



Nesterov Accelerated ADMM for Fast Diffeomorphic Image Registration

Alexander Thorley¹, Xi Jia¹, Hyung Jin Chang¹, Boyang Liu²,
Karina Bunting², Victoria Stoll², Antonio de Marvao³, Declan P. O'Regan³,
Georgios Gkoutos⁴, Dipak Kotecha², and Jinming Duan¹(✉)

¹ School of Computer Science, University of Birmingham, Birmingham, UK
j.duan@cs.bham.ac.uk

² Institute of Cardiovascular Sciences, University of Birmingham, Birmingham, UK

³ MRC London Institute of Medical Sciences, Imperial College London, London, UK

⁴ Institute of Cancer and Genomic Sciences, University of Birmingham,
Birmingham, UK

Abstract. Deterministic approaches using iterative optimisation have been historically successful in diffeomorphic image registration (DiffIR). Although these approaches are highly accurate, they typically carry a significant computational burden. Recent developments in stochastic approaches based on deep learning have achieved sub-second runtimes for DiffIR with competitive registration accuracy, offering a fast alternative to conventional iterative methods. In this paper, we attempt to reduce this difference in speed whilst retaining the performance advantage of iterative approaches in DiffIR. We first propose a simple iterative scheme that functionally composes intermediate non-stationary velocity fields to handle large deformations in images whilst guaranteeing diffeomorphisms in the resultant deformation. We then propose a convex optimisation model that uses a regularisation term of arbitrary order to impose smoothness on these velocity fields and solve this model with a fast algorithm that combines Nesterov gradient descent and the alternating direction method of multipliers (ADMM). Finally, we leverage the computational power of GPU to implement this accelerated ADMM solver on a 3D cardiac MRI dataset, further reducing runtime to less than 2s. In addition to producing strictly diffeomorphic deformations, our methods outperform both state-of-the-art deep learning-based and iterative DiffIR approaches in terms of dice and Hausdorff scores, with speed approaching the inference time of deep learning-based methods.

Keywords: Image registration · Diffeomorphism · ADMM

1 Introduction

Over the past two decades, diffeomorphic image registration (DiffIR) has become a powerful tool for deformable image registration. The goal of DiffIR is to find a

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-87202-1_15) contains supplementary material, which is available to authorized users.

smooth and invertible spatial transformation between two images, such that every point in one image has a corresponding point in the other. Such transformations are known as diffeomorphisms and are fundamental inputs for computational anatomy in medical imaging.

Beg et al. pioneering work developed the large-deformation diffeomorphic metric mapping (LDDMM) framework [5], formulated as a variational problem where the diffeomorphism is represented by a geodesic path (flow) parameterised by time-dependent (non-stationary) smooth velocity fields through an ordinary differential equation. Vialard et al. proposed a geodesic shooting algorithm [24] in which the geodesic path of diffeomorphisms can be derived from the Euler-Poincaré differential (EPDiff) equation given only the initial velocity field. In this case, geodesic shooting requires only an estimate of the initial velocity rather than the entire sequence of velocity fields required by LDDMM [5], thus reducing the optimisation complexity. Recently, Zhang et al. developed the Fourier-approximated Lie algebras for shooting (FLASH) [28] for DiffIR, where they improved upon the geodesic shooting algorithm [22, 24] by speeding up the calculation of the EPDiff equation in the Fourier space. However, solutions of these approaches [5, 24, 28] rely on gradient descent, which requires many iterations to converge and is therefore very slow. To reduce the computational cost of DiffIR, Ashburner represented diffeomorphisms by stationary velocity fields (SVFs) [1], and proposed a fast DiffIR algorithm (DARTEL) [2] by composing successive diffeomorphisms using scaling and squaring. However, SVFs do not provide geodesic paths between images which may be critical to statistical shape analysis. Another approach for fast DiffIR is demons [23], although this is a heuristic method and does not have a clear energy function to minimise.

