



---

## QUASICONFORMAL MODEL WITH CNN FEATURES FOR LARGE DEFORMATION IMAGE REGISTRATION

HO LAW

School of Mathematics  
Georgia Institute of Technology, Atlanta, GA 30332, USA

GARY P. T. CHOI

Department of Mathematics  
Massachusetts Institute of Technology, Cambridge, MA 02139, USA

KA CHUN LAM

Machine Learning Team  
National Institute of Mental Health, Bethesda, MD 20892, USA

LOK MING LUI

Department of Mathematics  
The Chinese University of Hong Kong, Shatin, NT, Hong Kong, China

(Communicated by Hao-Min Zhou)

**ABSTRACT.** Image registration has been widely studied over the past several decades, with numerous applications in science, engineering and medicine. Most of the conventional mathematical models for large deformation image registration rely on prescribed landmarks, which usually require tedious manual labeling. In recent years, there has been a surge of interest in the use of machine learning for image registration. In this paper, we develop a novel method for large deformation image registration by a fusion of quasiconformal theory and convolutional neural network (CNN). More specifically, we propose a quasiconformal energy model with a novel fidelity term that incorporates the features extracted using a pre-trained CNN, thereby allowing us to obtain meaningful registration results without any guidance of prescribed landmarks. Moreover, unlike many prior image registration methods, the bijectivity of our method is guaranteed by quasiconformal theory. Experimental results are presented to demonstrate the effectiveness of the proposed method. More broadly, our work sheds light on how rigorous mathematical theories and practical machine learning approaches can be integrated for developing computational methods with improved performance.

---

2020 *Mathematics Subject Classification.* 65D18, 68U05, 68U10, 68T07.

*Key words and phrases.* Image registration, convolutional neural networks, quasiconformal theory.

This work was supported in part by the National Science Foundation under Grant No. DMS-2002103 (to Gary P. T. Choi), and HKRGC GRF under project ID 14305919 (to Lok Ming Lui).

© 2022 The Author(s). Published by AIMS, LLC. This is an Open Access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

**1. Introduction.** Image registration aims at establishing a meaningful correspondence between two images based on a given metric. Since the 1990s, it has been extremely useful in medical imaging, computer graphics and computer vision. For instance, one can register two medical images of a patient at different time points for disease diagnosis. One can also utilize image registration for creating animations and for tracking objects in different video frames. In particular, it is usually desirable but also more challenging to consider *non-rigid* image registrations, which involve transformations beyond simple translation, rotation and scaling, such as affine maps, diffeomorphic maps, conformal maps and quasiconformal maps.

Conventional mathematical models for non-rigid image registration can be categorized into three major types: landmark-based methods, intensity-based methods, and hybrid methods. Landmark-based methods use prescribed feature points (also known as *landmarks*) to guide the registration, while intensity-based methods use only the intensity of the images to compute the registration. Hybrid methods take the advantages of the two methods by considering both the image intensity and the landmark correspondences, which make them particularly effective for the case where the source image and the target image are assumed to be with a large deformation. However, in practice, manual landmark labeling is time-consuming and easily affected by human errors.

In this work, we propose a novel variational model for large deformation image registration by combining quasiconformal theory and convolutional neural network (CNN). Specifically, instead of using a landmark-based fidelity term, we propose a novel fidelity term that uses feature vectors obtained from a truncated classification network to guide the registration. As the CNN features do not necessarily yield a smooth and bijective map, quasiconformal energy terms are included in the variational model to reduce quasiconformal distortion and ensure bijectivity. Our experiments show that even only using a common pre-trained classification network, the proposed variational model already works very well for registering a large variety of images.

The rest of the paper is organized as follows. In Section 2, we highlight the contributions of our work. In Section 3, we review the related works on image registration and machine learning. In Section 4, we introduce some notations used in this work. In Section 5, we introduce the theory of quasiconformal maps and convolutional neural networks. The proposed model and the implementation are then described in Section 6 and Section 7 respectively. In Section 8, we demonstrate the effectiveness of the proposed algorithm using various experiments. We conclude the paper and discuss possible future works in Section 9.

**2. Contributions.** The contributions of our work are as follows:

- (i) Our proposed variational model involves a novel combination of quasiconformal maps and deep learning for large deformation image registration. Specifically, we propose a fidelity term to incorporate the features extracted using a pre-trained classification CNN into our quasiconformal energy model.
- (ii) Unlike many prior image registration models, our method does not rely on any prescribed landmarks.
- (iii) By simply using a common pre-trained classification network for getting the CNN feature vectors, our proposed model is already capable of yielding accurate registration results.

- (iv) The bijectivity of the proposed method is theoretically guaranteed by quasiconformal theory.

### 3. Related works.

**3.1. Non-rigid image registration.** Over the past several decades, image registration has been extensively studied (see [7, 38, 45, 52] for detailed surveys). In [23], Horn and Schunck presented a method for registering images using optical flow. In [27], Johnson and Christensen proposed a method for consistent landmark and intensity-based image registration using thin-plate spline (TPS) [6]. In [5], Beg *et al.* proposed the large deformation diffeomorphic metric mapping (LDDMM) algorithm for image warping. Joshi and Miller [28] used large deformation diffeomorphisms for matching landmarks. Other publicly available image registration tools include Elastix [30], deformable image registration using discrete optimization (DROP) [19], flexible algorithms for image registration (FAIR) [40] etc.

Diffeomorphic Demons (DDemons), developed by Vercauteren *et al.* [47], is a well-known non-parametric diffeomorphic image registration method stemming from the work of Thirion [46]. By combining DDemons and quasiconformal theory, Lam and Lui [34] proposed the quasiconformal hybrid registration (QCHR) method that reduces the local geometric distortion of the registration map. The QCHR method has been successfully applied to different registration and shape analysis tasks [12, 50]. However, in case manual landmark labeling is not available and the main features in the input images do not overlap, these methods may not work well. More specifically, if there is little or no overlap of the features in the source image and the target image, these methods will only shrink the features and yield a meaningless registration map. Therefore, these methods can only handle large deformation image registration with the presence of prescribed landmarks.

**3.2. Imaging and deep learning.** In recent years, the use of neural networks has become increasingly popular for imaging and computer vision [4, 21, 25, 26, 42, 51]. In particular, there have been a number of works that use deep learning framework for image and surface registration, such as DLIR [16], VoxelMorph [3], FAIM [33], Quicksilver [49] and cycle-consistent training [32]. These methods usually require carefully designing the neural network architecture and performing extensive training for achieving the registration. By contrast, the goal of our work is to develop an energy model for image registration without having to build any registration network or perform any training. Our method only uses feature vectors extracted by a pre-trained classification network and solves an energy minimization problem to obtain the registration.

**4. Notations.** Here we introduce some notations used in this work.

- $z = x + iy$ : A complex number with the real part  $x$  and the imaginary part  $y$ , where  $i^2 = -1$ .
- $\bar{z} = x - iy$ : The complex conjugate of a complex number  $z = x + iy$ .
- $|z| = \sqrt{x^2 + y^2}$ : The modulus of a complex number  $z = x + iy$ .
- $\|\mu\|_\infty = \sup_{z \in \mathbb{C}} |\mu(z)|$ : The supremum norm of a function  $\mu : \mathbb{C} \rightarrow \mathbb{C}$ .
- $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})^T$ : The gradient operator for functions  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ .
- $\|\mathbf{x}\|_p = (x_1^p + x_2^p + \dots + x_n^p)^{1/p}$ : The  $L^p$ -norm of a vector  $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ , where  $p$  is a positive integer.




FIGURE 1. An illustration of how the Beltrami coefficient  $\mu$  determines the conformality distortion. Under a quasiconformal map  $f$ , an infinitesimal circle around a point  $p$  is mapped to an infinitesimal ellipse centered at  $f(p)$ , where the major axis length and the minor axis length are given by  $|f_z(p)|(1 + |\mu(p)|)$  and  $|f_z(p)|(1 - |\mu(p)|)$ , respectively. Therefore, the maximal dilation of  $f$  is  $K(f) = \frac{1 + \|\mu\|_\infty}{1 - \|\mu\|_\infty}$ . Also, the orientation change of the major axis of the ellipse is given by  $\arg(\mu(p))/2$ .

- $\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}$ : The Frobenius norm of an  $m \times n$  matrix  $A$  with its  $(i, j)$ -th entry being  $a_{ij}$ .

## 5. Theoretical background.

**5.1. Quasiconformal theory.** In this work, we use quasiconformal maps to obtain diffeomorphic image registrations with large deformations. In this section, we describe some related concepts in quasiconformal theory. Readers are referred to [17, 36] for more details.

Let  $f : \mathbb{C} \rightarrow \mathbb{C}$  be a holomorphic function on the complex plane with  $f(z) = f(x, y) = u(x, y) + iv(x, y)$ , where  $z = x + iy$ ,  $u, v$  are real-valued functions, and the derivative of  $f$  is non-zero everywhere.  $f$  is said to be *conformal* if it satisfies the Cauchy–Riemann equations

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y} \quad \text{and} \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}. \quad (1)$$

If we denote

$$\frac{\partial f}{\partial \bar{z}} = \frac{1}{2} \left( \frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right) \quad \text{and} \quad \frac{\partial f}{\partial z} = \frac{1}{2} \left( \frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right), \quad (2)$$

then Equation (1) can be rewritten as

$$\frac{\partial f}{\partial \bar{z}} = 0. \quad (3)$$

Quasiconformal maps are a generalization of conformal maps. Intuitively, the first order approximations of conformal maps take small circles to small circles, while the first order approximations of quasiconformal maps take small circles to small ellipses of bounded eccentricity [17]. Mathematically, an orientation-preserving homeomorphism  $f : \mathbb{C} \rightarrow \mathbb{C}$  is said to be *quasiconformal* if it satisfies the Beltrami equation:

$$\frac{\partial f}{\partial \bar{z}} = \mu(z) \frac{\partial f}{\partial z} \quad (4)$$

for some complex-valued function  $\mu$  (called the *Beltrami coefficient*) with  $\|\mu\|_\infty < 1$ . Note that  $\mu$  measures how far the map at each point is deviated from a conformal map. In particular, the map  $f$  is conformal around a small neighborhood of  $p$  if and only if  $\mu(p) = 0$ . One may express  $f$  around a point  $p$  with respect to its local parameter as follows (see Fig. 1):

$$f(z) = f(p) + f_z(p)(z - p) + f_{\bar{z}}(p)\overline{z - p} = f(p) + f_z(p)(z - p + \mu(p)\overline{z - p}). \quad (5)$$

In other words, locally  $f$  can be considered as a map composed of a translation by  $f(p)$  together with a stretch map  $S(z) = z - p + \mu(p)\overline{z - p}$  with a multiplication of  $f_z(p)$ . Because of the factor  $\mu(p)$  in  $S(z)$ ,  $f$  maps a small circle to a small ellipse. More specifically, the maximal magnification factor is  $|f_z(p)|(1 + |\mu(p)|)$  and the maximal shrinkage factor is  $|f_z(p)|(1 - |\mu(p)|)$ . The maximal dilation of  $f$  is then given by

$$K(f) = \frac{1 + \|\mu\|_\infty}{1 - \|\mu\|_\infty}. \quad (6)$$

Also, the orientation change of the major axis of the ellipse is given by  $\arg(\mu(p))/2$ . This shows that the Beltrami coefficient  $\mu$  provides us with useful information of the quasiconformality of the mapping  $f$ .

