



**HAL**  
open science

# Learning a Probabilistic Model for Diffeomorphic Registration

Julian Krebs, Hervé Delingette, Boris Mailhé, Nicholas Ayache, Tommaso Mansi

► **To cite this version:**

Julian Krebs, Hervé Delingette, Boris Mailhé, Nicholas Ayache, Tommaso Mansi. Learning a Probabilistic Model for Diffeomorphic Registration. IEEE Transactions on Medical Imaging, 2019, pp.2165-2176. 10.1109/TMI.2019.2897112 . hal-01978339

**HAL Id: hal-01978339**

**<https://hal.science/hal-01978339>**

Submitted on 11 Feb 2019

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

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

# Learning a Probabilistic Model for Diffeomorphic Registration

Julian Krebs, Hervé Delingette, Boris Mailhé, Nicholas Ayache and Tommaso Mansi

**Abstract**—We propose to learn a low-dimensional probabilistic deformation model from data which can be used for registration and the analysis of deformations. The latent variable model maps similar deformations close to each other in an encoding space. It enables to compare deformations, generate normal or pathological deformations for any new image or to transport deformations from one image pair to any other image.

Our unsupervised method is based on variational inference. In particular, we use a conditional variational autoencoder (CVAE) network and constrain transformations to be symmetric and diffeomorphic by applying a differentiable exponentiation layer with a symmetric loss function. We also present a formulation that includes spatial regularization such as diffusion-based filters. Additionally, our framework provides multi-scale velocity field estimations. We evaluated our method on 3-D intra-subject registration using 334 cardiac cine-MRIs. On this dataset, our method showed state-of-the-art performance with a mean DICE score of 81.2% and a mean Hausdorff distance of 7.3mm using 32 latent dimensions compared to three state-of-the-art methods while also demonstrating more regular deformation fields. The average time per registration was 0.32s. Besides, we visualized the learned latent space and show that the encoded deformations can be used to transport deformations and to cluster diseases with a classification accuracy of 83% after applying a linear projection.

**Index Terms**—deformable registration, probabilistic encoding, deep learning, conditional variational autoencoder, latent variable model, and deformation transport.

## I. INTRODUCTION

**D**EFORMABLE image registration, the process of finding voxel correspondences in a pair of images, is an essential task in medical image analysis [1]. This mapping – the deformation field – can be used for example in pre-op / post-op studies, to find the same structures in images from different modalities or to evaluate the progression of a disease. The analysis of geometric changes in successive images is important for instance for diagnosing cardiovascular diseases and selecting the most suited therapies. A possible approach is to register sequential images and analyze the extracted deformations for example by parallel transport [2] or by creating an adapted low-dimensional subspace [3].

We propose a registration algorithm that learns a deformation model directly from training images. Inspired by recent

generative latent variable models, our method learns a low-dimensional probabilistic deformation encoding in an unsupervised fashion. This latent variable space encodes similar deformations close to each other and allows the generation of synthetic deformations for a single image and the comparison and transport of deformations from one case to another.

### A. Deformable Image Registration

Traditionally, deformable registration is solved by numerical optimization of a similarity metric which measures the distance between the fixed and the deformed moving image. The moving image is warped given a predefined deformation model in order to get closer to the fixed image. Unfortunately, this results in an ill-posed problem which requires further regularization based on prior assumptions [1]. Various regularization energies have been proposed including elastic- [4], [5] or diffusion-based methods [6]–[8] (cf. [1]). Diffeomorphic transforms are folding-free and invertible. The enforcement of these properties in many medical applications has led to the wide use of diffeomorphic registration algorithms. Popular parametrizations of diffeomorphisms include the Large Deformation Diffeomorphic Metric Mapping (LDDMM) [9]–[11], a symmetric normalization approach [12] or stationary velocity fields (SVF) [13], [14].

In recent years, learning-based algorithms – notably Deep Learning (DL) – have been proposed to avoid long iterative optimization at test time. In general, one can classify these algorithms as supervised and unsupervised. Due to the difficulty of finding ground truth voxel mappings, supervised methods need to rely on predictions from existing algorithms [15], [16], simulations [17]–[19] or a combination of both [20], [21]. The latter can be achieved for example by projecting *B-spline* displacement estimations in the space of a statistical deformation model from which one can extract simulations by sampling of its components [20]. Diffeomorphic approaches predict patches of the initial momentum of LDDMMs [15] or dense SVFs [16]. Supervised methods are either limited by the performance of existing algorithms or the realism of simulations. Furthermore, retrieving deformations from existing algorithms on a large database is time-consuming and increases the training complexity.

Unsupervised approaches to registration aim to optimize an image similarity, often combined with a penalization or smoothing term (regularization). These learning approaches first appeared in the computer vision community [22], [23] and were recently applied to medical image registration [24]–[28]. Unlike traditional methods, learning-based approaches also

J. Krebs is with the Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, 06902 France and also Siemens Healthineers, Digital Services, Digital Technology and Innovation, Princeton, NJ, 08540 USA.

H. Delingette and N. Ayache are with the Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, 06902 France.

B. Mailhé and T. Mansi are with Siemens Healthineers, Digital Services, Digital Technology and Innovation, Princeton, NJ, 08540 USA.

can include task-specific information such as segmentation labels during training while not requiring those labels at test time. Instead of using an image similarity, Hu et al. [29] proposed to optimize the matching of labels based on a multi-scale DICE loss and a deformation regularization. Fan et al. [26] proposed to jointly optimize a supervised and unsupervised objective by regressing *ground-truth* deformation fields (from an existing algorithm), while simultaneously optimizing an intensity-based similarity criterion. The disadvantage of these *semi-supervised* approaches is that their training complexity is higher since label information needs to be collected, and for example deformations outside the segmented areas are not guaranteed to be captured. Most unsupervised approaches use B-spline grids or dense deformation fields, realized with spatial transformer layers (STN [30]) for an efficient and differentiable linear warping of the moving image. However, it has not been shown yet that these approaches lead to sufficiently regular and plausible deformations.

### B. Deformation Analysis and Transport

Understanding the deformation or motion of an organ goes beyond the registration of successive images. Therefore, it has been proposed to compare and characterize shape and motion patterns by normalizing deformations in a common reference frame [2], [31] and for example by applying statistical methods to study the variation of cardiac shapes [32]. In the diffeomorphic setting, various dimensionality reduction methods have been proposed. Vaillant et al. [33] modeled shape variability by applying PCA in the tangent space to an atlas image. Qiu et al. used a shape prior for surface matching [34]. While these methods are based on probabilistic inference, dimensionality reduction is done after the estimation of diffeomorphisms. Instead Zhang et al. [35] introduced a latent variable model for principle geodesic analysis that estimates a template and principle modes of variation while inferring the latent dimensionality from the data. Instead of having a general deformation model capable of explaining the deformations of any image pair in the training data distribution, this registration approach still depends on the estimation of a smooth template. Using the SVF parametrization for cardiac motion analysis, Rohé et al. [3] proposed to build affine subspaces on a manifold of deformations, the barycentric subspaces, where each point on the manifold represents a 3-D image and the geodesic between two points describes the deformation.

For uncertainty quantification, Wassermann et al. [36] used a probabilistic LDDMM approach applying a stochastic differential equation and Wang et al. [37] employed a low-dimensional Fourier representation of the tangent space of diffeomorphisms with a normal assumption. While both approaches contain probabilistic deformation representations, they have not been used for sampling and the representations have not been learned from a large dataset.

In the framework of diffeomorphic registration, *parallel transport* is a promising normalization method for the comparison of deformations. Currently used *parallel transport* approaches are the Schild's [38] or pole ladder [2], [39] using the SVF parametrization or approaches based on Jacobi fields

using the LDDMM parametrization [40], [41]. In general, these approaches aim to convert and apply the temporal deformation of one subject to another subject. However, this *transport* process typically requires multiple registrations, including difficult registrations between different subjects.

### C. Learning-based Generative Latent Variable Models

Alternatively and inspired by recently introduced learning-based generative models, we propose to learn a latent variable model that captures deformation characteristics just by providing a large dataset of training images. In the computer vision community, such generative models as Generative adversarial networks (GAN) [42], stochastic variational autoencoders (VAE) [43] and adversarial autoencoders (AAE) [44] have demonstrated great performance in learning data distributions from large image training sets. The learned models can be used to generate new synthetic images, similar to the ones seen during training. In addition, probabilistic VAEs are latent variable models which are able to learn continuous latent variables with intractable posterior probability distributions (encoder). Efficient Bayesian inference can be used to deduce the posterior distribution by enforcing the latent variables to follow a predefined (simple) distribution. Finally, a decoder aims to reconstruct the data from that representation [43]. As an extension, conditional variational autoencoders (CVAE) constrain the VAE model on additional information such as labels. This leads to a latent variable space in which similar data points are mapped close to each other. CVAEs are for example used for semi-supervised classification tasks [45].

Generative models also showed promising results in medical imaging applications such as in classifying cardiac diseases [46] or predicting PET-derived myelin content maps from multi-modal MRI [47]. Recently, unsupervised adversarial training approaches have been proposed for image registration [21], [28], [48]. Dalca et al. [27] developed a framework which enforces a multivariate Gaussian distribution on each component of the velocity field for measuring uncertainty. However, these approaches do not learn global latent variable models which map similar deformations close to each other in a probabilistic subspace of deformations. To the best of the authors' knowledge, generative approaches for registration which allow the sampling of new deformations based on a learned low-dimensional encoding have not been proposed yet.

### D. Probabilistic Registration using a Generative Model

We introduce a generative and probabilistic model for diffeomorphic image registration, inspired by generative latent variable models [43], [45]. In contrast to other probabilistic approaches such as [15], [27], we learn a low-dimensional global latent space in an encoder-decoder neural network where the deformation of a new image pair is mapped to and where similar deformations are close to each other. This latent space, learned in an unsupervised fashion, can be used to generate an infinite number of new deformations for any single image from the data distribution and not only for a unique template as in the Bayesian inference procedure for model parameter estimation in [35]. From this abstract representation

of deformations, diffeomorphic deformations are reconstructed by decoding the latent code under the constraint of the moving image. To the best of the author’s knowledge, this method describes the first low-dimensional probabilistic latent variable model that can be used for deformation transport from one subject to another. Through applying a latent deformation code of one image pair on a new constraining image, deformation transport (and sampling from the latent space) is useful for instance for simulating cardiac pathologies or synthesizing a large number of pathological and healthy heart deformations for data augmentation purposes.

