LDDMEm: Large Deformation Diffeomorphic Metric Embedding

Decoupling Shape Analysis from Image Registration

Greg M. Fleishman[1][0000−0003−1879−6521] and P. Thomas Fletcher[2][0000−0003−3417−2380]

[1] HHMI, Janelia Research Campus fleishmang@janelia.hhmi.org
[2] University of Virginia

Abstract. We present a method, open-source software, and experiments which embed arbitrary deformation vector fields produced by any method (e.g., ANTs or VoxelMorph) in the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework. This decouples formal diffeomorphic shape analysis from image registration, which has many practical benefits. Shape analysis can be added to study designs without modification to already chosen image registration methods and existing databases of deformation fields can be reanalyzed within the LDDMM framework without repeating image registrations. Pairwise time series studies can be extended to full time series regression with minimal added computing. The diffeomorphic rigor of image registration methods can be compared by embedding deformation fields and comparing projection distances. Finally, the added value of formal diffeomorphic shape analysis can be more fairly evaluated when it is derived from and compared to a baseline set of deformation fields. In brief, the method is a straightforward use of geodesic shooting in diffeomorphisms with a deformation field as the target, rather than an image. This is simpler than the image registration case which leads to a faster implementation that requires fewer user derived parameters.

Keywords: LDDMM · Diffeomorphisms · Image Registration · Shape Analysis · Computational Anatomy · Manifold Statistics

1 Motivation

The Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework[5] is a beautiful example of advanced mathematics applied to a practical problem. In addition to registering two images it provides a rigorous guarantee of transform smoothness. Additionally, it can model real continuum mechanical properties of imaged materials[15], shape interpolate image time series of arbitrary length[14], and provide a rigorous setting in which to conduct shape statistics[8]. While LDDMM has been substantially developed in the literature, its adoption for practical use has not kept pace due to several barriers. For practical studies, robust and accurate estimation of correspondence between noisy images is of

primary importance. Non-LDDMM methods are less encumbered by advanced theoretical foundations, are less computationally complex, and their software development has focused on features that strengthen the capacity to accurately register noisy image data. Such features include multiple transformation models, image matching functionals, regularization schemes, multi-scale optimization, and fast deep learning models[18,2]. Even for studies that may benefit from using LDDMM, the greater availability, open-source developer mind share, and documentation of non-LDDMM methods may be more appealing than the benefits of formal shape analysis. Furthermore, deep learning has provided the community with a new set of image registration tools that are often very fast which is an appealing property for practical applications[6,3]. However, when test data does not match training data these methods can perform poorly with little ability to change their behavior other than more training, which may require a substantial investment. Deep learning methods also rarely come with theoretical guarantees or modeling properties such as those provided by LDDMM. This paper proposes a method to connect the theoretical benefits of LDDMM to the practical benefits of non-LDDMM image registration including deep learning.

Large Deformation Diffeomorphic Metric Embedding (LDDMEm) takes any displacement vector field and solves for a nearest approximation within the LDDMM framework. This capability effectively decouples image registration of noisy data from shape analysis. Images can be registered with any method. Later, shape analysis can be added to the study design without repeating the comparatively costly and challenging to parameterize image registrations. The effects of shape analysis can also be more fairly evaluated, as rather than requiring a different method entirely, they are derived from the existing study deformations. After a brief background, the theoretical framework for LDDMEm is developed. Synthetic and real data experiments are presented. Finally, open source LDDMEm code is described.

