Bayesian Mixed Effect Atlas Estimation with a Diffeomorphic Deformation Model

Stéphanie Allassonnière, S Durrleman, E Kuhn

▶ To cite this version:

Bayesian Mixed Effect Atlas Estimation with a Diffeomorphic Deformation Model[*]

S. Allassonnière[†], S. Durrleman[‡], and E. Kuhn[§]

Abstract. In this paper we introduce a diffeomorphic constraint on the deformations considered in the deformable Bayesian mixed effect template model. Our approach is built on a generic group of diffeomorphisms, which is parameterized by an arbitrary set of control point positions and momentum vectors. This enables us to estimate the optimal positions of control points together with a template image and parameters of the deformation distribution which compose the atlas. We propose to use a stochastic version of the expectation-maximization algorithm where the simulation is performed using the anisotropic Metropolis adjusted Langevin algorithm. We propose also an extension of the model including a sparsity constraint to select an optimal number of control points with relevant positions. Experiments are carried out on the United States Postal Service database, on mandibles of mice, and on three-dimensional murine dendrite spine images.

Key words. deformable template model, atlas estimation, diffeomorphic deformations, stochastic algorithm, anisotropic MALA, control point optimization, sparsity

AMS subject classifications. 60J22, 62F10, 62F15, 62F30, 62M40

DOI. 10.1137/140971762

1. Introduction. In this paper, we are interested in summarizing the variability observed in a collection of images by a small number of characteristics. Each image is considered as a different instance of the same "shape" or "object," for instance, the same digit written by different people or scans of the same bone observed in different individuals. This problem can be addressed in the framework of computational anatomy as introduced in [24]. The goal is to find a representative image of the object, called a template, and to quantify the observed variability in shape of this object by template-to-subject deformations. This geometric characteristic together with the template form the atlas. The corresponding model assumes that each observed image is a smooth deformation of the template plus an additive noise. The deformation is defined in the underlying space included in $\mathbb{R}^d$, where d equals 2 or 3 in applications. The template-to-subject deformations are used to quantify the geometric variability of the population via a metric on the set of characteristic mappings.

The study of the mathematical foundations of this deformable template model have been initiated in [23]. The characterization of the geometric variability has been addressed in different ways, for example, in [39] or [38]. This model could be considered from either a deterministic or stochastic point of view based on the idea that observations come from random

[†]Ecole Polytechnique, CMAP, 91128 Palaiseau, France (stephanie.allassonniere@polytechnique.edu).

[‡]Aramis lab, INRIA Paris Roquencourt, and Sorbonne Universités, UPMC univ Paris 06 UMR 1127, and CNRS UMR 7225 and ICM, 75013 Paris, France (stanley.durrleman@inria.fr). The research of this author was supported by ANR IRMGroup and the program "Investissements d'avenir" ANR-10-IAIHU-06.

[§]INRA Unité MIA, 78352 Jouy-en-Josas, France (estelle.kuhn@jouy.inra.fr).

deformations of the template image. Such approaches were developed among others in [32, 29, 28, 31, 13], and have demonstrated great impact in the field of image analysis. Models of deformations usually differ in the smoothness constraint of the mappings, which has to be adapted to the observations. To deal with the large geometric variability observed in real data, one could not restrict deformations to be rigid body and should consider nonrigid deformations that may have up to an infinite number of degrees of freedom. A simple model of deformations may be the so-called "linearized deformation model." A linearized deformation ϕ is defined by the displacement field v of each point in the domain $D \subset \mathbb{R}^d$: for all $r \in D, \quad \phi(r) = r + v(r)$. The main advantage of this class of deformations is its numerical simplicity since it parameterizes the deformation by a single vector field v. Nevertheless, even with regularity conditions on v, there is no guarantee that the deformation is invertible, which means that the deformation may create holes or overlapping regions in the domain. To avoid such unrealistic behaviors, one wants to consider diffeomorphic maps which preserve the topology of the shapes in the image set. This amounts to assuming that every sample has the same topology or, equivalently, that differences within sample shapes do not rely on changes of topology.

Diffeomorphic deformations can be built on linearized deformations in the framework of the large diffeomorphic deformation metric mapping (LDDMM), which has been introduced in [39, 16] and further developed among others in [26, 34, 12, 22, 25, 7]. In this framework, the above linearized deformations are considered as infinitesimal deformations, and the vector field v is seen as an instantaneous velocity field. The composition of such deformations creates a flow of diffeomorphisms, which can be written at the limit as the solution of a differential equation. The set of such diffeomorphisms can be equipped with a group structure and a right-invariant metric, providing regularity on the driving velocity fields. It follows that the set of images is given the structure of an infinite-dimensional manifold, on which distances are computed as the geodesic length in the deformation group between the identity map and the diffeomorphism that maps one image to another one.

It has been shown in [19] that this infinite-dimensional deformation set can be efficiently approximated by a finite control point parametrization carrying momentum vectors. This finite dimension reduction is a key aspect for statistical analysis. In [20], the authors have enforced the velocity fields that are defined everywhere in the domain to be parameterized by a finite set of control points. Positions of control points are not given as a prior but optimized as parameters of the statistical model. As a consequence, control points tend to move to the regions showing the largest variability among samples while optimizing a least-squares criterion. This optimization allows us, at the same time, to reduce the number of control points for the same matching accuracy, compared to the case where control points are fixed as the nodes of a regular lattice.

Once the deformation model has been fixed, one needs to estimate the parameters of the associated statistical model, including, in particular, the template image. Different algorithms have been proposed to solve the template estimation. Most of them are based on a deterministic gradient descent. In particular, in [20], the authors manage simultaneously the optimization in control point positions and momentum vectors thanks to a joint gradient descent. Although providing visually interesting results in several practical cases, the nature of the limit is not identified. Moreover, this type of method fails in specific cases, in particular, using noisy training data. Another point of view is to consider stochastic algorithms.

For example, in [42] the authors used a Hamiltonian Monte Carlo sampler in a Monte Carlo expectation maximization algorithm in the dense LDDMM setting; however, there is no theoretical convergence property proved for this algorithm. In the linearized deformation setting a convergent algorithm has been proposed in [6] to solve the atlas estimation issue. It is based on the stochastic approximation expectation maximization (SAEM) introduced in [17], and further extended using Monte Carlo Markov chain (MCMC) methods in [27, 6], thus allowing for a wider scope of applications. Concerning the geometric variability, one can adapt the kernel of the deformations using multiscale kernels such as in [36, 37] or learn the metric of the deformation group as in [41]. These methods are interesting as prior knowledge on the deformations. However, they do not take advantage of the simultaneous estimation of the template image and its variability.

In this paper, we aim at estimating an atlas of a population of images which is composed of a template image and a quantification of the geometric variability using the deformable template framework. We consider the LDDMM setting where the deformations are parameterized by a finite number of initial control point positions and momenta such as in [20] and extend the Bayesian mixed effect (BME) template model introduced in [2]. It is a statistical generative model which enables us to regenerate new synthetic samples that can highlight subtle behaviors in a given population. The Bayesian framework is particularly adapted to the high-dimensional low sample size paradigm, very common in biomedical imaging. In this model, the deformations are modeled as random variables which are not observed. This enables us to estimate the representative parameters of their distribution which will characterize the geometric variability. Moreover, we introduce the control point positions as population parameters into the model, so that they can be optimized in the estimation process. This enables us to better fit the deformation model leading to a more accurate estimation of the geometric parameters.

From an algorithmic point of view, we propose to use the anisotropic Metropolis adjusted Langevin algorithm (AMALA) within the SAEM algorithm introduced in [3] to estimate the model parameters. This algorithm has shown very interesting theoretical and numerical properties. Indeed, the AMALA sampler enables us to better explore the target distribution support in very high dimension space compared to other samplers. It also increases the speed of convergence of the estimation algorithm. Moreover, we take advantage in our sampler of the efficient computation used in the joint gradient descent in [20] so that the optimization of control point positions is of no additional cost at each iteration.

Another interesting question is how to optimize the number of control points required to parameterize the deformations. Indeed, the number of control points essentially depends on the variability in the data: it should be estimated rather than fixed by the user. In the geometric approach given in [20], control points were automatically selected using an L^1-type penalty that tends to zero out momentum vectors of small magnitude. Numerically it is solved by an adapted gradient descent called FISTA (see [11]). However, this penalty acts on each observation separately, meaning that a control point that is needed to match only a single observation will be kept in the final set of control points. From a statistical point of view, one may think about this control point as an outlier and would like to remove it from the basis. The L^1 penalty is also not suitable for statistical purposes, since its associated distribution, namely, the Laplace prior, does not generate sparse variables. In other words, the criterion

with L^1 penalty that is minimized in [20] could not be interpreted as the log-likelihood of a statistical model generating sparse solutions.

In this paper, we propose to include a sparsity constraint in the parameter space of our statistical model through a thresholding step, borrowing ideas from the Group LASSO literature initiated in [10]. This has the advantage of selecting control points based on their importance for the description of the variability of the whole population, and not only of one single sample. The thresholding step is then included in the maximization step, so that the same AMALA-SAEM algorithm can be used for the estimation process. We also exhibit a criterion to select an optimal threshold leading to an optimal number of control points.

This paper is organized as follows. We first recall the LDDMM setting in the case of control point parametrization in section 2. The BME template model derived for the atlas estimation issue is presented in section 3. The AMALA-SAEM algorithm is detailed in section 4. The extension toward sparsity is presented in section 5. Section 6 is devoted to experiments on handwritten digits, mouse mandible images, and three-dimensional (3D) murine dendrite spine images. Section 7 gives conclusions and perspectives for this work.

2. Model of diffeomorphic deformations.

2.1. Large deformation diffeomorphic metric mapping.
The model of diffeomorphic deformations we choose is derived from the LDDMM framework (see [39, 18, 32]), which generalizes the linearized deformation setting that has been used for the statistical estimation of atlases in [2]. In the linearized deformation setting, the deformation ϕ is given by

$$(2.1) \qquad \phi(r) = r + v(r) \quad \forall r \in D,$$

with $d = 2$ or 3, and v a vector field on $\mathbb{R}^d$.

It is convenient to assume that v belongs to a reproducible kernel Hilbert space (RKHS) in order to control its spatial smoothness. Essentially, the RKHS V is the set of the square integrable functions regularized by the convolution with a known kernel $K_g : V = L^2(D) * K_g$ (see [9] for more theoretical details.)