We use a variational inference method (a CVAE [45]) with the objective of *reconstructing* the fixed image by warping the moving image. The decoder of the CVAE is conditioned on the moving image to ease the encoding task: by making appearance information easily accessible in the decoder (in the form of the moving image), the latent space is more likely to encode deformation rather than appearance information. This implicit decoupling of deformation and appearance information allows to transport deformations from one case to another by pairing a latent code with a new conditioning image. The framework provides multi-scale estimations where velocities are extracted at each scale of the decoding network. We use the SVF parametrization and diffeomorphisms are extracted using a vector field exponentiation layer, based on the *scaling and squaring* algorithm proposed in [13]. This algorithm has been successfully applied in neural networks in our previous work [49] and in [27]. The framework contains a dense spatial transformer layer (STN) and can be trained end-to-end with a choice of similarity metrics: to avoid asymmetry, we use a symmetric and normalized local cross correlation criterion. In addition, we provide a generic formulation to include regularization terms to control the deformation appearance (if required), such as diffusion regularization in form of Gaussian smoothing [8]. During training, similarity loss terms for each scale and a loss term enforcing a prior assumption on the latent variable distribution are optimized by using the concept of *deep supervision* (cf. [50]). During testing, the low-dimensional latent encoding, multi-scale estimations of velocities, deformation fields and deformed moving image are retrieved in a single forward path of the neural network.

We evaluate our framework on the registration of cardiac MRIs between end-diastole (ED) and end-systole (ES) and provide an intensive analysis on the structure of the latent code and evaluate its application for transporting encoded deformations from one case to another.

This paper extends our preliminary work [49] by adding:

- Detailed derivations of the probabilistic registration framework including a generic regularization model.
- Deep supervision, multi-scale estimations and a normalized loss function to improve registration performances.
- Analysis of size and structure of the latent variable space.
- Evaluation of the deformation transport by comparing it to a state-of-the-art algorithm [2].

## II. METHODS

In image registration, the goal is to find the spatial transformation  $\mathcal{T}_z : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  which is parametrized by a  $d$ -

dimensional vector  $z \in \mathbb{R}^d$ . The optimal values of  $z$  are the ones which best warp the moving image  $M$  in order to match the fixed image  $F$  given the transformation  $\mathcal{T}_z$ . Both images  $F$  and  $M$  are defined in the spatial domain  $\Omega \in \mathbb{R}^3$ . Typically,  $z$  is optimized by minimizing an objective function of the form:  $\arg \min_z \mathcal{F}(z, M, F) = \mathcal{D}(F, M \circ \mathcal{T}_z) + \mathcal{R}(\mathcal{T}_z)$ , where  $\mathcal{D}$  is a metric measuring the similarity between fixed  $F$  and warped moving image  $M \circ \mathcal{T}_z$ .  $\mathcal{R}$  is a spatial regularizer [1]. Recent unsupervised DL-based approaches [22], [24], [25] try to learn to maximize such a similarity metric  $\mathcal{D}$  using stochastic gradient descent methods and a spatial transformer layer (STN [30]) for warping the moving image  $M$ .

In extension, we propose to model registration by learning a probabilistic deformation parametrization vector  $z$  from a set of example image pairs  $(M, F)$ . Thereby, we constrain the low-dimensional  $z$  to follow a prior distribution  $p(z)$ . In other words, our approach contains two key parts: a latent space encoding to model deformations and a decoding function that aims to *reconstruct* the fixed image  $F$  from this encoded transformation – by warping the moving image  $M$ . In addition, this decoding function is generative as it allows to sample new deformations based on  $p(z)$ .

### A. Probabilistic model for multi-scale registration

We assume a generative probabilistic distribution for registration  $p_{true}(F|M)$ , capturing the deformation from  $M$  towards  $F$ . We aim at learning a parameterized model  $p_\theta(F|M)$  with parameters  $\theta$  which allows us to sample new  $F$ ’s that are similar to samples from the unknown distribution  $p_{true}$ . To estimate the posterior  $p_\theta$  we use a latent variable model parameterized by  $z$ . Following the methodology of a VAE [43], we assume the prior  $p(z)$  to be a multivariate unit Gaussian distribution with spherical covariance  $I$ :

$$p(z) \sim \mathcal{N}(0, I). \quad (1)$$

Using multivariate Gaussians offers a closed-form differentiable solution, however,  $p(z)$  could take the form of other distributions. In this work, we parameterize deformation fields  $\phi$  by stationary velocity fields (SVF), denoted by velocities  $v$ :  $\phi = \exp(v)$  [13]. These transformation maps  $\phi$  are given as the sum of identity and displacements  $u(x)$  for every position  $x \in \Omega$ :  $\phi(x) = x + u(x)$ . In the multi-scale approach, we define velocities  $v^s$  at scale  $s \in \mathcal{S}$  where  $\mathcal{S}$  is the set of different image scales ( $s = 1$  describes the original scale for which we omit writing  $s$  and  $s = 2, 3, \dots$  the scale, down-sampled by a factor of  $2^{s-1}$ ). For each scale  $s$ , a family of *decoding* functions  $f_v^s$  is defined, parameterized by a fixed  $\theta^s \subset \theta$  and dependent on  $z$  and the moving image  $M^s$ :

$$v^s = f_v^s(z, M^s; \theta^s). \quad (2)$$

In the training, the goal is to optimize  $\theta^s$  such that all velocities  $v^s$  are likely to lead to warped moving images  $M^{*s}$  that will be similar to  $F^s$  in the training database.  $M^{*s}$  is obtained by exponentiation of  $v^s$  and warping of the moving image. Using Eq. 2, we can define the families of functions  $f^s$ :

$$M^{*s} := f^s(z, M^s; \theta^s) = M^s \circ \exp(f_v^s(z, M^s; \theta^s)). \quad (3)$$






Fig. 1. (a) Generative process for registration representing the likelihood of the fixed image  $F$  given the latent variable vector  $z$  and moving  $M$ :  $p_\theta(F|z, M)$ , where  $\omega$  and  $\theta$  are fixed parameters. (b) Generative process for regularized image registration where the likelihood depends on the regularized velocities  $p_\theta(F|v^*, M)$ .

In order to express the dependency of  $f^s$  on  $z$  and  $M^s$  explicitly, we can define a distribution  $p(F^s|z, M^s; \theta^s)$ . The product over the different scales gives us the output distribution:

$$p_\theta(F|z, M) = \prod_{s \in \mathcal{S}} p(F^s|z, M^s; \theta^s). \quad (4)$$

By using the law of total probability, this leads to the following stochastic process for computing  $p_\theta(F|M)$  which is also visualized in Fig. 1a (cf. [45]):

$$p_\theta(F|M) = \int_z p_\theta(F|z, M)p(z)dz. \quad (5)$$

The likelihood  $p_\theta(F|z, M)$  can be any distribution that is computable and continuous in  $\theta$ . In VAEs, the choice is often Gaussian, which is equivalent to adopting a sum-of-squared differences (SSD) criterion (cf. [43]). We propose instead to use a local cross-correlation (LCC) distribution due to its robustness properties and superior results in image registration compared to SSD (cf. [8], [51]). Thus, we use the following Boltzmann distribution as likelihood:

$$p_\theta^s(F^s|z, M^s) \sim \exp(-\lambda \mathcal{D}_{LCC}(F^s, M^s, v^s)), \quad (6)$$

where  $v^s = f_v^s(z, M^s; \theta^s)$  are the velocities and  $\lambda$  is a scalar hyperparameter. The symmetric  $\mathcal{D}_{LCC}$  is defined as:

$$\mathcal{D}_{LCC}(F^s, M^s, v^s) = \frac{1}{P} \sum_{x \in \Omega} \frac{\sum_i ((F_{x_i}^{*s} - \bar{F}_x^{*s})(M_{x_i}^{*s} - \bar{M}_x^{*s}))^2}{\left(\sum_i (F_{x_i}^{*s} - \bar{F}_x^{*s})^2\right) \left(\sum_i (M_{x_i}^{*s} - \bar{M}_x^{*s})^2\right) + \tau} - 1, \quad (7)$$

with  $P$  pixels  $x \in \Omega$ , the symmetrically warped images  $M^{*s} = M^s \circ \exp(v^s/2)$  and  $F^{*s} = F^s \circ \exp(-v^s/2)$ . The bar  $\bar{F}_x$  symbolizes the local mean grey levels of  $F_x$  derived by mean filtering with kernel size  $k$  at position  $x$ .  $i$  is iterating through this  $k \times k$ -window. A small constant  $\tau$  is added for numerical stability ( $\tau = 1e^{-15}$ ).

1) *Learning the constrained deformation encoding*: In order to optimize the parameterized model over  $\theta$  (Eq. 5), two problems must be solved: First, how to define the latent variables  $z$ , for example decide what information these variables

represent. VAEs assume there is no simple interpretation of the dimensions of  $z$  but instead assert that samples of  $z$  are drawn from a simple distribution  $p(z)$ .

Second, the integral over  $z$  is intractable since one would need to sample a too large number of  $z$ 's to get an accurate estimate of  $p_\theta(F|M)$ . Instead of sampling a large number of  $z$ 's, the key assumption behind VAEs is to sample only  $z$ 's that are likely to have produced  $F$  and compute  $p_\theta(F|M)$  only from those. To this end, one needs to compute the intractable posterior  $p(z|F, M)$ . Due to this intractability, in VAEs [43], the posterior is approximated by learning an *encoding* distribution  $q_\omega(z|F, M)$ , using a neural network with parameters  $\omega$  (the encoder). This approximated distribution can be related to the true posterior using the Kullback-Leibler divergence (KL) which leads (after rearranging the terms) to the evidence lower bound (ELBO) of the log marginalized likelihood  $\log p_\theta(F|M)$  (cf. [43], [45]):

$$\log p_\theta(F|M) - \text{KL}[q_\omega(z|F, M)||p(z|F, M)] = \mathbb{E}_{z \sim q} [\log p_\theta(F|z, M)] - \text{KL}[q_\omega(z|F, M)||p(z)]. \quad (8)$$

The KL-divergence on the left hand side gets smaller the better  $q_\omega(z|F, M)$  approximates  $p(z|F, M)$  and ideally vanishes if  $q_\omega$  is of enough capacity. Thus, maximizing  $\log p_\theta(F|M)$  is equivalent to maximizing the ELBO on the right hand side of the equation consisting of encoder  $q_\omega$  and decoder  $p_\theta$  which can be both optimized via stochastic gradient descent.