2 Background: LDDMM

LDDMM has been thoroughly described in the literature[7,12,13]. The most relevant background materials for the present work are Vialard et al.[17], which describes LDDMM from a Hamiltonian physics perspective and formulates an optimal control algorithm on the initial conditions of the dynamical system, and Singh et al.[16], which equivalently describes LDDMM from a Lie group perspective and relaxes the Vialard shooting algorithm to optimize vector instead of scalar valued initial momenta. LDDMM models the deformation of image $I \in L^2(\Omega, R)$ onto image J as a time dependent flow of diffeomorphisms. The following augmented Hamiltonian describes the entire system, and we will go through it term by term:

$$\begin{aligned} \widetilde{\mathcal{E}} \;&=\; \frac{1}{2}||v_0||_V^2 \;+\; \frac{1}{2\sigma^2}||I(\phi_{1,0}) - J||_{L^2}^2 \\ +\; \int_0^1 \langle \hat{v}, v - Km \rangle_{L^2} dt \;&+\; \int_0^1 \langle \hat{m}, \dot{m} + \mathrm{ad}_v^* m \rangle_{L^2} dt \;+\; \int_0^1 \langle \hat{I}, \dot{I} + \nabla I \cdot v \rangle_{L^2} dt \end{aligned} \tag{1}$$

It is appropriate to think of the first two terms as kinetic and potential energy thus comprising the Hamiltonian itself. The kinetic energy is the norm of velocity vector field v_0 in the space V, which is defined by its inner product: $||v_0||_V = \langle v_0, Lv_0 \rangle_{L^2} = \langle v_0, m_0 \rangle_{L^2} = \langle Km_0, m_0 \rangle_{L^2}$. Here, $m_0 = Lv_0$ is the momentum corresponding to velocity v_0. We have also used the common notation convention: $L^{-1} = K$. The operator L describes the assumed continuum mechanical properties of the spatial domain. It can be derived from real physical properties (e.g., elasticity or viscosity) or assumed, and typically has the form $L = (a\nabla^2 + b\nabla(\nabla\cdot) + c)^d$; the only requirement is that it be self-adjoint. Intuitively, L penalizes the derivatives of v_0, meaning smoother fields have smaller magnitudes.

The potential energy term quantifies the disagreement between the image I warped by a transform $\phi_{1,0}$, and the image J. Here, we write the sum of squared differences, but any differentiable measure is acceptable. The transform $\phi_{1,0}$ arises from a time dependent velocity flow, the integral of which establishes particle position relationships both forward and backward in time:

$$(a)\ \ \frac{d\phi_{0,t}}{dt} = v(\phi_{0,t}) \qquad (b)\ \ \frac{d\phi_{t,0}}{dt} = -\nabla\phi_{t,0} \cdot v \tag{2}$$

The image I is said to be observed at time $t = 0$. Its spatial domain flows according to $v(t)$. Consequently, the appearance of I at time $t = 1$, at the spatial positions upon which J is sampled, is given by $\phi_{1,0}$, which maps positions from time $t = 1$ to $t = 0$. The Hamiltonian terms are what we wish to optimize. We would like the velocity flow $\tilde{v}(t)$ that yields the transform $\tilde{\phi}_{1,0}$ such that the sum of kinetic and potential energy terms is minimal. Soon, we will constrain the problem such that only the initial velocity $\tilde{v}_0$ must be found.

The terms on the second line of equation (1) are augmentations that enforce constraints placed upon the system. The Lagrange multipliers $\hat{v}, \hat{m}$, and $\hat{I}$ are unconstrained. To recover the Hamiltonian, the right hand side of each inner product must be zero at all points in space and time, i.e., those equations must be satisfied. The first constraint is $v = Km$, which is always true by definition, but including this term simplifies later calculations so it is added for convenience. The second constraint, $\partial_t m = -\mathrm{ad}^*_v m$, is the most interesting and is a defining characteristic of LDDMM. Suppose the transform $\tilde{\phi}_{1,0}$ optimally matches image I to image J. We would like the arc length $\int_0^1 ||\tilde{v}(t)||_V dt$, which via equation (2b) terminates at the point $\tilde{\phi}_{1,0}$, to be minimal. This is precisely the definition of a geodesic, and a geodesic is uniquely determined by its tangent vector at $t = 0$. Thus, we are only required to specify $\tilde{v}_0$, so long as we constrain $\tilde{v}(t)$ to evolve without extrinsic acceleration. Naively, one might assume that the correct constraint would be $\partial_t v = 0$. However, recall that the flow is assumed to be in a continuum with mechanical properties determined by L. As particles displace, intrinsic strain or viscous forces develop and the particles accelerate. Such forces are proportional to the spatial derivatives of the displacement and the velocity. The operator is thusly defined $\mathrm{ad}^*_v m = (\nabla v)^T m + (\nabla m)v + (\nabla \cdot v)m$, and recall $m = Lv$. Intuitively, we can see that $\partial_t m = -\mathrm{ad}^*_v m$ accounts for such strain

