3D-2D Registration using X-ray Simulation and CMA-ES

Tianci Wen[1][0000−0002−3524−3569], Radu P. Mihail[2][0000−0002−3682−3485], and Franck P. Vidal[1][0000−0002−2768−4524]

[1] School of Computer Science and Electronic Engineering, Bangor University, UK
{t.wen,f.vidal}@bangor.ac.uk
[2] Department of Computer Science,Valdosta State University, Georgia, USA
rpmihail@valdosta.edu

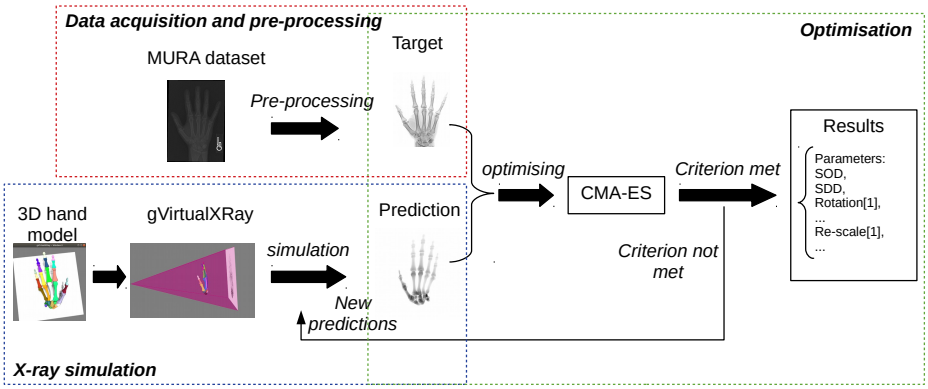

Fig. 1: Registration pipeline based on X-ray simulation and CMA-ES.

Abstract. Radiographs of the hand are useful in diagnosing and staging diseases such as rheumatoid arthritis (RA) and other musculoskeletal diseases. Radiographs are projections of the 3D anatomy, with the useful information such as pose and pathology becoming lost in the process. We propose a 3D hand pose recovery method for radiographs of hands using a novel hybrid image registration method. Our pose recovery pipeline consists of aligning a simulated X-ray (digitally reconstructed radiograph) of an articulated phantom mesh model to a real hand radiograph using Covariance Matrix Adaptation Evolution Strategy. Early results demonstrate that our approach works well. Further inquiry is required to evaluate the applicability of our registration approach to other articulated musculoskeletal anatomy.

Keywords: 2D–3D registration · X-ray simulation · artificial evolution · evolutionary computing · CMA-ES · DRRs.

1 Introduction

Computational methods of disease tracking and progression prediction based on the analysis of medical imagery is receiving heightened attention in recent years.

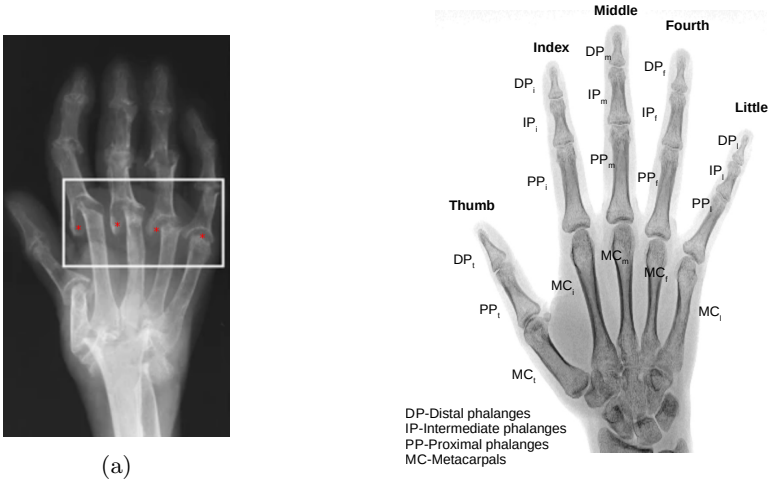

(a) (b)

Fig. 2: a) Erosions induced by RA inflammatory processes visible around the red asterisks. Image source [19]. b) Naming of hand fingers and bones. Source: The MURA dataset.

Chronic diseases of the human musculoskeletal system caused by autoimmune processes lead to progressive, irreversible anatomical changes over time. In the case of rheumatoid arthritis (RA), a chronic inflammatory disorder with largely unknown pathogenesis, patients often present to the clinician with swelling of the hands. If left untreated, the disease progresses in distinct stages, from joint pain, swelling, stiffness to cartilage loss, bone erosion, deformities and total loss of joint function [15].

Plain radiographic imaging (X-rays) of the hands is done routinely for diagnostic and tracking purposes, as routine care for RA patients. Since hand radiographs are relatively inexpensive and low-risk, they provide clinicians with baselines, and visible changes over time. The rate of disease progression is modulated by treatment and lifestyle choices, but distinct deformations have been documented [17,19]. Typical deformities include boutonnière, swan-neck, hitchhiker's thumb and claw toe, Other, less obvious changes include bone erosions, induced by the inflammation of the synovial membrane, as shown in Figure 2a.

Radiographs are projections of 3D structures, hence much information is lost. The anatomy of the hand varies among different individuals, e.g. the ratio of the lengths of the long bones is not always consistent. Using this observation and the ability to speedily create digitally reconstructed radiographs (DRRs), we propose a method to register a 3D mesh model of a hand to Posterior Anterior (PA) view hand radiograph.