2) *Optimizing the ELBO*: According to the right-hand side of Eq. 8, there are two terms to optimize, the KL-divergence of prior  $p(z)$  and encoder distribution  $q_\omega(z|F, M)$  and the expectation of the reconstruction term  $\log p_\theta(F|z, M)$ . Since the prior is a multivariate Gaussian, the encoder distribution is defined as  $q_\omega(z|F, M) = \mathcal{N}(z|\mu_\omega(F, M), \Sigma_\omega(F, M))$ , where  $\mu_\omega$  and  $\Sigma_\omega$  are deterministic functions learned in an encoder neural network with parameters  $\omega$ . The KL-term can be computed in closed form as follows (constraining  $\Sigma_\omega$  to be diagonal):

$$\text{KL}[\mathcal{N}(\mu_\omega(F, M), \Sigma_\omega(F, M))||\mathcal{N}(0, I)] = \frac{1}{2} \left( \text{tr}(\Sigma_\omega(F, M)) + \|\mu_\omega(F, M)\| - k - \log \det(\Sigma_\omega(F, M)) \right),$$

where  $k$  is the dimensionality of the distribution.

The expected log-likelihood  $\mathbb{E}_{z \sim q} [\log p_\theta(F|z, M)]$ , the reconstruction term, could be estimated by using many samples of  $z$ . To save computations, we treat  $p_\theta(F|z, M)$  as  $\mathbb{E}_{z \sim q} [\log p_\theta(F|z, M)]$  by only taking one sample of  $z$ . This can be justified as optimization is already done via stochastic gradient descent, where we sample many image pairs  $(F, M)$  from the dataset  $\mathcal{X}$  and thus witness different values for  $z$ . This can be formalized with the expectation over  $F, M \sim \mathcal{X}$ :

$$\mathbb{E}_{F, M \sim \mathcal{X}} \left[ \mathbb{E}_{z \sim q} [\log p_\theta(F|z, M)] - \text{KL}[q_\omega(z|F, M)||p(z)] \right].$$

To enable back-propagation through the sampling operation  $q_\omega(z|F, M)$ , the *reparametrization* trick [43] is used in practice, where  $z = \mu_\omega + \epsilon \Sigma_\omega^{1/2}$  (with  $\epsilon \sim \mathcal{N}(0, I)$ ). Thus,

for image pairs  $(F, M)$  from a training dataset  $\mathcal{X}$  the actual objective becomes:

$$\mathbb{E}_{F, M \sim \mathcal{X}} \left[ \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} \left[ \log p_\theta(F|z = \mu_\omega(F, M) + \Sigma_\omega^{1/2}(F, M) * \epsilon, M) \right] - \text{KL}[q_\omega(z|F, M)||p(z)] \right]. \quad (9)$$

After insertion of Eq. 4, the log of the product over the scales  $s \in \mathcal{S}$  results in the sum of the log-likelihood distributions:

$$\mathbb{E}_{F, M \sim \mathcal{X}} \left[ \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} \left[ \sum_{s \in \mathcal{S}} \log p_{\theta^s}(F^s|z = \mu_\omega(F, M) + \Sigma_\omega^{1/2}(F, M) * \epsilon, M^s) \right] - \text{KL}[q_\omega(z|F, M)||p(z)] \right]. \quad (10)$$

### B. Introducing regularization on velocities

So far, we have considered that at each scale  $s$ , a velocity field  $v^s$  is generated by a decoding function  $f_v^s(z, M^s; \theta^s)$  through a neural network. To have a better control of its smoothness, we propose to regularize spatially  $v^s$  through a Gaussian convolution with standard deviation  $\sigma_G$ :

$$\hat{v}^s = G_{\sigma_G} * v^s \quad (11)$$

Gaussian smoothing was applied here, but it could be replaced by any quadratic Tikhonov regularizer or by any functional enforcing prior knowledge about the velocity field.

In the remainder, we show how the regularization of velocities can be inserted into the proposed probabilistic framework. To make the notation less cluttered, we drop the scale  $s$  superscript in the velocity notations. Until now, the velocities  $v$  have been handled as fixed parameters  $v = f_v(z, M; \theta)$ . We can equivalently assume that velocities  $v$  are random variables with a Dirac posterior probability:  $p_\theta(v|z, M) \sim \delta_{f_v(z, M; \theta)}(v)$ .

We now introduce the random variable  $v^*$  which represents the regularized velocities as shown in Fig. 1b. This quantity is linked to the regular velocities  $v$  through a Gaussian distribution  $p(v|v^*) = G(v^*, 1)$  such that  $v$  is close to  $v^*$  in terms of  $L_2$  norm. Furthermore, we define a diffusion-like regularization prior on  $v^*$  [52]:

$$\log p(v^*) \propto \int_{\Omega} \sum_{i=1}^{\infty} \frac{\sigma_G^{2i}}{2^i i!} \left[ \frac{\partial^i}{\partial \Omega^i} v^* \right]^2 d\Omega,$$

taking into account the Taylor expansion of the Fourier transform of the Gaussian. The maximum *a posteriori* of the regularized velocities  $\hat{v}$  is then obtained through Bayes law:

$$\hat{v} = \arg \max_{v^*} \log p(v^*|v) = \arg \max_{v^*} \log p(v|v^*) + \log p(v^*)$$

which in this case is equivalent to solving the Heat equation [52] and leads to a Gaussian convolution:  $\hat{v} = G_{\sigma_G} * v$ .

Finally, we conveniently assume that the posterior probability of  $v^*$  is infinitely peaked around its mode, *i.e.*  $p(v^*|v) \sim \delta_{\hat{v}}(v^*)$  (assumption sometimes made for the Expectation-Maximization algorithm [53]). In the decoding process, we

can now marginalize out the velocity variables  $v$  and  $v^*$  by integrating over both such that only  $\hat{v}$  remains:

$$\begin{aligned} p_\theta(F|M) &= \int_z \int_v \int_{v^*} p(F|v^*, M) p(v^*|v) p_\theta(v|z, M) p(z) \\ &\quad dv dv^* dz \\ &= \int_z p(F|\hat{v}, M) p(z) dz. \end{aligned} \quad (12)$$

Thus, the proposed graphical model leads to a decoder working with the regularized velocity field  $\hat{v}$  instead of the  $v$  generated by the neural network. When combining regularized velocities  $\hat{v}^s$  at all scales, we get:

$$p_\theta(F|M) = \int_z \prod_{s \in \mathcal{S}} p(F^s|\hat{v}^s, M^s) p(z) dz. \quad (13)$$

This can be optimized as before and leads to Gaussian convolutions at each scale if considering diffusion-like regularization. Thus, the multi-scale loss function per training image pair  $(F, M)$  for one sample  $\epsilon$  is defined as (cf. Eq. 10):

$$\begin{aligned} \arg \min_{\omega, \theta} \frac{1}{2} \left( \text{tr}(\Sigma_\omega) + \mu_\omega^\top \mu_\omega - k - \log \det(\Sigma_\omega) \right) \\ - \lambda \sum_{s \in \mathcal{S}} \mathcal{D}_{LCC}(F^s, M^s, \hat{v}^s), \end{aligned} \quad (14)$$

where  $\hat{v}^s$  depends on  $v^s$  and therefore on  $\theta$  (cf. Eq. 2 and 11).

### C. Network architecture

The encoder-decoder neural network takes the moving and the fixed image as input and outputs the latent code  $z$ , velocities  $v$ , the deformation field  $\phi$  and the warped moving image  $M^*$ . The last three are returned at the different scales  $s$ . The encoder consists of strided convolutions while the bottleneck layers  $(\mu, \sigma, z)$  are fully-connected. The deconvolution layers in the decoder were conditioned by concatenating each layer's output with sub-sampled versions of  $M$ . Making appearance information of the moving image easily accessible for the decoder, allows the network to focus on deformation information – the differences between moving and fixed image – that need to pass through the latent bottleneck. While it is not guaranteed that the latent representation contains any appearance information, it comes at a cost to use the *small* bottleneck for appearance information. At each decoding scale, a convolutional layer reduces the number of filter maps to three. Then, a Gaussian smoothing layer (cf. Eq. 11) with variance  $\sigma_G^2$  is applied on these filter maps. The resulting velocities  $v^s$  (a SVF) are exponentiated by the *scaling and squaring* layer [49] in order to retrieve the diffeomorphism  $\phi^s$  which is used by a dense STN to retrieve the warped image  $M^{*s}$ . The latent code  $z$  is computed according to the reparametrization trick. During training, the network parameters are updated through back-propagation of the gradients with respect to the objective Eq. 10, defined at each multi-scale output. Finally during testing, registration is done in a single forward path where  $z$  is set to  $\mu$  since we want to execute registration deterministically. One can also think of drawing several  $z$  using  $\sigma$  and use the different outputs for uncertainty estimation as in [27] which we do not further pursue in this




Fig. 2. (a) Probabilistic multi-scale registration network based on a CVAE. An encoder maps deformations to latent variables  $z \in \mathbb{R}^d$  (with for example  $d = 32$ ) from which a decoder extracts velocities and diffeomorphisms at different scales while being conditioned on the moving image  $M$ . (b) After training, the decoder network can be also used to sample and transport deformations: Apply  $z$ -code on any new image  $M$ .

work. The network architecture can be seen in Fig. 2a. Besides registration, the trained probabilistic framework can be also used for the sampling of deformations as shown in Fig. 2b.