One may further assume that v is given as a finite linear combination of the RKHS basis vector fields, yielding

$$(2.2) \qquad v(r) = \sum_{k=1}^{k_g} K_g(r, c_k)\alpha_k,$$

where $(c_k)_k$ is a set of k_g geometric control points, $(\alpha_k)_k$ the vector weights attached to the control points (called momenta in this context), and K_g is the geometric interpolation kernel. The value of the vector field at any position r is obtained by interpolating the momenta located at the control points.

The advantage of this formulation is restricting the vector v to belong to a finite-dimensional subspace of the RKHS, which allows the straightforward definition and estimation of statistics such as means and covariance matrices. Another advantage is that v depends linearly on the momentum vectors, which greatly eases the derivation of the statistical model equations.

However, if the magnitude of the vector $v(r)$ or if the Jacobian of the vector field v becomes too large, then the linearized deformations are not invertible, thus leading to foldings or holes that may be undesirable for applications.

The LDDMM framework offers a generic way to define diffeomorphic maps, which guarantees their smoothness and invertibility. The approach introduced in [20] and [21] is a direct extension of the linearized deformation setting. It is still built on parametrization of the diffeomorphic maps lying in a finite-dimensional subspace of a given RKHS. However, the dependence of the deformations on their parameters will not be linear anymore.

2.2. Parametric LDDMM construction. To build a diffeomorphic map, we use the linearized deformations given in (2.1) as infinitesimal steps, and consider the corresponding vector field as an instantaneous velocity field. More precisely, we consider time-dependent velocity fields $(v_t)_t$ for a time parameter t varying in $[0, 1]$. The motion of a point r_0 in the domain of interest D describes a curve $t \mapsto r(t)$ which is the integral curve of the following ordinary differential equation (ODE) called the flow equation:

(2.3)
$$\begin{cases} \dfrac{\mathrm{d}r(t)}{\mathrm{d}t} = v_t(r(t)), \\ r(0) = r_0 \,. \end{cases}$$

The deformation ϕ_1 is defined as follows:

$$\forall r_0 \in D, \quad \phi_1(r_0) = r(1) \,.$$

Conditions under which this map ϕ_1 is diffeomorphic can be found in [12]. In particular, the existence, uniqueness, and diffeomorphic property of the solution are satisfied if the velocity v_t belongs to an RKHS at all time t and is square integrable in time.

Under these conditions, the model builds a flow of diffeomorphic deformations $\phi_t : r_0 \mapsto r(t)$ for all $t \in [0, 1]$. The flow describes a curve in a subgroup of diffeomorphic deformations starting at the identity map. The RKHS V plays the role of the tangent space of such an infinite-dimensional Riemannian manifold at the identity map Id. We can provide this group of diffeomorphisms with a right-invariant metric, where the square distance between the identity map Id $= \phi_0$ and the final deformation ϕ_1 is given as the total kinetic energy used along the path: $d(\mathrm{Id}, \phi_1)^2 = \int_0^1 \|v_t\|_V^2 dt$, where $\| \cdot \|_V$ is the norm in the RKHS. The existence and uniqueness of minimizing paths under regularity conditions on the velocity vector field have been shown in [32].

According to mechanical principles, one can show that the kinetic energy is preserved along the geodesic paths, namely, for all $t \in [0, 1]$, $\|v_t\|_V = \|v_0\|_V$. Moreover, the velocity fields (v_t) along such paths satisfy Hamiltonian equations, meaning that the geodesic is fully parameterized by the initial velocity field v_0. This velocity field plays the role of the Riemannian logarithm of the final diffeomorphism ϕ_1. Therefore, it belongs to a vector space and allows the definition of tangent-space statistics in the spirit of [40] and [35].

Following [21] and [20], we further assume that v_0 is the interpolation of momentum vectors $(\alpha_{0,k})_k$ at control point positions $(c_{0,k})_k$:

(2.4)
$$v_0(r) = \sum_{k=1}^{k_g} K_g(r, c_{0,k}) \alpha_{0,k} \,,$$

where K_g is the kernel associated with the RKHS V. In this context, it has been shown in [33] that the velocity fields $(v_t)_t$ along the geodesic path starting at the identity map in the

direction of v_0 keep the same form:

$$(2.5) \qquad v_t(r) = \sum_{k=1}^{k_g} K_g(r, c_k(t)) \alpha_k(t),$$

where the control point positions $(c_k(t))_k$ and the momentum vectors $(\alpha_k(t))_k$ satisfy the Hamiltonian equations

$$(2.6) \qquad \begin{cases} \dfrac{\mathrm{d}c_k(t)}{\mathrm{d}t} = \sum_{l=1}^{k_g} K_g(c_k(t), c_l(t)) \alpha_l(t), \\ \dfrac{\mathrm{d}\alpha_k(t)}{\mathrm{d}t} = -\dfrac{1}{2} \left(\sum_{l=1}^{k_g} d_{c_k(t)} (K_g(c_k(t), c_l(t)) \alpha_l(t)) \right)^t \alpha_k(t) \end{cases}$$

with initial conditions $c_k(0) = c_{0,k}$ and $\alpha_k(0) = \alpha_{0,k}$ for all $1 \le k \le k_g$. This is similar to the equations of motion of a set of k_g self-interacting particles, with K_g modeling the interactions. One can easily verify that the Hamiltonian defined as $H_t = \|v_t\|_V^2 = \sum_{k=1}^{k_g} \sum_{l=1}^{k_g} \alpha_l(t)^t K_g(c_l(t), c_k(t)) \alpha_k(t)$ is constant in time when control point positions and momentum vectors satisfy the system (2.6).

This model defines a finite-dimensional subgroup of the group of diffeomorphisms. For a given set of initial control points, the diffeomorphisms are parameterized by the momentum vectors attached to them. For one instance of the initial momentum vectors, one builds the motion of the control points and of the momentum vectors by integrating the Hamiltonian system (2.6). Then, they define a dense velocity field at each time t according to (2.5). Finally, one can find the motion $\phi_t(r_0)$ of any point r_0 in the domain D by integrating the flow equation (2.3). In this framework, the tangent-space representation of the diffeomorphic deformation ϕ_1 is given by the initial velocity field v_0 parameterized by $z = ((c_{0,k}, \alpha_{0,k}))_k$, called the initial state of the particle system. The position $\phi_1(r)$ depends on the parameters $((c_{0,k}, \alpha_{0,k}))_k$ via the integration of two nonlinear differential equations in (2.6) and (2.3).

Remark 1. The LDDMM framework formulation involves a coupling on the control point and the momentum evolutions along the geodesic path, which is not the case in the linearized deformation setting. This joint equation introduces more constraints reducing the dimension of the solution space. Therefore, the identifiability of the control point positions may be expected in our LDDMM framework. This property would most probably fail in the linearized deformation setting where the momenta and the control points are not coupled.

In section 3, we define the stochastic model of deformations based on parametric distributions of the initial state of particles.

3. Statistical model and parameter estimation. As pointed out in [2], the gradient descent optimization with respect to the template together with the momenta does not necessarily converge if the training set is noisy. To solve this problem, we propose here a statistical estimation algorithm in the spirit of [6] extending the BME template model to the LDDMM setting. This leads us to consider the deformations as well as the control point positions as nonobserved random variables.

3.1. Statistical generative model. We choose to model our data by a generative hierarchical model. This allows us to generate images from the atlas. The ability to generate synthetic images is important for interpreting the features captured by the model. It may highlight variability patterns that could not be perceived by simple visual inspection of the training images. In this model, the distribution of the deformations in the diffeomorphism group is parameterized. In a statistical approach, these parameters are estimated from the data, thus providing a metric in the shape space which is adapted to the data and takes into account the deformation constraints. This is in contrast to geometric approaches that estimate the template using a fixed metric.

More precisely, let I_0 be a template image $I_0 : \mathbb{R}^d \to \mathbb{R}$. We consider an observation, namely, an image y, as a noisy discretization on a regular grid Λ of a diffeomorphic deformation of the template image. Let ϕ_1^z be the solution of both the flow equation (2.3) and the Hamiltonian system (2.6) with initial condition $z = ((c_{0,k}, \alpha_{0,k}))_k$. Then, for all $s \in \Lambda$,

$$(3.1) \qquad\qquad y(s) = I_0((\phi_1^z)^{-1}(r_s)) + \sigma\epsilon(s),$$

where $\sigma\epsilon$ denotes an additive zero mean Gaussian random white noise on the grid Λ with covariance matrix $\sigma^2\mathrm{Id}$ and which is independent of all other variables, and r_s is the coordinate of the voxel s in the continuous domain D.

We are provided with n images $\mathbf{y} = (y_i)_{1 \leq i \leq n}$ in a training set. We assume that each of them follows the probabilistic model (3.1) and that they are independent.

We consider the initial state of particles, namely, the control point positions and the momentum vectors, as random variables and estimate their probabilistic distributions, restricting ourselves to the case of parametric distributions. We assume that control points live in the template domain D and that they are the same for all observations. By contrast, the momentum vectors attached to them are specific to each observation, as they parameterize the matching of the template with each sample image.

Therefore, we propose the following probabilistic model: we assume that the initial control point positions $\mathbf{c}_0 = (c_{0,k})_{1 \leq k \leq k_g}$ are drawn from a Gaussian distribution with mean $\bar{c}_0$ and covariance $a_c\mathrm{Id}$, where Id is the identity matrix of dimension dk_g. We choose a small value of a_c so that the control points are concentrated around their mean. We define the initial momenta $\boldsymbol{\alpha}_0 = (\boldsymbol{\alpha}_0^i)_{1 \leq i \leq n}$ with $\boldsymbol{\alpha}_0^i = (\alpha_{0,k}^i)_{1 \leq k \leq k_g}$. We assume that the variables $(\boldsymbol{\alpha}_0^i)_{1 \leq i \leq n}$ are independent identically distributed and follow a Gaussian distribution with mean 0 and covariance matrix Γ_g. Note that this covariance matrix depends on the initial control point positions, since the momenta are attached to them. Moreover the momenta $\boldsymbol{\alpha}_0$ are assumed to be independent of the control point positions $\mathbf{c}_0$ given Γ_g.

