Selective Metamorphosis for Growth Modelling with Applications to Landmarks

Andreas Bock, Alexis Arnaudon(✉), and Colin Cotter

Department of Mathematics, Imperial College London, London SW7 2AZ, UK
alexis.arnaudon@imperial.ac.uk

Abstract. We present a framework for shape matching in computational anatomy allowing users control of the degree to which the matching is diffeomorphic. The control is a function defined over the domain describing where to violate the diffeomorphic constraint. The location can either be specified from prior knowledge of the growth location or learned from data. We consider landmark matching and infer the distribution of a finite dimensional parameterisation of the control via Markov chain Monte Carlo. Preliminary analytical and numerical results are shown and future paths of investigation are laid out.

Keywords: LDDMM · Computational anatomy · Metamorphosis · MCMC

1 Introduction

In computational anatomy [10,11] one of the most fundamental problems is to continuously deform an image or shape into another and thereby obtain a natural notion of distance between them as the energy required for such a deformation. Common methods to compute image deformations are based on diffeomorphic deformations which assume that the images are continuously deformed into one another with the additional property that the inverse deformation is also continuous. This is a strong requirement for images which implies that the 'mass' of any part of the image is conserved: we cannot create or close 'holes'. This is also a crucial property in fluid mechanics and in fact the theory of diffeomorphic matching carrying the moniker *Large Deformation Diffeomorphic Metric Mapping* (LDDMM) [5,23] has been inspired by fluid mechanics. Indeed, Arnold [4] made the central observation that the geodesic equations for the diffeomorphism group induced by divergence-free vector fields corresponded to that of incompressible flows. If a strictly diffeomorphic matching is not possible or necessary, an extension of LDDMM called metamorphosis [14,25] is available which introduces a parameter σ^2 parameterising the deviation from diffeomorphic matching

A. Arnaudon acknowledges EPSRC funding through award EP/N014529/1 via the EPSRC Centre for Mathematics of Precision Healthcare.

F. Nielsen and F. Barbaresco (Eds.): GSI 2019, LNCS 11712, pp. 39–48, 2019.
https://doi.org/10.1007/978-3-030-26980-7_5

allowing for topological variations e.g. growth via image intensity. In particular, if $\sigma^2 = 0$ the deformation is purely diffeomorphic as in LDDMM. See [18,22,24] for technical details pertaining to the construction of the metamorphosis problem. While diffeomorphic paths always exist for landmark problems [12] this theory allows one to match images of shapes with different topological features, which is ill-conditioned for standard LDDMM. Indeed, even inexact matching in LDDMM for such problems yields large energies and spurious geodesics that do not contribute to an intuitive matching, see Fig. 1. As observed here, introducing $\sigma^2 > 0$ regularises the problem and qualitative improves the matching.