Besides, given a Beltrami coefficient  $\mu : \mathbb{C} \rightarrow \mathbb{C}$  with  $\|\mu\|_\infty < 1$ , there is always a quasiconformal mapping from  $\mathbb{C}$  onto itself which satisfies the Beltrami equation (4) in the distribution sense [17]. More precisely, we have the following theorem:

**Theorem 5.1** (Measurable Riemann Mapping Theorem). *Suppose  $\mu : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$  is Lebesgue measurable and satisfies  $\|\mu\|_\infty < 1$ . Then there exists a quasiconformal homeomorphism  $\phi$  from  $\overline{\mathbb{C}}$  onto itself, which is in the Sobolev space  $W^{1,2}(\overline{\mathbb{C}})$  and satisfies the Beltrami equation (4) in the distribution sense. Furthermore, by fixing  $0, 1$  and  $\infty$ , the quasiconformal homeomorphism  $\phi$  is uniquely determined for any given  $\mu$ .*

Theorem 5.1 suggests that under suitable normalization, a homeomorphism from  $\mathbb{C}$  onto itself can be uniquely determined by its associated Beltrami coefficient.

**5.2. Linear Beltrami solver.** The Linear Beltrami solver (LBS) [34] provides us with an efficient way for reconstructing a quasiconformal map  $f(z) = u(x, y) + iv(x, y)$  given a Beltrami coefficient  $\mu_f$ . First, note that the Beltrami equation (4) can be rewritten as

$$\mu(f) = \frac{(u_x - v_y) + i(v_x + u_y)}{(u_x + v_y) + i(v_x - u_y)}. \quad (7)$$

Now, let  $(\mu \circ f)(z) = \rho(z) + i\tau(z)$  where  $\rho, \tau$  are real-valued functions. We can then express  $v_x$  and  $v_y$  as linear combinations of  $u_x$  and  $u_y$ :

$$\begin{aligned} -v_y &= \alpha_1 u_x + \alpha_2 u_y; \\ v_x &= \alpha_2 u_x + \alpha_3 u_y, \end{aligned} \quad (8)$$

where  $\alpha_1 = \frac{(\rho-1)^2 + \tau^2}{1 - \rho^2 - \tau^2}$ ,  $\alpha_2 = -\frac{2\tau}{1 - \rho^2 - \tau^2}$ , and  $\alpha_3 = \frac{1 + 2\rho + \rho^2 + \tau^2}{1 - \rho^2 - \tau^2}$ . Similarly, we can express  $u_x$  and  $u_y$  as linear combinations of  $v_x$  and  $v_y$ :

$$\begin{aligned} u_y &= \alpha_1 v_x + \alpha_2 v_y; \\ -u_x &= \alpha_2 v_x + \alpha_3 v_y. \end{aligned} \quad (9)$$




FIGURE 2. An illustration of receptive field.

Since  $\nabla \cdot \begin{pmatrix} -v_y \\ v_x \end{pmatrix} = 0$  and  $\nabla \cdot \begin{pmatrix} u_y \\ -u_x \end{pmatrix} = 0$ , we have

$$\nabla \cdot \left( A \begin{pmatrix} u_x \\ u_y \end{pmatrix} \right) = 0 \quad \text{and} \quad \nabla \cdot \left( A \begin{pmatrix} v_x \\ v_y \end{pmatrix} \right) = 0, \quad (10)$$

where  $A = \begin{pmatrix} \alpha_1 & \alpha_2 \\ \alpha_2 & \alpha_3 \end{pmatrix}$ .

In the discrete case, the elliptic PDEs (10) can be discretized as sparse positive definite linear systems and hence can be efficiently solved. Readers are referred to [13, 34] for more details.

**5.3. Convolutional neural networks (CNNs).** Another important component in our proposed method is the use of features extracted by CNNs. A comprehensive introduction to CNN can be found in [20]. Here, we introduce the concept of *receptive field* in CNN, which is particularly related to our work.

In the context of neural network for imaging, a receptive field is an area that is being read by the network at a specific layer. More specifically, the network takes in an image pixel by pixel initially and treats each pixel as a 3-dimensional vector for RGB images. Then after certain layers of convolution and pooling, the network starts to read the image region by region and it does so by assigning each region a high-dimensional vector. This region is called a receptive field. In general, at the end of a CNN, the size of a receptive field would usually outgrow the size of input image, and different receptive fields overlap each other.

An illustration of receptive field is shown in Fig. 2. Suppose the input image is of size  $5 \times 5$ , kernel of size  $3 \times 3$ , padding of size 1 and stride of size  $2 \times 2$ . After two convolutions, the purple pixels, including the padding, on the rightmost grid contribute to the value of the top left corner of the  $2 \times 2$  grid. Thus, these purple pixels in the original input image form a receptive field of that top left value.

**6. Proposed method.** In this section, we describe our proposed quasiconformal energy model with CNN features for large deformation image registration.

**6.1. Measuring the correlation of image patches using a pre-trained network and receptive fields.** In [41], Rocco *et al.* proposed a method for obtaining similarity information from two images using some typical image classification networks. The main idea is to extract a feature vector from each receptive field of an image using a truncated classification network, and then compute the similarities between each receptive field of the source image and the target image using inner

products. Since receptive fields are large and usually overlap with each other, this gives a way of roughly evaluating the interaction between patches on a single image. Motivated by this approach, in this work, we propose to partition the images and then evaluate the correlation of the image patches with the aid of a classification network. We will then incorporate the similarity information into a quasiconformal energy model for computing the registration. Unlike the work by Rocco *et al.*, our approach requires a higher level of preciseness as we will be defining a bijective correspondence between two images. Thus, instead of allowing the receptive field to grow with little control, we partition the images into patches to feed into network, so as to gain back the control of how large and where the network reads for generating a feature vector.

If we feed an image into a typical classification neural network such as ResNet [22], VGG [43] or DenseNet [24], then up to some layer in the middle, we can obtain a high-dimensional vector for each receptive field and stack them vertically to yield a high-dimensional vector. Fig. 3 illustrates the process of producing such a high-dimensional vector. We first partition both the moving image and the fixed image into smaller patches, and then feed each patch through a truncated classification network. Then, a 3D array of size  $m \times n \times d$  is generated, where  $m, n$  are respectively the number of receptive fields along the width and height of the input image depending on stride, kernel and padding size, and  $d$  is the dimension of the feature vector depending on the architecture of the network. We then transform each of these 3D arrays by vertically stacking along one direction, thereby producing the vectors as shown in Fig. 3. Mathematically, this stacking can be viewed as an isomorphism that maps the output array to a vector in  $\mathbb{R}^{mnd}$ . The inner product of this supreme feature vector of two images is then the sum of the correlation scores of the corresponding pairs of receptive fields. This provides us with a quantitative way for measuring the correlation of two images, which plays an important role in our proposed energy model described in the next section.

**6.2. Quasiconformal energy model with CNN features.** The classification network gives us a reliable quantitative measurement of the correlation between two images. The next crucial step is to incorporate this useful information into the registration model. Let  $\mathcal{M}$  and  $\mathcal{S}$  be the moving and static images respectively. Each image is uniformly partitioned into  $m$  small image patches. To find a diffeomorphism  $f : \mathcal{M} \rightarrow \mathcal{S}$  such that  $\|\mu(f)\|_\infty < 1$ , we propose the following model:

$$f = \arg \min_{g: \mathcal{M} \rightarrow \mathcal{S}} E_C(\mu(g)), \tag{11}$$

where

$$E_C(\mu(g)) = \int_{\mathcal{M}} (|\nabla \mu|^2 + \alpha |\mu|^2 + \beta (I_{\mathcal{M}} - I_{\mathcal{S}}(g))^2) dz + \gamma \|C \otimes D(g) - C\|_F^2. \tag{12}$$

Here,  $\alpha, \beta, \gamma$  are some positive real numbers,  $I_{\mathcal{M}}$  and  $I_{\mathcal{S}}$  denote the intensity of the images,  $C$  is an  $m \times m$  matrix where  $m$  is the number of image patches for each image,  $D : \{g : \mathcal{M} \rightarrow \mathcal{S}\} \rightarrow [0, 1]^{m \times m}$  maps  $g$  to an  $m \times m$  matrix, and  $\otimes$  is the Hadamard product with

$$(A \otimes B)_{ij} = (A)_{ij}(B)_{ij} \tag{13}$$

for any matrices  $A, B$  of the same dimension. The key idea of the proposed model is to use the features extracted by a robust CNN to guide the registration process via the last term in the energy (12) while maintaining the diffeomorphic property of the deformation using the first two terms of the integral based on quasiconformal theory.




FIGURE 3. The process of obtaining feature vectors from the images. The two images on the left are the moving image and the fixed image. We first partition both images into smaller patches, and then feed each patch into the truncated classification network to obtain a 3D array of size  $m \times n \times d$ , where  $m, n$  are the number of receptive fields along the width and height of the input image depending on stride, kernel and padding size, and  $d$  is the dimension of the feature vector depending on the architecture of the network. We can then transform this 3D array into a vector by vertically stacking along one direction. Here we partition each image into  $3 \times 3 = 9$  patches for illustrative purposes, and hence 9 vectors in  $\mathbb{R}^{mnd}$  are produced for each of the two images as shown on the right. In practice, a finer partition is often used.

More specifically, the first term  $|\nabla\mu|^2$  of the integral in the energy (12) is used for enhancing the smoothness of  $f$ . Recall that by Theorem 5.1, there exists a correspondence between quasiconformal mappings and the Beltrami coefficients, which measure the local geometric distortion of the mappings. Therefore, a smaller  $|\nabla\mu|^2$  gives a smoother quasiconformal map  $f$ . The second term  $|\mu|^2$  of the integral is used for reducing the conformality distortion of the registration. As described in Section 5, a map is conformal if and only if  $\mu = 0$ . A smaller  $|\mu|^2$  gives a map with a lower local geometric distortion, which is desirable. The third term  $(I_{\mathcal{M}} - I_S(g))^2$  of the integral aims to reduce the intensity difference of the images.

The last term  $\|C \otimes D(g) - C\|_F^2$  in the proposed energy (12) is a novel fidelity term that gives a descent direction to guide the registration based on the correlation of different regions of two images, which plays an important role in our proposed method. More specifically, the feature vectors encode important hidden information extracted by a classification CNN from big image datasets, which may not be easily represented using conventional mathematical quantities such as curvature or gradient. With this novel fidelity term, we can utilize such information and combine it with quasiconformal maps for achieving a meaningful registration. One key feature of our proposed model is that it allows us to save the effort of going through the training of a neural network for registration, as there are already many publicly available classification networks well-trained on giving meaningful feature vectors



that we can re-use. Another advantage of this approach is the generality: for most of the common images, the CNN should give fairly good guidance for rudimentary correspondence on two images even if it is not trained on a specific type of images. We remark that in case we want to focus on a more specific type of images, it is also possible to use a pre-trained classification network on such type of images to get the feature vectors. Below, we first introduce the definition of  $C$  and  $D(g)$ , and then explain the use of the Hadamard product.

Roughly speaking, the matrix  $C$  is used for deciding how trustworthy the correlation between a pair of patches is when compared with the others. To define  $C$ , we first denote  $H : P \rightarrow \mathbb{R}^d$  as a truncated classification network that sends an image patch  $P$  to a  $d$ -dimensional feature vector. Define a matrix  $\tilde{C} \in \mathbb{R}^{m \times m}$ , where  $m$  is the number of patches we partition on each image [41]:

$$(\tilde{C})_{ij} = \left\langle \frac{H(P_i^{\mathcal{M}})}{\|H(P_i^{\mathcal{M}})\|_2}, \frac{H(P_j^{\mathcal{S}})}{\|H(P_j^{\mathcal{S}})\|_2} \right\rangle, \tag{14}$$

where  $P_i^{\mathcal{M}}$  is the  $i$ -th image patch of the moving image  $\mathcal{M}$  and  $P_j^{\mathcal{S}}$  is the  $j$ -th image patch of the static image  $\mathcal{S}$ . It is noteworthy that passing an image patch into a truncated network would yield a 3D matrix of size  $m \times n \times d$ , which is dependent on the architecture of the truncated classification network, that can be viewed as  $\mathbb{R}^{mnd}$ . This allows us to define the inner product of 3D matrices as the usual inner product of vectors.