Following the same lines as in [2], we parameterize the template function I_0 as a linear combination of gray level values of fixed voxels $(b_k)_{1 \leq k \leq k_p}$ equidistributed on the domain D. The interpolation kernel is denoted by K_p (the subscript p refering to the photometric part of the model in contrast to g for the geometric one) and the combination weights are denoted by $\mathrm{w} \in \mathbb{R}^{k_p}$. Thus we have for all $r \in D$,

$$(3.2) \qquad\qquad I_0(r) = \sum_{k=1}^{k_p} K_p(r, b_k)\mathrm{w}_k \, .$$

The action of a diffeomorphism on this template is the linear combination of the deformed kernel with the same weights: for all $r \in D$,

$$(3.3) \qquad \mathbf{K}_{\mathrm{p}}^z \mathrm{w}(r) = I_0 \circ (\phi_1^z)^{-1}(r) = \sum_{k=1}^{k_p} K_p \left((\phi_1^z)^{-1}(r), b_k \right) \mathrm{w}_k \,.$$

The parameters of the model are $\theta = (\mathrm{w}, \sigma^2, \Gamma_g, \bar{\boldsymbol{c}}_{\mathbf{0}})$ and the random variables $(\boldsymbol{\alpha}_0, \mathbf{c}_0)$ are considered as hidden random variables. As we often deal with small sample size in practice, we restrict our inference to a Bayesian setting. Some of the priors can be informative like the one of Γ_g. Other priors may be noninformative, for instance, the expectation of the control point positions for which no additional information is available. The complete model is therefore written as

$$(3.4) \qquad \begin{cases} \theta = (\mathrm{w}, \sigma^2, \Gamma_g, \bar{\boldsymbol{c}}_{\mathbf{0}}) \sim \nu_p \otimes \nu_g, \\[2mm] \mathbf{c}_0 \sim \mathcal{N}_{dk_g}(\bar{\boldsymbol{c}}_{\mathbf{0}}, a_c \mathrm{Id})| \ \theta, \\[2mm] \boldsymbol{\alpha}_0^i \sim \mathcal{N}_{dk_g}(0, \Gamma_g)| \ \theta, \ \forall 1 \le i \le n \,, \\[2mm] y_i \sim \mathcal{N}_{|\Lambda|}(\mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \mathrm{w}, \sigma^2 \mathrm{Id}) \mid (\mathbf{c}_0, \boldsymbol{\alpha}_0^i), \ \theta, \ \forall 1 \le i \le n \,. \end{cases}$$

We define the prior distributions as follows:

$$\begin{cases} \nu_g(\mathrm{d}\Gamma_g, \mathrm{d}\bar{\boldsymbol{c}}_{\mathbf{0}}) \propto \left(\exp(-\langle \Gamma_g^{-1}, \Sigma_g \rangle_F / 2) \frac{1}{\sqrt{|\det(\Gamma_g)|}} \right)^{a_g} \\[2mm] \qquad\qquad\qquad \times \exp \left(-\frac{1}{2}(\bar{\boldsymbol{c}}_{\mathbf{0}} - \mu_c)^t \Sigma_c^{-1}(\bar{\boldsymbol{c}}_{\mathbf{0}} - \mu_c) \right) \mathrm{d}\Gamma_g \mathrm{d}\bar{\boldsymbol{c}}_{\mathbf{0}}, \\[2mm] \nu_p(\mathrm{dw}, \mathrm{d}\sigma^2) \propto \ \exp \left(-\frac{1}{2}\mathrm{w}^t \Sigma_p^{-1} \mathrm{w} \right) \cdot \left(\exp \left(-\frac{\sigma_0^2}{2\sigma^2} \right) \frac{1}{\sqrt{\sigma^2}} \right)^{a_p} \mathrm{dw}\mathrm{d}\sigma^2 \,, \end{cases}$$

where $\langle .,. \rangle_F$ designates the Frobenius scalar product and the hyperparameters satisfy $a_g \ge 4k_g + 1$, $\Sigma_g = \mathrm{Id}$, $\sigma_0^2 > 0$, $a_p \ge 3$, and Σ_p is derived from the interpolation kernel K_p and the photometric grid $(b_k)_{1 \le k \le k_p}$ (see [2] for more details). Concerning the hyperparameters of the control point prior (μ_c, Σ_c), we choose μ_c to be the vector of the equidistributed grid coordinates. The covariance matrix Σ_c is assumed noninformative. All priors are the natural conjugate priors and are assumed independent to ease derivations.

Remark 2. From a modeling point of view, the positions of the control points could have been considered as parameters of our model since they are fixed effects of the whole population as well as the template. However considering control points as parameters does not lead to a model belonging to the exponential family. Thus, we could not benefit from the convergence properties and efficient implementation of the SAEM algorithm for this class of models. Therefore, we model the control point positions as random variables following a Gaussian distribution.

3.2. Parameter estimation. Let us define $\mathbf{y} = (y_1, \dots, y_n)$. We consider the maximum a posteriori (MAP) estimator denoted by $\hat{\theta}_n$ obtained by maximizing the posterior density of θ

conditional to **y** as follows:

(3.5)
$$\hat{\theta}_n = \underset{\theta}{\mathrm{argmax}}\, q(\theta|\mathbf{y}).$$

For the sake of simplicity all likelihoods will be denoted by q. We will state the existence and the consistency of this MAP estimator in the next paragraphs.

We first show that for any finite sample the maximum a posteriori will lie in the parameter set Θ; this is nontrivial due to the highly nonlinear relationship between parameters and observations in the model.

Theorem 3.1. *For any sample* **y**, *there exists* $\hat{\theta}_n \in \Theta$ *such that* $q(\hat{\theta}_n|\mathbf{y}) = \sup_{\theta \in \Theta} q(\theta|\mathbf{y})$.

The proof is provided in the appendix in section 8.

We are also interested in the consistency properties of the MAP estimator without making strong assumptions on the distribution of the observations **y** denoted by P. We prove the convergence of the MAP estimator to the set Θ_* defined by

$$\Theta_* = \left\{\, \theta_* \in \Theta \;\middle|\; E_P(\log q(y|\theta_*)) = \sup_{\theta \in \Theta} E_P(\log q(y|\theta)) \right\}.$$

Theorem 3.2. *Assume that* Θ_* *is nonempty. Then, for any compact set* $K \subset \Theta$, *for all* $\varepsilon > 0$,

$$\lim_{n \to +\infty} P(\, \delta(\hat{\theta}_n, \Theta_*) \geq \varepsilon \,\wedge\, \hat{\theta}_n \in K\,) = 0\,,$$

where δ *is any metric compatible with the usual topology on* Θ.

The proof follows the lines of [2]. Indeed, the observed likelihood of our diffeomorphic BME template model has the same regularity properties and asymptotic behaviors in the parameters as the linearized one.

Remark 3. In [2], the authors have proved that under a weak additional condition, Θ_* is not empty. This makes the use of an important property of the linearized deformations: the amplitude of the deformation increases as the amplitude of its coefficients increases. This enables us to prove that large amplitude deformations would not be suitable to optimize the observed likelihood. In the LDDMM setting, this property cannot be guaranteed anymore. The relation between the range of the deformation and its momenta depends on the curvature of the diffeomorphisms space (which is flat in the linearized deformation framework). Therefore, proving that Θ_* is not empty will require us to know the curvature of the deformation space. This is unfortunately not known except in very simple cases (see [30]) preventing a direct generalization of the previous proof.

4. Algorithmic method. In this section, we detail the estimation algorithm chosen to maximize the posterior distribution $q(\theta|\mathbf{y})$ in the parameter θ. We use the SAEM algorithm introduced in [17] coupled with a MCMC method as suggested in [27] and [6]. Let us first recall the principles of the SAEM-MCMC algorithm in the general case of a model belonging to the curved exponential family. This algorithm is iterative, each iteration consisting in four steps:

Simulation step. The missing data (here positions of initial control points and momentum vectors) are drawn using a transition probability of a convergent Markov chain Π_θ having the conditional distribution $\pi_\theta(.) = q(.|\mathbf{y}, \theta)$ as stationary distribution.

Stochastic approximation step. A stochastic approximation is carried out on the sufficient statistics of the model using the simulated value of the missing data and a decreasing sequence of positive step sizes.

Projection on random boundaries. If the result of the stochastic approximation falls outside a compact set of increasing size, it is projected back to a fixed compact set.

Maximization step. The parameters θ are updated by maximizing the complete log-likelihood evaluated in the projected sufficient statistics.

Due to the high dimension of the hidden variables, we need to pay attention to the sampling step, as detailed in subsection 4.1. The details of the other steps of the algorithm are presented in subsection 8.2 in the appendix. The theoretical properties of this estimation algorithm are discussed in subsection 4.2.

4.1. Simulation step of the stochastic expectation maximization algorithm.

4.1.1. AMALA sampler. In our applications, the missing variables composed of the initial momenta and positions of control points $\mathbf{z} = (\mathbf{c}_0, \boldsymbol{\alpha}_0)$ are of very high dimension. In this case, the AMALA sampler proposed in [3] seems better suited for our stochastic extectation maximization algorithm than more standard samplers. For example, the Gibbs sampler solves the problems of low numerical acceptation rate and trapping states by looping over each coordinate to better stride the target density support. However, this involves a huge number of loops and heavy computations in the acceptation ratio preventing from any use in very high dimensions. By contrast, the AMALA sampler is more performant in terms of computational time while exploring the target support as well as the Gibbs sampler.

To be more precise, the AMALA sampler is an anisotropic version of the well-known Metropolis adjusted Langevin algorithm, where the covariance matrix of the proposal is optimized to take into account the anisotropy and the coordinate correlations of the target distribution. Using our previous notation, the drift vector denoted by $D_\theta(\mathbf{z})$ is equal to

$$(4.1) \qquad D_\theta(\mathbf{z}) = \frac{b}{\max(b, |\nabla \log \pi_\theta(\mathbf{z})|)} \nabla \log \pi_\theta(\mathbf{z}),$$

with $b > 0$ a truncation boundary. This vector $D_\theta(\mathbf{z})$ is the concatenation of the truncated gradients with respect to $\mathbf{c}_0$ and $(\boldsymbol{\alpha}_0^i)_{1 \leq i \leq n}$ denoted, respectively, by $D_\theta^0(\mathbf{z}), D_\theta^1(\mathbf{z}), \ldots, D_\theta^n(\mathbf{z})$.

Starting from the current value $\mathbf{z}_k$ of the Markov chain, the candidate $\mathbf{z}_c$ is sampled from the Gaussian distribution with expectation $\mathbf{z}_k + \delta D_\theta(\mathbf{z}_k)$ and covariance matrix $\delta \Sigma_\theta(\mathbf{z}_k)$, where $\Sigma_\theta(\mathbf{z})$ is given as

$$(4.2) \qquad \Sigma_\theta(\mathbf{z}) = \varepsilon \mathrm{Id}_{(n+1)dk_g} + diag\left(D_\theta^0(\mathbf{z})D_\theta^0(\mathbf{z})^t, D_\theta^1(\mathbf{z})D_\theta^1(\mathbf{z})^t, \ldots, D_\theta^n(\mathbf{z})D_\theta^n(\mathbf{z})^t\right)$$

with $\varepsilon > 0$ a small regularization parameter and $\delta > 0$.

