Registration of 3D triangular models to 2D X-ray projections using black-box optimisation and X-ray simulation

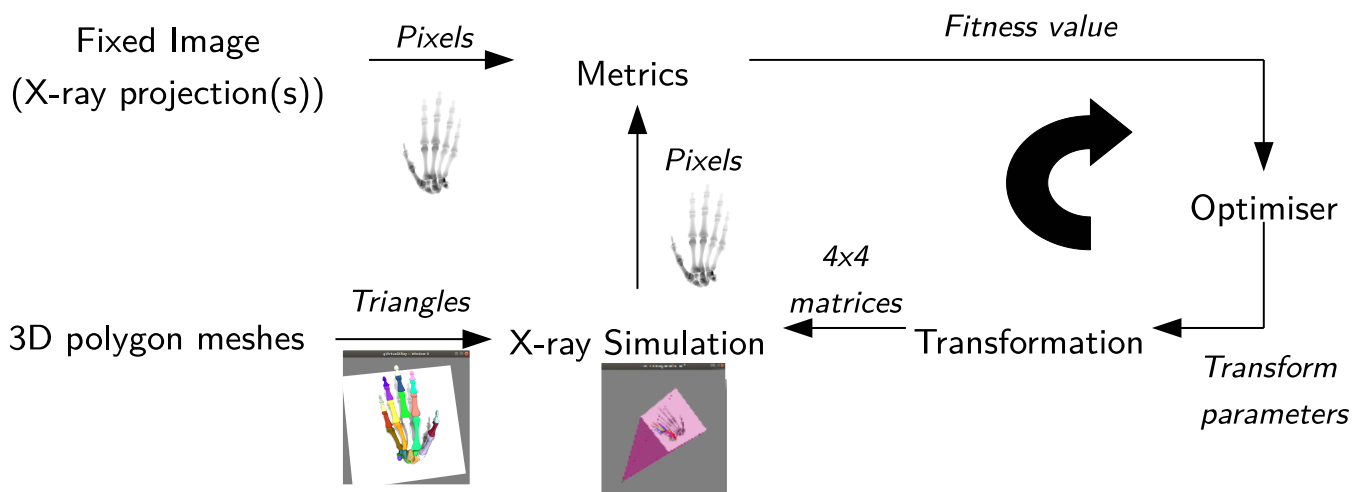T. Wen[1] , R. P. Mihail[2] , S. F. Al-Maliki[1,3] , J.-M. Létang[4] , and F. P. Vidal[†1]

[1] School of Computer Science and Electronic Engineering, Bangor University, UK
[2] Department of Computer Science, Valdosta State University, GA, USA
[3] University of Basrah, Iraq
[4] Université de Lyon, CREATIS; CNRS UMR5220; Inserm U1044; INSA-Lyon; Université Lyon 1; Centre Léon Bérard, France

Figure 1: *Registration pipeline based on X-ray simulation and black-box optimisation techniques.*

Abstract

Registration has been studied extensively for the past few decades. In this paper we propose to solve the registration of 3D triangular models onto 2D X-ray projections. Our approach relies extensively on global optimisation methods and fast X-ray simulation on GPU. To evaluate our pipeline, each optimisation is repeated 15 times to gather statistically meaningful results, in particular to assess the reproducibility of the outputs. We demonstrate the validity of our approach on two registration problems: i) 3D kinematic configuration of a 3D hand model, i.e. the recovery of the original hand pose from a postero-anterior (PA) view radiograph. The performance is measured by Mean Absolute Error (MAE). ii) Automatic estimation of the position and rigid transformation of geometric shapes (cube and cylinders) to match an actual metallic sample made of Ti/SiC fibre composite with tungsten (W) cores. In this case the performance is measured in term of F-score (86%), accuracy (95%), precision (75%), recall (100%), and true negative rate (94%). Our registration framework is successful for both test-cases when using a suitable optimisation algorithm.

1. Introduction

In this paper, we propose to deform polygon meshes so that their corresponding X-ray projection(s) finely match an input image.

Our framework heavily relies on optimisation techniques for the tuning of deformation parameters and on fast X-ray simulation on graphics processing unit (GPU) for the generation of X-ray projections.

[†] Secretary UK Chapter of the Eurographics Association

This research is related to registration where a moving dataset

(called *source*) is deformed to match another one, which is fixed (called *target*). In this article we focus on 3D/2D registration. In most 3D/2D registration problems, the source corresponds to a 3D volume (e.g. voxels from a computed tomography (CT) or magnetic resonance imaging (MRI) dataset) and the target is a 2D image (e.g. an X-ray radiograph) [vPT*05,MTLP12]. For example, Birkfellner *et al.* used digitally reconstructed radiographs (DRRs) (reformatting of CT voxel data to generate images equivalent to X-ray projections) in a rigid registration of a 3D volume from CT voxel data onto a 2D X-ray projection [BWB*03]. It is also possible to register polygon meshes onto 2D images. For example Deligianni *et al.* registered polygon meshes onto 2D endoscopic images [DCY06].

In our application the source is a scenegraph containing multiple 3D triangular models, and the target corresponds to an X-ray radiograph or a sinogram. Every 3D model can be deformed in nonrigid framework. Our hypothesis is that fast X-ray simulation from polygons on the GPU can be embedded in objective functions to transform 3D surface models using black-box optimisation. Two test-cases are investigated to validate the feasibility of our approach: i) the registration of a 3D hand model on a synthetic 2D X-ray radiograph, and ii) the registration of titanium/silicon carbide (Ti/SiC) fibre composite with tungsten (W) cores in 3D X-ray microtomography (μ-CT). The first one is a proof-of-concept for a medical application in rheumatoid arthritis (RA). The second one uses actual data from a material science application to study the structure of tiny metal samples.