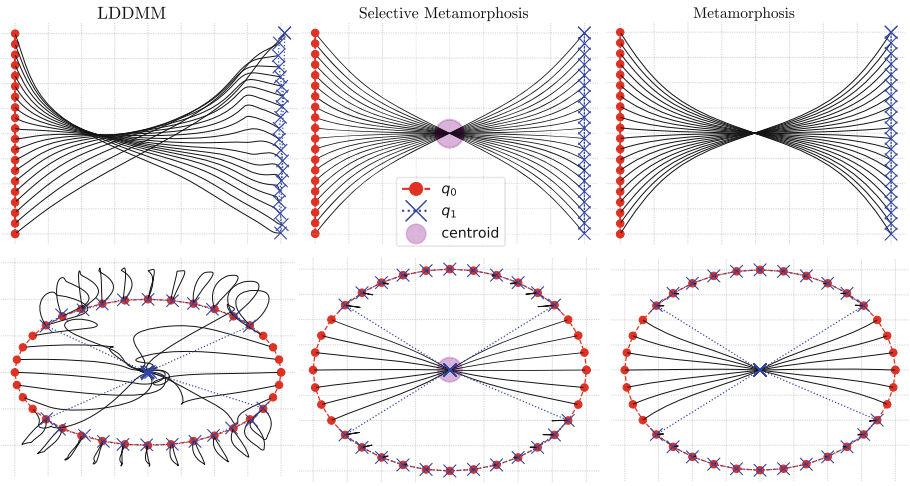

Fig. 1. This figure illustrates landmark matching with classical LDDMM (left column), metamorphosis (right column) and our selective metamorphosis approach (middle column). We perform a matching between two landmark configurations q_0 (circles) and q_1 (crosses), with the continuous lines between them describing trajectories. LDDMM fails to perform the matching and we observe unnatural landmark trajectories whereas metamorphosis achieves a more intuitive matching. Selective metamorphosis has the additional advantage of only breaking the diffeomorphic property where needed in along the matching, thus preserving more of the desired diffeomorphic property of the matching. These simulations where done for landmarks with Gaussian kernel of variance 0.5, 100 timesteps from $t = 0$ to $t = 1$, and a metamorphosis kernel of variance 0.2.

In this work, we modify metamorphosis to include a spatially dependent control parameter $x \mapsto \nu(x)$ in order to selectively allow non-diffeomorphic (*metamorphic*) matching in parts of the domain. For $\nu(\cdot) = \sigma^2$ our theory recovers the standard metamorphosis model. However, with a localised control (e.g. a Gaussian centred at a point in $\mathbb{R}^d$), we can selectively introduce metamorphosis in an image and model local topological effects such as growth phenomena. The difficulty of this problem is to infer the function ν without prior knowledge of

the location of the topological effects. This problem is similar to the one treated in [3], where such functions were parameterising the randomness in LDDMM matching of shapes. We will use a Markov chain Monte Carlo (MCMC) approach to infer appropriate functions ν, such that the topological effects are well described and a large part of the matching remains diffeomorphic. In this paper we focus on landmark matching but aim to extend the theory to data images.

2 Metamorphosis for Landmarks

In this paper we are concerned with diffeomorphometric approaches to shape matching. To this end, we use time-dependent velocity fields $u_t \in V$, where the Hilbert space V is continuously embedded in $\mathsf{C}_0^k(\mathbb{R}^d)$, $k \geq 1$. It induces a curve φ_t on a subgroup $\mathrm{Diff}_V(\mathbb{R}^d)$ of diffeomorphisms [4,26] via the equation

$$\dot{\varphi}_t = u_t \circ \varphi_t, \qquad \varphi_0 = \mathrm{id}. \tag{1}$$

This is used in the matching problem of two images I_0 and I_1 with cost

$$S(u) = \int_0^1 \frac{1}{2}\|u_t\|_V^2 \mathrm{d}t + \frac{1}{2\lambda^2} F(I_0 \circ \varphi_1^{-1}, I_1) \longrightarrow \text{min. subject to (1)}, \tag{2}$$

where F denotes a similarity measure between the deformed initial image $I_0 \circ \varphi_1$ and the target image I_1 to allow inexact matching parameterised by λ^2. The LDDMM approach takes F as an L^2 norm of the difference between its arguments. In this work, we will consider singular solutions for M landmarks with positions $q_t^i \in \mathbb{R}^d$ and momenta $p_t^i \in \mathbb{R}^d$ for $i = 1..M$ so that the velocity is

$$u_t(x) = \sum_{i=1}^{M} p_t^i K(x - q_t^i)\,, \tag{3}$$

where $K : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ is the kernel associated to the norm $\|\cdot\|_V$. For metamorphosis, we follow the notation of [14, Definition 1], and in addition to the deformation φ_t, we introduce a template variable $\boldsymbol{\eta}_t$ such that the positions $\mathbf{q}_t$ of a set of landmarks and the template velocity $\mathbf{z}_t$ are given by $\mathbf{z}$ as

$$\mathbf{q}_t = \varphi_t \boldsymbol{\eta}_t \qquad \text{and} \qquad \mathbf{z}_t = \varphi_t \dot{\boldsymbol{\eta}}\,. \tag{4}$$

