is required to automatically produce an accurate CAD model of a material from X-ray tomography when CT data is so corrupted by artefacts that geometric analysis is challenging. Our model takes into account geometrical properties as well as beam hardening, impulse response of the detector, phase contrast, and photon noise. Every unknown of our model was automatically tuned using a stochastic optimisation algorithm: CMA-ES. It minimises the discrepancies between the projection data acquired at ESRF and the projections simulated using our model. The code was implemented on GPU using gVirtualXRay, an open source tool for simulating X-ray images. 14 different objectives functions were evaluated on a supercomputer. Our results showed that the CT image reconstructed from the simulated data is $93.68\% \pm 0.87\%$ correlated with the experimental CT slice. All the features seen in the experimental CT image, including artefacts, were accurately replicated. Our framework enabled an automatic, accurate and stable geometric analysis of the scanned material.

Being able to accurately model the imaging chain opens up new perspective, e.g. the inclusion of such models in iterative reconstruction algorithms. In future work, we aim to take advantage of the knowledge discovered with our registration framework. One topic of interest is to correct imaging artefacts and improve the value of linear attenuation coefficients.

Ethical approval

Ethical approval for the user-study was obtained from the School of Computer Science and Electronic Engineering's

Ethics Committee at Bangor University (APPROVAL CSEE-R-2001-CG-003).

Acknowledgements

We thank Jean-Yves Buffière and Ce Xiao of MATEIS laboratory (Lyon), and Wolfgang Ludwig of ESRF (Grenoble) for the new projection and CT data. We also thank NVIDIA Corporation for the donation of the NVIDIA TITAN Xp GPU used in the development and validation of gVirtualXRray, and Supercomputing Wales for the use of its supercomputer.