Anomaly detection in hand radiographs is important for disease staging and monitoring. Our registration method is a pre-processing step for algorithms

that modify the mesh model using domain-specific knowledge to better track disease-induced changes without expensive volumetric scans that clinicians may not be equipped with or are cost-prohibitive. Time-series, patient-specific information regarding the progression of a disease is critical for treatment planning and drug effectiveness monitoring.

Our main contribution consists of a novel registration method of a highly articulated and anatomically correct 3D mesh model, hereby referred to as the phantom model, to a real radiograph using a DRR software (`https://sourceforge.net/p/gvirtualxray/`) [21,22]. Our proposed optimisation process (inside green box in Figure 1) consists of using an evolutionary algorithm to solve for the articulated 3D pose of the virtual hand that best fits the real radiograph. Similar works on 3D/2D registration using Evolutionary Algorithms (EAs) are presented by Gomez et al. [9,8]. We use Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [12] from Pymoo [3] as the numerical optimisation algorithm. This paper is based on previous work where synthetic data were used [25]. The paper is organised as follows: Section 2 provides an overview of related work, Section 3 describes our method, including the data used and data pre-processing, Section 4 describes our results as well as quantitative and qualitative evaluation, Section 5 provides conclusions and describes future work.

2 Background

Medical image registration is important for pathology modeling, treatment and surgery planning, diagnosis and prognosis, among other tasks. Work in this area has focused on all imaging modalities, including projective (X-rays) and volumetric, either from computed tomography (CT) or magnetic resonance imaging (MRI), and a combination 2D to 3D and 3D to 3D registration. The medical image literature is vast, and we refer the interest reader to several review papers [23,24,1].

Registration of DRRs of articulated objects to real radiographs is scarcely investigated. Related work in this area falls into several categories: 2D radiograph to 3D volumetric scan registration [7,5], biplanar radiograph to model registration [14,4,2].

Kanhonou et al. [14] propose a method for automatic registration of phantom tibia and femur to biplanar radiographs for pathology detection using rigid transformations as the optimisation variables. Englander et al. [4] propose a vision-based method for in-vivo registration of phantom tibia and femur models to high-speed biplanar radiographs to study Anterior Cruciate Ligament (ACL) length and strain during dynamic activity. Aubert et al. [2] propose a spine reconstruction method using deep convolutional neural networks (CNNs) from biplanar radiographs.

Our work fits in a hybrid paradigm: the registration of a DRR from a phantom model (3D triangular mesh model) using simulated radiography to a real radiograph (a single 2D image). Our work aligns most closely with that of

Gong et al. [10] who register phantom models in DRRs to fluoroscopy imagery for treatment planning in complex bone fractures. In contrast with [10] where the registration involves a rigid body transform, we register a highly articulated mesh model of a hand.

We also propose an automatic method for the registration of a hand phantom to a single radiograph using rigid transformations (translation, rotation and scaling) as the optimisation variables. Due to the repetitive use of DRR calculations by the optimisation algorithm, an extremely fast implementation is needed. DRR calculation codes often rely on CT scans data as input. They can be implemented using graphics processing unit (GPU) programming to speedup computations [6]. Polygon meshes can be converted into voxel data; this is called voxelisation [13]. However, as the hand model is deformed by the optimisation algorithm for each DRR, this approach will be far too costly and can actually be avoided. The alternative is to use a fast X-ray simulation model that can support polygon meshes [16]. It can be efficiently implemented on GPU using real-time Computer Graphics (CG) techniques. This approach is the most suitable one in our application context as real-time CG techniques are designed for polygon meshes. The Virtual X-ray Imaging Library on GPU (gVirtualXRay) provides an open-source implementation that can run on laptops, desktop computers, and large supercomputers. gVirtualXRay is a C++ library to simulate X-ray imaging [22]. It also provides a Python 3 wrapper that we used to prototype our framework in this study.

3 Methodology

3.1 The MURA Dataset

The MURA dataset, freely available on GitHub at `https://stanfordmlgroup.github.io/competitions/mura/`, contains 40,561 musculoskeletal radiographs from 14,863 clinical studies of 12,173 different patients [18]. It focus on upper extremities of human body, including elbow, finger, forearm, hand, humerus, shoulders and wrist. Each radiograph is manually labelled as normal or abnormal. 15 radiographs (Figure 3) are selected to test the performance of our approach (Section 4.1), which only require a single radiograph.

Unlike typical medical images such as Digital Imaging and Communications in Medicine (DICOM) [11], radiographs in the MURA dataset are in Portable Network Graphics (PNG) format which a lot of information is missing (Table 1). To adapt data to use in our experiments, those radiographs are pre-processed to improve the structural difference between the object and other areas. More discussion about image pre-processing in Section 3.2.