The rest of the paper is organised as follows: Section 2 discusses background material, Section 3 describes the methodology, Section 4 discusses the results and evaluation, and we conclude the work and give future directions in this research in Section 5.

2. Background

2.1. Hand registration for rheumatoid arthritis

Radiograph (X-ray) views of hands help clinicians diagnose a wide range of conditions from simple fractures to skeletal maturity and bone and joint damage due to autoimmune disorders. Simple radiograph imaging of the hand and wrist are often preferred over volumetric imaging modalities, such as CT and MRI, due to lower patient radiation doses and wide availability of equipment.

Hand radiographs are 2D projections of 3D anatomy, therefore one view is often not enough for clinical diagnoses, so radiologists rely heavily on experience and atlas matching of multiple views to infer often small, but relevant irregularities. Such is the case for RA, an autoimmune disease that affects 0.5% to 1% of the population in the Northern hemisphere [SBB*16]. RA manifests with the presence of acute inflammation sites, located primarily in the small joints of the hands and feet. Inflammation causes swelling and pain. Over time, joint damage and deformity through erosions and misalignment occur, which lead to functional loss of fine motor ability. If detected early, RA can be treated with disease modifying antirheumatic drugs (DMARDs). Treatments can induce remission, however, treatment efficacy
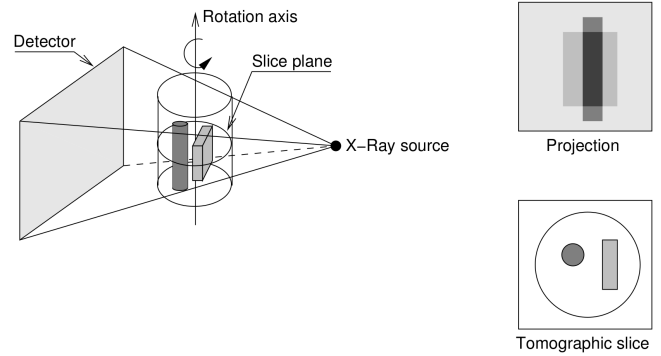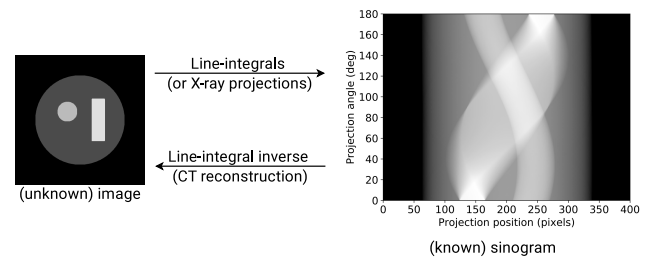

Figure 2: *X-ray tomography.*

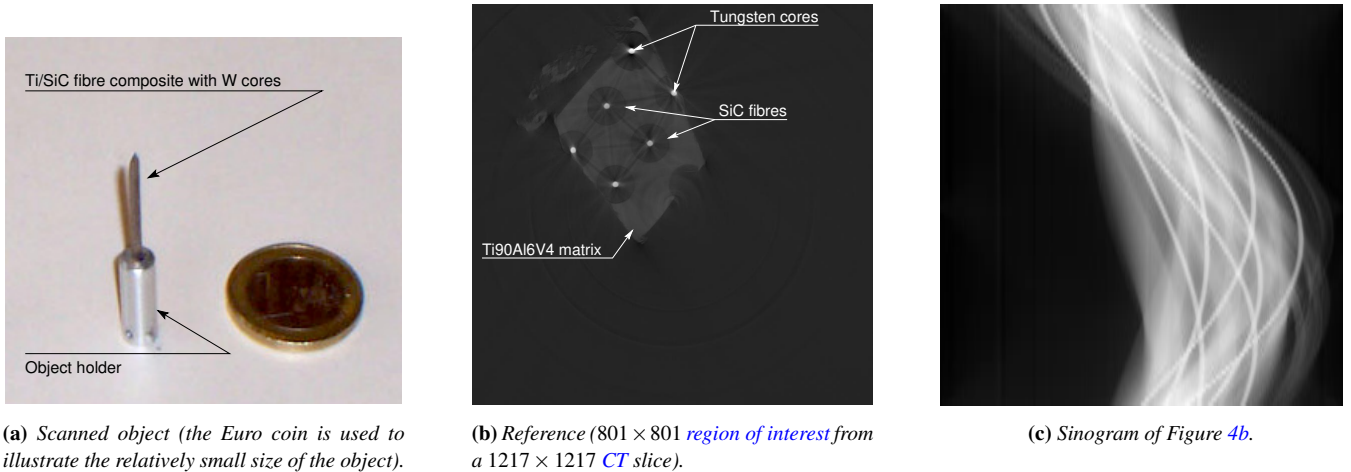**Figure 3:** *Inverse problem: reconstruction of an unknown image from a known sinogram.*

varies between patients, and often only slow the radiographically visible progression [SSG18].

2.2. Nondestructive testing 3D X-ray microtomography

Computed tomography is a 3D imaging modality that takes as input a large set of radiographs taken at different angles around the patient or an object in material science (see Figures 2 and 4). This set of projections expressed as a function of the angle (see Figure 3) is called the sinogram. Each pixel of this sinogram can be seen as a line-integral (also referred to in 2D as the Radon transform) of the local attenuation coefficients (see μ_i in Eq. 1) along the X-ray path. The reconstruction process is based on the inversion of this set of line-integrals and is usually analytically implemented by filtered back-projection (FBP). The output is the 3D map of the attenuation coefficients expressed as a CT number, namely the Hounsfield Unit (HU) in medical science, or as a linear attenuation coefficient in material science.

3D X-ray microtomography is standard practise for qualitative and quantitative testing of materials [MW14]. Virtual X-ray imaging is now almost systematically associated with 3D tomography for multiple applications, such as the investigation of sources of artefacts [VLPC05], the estimation of scatter contribution using deep convolutional network [MSKK18], or the geometric calibration via iterative registration [BTKB*13].