or viscous forces through the spatial derivatives of v and m. More formally, the definition of $\mathrm{ad}^*_v m$ is a special case of the Euler-Poincaré equation for the Lie group of diffeomorphisms[1], the derivation of which is beyond the scope of this paper. The third constraint is $\partial_t I = -\nabla I \cdot v$. This is a consequence of equation (2b) and describes how the given image I flows along the geodesic $\phi_{t,0}$.

In summary, the augmented Hamiltonian equation (1) describes the energy of a dynamical system parameterized by an initial velocity v_0. The system is constrained to evolve along a geodesic path ending at the transform $\phi_{1,0}$. The stationary points of this functional optimally balance the magnitude of v_0 (with respect to L) with the image matching accuracy of $\phi_{1,0}$. Given I, J, and choosing values (a, b, c, d, σ^2), which determine L and the balance between kinetic and potential energies, we require an algorithm to determine $\tilde{v}_0$, or equivalently $\tilde{m}_0$, which minimizes (1). A gradient descent algorithm can be obtained by taking the variation of (1) with respect to the flows v, I, $\hat{v}$, $\hat{m}$, and $\hat{I}$ with the following result:

$$\partial_{m_0} \tilde{\mathcal{E}} = K m_0 - \hat{m}_0 \tag{3}$$

$$\begin{cases} \dot{\hat{m}} = (\nabla v)\hat{m} - (\nabla \hat{m})v + \hat{v} & \hat{m}_1 = 0 \\ \dot{\hat{I}} = -\nabla \cdot (\hat{I} v) & \hat{I}_1 = I(\phi_{1,0}) - J \\ L\hat{v} = \hat{I}\nabla I - \mathrm{ad}^*_{\hat{m}} m \end{cases} \tag{4}$$

Equation (3) requires $\hat{m}_0$ which can be obtained by integrating the system (4) backward in time.

3 Methods: LDDMEm

LDDMEm replaces image matching with displacement vector field matching. The potential energy term $||I(\phi_{1,0}) - J||^2_{L^2}$ is replaced with a similar term acting directly on vector fields: $||\phi_{0,1} - \phi^*||^2_{L^2}$ for the given field ϕ^*. Also, the third Lagrangian constraint $\int_0^1 \langle \hat{I}, \dot{I} + \nabla I \cdot v \rangle_{L^2} dt$ is replaced with the construction of the flow of transforms itself, namely equation (2a): $\int_0^1 \langle \hat{\phi}, \dot{\phi} - v(\phi) \rangle_{L^2} dt$. This results in the augmented Hamiltonian:

$$\begin{aligned} \widetilde{\mathcal{E}} \;&=\; \frac{1}{2}||v_0||^2_V \;+\; \frac{1}{2\sigma^2}||\phi_{0,1} - \phi^*||^2_{L^2} \\ +\; \int_0^1 \langle \hat{v}, v - Km \rangle_{L^2} dt \,+\, \int_0^1 \langle \hat{m}, \dot{m} + \mathrm{ad}^*_v m \rangle_{L^2} dt \;+\; \int_0^1 \langle \hat{\phi}, \dot{\phi} - v(\phi) \rangle_{L^2} dt \end{aligned} \tag{5}$$