### III. EXPERIMENTS

We evaluate our framework on cardiac intra-subject registration. End-diastole (ED) frames are registered to end-systole (ES) frames from cine-MRI of healthy and pathological subjects. These images show large deformations. Additionally, we evaluate the learned encoding of deformations by visualizing the latent space and transporting encoded deformations from one patient to another. All experiments are in 3-D.

1) *Data*: We used the 334 ED-ES frame pairs of short-axis cine-MRI sequences. 184 cases were acquired from different hospitals and 150 cases were used from the Automatic Cardiac Diagnosis Challenge (ACDC) at STACOM 2017 [54], mixing congenital heart diseases with images from adults. We used 234 cases for training and for testing 100 cases from ACDC, that contain segmentation and disease class information. The testing set contained 20 cases of each of the following cardiac diseases: dilated cardiomyopathy (DCM), hypertrophic cardiomyopathy (HCM), previous myocardial infarction (MINF), abnormal right ventricle (RV) and healthy (Normal). All images were resampled with a spacing of  $1.5 \times 1.5 \times 3.15$  mm and cropped to a size of  $128 \times 128 \times 32$  voxels, by equally removing voxels from all sides. These dimensions were chosen to save computation time and are not a limitation of the framework.

2) *Implementation details*: Our neural network consisted of four encoding convolutional layers with strides (2, 2, 2, 1) and three decoding deconvolutional layers. Each scale contained two convolutional layers and a convolutional Gaussian layer with  $\sigma_G = 3mm$  (kernel size 15) in front of an exponentiation

and a spatial transformer layer using trilinear interpolation (cf. Fig. 2a). The dimensionality of the latent code  $z$  was set to  $d = 32$  as a compromise of registration quality and generalizability (cf. experiment on latent vector dimensionality). The number of trainable parameters in the network was  $\sim 420k$ . LeakyReLU activation functions and L2 weight decay of  $1 * 10^{-4}$  were applied on all layers except the last convolutional layer in each scale where a  $\tanh$  activation function was used in order to avoid extreme velocity values during training. All scales were trained together, using linearly down-sampled versions of the input images for the coarser scales. In all experiments, the number of iterations in the exponentiation layer was set to  $N = 4$  (evaluated on a few training samples according to the formula in [13]). During the training, the mean filter size of the LCC criterion was  $k = 9$ . The loss hyper parameter was empirically chosen as  $\lambda = 5000$  such that the similarity loss was optimized while the latent codes roughly had zero means and variances of one. We applied a learning rate of  $1.5 * 10^{-4}$  with the Adam optimizer and a batch size of one. For augmentation purposes, training image were randomly shifted, rotated, scaled and mirrored. The framework has been implemented in *Tensorflow* using *Keras*<sup>1</sup>. Training took  $\sim 24$  hours and testing a single registration case took 0.32s on a *NVIDIA GTX TITAN X* GPU.

3) *Registration*: We compare our approach with the LCC-demons (Dem, [8]) and the ANTs software package using Symmetric Normalization (SyN, [12]) with manually tuned parameters (on a few training images) and the diffeomorphic DL-based method VoxelMorph [27] (VM) which has been trained using the same augmentation techniques as our algorithm. For the latter, we set  $\sigma = 0.05$ ,  $\lambda = 50000$  and applied a reduced

<sup>1</sup><https://keras.io/>




Fig. 3. Boxplots of registration results comparing the undeformed (Undef) case to the different algorithms: lcc-demons (Dem), SyN, voxelmorph (VM) and our single scale (S1) respectively multi-scale (S3) using RMSE, gradient of the determinant of the Jacobian, DICE scores (logit-transform) and Hausdorff distances (HD in mm). Mean values are denoted by red bars. Higher values are better.

learning rate of  $5 \times 10^{-5}$  for stability reasons while using more training epochs. Higher values for  $\lambda$  led to worse registration accuracy. We also show the improvement of using a multi-scale approach (with 3 scales, S3) compared to a single-scale objective (S1).

We measure registration performance with the following surrogates: intensity root mean square error (RMSE), DICE score, 95%-tile Hausdorff distance (HD in mm). To quantify deformation regularity, we show the determinant of the Jacobian qualitatively, while we also computed the mean magnitude of the gradients of the determinant of the Jacobian (Grad Det-Jac). We decided to report this second-order description of deformations to better quantify differences in smoothness among the different methods, which are not obvious by taking the mean of the determinant of the Jacobian as bigger and smaller values tend to cancel each other out. DICE and HD scores were evaluated on the following anatomical structures: myocardium (LV-Myo) and epicardium (LV) of the left ventricle, left bloodpool (LV-BP), right ventricle (RV) and LV+RV (Fig. 4).




Fig. 4. Cardiac structures used only for measuring registration accuracy.

Table I shows the mean results and standard deviations of all algorithms. In terms of DICE scores, our algorithm using three scales (Our S3) shows the best performances on this dataset while the single-scale version (Our S1) performed similarly compared to the LCC-demons and the SyN algorithm. Hausdorff distances were significantly improved using both of our algorithms. Detailed registration results are shown in Fig. 3. Interestingly, we found that the SyN algorithm showed marginally better DICE scores than the LCC-demons which has been also reported on brain data [8].

Qualitative registration results of a pathological (HCM) and a healthy case (Normal) are presented in Fig. 5a<sup>2</sup>. The warped moving image (with and without grid overlay) and the determinant of the Jacobian (Det. Jac.) are shown. Displacements are visualized using the color encoding as typical for the optical

flow in computer vision tasks. Middle and coarse scale outputs of our multi-scale method are shown in Fig. 5b. We computed the determinant of the Jacobian using *SimpleITK*<sup>3</sup> and found that for all methods no negative values were observed on our test dataset. Compared to the other algorithms, our approach produced smoother and more regular deformations as qualitatively shown by the determinant of the Jacobian in Fig. 5a and quantitatively by the significantly smaller mean gradients of the determinant of the Jacobian (Table I)<sup>4</sup>. Despite the fact of being diffeomorphic, the voxelmorph algorithm produced more irregular deformation fields compared to all other algorithms. Our single-scale approach resulted in slightly smoother deformations which is probably due to the fact that it performed less accurately in compensating large deformations.

TABLE I  
REGISTRATION PERFORMANCE WITH MEAN AND STANDARD DEVIATION SCORES (IN BRACKETS) OF RMSE, DICE, HAUSDORFF DISTANCE (HD IN MM) AND THE MEAN GRADIENT OF THE DETERMINANT OF JACOBIANS (GRAD DET-JAC,  $\times 10^{-2}$ ) COMPARING THE UNDEFORMED CASE (UNDEF), LCC-DEMONS (DEM), SYN, VOXELMORPH (VM) AND OUR METHOD.

Method	RMSE	DICE	HD	Grad Det-Jac
Undef	0.37 (0.17)	0.707 (0.145)	10.1 (2.2)	–
Dem	0.29 (0.16)	0.799 (0.096)	8.3 (2.7)	2.9 (1.0)
SyN	0.32 (0.16)	0.801 (0.091)	8.1 (3.6)	3.4 (0.5)
VM	<b>0.24 (0.08)</b>	0.790 (0.096)	8.4 (2.6)	9.2 (0.5)
Our S1	0.31 (0.15)	0.797 (0.093)	7.9 (2.6)	<b>1.2 (0.3)</b>
Our S3	0.30 (0.14)	<b>0.812 (0.085)</b>	<b>7.3 (2.7)</b>	1.4 (0.3)

We applied the Wilcoxon signed-rank test with  $p < 0.001$  to evaluate statistical significance of the differences in the results of Fig. 3. This method is chosen as a paired test without the assumption of normal distributions. For all metrics, the results of our multi-scale algorithm (Our S3) showed significant differences compared to the results of all other methods (including Our S1). With respect to our single-scale algorithm (Our S1), only the differences in DICE scores were not statistically significant in comparison with the LCC-demons (Dem).

<sup>3</sup><http://www.simpleitk.org/>

<sup>4</sup>Visualization of the gradients of the Jacobian determinant are presented in Fig. 13 (available in the supplementary files /multimedia tab).

<sup>2</sup>Qualitative registration results for all five diseases are also presented in Fig. 10-12 (available in the supplementary files /multimedia tab).






Fig. 5. (a) Qualitative registration results showing a pathological (hypertrophy) and a normal case. Warped moving image  $M^*$ , displacements  $u$ , warped moving image with grid overlay and Jacobian determinant are shown for LCC-demons (Dem), SyN, voxelmorph (VM) and our approach using 3 scales (Our S3). (b) Middle and coarse scale predictions of our multi-scale method (Our S3).

Note, that higher DICE and HD scores can be achieved by choosing a higher latent dimensionality (cf. Experiment III-4a), which however comes at the cost of a more complex encoding space, making analysis tasks more difficult. We also tested the first version of voxelmorph [25] on our dataset. We chose to show the results of the latest version [27] due to the fact that this version is diffeomorphic and that its DICE and HD results were better (cf. [49]).

4) *Deformation encoding*: For evaluating the learned latent space, we investigated (a) the effects of the size of the latent vector on the registration accuracy, (b) the structure of the encoded space by visualizing the distribution of cardiac

diseases and showing simulated deformations along the two main axes of variations and (c) we applied our framework on deformation transport and compare its performance with a state-of-the-art algorithm.

a) *Latent Vector Size*: In Fig. 6 we analyzed the influence of the size of the latent code vector with respect to registration accuracy in terms of DICE and HD scores. With a relatively small latent size of  $d = 8$ , competitive accuracy is achieved. With an increasing dimensionality, performance increases but reaches a plateau eventually. This behavior is expected, since CVAEs tend to ignore components if the dimensionality of the latent space is too high [45]. For the cardiac use case, we

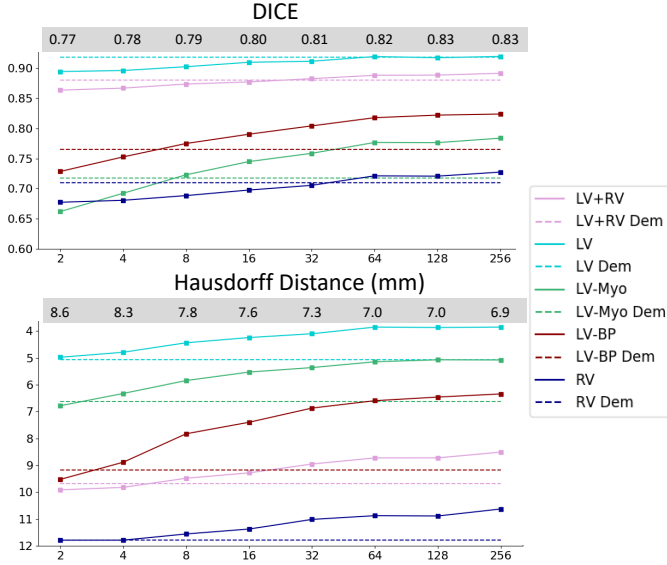


Fig. 6. Showing the influence of the latent vector size  $d$  on the registration accuracy in terms of DICE and Hausdorff distances in mm of the different anatomical structures with the mean of all structures shown in the grey boxes. The performance of the LCC-demons (Dem) is shown as reference with dashed lines.