We then extend the action functional (2) to account for the template variable as

$$S_m(\mathbf{q}_t, \mathbf{p}_t, \mathbf{z}_t) = \int_0^1 \frac{1}{2}\left(\|u_t\|_V^2 + \frac{1}{\sigma^2}\sum_{i=1}^{M} |z_t^i|^2\right) \mathrm{d}t, \tag{5}$$

where now the reconstruction relation is

$$\dot{\mathbf{q}}_t = u_t(\mathbf{q}_t) + \mathbf{z}_t, \tag{6}$$

see [14,25] for more details. By taking variations carefully, we obtain the equations of motion

$$\begin{aligned} \dot{\mathbf{p}}_t &= -\nabla u_t(\mathbf{q}_t)^T \mathbf{p}_t \\ \dot{\mathbf{q}}_t &= u_t(\mathbf{q}_t) + \sigma^2 \mathbf{p}_t, \end{aligned} \tag{7}$$

where $\mathbf{z}_t = \sigma^2 \mathbf{p}_t$ and u_t is defined in (3).

3 Selective Metamorphosis for Landmarks

We can now extend the metamorphosis setting to be able to locally control the amount of non-diffeomorphic evolution. For this, we introduce a function $\nu : \mathbb{R}^d \to \mathbb{R}$ replacing the parameter σ^2 such that $\nu(x) = \sigma^2$ corresponds to the classic landmark metamorphosis. The action for selective metamorphosis thus becomes

$$S^{\nu}_{sm}(\mathbf{q}_t, u_t, \mathbf{z}_t) = \int_0^1 \frac{1}{2} \left(\|u_t\|_V^2 + \sum_{i=1}^{M} \frac{1}{\nu(q_t^i)} |z_t^i|^2 \right) \mathrm{d}t, \tag{8}$$

which we minimise subject to the reconstruction Eq. (4) and the boundary conditions $\mathbf{q}_0$ and $\mathbf{q}_1$ at time $t = 0, 1$. In the case of landmarks we have as before that $\mathbf{z}_t = \nu(\mathbf{q}_t)\mathbf{p}_t$ so we can eliminate the template variable $\mathbf{z}_t$ and write

$$S^{\nu}_{sm}(\mathbf{q}_t, u_t, \mathbf{p}_t) = \int_0^1 \frac{1}{2} \left(\|u_t\|_V^2 + \sum_{i=1}^{M} \nu(q_t^i) |p_t^i|^2 \right) \mathrm{d}t. \tag{9}$$

The problem defined by (9) yields the following equations for selective metamorphosis for landmarks:

$$\begin{aligned} \dot{\mathbf{p}}_t &= -\nabla u_t(\mathbf{q}_t)^T \mathbf{p}_t - \frac{1}{2} \nabla \nu(\mathbf{q}_t) |\mathbf{p}_t|^2 \\ \dot{\mathbf{q}}_t &= u_t(\mathbf{q}_t) + \nu(\mathbf{q}_t)\mathbf{p}_t, \end{aligned} \tag{10}$$

with $\mathbf{q}_0$, $\mathbf{q}_1$ fixed. Again, the velocity is fully described by $\mathbf{p}$ and $\mathbf{q}$ via (3). As we see from these equations, our approach offers a granularity not attainable via classical inexact landmark matching or metamorphosis. Namely, with ν it is possible to specify where in the image growth is allowed. As an example, a medical expert may want to allow for metamorphic growth near a tumour-prone area of the brain whilst allowing for purely diffeomorphic growth of the skull of the patient.

A practical procedure for solving (10) with the velocity defined in (3) is called *shooting*, where we replace the end-point condition $\mathbf{q}_1$ with a guess for $\mathbf{p}_0$, and iteratively update $\mathbf{p}_0$ using automatically computed adjoint (or backward) equations until $\mathbf{q}_1$ compares to $\mathbf{q}(1)$ below a certain tolerance. We will perform this procedure directly with an automatic differentiation package Theano [21], see [16,17] for more details on the implementation.