We denote by p_θ the probability density function of this proposal distribution and by $\rho_\theta(\mathbf{z}_k, \mathbf{z}_c)$ the acceptance rate defined as

$$(4.3) \qquad \rho_\theta(\mathbf{z}_k, \mathbf{z}_c) = \min\left(1, \frac{\pi_\theta(\mathbf{z}_c)p_\theta(\mathbf{z}_c, \mathbf{z}_k)}{p_\theta(\mathbf{z}_k, \mathbf{z}_c)\pi_\theta(\mathbf{z}_k)}\right).$$

Then, the new value $\mathbf{z}_{k+1}$ of the Markov chain equals $\mathbf{z}_c$ with probability $\rho_\theta(\mathbf{z}_k, \mathbf{z}_c)$, and $\mathbf{z}_k$ otherwise.

Remark 4. For numerical efficiency, we do not take into account correlations between the momenta and the control point positions in the proposal. Moreover, the observations being independent, the covariance matrix of the momenta is block diagonal since the momenta are independent conditionally of the control point positions.

We now move to the computation of the gradient of the conditional distribution logarithm which appears in the drift D_θ. It happens that the conditional distribution logarithm is actually equaled to minus the usual energy used to compute the best match between images in the LDDMM framework. Therefore, we pay attention to the computation of the gradient of this quantity in the following paragraph.

4.1.2. Gradient computation in the LDDMM deformation model. We recall here the result established in [20]. For clarity purposes, we adopt compact matrix notation. The initial state of the system, which consists of the initial positions of control points $\mathbf{c}_0$ and their n associated momentum vectors $\boldsymbol{\alpha}_0$ for all n subjects is denoted by $\mathbf{z} = (\mathbf{c}_0, \boldsymbol{\alpha}_0)$. The position of this set of particles at later time t is denoted by $\mathbf{z}(t) = (\mathbf{c}_0(t), \boldsymbol{\alpha}(t))$, and satisfies the set of coupled ODEs (2.6). This system of ODEs can be rewritten in short as

$$(4.4) \qquad \begin{cases} \dot{\mathbf{z}}(t) = F(\mathbf{z}(t)), \\ \mathbf{z}(0) = \mathbf{z}. \end{cases}$$

Let $X^z(t,.)$ denote the mapping $r \in D \to X^z(t,r) = \phi_t^z((\phi_1^z)^{-1}(r))$. For $t = 1$, $X^z(1,.) = \mathrm{Id}_{L^2(D)}$ is the identity map. For $t = 0$, $X^z(0,.) = (\phi_1^z)^{-1}(.)$ is the inverse mapping of the domain D that is needed to deform the images. The interest in using the flow $\phi_t^z \circ (\phi_1^z)^{-1}$ (and not $(\phi_t^z)^{-1}$, for instance) is that the trajectory of any pixel under this flow in D is exactly the same as for the direct flow, but in the reverse direction. More precisely, $X^z(t,.)$ is a solution of the following ODE integrated backward from $t = 1$ to $t = 0$:

$$(4.5) \qquad \begin{cases} \dfrac{\partial X^z(t,.)}{\partial t} = -v_t(X^z(t,.)) = -\sum_{k=1}^{k_g} K_g(X^z(t,.), c_k(t)) \alpha_k(t), \\ X^z(1,.) = \mathrm{Id}_{L^2(D)}, \end{cases}$$

which can be rewritten in short as

$$(4.6) \qquad \begin{cases} \dfrac{\mathrm{d} X^z(t,r)}{\mathrm{d} t} = G(X^z(t,r), z(t)) \\ X^z(1,r) = r \end{cases}$$

for all $r \in D$.

The solution at $t = 0$, $X^z(0,.)$, is used to deform the template image I_0:

$$(4.7) \qquad I_0((\phi_1^z)^{-1}(r)) = I_0(X^z(0,r)).$$

From a numerical point of view, we discretize the image domain D into an array of $|\Lambda|$ pixels. The map $X^z(t,.)$ is therefore a vector of dimension $d|\Lambda|$: $\{X^z(t,r_s)\}_{s=1,\ldots,|\Lambda|}$, which

gives the trajectory of any pixel r_s under the flow equation (4.5) (where G is a map from $\mathbb{R}^{d|\Lambda|} \times \mathbb{R}^{dn_g}$ to $\mathbb{R}^{d|\Lambda|}$). The gray value of the deformed image at pixel r_s is computed by interpolating the gray values in I_0 located at pixels around position $X^z(0, r_s)$ using (3.2).

Proposition 1. *Let us denote by* $\mathbf{z} = (\mathbf{c}_0, \boldsymbol{\alpha}_0)$ *the* $(n+1)dk_g$ *parameters of the following criterion E_θ:*

$$E_\theta(\mathbf{z}) = A(X^{\mathbf{z}}(0, .)) + L(\mathbf{z}),$$

where

(4.8)
$$A(X^{\mathbf{z}}(0, .)) = \frac{1}{\sigma^2} \sum_{i=1}^{n} \|y_i - I_0((\phi_1^{z_i})^{-1})\|^2,$$

$$L(\mathbf{z}) = \sum_{i=1}^{n} (\boldsymbol{\alpha}_0^i)^t \Gamma_g^{-1} \boldsymbol{\alpha}_0^i,$$

$$\dot{\mathbf{z}}(t) = F(\mathbf{z}(t)), \qquad\qquad \mathbf{z}(0) = \mathbf{z},$$
$$\dot{X}^{\mathbf{z}}(t, .) = G(X^{\mathbf{z}}(t, .), \mathbf{z}(t)), \qquad\qquad X^{\mathbf{z}}(1, .) = Id_{L^2(D)},$$

and $X^{\mathbf{z}}(t, .) \in L^2(D, \mathbb{R}^d)$ for all t and A is a differentiable map from $L^2(D, \mathbb{R}^d)$ to $\mathbb{R}$.

Then, the gradient of E_θ is given by

(4.9)
$$\nabla_{\mathbf{z}} E_\theta = \xi(0) + \nabla_{\mathbf{z}} L,$$

where two auxiliary variables $\xi(t)$ (in $\mathbb{R}^{(n+1)dk_g}$) and $\eta(t, .)$ (in $L^2(D, \mathbb{R}^d)$) satisfy the following linear ODEs:

(4.10)
$$\begin{cases} \dot{\eta}(t, .) = -(\partial_1 G(X^{\mathbf{z}}(t, .), \mathbf{z}(t)))^* \eta(t, .), \\ \eta(0, .) = -\nabla_{X^{\mathbf{z}}(0, .)} A, \end{cases}$$

(4.11)
$$\begin{cases} \dot{\xi}(t) = -\partial_2 G(X^{\mathbf{z}}(t, .), \mathbf{z}(t))^* \eta(t, .) - d_{\mathbf{z}(t)} F^t \xi(t), \\ \xi(1) = 0, \end{cases}$$

where * denotes the adjoint operator in $L^2(D, \mathbb{R}^d)$.

This proposition states that the gradient is computed by integrating two linear ODEs that couple the information in the initial momenta and in the initial control points. Computing the gradient only with respect to the initial momenta does not decrease the computation time. The coupling implies that the gradients with respect to each coordinate of the hidden variables are computed simultaneously. The expression in coordinates of the terms in Proposition 1 as well as its proof can be found in [20].

Remark 5. The gradient of E_θ is indeed equal to $-\nabla_{\mathbf{z}} \log \pi_\theta$. This vector belongs to the tangent space of the Riemannian manifold of the data at the template point. Thus it provides crucial information about the local shape of this manifold. Therefore, it is of great interest to include this quantity into the estimation algorithm. It has been done from a deterministic point of view in [20] through a gradient descent on E_θ. Nevertheless, this algorithm may be stuck into local minima of this energy E_θ. To avoid such behaviors, stochastic algorithms are

well-known powerful tools. In particular, the AMALA-SAEM combines both advantages: it is a stochastic algorithm whose samples are based on this gradient direction.

Remark 6. The AMALA-SAEM requires the computation of this gradient at each iteration of the algorithm. Its numerical cost has to be compared with the cost when using other samplers. The hybrid Gibbs sampler that was chosen in [6] for the linearized deformation may also be used here. Although it does not require us to compute the gradient, it needs to loop over each coordinate of the hidden variables. This loop in our generalized LDDMM model would involve an integration of the Hamiltonian system (2.6) and the flow equation (2.3) for each coordinate. This scheme would be particularly inefficient due to the coupling between all control points and initial momenta in these equations: one would need to compute the full set of coupled equations each time one updates a coordinate. For k_g control points in dimension d and n observations, the hidden variable is of dimension $(n + 1)dk_g$. The Gibbs sampler needs then to integrate $(n + 1)dk_g$ times the Hamiltonian system (2.6) and the flow equation (2.3) which are differential systems in dimension $(n + 1)dk_g$ and $d|\Lambda|$. The AMALA only requires a single one of this step and the gradient computation which involves two differential equations (4.10) and (4.11), both in dimension $(n + 1)dk_g$. Although the differential equations in the gradient are more complex than those in the Hamiltonian system (as they require the Hessian of the kernel and not only its gradient), the AMALA sampler is still much more efficient for random variables of very high dimension than the Gibbs sampler.

We summarize the complete estimation procedure in Algorithm 1.

4.2. Discussion on theoretical properties. The AMALA-SAEM algorithm has already been applied to the BME template model in the context of linearized deformations (see [3]). In that paper, the almost sure convergence of the parameter sequence as well as its asymptotic normality (Theorems 1 and 2 in [3]) have been proven under the usual assumptions on both the model and the step size sequences. Thus, we can wonder whether our LDDMM BME template model fits into these assumptions. First of all, we note that our model belongs to the curved exponential family. Moreover, it satisfies the regularity and integrability conditions required in assumptions (M1)–(M6) and (B2) of Theorem 1 in [3]. However, due to the very complex dependencies of the LDDMM model, the superexponential property (B1) of the conditional density and, related to it, its polynomial upper bound (M8) cannot be guaranteed. Nevertheless, both assumptions sound reasonable in the applications that we are targeting. In the following experiments, the convergence of the algorithm is demonstrated, thus corroborating our hypothesis.

5. Extension toward sparse representation of the geometric variability. Obviously, the number of degrees of freedom needed to describe the variability of a given shape should be adapted to this shape. Therefore, the number of control points in our model should be estimated as a parameter of the model and not fixed by the user. This leads to optimizing automatically the dimension of the deformation model. We propose here to optimize the positions of the control points and to select a subset of the most relevant ones for the description of the variability simultaneously.