Registration is a generic issue in computer vision and goes well beyond the medical field. In nondestructive testing the quality of part-to-part or part-to-CAD comparisons depends largely on the

(a) *Scanned object (the Euro coin is used to illustrate the relatively small size of the object).*

(b) *Reference (*801 × 801 *region of interest from a* 1217 × 1217 *CT slice).*

(c) *Sinogram of Figure 4b.*

Figure 4: *CT slice of a Ti/SiC fibre composite with tungsten cores. Experiment was carried out at ESRF using synchrotron radiation. The fibre diameter is 140µm and its core diameter 30µm.*

accuracy of the registration module. Moreover, severe artefacts are observed in 3D images of composite materials reconstructed using standard tomographic methods when dense inclusions are present [VLPC05]. As a consequence, the reconstruction is no longer a quantitative map of the linear attenuation distribution in the sample. X-ray simulation coupled to registration makes it possible to reduce the imaging artefacts.

2.3. Contributions

This paper proposes a novel idea, to the best of the author's knowledge never investigated prior to the publication of this manuscript. We show and evaluate the feasibility of several optimisation methods to automatically solve two registration problems:

- 3D kinematic configuration of a hand model, i.e. the recovery of the original hand pose in a postero-anterior (PA) view radiograph. Due to the large variability in clinical practice, no assumption is made in the registration on the source-detector distance (SDD), source-object distance (SOD) and X-ray beam spectrum. The SDD, SOD, rotation angles of all the joints in the hand are automatically estimated by our registration framework. The knowledge of the 3D pose of a hand in a radiograph will allow us to trivially and accurately solve bone segmentation and flag clinically relevant (e.g. rheumatoid) differences from a baseline.
- Automatic estimation of the position and rigid transformation of geometric shapes (cube and cylinders) to match an actual metallic sample made of Ti/SiC fibre composite with tungsten cores. The fibre diameter is 140µm and its core diameter 30 µm. Figure 4 shows the corresponding CT reconstruction and its projections as a sinogram. The experiment was carried out at the European Synchrotron Radiation Facility (ESRF) on the BM05 beamline (distance source/detector: 55 m, i.e. parallel beam). The SOD is 80 mm. The X-ray detector is a CCD camera of 1024 × 1024 pixels resolution, whose effective pixel

size is 1.9 µm. Synchrotron radiation of 33 keV was used. 900 projections over an 180° angular span were acquired. This *a priori* knowledge can be modelled in the X-ray simulation environment.

3. Methodology

Figure 1 provides a summary of our approach. It corresponds to a 3D/2D registration problem where a moving dataset (called source) is deformed to match another one, which is fixed (called target). In most 3D/2D registration problems, the source corresponds to a volume dataset (e.g. from a CT or MRI) and the target is a 2D image. In our application the source is a scenegraph containing 3D triangular models, and the target is a simulated 2D image (X-ray radiograph or sinogram). The 3D models are deformed using rigid body transformations applied on the triangular meshes. The parameters of the transformations are finely tuned to minimise the difference between the target and the image simulated using the deformed models. This parameter tuning is achieved using black-box optimisation. Note that it is also possible to maximise the correlation (or any other measurement of similarity) between the two images.

3.1. X-ray simulation toolbox

X-ray simulation is extensively studied in physics; with application in medicine and material science. Physically-based simulation codes are available [A*03, B*04, BSFVS95, FDLB06]. The most popular technique relies on the Monte Carlo (MC) method. It is often used in dosimetry for radiotherapy due to its high level of accuracy. MC-based simulation codes implement probabilistic X-ray interaction models for the transport of photons in matter. X-photons cross matter. During their path into the object, they can interact with matter. Photons can reach the detector without any interaction (see '1' in Figure 5). They can be absorbed, in which case they do not reach the detector and do not contribute to the X-ray projection (see '2'). Scattered

photons decrease the image quality with noise and blur (see '3'). Photons can also be scattered then absorbed (see '4'), or even be scattered multiple times. At each iteration of the MC algorithm, interaction events may occur between each photon and the material it is crossing based on these probabilistic models. For some applications this approach leads to excessive computation time due to the stochastic nature of MC methods. Weeks may be required to simulate a single X-ray projection with an acceptable signal-to-noise ratio (SNR). For transmission imaging (inc. radiography, CT, cone beam computed tomography (CBCT), and fluoroscopy), deterministic calculation using the Beer-Lambert law (also known as attenuation law) might be considered as a sufficient description for the expected value of the number of photons that did not interact with the objects:

$$N_{out}(E) = N_{in}(E) \exp\left(-\sum_i \mu_i(E, \rho, Z) L_p(i)\right) \qquad (1)$$

with $N_{in}(E)$ the number of incident photons at energy E, $N_{out}(E)$ the number of transmitted photons of energy E, μ_i the linear attenuation coefficient (in cm^{-1}) of the i^{th} object and $L_p(i)$ the path length of the ray in the i^{th} object. μ_i depends on E the energy of incident photons, ρ the material density of the object, and Z the atomic number of the object material.