An alternative approach to improve DiffIR speed is to leverage deep learning, usually adopting convolutional neural networks (CNNs) to learn diffeomorphic transformations between pairwise images in a training dataset. Registration after training is then achieved efficiently by evaluating the network on unseen image pairs with or without further (instance-specific) optimisation [4]. Deep learning in DiffIR can be either supervised or unsupervised. Yang et al. proposed Quicksilver [26] which utilises a supervised encoder-decoder network as the patch-wise prediction model. Quicksilver is trained with the initial momentum¹ of LDDMM as the supervision signal, which does not need to be spatially smooth and hence facilitates a fast patch-wise training strategy. Wang et al. extended FLASH [28] to DeepFLASH [25] in a learning framework, which predicts the initial velocity field in the Fourier space (termed the low dimensional bandlimited space in their paper). DeepFLASH is more efficient in general as the backbone network consists of only a decoder. As Quicksilver and DeepFLASH need to be trained using the numerical solutions of LDDMM [22] as ground truth, their performance may be bounded by that of LDDMM [22]. Dalca et al. proposed diffeomorphic VoxelMorph [9] that leverages the U-Net [20] and the spatial transformer network [15]. Mok et al. developed a fast symmetric method (SymNet) [18] that guarantees topology preservation and invertibility of the transformation. Both

¹ Momentum is the dual form of velocity. They are connected by a symmetric, positive semi-definite Laplacian operator as defined in Eq. (6).

Voxelmorph and SymNet impose diffeomorphisms using SVFs in an unsupervised fashion. Whilst the fast inference speed of deep learning is a clear benefit, the need for large quantities of training data and the lack of robust mathematical formulation in comparison with deterministic iterative approaches means both paradigms still have their advantages.

The goal of this study is to reduce the speed gap between learning-based and iterative approaches in DiffIR, while still retaining the performance advantage of iterative approaches. We first propose an iterative scheme to decompose a diffeomorphic transformation into the compositions of a series of small non-stationary velocity fields. Next, we define a clear, general and convex optimisation model to compute such velocity fields. The data term of the model can be either \mathcal{L}^1 or \mathcal{L}^2 norm and the \mathcal{L}^2 regularisation term uses a derivative of arbitrary order whose formulation follows the multinomial theorem. We then propose a fast solver for the proposed general model, which combines Nesterov acceleration [19] and the alternating direction method of multipliers (ADMM) [6, 13, 17]. By construction, all resulting subproblem solutions are point-wise and closed-form without any iteration. As such, the solver is very accurate and efficient, and can be implemented using existing deep learning frameworks (e.g. PyTorch) for further acceleration via GPU. We show on a 3D cardiac MRI dataset that our methods outperform both state-of-the-art learning and optimisation-based approaches, with a speed approaching the inference time of learning-based methods.

2 Diffeomorphic Image Registration

Let $\phi : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ denote the deformation field that maps the coordinates from one image to those in another image. Computing a diffeomorphic deformation can be treated as modelling a dynamical system [5], given by the ordinary differential equation (ODE): $\partial\phi/\partial t = \mathbf{v}_t(\phi_t)$, where $\phi_0 = \text{Id}$ is the identity transformation and \mathbf{v}_t indicates the velocity field at time t ($\in [0, 1]$). A straightforward approach to numerically compute the ODE is through Euler integration, in which case the final deformation field ϕ_1 is calculated from the compositions of a series of small deformations. More specifically, a series of N velocity fields are used to represent the time varying velocity field \mathbf{v}_t . For N uniformly spaced time steps $(0, t_1, t_2, \dots, t_{N-2}, t_{N-1})$, ϕ_1 can be achieved by:

$$\phi_1 = \left(\text{Id} + \frac{\mathbf{v}_{t_{N-1}}}{N}\right) \circ \left(\text{Id} + \frac{\mathbf{v}_{t_{N-2}}}{N}\right) \circ \dots \circ \left(\text{Id} + \frac{\mathbf{v}_{t_1}}{N}\right) \circ \left(\text{Id} + \frac{\mathbf{v}_0}{N}\right), \quad (1)$$

where \circ denotes function composition. This greedy composite approach is seen in some DiffIR works [7, 23] and if the velocity fields $\mathbf{v}_{t_i}, \forall i \in \{0, \dots, N-1\}$ are sufficiently small whilst satisfying some smoothness constraint, the resulting compositions should result in a deformation that is diffeomorphic [8].

The fact that the numerical solver (1) requires small and smooth velocity fields motivates us to consider classical optical flow models, among which a simple yet effective method was proposed by Horn and Schunck (HS), who computed velocity \mathbf{v} through $\min_{\mathbf{v}} \left\{ \frac{1}{2} \|\langle \nabla I_1, \mathbf{v} \rangle + I_1 - I_0 \|^2 + \frac{\lambda}{2} \|\nabla \mathbf{v}\|^2 \right\}$ [14]. In the minimisation problem, the first data term imposes brightness constancy, where

$I_0 \in \mathbb{R}^{MNH}$ (MNH is the image size) and $I_1 \in \mathbb{R}^{MNH}$ are a pair of input images and $\langle \cdot, \cdot \rangle$ denotes inner product. The second regularisation term induces spatial smoothness on $\mathbf{v} \in (\mathbb{R}^{MNH})^3$, where λ is a hyperparameter and ∇ denotes the gradient operator. As the HS data term is based on local Taylor series approximations of the image signal, it is only capable of recovering small velocity fields, which together with the smooth regularisation makes the HS model perfect to combine with (1) for effective diffeomorphic registration.

