

Wavelet-Based Multiscale Flow For Realistic Image Deformation in the Large Diffeomorphic Deformation Model Framework

Fleur Gaudfernau^{1*}, Eléonore Blondiaux², Stéphanie Allasonnière¹,
Erwan Le Pennec³

¹CRC, Université de Paris, INRIA EPI HeKa, INSERM UMR 1138, Sorbonne Université.

²Service de Radiologie, Hôpital Armand-Trousseau, APHP.

³CMAP, Ecole polytechnique, Institut Polytechnique de Paris.

*Corresponding author(s). E-mail(s): fleur.gaudfernau@etu.u-paris.fr;
Contributing authors: eleonore.blondiaux@aphp.fr;
stephanie.allasonniere@u-paris.fr; erwan.le-pennec@polytechnique.edu;

Abstract

Estimating accurate high-dimensional transformations remains very challenging, especially in a clinical setting. In this paper, we introduce a multiscale parameterization of deformations to enhance registration and atlas estimation in the Large Deformation Diffeomorphic Metric Mapping framework. Using the Haar wavelet transform, a multiscale representation of the initial velocity fields is computed to optimize transformations in a coarse-to-fine fashion. This additional layer of spatial regularization does not modify the underlying model of deformations. As such, it preserves the original kernel Hilbert space structure of the velocity fields, enabling the algorithm to perform efficient gradient descent. Numerical experiments on several datasets, including abnormal fetal brain images, show that compared to the original algorithm, the coarse-to-fine strategy reaches higher performance and yields template images that preserve important details while avoiding unrealistic features. This highly versatile strategy can easily be applied to other mathematical frameworks for almost no additional computational cost.

Keywords: Deformable template model, atlas estimation, diffeomorphic deformations, Haar Wavelet, coarse-to-fine algorithm

1 Introduction

Although the quantitative analysis of anatomical images is an old problem [1], to this day it is still a challenging one, especially given that datasets of clinical images are often small in number and large in dimension. Estimating the transformation that warps an object onto another provides an efficient way of quantifying shape differences, which is the cornerstone idea of Computational Anatomy [2].

The choice of the function describing the template-to-subject transformations is of prime importance [3]. To account for the intra- and inter-subject anatomical variability in clinical images, non-linear deformations are mandatory. The Large Deformation Diffeomorphic Metric Mapping (LDDMM) setting [4; 5; 6] is a mathematical framework in which objects are warped through diffeomorphic transformations of the ambient space, i.e. high-dimensional, smooth and invertible functions with smooth inverse

that preserve image topology. The group of possible transformations forms a Riemannian manifold of infinite dimension and parameterizes a flexible representation of deformations.

The main advantage of the LDDMM lies in the fact that it constrains diffeomorphisms to be geodesic flows, i.e. the shortest paths between the objects according to a metric regularizing transformations [4]. Therefore, one can compute distances between shapes and perform meaningful statistical analyses on populations of shapes, such as Principal Component Analysis [7], geodesic regression [8] and atlas estimation [9].

In this paper, we are interested in atlas estimation, a method to model the mean and variability over a collection of images that are instances of the same anatomical object. An estimate of the average shape is given in the form of a template image, which represents the invariants across the population, i.e. shared anatomical features, and the variability is given by deformations from the template space to each subject’s space, which express how these common features vary within the population [10].

Atlas estimation has many applications in the field of medical image analysis. The template image can be used as reference to describe average anatomical structures or serve as a tool to automatically segment new subjects. Within the LDDMM framework, atlases have been used to characterize pathological deviations from normality [11], isolate subgroups in a population [12], and in a spatio-temporal fashion to characterize pathological changes, such as hippocampal reduction in Alzheimer’s disease [12].