In this case, methods based on the ray-tracing principle are often used as a fast alternative to MC methods [FDLB06]. To speed-up computations, GPU can be used, but often focusing on radiotherapy and voxelised data [B*12, JZJ14]. For this 3D/2D registration problem, we rely on the simulation of X-ray images from polygon meshes. We make use of gVirtualXRay, a cross-platform Open-Source library [Vid13, VV16]. It implements a multi-pass renderer to solve Eq. 1. It is written in C++ using OpenGL and OpenGL Shading Language (GLSL) to provide real-time performance. Wrappers to other popular languages are also provided. Here we use the Python3 wrapper. Technical details can be found in [VGF*09, VGF*10, VV16]. gVirtualXRay has been used in various medical applications where speed and accuracy are both requirements, including radiograph generation from dynamic polygon meshes (to simulate patient posing and respiration) [SMV*17, VGF*10] and the simulation of CT data acquisition and reconstruction including beam hardening or motion artefacts due to the respiration [VV15].
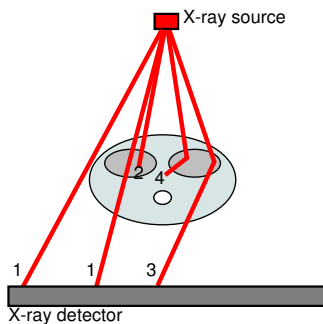

Figure 5: *X-photons/matter interactions.*

3.2. Optimisation

The implementation of the registration is written in Python using the *minimize* package provided by SciPy [JOP*01] and our own implementation of stochastic optimisation methods available on GitHub [AMV19].

We model the registration as a black-box optimisation problem where a simulated image (called *source*) should accurately match a given input image (called *target*). These image correspond to one or several X-ray projections depending on the problem considered: a radiograph for the hand problem, a sinogram for the μ-CT problem. The parameters used to simulate the X-ray projections are finely tuned by an optimisation algorithm using an objective function. In minimisation, the objective function corresponds to a distance metrics between the *source* and *target*. The optimisation algorithm aims to lower its value as much as possible. In maximisation, similarity or correlation metrics are used instead and the optimisation algorithm aims to increase its value. In black-box optimisation, no assumption is made on the shape of the objective function and no gradient information is used, which is the case in our registration problems. The objective function that we minimise corresponds to the mean absolute error (MAE):

$$\text{MAE}(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{w \times h} \sum_j^h \sum_i^w \left| \hat{\mathbf{Y}}(i,j) - \mathbf{Y}(i,j) \right| \qquad (2)$$

where $\mathbf{Y}$ is the target, $\hat{\mathbf{Y}}$ is the source, w and h are the number of pixels in $\mathbf{Y}$ and $\hat{\mathbf{Y}}$ along the horizontal and vertical axis respectively. For the purpose of comparison, lower MAE values can typically be interpreted as better predictions: MAE is equal to zero when $\mathbf{Y}$ and $\hat{\mathbf{Y}}$ are strictly identical.

The hand registration problem and the fitting the Ti-6Al-4V matrix in the μ-CT registration problem are solved by minimising Eq. 2 with two well-known black-box optimisation methods that can be used when the derivative of the objective function is unknown. We selected the downhill simplex method (also known as Nelder–Mead method [NM65] and the Conjugate Gradient (CG) method [HS52] due to their popularity. We rely on the *minimize* functions provided by the SciPy Python package. Such methods are however known to be prone to get stuck in local optima and that the final solution may depend on the starting solution used to initialise the algorithm. To address these possible deficiencies we also used three stochastic search methods: pure random search (PRS), Simulated Annealing (SA) and Evolutionary Algorithm (EA). The silicon carbide (SiC) fibres and W cores are registered using a Fly algorithm, a cooperative co-evolution algorithm.

3.2.1. Pure random search

The idea of PRS is simple: i) generate N random sample points uniformly distributed within given boundaries, and ii) select the parameters that yield to the best result. PRS is used as a baseline to compare with other methods. An exhaustive search (or grid search) can also be used as baseline method but it has been discarded due to the high-dimensionality of the registration problems considered here.

3.2.2. Simulated Annealing

With Downhill Simplex, CG and SA, a unique current solution is used at any one time. However, the introduction of an acceptance probability enables the SA algorithm to accept solutions worse than the current one to escape local optima [MRR*53].

Simulated Annealing is based on annealing in metallurgy where a material is heated and cooled down in a controlled environment. The aim is to reduce defects in the treated material. SA mimics this behaviour with:

- A high temperature is first used. The aim is to widely explore the search space.
- The temperature is progressively reduced to restrict the search space, hence refine the results.

To further improve the solution, some implementations include a restart mechanism to repeat the annealing process several times.

3.2.3. Evolutionary Computing

In PRS and EAs, concurrent solutions are used at each iteration of the algorithm [BS93]. An EA relies on the theory of biological evolution proposed by Charles Darwin [Dar59]. A solution is called 'individual'. An individual is defined by a sequence of genes. The performance of each individual is measured using a fitness function (i.e. the objective function to maximise). The set of individuals is called 'population'. Genetic operators (usually selection, crossover, mutation and elitism) are repetitively used on the individuals of the population. Fittest individuals (according to their fitness values) are given a higher probability to reproduce. At each iteration, a new population of offspring is generated from the current population (the parents). At the end of the algorithm, the best individual is extracted. To further improve the solution, some implementations include a restart mechanism that repeats the whole process several times.

3.2.4. Fly algorithm

To optimise the location of cylinders in Figure 4b, a dedicated Evolutionary Algorithm based on the Parisian approach is used: All the individuals of the population collaborate toward a common goal rather than compete [CLRS00]. It is called Fly algorithm [Lou01]. It was initially developed for real-time stereo vision in robotics. The Fly algorithm is implemented as any other EA (i.e. with all the common genetic operators such as selection, crossover and mutation) [GAMV18]. It evolves a population of individuals called *flies*, which correspond to 3D points in the problem space. Each fly is projected in the object space. The type of projection is problem dependent. The fitness function is computed from those projections. In the Parisian approach there are two levels of fitness function:

1. The *local fitness function* evaluates the performance of a given fly. It is used during the selection process. For a fly, improving its local fitness means increasing its chances of survive and reproduce.
2. The *global fitness function* assesses the performance of the whole population. Improving (maximising or minimising depending on the problem considered) the global fitness is the goal of the population.