Then, we normalize each row of  $\tilde{C}$  by

$$(\hat{C})_{ij} = \frac{\tilde{C}_{ij} - \delta_i}{\sigma_i}, \tag{15}$$

where  $\sigma_i$  and  $\delta_i$  are the standard deviation and mean of row  $i$  respectively. The reason for normalising each row of  $\tilde{C}$  is to spread out the similarity of patches. For example, if a row in  $\tilde{C}$  happens to have most entries fluctuate within a small range, it is not easy to assess the similarity between the patches. Normalising the matrix row-by-row is beneficial for finding descent direction.

In many applications, it is common that there is a large background region with a uniform color in the images to be registered. For instance, the X-ray, CT and MRI scans in medical imaging are usually with a completely black background. For such uniform background patches in the moving image, it is expected that their corresponding rows in  $\hat{C}$  would be exactly the same and hence are not useful for the registration process. Since it is expected that the feature vectors extracted from different background patches on the source image would give an inner product almost the same with the feature vectors extracted from any patches on the target image, there is a need to further prune the meaningless rows in  $\hat{C}$ . To neglect these background rows, we define an  $m \times m$  elimination matrix  $E$  with

$$(E)_{ij} = \begin{cases} 0, & \text{if } i \neq j, \\ 0, & \text{if } i = j \text{ and row } i \text{ is not unique in } \hat{C}, \\ 1, & \text{if } i = j \text{ and row } i \text{ is unique in } \hat{C}. \end{cases} \tag{16}$$

Now, we can remove the background rows in  $\hat{C}$  and obtain an updated matrix  $C^\dagger \in \mathbb{R}^{m \times m}$  by

$$C^\dagger = E\hat{C}. \tag{17}$$

Finally, the matrix  $C$  in the proposed energy (12) is constructed by sparsifying  $C^\dagger$ : For each row in  $C^\dagger$ , we keep the largest entry and set all other entries to be zero. The reason for sparsifying  $C^\dagger$  is essentially for computation efficiency. More specifically, after the sparsification, we can obtain a descent direction based on the most significant correlation of the patches for the energy minimization in the later steps. Without the sparsification step, the other non-zero entries would generate some misleading direction in finding the descent direction and hence affect the efficiency.

To define  $D(g)$  in the energy (12), we denote the center of image patch  $P_i^{\mathcal{M}}$  and  $P_j^{\mathcal{S}}$  on the moving and static images as  $x_i^{\mathcal{M}}$  and  $x_j^{\mathcal{S}}$  respectively. The mapping  $D : \{g : \mathcal{M} \rightarrow \mathcal{S}\} \rightarrow [0, 1]^{m \times m}$  is defined by

$$(D(g))_{ij} = \exp\left(-\frac{\|g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}\|_2^2}{\sigma^2}\right), \quad (18)$$

where  $\sigma$  is a small number to be chosen.

As for the purpose of the Hadamard product in the proposed fidelity term  $\|C \otimes D(g) - C\|_F^2$ , note that  $(D(g))_{ij} = 1$  if and only if  $g(x_i^{\mathcal{M}}) = x_j^{\mathcal{S}}$ , i.e.  $g$  maps the patch  $P_i^{\mathcal{M}}$  on the moving image  $\mathcal{M}$  to the patch  $P_j^{\mathcal{S}}$  on static image  $\mathcal{S}$  exactly. Now, note that the fidelity term can be expressed as a weighted sum of correspondences:

$$\|C \otimes D(g) - C\|_F^2 = \sum_{ij} C_{ij}^2 ((D(g))_{ij} - 1)^2. \quad (19)$$

Therefore, with the use of the Hadamard product,  $(C)_{ij}$  can serve as a weighting factor to determine how exact the mapping between  $x_i^{\mathcal{M}}$  and  $x_j^{\mathcal{S}}$  should be. More specifically, if  $(C)_{ij}$  is large, i.e. the correlation between the two patches is high as determined by the pre-trained network, then  $(D(g))_{ij}$  should be as close to 1 as possible; if  $(C)_{ij}$  is small, the requirement for  $(D(g))_{ij} = 1$  can be relaxed.

Moreover, in the optimization process, note that the proposed fidelity term can be handled using gradient descent. The descent direction (denoted as  $df_{\mathcal{W}} : \{x_i^{\mathcal{M}}\}_{i=1}^m \rightarrow \mathbb{R}^2$ ) can be explicitly written as

$$df_{\mathcal{W}}(x_i^{\mathcal{M}}) = \frac{4}{\sigma^2} \sum_j (C)_{ij}^2 \left( \exp\left(-\frac{\|g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}\|_2^2}{\sigma^2}\right) - 1 \right) (g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}). \quad (20)$$

Therefore, if  $(C)_{ij} \approx 0$ , i.e. the network determines that the patch  $P_i^{\mathcal{M}}$  on the moving image is not correlated to the patch  $P_j^{\mathcal{S}}$  on the target image, the corresponding descent direction for the  $(i, j)$  pair will be close to 0. In other words, the Hadamard product allows us to eliminate unwanted descent directions based on the information provided by the CNN encoded in the matrix  $C$ .

On the contrary, if we replace the proposed fidelity term with  $\|D(g) - C\|_F^2$ , then under the same condition of  $(C)_{ij} \approx 0$ , the descent direction will become

$$\begin{aligned} df(x_i^{\mathcal{M}}) &= \frac{4}{\sigma^2} \sum_j \left( \exp\left(-\frac{\|g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}\|_2^2}{\sigma^2}\right) - (C)_{ij} \right) (g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}) \\ &\approx \frac{4}{\sigma^2} \sum_j \exp\left(-\frac{\|g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}\|_2^2}{\sigma^2}\right) (g(x_i^{\mathcal{M}}) - x_j^{\mathcal{S}}) \\ &\neq 0 \quad \text{if } g(x_i^{\mathcal{M}}) \neq x_j^{\mathcal{S}}. \end{aligned} \quad (21)$$

This shows that the descent direction is noisy in case the Hadamard product is not used.

After justifying our proposed energy model (12), it is natural to ask whether the existence of the minimizer (11) is guaranteed. We have the following result:

**Proposition 6.1** (Existence of minimizer). *Let*

$$\mathcal{A} = \{\nu \in C^1(\Omega_1) : \|D\nu\|_\infty \leq L; \|\nu\|_\infty \leq 1 - \epsilon\} \tag{22}$$

for some  $L > 0$  and some small  $\epsilon > 0$ . Then the proposed energy  $E_C$  has a minimizer in  $\mathcal{A} \subset C^1(\Omega_1)$  if  $I_M, I_S$  are  $L^2$  functions.

*Proof.* First, note that if  $I_M, I_S$  are  $L^2$  functions, then  $E_C$  is continuous. Therefore, to prove that  $E_C$  has a minimizer in  $\mathcal{A}$ , it suffices to prove that  $\mathcal{A}$  is compact. The proof follows from the proof of Proposition 4.1 in [34] and is outlined below.

Note that the Beltrami coefficient  $\mu^*$  associated with the unique Teichmüller map is in  $\mathcal{A}$  by choosing  $L = \|D\mu^*\|_\infty$  and  $\epsilon \in (0, 1 - \|\mu^*\|_\infty)$ . Therefore,  $\mathcal{A} \neq \emptyset$ .

Now, we prove that  $\mathcal{A}$  is complete. Let  $\{\nu_n\}_{n=1}^\infty$  be a Cauchy sequence in  $\mathcal{A}$  under the norm  $\|\cdot\|_s := \|\cdot\|_\infty + \|D\cdot\|_\infty$ . Then  $\{D\nu_n\}_{n=1}^\infty$  is also a Cauchy sequence under the norm  $\|\cdot\|_s$ . As  $\mathcal{A} \subset W^{1,\infty}(\Omega_1)$  and  $W^{1,\infty}$  is complete, it follows that  $\{D\nu_n\}_{n=1}^\infty$  converges to some  $g \in W^{1,\infty}$  uniformly. As  $D\nu_n$  is continuous,  $g$  is also continuous. Also,  $\{\nu_n\}_{n=1}^\infty$  is convergent. Hence, we have  $\nu_n \rightarrow \nu$  and  $D\nu_n \rightarrow D\nu$  uniformly for some  $\nu \in C^1(\Omega_1)$ . Moreover, as  $\|\nu_n\|_\infty \leq 1 - \epsilon$  for all  $n$ , we have  $\|\nu\|_\infty \leq 1 - \epsilon$ . Similarly, as  $\|D\nu_n\|_\infty \leq L$  for all  $n$ , we have  $\|D\nu\|_\infty \leq L$ . Therefore,  $\nu \in \mathcal{A}$  and hence  $\mathcal{A}$  is Cauchy complete. We can also show that  $\mathcal{A}$  is totally bounded using the argument in the proof of Proposition 4.1 in [34]. Hence,  $\mathcal{A}$  is compact.

As  $E_C$  is continuous in  $\mathcal{A}$  and  $\mathcal{A}$  is compact, we conclude that  $E_C$  has a minimizer in  $\mathcal{A}$ . □

As a remark, a major difference between the proposed energy model and the landmark- and intensity-based quasiconformal hybrid registration (QCHR) method [34] is that QCHR uses the prescribed feature pairs as hard landmark constraints for the energy minimization, while in our method we do not impose any landmark constraints. Using a common classification network pre-trained on very wide-ranging datasets, we obtain multiple correlated pairs of image patches. Instead of enforcing them as landmark constraints, we encode the correlation information in the proposed fidelity term in the energy (12), which helps supply a good descent direction for yielding a good registration. This does not only avoid any potential non-bijectionality due to incorrectly correlated pairs but also allow the energy model to utilize the correlation information for the registration of more general images that it has not seen before.

**6.3. Energy minimization using the penalty splitting method.** The penalty splitting method is used for solving the energy minimization problem (11). More specifically, instead of directly minimizing the energy (12), we minimize the following energy

$$E_C(\nu, f) = \int_{\mathcal{M}} (|\nabla\nu|^2 + \alpha|\nu|^2 + \rho|\nu - \mu(f)|^2 + \beta(I_M - I_S(f))^2) dz + \gamma \|C \otimes D(f) - C\|_F^2. \tag{23}$$

Here, the new term  $\rho|\nu - \mu(f)|^2$  forces  $\mu(f)$  to closely resemble  $\nu$ . This allows us to consider the Beltrami coefficient  $\mu$  and the mapping  $f$  separately, so that the optimization can be performed more easily.

Minimizing over  $\nu$ . If we fix  $f = f_n$ , it suffices to minimize

$$E_C(\nu, f_n) = \int_{\mathcal{M}} (|\nabla \nu|^2 + \alpha |\nu|^2 + \rho |\nu - \mu(f_n)|^2) dz. \quad (24)$$

As discussed in [34], the minimizer of the above energy can be obtained by solving the Euler–Lagrange equation:

$$(-\Delta + 2\alpha I + 2\rho I)\nu_{n+1} = 2\rho\mu(f_n). \quad (25)$$

In the discrete case, Equation (25) can be discretized as a sparse linear system and be solved efficiently.

Minimizing over  $f$ . If we fix  $\nu = \nu_n$ , it suffices to minimize

$$E_C(\nu_n, f) = \int_{\mathcal{M}} (\rho |\nu_n - \mu(f)|^2 + \beta (I_{\mathcal{M}} - I_{\mathcal{S}}(f))^2) dz + \gamma \|C \otimes D(f) - C\|_F^2. \quad (26)$$

This can be done by performing gradient descent on the Beltrami coefficient  $\mu = \mu_f$  associated with  $f$ .