With these changes we proceed taking variations as before to obtain a gradient descent algorithm with the following result:

$$\partial_{m_0} \tilde{\mathcal{E}} = K m_0 - \hat{m}_0 \tag{6}$$

$$\begin{cases} \dot{\hat{m}} = (\nabla v)\hat{m} - (\nabla \hat{m})v + \hat{v} & \hat{m}_1 = 0 \\ \dot{\hat{\phi}} = -(\nabla v)^T \hat{\phi} & \hat{\phi}_1 = \phi_{0,1} - \phi^* \\ L\hat{v} = -|\nabla \phi_{t,0}| \cdot \hat{\phi}(\phi_{t,0}) - \mathrm{ad}^*_{\hat{m}} m \end{cases} \tag{7}$$

Comparing (7) to (4) we see the differences. The adjoint variable $\hat{I}$ is replaced with $\hat{\phi}$ whose backward integration is an advection along v. In the $\hat{v}$ equation, $\hat{I}$ is replaced with $\hat{\phi}$ pulled from $t = 0$ to the current time. This pull of $\hat{\phi}$ is a reflection of the choice to model ϕ^* as $\phi_{0,1}$ instead of $\phi_{1,0}$. Since LDDMM entails the construction of both the forward and inverse transform flows, but we are given only one target transform, we are free to choose between the forward and inverse transform to compare to the given field. We chose $\phi_{0,1}$ because equation (2a) is more numerically stable than equation (2b). This system is optimized in the same way as the image matching case returning an initial momentum $\tilde{m}_0$ which determines a geodesic of transforms that passes through ϕ^*. With this embedding, the geodesic trajectory implied by ϕ^* can be interpolated, extrapolated, or averaged with other geodesics to study the deformation(s) in a formal shape space.

One interesting property of this formulation is that, unlike the noisy image matching case, the use of sum of squared differences to compare $\phi_{0,1}$ and ϕ^* is not only the simplest option but also sufficiently general since by definition the values of $\phi_{0,1}$ and ϕ^* are directly comparable. Another interesting property is the comparison to point matching formulations[11]. If ϕ^* is viewed as a dense set of coordinate matches, then LDDMEm is the generalization of point match based LDDMM methods to an infinite number of points, or more practically, a point match for every sample location.

4 Experiments

4.1 Synthetic

Two-dimensional synthetic experiments were designed to validate the ability of LDDMEm to embed displacement vector fields from multiple sources. Additionally, these experiments demonstrate one potential advantage of formal shape modeling: extrapolation accuracy. We begin with a smoothed image of an ellipse. A synthetic initial momentum field was generated and integrated to produce a geodesic of shape deformation sampled at time points $t = \{0, 1, 2, 16\}$. The shape at $t = 0$ (the ellipse) was registered as the moving image to the shape at $t = 2$ using three different commonly used packages: simpleITK[18], ANTs[2], and the pre-trained VoxelMorph SynthMorph shapes model[9]. These transforms were naively interpolated/extrapolated to the unseen time points $t = \{1, 16\}$ by multiplying the whole field by the appropriate scaling factor. The transforms were also embedded in the LDDMM framework with LDDMEm yielding initial momenta that integrate to close approximations of the given transforms at $t = 2$. These geodesics were integrated to the unseen time points $t = \{1, 16\}$ and compared to the naive interpolation/extrapolation results from scaling the given transforms directly. The results are presented in figure 1.

All three registration methods aligned the ellipse to the $t = 2$ shape accurately, with only a slight residual remaining in the VoxelMorph case. LDDMEm embedded all transforms accurately, reproducing them at sub-voxel accuracy.

For interpolation to $t = 1$, both the naive scaling and LDDMM geodesic integration methods were accurate. However, for extrapolation to $t = 16$, integration of the initial momentum produced a closer shape match than naive scaling in all three cases.