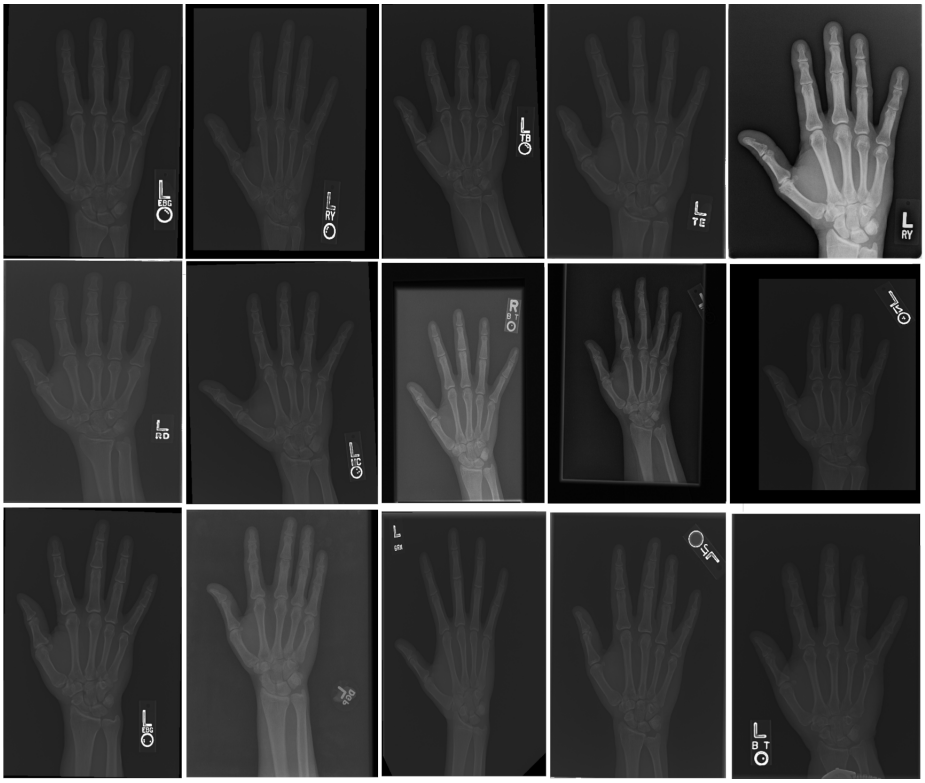

Fig. 3: 15 selected hand radiographs from MURA dataset. Images are numbered from image 1 (top left) to image 15 (bottom right). Top row: image 1-5, middle row: image 6-10, and bottom row: image 11-15.

3.2 Image pre-processing

Although the MURA dataset was created for the purpose of abnormality detection, it is also useful for testing the performance of our registration framework. We focus on hand radiographs with PA views. We manually pre-processed 15 different radiographs as follow: i) each radiograph is cropped so that only the hand part remains, ii) the left or right marker is removed and the area surrounding the hand is "cleaned" (to have same pixel values), iii) the skin around finger is removed as much as possible, iv) the radiographs, which are negative images, are inverted to match the positive images generated by X-ray simulation, and v) all images are re-scaled to the same size of simulated X-ray images. Corresponding images are shown in Figure 4. In each image, we define the name of each finger (from left to right): thumb, index finger, middle finger, the fourth finger and little finger, as shown in Figure 2b.

Table 1: Information needed to simulate X-ray radiographs that are present in typical medical imaging file formats such as DICOM but missing in the PNG files from the MURA dataset.

Properties	**PNG**	**DICOM**
X-ray tube voltage	plausible energy beam	known
Quantisation	8-bit	16-bit
source-to-object distance (SOD)	need to estimate	known
source-to-detector distance (SDD)	need to estimate	known
Pixel spacing (mm)	need to estimate	known
Object location	need to estimate	known
Object orientation	need to estimate	known

3.3 gVirtualXRay

In order to simulate an X-ray image (inside blue box of Figure 1), some input parameters are needed:

- An incident energy beam,
- A 3D object to scan, including its geometry, position, and orientation.
- A virtual X-ray detector, including its pixel resolution, pixel spacing, and orientation.
- The source-to-object distance (SOD),
- The source-to-detector distance (SDD).

All these parameters must be set before a DRR can be generated. Some of them can be set once for all, such as the incident energy beam, the detector's resolution and orientation, and pixel spacing. As mentioned above, the actual values are unknown due to the use of PNG files instead of DICOM files in the MURA dataset. We used plausible values. The registration consists in tuning all the other parameters.

3.4 Optimisation

CMA-ES [12] is a widely used optimisation algorithm which provides a great baseline result. Other optimisation algorithms could be used for further evaluations of our approaches such as multi-objective optimisation algorithms [20]. CMA-ES is a special evolution strategy with adaption of covariance matrix. It is used to solve complex problems that require derivative-free optimisation. An evolution strategy is inspired by biological evolution. The idea is: initialising individuals (a set of solutions), recombination and mutation is used to create new individuals, best individuals are then selected based on their fitness value to become the parents of next generation of individuals. This process is repeated until satisfactory results (set termination criterion) are found. In Evolution Strategies (ES), new individuals are created by sampling from the probability distribution. In CMA-ES, however, sampling