Fig. 7. Cardiac disease distribution after projecting the latent variables  $z$  of the test images on a 2-D CCA (canonical correlation analysis) space. Using an 8-D CCA and applying SVM with 10-fold cross-validation leads to a classification accuracy of 83%

chose  $d = 32$  components as a trade-off between accuracy and latent variable size.

#### b) Disease Distribution and Generative Latent Space:

In this experiment, we used disease information and encoded  $z$ -codes of the test images to visualize the learned latent space. Using linear CCA (canonical correlation analysis), we projected the  $z$ -codes (32-D) to a 2-D space by using the two most discriminative CCA components. Fig. 7 shows that the 100 test sets are clustered by classes in this space. Taking the 8 most discriminative CCA components into account, the classification accuracy of the five classes is **83%** with 10-fold cross-validation using support vector machines (SVM). In a second experiment, we applied principal component analysis (PCA) on the  $z$ -codes of the training dataset. We simulated deformations by sampling equally distributed values in the

range of  $\pm 2.5$  standard deviations of the two largest principal components and extracting the  $z$ -codes through inverse projections. Fig. 8 shows reconstructed displacements and deformed images when applying these generated  $z$ -codes on a random test image. One can see the different influences of the two eigenvalues. The first eigenvalue (horizontal) focuses on large deformations while the second one focuses on smaller deformations as the right ventricle. The results of these two experiments which are solely based on applying simple linear transformations, suggest that deformations that are mapped close to each other in the deformation latent space have similar characteristics.

c) *Deformation Transport*: Pathological deformations can be transported to healthy subjects by deforming the healthy ES frames using pathological ED-ES deformations. Our framework allows for deformation transport by first registering the ED-ES frames of a given pathological case (Step 1), which we call *prediction* in this experiment. Secondly, we use the  $z$ -codes from these predictions along with the ED frame of a healthy subject to transport the deformations (Step 2, cf. Fig. 2b). Note, that this procedure does not require any inter-subject registrations.

We compare our approach with the pole-ladder algorithm (PL [2]). All intra- and inter-subject registrations required by the pole ladder were performed using the LCC-demons [8]. For the inter-subject pairs, we aligned the test data with respect to the center of mass of the provided segmentation and rotated the images manually for rigid alignment. This alignment step was done only for the pole ladder experiment<sup>5</sup>.

Qualitative results are shown in Fig. 9 where the predicted deformations of one hypertrophy (A, HCM) and one cardiomyopathy (B, DCM) case (step 1) were transported to two healthy (Normal) subjects (step 2, targets C and D). Note that our algorithm automatically determines orientation and location of the heart. In Table II, we evaluated the average ejection fraction (EF) of the ED-ES deformation prediction of the pathologies (step 1) and the average EF after transport to normal subjects (step 2). Hereby, we assume that EFs, as a relative measure, stay similar after successful transport (such that the absolute difference, EF step 1 - EF step 2, is small). The table shows the average of transporting 5 HCM

TABLE II  
MEAN EJECTION FRACTION (EF IN % WITH STANDARD DEVIATION IN PARENTHESES) OF PATHOLOGICAL DEFORMATION PREDICTIONS (STEP 1) SHOULD STAY SIMILAR TO THE MEAN EF AFTER THE TRANSPORT TO HEALTHY/NORMAL SUBJECTS (STEP 2). OUR ALGORITHM SHOWS SMALLER ABSOLUTE DIFFERENCES COMPARED TO THE POLE LADDER (PL).

Path.	Step 1: Prediction		Step 2: Transport		Difference	
	PL (Dem)	Our	PL (Dem)	Our	PL	Our
HCM	29.4 (6)	44.1 (7)	35.5 (8)	38.6 (13)	6.1	<b>5.5</b>
DCM	10.8 (2)	12.7 (7)	16.7 (4)	13.9 (7)	5.9	<b>1.2</b>

and 5 DCM cases to 20 normal cases (200 transports). For our algorithm, the absolute differences in EFs are much smaller for

<sup>5</sup>The pipeline for the parallel transport experiment using the pole ladder algorithm is presented in Fig. 14 (available in the supplementary files /multimedia tab).






Fig. 8. Reconstruction of simulated displacements and an accordingly warped random test image after generating  $z$ -codes by equally sampling along the two largest principal components within a range of  $\pm 2.5$  sigma around their mean values (red box). The PCA was fitted using all training  $z$ -codes. The blue box indicates the image closest to the identity deformation. One can see that the horizontal eigenvalue influences large deformations while the vertical eigenvalue focuses on smaller ones, for example the right ventricle.

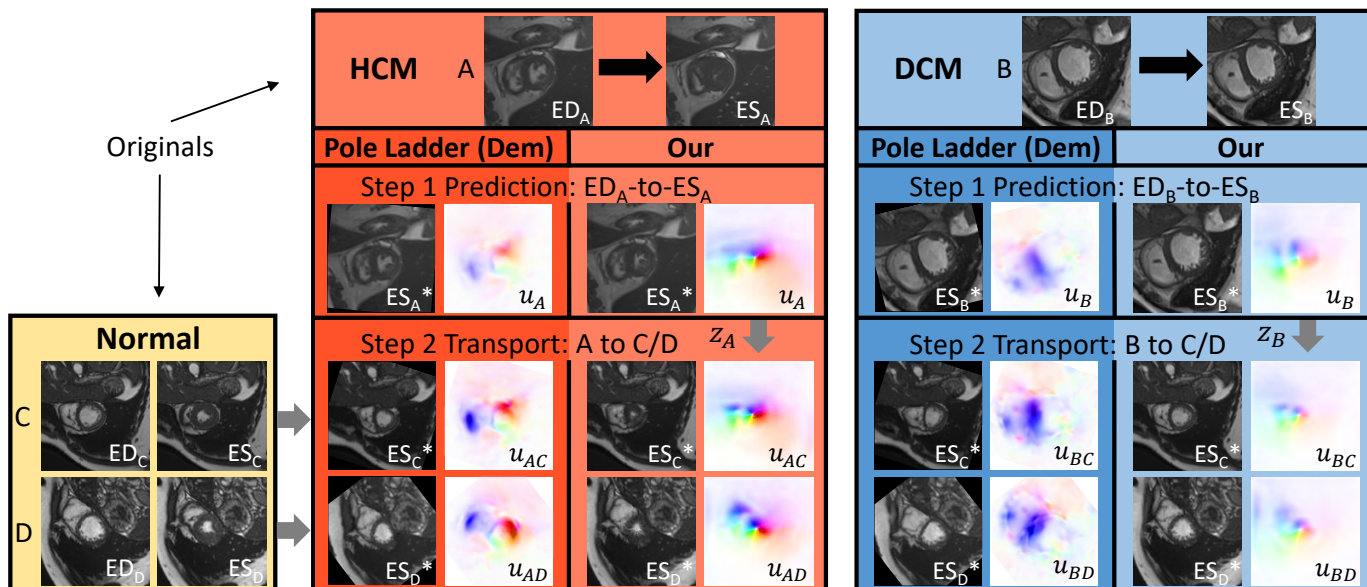


Fig. 9. Transport pathological deformation predictions (Step 1, hypertrophy HCM, myopathy DCM) to healthy (Normal) subjects by using the pole ladder (with LCC-demons) and our probabilistic method (Step 2). Note that the pole ladder algorithm requires the registration between pathological and normal subjects while our approach is able to rotate and translate deformations encoded in the latent space  $z$  automatically.

DCM cases and similarly close in HCM cases in comparison to the pole ladder. All test subjects were not used during training. The EF is computed based on the segmentation masks (warped with the resulting deformation fields). Besides, it can be seen, that predictions done by the demons are underestimating the EFs for HCM cases which should be  $>40\%$  according to the ACDC data set specifications.

#### IV. DISCUSSION AND CONCLUSIONS

We presented an unsupervised multi-scale deformable registration approach that learns a low-dimensional probabilistic deformation model. Our method not only allows the accurate registration of two images but also the analysis of deformations. The framework is generative, as it is able to simulate

deformations given only one image. Furthermore, it provides a novel way of deformation transport from one subject to another by using its probabilistic encoding. In the latent space, similar deformations are close to each other. The method enables the addition of a regularization term which leads to arbitrarily smooth deformations that are diffeomorphic by using an exponentiation layer for stationary velocity fields. The multi-scale approach, providing velocities, deformation fields and warped images in different scales, leads to improved registration results and a more controlled training procedure compared to a single-scale approach.

We evaluated the approach on end-diastole to end-systole cardiac cine-MRI registration and compared registration performance in terms of RMSE, DICE and Hausdorff distances to

two popular algorithms [8], [12] and a learning-based method [27], which are all diffeomorphic. While the performance of our single-scale approach was comparable to the LCC-demons and the SyN algorithm, our multi-scale approach (using 32 latent dimensions) showed statistically significant improvements in terms of registration accuracy. Generally, our approach produced more regular deformation fields, which are significantly smoother than the DL-based algorithm. Using our method with a non-generative U-net style network [55] without a deformation encoding performed similarly compared to the proposed generative model. Adding supervised information such as segmentation masks in the training procedure as in [26], [29] led to a marginal increase in terms of registration performance ( $\sim 1\text{-}2\%$  in DICE scores), so we decided that the performance gain is not large enough in order to justify the higher training complexity. Theoretically, our method allows measurement of registration uncertainty as proposed in [27] which we did not further investigate in this work.