In [20], the control point selection is done adding an L^1 penalty on the momenta to the energy E_θ and performing an adapted gradient descent called FISTA (see [11]). The effect of this penalty is to zero out momenta of small magnitude and to slightly decrease the magnitude

ALGORITHM 1. AMALA-SAEM WITH TRUNCATION ON RANDOM BOUNDARIES.

Set $\kappa_0 = 0$, $s_0 \in \mathcal{K}_0$ and $\mathbf{z}_0 \in \mathrm{K}$.
for all $k \geq 1$ **do**
  Sample $\bar{\mathbf{z}}$ from the AMALA transition kernel :
     Sample $\mathbf{z}_c \sim \mathcal{N}(\mathbf{z}_k + \delta D_{\theta_k}(\mathbf{z}_k), \delta \Sigma_{\theta_k}(\mathbf{z}_k))$ where

$$D_\theta(\mathbf{z}) = \frac{b}{\max(b, |\nabla \log \pi_\theta(\mathbf{z})|)} \nabla \log \pi_\theta(\mathbf{z}),$$

  and

$$\Sigma_\theta(\mathbf{z}) = \varepsilon \mathrm{Id}_{(n+1)dk_g} + diag\left(D_\theta^0(\mathbf{z})D_\theta^0(\mathbf{z})^t, D_\theta^1(\mathbf{z})D_\theta^1(\mathbf{z})^t, \ldots, D_\theta^n(\mathbf{z})D_\theta^n(\mathbf{z})^t\right),$$

  with $\varepsilon > 0$ a small regularization parameter and $\delta > 0$.
     Then set $\bar{\mathbf{z}} = \mathbf{z}_c$ with probability $\rho_{\theta_k}(\mathbf{z}_k, \mathbf{z}_c)$ and $\bar{\mathbf{z}} = \mathbf{z}_k$ otherwise, where ρ_θ
  is given in (4.3).

  Compute $\bar{s} = s_{k-1} + \Delta_{k-1}(S(\bar{\mathbf{z}}) - s_{k-1})$ where $(\Delta_k)_{k \in \mathbb{N}}$ is a decreasing positive step
  size sequence.
  if $\bar{s} \in \mathcal{K}_{\kappa_{k-1}}$ and $|\bar{s} - s_{k-1}| \leq \varepsilon_{k-1}$ **then**
    set $(s_k, \mathbf{z}_k) = (\bar{s}, \bar{\mathbf{z}})$ and $\kappa_k = \kappa_{k-1}$,
  else
    set $(s_k, \mathbf{z}_k) = (\tilde{s}, \tilde{\mathbf{z}}) \in \mathcal{K}_0 \times \mathrm{K}$ and $\kappa_k = \kappa_{k-1} + 1$,
    where $(\tilde{s}, \tilde{\mathbf{z}}) \in \mathcal{K}_0 \times \mathrm{K}$.
  end if
  $\theta_k = \hat{\theta}(s_k)$
end for

of the other ones. A control point which does not contribute to *at least* one of the template-to-observation deformations at the convergence of the algorithm is called inactive. Note that since control points move in the domain, inactive control points may become active during the optimization process, and vice versa.

This method suffers from three main limitations. First, the Laplace prior associated with the L^1 penalty does not generate sparse observations. Second, the method keeps active control points that may contribute to only a few template-to-observation deformations. Last, L^1 penalty implies a soft thresholding step on the momentum vectors, thus reducing the norm of these vectors keeping the direction and therefore the local curvature. As a consequence, important momenta for the description of the variability will also be penalized. In the following, we propose to select control points given their importance to describe the variability of the *whole* population, and not of outliers. The idea is to inactivate a control point if the distribution of the momenta attached to it is not strongly correlated with the momentum distribution of other control points. Therefore our procedure selects control point positions and their number, relevant with regards to the whole population.

This constraint on the momenta is taken into account in the model by modifying the update of the geometrical parameter in (8.5) in the maximization step. We assume that the geometric covariance matrix Γ_g is no longer the barycenter between the stochastic approximation of the

empirical covariance matrix and the prior but is of the form $\Gamma_g = A_g + \varepsilon_g \text{Id}$, where ε_g is a small positive real number and A_g is a sparse symmetric positive matrix. To construct A_g, we modify the empirical covariance term in the third equation of (8.5). Let c_k^g be one of the control points. We compute the sum of the Frobenius norms of the submatrices of the stochastic approximation s_3 of the sufficient statistic S_3 (see (8.3)) which is the empirical covariance of this control point with all others:

$$(5.1) \qquad t_k = \sum_{j=1}^{k_g} \| s_3(c_k^g, c_j^g) \|_F \,,$$

where $s_3(c_k^g, c_j^g)$ is the $d \times d$ submatrix of s_3 corresponding to the control points c_k^g and c_j^g. Let us fix a positive threshold λ. The control point c_k^g is said to be active if

$$(5.2) \qquad t_k \geq \lambda \,.$$

Let us denote $\mathcal{A}$ the set of all active points. Then, we define the sparse matrix A_g as follows:

$$(5.3) \qquad \forall (k,j) \in \{1, \dots k_g\}, \quad A_g(c_k^g, c_j^g) = s_3(c_k^g, c_j^g) \mathbb{1}_{c_k \in \mathcal{A}} \mathbb{1}_{c_j \in \mathcal{A}} \,.$$

By analogy with (8.5), the matrix Γ_g is updated as follows:

$$(5.4) \qquad \Gamma_g = \frac{1}{n + a_g} (A_g + a_g \text{Id}_{dk_g}) \,,$$

which also corresponds to introduce a specific prior on Γ_g.

This update is performed at each iteration of the estimation algorithm in the M-step. The threshold λ in our approach plays an equivalent role as the weight of the L^1 penalty in the criterion optimized in [20]. The larger λ, the sparser the solution.

In order to be self-adapted to the data, it could be a benefit to fix the threshold λ as a ratio of the maximum correlations between control points instead of setting a fixed value as in (5.2). Thus, a control point c_k^g is now active if

$$(5.5) \qquad t_k \geq \lambda \max_{1 \leq j \leq k_g} t_j \,.$$

Moreover, starting to threshold before the Markov chain reaches its stationarity can lead to poor covering of the target distribution support. Therefore, in practice, we start the threshold process after the burn-in period of the estimation algorithm.

To go one step further, we propose to automatically select an optimal threshold λ. We consider a criterion based on two relevant quantities, namely, the data attachment residual over the training images and the number of active control points. Indeed, the larger the threshold, the larger the residual and the lower the number of active control points. These quantities are computed for different values of the threshold. These sequences are then normalized to 1. The optimal threshold is chosen to be the point where the two normalized sequences intersect.

Figure 1. *Training set of the USPS database (20 images per digit—inverse video).*


Figure 2. *Estimated templates with varying numbers of control points:* 4 *(left),* 9 *(middle), and* 16 *(right) with either fixed (top) or estimated (bottom) control points positions.*

6. Experiments.

6.1. Handwritten digit experiments. Our first experiments are run on the United States Postal Service (USPS) handwritten digit database which is a traditional benchmark for quantitative performance evaluation of template estimation. Twenty images of each digit are used as the training sample which is presented in Figure 1. This sample shows a large geometric and photometric variability. We consider the model with random control points presented in (3.4) as well as its simplified version where the control points are fixed. The number of control points is chosen equal to 4, 9, or 16 depending on the experiments. We infer the atlas of each digit independently using our stochastic estimation algorithm for the two models. Note that the values of the hyperparameters, namely, the parameters of the interpolation kernels and prior distributions on the parameters are chosen as in [2]. In addition $a_g = 0.5$ in the following experiments.

We present the estimated templates obtained with both models and varying numbers of control points in Figure 2. The first row shows the template images estimated with fixed control points. The second one provides the estimated templates together with the estimated control point positions.

Figure 3. *Synthetic samples from the generative model with estimated control point positions: 4 (top) and 9 (bottom) control points for digit 8.*


Figure 4. *Synthetic samples from the generative model with either fixed (top) or estimated (bottom) control point positions for digit 6.*

As expected, the contours in the template image become sharper in both cases as the number of control points increases. Moreover, the number of control points being fixed, the sharpness of the estimated template is improved by allowing the control points to move toward optimized positions. We can also note that the estimated control points are informative as they tend to move toward the contours of the digits, and, in particular, toward those that correspond to the regions of highest variability among samples. It is particularly noticeable on digits 5 and 6, for example.

Note that we checked empirically the identifiability of the control point positions by running several times the same experiment with different random initializations.

In Figure 3, we compare the geometry captured with 4 (top) and 9 (bottom) estimated control points for digit 8. As shown in Figure 2, the contours of the template image with only 4 control points is less sharp particularly in its upper part than the one with 9 control points. For the 4 point model, there are only two close control points in the lower part of the shape, whereas there are three of them spread around the loop with 9 control points. These additional degrees of freedom make the deformation model more flexible as highlighted in Figure 3. Not only the template looks better with an increasing number of control points but the captured geometric variability is also improved.

We evaluate the relevance of the estimated covariance matrix via the generation of synthetic samples. In Figure 4, we compare the geometry captured with 9 control points using fixed (top) and estimated (bottom) control point models. Although the template of the digit 6 looks similar in both cases, this experiment shows that the geometric variabilities captured by the models are rather different. The model with equidistributed fixed control points generates an unrealistic shape of the digit 6 and therefore does not reflect well the geometric variability observed in the training set. Optimizing for control point positions enables us to retrieve a much more natural geometric variability. This optimization increases the number

Figure 5. *Evolution of the classification score for varying numbers of control points either fixed (blue) or estimated (red).*

of hidden variables to sample, although the dimension of the covariance matrix remains the same, namely, $dk_g \times dk_g$. Updates in control point positions optimize the subgroup of diffeomorphisms of fixed dimension that is the most adapted to describe the variability of a given data set.