The local fitness can be implemented as a marginal fitness (F_m) using the global fitness with the leave-one-out cross-validation principle:

$$F_m(i) = \text{MAE}(\mathbf{Y}, \hat{\mathbf{Y}} \setminus \{i\}) - \text{MAE}(\mathbf{Y}, \hat{\mathbf{Y}}) \tag{3}$$

where where $\hat{\mathbf{Y}}$ is the projections estimated from the whole population (e.g. it is the sinogram simulated from the flies in the μ-CT registration), $\hat{\mathbf{Y}} \setminus \{i\}$ is the projections estimated from the population without Fly i. The idea behind the leave-one-out cross-validation is to assess the error metric twice: once with Fly i in the population, and once without it. By comparing the two values (the subtraction in Eq. 3) we can determine if having Fly i is beneficial or not for the population:

- If F_m is positive, the error is smaller when the fly is included: the fly has a positive impact on the population's performance. It is a good fly, i.e. a good candidate for reproduction.
- If F_m is negative, the error is larger when the fly is included: the fly has a negative impact on the population's performance. It is a bad fly, i.e. a good candidate for death.

F_m is therefore a measure maximised by the algorithm whereas $\text{MAE}(\mathbf{Y}, \hat{\mathbf{Y}})$ is minimised.

4. Results

The whole registration framework is cross-platform and relies on open-source technologies. It has been tested on Linux and Microsoft Windows.

This is important to test the Nelder–Mead, Conjugate Gradient, and Simulated Annealing methods several times with different initial solutions, as well as gathering enough data to be statistically meaningful with the stochastic methods. To assess the performance of our registration framework, each optimisation algorithm is therefore run 15 times with different starting solutions. The same starting solution is used for each optimisation algorithms of the same run.

To evaluate the solution given by each algorithm for each run, the prediction and target are compared with MAE. MAE is a distance metrics used in the objective functions.
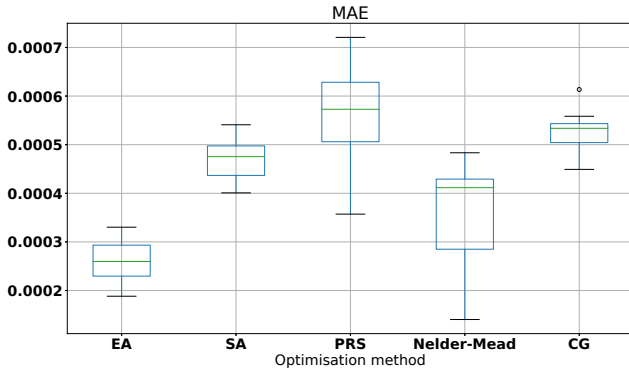
4.1. Hand registration

The hand registration problem is a toy problem to assess the feasibility of our framework. The scenegraph of a 3D hand model is used (see Figure 6). Some of the nodes, including the root node, are arbitrary rotated so that the fingers are now close to each other and the whole hand is tilted. We chose this pose to check the ability of the optimisation algorithms to retrieve subtle changes. The SDD and SOD are set. The corresponding X-ray radiograph is then computed: it is the target image in the registration. The transformations mentioned above, including SDD and SOD, are now 'forgotten'. The aim of the registration is to retrieve all these parameters. For fair comparison, each methods are stopped after the objective function is called 500 times, except for SA where the objective function is called 508 times. In EA, 25 individuals are generated and running for 20 generations.

Figure 7 shows the performance of each method performing on
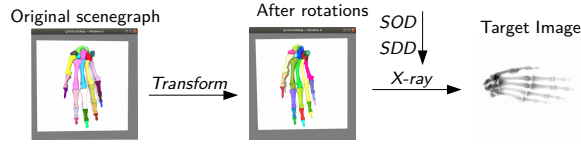
Figure 6: *Original and target hand poses for registration.*

Figure 7: *Overall performance over 15 runs of each method on the hand registration problem. Each optimisation algorithm called the objective function about 500 times.*

Methods	X-ray image	Error map
Nelder-mead MAE = 3.45E-4		
CG MAE = 5.41E-4		
PRS MAE = 5.62E-4		
SA MAE = 3.85E-4		
EA MAE = 3.17E-4		

Figure 8: *Comparing individual result for each method on the hand registration problem (median solution over 15 runs, i.e. solution corresponding to the green bars in Figure 7).*

the hand registration problem over 15 runs. EA outperforms others with lowest median value of distance. It also has good stability, which is at an advantage in medical field. Nelder-Mead provides good results but comes with a higher standard deviation, which means it can fail to find the global optimum when it gets stuck in a local optimum. SA performs close to Nelder-Mead but comes with narrow spread. It might be a better alternative for variation-sensitive problems. Both PRS and CG do not provide very good results in this problem.

To further refine our analysis, the visualisation of the error map between target and prediction (median solution) for each method is provided in Figure 8. The error maps are normalised between 0 and 1, and visualised using the 'thermal' colour look-up table provided by Fiji [SACF*12]. Our aim is to highlight in each figure, independently from one another, where the error is the highest. EA and Nelder-Mead provide the best median results in term of MAE. EA's error map shows that the error is concentrated in the thumb. The fingers are now close to each other as in the target image. For Nelder-Mead, the error is concentrated in the middle and ring fingers. SA provides acceptable results. CG and PRS clearly provide the worse results, which are in line with our expectations.

4.2. X-ray microtomography registration

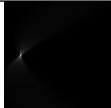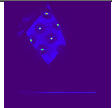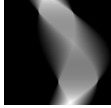The registration is performed into three successive steps:

1. The Ti90-Al6-V4 matrix is first registered using the same method as in Section 4.1.
2. The fibres and cores are then registered using a cooperative co-evolution algorithm (namely the Fly algorithm).
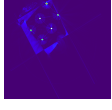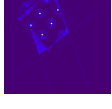3. The registered matrix, fibres and cores are eventually combined.