Based on this observation, we propose a simple iterative scheme for DiffIR. For a pair of images I_0 (target image) and I_1 (source image), first we solve the HS model for $\mathbf{v}_{t_{N-1}}^*$. Of note, as both terms in the HS model are linear and quadratic, if we derive the first-order optimality condition with respect to \mathbf{v} from the model, we end up with a system of linear equations (i.e. regularised normal equation), which can be solved analytically by matrix inversion. We then warp the source image I_1 to I_1^ω using $I_1^\omega = I_1 \circ (\text{Id} + \mathbf{v}_{t_{N-1}}^*)$. Note that we set the denominator coefficient N to 1 in this paper. Next, we pass I_1^ω and I_0 to the HS model again and solve for $\mathbf{v}_{t_{N-2}}^*$, with which we update the warped source image I_1^ω with $I_1^\omega = I_1 \circ (\text{Id} + \mathbf{v}_{t_{N-1}}^*) \circ (\text{Id} + \mathbf{v}_{t_{N-2}}^*)$. We repeat this iterative process until that the final warped source image $I_1^\omega = I_1 \circ (\text{Id} + \mathbf{v}_{t_{N-1}}^*) \circ (\text{Id} + \mathbf{v}_{t_{N-2}}^*) \circ \dots \circ (\text{Id} + \mathbf{v}_{t_1}^*) \circ (\text{Id} + \mathbf{v}_0^*)$ is close to the target image I_0 within a pre-defined, small and positive tolerance. We note that unlike [5, 22, 24, 28] there is no need for the proposed iterative scheme to pre-define the number of velocity fields as it is automatically computed based on the similarity between the warped and target images, however we can and will impose a restriction on the number of iterations to reduce computation time.

The HS model however suffers from the fact that the quadratic data term is not robust to outliers and that the first-order diffusion regularisation performs poorly against large deformations (which often requires an additional affine linear pre-registration to be successful [12]). To circumvent these shortcomings of the HS model, we propose a new, general variational model, given by

$$\min_{\mathbf{v}} \left\{ \frac{1}{s} \|\rho(\mathbf{v})\|^s + \frac{\lambda}{2} \|\nabla^n \mathbf{v}\|^2 \right\}, \quad (2)$$

where $\rho(\mathbf{v}) = \langle \nabla I_1, \mathbf{v} \rangle + I_1 - I_0$, $s = \{1, 2\}$ and $n > 0$ is an integer number. ∇^n denotes an arbitrary order gradient. For a scalar-valued function $f(x, y, z)$, $\nabla^n f = [(k_1, k_2, k_3) \frac{\partial^n f}{\partial x^{k_1} \partial y^{k_2} \partial z^{k_3}}]^T$, where $k_i = 0, \dots, n$, $i \in \{1, 2, 3\}$ and $k_1 + k_2 + k_3 = n$. Moreover, (k_1, k_2, k_3) are known as multinomial coefficients which are computed by $\frac{n!}{k_1! k_2! k_3!}$. Note that if we set $s = 2$ and $n = 1$, the proposed model is recovered to the original HS model. However, if we set $s = 1$, the data term will be the sum absolute differences (SAD) [27] which is more robust to outliers in images. On the other hand, the proposed regularisation has a general form, which allows reconstruction of a smooth polynomial function of arbitrary degree for the velocity. Unfortunately, solving the proposed model becomes non-trivial due to the non-differentiability in the data term and the generality in the derivative order. We note that even if we only solve the original HS model, we still need to do matrix inversion which is impractical for 3D high-dimensional data used in this paper. To address these issues, we propose an effective algorithm in the following section to minimise the proposed convex model efficiently.