The analysis of the deformation encoding showed that the latent space projects similar deformations close to each other such that diseases can be clustered. Disease classification could be potentially enforced in a supervised way as in [46]. Furthermore, our method showed comparable quantitative and qualitative results in transporting deformations with respect to a state-of-the-art algorithm which requires the difficult step of inter-subject registration that our algorithm does not need.

It is arguable if the simple assumption of a multivariate Gaussian is the right choice for the prior of the latent space (Eq. 1). Possible other assumptions such as a mixture of Gaussians are subject to future work. The authors think that the promising results of the learned probabilistic deformation model could be also applicable for other tasks such as evaluating disease progression in longitudinal studies or detecting abnormalities in subject-to-template registration. An open question is how the optimal size of the latent vector changes in different applications. In future work, we plan to further explore generative models for learning probabilistic deformation models.

#### ACKNOWLEDGMENTS

Data used in preparation of this article were obtained from the EU FP7-funded project MD-Paedegree and the ACDC STACOM challenge 2017 [54]. This work was supported by the grant AAP Santé 06 2017-260 DGA-DSH, and by the Inria Sophia Antipolis - Méditerranée, "NEF" computation cluster. The authors would like to thank Xavier Pennec for the insightful discussions and Adrian Dalca for the help with the Voxelmorph [27] experiments.

**Disclaimer:** This feature is based on research, and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.

#### REFERENCES

- [1] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: A survey," *IEEE Transactions on Medical Imaging*, vol. 32, no. 7, pp. 1153–1190, 2013.
- [2] M. Lorenzi and X. Pennec, "Efficient parallel transport of deformations in time series of images: from Schild's to pole ladder," *Journal of mathematical imaging and vision*, vol. 50, no. 1-2, pp. 5–17, 2014.
- [3] M.-M. Rohé, M. Sermesant, and X. Pennec, "Low-dimensional representation of cardiac motion using barycentric subspaces: A new group-wise paradigm for estimation, analysis, and reconstruction," *Medical image analysis*, vol. 45, pp. 1–12, 2018.
- [4] C. Davatzikos, "Spatial transformation and registration of brain images using elastically deformable models," *Computer Vision and Image Understanding*, vol. 66, no. 2, pp. 207–222, 1997.
- [5] M. Burger, J. Modersitzki, and L. Ruthotto, "A hyperelastic regularization energy for image registration," *SIAM Journal on Scientific Computing*, vol. 35, no. 1, pp. B132–B148, 2013.
- [6] J.-P. Thirion, "Image matching as a diffusion process: an analogy with Maxwell's demons," *Medical image analysis*, vol. 2, no. 3, pp. 243–260, 1998.
- [7] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Non-parametric diffeomorphic image registration with the demons algorithm," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2007, pp. 319–326.
- [8] M. Lorenzi, N. Ayache, G. B. Frisoni *et al.*, "LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm," *NeuroImage*, vol. 81, pp. 470–483, 2013.
- [9] M. F. Beg, M. I. Miller, A. Trounev *et al.*, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *International journal of computer vision*, vol. 61, no. 2, pp. 139–157, 2005.
- [10] Y. Cao, M. I. Miller, R. L. Winslow, L. Younes *et al.*, "Large deformation diffeomorphic metric mapping of vector fields," *IEEE transactions on medical imaging*, vol. 24, no. 9, pp. 1216–1230, 2005.
- [11] M. Zhang and P. T. Fletcher, "Finite-dimensional Lie algebras for fast diffeomorphic image registration," in *International Conference on Information Processing in Medical Imaging*. Springer, 2015, pp. 249–260.
- [12] B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain," *Medical image analysis*, vol. 12, no. 1, pp. 26–41, 2008.
- [13] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache, "A log-euclidean framework for statistics on diffeomorphisms," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2006, pp. 924–931.
- [14] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Symmetric log-domain diffeomorphic registration: A demons-based approach," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2008, pp. 754–761.
- [15] X. Yang, R. Kwitt, M. Styner, and M. Niethammer, "Quicksilver: Fast predictive image registration—a deep learning approach," *NeuroImage*, vol. 158, pp. 378–396, 2017.
- [16] M.-M. Rohé, M. Datar, T. Heimann, M. Sermesant, and X. Pennec, "SVF-net: Learning deformable image registration using shape matching," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 266–274.
- [17] H. Sokooti, B. de Vos *et al.*, "Nonrigid image registration using multi-scale 3D convolutional neural networks," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 232–239.
- [18] H. Uzunova, M. Wilms, H. Handels, and J. Ehrhardt, "Training CNNs for image registration from few samples with model-based data augmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 223–231.
- [19] K. A. Eppenhof, M. W. Lafarge, P. Moeskops, M. Veta, and J. P. Pluim, "Deformable image registration using convolutional neural networks," in *Medical Imaging 2018: Image Processing*, vol. 10574. International Society for Optics and Photonics, 2018, p. 105740S.
- [20] J. Krebs, T. Mansi, H. Delingette *et al.*, "Robust non-rigid registration through agent-based action learning," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 344–352.
- [21] D. Mahapatra, B. Antony, S. Sedai, and R. Garnavi, "Deformable medical image registration using generative adversarial networks," in *Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on*. IEEE, 2018, pp. 1449–1453.
- [22] J. Y. Jason, A. W. Harley, and K. G. Derpanis, "Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness," in *European Conference on Computer Vision*. Springer, 2016, pp. 3–10.
- [23] X. Liang, L. Lee, W. Dai *et al.*, "Dual motion GAN for future-flow embedded video prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1744–1752.



- [24] B. D. de Vos, F. F. Berendsen, M. A. Viergever, M. Staring, and I. Išgum, "End-to-end unsupervised deformable image registration with a convolutional neural network," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2017, pp. 204–212.
- [25] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "An unsupervised learning model for deformable medical image registration," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9252–9260.
- [26] J. Fan, X. Cao, P.-T. Yap, and D. Shen, "Birnet: Brain image registration using dual-supervised fully convolutional networks," *arXiv preprint arXiv:1802.04692*, 2018.
- [27] A. V. Dalca, G. Balakrishnan, J. Guttag, and M. Sabuncu, "Unsupervised learning for fast probabilistic diffeomorphic registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 729–738.
- [28] C. Tanner, F. Ozdemir, R. Profanter, V. Vishnevsky, E. Konukoglu, and O. Goksel, "Generative adversarial networks for MR-CT deformable image registration," *arXiv preprint arXiv:1807.07349*, 2018.
- [29] Y. Hu, M. Modat, E. Gibson, W. Li, N. Ghavami, E. Bonmati, G. Wang, S. Bandula, C. M. Moore, M. Emberton *et al.*, "Weakly-supervised convolutional neural networks for multimodal image registration," *Medical Image Analysis*, 2018.
- [30] M. Jaderberg, K. Simonyan, A. Zisserman *et al.*, "Spatial transformer networks," in *Advances in neural information processing systems*, 2015, pp. 2017–2025.
- [31] N. Duchateau, M. De Craene, G. Piella, E. Silva, A. Doltra, M. Sitges, B. H. Bijnens, and A. F. Frangi, "A spatiotemporal statistical atlas of motion for the quantification of abnormal myocardial tissue velocities," *Medical image analysis*, vol. 15, no. 3, pp. 316–328, 2011.
- [32] W. Bai, W. Shi, A. de Marvao, T. J. Dawes, D. P. O'Regan, S. A. Cook, and D. Rueckert, "A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion," *Medical image analysis*, vol. 26, no. 1, pp. 133–145, 2015.
- [33] M. Vaillant, M. I. Miller, L. Younes, and A. Trounev, "Statistics on diffeomorphisms via tangent space representations," *NeuroImage*, vol. 23, pp. S161–S169, 2004.
- [34] A. Qiu, L. Younes, and M. I. Miller, "Principal component based diffeomorphic surface mapping," *IEEE transactions on medical imaging*, vol. 31, no. 2, pp. 302–311, 2012.
- [35] M. Zhang and P. T. Fletcher, "Bayesian principal geodesic analysis in diffeomorphic image registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2014, pp. 121–128.
- [36] D. Wassermann, M. Toews, M. Niethammer, and W. Wells, "Probabilistic diffeomorphic registration: Representing uncertainty," in *International Workshop on Biomedical Image Registration*. Springer, 2014, pp. 72–82.
- [37] J. Wang, W. M. Wells, P. Golland, and M. Zhang, "Efficient laplace approximation for bayesian registration uncertainty quantification," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 880–888.
- [38] M. Lorenzi, N. Ayache, and X. Pennec, "Schild's ladder for the parallel transport of deformations in time series of images," in *Biennial International Conference on Information Processing in Medical Imaging*. Springer, 2011, pp. 463–474.
- [39] S. Jia, N. Duchateau, P. Mocerri, M. Sermesant, and X. Pennec, "Parallel transport of surface deformations from pole ladder to symmetrical extension," in *ShapeMI MICCAI 2018: Workshop on Shape in Medical Imaging*, 2018.
- [40] L. Younes, "Jacobi fields in groups of diffeomorphisms and applications," *Quarterly of applied mathematics*, pp. 113–134, 2007.
- [41] M. Louis, A. Bône, B. Charlier, S. Durrleman, A. D. N. Initiative *et al.*, "Parallel transport in shape analysis: a scalable numerical scheme," in *International Conference on Geometric Science of Information*. Springer, 2017, pp. 29–37.
- [42] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville *et al.*, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [43] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [44] A. Makhzani, J. Shlens, N. Jaitly, and I. Goodfellow, "Adversarial auto-encoders," in *International Conference on Learning Representations*, 2016. [Online]. Available: <http://arxiv.org/abs/1511.05644>
- [45] D. P. Kingma *et al.*, "Semi-supervised learning with deep generative models," in *Advances in Neural Information Processing Systems*, 2014, pp. 3581–3589.
- [46] C. Biffi, O. Oktay, G. Tarroni, W. Bai, A. De Marvao, G. Doumou, M. Rajchl, R. Bedair *et al.*, "Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 464–471.
- [47] W. Wei, E. Poirion, B. Bodini, S. Durrleman, N. Ayache, B. Stankoff, and O. Colliot, "Learning myelin content in multiple sclerosis from multimodal MRI through adversarial training," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 514–522.
- [48] J. Fan, X. Cao, Z. Xue, P.-T. Yap, and D. Shen, "Adversarial similarity network for evaluating image alignment in deep learning based registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 739–746.
- [49] J. Krebs, T. Mansi, B. Mailhé, N. Ayache, and H. Delingette, "Unsupervised probabilistic deformation modeling for robust diffeomorphic registration," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 101–109.
- [50] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in *Artificial Intelligence and Statistics*, 2015, pp. 562–570.
- [51] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein *et al.*, "A reproducible evaluation of ants similarity metric performance in brain image registration," *Neuroimage*, vol. 54, no. 3, pp. 2033–2044, 2011.
- [52] M. Nielsen, L. Florack, and R. Deriche, "Regularization, scale-space, and edge detection filters," *Journal of Mathematical Imaging and Vision*, vol. 7, no. 4, pp. 291–307, 1997.
- [53] K. Kurihara *et al.*, "Bayesian  $k$ -means as a "maximization-expectation" algorithm," *Neural Computation*, vol. 21, no. 4, pp. 1145–1172, 2009.
- [54] O. Bernard, A. Lalonde, C. Zotti, F. Cervenansky, X. Yang, P.-A. Heng, I. Cetin, K. Lekadir, O. Camara, M. A. G. Ballester *et al.*, "Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved?" *IEEE Transactions on Medical Imaging*, vol. 37, no. 11, pp. 2514–2525, 2018.
- [55] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