The experimental data included 900 angles over 180°. To limit the computation time, we reduce it to 90 angles only. To further limit the computation time, the objective function is computed using sinograms rather than CT slices (i.e. no FBP reconstruction required during the registration). However, the accuracy of the final solution is best evaluated using the CT slices reconstructed from the registered sinograms. Some of the code is freely available on GitHub [Vid19]. The results below show that it leads to acceptable results.

Ti90-Al6-V4 matrix: Figure 10 shows the performance of the cube registration with each optimisation algorithm over 15 runs in term of objective function (distance of sinograms) between the reference CT slice (Figure 4b) and the FBP of the sinogram used in the corresponding objective function. Our observations are similar to those made in Section 4.1. CG systematically leads to a large distance between the reference and the predictions. Nelder-Mead can provide relatively good median solution over 15 runs. However the results are excessively dependant on the starting solution used to initialise the algorithm (see larger standard deviation and presence of outliers). In two runs the prediction is totally uncorrelated to the reference CT slice, which indicates that the algorithm was permanently stuck in poor local optima. In this respect, a naive approach such as pure random search can provide a valid alternative to some popular optimisation algorithms. EA outperforms other methods with lowest mean value of distance. To ascertain this analysis, the median result in term of objective function for each optimisation algorithm over 15 runs are visualised in Figure 9. Error maps are also shown. Only Nelder-Mead and Evolutionary Algorithm provided a near perfect match, but again, the results of the Nelder-Mead algorithm can be unreliable, whereas the results of the Evolutionary Algorithm and Simulated Annealing are more consistent.

| Methods	Sinogram	**FBP**	**Error map (FBP)**
CG	MAE = 2474.41		
Nelder-Mead	MAE = 695.75		
PRS	MAE = 1081.54		
SA	MAE = 837.10		
EA	MAE = 739.19		

Figure 9: *Simulated images for each optimisation method (median solution in term of MAE of FBP over the 15 runs). The MAE of sinograms is the corresponding objective function value.*
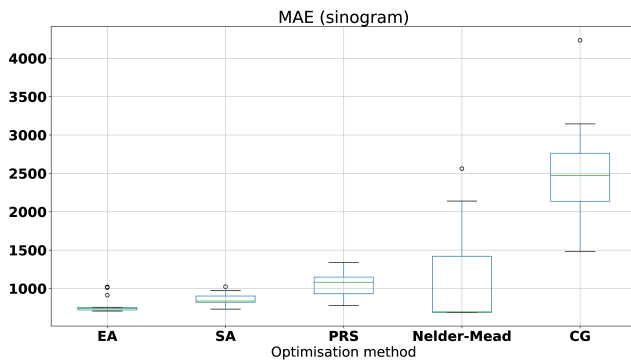
Figure 10: *Performance comparison of the optimisation algorithms over 15 runs for the cube registration.*

SiC fibres and W cores: Once the Ti90-Al6-V4 matrix is registered, the silicon carbide fibres and tungsten cores can be registered using the Fly algorithm. We used 40 individuals, i.e. 40 candidate fibres. 6 fibres and cores are present in the reference CT slice (see Figure 4b). At the end of the optimisation process, the worst fly is removed from the population if its fitness is negative or null. This process is repeated until the fitness of all the remaining flies is strictly positive. The aim is to classify the flies as 'fibre' or 'non-fibre'. Figure 11 shows the predicted CT slice. Table 1 is its corresponding confusion table. All the fibres have been successfully identified, i.e. there are 6 true positives (tp) and 0 false negative (fn). There are 2 false positives (fp) and 32 true negatives (tn). The *accuracy* measures the ratio of
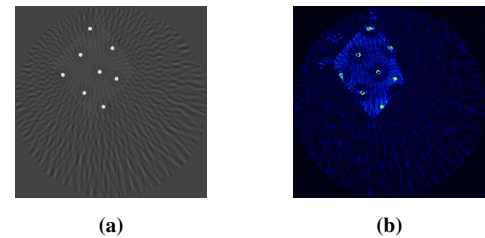
Figure 11: *Predicted CT slice after registration of the fibres.* **(a)** *FBP of the sinogram simulated with the SiC fibres and W cores),* **(b)** *absolute error map between Figures 4b and 11a.*

correctly predicted observations ($tp + tn$) to the total observations ($tp + tn + fp + fn$). It is 0.95 in our case. The *precision* is the

Table 1: *Confusion table corresponding to the optimisation results produced by the Fly algorithm in Figure 11.*

	Fibre (predicted)	Non-Fibre (predicted)
Fibre (actual)	6	0
Non-Fibre (actual)	2	32

ratio of correctly predicted positive observations (tp) to the total predicted positive observations ($tp + fp$). It is 0.75 in our case. *Recall* (also called sensitivity) corresponds to the true positive rate. It is the ratio of correctly predicted positive observations (tp) to all observations in the actual class ($tp + fn$). It is 1 in our case. Similarly the true negative rate is ratio of correctly predicted negative observations (tn) to all observations in the actual class ($tn + fp$). It is 0.94. We also computed the F-score as it is a popular measure for the analysis of binary classification:

$$F = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} = 0.8571 \qquad (4)$$

The five values are close to 1, which demonstrates the validity of our approach.

5. Conclusion

Our research hypothesis was that 3D/2D registration of 3D triangular meshes onto a 2D image can be performed using optimisation and fast X-ray simulation on GPU. We have implemented a registration framework in Python using open-source technologies, namely SciPy (Conjugate Gradient and Nelder-Mead) and our own implementation of pure random search, Simulated Annealing, Evolutionary Algorithm and Fly algorithm for the optimisation, and gVirtualXRay for the X-ray simulation on GPU. The framework has been sucessfully tested using two test-cases, i) hand registration for a rheumatoid arthritis application using a synthetic target image generated with known parameters, and ii) the registration of a cube and cylinders with the CT slice of an actual Ti/SiC fibre composite with tungsten cores. To evaluate our pipeline, each optimisation has been repeated 15 times to gather statistically meaningful results. The outcome shows that:

- Care must be given when choosing an optimisation algorithm as some traditional and popular techniques may lead to unreliable results due to local optima;
- Stochastic algorithms such as Simulated Annealing and Evolutionary Algorithm can outperform Conjugate Gradient and Nelder-Mead in term of accuracy and/or reliability;
- X-ray simulations can be repeatedly computed in objective functions thanks to the computational power provided by modern GPUs.