Theorem 1. *Let ν be bounded from below away from zero by $\nu_{inf} \in \mathbb{R}$ and from above by $0 < \sigma^2 \in \mathbb{R}$. Then there exists a minimiser of (9) admissible to (6).*

Proof. The functional in (9) is not convex so we work with a reformulation to ensure the required lower semi-continuity. Define a variable $w_t^i = \sqrt{\nu(q_t^i)} p_t^i$ in the problem:

$$\inf_{\substack{u\in L^2([0,1],V)\\ \mathbf{q}\in H^1([0,1],\mathbb{R}^{d\times M})\\ \mathbf{w}\in L^2([0,1],\mathbb{R}^{d\times M})}} \int_0^1 \frac{1}{2}\left(\|u_t\|_V^2 + \sum_{i=1}^M |w_t^i|^2\right) \mathrm{d}t$$

$$\dot{\mathbf{q}}_t^i = u_t(\mathbf{q}_t) + \sqrt{\nu(\mathbf{q}_t)}\mathbf{w}_t$$
$$\mathbf{q}_0,\ \mathbf{q}_1 \text{ fixed}$$

First, note that owing to the constraint effectively being a boundary value problem, we cannot always find a $\mathbf{q}$ for arbitrary pairs of $(u,\ \mathbf{w})$. We define a bounded operator $(\mathbf{q},\ u_t) \mapsto \frac{\dot{\mathbf{q}}_t - u_t(\mathbf{q}_t)}{\sqrt{\nu(\mathbf{q}_t)}} \triangleq \mathbf{w}$:

$$\Big(\sum_{i=1}^M |w_t^i|^2\Big)^{\frac{1}{2}} = \|\mathbf{w}\|_2 = \|\frac{\dot{\mathbf{q}}_t - u_t(\mathbf{q}_t)}{\sqrt{\nu(\mathbf{q}_t)}}\|_2 \lesssim \nu_{\inf}^{-1}\Big(\|\dot{\mathbf{q}}_t\|_2 + \|u_t(\mathbf{q}_t)\|_V\Big)\,.$$

From this we generate a minimising sequence $(\mathbf{q}^n, u^n, \mathbf{w}^n)_{n\geq 0}$ admissible to (11). The rest of the proof is standard, see e.g. [26]. We show the constraint equation is continuous with respect to the weak topology on $X \triangleq H^1([0,1],\mathbb{R}^{d\times M}) \times L^2([0,1],\ V) \times L^2([0,1],\mathbb{R}^{d\times M})$ i.e. $e(\mathbf{q}_t^n,\ \mathbf{w}_t^n,\ u_t^n) \rightharpoonup e(\mathbf{q}_t,\ \mathbf{w}_t,\ u_t)$ where $e(q,\ w,\ u) \triangleq \dot{q} - u(q) - \sqrt{\nu(q)}w$. Then,

$$\langle\sqrt{\nu(\mathbf{q}_t)}\mathbf{w}_t - \sqrt{\nu(\mathbf{q}_t^n)}\mathbf{w}_t^n,\ \phi\rangle \lesssim \nu_{\inf}\langle\mathbf{w}_t - \mathbf{w}_t^n,\ \phi\rangle \to 0\,,\quad \forall\phi\in L^2([0,1],\mathbb{R}^{d\times M})\,.$$

Further, for $\phi \in L^2([0,1],\ V)$,

$$\langle u_t(\mathbf{q}_t) - u_t^n(\mathbf{q}_t^n),\ \phi\rangle = \langle u_t(\mathbf{q}_t) - u_t^n(\mathbf{q}_t),\ \phi\rangle + \langle u_t^n(\mathbf{q}_t) - u^n(\mathbf{q}_t^n),\ \phi\rangle\,.$$

The first term vanishes trivially, while for the second we see

$$\langle u_t^n(\mathbf{q}_t) - u_t^n(\mathbf{q}_t^n),\ \phi\rangle \leq \mathrm{Lip}(u_t^n)\langle\mathbf{q}_t - \mathbf{q}_t^n,\ \phi\rangle \to 0$$

Since linear operators are naturally compatible with the weak topology the required continuity follows. Passing to subsequences where necessary we can by classic results extract bounded subsequences converging to weak limits where necessary to obtain a minimiser. Convexity of S implies weak lower semi-continuity concluding the proof. □