References

- [1] Allison, J., Amako, K., Apostolakis, J., Arce, P., Asai, M., Aso, T., Bagli, E., Bagulya, A., Banerjee, S., Barrand, G., Beck, B., Bogdanov, A., Brandt, D., Brown, J., Burkhardt, H., Canal, P., Cano-Ott, D., Chauvie, S., Cho, K., Cirrone, G., Cooperman, G., Cortés-Giraldo, M., Cosmo, G., Cuttone, G., Depaola, G., Desorgher, L., Dong, X., Dotti, A., Elvira, V., Folger, G., Francis, Z., Galoyan, A., Garnier, L., Gayer, M., Genser, K., Grichine, V., Guatelli, S., Guèye, P., Gumplinger, P., Howard, A., Hřivnáčová, I., Hwang, S., Incerti, S., Ivanchenko, A., Ivanchenko, V., Jones, F., Jun, S., Kaitaniemi, P., Karakatsanis, N., Karamitros, M., Kelsey, M., Kimura, A., Koi, T., Kurashige, H., Lechner, A., Lee, S., Longo, F., Maire, M., Mancusi, D., Mantero, A., Mendoza, E., Morgan, B., Murakami, K., Nikitina, T., Pandola, L., Paprocki, P., Perl, J., Petrović, I., Pia, M., Pokorski, W., Quesada, J., Raine, M., Reis, M., Ribon, A., Ristić Fira, A., Romano, F., Russo, G., Santin, G., Sasaki, T., Sawkey, D., Shin, J., Strakovsky, I., Taborda, A., Tanaka, S., Tomé, B., Toshito, T., Tran, H., Truscott, P., Urban, L., Uzhinsky, V., Verbeke, J., Verderi, M., Wendt, B., Wenzel, H., Wright, D., Wright, D., Yamashita, T., Yarba, J., Yoshida, H., 2016. Recent developments in Geant4. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 835, 186–225. doi:10.1016/j.nima.2016.06.125.
- [2] Als-Nielsen, J., Mcmorrow, D., 2011. Elements of Modern X-Ray Physics. John Wiley & Sons. URL: https://www.ebook.de/de/product/14773530/des_mcmorrow_als_nielsen_elements_of_modern_x_ray_physics.html.
- [3] Auger, A., Hansen, N., 2005. A restart CMA evolution strategy with increasing population size, in: McKay, B., et al. (Eds.), The 2005 IEEE International Congress on Evolutionary Computation (CEC'05), pp. 1769–1776.
- [4] Berger, M.J., Hubbell, J.H., 1987. XCOM: Photon cross sections on a personal computer. Technical Report. National Institute of Standards and Technology. URL: <https://doi.org/10.2172/6016002>.
- [5] Berger, M.J., Hubbell, J.H., Seltzer, S.M., Chang, J., Coursey, J.S., Sukumar, R., Zucker, D.S., Olsen, K., 1998. XCOM: Photon cross section database. URL: <https://www.physics.nist.gov/PhysRefData/Xcom/html/xcom1.html>.
- [6] Endrizzi, M., 2018. X-ray phase-contrast imaging. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 878, 88–98. doi:10.1016/j.nima.2017.07.036. radiation Imaging Techniques and Applications.
- [7] Freud, N., Duvauchelle, P., Létang, J., Babot, D., 2006. Fast and robust ray casting algorithms for virtual x-ray imaging. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 248, 175–180. doi:10.1016/j.nimb.2006.03.009.
- [8] Haiderbhai, M., Ledesma, S., Lee, S.C., Seibold, M., Fürnstahl, P., Navab, N., Fallavollita, P., 2020. pix2xray: converting RGB images into x-rays using generative adversarial networks. International Journal of Computer Assisted Radiology and Surgery 15, 973–980. doi:10.1007/s11548-020-02159-2.
- [9] Hansen, N., 2009. Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed, in: Workshop Proceedings of the GECCO Genetic and Evolutionary Computation Conference, ACM. pp. 2389–2395.
- [10] Hansen, N., Ostermeier, A., 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. Evolutionary Computation 9, 159–195. doi:10.1162/106365601750190398.
- [11] Hehn, L., Gradl, R., Dierolf, M., Morgan, K.S., Paganin, D.M., Pfeiffer, F., 2020. Model-based iterative reconstruction for propagation-based phase-contrast x-ray CT including models for the source and the detector. IEEE Transactions on Medical Imaging 39, 1975–1987. doi:10.1109/TMI.2019.2962615.
- [12] Hiller, J., Hornberger, P., 2016. Measurement accuracy in x-ray computed tomography metrology: Toward a systematic analysis of interference effects in tomographic imaging. Precision Engineering 45, 18–32. doi:10.1016/j.precisioneng.2015.12.003.
- [13] Illingworth, J., Kittler, J., 1987. The adaptive hough transform. IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-9, 690–698. doi:10.1109/TPAMI.1987.4767964.
- [14] ISO Central Secretary, 2021. Geometrical product specifications (GPS) – Acceptance and reverification tests for coordinate measuring systems (CMS) – Part 11: CMSs using the principle of X-ray computed tomography (CT). Standard ISO/IEC 10360-11:2021. International Organization for Standardization, Geneva, CH.
- [15] Jailin, C., Buffière, J.Y., Hild, F., Poncelet, M., Roux, S., 2017. On the use of flat-fields for tomographic reconstruction. Journal of Synchrotron Radiation 24, 220–231. doi:10.1107/S1600577516015812.
- [16] Johansson, J., Forsell, C., 2016. Evaluation of parallel coordinates: Overview, categorization and guidelines for future research. IEEE Transactions on Visualization and Computer Graphics 22, 579–588. doi:10.1109/TVCG.2015.2466992.
- [17] Joskowicz, L., Cohen, D., Caplan, N., Sosna, J., 2018. Automatic segmentation variability estimation with segmentation priors. Medical Image Analysis 50, 54–64. doi:10.1016/j.media.2018.08.006.
- [18] Kovács, I., Várady, T., Salvi, P., 2015. Applying geometric constraints for perfecting CAD models in reverse engineering. Graphical Models 82, 44–57. doi:10.1016/j.gmod.2015.06.002.
- [19] Liu, Z., Bicer, T., Kettimuthu, R., Gursoy, D., De Carlo, F., Foster, L., 2020. TomoGAN: low-dose synchrotron x-ray tomography with generative adversarial networks: discussion. J. Opt. Soc. Am. A 37, 422–434. doi:10.1364/JOSAA.375595.
- [20] Monnin, P., Bulling, S., Hoszowska, J., Valley, J.F., Meuli, R., Verdun, F.R., 2004. Quantitative characterization of edge enhancement in phase contrast x-ray imaging. Medical Physics 31, 1372–1383. doi:10.1118/1.1755568.
- [21] Oliveira, F.P., Tavares, J.M.R., 2014. Medical image registration: a review. Computer Methods in Biomechanics and Biomedical Engineering 17, 73–93. doi:10.1080/10255842.2012.670855. pMID: 22435355.
- [22] Peterzol, A., Olivo, A., Rigon, L., Pani, S., Dreossi, D., 2005. The effects of the imaging system on the validity limits of the ray-optical approach to phase contrast imaging. Medical Physics 32, 3617–3627. doi:10.1118/1.2126207.
- [23] Razlighi, Q.R., Kehtarnavaz, N., Yousefi, S., 2013. Evaluating similarity measures for brain image registration. Journal of Visual Communication and Image Representation 24, 977–987. doi:10.1016/j.jvcir.2013.06.010.
- [24] Rogelj, P., Kovacic, S., 2001. Similarity measures for nonrigid registration, in: Sonka, M., Hanson, K.M. (Eds.), Medical Imaging 2001: Image Processing, International Society for Optics and Photonics. SPIE. pp. 569–578. doi:10.1117/12.431131.
- [25] Stayman, J.W., Otake, Y., Prince, J.L., Khanna, A.J., Siewerdsen, J.H., 2012. Model-based tomographic reconstruction of objects containing known components. IEEE Transactions on Medical Imaging 31, 1837–1848. doi:10.1109/tmi.2012.2199763.
- [26] Tekawade, A., Sforzo, B.A., Matusik, K.E., Kastengren, A.L., Powell,