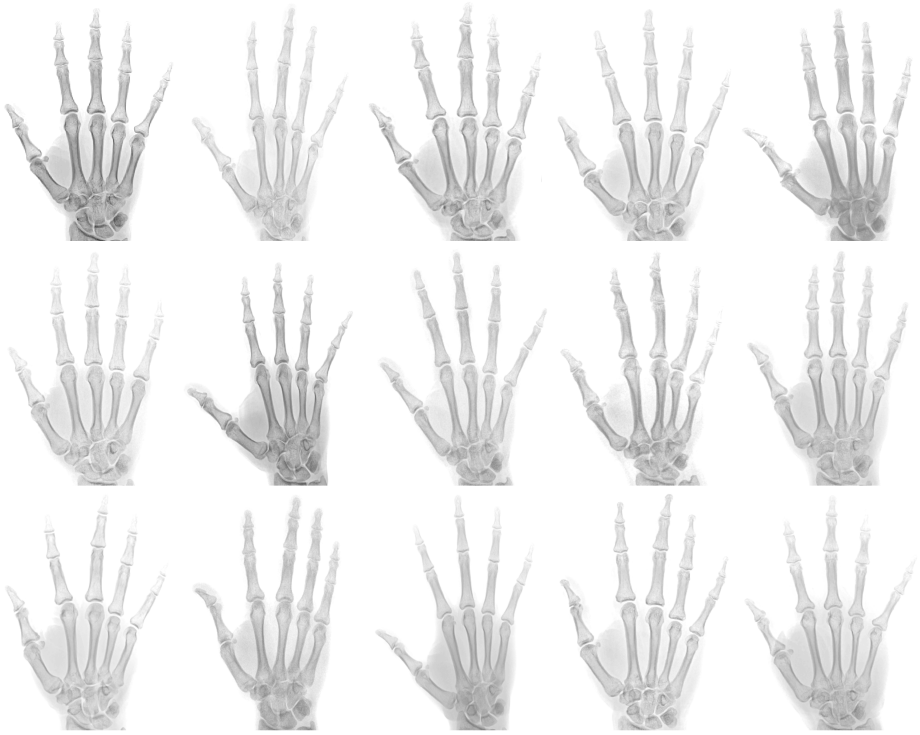

Fig. 4: 15 selected hand radiographs from the MURA dataset after pre-processing. Images are numbered from Image 1 (top left) to Image 15 (bottom right). Top row: Images 1-5, middle row: Images 6-10, and bottom row: Images 11-15.

is achieved through the use of a covariance matrix of the distribution. This gives CMA-ES great advantages in the ill-conditioned problems which small changes of input variables result in large change of output. The optimisation process is shown in Figure 1 (green box).

Before thoptimisation, all images are normalised to have zero-mean and unit-variance. This is to prevent that some features becoming too dominant during optimisation while other features would be less relevant. We use Mean absolute error (MAE) to construct the objective function, i.e. as the fitness function to be minimised by CMA-ES. MAE is the sum of absolute errors between samples and then divided by total number of samples (Eq. 1). The best number can be achieved is zero. Typically, lower MAE value indicates better optimising result.

$$\mathrm{MAE}(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{w \times h} \sum_{j}^{h} \sum_{i}^{w} \left| \hat{\mathbf{Y}}(i,j) - \mathbf{Y}(i,j) \right| \tag{1}$$

Table 2: Rotation and re-scaling parameters to be optimised and corresponding ranges.

Bones / Parameters	**Whole hand**	**Thumb**	**Index**	**Middle**	**Fourth**	**Little**
Rotation range (degrees)	[-20, 20]	MC_t: $[-10, 10]$ $[-20, 0]$ PP_t: $[-10, 10]$	PP_i: $[-10, 10]$ $[-20, 0]$ IP_i: $[-20, 0]$ DP_i: $[-20, 0]$	PP_m: $[-10, 10]$ $[-20, 0]$ IP_m: $[-20, 0]$ DP_m: $[-20, 0]$	PP_f: $[-10, 10]$ $[-20, 0]$ IP_f: $[-20, 0]$ DP_f: $[-20, 0]$	PP_l: $[-10, 10]$ $[-20, 0]$ IP_l: $[-20, 0]$ DP_l: $[-20, 0]$
Rescaling ratio	-	$[0.9, 1.1]$	$[0.9, 1.1]$	$[0.9, 1.1]$	$[0.9, 1.1]$	$[0.9, 1.1]$

where $\mathbf{Y}$ is the target image, $\hat{\mathbf{Y}}$ is the predicted image, w and h are width and height of target and predicted images, respectively.

Our registration problem is considered to be complex both in terms of number of parameters and the corresponding data range. There are 38 parameters that need to be optimised including 2 distance parameters: SOD and SDD. SOD is a ratio of SDD with a value between 0.7 and 0.95. By using ratios, we make sure that the distance between the source and the object, is always less than the distance between the source and the detector. SDD ranges between 10 and 1000 centimetres. There are 22 rotating angles and 14 rescaling factors, which are shown in Table 2. The rotation range is determined based on the modelling of rotations of the real hand except the whole hand, which is determined by a priori knowledge of the PA pose but adding some degrees of complexity. There are no constraint handlers implemented in our framework since our work is still at preliminary stage. Further improvements of registration results would certainly involves modelling restrictions among parameters.

Table 3: Registration results for 15 different target images along with corresponding metric values.