Nonetheless, one can notice that beyond a certain number of control points, the improvement is less obvious (see Figure 2). This suggests that there may be an *intrinsic dimension* of the deformation space that is optimal (neither too small nor too redundant) for the description of the variability of a given data set. This is also highlighted by the following classification experiment. To perform the classification, we use the test set available in the USPS database. It contains 2007 digit images. The allocated class for a test image is calculated as follows: we approximate the posterior distribution $q(c|y; \theta_c)$ of the class c given this image y using

$$(6.1) \qquad q(c|y; \theta_c) \simeq Cq(y|(\boldsymbol{\alpha}_0^*)_c, (\bar{\boldsymbol{c}}_0)_c, \theta_c)q((\boldsymbol{\alpha}_0^*)_c|\theta_c)q((\bar{\boldsymbol{c}}_0)_c|\theta_c)q(\theta_c) \,,$$

where $(\bar{\boldsymbol{c}}_0)_c$ is the estimated vector of control points for the class c and

$$(\boldsymbol{\alpha}_0^*)_c = \underset{\boldsymbol{\alpha}_0}{\operatorname{argmax}}\, q(\boldsymbol{\alpha}_0|y, (\bar{\boldsymbol{c}}_0)_c, \theta_c) \,.$$

This approximation has already been used in [2].

Classification results are presented in Figure 5 for 4, 9, 16, 25, 36, 49, and 64 control points using both fixed (blue) and random (red) control point models. The scale of the Gaussian interpolation kernel K_g is fixed such that considering 36 control points leads to one point every kernel scale.

With no control point, the model classifies according to the L^2 similarity with the gray level average image. This mean image, though very fuzzy, is still informative and leads to a classification score of about 85%. If the number of control points is increased, the model incorporates deformations. The template images become less fuzzy (see Figure 2), deformations explain part of the shape variability in an interpretable way (see Figure 4),

Figure 6. *Five training images from the mouse mandibles.*

Figure 7. *Estimated templates of the mouse mandible images obtained with* 260 *fixed control points (left), with* 117 *(middle), and* 70 *(right) estimated control points.*

and the classification scores increase (see Figure 5). Near the maximal classification score, models with estimated control points perform better. As already noted, the slight increase in classification score goes with a much more realistic and interpretable representation of the variability (see Figures 3 and 4). If the number of control points is drastically increased, overall classification scores drop down, as we fall typically into an overfitting situation. Allowing control point positions to be optimized further increases the dimension of parameters. In this case, the deformation model becomes so flexible that it can accommodate for any small differences in shapes, and does not generalize well.

The best performances are reached for in-between numbers of control points. In this region, estimating the positions of these control points allows us to reach higher classification scores. This confirms the idea of the existence of an intrinsic dimension of the deformation space. How to find such dimension is the purpose of the sparse extended model presented in section 5 and experimented on in section 6.3.

6.2. Mouse mandible experiment. We consider a second training set composed of 36 X-ray scans of mouse mandibles. Five of them are presented in Figure 6. The estimated template images resulting from three different experiments are shown in Figure 7. The image on the left shows the template estimated using 260 fixed equidistributed control points. The image on the middle (resp., right) shows the estimated template using 117 (resp., 70) estimated control points. These templates look similar, thus showing that the same photometric invariants have been captured in each experiment. These invariants include the main bones of the mandibles (i.e., the brightest areas in the image). The decrease in number of control points is balanced by the optimization of their optimal positions. Control points in the right image are noticeably located on the edges of the shape in order to drive the dilation, contraction, and opening of

Figure 8. *Synthetic training sample:* 10 *examplars among the* 20 *regular base and the outlier on the right end.*

Figure 9. *Estimated templates. Left: template estimated with the deterministic gradient descent of* [20] *for the data set including the outlier. The colored arrows represent the initial momenta for different subjects (in different colors) which are given as output of the algorithm. Middle and right: templates estimated with the stochastic algorithm for the* 20 *regular images (middle) and the* 20 *regular images plus the outlier (right).*

the mandible. Depending on the desired precision of the atlas, we can reduce even more the number of control points. This enables a faster estimation task at the cost of providing less information about the data.

6.3. Toward sparse representation.

6.3.1. Efficiency of control point selection.
We test our estimation procedure for the sparse constrained model on a toy example in order to exhibit the stability of the estimated geometry with respect to outliers. We compare these results to the atlas estimation given by the procedure in [20] where the selection of control points is based on individual contributions. Note that the same deformation model is used in both approaches. We create a data set of 20 images which are composed of vertical translations and vertical dilations of a given rectangle. An outlier image is then introduced into this database. This outlier has an excrescence on its left border (see Figure 8). We run first our algorithm with the 20 regular training images. Then we run our algorithm and the one proposed in [20] with these 20 images together with the outlier. The three estimated templates with their respective optimized control point positions are shown in Figure 9. The gray level differences are negligible. Runing our estimation procedure for the training images with or without the outlier leads to very similar estimates of the control point positions, thus showing the robustness of our estimate with respect to outliers. Samples generated from both estimated models are presented in Figure 10. They show only vertical deformations up to the isotropy of the Gaussian interpolation kernel. This confirms the ability of the threshold process to limit the effect of the outlier in the data set. By contrast, the method proposed in [20], run on the same data set including the outlier, exhibits a very different result. On the left border of the shape, two control points, with momenta magnitude of the same order as the other ones, play an important role in this model although only explaining the variability of the outlier.

Figure 10. *Estimated geometry: samples generated from the 20 regular image estimated model (top) and from the 20 regular images plus the outlier estimated model (bottom).*

Figure 11. *Evolution of the estimated templates and of their number of active control points with respect to the threshold parameter. From left to right: λ equals $0.3, 0.45, 0.6, 0.75,$ and 0.8.*

6.3.2. Results on the USPS database.

We then run the estimation algorithm presented above with the extension described in section 5 and the threshold rule (5.5) on the USPS database presented above. We conduct different experiments with different thresholds λ between 0.3 and 0.8 in order to see the evolution of the sparsity with respect to this parameter and also to capture the most interesting one (depending on the training digit). The initial number of control points is set to 16. The results of these experiments are presented in Figures 11 and 12.

As expected, increasing the threshold λ decreases the final number of selected control points, whose effects on template sharpness and description of variability have been presented in Figures 2, 3, and 4. Using the modified prior given in (5.4) to enforce sparsity allows us to automatically select a subset of control points leading to estimation results of the same

Figure 12. *Synthetic samples of digit* 2 *from the generative model using the estimated parameters for thresholds* 0.3 *(top),* 0.6 *(middle), and* 0.8 *(bottom).*

accuracy (see Figures 11 and 12). Contrary to the L^1 prior used in [20], our sparsity prior selects a small number of control points without penalizing the magnitude of the momenta. Hence the variability of the model is not underestimated. In this respect, our thresholding process has an effect which is closer to the expected L^0 norm than its surrogate L^1 norm.

Furthermore, considering the optimal threshold, we compare the estimated geometry of our model with the posterior geometry deduced from the estimated momenta of [20]. Indeed, whereas we estimate the matrix Γ_g along with the template, it is a fixed parameter in [20]. In order to evaluate the accuracy of the momenta, the authors used the posterior empirical matrix to generate a synthetic sample (see [20, Fig. 9]). This compares with our Figure 12. Learning Γ_g enables us to better catch the geometric variability of the sample. Moreover it also allows us to reduce the optimal number of control points to better fit the data.

Independently of the threshold λ, control points move in areas where the shape is the most variable. This can be noticed in the loop of the digit 2 which is highly variable, especially in contrast to the loop of the digit 6 which is much more stable in shape across observations. This can be seen by a fastest decrease in number of control points when the threshold λ is increased for the digit 6 compared to digit 2. It is also interesting to notice how our model deals with a mixture of 2 that could be written with or without a loop. Such variability violates the hypothesis of our model, which assumes that observations derive from *diffeomorphic* deformations of the template image. In this situation, the model estimates a template image that is fuzzy in the region of the loop: the nondiffeomorphic variability has been interpreted as a photometric variation. To overcome this problem, one may investigate the use of several template images in the atlas along the lines of [4].

The optimal threshold is chosen applying the criterion described in section 5. Figure 13 shows the estimated templates with their control points corresponding to the optimal threshold. The number of control points reflects the variability of the digits. In particular, very constrained shapes (see digits 1 and 9) require fewer control points than very complex irregular forms (see digits 3 and 8). Note that in most of the cases (nine digits out of ten) the selection criterion enables us to select thresholds between 50 and 60%.

6.4. 3D experiments on murine dendrite spines. We run our stochastic algorithm on a set of murine dendrite spine images (see [1, 14, 15]). The data set consists of 30 binary images of microscopic structures, tiny protuberances found on many types of neurons termed dendrite

Figure 13. *Estimated templates with their optimal numbers and positions of control points.*


Figure 14. $3D$ *views of* 8 *examples of the* 30 *elements of the dendrite spine database.*

spines. The images are from control mice and knockout mice which have been genetically modified to mimic human neurological pathologies like Parkinson's disease. The acquisition process consisted of electron microscopy after injection of Lucifer yellow and subsequent photo-oxidation. The shapes were then manually segmented on the tomographic reconstruction of the neurons. A representative set of these images are presented in Figure 14. Each image is a binary (background $= 0$, object $= 2$) cubic volume of size 28^3. We can notice here the large geometrical variability of this population of images.

The estimated template is presented in Figure 15. This average shape is similar to the one presented in [5] and [3] and captures well the geometrical invariance present in the training set. In order to visualize the estimated geometry, we have drawn synthetic samples thanks to the generative model with the estimated parameters. Eight of these samples are shown in Figure 16. As expected, the geometry which is captured by the diffeomorphic deformation

Figure 15. *Estimated template using the proposed algorithm with* 30 *training examples of the database.*

Figure 16. 3D *synthetic samples sampled with the generative model using the estimated parameters.*

model is more constrained than the one with the linearized deformation model (see Figure 4 in [5]). Although this generates smoother shapes, it looks realistic with regard to the training samples. The selection of control points has been performed in the same way as for the USPS database. This reduces the number of geometrical control points to 8 out of 15^3.

Another great advantage of this estimation process is its computational performance. Compared to the stochastic estimation algorithm using the Gibbs sampler in the linearized deformation model, it runs about 15 times faster. The finite parametrization of the LDDMM deformations enables to reduce the computation time although it now involves costly gradient and shooting steps. This allows us to compute an estimated atlas in 36 h for this data set.

7. Conclusion and perspectives. In this paper, we presented a generalization of the BME template model, in which a diffeomorphic constraint has been added on the deformations. Moreover, a finite-dimensional parametrization of these deformations via control points has been used, which enabled us to include the positions of these points as parameters of the model. The AMALA-SAEM algorithm appears to be particularly well suited to estimate the parameters of such models, especially in comparison with the Gibbs-SAEM. This opened up the possibility to root the usual atlas estimation using large diffeomorphic deformations into a rigorous statistical framework, and to propose a tractable stochastic algorithm to estimate its parameters. The results on handwritten digits, mouse mandibles, and murine dendrite spine images show the interest of this model. Moreover, the issue of the optimal number of control points has been addressed including a selection step to only keep the most informative points. We proposed an empirical criterion to optimize the threshold leading to a model selection. This choice is usually done by cross validation, which requires a large training database and is computationally costly.