3 Nesterov Accelerated ADMM

In order to minimise the convex problem (2) efficiently, we utilise the alternating direction method of multipliers (ADMM) [6, 13, 17] accelerated by Nesterov's approach for gradient descent methods [19]. First, we introduce an auxiliary primal variable $\mathbf{w} \in (\mathbb{R}^{MNH})^3$, converting (2) into the following equivalent form

$$\min_{\mathbf{v}} \left\{ \frac{1}{s} \|\rho(\mathbf{v})\|^s + \frac{\lambda}{2} \|\nabla^n \mathbf{w}\|^2 \right\} \text{ s.t. } \mathbf{w} = \mathbf{v}. \quad (3)$$

The introduction of the constraint $\mathbf{w} = \mathbf{v}$ decouples \mathbf{v} in the regularisation from that in the data term so that we can derive closed-form, point-wise solutions with respect to both variables regardless of the non-differentiability and generality of the problem (2). To guarantee an optimal solution, the above constrained problem can be converted to a saddle problem solved by the proposed fast ADMM. Let $\mathcal{L}_{\mathcal{A}}(\mathbf{v}, \mathbf{w}; \mathbf{b})$ be the augmented Lagrange functional of (3), defined as

$$\mathcal{L}_{\mathcal{A}}(\mathbf{v}, \mathbf{w}; \mathbf{b}) = \frac{1}{s} \|\rho(\mathbf{v})\|^s + \frac{\lambda}{2} \|\nabla^n \mathbf{w}\|^2 + \frac{\theta}{2} \|\mathbf{w} - \mathbf{v} - \mathbf{b}\|^2, \quad (4)$$

where $\mathbf{b} \in (\mathbb{R}^{MNH})^3$ is the augmented Lagrangian multiplier (aka. dual variable), and $\theta > 0$ is the penalty parameter which has impact on how fast the minimisation process converges. It is known that one of the saddle points for (4) will give a minimiser for the constrained minimisation problem (3). We use the Nesterov accelerated ADMM to optimise variables in (4), which is given as

$$\begin{cases} \mathbf{v}^k = \arg \min_{\mathbf{v}} \frac{1}{s} \|\rho(\mathbf{v})\|^s + \frac{\theta}{2} \|\widehat{\mathbf{w}}^k - \mathbf{v} - \widehat{\mathbf{b}}^k\|^2 \\ \mathbf{w}^k = \arg \min_{\mathbf{w}} \frac{\lambda}{2} \|\nabla^n \mathbf{w}\|^2 + \frac{\theta}{2} \|\mathbf{w} - \mathbf{v}^k - \widehat{\mathbf{b}}^k\|^2 \\ \mathbf{b}^k = \widehat{\mathbf{b}}^k + \mathbf{v}^k - \mathbf{w}^k \\ \alpha^{k+1} = \frac{1 + \sqrt{1 + 4(\alpha^k)^2}}{2} \\ \widehat{\mathbf{w}}^{k+1} = \mathbf{w}^k + \frac{\alpha^k - 1}{\alpha^{k+1}} (\mathbf{w}^k - \mathbf{w}^{k-1}) \\ \widehat{\mathbf{b}}^{k+1} = \mathbf{b}^k + \frac{\alpha^k - 1}{\alpha^{k+1}} (\mathbf{b}^k - \mathbf{b}^{k-1}) \end{cases}. \quad (5)$$

At the beginning of the iteration, we set $\widehat{\mathbf{w}}^{-1} = \widehat{\mathbf{w}}^0 = \mathbf{0}$, $\widehat{\mathbf{b}}^{-1} = \widehat{\mathbf{b}}^0 = \mathbf{0}$ and $\alpha^0 = 1$. k above denotes the number of iterations. As can be seen from (5), we decompose (4) into two convex subproblems with respect to \mathbf{v} and \mathbf{w} , update the dual variable \mathbf{b} , and then apply Nesterov's acceleration to both primal variable \mathbf{w} and dual variable \mathbf{b} . This accelerated iterative process is repeated until the optimal solution \mathbf{v}^* is found. The convergence of this accelerated ADMM has been proved in [13] for convex problems and the authors have also demonstrated the accelerated ADMM has a convergence rate of $\mathcal{O}(1/k^2)$, which is the optimal rate for first-order algorithms. Moreover, due to the decoupling we are able to derive a closed-form, point-wise solution for both subproblems of \mathbf{v} and \mathbf{w}

$$\begin{cases} \mathbf{v}^k = \widehat{\mathbf{w}}^k - \widehat{\mathbf{b}}^k - \frac{\hat{z}}{\max(|\hat{z}|, 1)} \frac{\nabla I_1}{\theta} & \text{if } s = 1 \\ (\mathbf{J}\mathbf{J}^T + \theta \mathbf{1}) \mathbf{v}^k = \theta (\widehat{\mathbf{w}}^k - \widehat{\mathbf{b}}^k) - \mathbf{J}(I_1 - I_0) & \text{if } s = 2 \\ \mathbf{w}^k = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{v}^k + \widehat{\mathbf{b}}^k)/(\lambda \mathcal{F}(\Delta^n) + \theta)) & \end{cases}. \quad (6)$$