Image number & Metrics	**Target**	**Prediction**	**Error map**	**Number of objective function calls**
1 & MAE=0.3937 ZNCC=0.7060	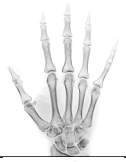	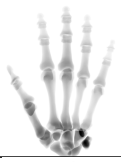	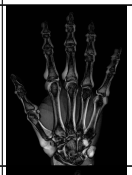	2050
2 & MAE=0.3497 ZNCC=0.7378	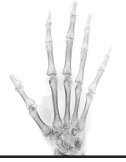	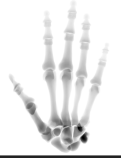	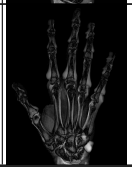	1964

Continued on next page

Table 3 – continued from previous page

Image number & Metrics	Target	Prediction	Error map	Number of objective function calls
3 & MAE=0.4060 ZNCC=0.7264	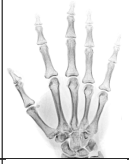	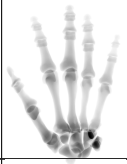	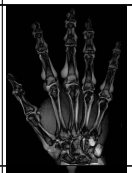	1934
4 & MAE=0.3953 ZNCC=0.7051	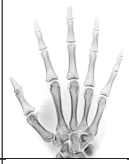	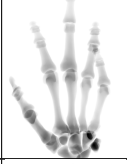	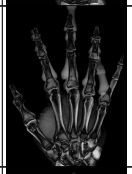	2056
5 & MAE=0.4398 ZNCC=0.6603	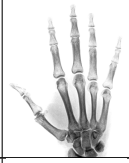	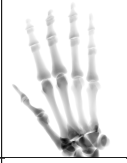	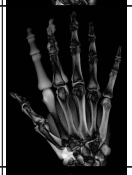	2050
6 & MAE=0.3947 ZNCC=0.7179	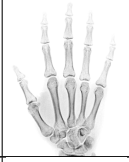	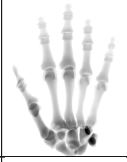	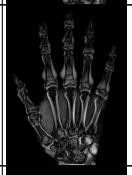	2056
7 & MAE=0.4029 ZNCC=0.7012	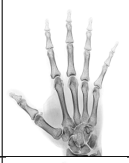	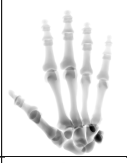	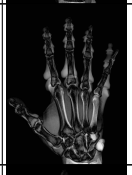	2055
8 & MAE=0.3765 ZNCC=0.7202	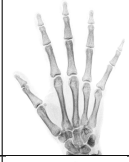	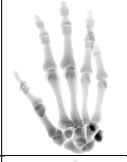	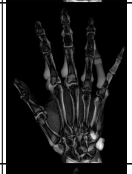	1886
9 & MAE=0.4356 ZNCC=0.6130	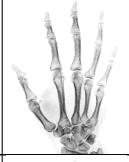	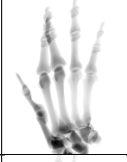	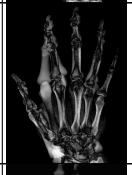	2119
10 & MAE=0.4790 ZNCC=0.6503	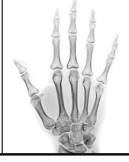	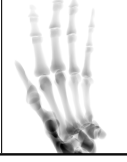	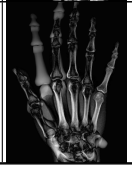	2308

Continued on next page

Table 3 – continued from previous page

Image number & Metrics	Target	Prediction	Error map	Number of objective function calls
11 & MAE=0.4167 ZNCC=0.7094	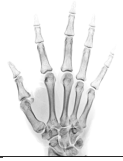		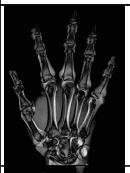	2186
12 & MAE=0.4073 ZNCC=0.7156	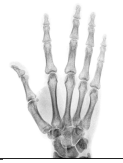		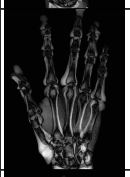	1817
13 & MAE=0.3903 ZNCC=0.7218	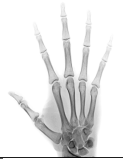		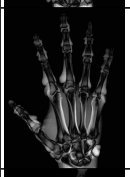	2016
14 & MAE=0.4464 ZNCC=0.6535	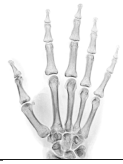	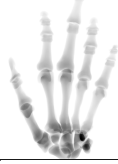	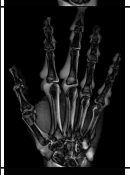	1898
15 & MAE=0.4206 ZNCC=0.7409	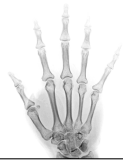		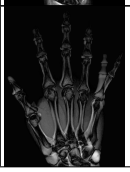	2027

4 Results

There are two ways to assess the effectiveness of our method in solving the registration problem. In any case, several runs must be performed to gather statistically meaningful data. In Section 4.1, we selected 15 different radiographs and tested our method once on each of the radiographs. The emphasis is on **data and simulation variability: For different input images, does the algorithm always provide outputs of similar quality?** In Section 4.2, we selected the images of the worse, median and best registrations of Section 4.1. The registration is then repeated 15 times for these three images. The emphasis is on **optimisation algorithm variability: For a given input image, does the algorithm always provide a similar output?** We also aim to determine if some images harder to register than others.

4.1 Data and simulation variability

Here, we aim to determine if the algorithm always provide outputs of similar quality on different input images. 15 registrations using 15 different real X-ray images were performed, i.e. one registration per image, due to computational demand (about 4 hours per registration on a single Intel Core i5-8400 (2.80GHz) central processing unit (CPU) and a single NVIDIA GeForce GTX 1070 Ti GPU). The 15 pre-processed images that we used are shown in Figure 4. MAE is used to compare target and predicted images during registration because it is relatively faster to compute. Zero mean normalised cross correlation (ZNCC) is used for visual analysis of the predicted images after registration. It is a measurement of similarity between two images. Since it is hard to interpret the value of MAE, ZNCC is very helpful to analyse the performance of the registrations.