A natural extension is to consider the mixture model introduced in [4] using the same LDDMM formulation with estimated control points. This model can easily be proposed, however, the difficulty stands in adapting the stochastic algorithm with the AMALA sampler so that it remains tractable.

Another remark concerns the choice of a Gaussian distribution for the momenta. In the linearized deformation model, this looks reasonable; indeed, in a flat manifold, the mean of the deformation (and thus the equivalent of the momentum α) starting from the ideal template to all the data should be close to 0. The global behaviors can be well approximated by some Gaussian behavior. But as soon as you consider large deformations, you are no longer in a flat manifold and the curvature has to be taken into account. In this manifold, matching one point (the template) to two close points (two different targets) will not necessarily imply that the two corresponding momenta are close to each other. This leads us to think that the Gaussian model should be changed to some other which would take the curvature of points (images) into account.

Finally, looking at the analytical expression of the observed log-likelihood, we recognize the terms coming from the Gaussian distributions on the observations and on the initial momenta as the two terms of a LDDMM registration energy (resp., data attachment and L^2 penalty terms). For this reason, it would be coherent with this setting to use the metric Γ_g both in the penalty term and the definition of the velocity field using interpolation matrix K_g. One further interest of this generalization will be to include a correlation between these two matrices, but this is not straightforward.

8. Appendix.

8.1. Proof of the existence of the MAP estimator. From (3.4) we have that for any $\theta = (\mathrm{w}, \sigma^2, \Gamma_g, \bar{\mathbf{c}}_\mathbf{0}) \in \Theta$

$$q(\mathbf{y}|\boldsymbol{\alpha}_0, \mathbf{c}_0, \mathrm{w}, \sigma^2) q(\boldsymbol{\alpha}_0|\Gamma_g) q(\mathbf{c}_0|\bar{\mathbf{c}}_\mathbf{0})$$

(8.1)
$$\leq (2\pi\sigma^2)^{-|\Lambda|/2} (2\pi)^{-k_g} |\det(\Gamma_g)|^{-1/2} \exp\left(-\frac{1}{2}\boldsymbol{\alpha}_0^t \Gamma_g^{-1} \boldsymbol{\alpha}_0\right)$$

$$\times (2\pi a_c)^{-dk_g/2} \exp\left(-\frac{1}{2a_c}\|\mathbf{c}_0 - \bar{\mathbf{c}}_\mathbf{0}\|^2\right)$$

so that integrating over $\boldsymbol{\alpha}_0$ and $\mathbf{c}_0$ and adding the priors on each parameter, we get

$$\log(q(\theta|\mathbf{y})) \leq -\frac{a_g}{2}\langle R_g, \Sigma_g\rangle_F + \frac{a_g}{2}\log(|\det(R_g)|) - \frac{a_p\sigma_0^2}{2\sigma^2} - \frac{n|\Lambda| + a_p}{2}\log(\sigma^2)$$
$$- \frac{1}{2}\mathrm{w}^t \Sigma_p^{-1}\mathrm{w} - \frac{1}{2}(\bar{\mathbf{c}}_\mathbf{0} - \mu_c)^t \Sigma_c^{-1}(\bar{\mathbf{c}}_\mathbf{0} - \mu_c) + \mathrm{C},$$

where $R_g = \Gamma_g^{-1}$, and C does not depend on the parameters. If we denote by η_g^0 the smallest eigenvalue of Σ_g and by $\|\cdot\|$ the operator norm, we get

$$\langle R_g, \Sigma_g\rangle_F \geq \eta_g^0 \|R_g\| \quad \text{and} \quad \log(|\det(R_g)|) \leq (2k_g - 1)\log(|\det(R_g)|) - \log(\|\Gamma_g\|)$$

so that

$$\lim_{\|R_g\| + \|\Gamma_g\| \to \infty} -\frac{a_g}{2}\langle R_g, \Sigma_g\rangle_F + \frac{a_g}{2}\log(|\det(R_g)|) = -\infty.$$

Similarly, we can show

$$\lim_{\sigma^2 + \sigma^{-2} \to \infty} -\frac{a_p \sigma_0^2}{2\sigma^2} - \frac{n|\Lambda| + a_p}{2} \log(\sigma^2) = -\infty \,,$$

$$\lim_{\|\mathbf{w}\| \to \infty} -\frac{1}{2} \mathbf{w}^t \Sigma_p^{-1} \mathbf{w} = -\infty \,,$$

and

$$\lim_{\|\bar{\mathbf{c}}_0\| \to \infty} -\frac{1}{2} (\bar{\mathbf{c}}_0 - \mu_c)^t \Gamma_c^{-1} (\bar{\mathbf{c}}_0 - \mu_c) = -\infty \,.$$

Now considering the Alexandrov one point compactification $\Theta \cup \{\infty\}$ of Θ, we have

$$\lim_{\theta \to \infty} \log q(\theta | \mathbf{y}) \to -\infty \,.$$

Since $\theta \to \log q(\theta | \mathbf{y})$ is smooth on Θ, we get the result.

8.2. Description of the stochastic approximation, projection, and maximization steps of the algorithm. We detail here these three steps of the algorithm. The generalized large deformation BME template model belongs to the curved exponential family. Indeed, the log-likelihood is written

$$\log q(\mathbf{y}, \mathbf{c}_0, \boldsymbol{\alpha}_0, \theta)$$

$$= \sum_{i=1}^{n} \left(-\frac{|\Lambda|}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \|y_i - \mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \mathbf{w}\|^2 \right)$$

(8.2)
$$+ \sum_{i=1}^{n} \left(-\frac{dk_g}{2} \log(2\pi) - \frac{1}{2} \log(|\det(\Gamma_g)|) - \frac{1}{2} (\boldsymbol{\alpha}_0^i)^t \Gamma_g^{-1} \boldsymbol{\alpha}_0^i \right) - \frac{1}{2a_c} \|\mathbf{c}_0 - \bar{\mathbf{c}}_0\|^2$$

$$+ a_g \left(-\frac{1}{2} \langle \Gamma_g^{-1}, \Sigma_g \rangle_F - \frac{1}{2} \log(|\det(\Gamma_g)|) \right) + a_p \left(-\frac{\sigma_0^2}{2\sigma^2} - \frac{1}{2} \log(\sigma^2) \right)$$

$$- \frac{1}{2} \mathbf{w}^t \Sigma_p^{-1} \mathbf{w} - \frac{1}{2} (\bar{\mathbf{c}}_0 - \mu_c)^t \Sigma_c^{-1} (\bar{\mathbf{c}}_0 - \mu_c) + \text{C} \,,$$

where for all $u \in \Lambda$, $\mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \mathbf{w}(r_u) = \sum_{k=1}^{k_p} K_p((\phi_1^{\mathbf{c}_0, \boldsymbol{\alpha}_0^i})^{-1}(r_u), b_k) \mathbf{w}_k$.

This enables us to exhibit the following sufficient statistics:

(8.3)
$$\begin{cases} S_0(\mathbf{z}) = \mathbf{c}_0 \,, \\ S_1(\mathbf{z}) = \sum_{1 \le i \le n} \left(\mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \right)^t y_i \,, \\ S_2(\mathbf{z}) = \sum_{1 \le i \le n} \left(\mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \right)^t \left(\mathbf{K}_p^{(\mathbf{c}_0, \boldsymbol{\alpha}_0^i)} \right) \,, \\ S_3(\mathbf{z}) = \sum_{1 \le i \le n} (\boldsymbol{\alpha}_0^i)^t \boldsymbol{\alpha}_0^i \,. \end{cases}$$

For simplicity, we denote $S(\mathbf{z}) = (S_0(\mathbf{z}), S_1(\mathbf{z}), S_2(\mathbf{z}), S_3(\mathbf{z}))$ for any $\mathbf{z} = (\mathbf{c}_0, \boldsymbol{\alpha}_0) \in \mathbb{R}^{(n+1)dk_g}$. We define the corresponding sufficient statistic space

$$\mathcal{S} = \left\{ (s_0, s_1, s_2, s_3) | s_0 \in \mathbb{R}^{dk_g}, \ s_1 \in \mathbb{R}^{k_p}, \ s_2 + \sigma_0^2 \Sigma_p^{-1} \in \Sigma_{k_p,*}^{+}(\mathbb{R}), \ s_3 + a_g \Sigma_g \in \Sigma_{2k_g,*}^{+}(\mathbb{R}) \right\} \,.$$

Identifying in what follows s_2 and s_3 with their lower triangular part, the set S can be viewed as an open convex set of $\mathbb{R}^{n_s}$ with $n_s = dk_g + k_p + \frac{k_p(k_p+1)}{2} + k_g(2k_g + 1)$.

We carry out a stochastic approximation on the sufficient statistics of the model as follows:

$$\bar{s} = s_{k-1} + \Delta_{k-1}(S(\bar{\mathbf{z}}) - s_{k-1}),$$

where $(\Delta_k)_{k\in\mathbb{N}}$ is a decreasing positive step size sequence and $\bar{\mathbf{z}}$ is the output of the simulation step.

Since we are not dealing with hidden variables with compact support, we introduce a usual projection of the sufficient statistics on random boundaries. Let $(\mathcal{K}_q)_{q\geq 0}$ be an increasing sequence of compact subsets of $\mathcal{S}$ such as $\cup_{q\geq 0}\mathcal{K}_q = \mathcal{S}$ and $\mathcal{K}_q \subset \text{int}(\mathcal{K}_{q+1})$ for all $q \geq 0$. Let $(\varepsilon_k)_{k\geq 0}$ be a monotone nonincreasing sequence of positive numbers and $\mathbb{R}^{n_z}$, where $n_z = (n+1)dk_g$ a subset of $\mathbb{R}^{n_s}$. We construct the sequence $((s_k, \mathbf{z}_k))_{k\geq 0}$ as follows. As long as the stochastic approximation does not fall out of the current compact set and is not too far from its previous value, we run the AMALA-SAEM algorithm. If one of the two previous conditions is no longer satisfied, we reinitialize the sequences of s and $\mathbf{z}$ using a projection $(\tilde{s}, \tilde{\mathbf{z}}) \in \mathcal{K}_0 \times \mathrm{K}$ (for more details, see [8]).