4.2 ADNI 2

To test LDDMEm on real samples we used a set of 350 subjects from the Alzheimer's Disease Neuroimaging Initiative phase 2 dataset (ADNI-2)[4]. For each subject we have four scans: screening, and 6-month, 12-month, and 24-month follow up times. The same three registration methods were used, however for these experiments we switched to the dense-brain-T1-3d-mse-32feat pretrained VoxelMorph model, which was specifically trained for registration of T1 MRI volumes like the ADNI-2 data we study here. For all samples, and all registration methods, we registered the screening scan as the moving image to the 12-month follow up scan. These transforms were naively interpolated/extrapolated by multiplying by the appropriate scaling factor to estimate the anatomy at the 6-month and 24-month time points. The transforms were also embedded using LDDMEm to obtain initial momenta parameterizing geodesic flows of transforms. These geodesics were integrated to estimate the anatomy at the 6-month and 24-month time points. The results are presented in Figure 2.

Curiously, for all three registration methods, the LDDMEm embedded transform matched the screening image to the 12-month image more accurately than the given transform, despite having never seen either image. We hypothesize this to be the effect of additional regularization, but have not yet determined an exact reason. The extrapolation effect observed in the synthetic experiments is only present for the ANTs transforms, albeit with a very small effect size. It may be that the time intervals used are too short to observe the improved extrapolation of geodesic modeling. We are interested to see how a registration to the 3-month image extrapolates with LDDMEm to longer time points such as 36-months. These experiments are ongoing and are not presented here.

5 Software

LDDMEm is available as an open source Python package at:
github.com/GFleishman/lddmem. At the time of writing, the package contains three modules: `main`, `epdiff`, and `cli`. The `main` module contains the function `lddmem` which is the primary method for embedding a transform. The function supports multi-resolution optimization in both space and time, masking to ignore regions of fields that should not be matched, automatic convergence criteria, and other helpful features. The `main` module also contains the function `forward_integration` which can be used to integrate an initial momentum to any time point. The `epdiff` module contains implementations of essential operators including L, K, $\mathrm{ad}_v m$, and $\mathrm{ad}^*_v m$. Finally, the `cli` module is a command line interface for calling `main.lddmem`. Future plans include offering Simple Geodesic Regression[10] and Principal Geodesic Analysis[8].