- C.F., 2019. High-fidelity geometry generation from CT data using convolutional neural networks, in: Müller, B., Wang, G. (Eds.), *Developments in X-Ray Tomography XII*, SPIE. doi:10.1117/12.2540442.
- [27] Thompson, A., Maskery, I., Leach, R.K., 2016. X-ray computed tomography for additive manufacturing: a review. *Measurement Science and Technology* 27, 072001. doi:10.1088/0957-0233/27/7/072001.
- [28] Vidal, F., Létang, J., Peix, G., Cloetens, P., 2005. Investigation of artefact sources in synchrotron microtomography via virtual x-ray imaging. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 234, 333–348. doi:10.1016/j.nimb.2005.02.003.
- [29] Vidal, F.P., 2021. *gVirtualXR*. URL: <http://gvirtualxray.sourceforge.net/>.
- [30] Vidal, F.P., Garnier, M., Freud, N., Létang, J.M., John, N.W., 2009. Simulation of X-ray attenuation on the GPU, in: *Proceedings of Theory and Practice of Computer Graphics 2009*, Eurographics Association, Cardiff, UK. pp. 25–32. doi:10.2312/LocalChapterEvents/TPCG/TPCG09/025-032.
- [31] Vidal, F.P., Mitchell, I.T., Létang, J.M., 2021a. GitHub repository: Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts. URL: <https://github.com/effepivi/CT2CAD-Fibres>.
- [32] Vidal, F.P., Mitchell, I.T., Létang, J.M., 2021b. Supplementary material: Use of fast realistic simulations on GPU to extract CAD models from microtomographic data in the presence of strong CT artefacts. *Mendeley Data* doi:10.17632/r8b5mcrj6j.1.
- [33] Vidal, F.P., Villard, P.F., 2016. Development and validation of real-time simulation of x-ray imaging with respiratory motion. *Computerized Medical Imaging and Graphics* 49, 1–15. doi:10.1016/j.compmedimag.2015.12.002.
- [34] Villarraga-Gómez, H., Herazo, E.L., Smith, S.T., 2019. X-ray computed tomography: from medical imaging to dimensional metrology. *Precision Engineering* 60, 544–569. doi:10.1016/j.precisioneng.2019.06.007.
- [35] Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 600–612. doi:10.1109/TIP.2003.819861.
- [36] Wen, T., Mihail, R., Al-maliki, s., Letang, J., Vidal, F., 2019. Registration of 3D Triangular Models to 2D X-ray Projections Using Black-box Optimisation and X-ray Simulation, in: Vidal, F.P., Tam, G.K.L., Roberts, J.C. (Eds.), *Computer Graphics and Visual Computing (CGVC)*, The Eurographics Association. pp. 105–113. doi:10.2312/cgvc.20191265.
- [37] Withers, P.J., Bouman, C., Carmignato, S., Cnudde, V., Grimaldi, D., Hagen, C.K., Maire, E., Manley, M., Du Plessis, A., Stock, S.R., 2021. X-ray computed tomography. *Nature Reviews Methods Primers* 1. doi:10.1038/s43586-021-00015-4.
- [38] Zitová, B., Flusser, J., 2003. Image registration methods: a survey. *Image and Vision Computing* 21, 977–1000. doi:10.1016/S0262-8856(03)00137-9.