In the LDDMM framework, diffeomorphisms are constructed by integrating time-dependent velocity fields, which results in high computational complexity. Conveniently, the resulting flow of diffeomorphisms is fully determined by the initial state of the system, enabling us to only optimize the initial velocity fields [13]. However, finding optimal diffeomorphisms involves solving expensive Partial Differential Equations (the “geodesic equations”) [14] and operating over infinite-dimensional velocity fields. To further improve optimization efficiency, a number of papers introduced sparse finite-dimensional parameterizations of the initial vector fields [15; 16; 17; 18]. In this paper, we will work within the framework of Durrleman

et al. [15], where velocity fields are parameterized by the convolution of momentum vectors attached to control points and regularized by a Gaussian kernel defining a Reproducing Kernel Hilbert Space (RKHS).

The choice of this regularizer is critical as it restricts the range of transformations defined by the model [19]. Specifically, it constrains the deformations occurring on the images to a single scale. A large kernel width is likely to produce smooth but less accurate matches, while a fine kernel will generate more accurate yet unnatural deformations. As clinical images often present high variability at several scales, one might be tempted to increase the number of parameters in the model, i.e. to use many control points and a small kernel. However, fine kernels make large displacements more expensive than small ones, and such over-parameterization will likely trap the optimization procedure in a local minimum, achieving a reasonable numerical solution that is qualitatively bad. To overcome such problem, hierarchical algorithms have been widely used in the field of image registration [3; 20]: after solving the registration problem at coarse scales, the solution is transferred to increasingly fine scales to refine the transformation. These strategies avoid more efficiently trapping the algorithm in local minima related to unrealistic transformations.

In this paper, we propose to enhance the outcomes of the conventional RKHS-based LDDMM algorithm by employing a coarse-to-fine optimization procedure based on a Haar-like wavelet representation of the initial velocity fields. Our reparameterization can be seen as an additional layer of regularization which does not modify the underlying model of deformation, making our algorithm highly versatile and transferable to other mathematical frameworks. We apply the multiscale strategy to the atlas estimation algorithm of Durrleman et al. [15], which uses a finite parameterization of the velocity fields as a linear combination of RKHS basis elements. This enables us to rely on an efficient numerical scheme to compute the gradients all the while controlling the smoothness of the deformations with our multiscale scheme. We evaluate our algorithm on a registration example and different atlas estimation tasks and show that our strategy generates more natural template images as well as higher stability regarding the initialization.

This paper is organized as follows. We first explore works related to coarse-to-fine registration and atlas estimation in Section 2, then we recall the LDDMM setting in the case of the finite parameterization of the velocity fields in Section 3 and introduce our coarse-to-fine atlas estimation method in Section 4; then, we conduct experiments on toy data and fetal brain images in Section 5 and discuss our results in Section 6.

2 Related work

Even though multiscale image registration has been studied repeatedly in the literature, it has rarely been evaluated in the setting of population analysis. As registration is a special case of atlas estimation with a fixed template image, in the following we will review both registration and atlas estimation methods that have a multiscale property.

Rooted in the idea of analyzing images across different scales [21; 22], hierarchical approaches have played a pivotal role in tasks involving image motion. Traditional optical flow algorithms [23; 24; 25; 26] were among the first to perform image matching in a coarse-to-fine manner, using pyramids of images of progressively increasing resolution to circumvent the difficulty of capturing large motions and speed up computations. Concurrently, the development of sophisticated tools like Spline pyramids [27] and wavelet transforms [28] provided novel methods for representing images and deformations at different scales. Many popular medical image registration algorithms now use multiresolution representations of deformations, most of which are combined with Gaussian smoothing of the images: the Demons algorithm [29] applies Gaussian filters to the vector fields, IRTK [30; 31] parameterized B-splines with grids of increasing resolution, ROMEO [32] uses an adaptive multigrid strategy to estimate optical flow, DARTEL [33] also uses a multigrid method to estimate single-flow velocity fields and Syn [34] estimates symmetric diffeomorphisms with Gaussian smoothing of velocity fields and later on B-spline regularization [35].

In the following, we will only review multiscale strategies related to our method, i.e. that are multiscale with regard to the deformation field.