These preliminary results demonstrated the validity of our approach as the registrations were performed successfully. However, some limitations are present in our work. For example, it is assumed that all hand skeletons have the same proportion. There is no skin model in the hand registration problem and a synthetic radiograph was used. The Ti/SiC fibre composite registration was performed using a single CT slice only. Once the location of each cylinder is roughly known, the centre positions could be refined. Care will be given to the parameters of the optimisation algorithms to provide the best possible performance in term of computation time and accuracy of the solutions. In future all these deficiencies will be addressed and a more thorough and quantitative evaluation will be conducted, e.g. using real clinical X-ray radiographs for the hand registration.

Acronyms

CBCT cone beam computed tomography.
CG Conjugate Gradient.
CT computed tomography.
DMARD disease modifying antirheumatic drug.
EA Evolutionary Algorithm.
ESRF European Synchrotron Radiation Facility.
FBP filtered back-projection.
GLSL OpenGL Shading Language.
GPU graphics processing unit.
HU Hounsfield Unit.
MAE mean absolute error.
MC Monte Carlo.
MRI magnetic resonance imaging.
μ-CT X-ray microtomography.
PA postero-anterior.
PRS pure random search.
RA rheumatoid arthritis.
ROI region of interest.
SA Simulated Annealing.
SDD source-detector distance.
SiC silicon carbide.
SNR signal-to-noise ratio.
SOD source-object distance.
Ti titanium.
Ti/SiC titanium/silicon carbide.
W tungsten.
DRR digitally reconstructed radiograph.

References

[A*03] AGOSTINELLI S., ET AL.: Geant4–a simulation toolkit. *Nucl Instrum Methods Phys Res A 506*, 3 (2003), 250 – 303. doi:10.1016/S0168-9002(03)01368-8. 3

[AMV19] AL-MALIKI S. F., VIDAL F. P.: Examples of stochastic optimisation algorithms, 2019. [Online; accessed 08/08/2019]. URL: https://github.com/Shatha1978/Optimisation-algorithm-examples. 4

[B*04] BIELAJEW A., ET AL.: *History, overview and recent improvements of EGS4*. Tech. rep., Stanford Linear Accelerator Center (SLAC), 2004. doi:10.2172/839781. 3

[B*12] BERT J., ET AL.: Hybrid GATE: A GPU/CPU implementation for imaging and therapy applications. In *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)* (Oct. 2012), pp. 2247–2250. doi:10.1109/NSSMIC.2012.6551511. 4

[BS93] BÄCK T., SCHWEFEL H.: An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation 1*, 1 (March 1993), 1–23. doi:10.1162/evco.1993.1.1.1. 5

[BSFVS95] BARÓ J., SEMPAU J., FERNÁNDEZ-VAREA J. M., SALVAT F.: PENELOPE: An algorithm for Monte Carlo simulation of the penetration and energy loss of electrons and positrons in matter. *Nucl Instrum Methods Phys Res B 100*, 1 (1995), 31–46. doi:10.1016/0168-583X(95)00349-5. 3

[BTKB*13] BEN TEKAYA I., KAFTANDJIAN V., BUYENS F., SEVESTRE S., LEGOUPIL S.: Registration-based geometric calibration of industrial X-ray tomography system. *IEEE Transactions on Nuclear Science 60*, 5 (2013), 3937–3944. doi:10.1109/TNS.2013.2279675. 2

[BWB*03] BIRKFELLNER W., WIRTH J., BURGSTALLER W., BAUMANN B., STAEDELE H., HAMMER B., GELLRICH N. C., JACOB A. L., REGAZZONI P., MESSMER P.: A faster method for 3D/2D medical image registration—a simulation study. *Physics in Medicine and Biology 48*, 16 (jul 2003), 2665–2679. doi:10.1088/0031-9155/48/16/307. 2

[CLRS00] COLLET P., LUTTON E., RAYNAL F., SCHOENAUER M.: Polar IFS+Parisian Genetic Programming=Efficient IFS Inverse Problem Solving. *Genetic Programming and Evolvable Machines 1*, 4 (Oct. 2000), 339–361. doi:10.1023/A:1010065123132. 5

[Dar59] DARWIN C.: *On the origin of species by means of natural selection*. John Murray (London), 1859. 5

[DCY06] DELIGIANNI F., CHUNG A. J., YANG G.: Nonrigid 2-D/3-D registration for patient specific bronchoscopy simulation with statistical shape modeling: Phantom validation. *IEEE Transactions on Medical Imaging 25*, 11 (Nov 2006), 1462–1471. doi:10.1109/TMI.2006.883452. 2

[FDLB06] FREUD N., DUVAUCHELLE P., LÉTANG J. M., BABOT D.: Fast and robust ray casting algorithms for virtual X-ray imaging. *Nucl Instrum Methods Phys Res B 248*, 1 (2006), 175–180. doi:10.1016/j.nimb.2006.03.009. 3, 4

[GAMV18] GRAY C., AL-MALIKI S., VIDAL F.: Data exploration in evolutionary reconstruction of pet images. *Genetic Programming and Evolvable Machines 19*, 3 (Sept. 2018), 391–419. doi:10.1007/s10710-018-9330-7. 5

[HS52] HESTENES M. R., STIEFEL E.: Methods of Conjugate Gradients for Solving Linear Systems. *Journal of Research of the National Bureau of Standards 49*, 6 (Dec. 1952), 409–436. doi:10.6028/jres.049.044. 4