First, note that the descent direction for the intensity term  $(I_{\mathcal{M}} - I_{\mathcal{S}}(f))^2$  in the space of mappings (denoted as  $df_{\mathcal{I}}$ ) is given by

$$df_{\mathcal{I}} = -2(I_{\mathcal{M}} - I_{\mathcal{S}}(f))\nabla I_{\mathcal{S}}(f). \quad (27)$$

Also, the descent direction  $df_{\mathcal{W}}$  for the fidelity term  $\|C \otimes D(f) - C\|_F^2$  in the space of mappings is given by Equation (20). The next step is to obtain the descent directions in the space of Beltrami coefficients (denoted as  $d\mu_{\mathcal{I}}$  and  $d\mu_{\mathcal{W}}$ ) from  $df_{\mathcal{I}}$  and  $df_{\mathcal{W}}$ . Note that  $df_{\mathcal{I}}$  and  $df_{\mathcal{W}}$  can be viewed as a perturbation:

$$\frac{\partial(f + df_i)}{\partial \bar{z}} = (\mu + d\mu_i) \frac{\partial(f + df_i)}{\partial z}, \quad i = \mathcal{I}, \mathcal{W}. \quad (28)$$

Then, following the approach in [34], we can obtain  $d\mu_i$  as follows:

$$d\mu_i = \left( \frac{\partial df_i}{\partial \bar{z}} - \mu \frac{\partial df_i}{\partial z} \right) / \frac{\partial(f + df_i)}{\partial z}, \quad i = \mathcal{I}, \mathcal{W}. \quad (29)$$

For the term  $|\nu_n - \mu(f)|^2$ , the descent direction in the space of Beltrami coefficients (denoted as  $d\mu_{\mathcal{D}}$ ) is given by

$$d\mu_{\mathcal{D}} = -2(\nu_n - \mu(f)). \quad (30)$$

Combining the three descent directions in Equation (29) and Equation (30), we can update  $\mu$  at every gradient descent iteration as follows, for some  $t_1, t_2, t_3$  chosen according to  $\rho, \beta, \gamma$ :

$$\mu_{n+1} = \mu_n + (t_1 d\mu_{\mathcal{I}} + t_2 d\mu_{\mathcal{W}} + t_3 d\mu_{\mathcal{D}}). \quad (31)$$

After obtaining the new  $\mu_{n+1}$ , we can then reconstruct a quasiconformal map  $f_{n+1}$  associated with  $\mu_{n+1}$  using the LBS method [34].

**6.4. Additional intensity-based registration.** In general, the input source and target images will be well registered after we perform the minimization on the energy (12). However, sometimes some details of the images may still not be perfectly matched. Therefore, we introduce an additional step to further improve the registration result by performing a fully intensity-based matching.

More specifically, at the  $n$ -th iteration, we use the DDemons method [47] to register the current mapping result  $f_n(\mathcal{M})$  and the target image  $\mathcal{S}$ , and denote the updated map as  $f_n + df$ . Then, we adopt the same strategy as described above to

compute a descent direction  $d\nu$  in the space of Beltrami coefficients. We can then obtain an updated Beltrami coefficient by

$$\nu_{n+1} = \nu_n + td\nu, \quad (32)$$

from which we can reconstruct a quasiconformal map  $f_{n+1}$  using the LBS method [34], with a proper truncation for enforcing that  $\|\nu_{n+1}\|_\infty < 1$ . We repeat the process until meeting the convergence condition.

Our proposed registration method is summarized in Algorithm 1.

---

**Algorithm 1** Quasiconformal image registration with CNN features

---

**Input:** A moving image  $\mathcal{M}$ , a static image  $\mathcal{S}$ , a pre-trained CNN with truncation, the partition parameter  $m$ , the convergence threshold  $\epsilon$ , and the maximum number of iterations  $n_{\max}$ .

**Output:** A quasiconformal map  $f : \mathcal{M} \rightarrow \mathcal{S}$ .

- 1: Partition each image into  $m$  pieces.
  - 2: Pass all  $2m$  pieces into the truncated network and obtain the matrix  $C$ .  
*Energy minimization:*
  - 3: Initialize  $\mu_0 = 0$  and  $\nu_0 = 0$ .
  - 4: **while**  $|\nu_{n+1} - \nu_n| > \epsilon$  and  $n \leq n_{\max}$  **do**
  - 5:   Fixing  $f = f_n$ , obtain  $\nu_{n+1}$  by solving the Euler–Lagrange equation (25).
  - 6:   Reconstruct a quasiconformal map with the Beltrami coefficient  $\nu_{n+1}$  using the LBS method [34].
  - 7:   Fixing  $\nu = \nu_{n+1}$ , obtain  $\mu(f_{n+1})$  by the gradient descent method in Equation (31).
  - 8:   Reconstruct a quasiconformal map  $f_{n+1}$  with the Beltrami coefficient  $\mu_{n+1}$  using the LBS method [34].
  - 9:   Update  $n \leftarrow n + 1$ .
  - 10: **end while**  
*Additional intensity-based registration:*
  - 11: **while**  $|\nu_{n+1} - \nu_n| > \epsilon$  and  $n \leq n_{\max}$  **do**
  - 12:   Use the DDemons method [47] to register  $f_n(\mathcal{M})$  and  $\mathcal{S}$  and denote the mapping as  $f_n + df$ .
  - 13:   Obtain  $d\nu$  using the same procedure as in Equation (28) and Equation (29).
  - 14:   Update  $\nu_{n+1} \leftarrow \nu_n + td\nu$ .
  - 15:   Reconstruct a quasiconformal map  $f_{n+1}$  from the Beltrami coefficient  $\nu_{n+1}$  using the LBS method [34], with a proper truncation to enforce  $\|\nu_{n+1}\|_\infty < 1$  if necessary.
  - 16:   Update  $n \leftarrow n + 1$ .
  - 17: **end while**
- 

**7. Numerical implementation.** In the discrete case, we discretize the input images in the form of triangular meshes. Let  $V^1 = \{v_i^1\}_{i=1}^n$ ,  $V^2 = \{v_i^2\}_{i=1}^n$  be the vertex sets and  $F^1 = \{T_j^1\}_{j=1}^n$ ,  $F^2 = \{T_j^2\}_{j=1}^n$  be the face sets of the moving image  $\mathcal{M}$  and the static image  $\mathcal{S}$  respectively.

**7.1. Discretization of the Euler–Lagrange equation.** The Beltrami coefficient  $\mu(T)$  is first discretized on each triangular face  $T$  (see [13, 34] for details). We can

then compute the Beltrami coefficient on a vertex by taking the average value of  $\mu$  on its one-ring neighboring faces:

$$\mu(v_i) = \frac{1}{N_i} \sum_{T \in N_i} \mu(T), \quad (33)$$

where  $N_i$  is the collection of all faces incident to  $v_i$ .

The discrete Laplacian operator  $\Delta$  is given by

$$\Delta(f(v_i)) = \sum_{T \in N_i} \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} (f(v_j) - f(v_i)), \quad (34)$$

where  $\alpha_{ij}$  and  $\beta_{ij}$  are the two angles opposite to a common edge  $[v_i, v_j]$ .

One can then solve the Euler–Lagrange equation (25) on the vertices and finally discretize the solution  $\nu$  on each triangular face by taking the average value at its three vertices:

$$\nu(T) = \frac{1}{3} \sum_{v_i \in T} \nu(v_i). \quad (35)$$

**7.2. Numerical techniques for intensity matching.** We use two different numerical optimization techniques for the intensity-based registration in the main energy minimization step and the additional intensity-matching step in Algorithm 1. For the main energy minimization step, we apply the modified Demon force by Wang *et al.* [48] to find the deformation:

$$u = \frac{(I_{\mathcal{M}} - I_{\mathcal{S}})\nabla I_{\mathcal{S}}}{|\nabla I_{\mathcal{S}}|^2 + \alpha^2(I_{\mathcal{M}} - I_{\mathcal{S}})^2} + \frac{(I_{\mathcal{M}} - I_{\mathcal{S}})\nabla I_{\mathcal{M}}}{|\nabla I_{\mathcal{M}}|^2 + \alpha^2(I_{\mathcal{M}} - I_{\mathcal{S}})^2}. \quad (36)$$

This modified Demon force is applied when finding  $d\mu_{\mathcal{I}}$  in the update rule (31). This force is good for maintaining the diffeomorphic property of the mapping, the convergence speed as well as the stability with gradient descent. However, for the additional fully intensity-based matching step, as there can be multiple local minima for the intensity function, the gradient descent method does not necessarily yield an optimal result. Therefore, we adopt the BFGS optimization scheme [31], which takes a longer time but achieves a more accurate intensity-based registration result.

**7.3. Gradient descent for the fidelity term.** For the proposed fidelity term, the descent direction  $df(x_i^{\mathcal{M}})$  is given by Equation (20). In practice, we notice that Equation (20) may sometimes give non-orientation preserving descent directions in some local neighborhoods so that there may be triangle fold-overs in the underlying mesh, which require additional procedures of truncating the associated Beltrami coefficients and hence affect the convergence of the computation. To alleviate this issue, we apply a Gaussian smoothing on  $df(x_i^{\mathcal{M}})$  around every point  $x_i^{\mathcal{M}}$  with the smoothing parameter set to be the side length of the image divided by 50. With the smoothing, we can ensure that the descent is in the right direction while there will not be an extremely large change at a point relative to the local neighborhood of it.

**7.4. The choice of the pre-trained network and the model parameters.** Many different CNNs such as ResNet [22], VGG [43], and DenseNet [24] can be used as the pre-trained network for the proposed algorithm. For consistency, in all of our experiments we use a well-known CNN *DenseNet-201* [24], which is a densely connected convolutional network with 201 layers. All layers after the third dense block are truncated. The images are partitioned into  $m = 10 \times 10$ ,  $12 \times 12$ ,  $14 \times 14$ ,  $16 \times 16$  or  $18 \times 18$  patches. For each registration, we consider all choices

of  $m$  and choose the one that gives the best registration performance in terms of the similarity of the warped image and the target image and the smoothness of the mapping. Also, in practice we keep only the top 6 to 50 values of the correlation matrix  $C$  depending on the size of the features and set all other values to be 0. The trained CNN parameters are obtained from the standard online library PyTorch. For all experiments, the parameters remain unchanged and no re-training is done throughout the optimization algorithm.

As for the weighting factors in the proposed energy model (23), in general we set  $\alpha = 5$ ,  $\rho = 50$ ,  $\beta = 25\rho$ , and  $\gamma = 5\rho$ . The parameter in Equation (18) is set to be  $\sigma = 1$ .

**7.5. Multiresolution scheme.** To reduce the computational cost of registering high resolution images, we adopt a multiresolution scheme for the registration procedure [34]. We first coarsen both the moving image  $\mathcal{M}$  and the static image  $\mathcal{S}$  by  $k$  layers, where  $I_j^0 = I_j$  is the densest image and  $I_j^k$  is the coarsest image of  $I_j$  ( $j = 1, 2$ ). The registration process starts with registering  $\mathcal{M}^k$  and  $\mathcal{S}^k$ . After obtaining a diffeomorphism  $f_k$ , we proceed to a finer scale. Specifically, we obtain  $f_{k-1}$  by a linear interpolation on  $f_k$ , which serves as the initial map for the registration at the next finer layer. We repeat the above process until obtaining the registration at the finest (original resolution) layer. Using this multiresolution scheme, the computation can be significantly accelerated.