2.1 Multiscale representations of deformations

Signal representation provides numerous mathematical tools to represent signals, and several registration algorithms modelled deformations using basis functions from the Fourier and Wavelet transforms. These multiscale representations typically aim to improve computational efficiency by playing on the relation between a representation and a regularization, which allows using lower dimensional parameterization of the velocity fields along the trajectory. Examples of this approach made use of the Fourier transform [36; 18; 37], Galerkin method [38], or orthogonal decomposition [39]. Our work falls into a different category, which aims to modify an initial method marginally, to reach better local minima without changing the considered representation or regularization of the LDDMM. The solution is to construct a sequence of nested problems: in our case, a multiscale representation of only the initial velocity field. The inherent scale information within the new parameters is leveraged to solve the optimization problem in a coarse-to-fine manner. This idea is already present in Christensen et al. [40] who parameterized displacement fields by a Fourier basis with an adapted regularization term and estimated frequency coefficients in a coarse-to-fine fashion. However, experiments showed that modelling deformations using a wavelet basis provides better spatial regularization compared to using a Fourier basis [41]. This advantage stems from the fact that wavelets not only capture the frequency of the signal but also its location and orientation, making them a more natural fit for hierarchical optimization strategies.

Wavelet-based deformation models cover various wavelet types (e.g. Haar, Cai Wang and (BV, L^2)) and deformations (e.g. displacement vectors, B-splines and elastic deformations). Displacement vectors [42; 43] and later free form deformations [44] were described and optimized in a multi-resolution fashion through the Cai Wang wavelet [45]. Gefen et al. [46] modelled elastic deformations using finite-supported, semi-orthogonal wavelet functions. Topology-preserving displacement fields were modelled by polynomial spline basis functions and controlling the Jacobian of the transformation [47; 48]. Recently, wavelets (BV, L^2) were employed to generate a hierarchical representation of diffeomorphisms in the hyperelasticity framework and perform coarse-to-fine optimization [49].

2.2 Multiscale registration in the LDDMM framework

In the LDDMM framework, the choice of the spatial regularizer restricts the range of possible deformations to those occurring at a single scale, which often proves unrealistic [50]. Thus, a variety of papers have focused on increasing the flexibility of the deformation model. These strategies can be broadly classified into two categories: the first one *simultaneously* estimates coexisting flows of different scales [50; 51; 19; 52; 53], and the second one composes multiple scale flows which are estimated *sequentially* [54; 55].

Risser et al. [50; 51] first introduced a multi-kernel extension of the LDDMM framework in which the deformation flow is defined by a weighted sum of Gaussian kernels whose widths are specified by the user. Weights are tuned in a semi-automatic manner during a pre-registration step. In this framework, the RKHS structure of the velocity fields is lost, and a new definition of the norm is used to ensure efficient computation of the flow. A spatially-varying version of this framework, the kernel bundle, used sparsity priors to allow the weights of the kernel mixture to vary across spatial locations [19]. Even though this algorithm proved efficient on the registration of landmark points, the increase in computational cost restricts its application to registration problems involving few parameters. Using an algorithm for the multiresolution decomposition of surfaces, the kernel bundle framework was also combined with a coarse-to-fine strategy wherein the resolution of cortical surface meshes is progressively increased along with that of the deformation field [53]. Multi-kernel approaches were further combined with deep learning optimization in order to learn a local regularizer from the data [56; 57]. However, these methods increase significantly the complexity of the mathematical model, and several optimization procedures are required to tune the networks parameters, the kernel pre-weights and the deformation parameters.

A related approach, based on modular deformations, enables the user to impose spatially-varying constraints on the deformation field [52]. Diffeomorphisms are built by superimposing deformations modules that encode local geometrical transformations. As in the kernel bundle framework, the space of vector fields is equipped with an adapted norm. The need for prior knowledge about the deformation modules limits the practical application of the algorithm.