To compute ZNCC, the target and predicted images are normalised first, which is subtracting all pixels by the mean value and divided by standard deviation. Normalised target and predicted images are then multiplied. Finally, all values are added and divided by total number of pixels (see Eq. 2). ZNCC primarily concentrates on template matching and completely different images might have very high scores.

$$\mathrm{ZNCC}\left(\hat{\mathbf{Y}}, \mathbf{Y}\right) = \frac{1}{w \times h} \sum_{j}^{h} \sum_{i}^{w} \frac{\left(\hat{\mathbf{Y}} - \overline{\hat{\mathbf{Y}}}\right)\left(\mathbf{Y} - \overline{\mathbf{Y}}\right)}{\sigma_{\hat{\mathbf{Y}}} \sigma_{\mathbf{Y}}} \tag{2}$$

where $\overline{\mathbf{Y}}$ and $\overline{\hat{\mathbf{Y}}}$ are mean pixel value of target and predicted images, and $\sigma_{\mathbf{Y}}$ and $\sigma_{\hat{\mathbf{Y}}}$ are standard deviation of the pixel values in target and predicted images.

ZNCC ranges from -1 to 1, where i) the value is close to 1, the two images are highly similar which implies high level of correlation, ii) the value is 0, two images are extremely different which implies there is no correlation, iii) the value is -1, one image is the negative of the other image which implies they are anti-correlated or inversely correlated.

Table 3 lists results from 15 registrations. By looking at predictions and associated error maps, there are 6 registrations that successfully recovered all 5 fingers, where ZNCC is all above 0.7. There are 6 registrations that successfully recovered 4 fingers, where ZNCC is all above 0.7 except image 14. There are 3 registrations that recovered 1 finger, where ZNCC is all below 0.7. It is clear from Table 3 that Images 5, 9, 10 and 14 are visually worse than the other images. This trend is not necessarily visible in Figure 5a (bar chart of the MAE for each registration). However, Figure 5b clearly show two groups: Images 5, 9, 10 and 14 exhibit a significantly lower ZNCC than the other images.

Results demonstrate that our approaches perform well. However, there are some problems that need to be addressed in future researches:

1. Some fingers are not within images. For example, Registration 4, 12 and 14 have the middle finger extended outside the image space. Registration 10 has middle finger and fourth finger extended outside the image space.

2. In 3 registrations (5, 9 and 10) only the thumb is recovered, the middle finger is matched with the target's index finger, the fourth finger is matched with the target's middle finger, and the little finger is matched with the target's fourth finger.
3. A finger is overlapped with another finger. In Registration 4 and 8, the little finger is overlapped with the fourth finger. In Registration 9, the index finger is overlapped with middle finger.

4.2 Optimisation algorithm variability

Here, we aim to determine if the algorithm always provides outputs of similar quality on the same input image. We selected the best, median and worse results from the previous subsection, i.e. Images 2, 3, and 10. We perform another 14 registrations on each image and included their previous results (i.e. a total of 15 results per image) to test the variability of the algorithm, CMA-ES, when the input data is the same. Then we compute the mean and standard deviations (STDEVs) of MAE, ZNCC and number of calls to objective function over the 15 runs. The data is summarised in Table 4. It shows that CMA-ES provides registrations of consistent quality, both in terms of MAE and ZNCC, for Images 2 and 3 (low standard deviations). However, the standard deviations are much higher for Image 10 (the worse registration of Section 4.1). The MAE is higher than for the other two images, and the ZNCC lower. It indicates that, somehow, Image 10 is a lot harder to register than Images 2 and 3. The scatter plot in Figure 6 shows the MAE (circles) and ZNCC (triangles) plotted as a function of the number of generations. Green and purple marks are aligned and form horizontal lines: CMA-ES produces consistent registrations for Images 2 and 3. Blue marks are scattered over the plot: CMA-ES does not produce consistent registrations for Image 10.

CMA-ES can produce registrations of good quality consistently for some images. For other images, CMA-ES may fail. Further research is needed to comprehend what makes Image 10 hard to register compared to Images 2 and 3 as the three images are visually similar.

Table 4: Results for 15 registrations on Images 2, 3 and 10.

Image number	**MAE (mean ± STDEV)**	**ZNCC (mean ± STDEV)**	**Number of objective function calls (mean ± STDEV)**
2 (best run in Table 3)	0.3460 ± 0.0055	0.7426 ± 0.0074	2105 ± 189
3 (median run in Table 3)	0.4076 ± 0.0019	0.7227 ± 0.0035	1954 ± 132
10 (worse run in Table 3)	0.4701 ± 0.0218	0.6597 ± 0.0184	2137 ± 162