**8. Experimental results.** To demonstrate the effectiveness of our proposed method for large deformation image registration, we test it on various synthetic and real medical images. We remark that in all examples below, we focus on the energy at the finest (original resolution) layer of the multiresolution scheme. We compare our method with DDEmons [47], LDDMM [5], Elastix [30] and DROP [19] (all with code or software available online [18, 29, 37, 44]). For all the experiments, we used the identity map for the initialization of our proposed algorithm, which corresponds to  $\mu \equiv 0$ .

**8.1. Synthetic images.** We first test our method on a synthetic example with the source image containing a letter ‘Z’ (Fig. 5a) and the target image containing a tilted number ‘2’ (Fig. 5b). Two major difficulties of this example are the large displacement between the ‘Z’ and the ‘2’ and the shape difference between them involving certain sharp angle changes. Although one may use an affine map to roughly align them, the boundary of the two images will be largely mismatched. As shown in Fig. 5c, our proposed registration method effectively morphs the ‘Z’ image to match the ‘2’ image. One can also visualize the registration via the deformed underlying grid as shown in Fig. 5d, from which it can be observed that the mapping is smooth and bijective. As shown in Fig. 5e, the energy decreases rapidly throughout the iterations. For comparison, we apply DDEmons [47] (Fig. 5f), LDDMM [5] (Fig. 5g), Elastix [30] (Fig. 5h), and DROP [19] (Fig. 5i) for registering the images. It can be observed that all these methods fail to register the ‘Z’ and the ‘2’ shape due to the large displacement and shape change.

We then consider another synthetic example with the source image being an eagle (Fig. 6a). For the target image (Fig. 6b), we manually deform the eagle shape such that the wings are expanded wider. We apply our proposed method to obtain a registration between the two images (see Fig. 6c for the warped image and Fig. 6d for the warped image with the deformed underlying grid). Under the registration,

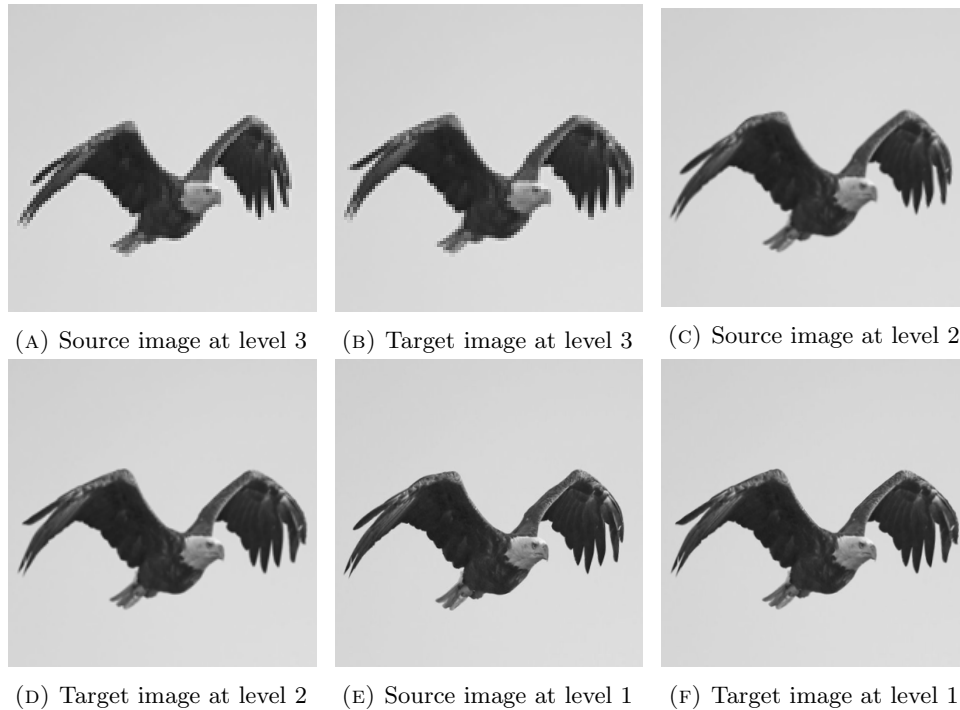


FIGURE 4. An illustration of the multiresolution Scheme. We first coarsen both input source and target images as shown in Fig. 4a and Fig. 4b. Then, we run our proposed algorithm on this coarsen pair. From the obtained mapping on this coarsest level, with linear interpolation, we warp the source image on second level to yield Fig. 4c and register it against Fig. 4d using our proposed algorithm. Finally, with the mapping from the last level, we interpolate it to warp the source image on the finest level as shown in Fig. 4e and register it against Fig. 4f. Refer to Fig. 6 for the final registration result.

the difference between the two images is significantly reduced, and the energy plot in Fig. 6e again shows that our method is highly effective. On the contrary, it can be observed that LDDMM [5] does not produce an accurate registration of the wings (Fig. 6f and Fig. 6g). While the registration result by Elastix [30] looks satisfactory (Fig. 6h), there are multiple overlaps in the deformed underlying grid (Fig. 6i). Overall, our method gives the best registration result with a more natural positional correspondence between different parts as shown by the deformation of the underlying grid and with bijectivity ensured.

In the next example, the source image is a rabbit (Fig. 7a), and the target image is the same rabbit with its ears bent (Fig. 7b). Our method successfully morphs the ears to the desired position as shown in Fig. 7c. From the deformed underlying grid in Fig. 7d, we can see that the mapping is smooth and bijective. Also, the energy plot in Fig. 7e shows that our method effectively reduces the energy within a small number of iterations. On the contrary, LDDMM [5] fails to match the bent ears (Fig. 7f and Fig. 7g). While Elastix [30] is able to deform the ears, the overall



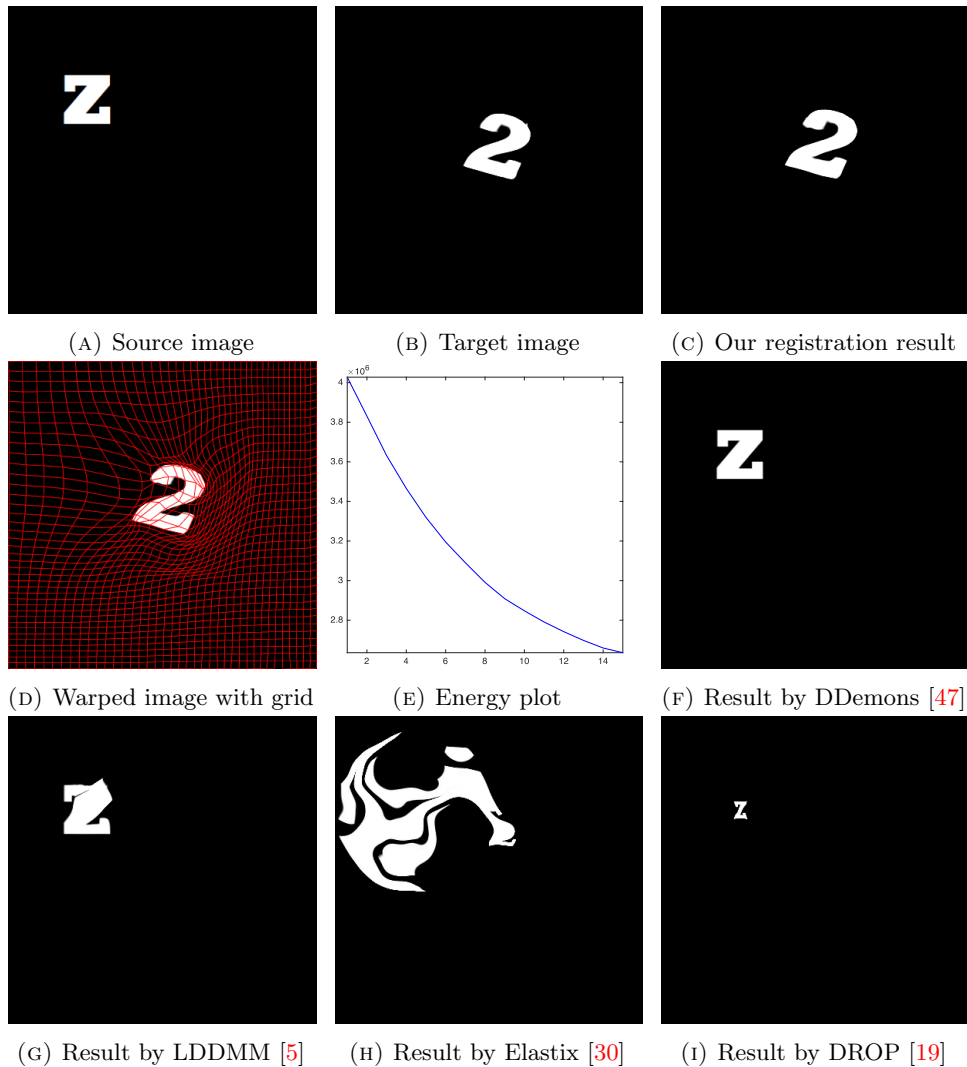


FIGURE 5. The ‘Z’ to ‘2’ example.

image shape is largely distorted (Fig. 7h and Fig. 7i). Fig. 7j and Fig. 7k show the results of using the model (12) without the second term involving the matrix  $C$ . We can see that the left ear of the rabbit shrinks, but fail to match the same ear in the target image. This is expected because without the second term, the algorithm could only compute the registration map based on the intensity difference. During the optimization, Equation (27) minimizes the intensity difference by shrinking the left ear. However, it is rather difficult to have the algorithm to elongate the ear again, since the overall energy has been decreased and a local minimum is reached. Therefore, the algorithm stops at this point.

**8.2. Real X-ray images.** After testing our proposed method on several synthetic images, we now consider applying it on real medical images. Here, we consider a hand X-ray image as the source image (Fig. 8a) and a deformed hand X-ray image

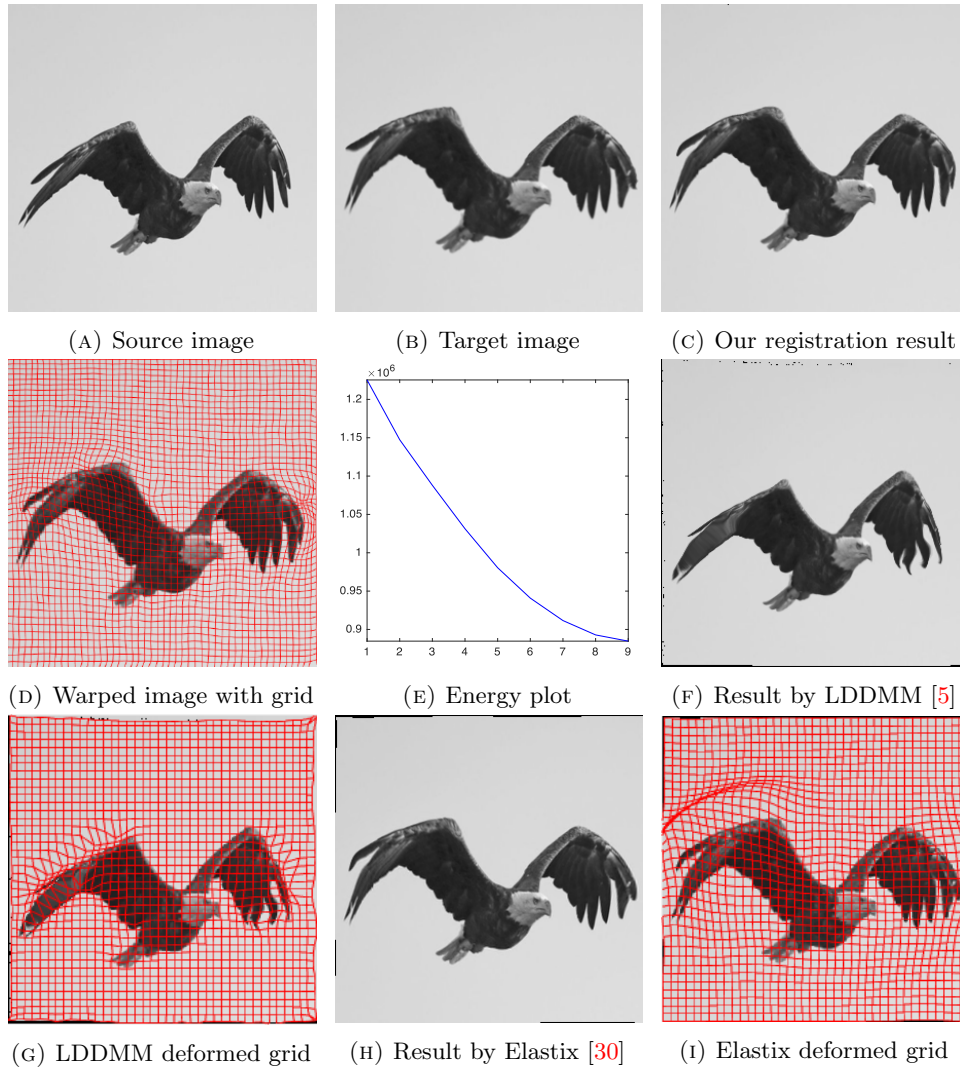


FIGURE 6. The eagle example.