Moreover, there exists a function $\hat{\theta}$ satisfying

(8.4) $$\forall \theta \in \Theta, \forall s \in \mathcal{S}, L(s; \hat{\theta}(s)) \geq L(s; \theta),$$

which yields a very simple maximization step given the sufficient statistics s of the model:

(8.5) $$\begin{cases} \hat{\mathbf{w}}(s) = \left(s_2 + \hat{\sigma}(s)^2(\Sigma_p)^{-1}\right)^{-1} s_1, \\[2mm] \hat{\sigma}^2(s) = \frac{1}{n|\Lambda|+a_p}\left(\|\mathbf{y}\|^2 + \hat{\mathbf{w}}(s)^t s_2 \hat{\mathbf{w}}(s) - 2\hat{\mathbf{w}}(s)^t s_1 + a_p \sigma_0^2\right), \\[2mm] \hat{\Gamma}_g(s) = \frac{1}{n+a_g}(s_3 + a_g \Sigma_g), \\[2mm] \hat{\bar{\mathbf{c}}}_0(s) = s_0. \end{cases}$$

REFERENCES

[1] G.M. ALDRIDGE, J.T. RATNANATHER, M.E. MARTONE, M. TERADA, M.F. BEG, L. FONG, E. CEY-HAN, A.E. KOLASNY, T.J.A. BROWN, E.L. COCHRAN, S.J. TANG, D.V. PISANO, M. VAILLANT, M.K. HURDAL, J.D CHURCHILL, W.T. GREENOUGH, M. I. MILLER, AND M.H. ELLISMAN, *Semiautomated Shape Analysis of Dendrite Spines from Animal Models of Fragilex and Parkinson's Disease using Large Deformation Diffeomorphic Metric Mapping*, in Society for Neuroscience Annual Meeting, Washington DC, 2005.

[2] S. ALLASSONNIÈRE, Y. AMIT, AND A. TROUVÉ, *Towards a coherent statistical framework for dense deformable template estimation*, J. R. Stat. Soc. Ser. B Stat. Methodol., 69 (2007), pp. 3–29.

[3] S. ALLASSONNIÈRE AND E. KUHN, *Convergent stochastic expectation maximization algorithm with efficient sampling in high dimension. Application to deformable template model estimation*, Comput. Statist. Data Anal., 91 (2015), pp. 4–19.

[4] S. ALLASSONNIÈRE AND E. KUHN, *Stochastic algorithm for Bayesian mixture effect template estimation*, ESAIM Probab. Statist., 14 (2010), pp. 382–408.

[5] S. ALLASSONNIÈRE, E. KUHN, AND A. TROUVÉ, *Bayesian consistent estimation in deformable models using stochastic algorithms: Applications to medical images*, J. Soc. Fran. Statist., 151 (2010), pp. 1–16.

[6] S. ALLASSONNIÈRE, E. KUHN, AND A. TROUVÉ, *Construction of Bayesian deformable models via a stochastic approximation algorithm: A convergence study*, Bernoulli, 16 (2010), pp. 641–678.

[7] S. ALLASSONNIRE, A. TROUVÉ, AND L. YOUNES, *Geodesic shooting and diffeomorphic matching via textured meshes*, in Energy Minimization Methods for Computer Vision and Pattern Recognition (EMMCVPR), A. Yuille A. Rangarajan, B. Vemuri, ed., Springer-Verlag, Berlin, 2005, pp. 365–381.

[8] C. ANDRIEU, É. MOULINES, AND P. PRIOURET, *Stability of stochastic approximation under verifiable conditions*, SIAM J. Control Optim., 44 (2005), pp. 283–312.

[9] N. ARONSZAJN, *Theory of reproducing kernels*, Trans. Amer. Math. Soc., 68 (1950), pp. 337–404.

[10] F.R. BACH, *Consistency of the group lasso and multiple kernel learning*, J. Mach. Learn. Res., 9 (2008), pp. 1179–1225.

[11] A. BECK AND M. TEBOULLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.

[12] M.F. BEG, M.I. MILLER, A. TROUVÉ, AND L. YOUNES, *Computing large deformation metric mappings via geodesic flows of diffeomorphisms*, Int J. Comp. Vision, 61 (2005), pp. 139–157.

[13] J. BIGOT, S. GADAT, AND J.-M. LOUBES, *Statistical M-estimation and consistency in large deformable models for image warping*, J. Math. Imaging Vision, 34 (2009), pp. 270–290.

[14] E. CEYHAN, L. FONG, T.N. TASKY, M.K. HURDAL, M.F. BEG, M.E. MARTONE, AND J.T. RATNANATHER, *Type-specific analysis of morphometry of dendrite spines of mice*, in 5th International Symposium on Image and Signal Processing and Analysis, 2007, IEEE, Piscataway, NJ, pp. 7–12.

[15] E. CEYHAN, R.. LKEN, L. FONG, T.N. TASKY, M.K. HURDAL, M.F. BEG, M.E. MARTONE, AND J.T. RATNANATHER, *Modeling metric distances of dendrite spines of mice based on morphometric measures*, in International Symposium on Health Informatics and Bioinformatics, Antalya, Turkey, 2007.

[16] G.E. CHRISTENSEN, R.D. RABBITT, AND M. I. MILLER, *Deformable templates using large deformation kinematics*, IEEE Trans. Image Proc., 5 (1996), pp. 1435–1447.

[17] B. DELYON, M. LAVIELLE, AND É. MOULINES, *Convergence of a stochastic approximation version of the EM algorithm*, Ann. Statist., 27 (1999), pp. 94–128.

[18] P. DUPUIS, U. GRENANDER, AND M.I. MILLER, *Variational problems on flows of diffeomorphisms for image matching*, Quart. Appl. Math., 56 (1998), pp. 587–600.

[19] S. DURRLEMAN, *Statistical Models of Currents for Measuring the Variability of Anatomical Curves, Surfaces and their Evolution*, Ph.D. thesis, Université Nice-Sophia Antipolis, France, 2010.

[20] S. DURRLEMAN, S. ALLASSONNIÈRE, AND S. JOSHI, *Sparse adaptive parameterization of variability in image ensembles*, Int. J. Comput. Vision, 101 (2013), pp. 161–183.

[21] S. DURRLEMAN, M. PRASTAWA, G. GERIG, AND S. JOSHI, *Optimal data-driven sparse parameterization of diffeomorphisms for population analysis*, in Information Processing in Medical Imaging, G. Székely and H. Hahn, eds., Lecture Notes in Comput. Sci. 6801, Springer, Berlin, 2011, pp. 123–134.

[22] J. GLAUNÈS, M. VAILLANT, AND M.I. MILLER, *Landmark matching via large deformation diffeomorphisms on the sphere*, J. Math. Imaging Vision, 20 (2004), pp. 179–200.

[23] U. GRENANDER, *General Pattern Theory*, Oxford Science Publications, Oxford, 1993.

[24] U. GRENANDER AND M.I. MILLER, *Computational anatomy: An emerging discipline*, Quart. Appl. Math., 56 (1998), pp. 617–694.

[25] D.R. HOLM, J.T. RATNANATHER, A. TROUVÉ, AND L. YOUNES, *Soliton dynamics in computational anatomy*, Neuroimage, 23 (2004), pp. 170–178.

[26] S. JOSHI AND M.I. MILLER, *Landmark matching via large deformation diffeomorphisms*, IEEE Trans. Image Proc., 9 (2000), pp. 1357–1370.

[27] E. KUHN AND M. LAVIELLE, *Coupling a stochastic approximation version of EM with an MCMC procedure*, ESAIM Probab. Statist., 8 (2004), pp. 115–131.

[28] J. MA, M.I. MILLER, A. TROUVÉ, AND L. YOUNES, *Bayesian template estimation in computational anatomy*, Neuroimage, 42 (2008), pp. 252–261.

[29] S. MARSLAND AND C. TWINING, *Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images*, IEEE Trans. Med. Imaging, 23 (2004), pp. 1006–1020.

[30] M. MICHELI, P.W. MICHOR, AND D. MUMFORD, *Sectional curvature in terms of the cometric, with applications to the Riemannian manifolds of landmarks*, SIAM J. Imaging Sci., 5 (2012), pp. 394–433.

[31] M. Miller, C.E. Priebe, A. Qiu, B. Fischl, A. Kolasny, T. Brown, Y. Park, J.T. Ratnanather, E. Busa, J. Jovicich, P. Yu, B.C. Dickerson, and R.L. Buckner, *Collaborative computational anatomy: An MRI morphometry study of the human brain via diffeomorphic metric mapping*, Human Brain Mapping, 30 (2009), pp. 2132–2141.

[32] M.I. Miller, A. Trouvé, and L. Younes, *On the metrics and Euler-Lagrange equations of computational anatomy*, Ann. Rev. Biomed. Engrg., 4 (2002), pp. 375–405.

[33] M.I. Miller, A. Trouvé, and L. Younes, *Geodesic shooting for computational anatomy*, J. Math. Imaging Vision, 24 (2006), pp. 209–228.

[34] M.I. Miller and L. Younes, *Group action, diffeomorphism and matching: A general framework*, Int. J. Comput. Vision, 41 (2001), pp. 61–84.

[35] X. Pennec, *Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements*, J. Math. Imaging Vision, 25 (2006), pp. 127–154.

[36] L. Risser, F.-X. Vialard, R. Wolz, M. Murgasova, D.D. Holm, and D. Rueckert, *Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping*, IEEE Trans. Med. Imaging, 30 (2011), pp. 1746–1759.

[37] S. Sommer, F. Lauze, M. Nielsen, and X. Pennec, *Sparse multi-scale diffeomorphic registration: The kernel bundle framework*, J. Math. Imaging Vision, 46 (2013), pp. 292–308.

[38] J.-P. Thirion, *Image matching as a diffusion process: An analogy with Maxwell's demons*, Med. Image Anal., 2 (1998), pp. 243–260.

[39] A. Trouvé, *Diffeomorphism groups and pattern matching in image analysis*, Int. J. Comput. Vision, 28 (1998), pp. 213–221.

[40] M. Vaillant, M.I. Miller, A. Trouvé, and L. Younes, *Statistics on diffeomorphisms via tangent space representations*, Neuroimage, 23 (2004), pp. 161–169.

[41] F.-X. Vialard and L. Risser, *Spatially-varying metric learning for diffeomorphic image registration: A variational framework*, in Medical Image Computing and Computer-Assisted Intervention - MICCAI 2014, Boston, 2014, Part I, Cham, Switzerland, 2014, pp. 227–234.

[42] M. Zhang, N. Singh, and P.T. Fletcher, *Bayesian estimation of regularization and atlas building in diffeomorphic image registration*, Information Processing in Medical Imaging, Lecture Notes in Comput. Sci. 7917, Springer, Berlin, 2013, pp. 37–48.