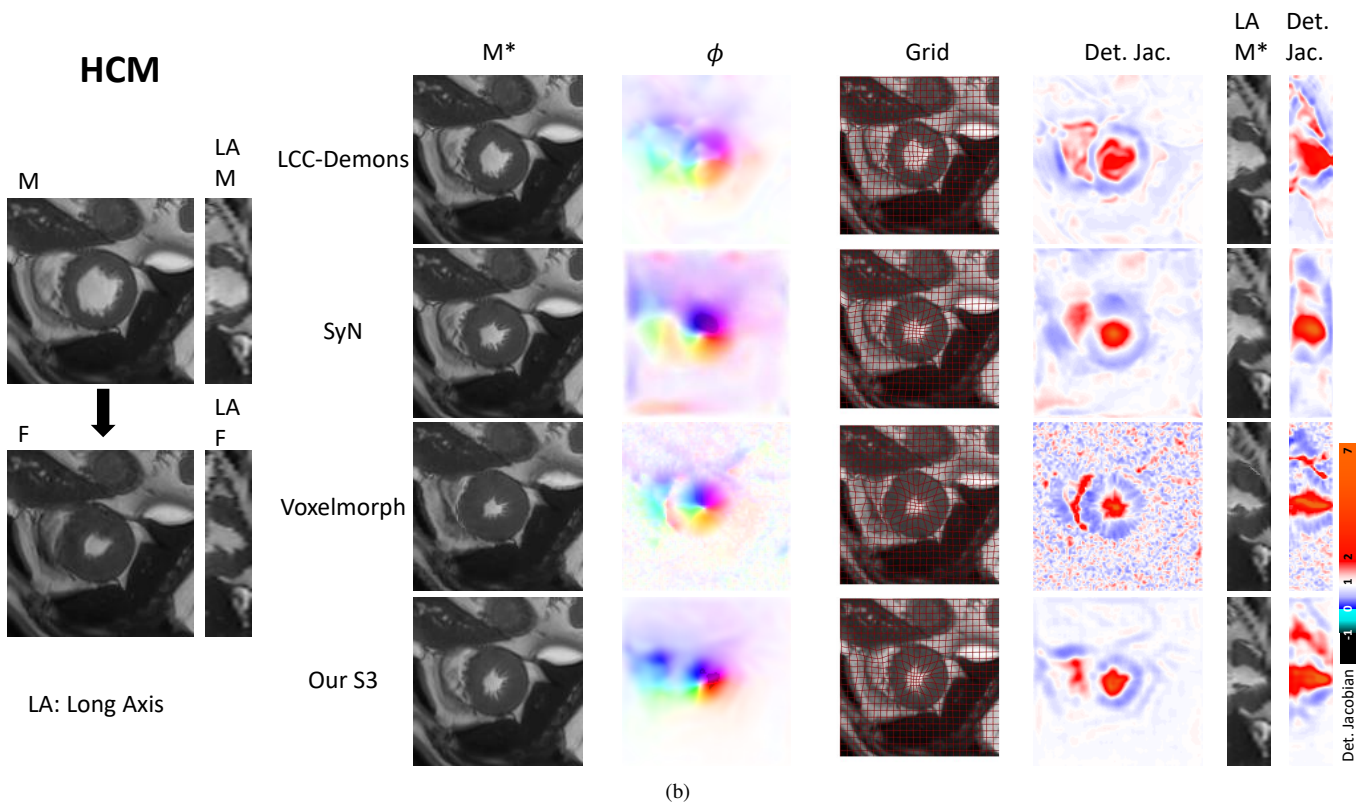
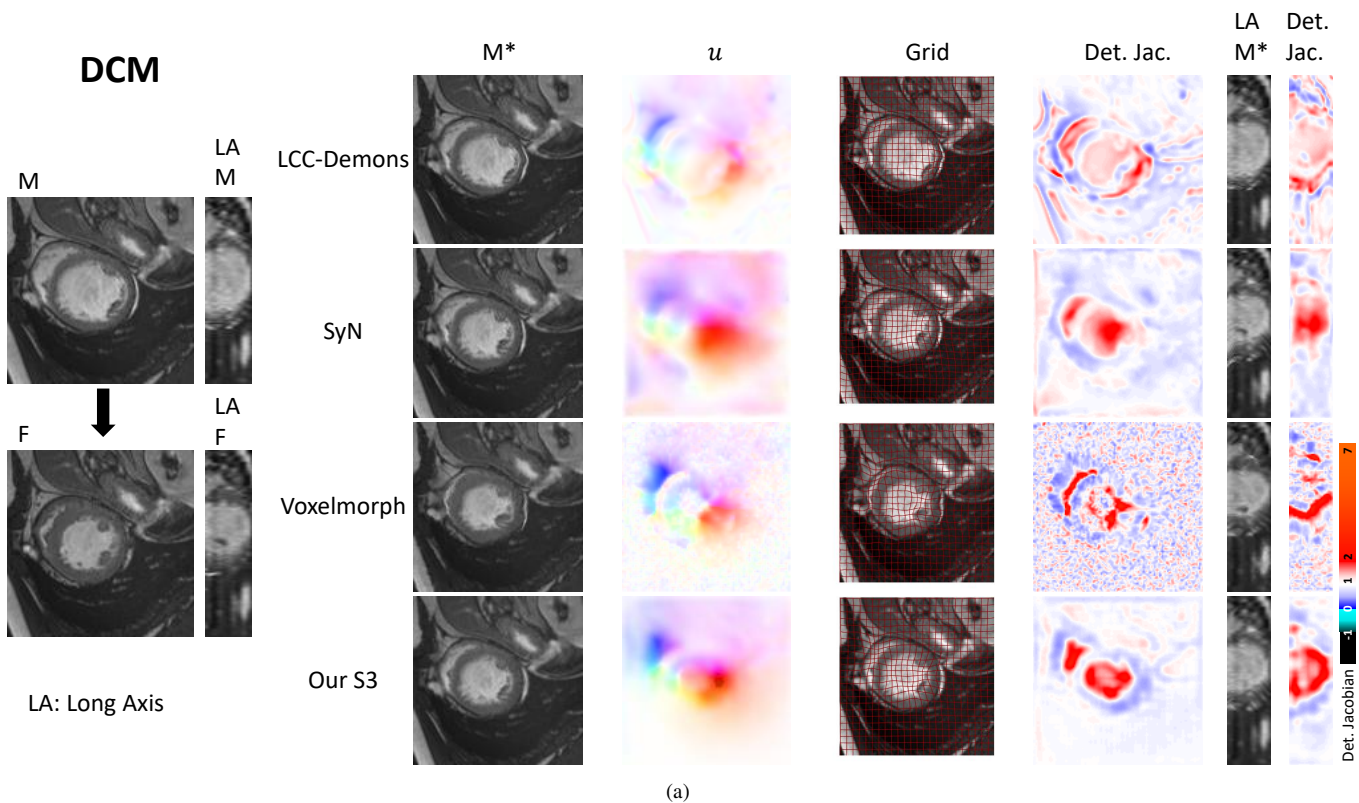
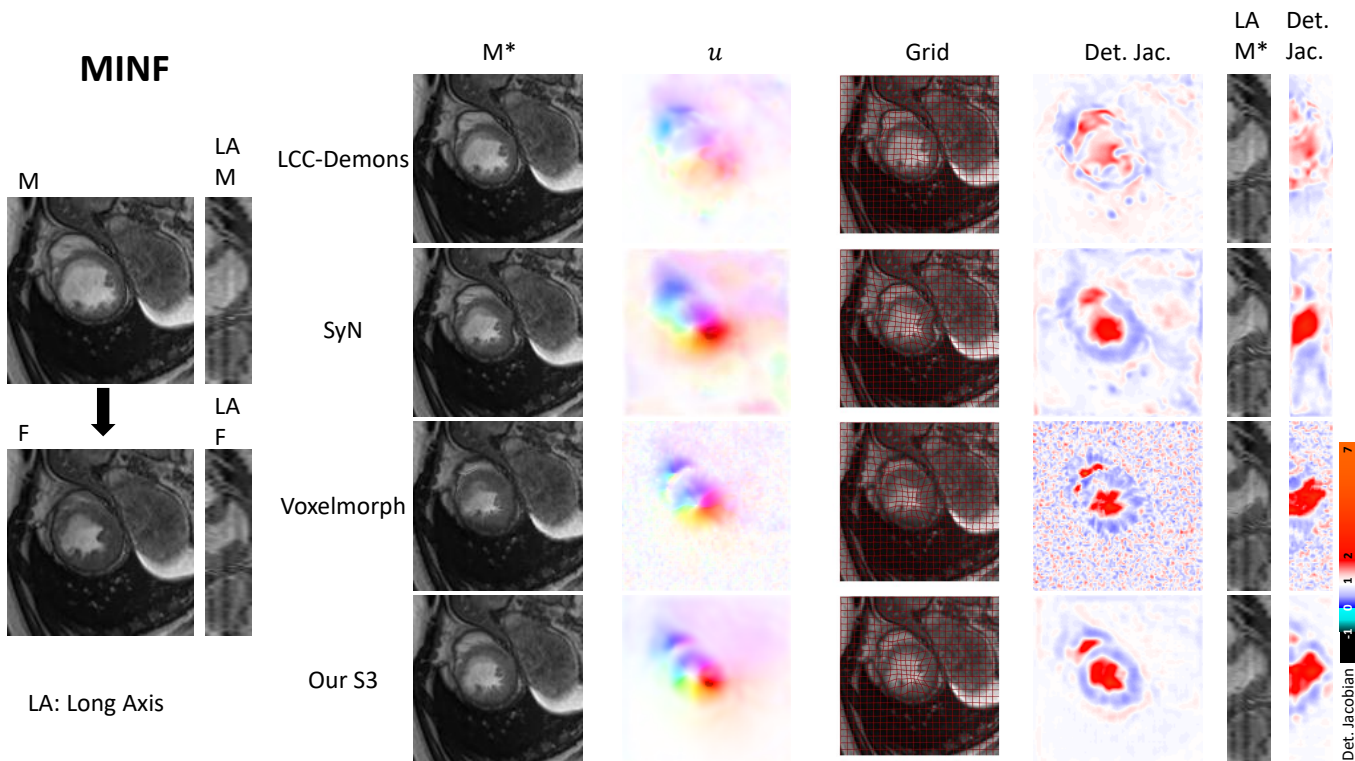
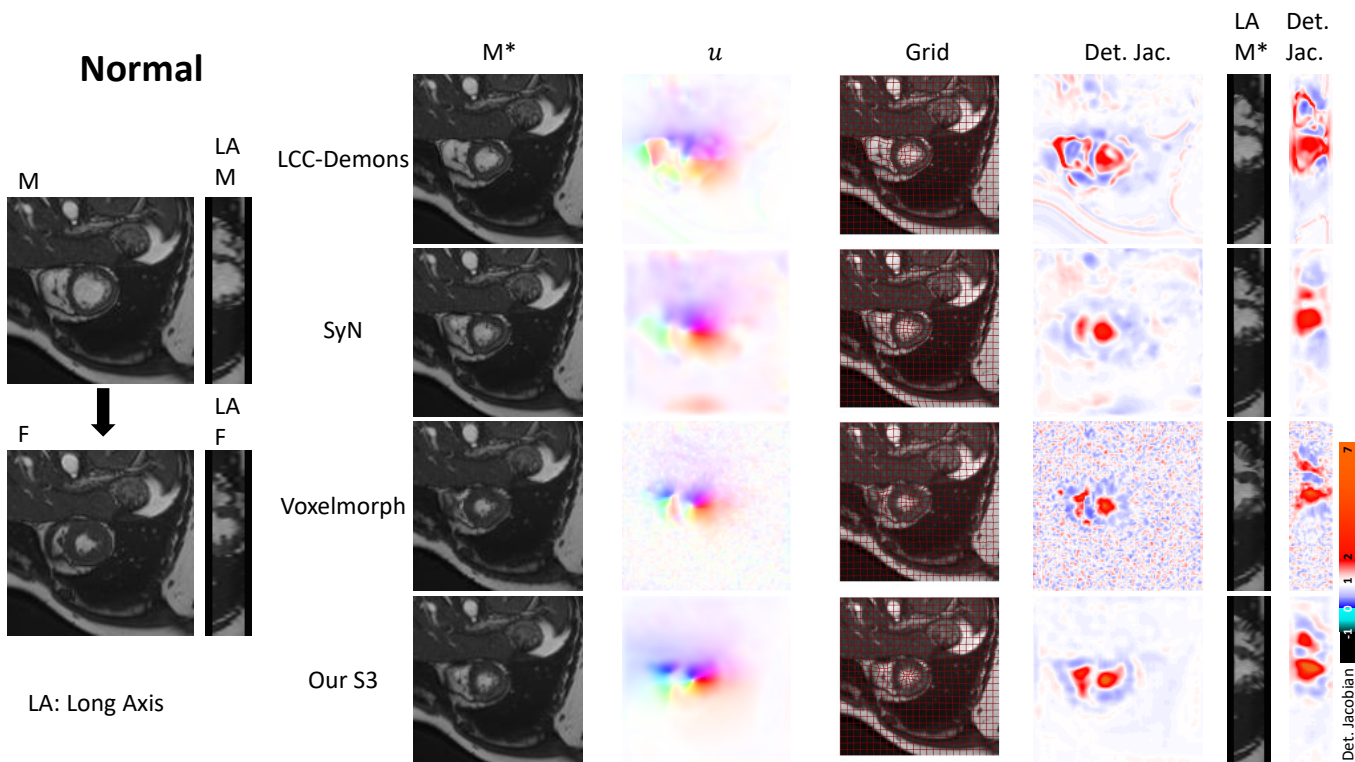