as the target image (Fig. 8b). Fig. 8c shows the original absolute intensity difference between the two images. It can be observed that different fingers are displaced in a nonuniform manner (for example, the displacement of the index finger is much larger than that of the little finger), while the wrist remains almost the same. Therefore, a simple rigid transformation is insufficient for yielding a good registration. As shown in Fig. 8d, our proposed method successfully deforms the source image to match the target image, and the final intensity difference is significantly smaller (see Fig. 8e). From the deformed underlying grid in Fig. 8f, it can be observed that the mapping is smooth and bijective. For comparison, both LDDMM [5] and DDEmons [47] fail to register the fingers and are non-bijective (see Fig. 8g, Fig. 8h, and Fig. 8i).

We then consider another example of registering two hand X-ray images with larger deformations (see Fig. 9a for the source image, Fig. 9b for the target image,

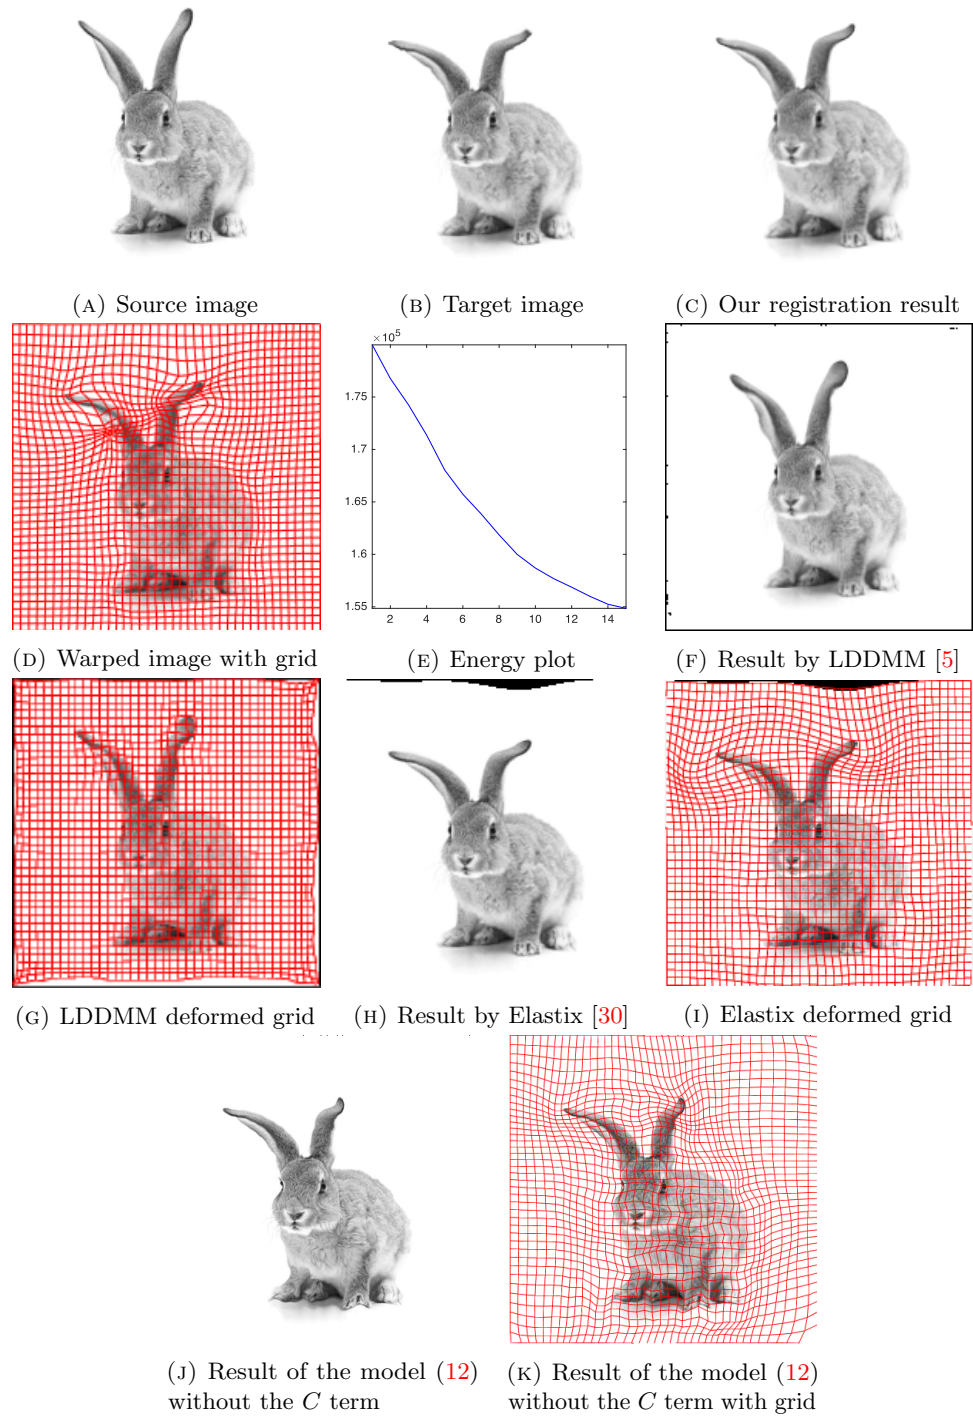


FIGURE 7. The rabbit example.

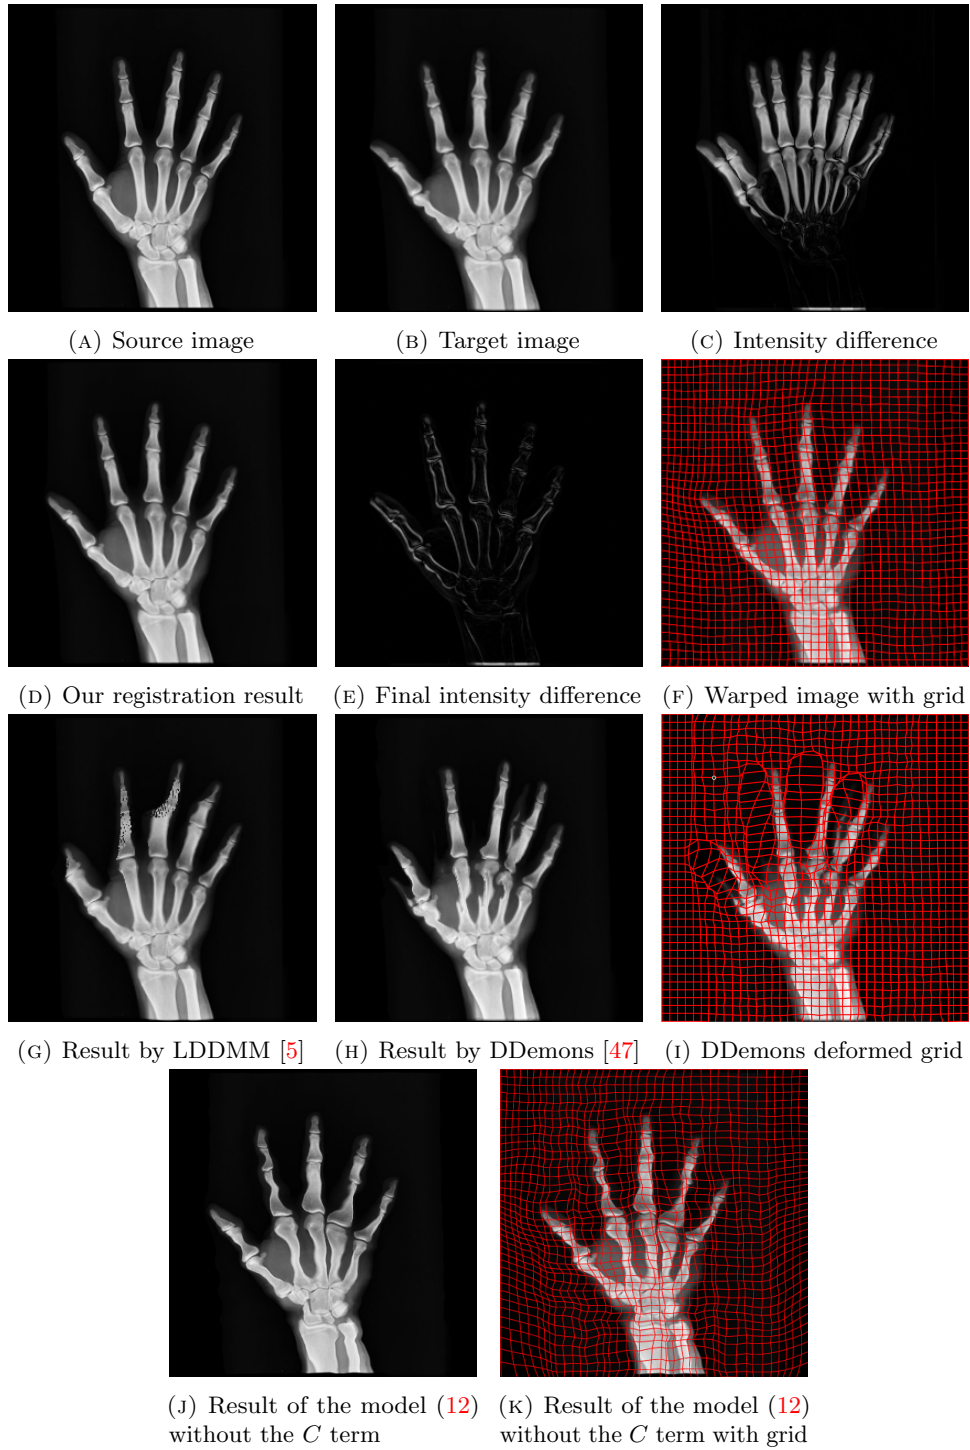


FIGURE 8. The first hand X-ray example.




FIGURE 9. The second hand X-ray example.

and Fig. 9c for their absolute intensity difference). The warped image produced by our proposed method (Fig. 9d) again closely resembles the target image with the intensity difference significantly reduced (see Fig. 9e). Fig. 9f shows that the mapping is smooth and bijective. For comparison, note that LDDMM [5] fails to match the fingers (Fig. 9g). While DROP [19] is capable of registering the fingers (Fig. 9h), it distorts the boundary shape of the overall image (Fig. 9i). The result of using the model (12) without the  $C$  term is also shown in Fig. 8j and Fig. 8k. We can observe that distortion at the index and ring finger is large.