A second and less explored axis of research constructed a hierarchical representation of deformations, based on non-coexisting vector flows of increasing resolution, which are estimated independently and then composed. In a theoretical paper, Modin et al. [54] extended (BV, L^2) wavelets to express diffeomorphisms as a composition of deformations of increasingly fine scales, which can be seen as a series of LDDMM steps. Despite the potential of this approach, the authors did not perform numerical experiments. A similar approach [55] constructed diffeomorphisms by composing a series of multiscale vector fields, which enables to progressively refine the deformation. Contrary to multi-kernel approaches, such strategies perform optimization in successive RKHS of increasingly finer resolution, in the spirit of coarse-to-fine strategies.

As we shall see in the following, our coarse-to-fine approach is more closely related to the one that composes multiple scale flows, in the sense that we perform optimization sequentially in sub-spaces of increasing resolution. However, our algorithm differs from the previous ones by the fact that our work brings changes to the optimization procedure rather than the deformation model: the multiscale structure is only used for the initial velocity field and so that the velocity fields are still defined, at core, by a single-scale RKHS. This simplifies the implementation of our algorithm while preserving the efficient optimization scheme of Durrleman et al. [15].

3 Model of diffeomorphic deformations

3.1 Large Deformation Diffeomorphic Metric Mapping

In the following, we consider a set of N images $(I_i)_{1 \leq i \leq N}$ of dimension d . We assume that each image I_i is a smooth deformation of a template I_{ref} plus an additive random white noise ϵ_i :

$$I_i = I_{ref} \circ \Phi_i^{-1} + \epsilon_i, \quad \forall i \in [1, n] \quad (1)$$

where Φ_i is the i^{th} template-to-subject deformation, and $I_{ref} \circ \Phi_i^{-1}$ denotes the action of the diffeomorphic deformation on the template.

In atlas estimation, one seeks to estimate the template image I_{ref} and the N template-to-subject deformations $(\Phi_i)_{1 \leq i \leq N}$. Note that registration is a specific case of atlas estimation where the template image is fixed and $N = 1$.

We choose to work in the LDDMM setting [4; 5], in which objects are deformed via deformations of the whole ambient space. This framework generalizes the linearized deformation setting in order to define diffeomorphic deformations that are invertible and smooth. Diffeomorphisms are constructed by integrating linearized deformations over time, which are considered as infinitesimal steps, according to the differential flow equation:

$$\begin{cases} \frac{dx(t)}{dt} = v_t(x(t)) \\ x(0) = x_0. \end{cases} \quad (2)$$

where $v_t \in V$ is an instantaneous velocity field belonging to a Hilbert space V and x can be seen as a particle moving along the curve $x(t)$ in the domain of interest D .

This model builds a flow of diffeomorphisms $\Phi_t : x_0 \rightarrow x(t) \forall t \in [0, 1]$. The diffeomorphism of interest Φ_1 is defined as the end point of the path $x(t)$, i.e.:

$$\forall x_0 \in D, \quad \Phi_1(x_0) = x(1).$$

Note that for any time t , Φ_t is indeed a diffeomorphism provided that the velocity field is regular enough, i.e. that it is continuous squared integrable.

3.2 Discrete parameterization of diffeomorphisms

Finally, we need to define an appropriate norm $\|\cdot\|_V$ for the Hilbert space V . To this end, we restrict ourselves to vector fields that belong to a RKHS [58] V defined by a kernel K_g . We also rely on the work of Durrleman et al. [15] to introduce a discrete parameterization of the velocity fields: we assume that the initial velocity field v_0 can be decomposed as a *finite* linear combination of the RKHS basis vector fields. The weights of the decomposition of a given deformation onto this basis are given by a set of momentum vectors $(\alpha_k(0))_k$ attached to k_g control points $(c_k(0))_k$.

$$v_{0,i}(x) = \sum_{k=1}^{k_g} K_g(x, c_{k,i}(0)) \alpha_{k,i}(0).$$

In this work, K_g is the Gaussian kernel:

$$K_g(x, y) = \exp\left(-\frac{\|x-y\|^2}{\sigma_g}\right) I_d, \text{ with } \sigma_g \text{ the kernel width and } I_d \text{ the identity matrix.}$$

The structure of the RKHS is such that the vector fields are continuous and squared integrable. As we shall see in Section 3.3, this will turned out to be important for performing optimization in a finite-dimensional setting.