In the case of $s = 1$, we have $\hat{z} = \theta \rho(\hat{\mathbf{w}}^k - \hat{\mathbf{b}}^k) / (|\nabla I_1|^2 + \epsilon)$, where $\epsilon > 0$ is a small, positive number to avoid division by zeros. In the case of $s = 2$, $\mathbf{J}\mathbf{J}^T$ (where $\mathbf{J} = \nabla I_1$) is a rank-1 outer product and $\mathbb{1}$ is an identity matrix. As the respective \mathbf{v} -subproblem in this case is differentiable, we can derive the Sherman–Morrison formula (i.e. see (6) middle equation) by directly differentiating this subproblem. Due to the identity matrix, the Sherman–Morrison formula [23] will lead to a closed-form, point-wise solution to \mathbf{v}^k . In the last equation, \mathcal{F} and \mathcal{F}^{-1} respectively denote the discrete Fourier transformation (DFT) and inverse DFT, and Δ^n is the n th-order Laplace operator which are discretised via the finite difference [11, 17]. In this case, $\mathcal{F}(\Delta^n) = 2^n(3 - \cos(\frac{2\pi p}{M}) - \cos(\frac{2\pi q}{N}) - \cos(\frac{2\pi r}{H}))^n$, where $p \in [0, M)$, $q \in [0, N)$ and $r \in [0, H)$ are grid indices. We note that all three solutions are closed-form and point-wise, and therefore can be computed very efficiently in 3D.

4 Experimental Results

Dataset: The dataset used in our experiments consists of 220 pairs of 3D high-resolution (HR) cardiac MRI images corresponding to the end diastolic (ED) and end systolic (ES) frames of the cardiac cycle. The raw volumes have a resolution of $1.2 \times 1.2 \times 2.0 \text{ mm}^3$. HR imaging requires only one single breath-hold and therefore introduces no inter-slice shift artifacts [10]. Class segmentation labels at ED and ES are also contained in the dataset for the right ventricle (RV), left ventricle (LV) and left ventricle myocardium (LVM). All images are resampled to $1.2 \times 1.2 \times 1.2 \text{ mm}^3$ resolution and cropped or padded to matrix size $128 \times 128 \times 96$. To train comparative deep learning methods and tune hyperparameters in different methods, the dataset is split into 100/20/100 corresponding to training, validation and test sets. We report final quantitative results on the test set only (Table 1).

Hyperparameter Tuning: We embed our approach into a three-level pyramid scheme to avoid convergence to local minima, resulting in a total of three loops in the final implementation. We tune two versions of hyperparameters for our model, the first with limited iterations to minimise runtime, and the second without such restrictions leading to a slower runtime but higher accuracy. For the first, we restrict the iteration number within the ADMM loop in (5) to ten in the first pyramid level and five in the following two levels. We also cap the number of warpings in (1) to ten in the first two pyramid levels and five in the final level. For the second, iterations are only terminated when a convergence threshold (range $[0.01, 0.13]$) for the ADMM loop is met and the difference between two warpings is below 2%. This range of thresholds ensures runtime is competitive with other iterative methods whilst increasing accuracy compared to our first version. For both versions, the validation set is used to tune hyperparameters λ (range $[0, 60]$) and θ (range $[0.01, 0.15]$) to maximise dice score for each combination of data term (\mathcal{L}^2 or \mathcal{L}^1) and regulariser of order 1 to 3 ($1^{st}O$, $2^{nd}O$ or $3^{rd}O$). See supplementary for the optimal values of λ and θ .

Comparative Methods: We compare our methods against a variety of state-of-the-art iterative and deep learning-based image registration methods. For iterative methods, we use the MIRTk implementation of free form deformations