Theorem 2. *Assume $\nu \in W^{2,\infty}(\mathbb{R}^d)$ and V is embedded in $C_0^k(\mathbb{R}^d)$, $k \geq 1$ (continuous functions with continuous derivatives to order k vanishing at infinity). Then, given $\mathbf{p}_0,\ \mathbf{q}_0, \in \mathbb{R}^{d\times M}$, (10) with (3) are integrable for all time.*

Proof. Establishing appropriate Lipschitz conditions implies integrability of the system akin to [6, Theorem 5]. We note that the kernel in (3) is Lipschitz in (p_t, q_t) by assumption, so the composition $(p,\ q) \mapsto u\circ q$ is also Lipschitz. $u(q) \mapsto \nabla u(q)^T$ consider $v,\ w \in V$ and $x,\ y \in \mathbb{R}^d$:

$$\|\nabla v(x) - \nabla w(y)\|_2 \lesssim \|v\|_V\|x-y\|_2 + \|v-w\|_V\|y\|_2 \tag{11}$$

so the mapping is Lipschitz in both the position and velocity. Given the conditions on ν the mappings

$$\begin{aligned}(q,\, p) &\mapsto \nu(q)p \\ (q,\, p) &\mapsto \nabla\nu(q)|p|^2\end{aligned} \tag{12}$$

are locally Lipschitz. Consequently we verify that for any $(\mathbf{p}_0,\, \mathbf{q}_0) \in B(0,r) \subset \mathbb{R}^{d\times M} \times \mathbb{R}^{d\times M}$, the system (10) is locally Lipschitz with constant L_{r,t_0} for some $t_0 > 0$. By the conservation of the Hamiltonian we can extend the existence of solutions to arbitary $t > t_0$. □

4 Bayesian Framework

We now place a stochastic model on ν inspired by the approach taken in [6] to infer most probable such functions. See also [1,2,20] for similar Bayesian approaches in computational anatomy. The goal is to develop an algorithm to infer ν from a given set of localised functions. We refer to [7,8] for an exposition of function space MCMC but we will consider a simpler case here. We consider ν as a sum of time-independent Gaussian functions

$$\nu_h(x) = \sum_{k=1}^{K} e^{-\sigma_k^{-2}\|h_k - x\|^2}. \tag{13}$$

This means that metamorphosis permitted in the neighbourhood of a point x (determined by the radius σ_k and *centroids* $h_k \in \mathbb{R}^2$ selected on the template) is proportional to the value of $\nu_h(x)$. As described in (10), ν_h follows the trajectory of the landmarks in the dynamics of $\mathbf{q}$. Note the number of landmarks, M, differ from the number of centroids K. For instance, we selected $K = 1$ in Fig. 1 due to our a priori knowledge of the trajectories (e.g. there is only a single point where landmarks cross or intersect).

Defining a density $p_{sm} \propto e^{-S^{\nu}_{sm}}$ over the space of triples $(\nu,\, \mathbf{q}_\nu,\, \mathbf{p}_\nu)$ leads to the preconditioned Crank-Nicholson Algorithm 1, see e.g. [13]. The parameter β has to be set such that the samples are un-correlated, which corresponds to an acceptance rate in the range 0.5–0.8.

In general, ν should accommodate the granularity of the deformation between two shapes and be able to resolve the topological changes necessary. This constitutes an interesting problem in and of itself, as it is *a priori* difficult to say what constitutes a good ν simply by inspecting the template and targets. Here we use Gaussians for their smoothness and simplicity, but we comment on extensions in Sect. 6.

5 Numerical Examples

This section displays some numerical results for our method to infer a distribution for the growth location using the landmark configurations seen in Fig. 1.