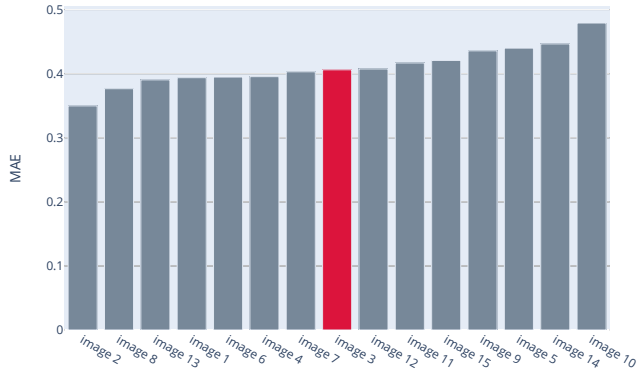

(a) MAE of each registration

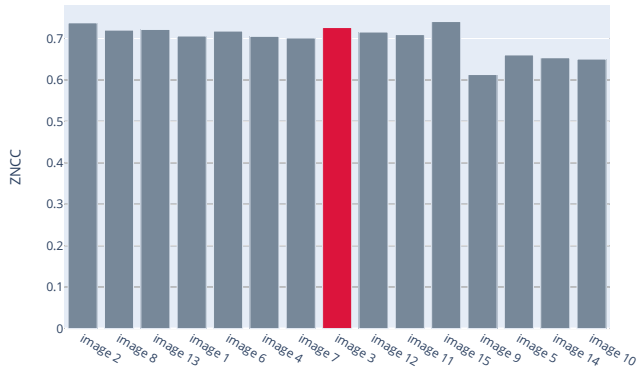

(b) ZNCC of each registration

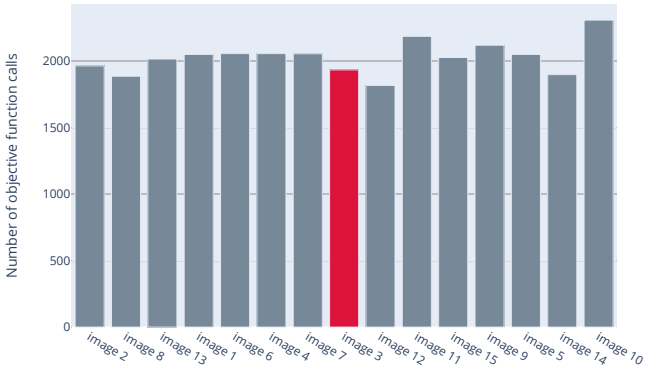

(c) Number of objective function calls of each registration

Fig. 5: Bar charts for quantitative results shown in Table 3. All data is sorted on MAE and the median result is highlighted in red.

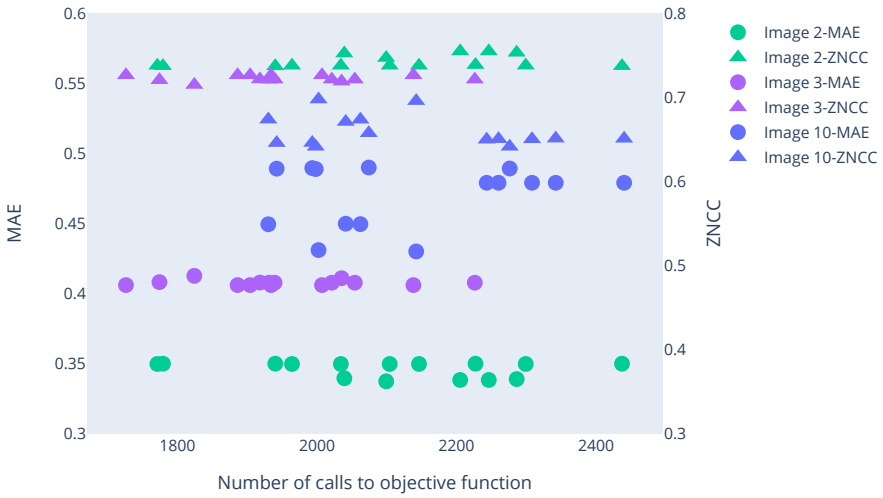

Fig. 6: Scatter plots of results for 15 registrations on Image 2, 3 and 10.

5 Conclusions and Future work

We have shown the feasibility of using DRRs to register an articulated phantom model to real hand radiographs without visible pathology. Our registration framework heavily relies on numerical optimisation. We used CMA-ES, a popular evolutionary algorithm for non-linear or non-convex continuous optimisation problems. Performing the registration on different X-ray images showed that results were not always of the same quality. However, CMA-ES produces similar results with low variability during 15 runs on the same image. This demonstrate that the stochastic nature of CMA-ES is not a concern in our case. It also indicates that some images are harder than others to register.

In the future, we will use this work as an initialisation step in a pipeline that tweaks the geometry of a phantom model to match the pathology seen in real radiographs, such as those afflicted by RA. This will allow clinicians access to clean (already segmented) volumetric rendering of pathological skeletal anatomy without the high doses of ionising radiation typically associated with CT scans, or where there is no access to volumetric scans.

We are also planning to address some of the problems discussed in previous section:

- We can pad white spaces (same pixel values as background) around images to make them bigger. The alternative way is to impose constraints on parameters where fingers are not allowed to extend outside images.
- There are some images (e.g. Image 5) where fingers are mis-matched or overlapped. We can look into multi-objective optimisation algorithms with

both MAE and ZNCC as objectives. ZNCC is very helpful for shape matching. The algorithm are working by minimising MAE and maximising ZNCC in the same time.