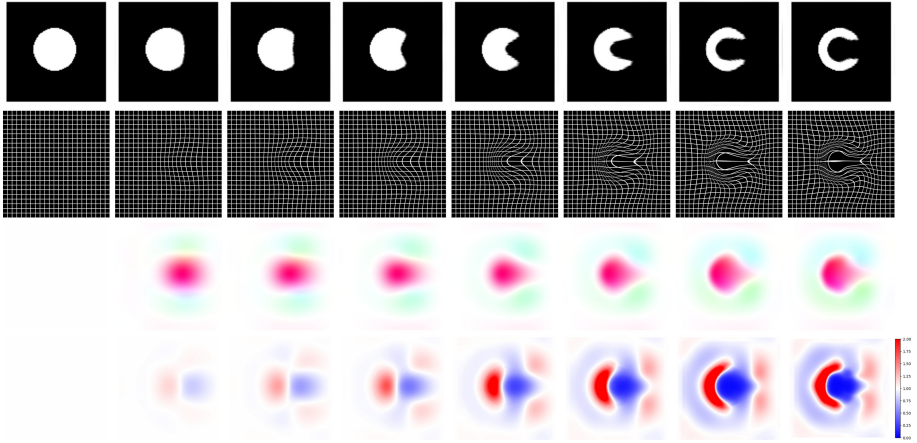


Fig. 1. A toy experiment showing evolution of diffeomorphism throughout iterations using our method. Rows one to four correspond to image, grid deformation, HSV deformation and jacobian determinant, respectively. Columns one to eight correspond to increasing iterations.

(FFD) [21] with a three-level pyramid scheme. Control point spacing is selected to maximise dice score on the validation set. We also compare with the popular diffeomorphic demons algorithm [23] implemented in SimpleITK [16], using a three-level pyramid scheme. The number of iterations and smoothing λ 's are optimised on the validation set. We also compare with ANTs SyN [3] using the official ANTs implementation with mean square distance and four registration levels. The hyper-parameters such as smoothing parameters, gradient step size, and the number of iterations are selected based on the validation set. For deep learning, we train both VoxelMorph [4] and its diffeomorphic version [9] on the training set. We adopt a lighter 5-level hierarchical U-shape network [18] than the original VoxelMorph paper [4] as our $128 \times 128 \times 96$ images require significant GPU memory. Both VoxelMorphy are optimised under an SAD loss defining image similarity, a diffusion regularisation loss to enforce smoothness of the velocity field and a local orientation consistency loss [18] that regularises the Jacobian determinant of the velocity field. For the diffeomorphic VoxelMorph, an additional 7 scaling and squaring layers are used to encourage diffeomorphic properties in the resulting deformation. The weights of smoothness regularisation and local orientation consistency are tuned to 0.1 and 1000 respectively on the validation set, and both trained for 30000 iterations with batch size 2 and learning rate 0.0001 on a GTX 1080Ti GPU. Training takes around 13 h.

Quantitative Results: To evaluate performance, we use the deformations produced by each method to warp the associated source (ES) class label to its corresponding target (ED) label. The warped source and target labels are then used to compute the dice similarity coefficient and Hausdorff distance between pairs, with Hausdorff distances averaged over each 2D slice within the 3D volumes. We

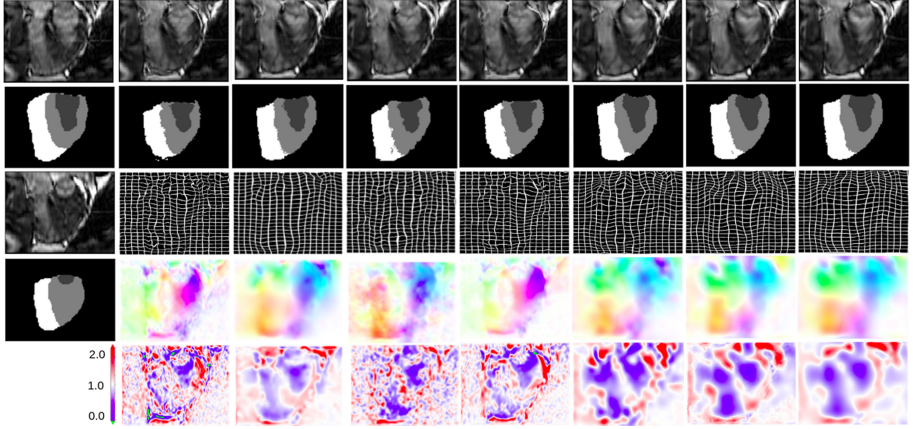


Fig. 2. Visual results of different methods. Column 1: target image, target label, source image, source label. Columns 2–8: warped source image, warped source label, grid, HSV_ϕ , $\det(\mathbf{J}_\phi)$ for (1) Voxelmorph, (2) SyN, (3) Demons, (4) diffeomorphic voxelmorph, (5) FFD, (6) $\mathcal{L}^1+3^{rd}O$ and (7) $\mathcal{L}^2+3^{rd}O$, respectively.

compute the percentage of non-positive voxels in the Jacobian \mathbf{J}_ϕ of each deformation ϕ to assess its diffeomorphic properties, for which a purely diffeomorphic deformation should be zero everywhere. We also compare runtime, measuring speed of iterative methods and inference time for two Voxelmorph models. All methods using a GPU are tested on a 16GB NVIDIA Tesla V100.