Algorithm 1. MCMC for selective metamorphosis

procedure MCMCSM(N, K, $\mathbf{q}_0$, $\mathbf{q}_1$, $\beta \in (0, 1]$)
  $j \leftarrow 1$
  $\nu^j \leftarrow$ initial guess in $\mathbb{R}^{d\times K}$
  Solve (10) with ν^j and $\mathbf{q}_0$, $\mathbf{q}_1$ to obtain $\omega^j = (\mathbf{q}^j, \mathbf{p}^j, u^j)$
  while $j < N$ **do**
    Sample a random point $\xi \in \mathcal{N}(0, \mathrm{Id}_{\mathbb{R}^d})^K$
    $\nu \leftarrow \beta\xi + \sqrt{1-\beta^2}\nu^j$
    Solve (10) with ν and $\mathbf{q}_0$, $\mathbf{q}_1$ to obtain $\omega = (\mathbf{q}, \mathbf{p}, u)$
    if RANDOMUNIT() $< \min(1, e^{-S^{\nu^j}_{sm}(\omega^j)+S^{\nu}_{sm}(\omega)})$ **then**
      $\nu^{j+1} \leftarrow \nu$
      $\omega^{j+1} \leftarrow \omega$
    else
      $\nu^{j+1} \leftarrow \nu^j$
      $\omega^{j+1} \leftarrow \omega^j$
    $j \leftarrow j + 1$
  return $\{\nu^j, \omega^j\}_{j=1}^N$

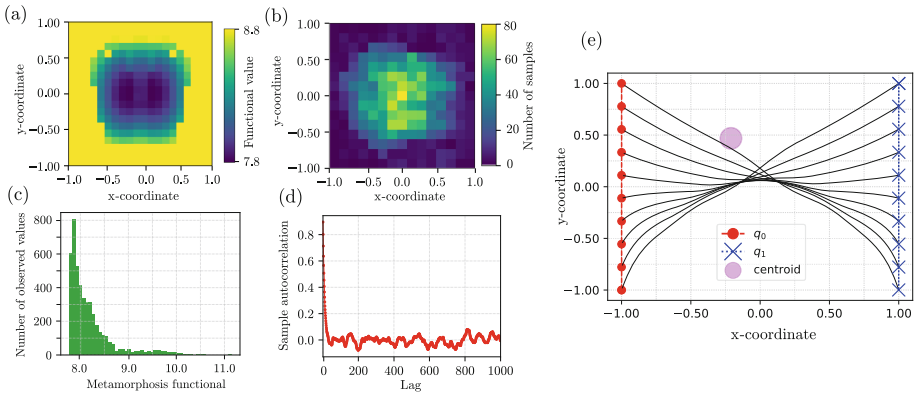

Fig. 2. We display the result of the MCMC Algorithm 1 applied to the inverted landmarks example of Fig. 1. (a) shows the analytical values for the functional (8) obtained for various positions of a single Gaussian ν. We observe a bimodal minimum near $(0, 0)$, which depends on the choice of the model parameters, and in particular on the landmark interaction length corresponding to the Gaussian kernel K and σ_ν. (b) displays a heat map for the sampled positions of the centroid from the MCMC method, where the bimodality is not clearly visible. (c) is histogram of the sampled values of the functional which rapidly decays, indicating a good sampling of the minimum value of the functional. (d) shows the autocorrelation function of the Markov chain, which decays rapidly to reach an un-correlated state after 50 iterations. (e) shows one of the MAP estimators where the centroid is near on the edge of one of the wells of the top left panel. The simulations parameters are set to $\sigma_\nu = 0.2$, and 0.7 for the velocity kernel, $K = 1$ and $\beta = 0.2$ across 5000 samples.

The parameters and results for the first configuration is shown in Fig. 2. These preliminary results demonstrate that even for a small number of samples the density of accepted samples corresponds at least heuristically to the analytical density histogram obtained by computing the value of the metamorphosis functional in (8).

We arrive at the same conclusion for the second example, for which the results are shown in Fig. 3. Moreover, we note that the geodesic equations for p and q are time-reversible meaning that the configuration in Fig. 3 corresponds to both particle collapse as well as hole creation. It is numerically relatively simple to control the behaviour of ν by simple scaling or by adding regularisation terms to (8) to e.g. penalise having ν's far away from the support of the images. Such cost can easily be added to the MCMC algorithm, depending on the prior information one can have on the shape matching problem.