References

1. Andrade, N., Faria, F.A., Cappabianco, F.A.M.: A practical review on medical image registration: from rigid to deep learning based approaches. In: 2018 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI). pp. 463–470. IEEE (2018)
2. Aubert, B., Vazquez, C., Cresson, T., Parent, S., de Guise, J.A.: Toward automated 3D spine reconstruction from biplanar radiographs using CNN for statistical spine model fitting. IEEE Trans Med Imaging **38**(12), 2796–2806 (2019)
3. Blank, J., Deb, K.: Pymoo: Multi-objective Optimization in python. IEEE Access **8**, 89497–89509 (2020)
4. Englander, Z.A., Martin, J.T., Ganapathy, P.K., Garrett, W.E., DeFrate, L.E.: Automatic registration of MRI-based joint models to high-speed biplanar radiographs for precise quantification of in vivo anterior cruciate ligament deformation during gait. J Biomech **81**, 36–44 (2018)
5. Esteban, J., Grimm, M., Unberath, M., Zahnd, G., Navab, N.: Towards fully automatic X-ray to CT registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 631–639. Springer (2019)
6. Folkerts, M., Jia, X., Gu, X., Choi, D., Majumdar, A., Jiang, S.: MO-FF-A4-05: Implementation and evaluation of various DRR algorithms on GPU. Medical Physics **37**(6Part6), 3367–3367 (2010)
7. Gao, C., Liu, X., Gu, W., Killeen, B., Armand, M., Taylor, R., Unberath, M.: Generalizing spatial transformers to projective geometry with applications to 2D/3D registration. arXiv preprint arXiv:2003.10987 (2020)
8. Gómez, O., Ibáñez, O., Valsecchi, A., Bermejo, E., Molina, D., Cordón, O.: Performance analysis of real-coded evolutionary algorithms under a computationally expensive optimization scenario: 3D–2D comparative radiography. Applied Soft Computing **97**, 106793 (2020)
9. Gómez, O., Ibanez, O., Valsecchi, A., Cordón, O., Kahana, T.: 3D-2D silhouette-based image registration for comparative radiography-based forensic identification. Pattern Recognition **83**, 469–480 (2018)
10. Gong, R.H., Stewart, J., Abolmaesumi, P.: Multiple-object 2-D–3-D registration for noninvasive pose identification of fracture fragments. IEEE Trans Biomed Eng **58**(6), 1592–1601 (2011)
11. Graham, R.N., Perriss, R.W., Scarsbrook, A.F.: DICOM demystified: a review of digital file formats and their use in radiological practice. Clinical radiology **60**(11), 1133–1140 (2005)
12. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. Evol Comput **9**(2), 159–195 (2001)
13. Huang, J., Yagel, R., Filippov, V., Kurzion, Y.: An accurate method for voxelizing polygon meshes. In: Kaufman, A.E., Yagel, R., Lorensen, W.E. (eds.) Proceeding of the 1998 IEEE Symposium on Volume Visualization, VVS 1998, Research Triangle Park, NC, USA, October 19-20, 1998. pp. 119–126. ACM / IEEE Comput Soc (1998)

14. Kanhonou, M., Cresson, T., Clement, J., Lavoie, F., Hagemeister, N., De Guise, J.: Registration and motion analysis. Int J CARS **9**(1), S31–S34 (2014)
15. Neumann, E., Lefèvre, S., Zimmermann, B., Gay, S., Müller-Ladner, U.: Rheumatoid arthritis progression mediated by activated synovial fibroblasts. Trends in molecular medicine **16**(10), 458–468 (2010)
16. N.Freud, P.Duvauchelle, Létang, J.M., D.Babot: 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**(1), 175 – 180 (2006)
17. Pelechas, E., Kaltsonoudis, E., Voulgari, P.V., Drosos, A.A.: Illustrated Handbook of Rheumatic and Musculo-Skeletal Diseases. Springer (2018)
18. Rajpurkar, P., Irvin, J., Bagul, A., Ding, D., Duan, T., Mehta, H., Yang, B., Zhu, K., Laird, D., Ball, R.L., et al.: Mura: Large dataset for abnormality detection in musculoskeletal radiographs. arXiv preprint arXiv:1712.06957 (2017)
19. Schett, G.: Erosive arthritis. Arthritis research & therapy **9**(1), 1–6 (2007)
20. Valsecchi, A., Bermejo, E., Damas, S., Cordón, O.: Metaheuristics for medical image registration. In: Martí, R., Panos, P., Resende, M.G.C. (eds.) Handbook of Heuristics, pp. 1–22. Springer International Publishing, Cham (2017)
21. 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. pp. 25–32. Eurographics Association, Cardiff, UK (Jun 2009)
22. 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**, 1–15 (Apr 2016)
23. Viergever, M.A., Antoine Maintz, J., Klein, S., Murphy, K., Staring, M., Pluim, J.P.: A survey of medical image registration–under review. Med Image Anal **33**, 140–144 (Oct 2016)
24. Wang, M., Li, P.: A review of deformation models in medical image registration. J Med Biol Eng **39**(1), 1–17 (2019)
25. Wen, T., Mihail, R., Al-maliki, S., Létang, J.M., Vidal, F.P.: Registration of 3D triangular models to 2D X-ray projections using black-box optimisation and X-ray simulation. In: Computer Graphics and Visual Computing (CGVC). pp. 105–113. The Eurographics Association (2019)