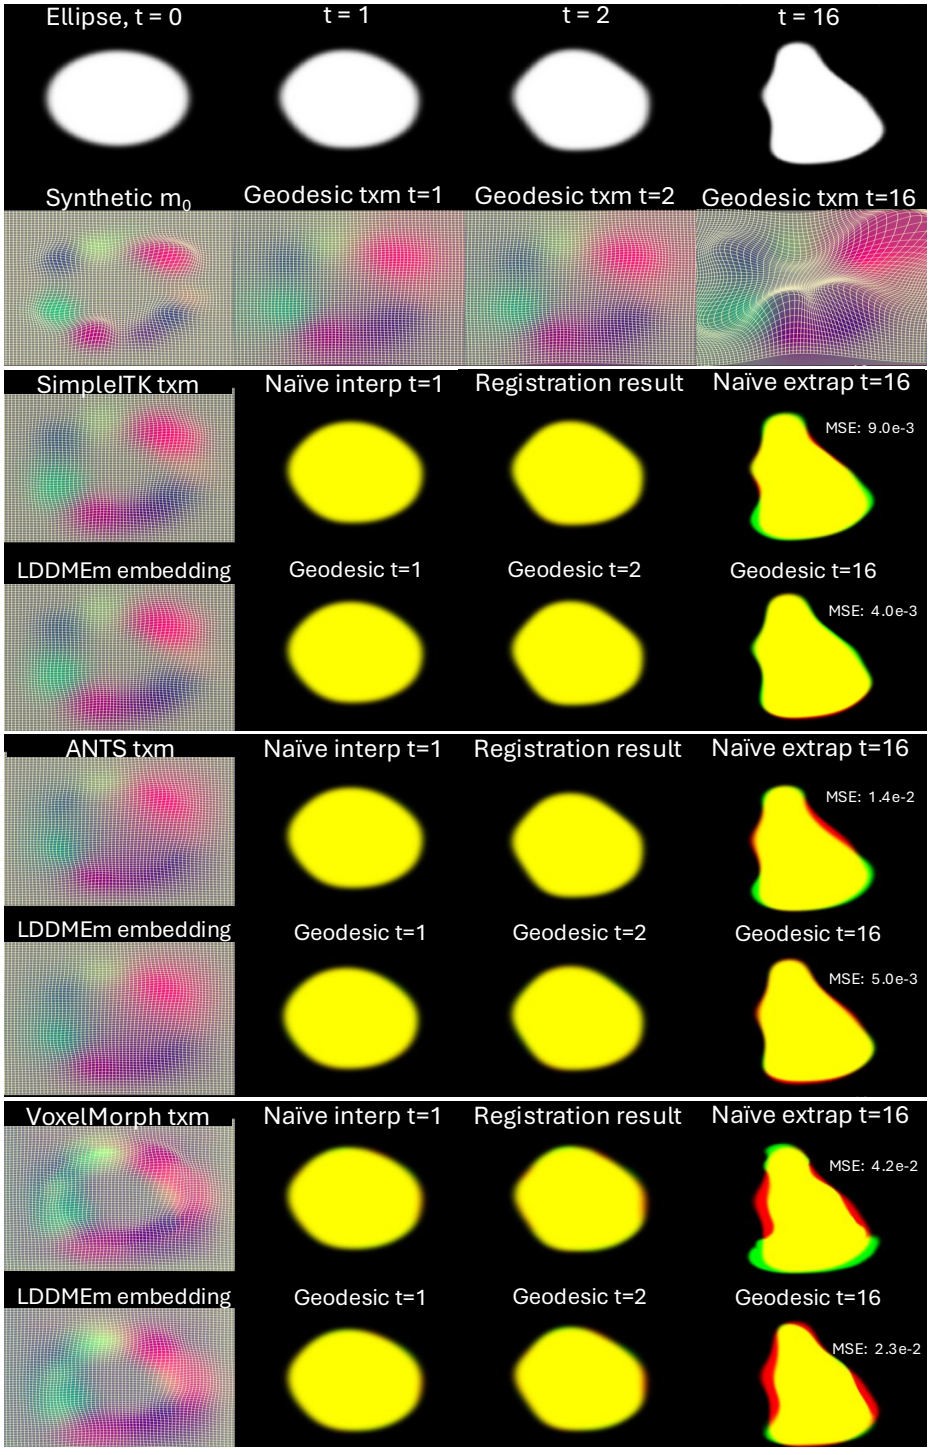

Fig. 1. Synthetic experiment results: LDDMEm successfully embeds transforms from three commonly used image registration packages and extrapolates shape change with greater accuracy. Target images in red, aligned images in green.