Fig. 10. Qualitative registration results showing a dilated cardiomyopathy (DCM) and a hypertrophic cardiomyopathy (HCM) case. Warped moving image  $M^*$ , displacements  $u$ , warped moving image with grid overlay and Jacobian determinant are shown for LCC-demons (Dem), SyN, voxelmorph (VM) and our approach using 3 scales (Our S3).



(a)



(b)

Fig. 11. Qualitative registration results showing a myocardial infarction (MINF) and healthy (Normal) case. Warped moving image  $M^*$ , displacements  $u$ , warped moving image with grid overlay and Jacobian determinant are shown for LCC-demons (Dem), SyN, voxelmorph (VM) and our approach using 3 scales (Our S3).



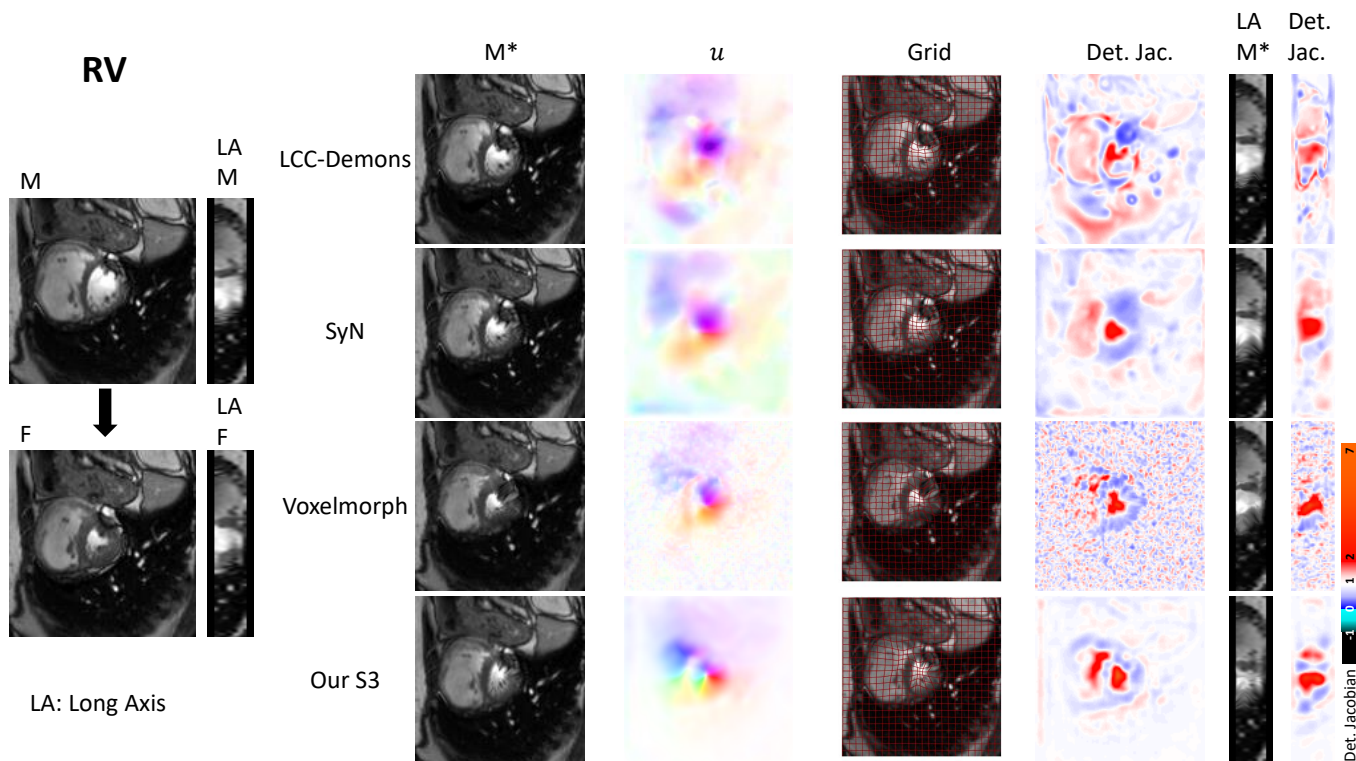


Fig. 12. Qualitative registration results showing an abnormal right ventricle case (RV). Warped moving image  $M^*$ , displacements  $u$ , warped moving image with grid overlay and Jacobian determinant are shown for LCC-demons (Dem), SyN, voxelmorph (VM) and our approach using 3 scales (Our S3).




Fig. 13. The gradient of the determinant of the Jacobian of a random test case for LCC-demons (Dem), SyN, voxelmorph (VM) and our approach using 1 and respectively 3 scales (Our S1, Our S3). Our single-scale approach shows the most regular deformation.





Fig. 14. (a) Symbolic pipeline for the parallel transport experiment using the pole ladder approach. (b) Visualization of all pipeline steps for one example.