Table 1. Statistics of each method averaged over test set. Dice and Hausdorff scores are denoted as an average across RV, LV and LVM. Standard deviations are shown after \pm . Hausdorff distance is measured in *mm* and runtime in seconds. Results of our methods using 2nd hyperparameter version are given in brackets.

Method	Dice	Hausdorff	% of $ \mathbf{J}_\phi \leq 0$	Runtime
Unreg	.493 \pm .043	8.404 \pm 0.894	–	–
Voxelmorph	.709 \pm .032	7.431 \pm 1.049	0.01 \pm 0.01	0.09 \pm 0.23
SyN	.721 \pm .051	6.979 \pm 1.233	0.00 \pm 0.00	77.39 \pm 9.57
Demons	.727 \pm .040	6.740 \pm 0.998	0.00 \pm 0.00	13.01 \pm 0.17
Diff-Voxelmorph	.730 \pm .031	6.498 \pm 0.975	0.01 \pm 0.01	0.09 \pm 0.21
FFD	.739 \pm .047	7.044 \pm 1.200	0.77 \pm 0.30	8.98 \pm 1.97
$\mathcal{L}^1+1^{st}O$.727 \pm .039 (.728 \pm .040)	6.949 \pm 1.056 (6.911 \pm 1.092)	0 \pm 0 (0 \pm 0)	1.63 \pm .26 (7.90 \pm .79)
$\mathcal{L}^2+1^{st}O$.719 \pm .039 (.726 \pm .041)	7.499 \pm 1.227 (7.105 \pm 1.199)	0 \pm 0 (0 \pm 0)	2.17 \pm .50 (4.69 \pm .34)
$\mathcal{L}^1+2^{nd}O$.749 \pm .042 (.761 \pm .041)	6.891 \pm 1.134 (6.364 \pm 1.036)	0 \pm 0 (0 \pm 0)	1.43 \pm .05 (5.02 \pm .57)
$\mathcal{L}^2+2^{nd}O$.735 \pm .043 (.751 \pm .045)	6.895 \pm 1.212 (6.855 \pm 1.231)	0 \pm 0 (0 \pm 0)	1.38 \pm .04 (5.00 \pm .41)
$\mathcal{L}^1+3^{rd}O$.753 \pm .042 (.768 \pm .042)	6.963 \pm 1.100 (6.515 \pm 1.077)	0 \pm 0 (0 \pm 0)	1.57 \pm .24 (4.49 \pm .56)
$\mathcal{L}^2+3^{rd}O$.736 \pm .042 (.757 \pm .046)	7.302 \pm 1.237 (6.927 \pm 1.264)	0 \pm 0 (0 \pm 0)	1.37 \pm .03 (4.72 \pm .41)

In Fig. 2, we show the visual comparison between all methods, where grid_ϕ shows the deformation applied to a unit grid, $\det(\mathbf{J}_\phi)$ the jacobian determinant and HSV_ϕ the final composition of each iterative velocity field, with hue depicting direction and saturation magnitude. In Table 1, we show the final results of each method averaged over the test set. Results from both versions of hyperparameters are displayed with the latter in brackets. It can be seen that our accelerated ADMM outperforms all other methods by at least 1% in terms of dice score despite limiting the number of iterations and warpings, increasing to nearly 3% for our second version. There is also a clear upward tendency in dice (see Fig. 3), further demonstrating the superiority of our methods. Moreover, our methods consistently maintain diffeomorphic properties with no negative elements present in \mathbf{J}_ϕ . Although the speed of our method is still slower than deep learning methods, the average runtime we achieved for our first hyperparameter version is barely over a second.

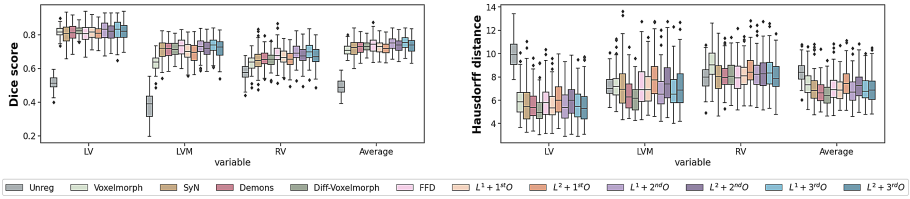


Fig. 3. Boxplots of Dice and Hausdorff for left ventricle (LV), left ventricle myocardium (LVM) and right ventricle (RV) of all methods on test set.