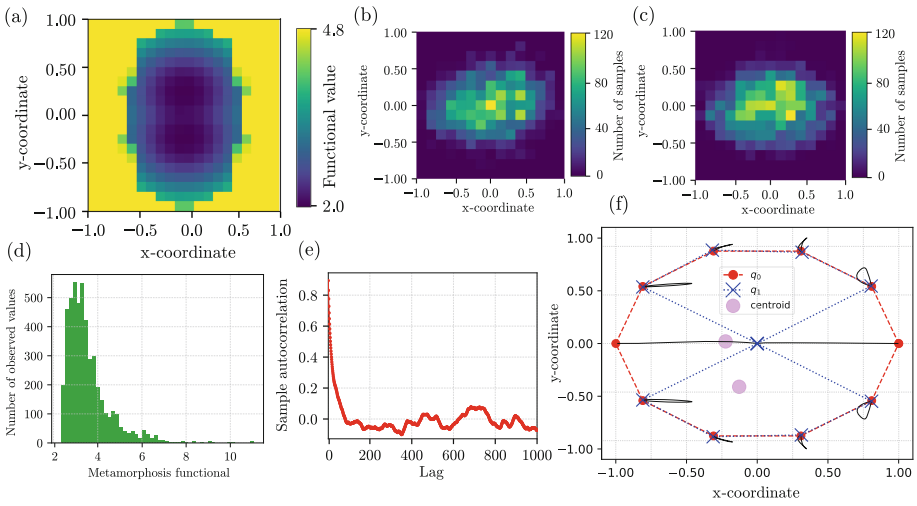

Fig. 3. Here we display the results for the second example (landmark collapse) of Fig. 1. Again, (a) shows the analytical values for a single ν field (8), which has also a bimodal structure, but in the other direction. For the MCMC we choose $K = 2$ Gaussian ν fields, and (b) and (c) displays two heat maps for the sampled positions of these centroids. (d) is a histogram of the sampled values of the functional, which has a peak at slightly higher values, possibly due to the redundant choice of two ν functions. (e) shows the autocorrelation function of the Markov chain which shows decorrelation after 100 steps. (f) shows the geodesics yielding one of the lowest functional values, where the two ν fields are close to each other, demonstrating the fact that only 1 would have been enough for this landmark configuration. The simulation parameters are the same as in Fig. 2 with the exception of $K = 2$.

6 Conclusion

We have presented a preliminary approach for selectively allowing photometric variation in a diffeomorphic image matching. We analysed the selective metamorphosis problem, the associated geodesic equations and demonstrated a proof of concept MCMC algorithm inferring a simple parameterisation of ν. This generalises LDDMM and metamorphosis and could provide a first-order exploratory tool for physicians to see if the development of a biological feature stems from a few violations of diffeomorphic evolution. This paper paves the way towards surgically investigating growth phenomena between topologically different images.

For future works we aim at extending the equations of Sect. 3 to images e.g. using the kernel framework in [19] or developing a space-time method.

In addition, there are many aspects of the probabilistic framework for the estimation of ν that need rigorous treatment and improvements. First, already in this simple setting, one would need to add additional penalties for the position of the centroids h_k to force them to remain for example near the centre of the image during the MCMC evolution. Second, natural extensions of our probabilistic approach by treating ν as a function could be considered and thus interpreting the resulting inverse problem through the appropriate measure-theoretical lens. Adding a time-dependency to ν can also be explored. Determining a truncated Fourier series of ν could lead to efficient numerical methods. Finally, we only used a simple MCMC algorithm, but a Metropolis-adjusted Langevin algorithm or Hamiltonian Monte-Carlo algorithm may be more appropriate to solve this problem.

To conclude, we hope that this framework could be used to model growth, in the spirit of the approaches of [15] or [9].