**8.3. Real CT images.** We now consider applying the proposed image registration method on real lung CT images retrieved from the National Lung Screening Trial (NLST) dataset [1]. Fig. 10a and Fig. 10b show two slices of lung CT images that we use as the source and the target (see Fig. 10c for the absolute intensity difference).




FIGURE 10. The lung CT example.

We remark that the CT images are originally with different intensity, and so we apply an intensity histogram matching before running the registration experiment. Fig. 10d shows the registration result obtained by our proposed method. It can be observed that our method successfully produces a large deformation on the right lung of the source image to match that of the target image (see also Fig. 10e for the final absolute intensity difference). On the contrary, DDemons [47] (Fig. 10f), LDDMM [5] (Fig. 10g), Elastix [30] (Fig. 10h) and DROP [19] (Fig. 10i) all fail to produce an accurate and bijective registration result. This shows that our method is more capable of handling large deformation image registration.

We then test our method on slices of chest CT images obtained from the Open Access Biomedical Image Search Engine [2]. Fig. 11a and 11b show the source image and target image respectively, and the intensity difference is shown in Fig. 11c. The

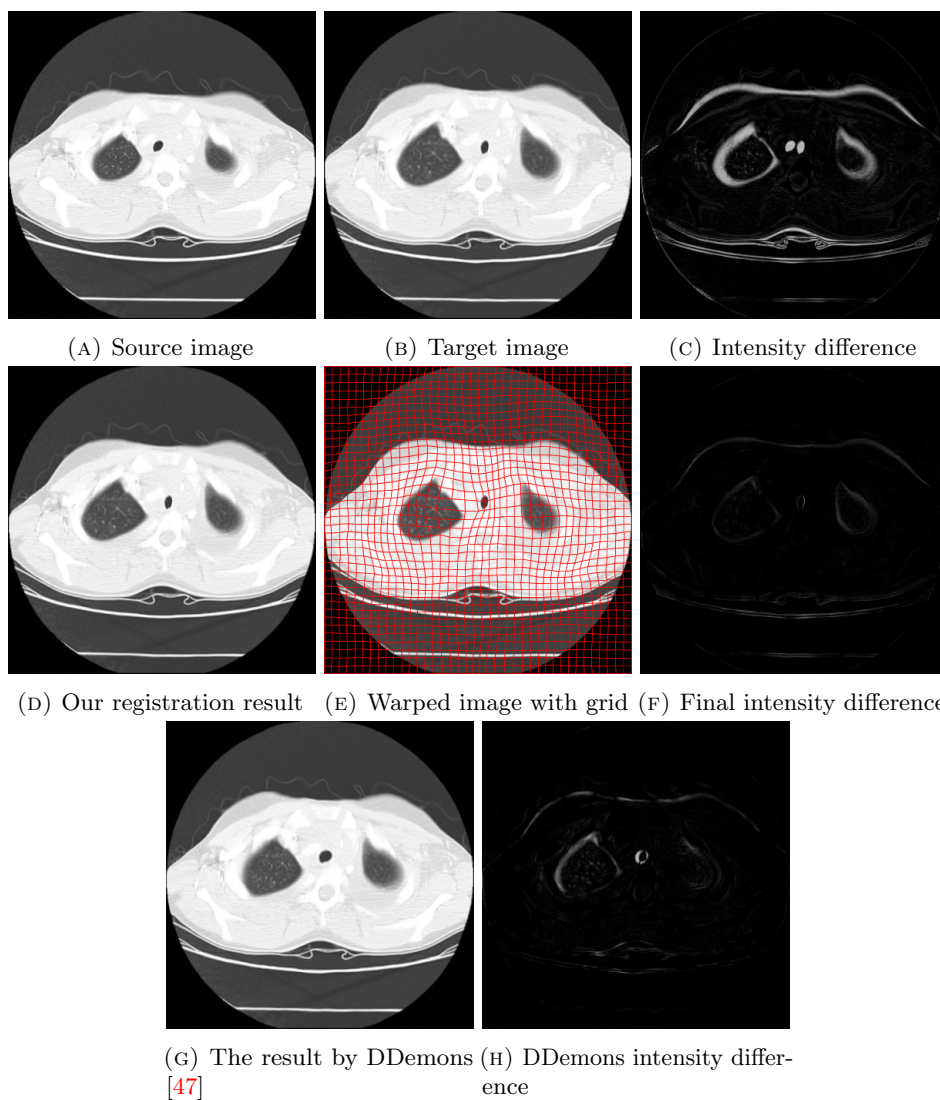


FIGURE 11. The chest CT example.

registration result obtained by our proposed method is shown in Fig. 11d (see also the result with the deformed underlying grid in Fig. 11e). From the final intensity difference plot in Fig. 11f, it is easily to see that our method matches not only the two large components but also the small dot at the center very well. On the contrary, DDemons [47] produces a suboptimal registration result with a significantly larger mismatch of the small component at the center (see Fig. 11g and Fig. 11h).

**8.4. Quantitative comparison between different image registration methods.** For a more quantitative analysis, we compare our proposed image registration method with DDemons [47], LDDMM [5], Elastix [30], and DROP [19] in terms of the accuracy, smoothness, overall quality, and bijectivity. Here, to assess the accuracy of a registration map  $f : \mathcal{M} \rightarrow \mathcal{S}$ , we consider the following similarity

energy

$$E_{\text{sim}}(f) = \frac{\sum \sum |I_{\mathcal{M}} - I_{\mathcal{S}}(f)|}{2} \left( \frac{1}{\sum \sum I_{\mathcal{M}}} + \frac{1}{\sum \sum (1 - I_{\mathcal{M}})} + \frac{1}{\sum \sum I_{\mathcal{S}}(f)} + \frac{1}{\sum \sum (1 - I_{\mathcal{S}}(f))} \right). \quad (37)$$

Note that  $E_{\text{sim}} = 0$  if and only if  $f(\mathcal{M}) = \mathcal{S}$ , i.e. the registration matches the two images exactly. To assess the smoothness of  $f$ , we consider the following smoothness energy

$$E_{\text{smooth}}(f) = \frac{1}{mn} \sqrt{\sum \sum \left( \left( \frac{\partial f}{\partial x} \right)^2 + \left( \frac{\partial f}{\partial y} \right)^2 \right)}, \quad (38)$$

where the size of  $\mathcal{M}$  is  $mn$ . The total energy  $E_{\text{total}}$  takes both the similarity and the smoothness into account for evaluating the overall quality of the registration:

$$E_{\text{total}} = E_{\text{sim}} + E_{\text{smooth}}. \quad (39)$$

Lastly, the bijectivity can be assessed by counting the number of flips in the mapping result. More specifically, we divide every image into  $h \times w$  small squares evenly, and further divide each small square into 2 triangles. After applying each registration method, we count the number of triangles with the orientation flipped. A registration is bijective if and only if there is no triangle flip.

Table 1 records the performance of the above-mentioned methods. It can be observed that our method always gives a bijective registration result, while the other methods may produce flips and hence are non-bijective for various examples. Also, our method achieves the lowest total energy among all methods for most examples. For the few examples that some other methods are with a lower total energy, their results all contain a large number of flips and hence are still undesirable. The experiments suggest that our proposed method is more desirable for large deformation image registration, with the bijectivity guaranteed and a good overall quality in terms of accuracy and smoothness achieved.

Overall, the advantage of our proposed method is attributed to the use of both the CNN features and quasiconformal theory. The use of the CNN features allows us to enforce a “weak” landmark correspondence without manual landmark labeling, and the minimization and truncation of the norm of the Beltrami coefficients ensures the bijectivity of the registration.

**9. Conclusion and future works.** In this work, we have proposed a novel method for large deformation image registration by taking the advantage of the mathematical properties of quasiconformal maps and the robustness of CNNs. With a novel fidelity term based on features extracted using a pre-trained classification CNN included in our proposed energy model, we are able to obtain meaningful descent directions for registering two images without imposing any landmark constraints. We also use quasiconformal theory for regularizing the registration process to ensure the bijectivity and reduce the local geometric distortion of the resulting mappings. Our experiments have shown that the proposed method outperforms the existing methods for the registration of both synthetic images and real medical images. While we currently do not have a convergence guarantee on our optimization algorithm, the existence of the minimizer of the proposed energy  $E_C$  is theoretically guaranteed as proved in Proposition 6.1. Also, from the experiments, the final mapping obtained by the algorithm is still in the set  $\mathcal{A}$ , satisfying the two constraints in Equation (22).



Method	Results ( $E_{sim}$ / $E_{smooth}$ / $E_{total}$ / # Flips)							
	'Z' to '2' (Fig. 5)	Eagle (Fig. 6)	Rabbit (Fig. 7)	Hand 1 (Fig. 8)	Hand 2 (Fig. 9)	Lung (Fig. 10)	Chest (Fig. 11)	
Our method	<b>0.3099</b> 0.4476 <b>0.7575 0</b>	<b>0.0916</b> 0.1558 <b>0.2474 0</b>	0.1953 0.1436 <b>0.3389 0</b>	0.1417 <b>0.1075</b> <b>0.2492 0</b>	<b>0.1317</b> 0.1536 <b>0.2853 0</b>	<b>0.2435</b> 0.4774 0.7210 <b>0</b>	<b>0.0368</b> 0.1414 <b>0.1783 0</b>	
DDemons [47]	2.0660 <b>0.0481</b> 2.1141 121	0.3724 0.2470 0.6194 3697	<b>0.1806</b> 0.1278 <b>0.3083</b> 174	0.4528 0.2651 0.7180 6036	0.4273 0.2718 0.6991 8925	0.6332 <b>0.2271</b> 0.8602 4191	0.0725 0.3565 0.4290 16316	
LDDMM [5]	2.0488 0.2941 2.3629 105	0.3570 0.2388 0.5968 26	0.5973 <b>0.0738</b> 0.6711 1	0.7018 0.1842 <b>0.8860 0</b>	0.9563 0.1902 1.1465 6	0.6843 0.2568 0.9410 209	0.2446 <b>0.0717</b> 0.3164 <b>0</b>	
Elastix [30]	3.6377 0.9185 4.5562 68280	0.1324 0.1779 0.3103 1158	0.1808 0.1679 <b>0.3487 0</b>	<b>0.1103</b> 0.1513 0.2617 <b>0</b>	0.1856 <b>0.1406</b> <b>0.3263 0</b>	0.3432 0.3203 <b>0.6635</b> 3579	0.0530 0.1408 0.1938 <b>0</b>	
DROP [19]	2.0554 0.3084 2.3638 2515	0.2509 <b>0.1518</b> 0.4027 <b>0</b>	0.1788 0.1634 <b>0.3423 0</b>	0.3021 0.1186 0.4207 <b>0</b>	1.2514 <b>0.1406</b> 1.3920 <b>0</b>	0.5065 0.6193 1.1259 3663	0.2415 0.1841 0.4256 <b>0</b>	

TABLE 1. The performance of different image registration methods for various synthetic and real medical images. Here,  $E_{sim}$  measures the accuracy of the registration mapping as described in Equation (37),  $E_{smooth}$  measures the smoothness of the mapping as described in Equation (38),  $E_{total}$  measures the overall quality of the mapping as described in Equation (39), and the number of flips reflects the bijectivity of the mapping. For each example and each measure, the best entry among all methods is in bold.

Based on the quantitative and qualitative assessments on the results, we believe that the final mapping obtained by the algorithm is close to the theoretical minimizer.