[JOP*01] JONES E., OLIPHANT T., PETERSON P., ET AL.: SciPy: Open source scientific tools for Python, 2001. [Online; accessed 08/08/2019]. URL: http://www.scipy.org/. 4

[JZJ14] JIA X., ZIEGENHEIN P., JIANG S. B.: GPU-based high-performance computing for radiation therapy. *Phys Med Biol 59*, 4 (2014), R151–R182. doi:10.1088/0031-9155/59/4/R151. 4

[Lou01] LOUCHET J.: Using an individual evolution strategy for stereovision. *Genetic Programming and Evolvable Machines 2*, 2 (June 2001), 101–109. doi:10.1023/A:1011544128842. 5

[MRR*53] METROPOLIS N., ROSENBLUTH A. W., ROSENBLUTH M. N., TELLER A. H., TELLER E.: Equation of state calculations by fast computing machines. *The Journal of Chemical Physics 21*, 6 (jun 1953), 1087–1092. doi:10.1063/1.1699114. 5

[MSKK18] MAIER J., SAWALL S., KNAUP M., KACHELRIESS M.: Deep scatter estimation (DSE): Accurate real-time scatter estimation for x-ray CT using a deep convolutional neural network. *Journal of Nondestructive Evaluation 37*, 3 (July 2018). doi:10.1007/s10921-018-0507-z. 2

[MTLP12] MARKELJ P., TOMAŽEVIČ D., LIKAR B., PERNUŠ F.: A review of 3D/2D registration methods for image-guided interventions. *Medical Image Analysis 16*, 3 (2012), 642 – 661. Computer Assisted Interventions. doi:10.1016/j.media.2010.03.005. 2

[MW14] MAIRE E., WITHERS P. J.: Quantitative X-ray tomography. *International Materials Reviews 59*, 1 (2014), 1–43. doi:10.1179/1743280413y.0000000023. 2

[NM65] NELDER J. A., MEAD R.: A Simplex Method for Function Minimization. *The Computer Journal 7*, 4 (Jan. 1965), 308–313. doi:10.1093/comjnl/7.4.308. 4

[SACF*12] SCHINDELIN J., ARGANDA-CARRERAS I., FRISE E., ET AL.: Fiji: an open-source platform for biological-image analysis. *Nature methods 9*, 7 (2012), 676–682. doi:10.1038/nmeth.2019. 6

[SBB*16] SMOLEN J. S., BREEDVELD F. C., BURMESTER G. R., BYKERK V., DOUGADOS M., EMERY P., KVIEN T. K., NAVARRO-COMPÁN M. V., OLIVER S., SCHOELS M., ET AL.: Treating rheumatoid arthritis to target: 2014 update of the recommendations of an international task force. *Annals of the rheumatic diseases 75*, 1 (2016), 3–15. doi:10.1136/annrheumdis-2015-207524. 2

[SMV*17] SUJAR A., MEULEMAN A., VILLARD P.-F., GARCÍA M., VIDAL F. P.: gVirtualXRay: Virtual X-Ray Imaging Library on GPU. In *Computer Graphics and Visual Computing (CGVC)* (2017), The Eurographics Association. doi:10.2312/cgvc.20171279. 4

[SSG18] SHIM J.-H., STAVRE Z., GRAVALLESE E. M.: Bone loss in rheumatoid arthritis: basic mechanisms and clinical implications. *Calcified tissue international 102*, 5 (2018), 533–546. doi:10.1007/s00223-017-0373-1. 2

[VGF*09] VIDAL F. P., GARNIER M., FREUD N., LÉTANG J. M., JOHN N. W.: Simulation of X-ray attenuation on the GPU. In *Proceedings of Theory and Practice of Computer Graphics 2009* (Cardiff, UK, June 2009), Eurographics Association, pp. 25–32. doi:10.2312/LocalChapterEvents/TPCG/TPCG09/025-032. 4

[VGF*10] VIDAL F. P., GARNIER M., FREUD N., LÉTANG J. M., JOHN N. W.: Accelerated deterministic simulation of X-ray attenuation using graphics hardware. In *Eurographics 2010 - Posters* (May 2010), Eurographics Association, p. Poster 5011. doi:10.2312/egp.20101026. 4

[Vid13] VIDAL F. P.: gVirtualXRay: Virtual X-Ray Imaging Library on GPU, 2013. [Online; accessed 08/08/2019]. URL: https://sourceforge.net/projects/gvirtualxray/. 4

[Vid19] VIDAL F. P.: Registration of polygon meshes onto X-ray projections: Application to 3D micro-tomography by synchrotron radiation, 2019. [Online; accessed 08/08/2019]. URL: https://github.com/effepivi/MicroTomoRegistration. 6

[VLPC05] VIDAL F., LÉTANG J., PEIX G., CLOETENS P.: 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*, 3 (2005), 333–348. doi:http://dx.doi.org/10.1016/j.nimb.2005.02.003. 2, 3

[vPT*05] VAN DE KRAATS E. B., PENNEY G. P., TOMAZEVIC D., VAN WALSUM T., NIESSEN W. J.: Standardized evaluation methodology for 2-d-3-d registration. *IEEE Transactions on Medical Imaging 24*, 9 (Sept. 2005), 1177–1189. doi:10.1109/TMI.2005.853240. 2

[VV15] VIDAL F. P., VILLARD P.-F.: Simulated motion artefact in computed tomography. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2015), The Eurographics Association, pp. 213–214. doi:10.2312/vcbm.20151228. 4

[VV16] VIDAL F. P., VILLARD P.-F.: Development and validation of real-time simulation of X-ray imaging with respiratory motion. *Computerized Medical Imaging and Graphics 49* (Apr. 2016), 1–15. doi:10.1016/j.compmedimag.2015.12.002. 4