5 Conclusion

In this paper, we proposed an accelerated algorithm that combines Nesterov gradient descent and ADMM, which tackles a general variational model with a regularisation term of arbitrary order. Through the compositions of a series of velocity fields produced by our accelerated ADMM, we model deformations as the solution to a dynamical system to ensure diffeomorphic properties. We implement our methods in the PyTorch framework, leveraging the power of a GPU to further accelerate our methods. In addition to reducing the difference in speed between deep learning to under 2s, our methods have achieved a new state-of-the-art performance for the DiffIR task.

Acknowledgements. The research is supported by the BHF Accelerator Award (AA/18/2/34218), the Ramsay Research Fund from the School of Computer Science at the University of Birmingham and the Wellcome Trust Institutional Strategic Support Fund: Digital Health Pilot Grant.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-Euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006). https://doi.org/10.1007/11866565_113
2. Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95–113 (2007)
3. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage* **54**(3), 2033–2044 (2011)
4. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**(8), 1788–1800 (2019)
5. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* **61**(2), 139–157 (2005)
6. Boyd, S., Parikh, N., Chu, E.: Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. Now Publishers Inc., Delft (2011)
7. Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.* **5**(10), 1435–1447 (1996)
8. Christensen, G.E., et al.: Topological properties of smooth anatomic maps. In: Information Processing in Medical Imaging, vol. 3, pp. 101–112. Kluwer Academic, Boston (1995)
9. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Med. Image Anal.* **57**, 226–236 (2019)
10. Duan, J., et al.: Automatic 3D bi-ventricular segmentation of cardiac images by a shape-refined multi-task deep learning approach. *IEEE Trans. Med. Imaging* **38**(9), 2151–2164 (2019)
11. Duan, J., Qiu, Z., Lu, W., Wang, G., Pan, Z., Bai, L.: An edge-weighted second order variational model for image decomposition. *Digit. Sig. Process.* **49**, 162–181 (2016)
12. Fischer, B., Modersitzki, J.: Curvature based image registration. *J. Math. Imaging Vis.* **18**(1), 81–85 (2003)
13. Goldstein, T., O'Donoghue, B., Setzer, S., Baraniuk, R.: Fast alternating direction optimization methods. *SIAM J. Imaging Sci.* **7**(3), 1588–1623 (2014)
14. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artif. Intell.* **17**(1–3), 185–203 (1981)
15. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Proceedings of the 28th International Conference on Neural Information Processing Systems, vol. 2, pp. 2017–2025 (2015)
16. Lowekamp, B.C., Chen, D.T., Ibáñez, L., Blezek, D.: The design of SimpleITK. *Front. Neuroinform.* **7**, 45 (2013)
17. Lu, W., Duan, J., Qiu, Z., Pan, Z., Liu, R.W., Bai, L.: Implementation of high-order variational models made easy for image processing. *Math. Methods Appl. Sci.* **39**(14), 4208–4233 (2016)
18. Mok, T.C., Chung, A.: Fast symmetric diffeomorphic image registration with convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4644–4653 (2020)

19. Nesterov, Y.E.: A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In: *Doklady Akademii Nauk*, vol. 269, pp. 543–547. Russian Academy of Sciences (1983)
20. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds.) *MICCAI 2015. LNCS*, vol. 9351. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
21. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging* **18**(8), 712–721 (1999)
22. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.T.: A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: *2013 IEEE 10th International Symposium on Biomedical Imaging*, pp. 1219–1222. IEEE (2013)
23. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S72 (2009)
24. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *Int. J. Comput. Vis.* **97**(2), 229–241 (2012)
25. Wang, J., Zhang, M.: DeepFlash: an efficient network for learning-based medical image registration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4444–4452 (2020)
26. Yang, X., Kwitt, R., Styner, M., Niethammer, M.: Quicksilver: fast predictive image registration—a deep learning approach. *NeuroImage* **158**, 378–396 (2017)
27. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime TV- L^1 optical flow. In: Hamprecht, F.A., Schnörr, C., Jähne, B. (eds.) *DAGM 2007. LNCS*, vol. 4713, pp. 214–223. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74936-3_22
28. Zhang, M., Fletcher, P.T.: Fast diffeomorphic image registration via Fourier-approximated lie algebras. *Int. J. Comput. Vis.* **127**(1), 61–73 (2019)