References

1. Allassonnière, S., Amit, Y., Trouvé, A.: Towards a coherent statistical framework for dense deformable template estimation. J. R. Stat. Soc. Ser. B (Stat. Methodol.) **69**(1), 3–29 (2007)
2. Allassonnière, S., Kuhn, E., Trouvé, A.: Map estimation of statistical deformable templates via nonlinear mixed effects models: deterministic and stochastic approaches. In: 2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy, pp. 80–91 (2008)
3. Arnaudon, A., Holm, D.D., Sommer, S.: A geometric framework for stochastic shape analysis. Found. Comput. Math. **19**, 653–701 (2018)
4. Arnold, V.I.: Sur la géométrie différentielle des groupes de lie de dimension infinie et ses applicationsa l'hydrodynamique des fluides parfaits. Ann. Inst. Fourier **16**(1), 319–361 (1966)
5. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. **61**(2), 139–157 (2005)
6. Cotter, C.J., Cotter, S.L., Vialard, F.-X.: Bayesian data assimilation in shape registration. Inverse Probl. **29**(4), 045011 (2013)

7. Cotter, S.L., Roberts, G.O., Stuart, A.M., White, D.: MCMC methods for functions: modifying old algorithms to make them faster. Stat. Sci. **28**, 424–446 (2013)
8. Dashti, M., Stuart, A.M.: The Bayesian approach to inverse problems. In: Ghanem, R., Higdon, D., Owhadi, H. (eds.) Handbook of Uncertainty Quantification, pp. 311–428. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-12385-1_7
9. Goriely, A.: The Mathematics and Mechanics of Biological Growth, vol. 45. Springer, New York (2017). https://doi.org/10.1007/978-0-387-87710-5
10. Grenander, U., Miller, M.I.: Representations of knowledge in complex systems. J. R. Stat. Society. Ser. B (Methodol.) **56**, 549–603 (1994)
11. Grenander, U., Miller, M.I.: Computational anatomy: an emerging discipline. Q. Appl. Math. **56**(4), 617–694 (1998)
12. Guo, H., Rangarajan, A., Joshi, S.: Diffeomorphic point matching. In: Paragios, N., Chen, Y., Faugeras, O. (eds.) Handbook of Mathematical Models in Computer Vision, pp. 205–219. Springer, Boston (2006). https://doi.org/10.1007/0-387-28831-7_13
13. Hairer, M., Stuart, A.M., Vollmer, S.J., et al.: Spectral gaps for a metropolis-hastings algorithm in infinite dimensions. Ann. Appl. Probab. **24**(6), 2455–2490 (2014)
14. Holm, D., Trouvé, A., Younes, L.: The Euler-Poincaré theory of metamorphosis. Q. Appl. Math. **67**(4), 661–685 (2009)
15. Kaltenmark, I.: Geometrical growth models for computational anatomy. Ph.D. thesis, Université Paris-Saclay (2016)
16. Kühnel, L., Arnaudon, A., Sommer, S.: Differential geometry and stochastic dynamics with deep learning numerics. arXiv preprint arXiv:1712.08364 (2017)
17. Kühnel, L., Sommer, S.: Computational anatomy in Theano. In: Cardoso, M.J., et al. (eds.) GRAIL/MFCA/MICGen 2017. LNCS, vol. 10551, pp. 164–176. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67675-3_15
18. Miller, M.I., Younes, L.: Group actions, homeomorphisms, and matching: a general framework. Int. J. Comput. Vis. **41**(1–2), 61–84 (2001)
19. Richardson, C.L., Younes, L.: Metamorphosis of images in reproducing kernel Hilbert spaces. Adv. Comput. Math. **42**(3), 573–603 (2016)
20. Schiratti, J.-B., Allassonnière, S., Colliot, O., Durrleman, S.: A Bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations. J. Mach. Learn. Res. **18**(1), 4840–4872 (2017)
21. Theano Development Team, Al-Rfou, R., et al.: Theano: a Python framework for fast computation of mathematical expressions. arXiv preprint arXiv:1605.02688 (2016)
22. Trouvé, A.: An infinite dimensional group approach for physics based models in pattern recognition. Preprint (1995)
23. Trouvé, A.: Diffeomorphisms groups and pattern matching in image analysis. Int. J. Comput. Vis. **28**(3), 213–221 (1998)
24. Trouvé, A., Younes, L.: Local geometry of deformable templates. SIAM J. Math. Anal. **37**(1), 17–59 (2005)
25. Trouvé, A., Younes, L.: Metamorphoses through Lie group action. Found. Comput. Math. **5**(2), 173–198 (2005)
26. Younes, L.: Shapes and Diffeomorphisms, vol. 171. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12055-8