A natural next step is to consider improving the performance of the registration using a more data-specific network. Currently, the CNN features are extracted using a common pre-trained classification network so that the proposed model can handle various types of images well. However, in many applications, the users are only interested in registering a specific type of images. For instance, ophthalmologists may only perform image registration on retina scans; pulmonologists may focus on registering lung scans of the patients; and neurologists may be most interested in registering brain MRI scans. By pre-training a more data-specific network, we can better include the prior information of a specific type of images and hence further improve the performance of our registration method for them.

Another possible future direction is to extend the proposed method for other registration problems. Using conformal parameterization algorithms for meshes [9, 11, 14] or point clouds [10, 39], we can effectively flatten two 3D mesh- or point-based surfaces  $\mathcal{M}, \mathcal{S}$  onto the 2D plane with low geometric distortion (denote the mappings by  $\varphi_{\mathcal{M}} : \mathcal{M} \rightarrow \mathbb{C}$  and  $\varphi_{\mathcal{S}} : \mathcal{S} \rightarrow \mathbb{C}$ ). Then, we can devise a similar energy minimization model for obtaining a quasiconformal map  $f$  between the flattened surfaces without imposing any landmark constraints. The composition map  $\varphi_{\mathcal{S}}^{-1} \circ f \circ \varphi_{\mathcal{M}}$  will then give a meaningful registration between the two surfaces  $\mathcal{M}, \mathcal{S}$ . As manual landmark labeling is not needed, such a surface registration method may largely facilitate 3D shape morphometry [8]. With the increasing importance of 3D image registration, it will also be interesting to consider combining 3D quasi-conformal maps [15, 35] and CNN features. These ideas will be further explored and validated in our future work.

REFERENCES

[1] *The National Lung Screening Trial (NLST)*, <https://cdas.cancer.gov/nlst/>.  
 [2] *Open Access Biomedical Image Search Engine*, <https://openi.nlm.nih.gov/>.

- [3] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag and A. V. Dalca, [VoxelMorph: A learning framework for deformable medical image registration](#), *IEEE Transactions on Medical Imaging*, **38** (2019), 1788–1800.
- [4] V. Balntas, E. Johns, L. Tang and K. Mikolajczyk, PN-Net: Conjoined triple deep network for learning local image descriptors, Preprint, [arXiv:1601.05030](#).
- [5] M. F. Beg, M. I. Miller, A. Trounev and L. Younes, [Computing large deformation metric mappings via geodesic flows of diffeomorphisms](#), *International Journal of Computer Vision*, **61** (2005), 139–157.
- [6] F. L. Bookstein, [Principal warps: Thin-plate splines and the decomposition of deformations](#), *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **11** (1989), 567–585.
- [7] L. G. Brown, [A survey of image registration techniques](#), *ACM Computing Surveys (CSUR)*, **24** (1992), 325–376.
- [8] G. P. T. Choi, H. L. Chan, R. Yong, S. Ranjitkar, A. Brook, G. Townsend, K. Chen and L. M. Lui, [Tooth morphometry using quasi-conformal theory](#), *Pattern Recognition*, **99** (2020), 107064.
- [9] G. P. T. Choi, Y. Leung-Liu, X. Gu and L. M. Lui, [Parallelizable global conformal parameterization of simply-connected surfaces via partial welding](#), *SIAM J. Imaging Sci.*, **13** (2020), 1049–1083.
- [10] G. P. T. Choi, Y. Liu and L. M. Lui, [Free-boundary conformal parameterization of point clouds](#), *J. Sci. Comput.*, **90** (2022), Paper No. 14, 26 pp.
- [11] G. P.-T. Choi and L. M. Lui, [A linear formulation for disk conformal parameterization of simply-connected open surfaces](#), *Adv. Comput. Math.*, **44** (2018), 87–114.
- [12] G. P. T. Choi, D. Qiu and L. M. Lui, [Shape analysis via inconsistent surface registration](#), *Proc. A.*, **476** (2020), 20200147, 21 pp.
- [13] P. T. Choi, K. C. Lam and L. M. Lui, [FLASH: Fast landmark aligned spherical harmonic parameterization for genus-0 closed brain surfaces](#), *SIAM J. Imaging Sci.*, **8** (2015), 67–94.
- [14] P. T. Choi and L. M. Lui, [Fast disk conformal parameterization of simply-connected open surfaces](#), *J. Sci. Comput.*, **65** (2015), 1065–1090.
- [15] Z. Daoping and K. Chen, [3D orientation-preserving variational models for accurate image registration](#), *SIAM J. Imaging Sci.*, **13** (2020), 1653–1691.
- [16] B. D. de Vos, F. F. Berendsen, M. A. Viergever, H. Sokooti, M. Staring and I. Išgum, A deep learning framework for unsupervised affine and deformable image registration, *Medical Image Analysis*, **52** (2019), 128–143.
- [17] F. P. Gardiner and N. Lakic, [Quasiconformal Teichmüller Theory](#), 76, American Mathematical Society, 2000.
- [18] B. Glocker, Drop - Deformable registration using discrete optimization, <http://campar.in.tum.de/Main/Drop>.
- [19] B. Glocker, A. Sotiras, N. Komodakis and N. Paragios, [Deformable medical image registration: Setting the state of the art with discrete methods](#), *Annual Review of Biomedical Engineering*, **13** (2011), 219–244.
- [20] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [21] X. Han, T. Leung, Y. Jia, R. Sukthankar and A. C. Berg, MatchNet: Unifying feature and metric learning for patch-based matching, In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2015), 3279–3286.
- [22] K. He, X. Zhang, S. Ren and J. Sun, [Deep residual learning for image recognition](#), In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2016), 770–778.
- [23] B. K. P. Horn and B. G. Schunck, [Determining optical flow](#), In *Artificial Intelligence*, **17** (1981), 185–203.
- [24] G. Huang, Z. Liu, L. Van Der Maaten and K. Q. Weinberger, [Densely connected convolutional networks](#), In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2017), 4700–4708.
- [25] M. Jahrer, M. Grabner and H. Bischof, Learned local descriptors for recognition and matching, In *Computer Vision Winter Workshop*, **2** (2008).
- [26] F. Jia, J. Liu and X.-C. Tai, [A regularized convolutional neural network for semantic image segmentation](#), *Anal. Appl.*, **19** (2021), 147–165.
- [27] H. J. Johnson and G. E. Christensen, [Consistent landmark and intensity-based image registration](#), *IEEE Transactions on Medical Imaging*, **21** (2002), 450–461.

- [28] S. C. Joshi and M. I. Miller, [Landmark matching via large deformation diffeomorphisms](#), *IEEE Transactions on Image Processing*, **9** (2000), 1357–1370.
- [29] S. Klein and M. Staring, [Elastix: A toolbox for rigid and nonrigid registration of images](#), <https://elastix.lumc.nl/>.
- [30] S. Klein, M. Staring, K. Murphy, M. A. Viergever and J. P. Pluim, [Elastix: A toolbox for intensity-based medical image registration](#), *IEEE Transactions on Medical Imaging*, **29** (2010), 196–205.
- [31] D.-J. Kroon, [Multimodality non-rigid demon algorithm image registration](#), MATLAB Central File Exchange, <https://www.mathworks.com/matlabcentral/fileexchange/21451-multimodality-non-rigid-demon-algorithm-image-registration>.
- [32] D. Kuang, [Cycle-consistent training for reducing negative jacobian determinant in deep registration networks](#), In *International Workshop on Simulation and Synthesis in Medical Imaging*, Springer, **11827** (2019), 120–129.
- [33] D. Kuang and T. Schmah, [FAIM—a ConvNet method for unsupervised 3D medical image registration](#), In *International Workshop on Machine Learning in Medical Imaging*, Springer, **11861** (2019), 646–654.
- [34] K. C. Lam and L. M. Lui, [Landmark- and intensity-based registration with large deformations via quasi-conformal maps](#), *SIAM J. Imaging Sci.*, **7** (2014), 2364–2392.
- [35] Y. T. Lee, K. C. Lam and L. M. Lui, [Landmark-matching transformation with large deformation via n-dimensional quasi-conformal maps](#), *J. Sci. Comput.*, **67** (2016), 926–954.
- [36] O. Lehto and K. I. Virtanen, *Quasiconformal Mappings in the Plane*, vol. 126, 2<sup>nd</sup> edition, Springer-Verlag Berlin Heidelberg, 1973.
- [37] H. Lombaert, [Diffeomorphic log demons image registration](#), MATLAB Central File Exchange, <https://www.mathworks.com/matlabcentral/fileexchange/39194-diffeomorphic-log-demons-image-registration>.
- [38] J. B. A. Maintz and M. A. Viergever, [A survey of medical image registration](#), *Medical Image Analysis*, **2** (1998), 1–36.
- [39] T. W. Meng, G. P.-T. Choi and L. M. Lui, [TEMPO: Feature-endowed Teichmüller extremal mappings of point clouds](#), *SIAM J. Imaging Sci.*, **9** (2016), 1922–1962.
- [40] J. Modersitzki, *FAIR: Flexible Algorithms for Image Registration*, SIAM, 2009.
- [41] I. Rocco, R. Arandjelovic and J. Sivic, [Convolutional neural network architecture for geometric matching](#), In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2017), 6148–6157.
- [42] E. Simo-Serra, E. Trulls, L. Ferraz, I. Kokkinos, P. Fua and F. Moreno-Noguer, [Discriminative learning of deep convolutional feature point descriptors](#), In *Proceedings of the IEEE International Conference on Computer Vision*, (2015), 118–126.
- [43] K. Simonyan and A. Zisserman, [Very deep convolutional networks for large-scale image recognition](#), Preprint, [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- [44] S. Sommer, [Segframe image registration](#), <https://github.com/nefan/segframe>.
- [45] A. Sotiras, C. Davatzikos and N. Paragios, [Deformable medical image registration: A survey](#), *IEEE Transactions on Medical Imaging*, **32** (2013), 1153–1190.
- [46] J.-P. Thirion, [Image matching as a diffusion process: An analogy with Maxwell’s demons](#), *Medical Image Analysis*, **2** (1998), 243–260.
- [47] T. Vercauteren, X. Pennec, A. Perchant and N. Ayache, [Diffeomorphic demons: Efficient non-parametric image registration](#), *NeuroImage*, **45** (2009), 61–72.
- [48] H. Wang, L. Dong, J. O’Daniel, R. Mohan, A. S. Garden, K. K. Ang, D. A. Kuban, M. Bonnen, J. Y. Chang and R. Cheung, [Validation of an accelerated ‘Demons’ algorithm for deformable image registration in radiation therapy](#), *Physics in Medicine & Biology*, **50** (2005), 2887.
- [49] X. Yang, R. Kwitt, M. Styner and M. Niethammer, [Quicksilver: Fast predictive image registration—a deep learning approach](#), *NeuroImage*, **158** (2017), 378–396.
- [50] C. P. Yung, G. P. T. Choi, K. Chen and L. M. Lui, [Efficient feature-based image registration by mapping sparsified surfaces](#), *Journal of Visual Communication and Image Representation*, **55** (2018), 561–571.
- [51] S. Zagoruyko and N. Komodakis, [Learning to compare image patches via convolutional neural networks](#), In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (2015), 4353–4361.
- [52] B. Zitova and J. Flusser, [Image registration methods: A survey](#), *Image and Vision Computing*, **21** (2003), 977–1000.

Received June 2021; 1st revision November 2021; 2nd revision January 2022; early access March 2022.

*E-mail address:* [hlaw@gatech.edu](mailto:hlaw@gatech.edu)

*E-mail address:* [ptchoi@mit.edu](mailto:ptchoi@mit.edu)

*E-mail address:* [kachun.lam@nih.gov](mailto:kachun.lam@nih.gov)

*E-mail address:* [lmlui@math.cuhk.edu.hk](mailto:lmlui@math.cuhk.edu.hk)