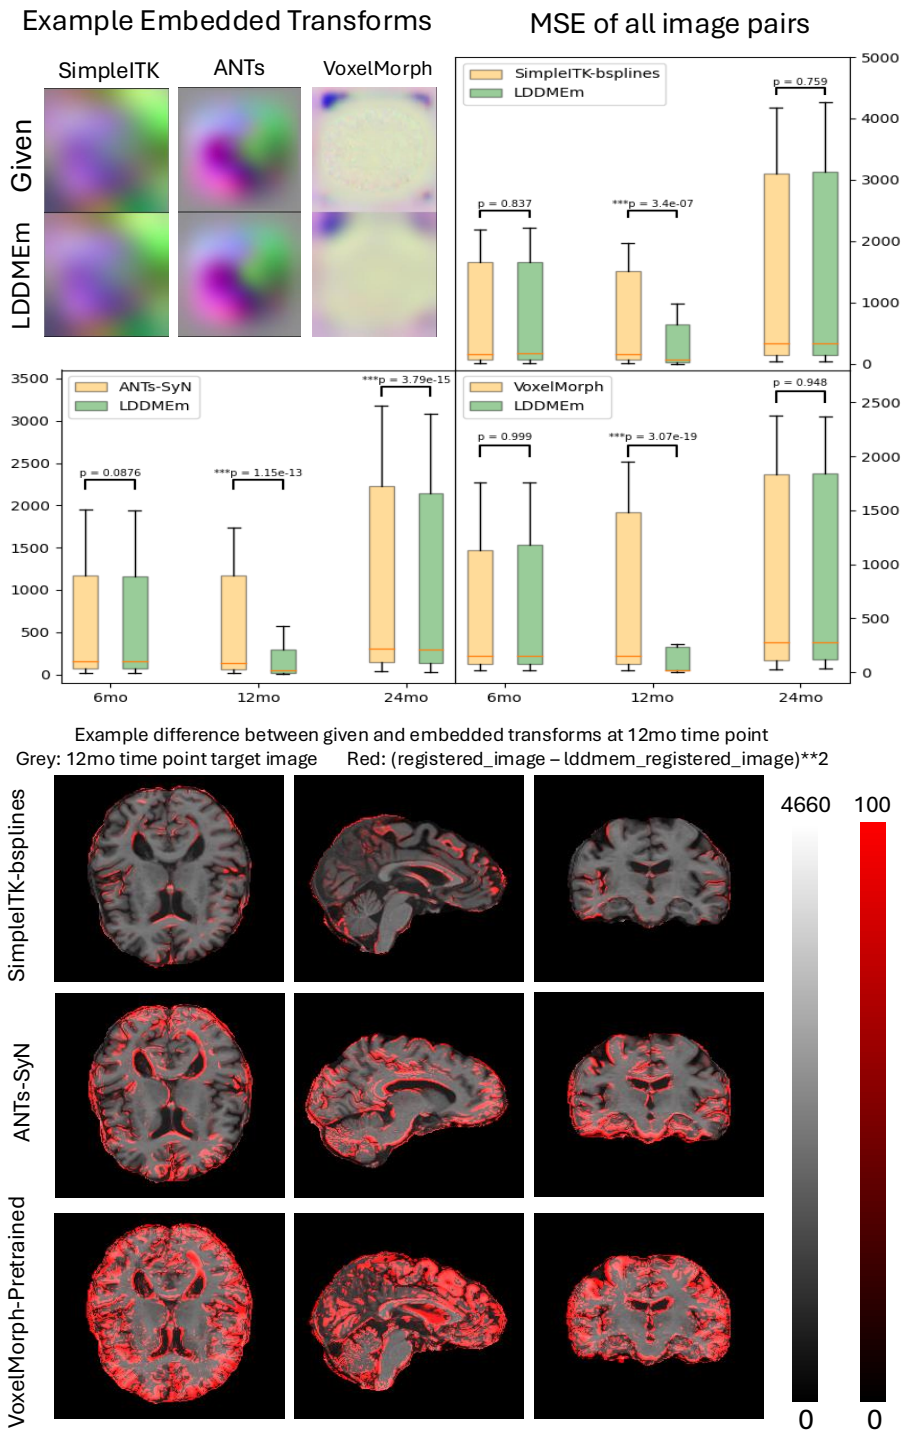

Fig. 2. ADNI2 experiment results: LDDMEm successfully embedded transforms generated from real data (upper left panel). Curiously, for all three registration methods, the LDDMEm embedded transforms align the original two images more accurately, with the largest differences observed for VoxelMorph (perhaps due to increased regularization, which can be seen visually in the upper left panel). The improved geodesic extrapolation observed for the synthetic data is only observed here for ANTs, albeit with a very small effect size.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arnold, V.: Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. Annales de l'Institut Fourier **16**(1), 319–361 (1966). https://doi.org/10.5802/aif.233, http://www.numdam.org/articles/10.5802/aif.233/
2. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. Medical Image Analysis **12**(1), 26–41 (2008). https://doi.org/https://doi.org/10.1016/j.media.2007.06.004, https://www.sciencedirect.com/science/article/pii/S1361841507000606
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Dalca, A.V., Guttag, J.: An unsupervised learning model for deformable medical image registration. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9252–9260 (2018). https://doi.org/10.1109/CVPR.2018.00964, http://doi.ieeecomputersociety.org/10.1109/CVPR.2018.00964
4. Beckett, L.A., Donohue, M.C., Wang, C., Aisen, P., Harvey, D.J., Saito, N.: The alzheimer's disease neuroimaging initiative phase 2: Increasing the length, breadth, and depth of our understanding. Alzheimers Dement **11**(7), 823–831 (Jul 2015). https://doi.org/10.1016/j.jalz.2015.05.004
5. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision **61**(2), 139–157 (2005). https://doi.org/10.1023/B:VISI.0000043755.93987.aa, https://doi.org/10.1023/B:VISI.0000043755.93987.aa
6. Ding, Z., Fleishman, G., Yang, X., Thompson, P., Kwitt, R., Niethammer, M.: Fast predictive simple geodesic regression. Medical Image Analysis **56**, 193–209 (2019). https://doi.org/https://doi.org/10.1016/j.media.2019.06.003, https://www.sciencedirect.com/science/article/pii/S1361841518301099
7. Dupuis, P., Grenander, U., Miller, M.: Variational problems on flows of diffeomorphisms for image matching. Quarterly of Applied Mathematics **56**(3), 587–600 (Sep 1998). https://doi.org/10.1090/qam/1632326
8. Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. IEEE Transactions on Medical Imaging **23**(8), 995–1005 (2004). https://doi.org/10.1109/TMI.2004.831793
9. Hoffmann, M., Billot, B., Greve, D.N., Iglesias, J.E., Fischl, B., Dalca, A.V.: Synthmorph: Learning contrast-invariant registration without acquired images. IEEE Trans Med Imaging **41**(3), 543–558 (Mar 2022). https://doi.org/10.1109/TMI.2021.3116879
10. Hong, Y., Shi, Y., Styner, M., Sanchez, M., Niethammer, M.: Simple geodesic regression for image time-series. In: Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D. (eds.) Biomedical Image Registration. pp. 11–20. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
11. Joshi, S., Miller, M.: Landmark matching via large deformation diffeomorphisms. IEEE Transactions on Image Processing **9**(8), 1357–1370 (2000). https://doi.org/10.1109/83.855431

12. Miller, M.I., Trouve, A., Younes, L.: On the metrics and euler-lagrange equations of computational anatomy. Annu Rev Biomed Eng **4**, 375–405 (2002). https://doi.org/10.1146/annurev.bioeng.4.092101.125733
13. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. J Math Imaging Vis **24**(2), 209–228 (Jan 2006). https://doi.org/10.1007/s10851-005-3624-0
14. Niethammer, M., Huang, Y., Vialard, F.X.: Geodesic regression for image time-series. In: Fichtinger, G., Martel, A., Peters, T. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011. pp. 655–662. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
15. Shen, Z., Vialard, F.X., Niethammer, M.: Region-specific diffeomorphic metric mapping. Adv Neural Inf Process Syst **32**, 1098–1108 (Dec 2019)
16. 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 (2013). https://doi.org/10.1109/ISBI.2013.6556700
17. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3d image registration via geodesic shooting using an efficient adjoint calculation. International Journal of Computer Vision **97**(2), 229–241 (2012). https://doi.org/10.1007/s11263-011-0481-8, https://doi.org/10.1007/s11263-011-0481-8
18. Yaniv, Z., Lowekamp, B.C., Johnson, H.J., Beare, R.: Simpleitk image-analysis notebooks: a collaborative environment for education and reproducible research. Journal of Digital Imaging **31**(3), 290–303 (2018). https://doi.org/10.1007/s10278-017-0037-8, https://doi.org/10.1007/s10278-017-0037-8