It has been proved [13] that the vector fields that define geodesic deformations with respect to the norm $\int_0^1 \|v_{t,i}\|_V^2 dt$ keep the same structure along time and write according to:

$$v_{t,i}(x) = \sum_{k=1}^{k_g} K_g(x, c_{k,i}(t)) \alpha_{k,i}(t), \quad (3)$$

where for any time t , $\alpha_{k,i}(t)$ is the k^{th} momentum vector related to subject i and attached to the control point $c_{k,i}(t)$.

Furthermore, the trajectory of the control points $(c_{k,i}(t))_k$ and momentum vectors $(\alpha_{k,i}(t))_k$ is described by the Hamiltonian system equations [13]:

$$\begin{cases} \frac{dc_k(t)}{dt} = \sum_{l=1}^{k_g} K_g(c_k(t), c_l(t)) \alpha_l(t) \\ \frac{d\alpha_k(t)}{dt} = - \sum_{l=1}^{k_g} d_{c_k(t)}(K_g(c_k(t), c_l(t)) \alpha_l(t))^t \alpha_k(t) \end{cases} \quad (4)$$

with initial conditions $c_k(0)$ and $\alpha_k(0)$ for all $1 \leq k \leq k_g$.

Finally, one verifies that the kinetic energy along geodesic paths is preserved over time, i.e. $\forall t \in [0, 1], \|v_{t,i}\|_V = \|v_{0,i}\|_V$. This implies that a geodesic transformation is fully parameterized by the initial velocity field $v_{0,i}$. Hence, estimation of the diffeomorphism $\Phi_{1,i}$ boils down to a geodesic shooting problem. The system is deterministic, and we only need to optimize the initial conditions $\alpha_{0,i} = (\alpha_{k,i}(0))_k$ and $c_{0,i} = (c_{k,i}(0))_k$ for each subject i along with the template image I_{ref} .

3.3 Optimization

In this work, the position of the control points is fixed on a regular grid of spacing σ_g . We seek to minimize a cost function expressing a trade-off between matching accuracy and regularity of the transformation(s). With the parameterization introduced in the previous section, the

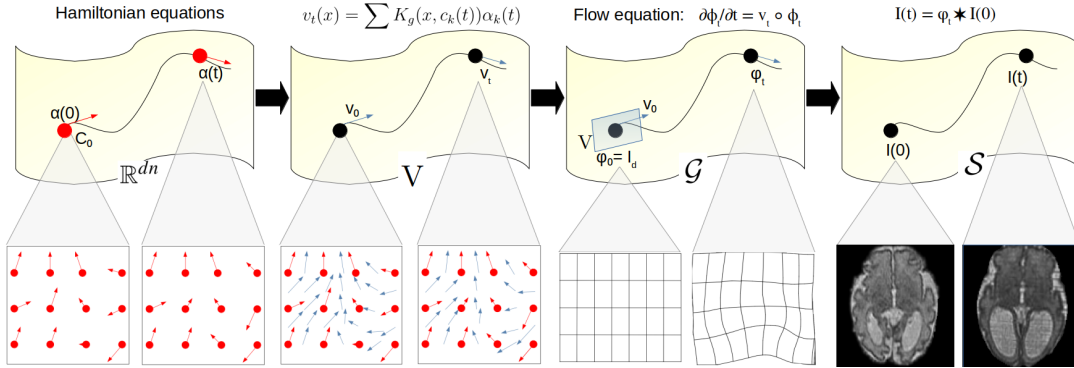


Fig. 1 Computing shape deformations from an initial image $I(0)$ and a set of initial momentum vectors $\alpha(0)$. 1- Given $\alpha(0)$ and the initial control points c_0 , integration of the Hamiltonian (Equation (4)) gives the trajectory of the momentum vectors $\alpha(t)$. 2- The velocity field v_t is computed by interpolating the momentum vectors with Equation (3). 3- Integration of the flow equation (Equation (2)) gives a flow of diffeomorphisms $(\Phi_t)_{t \in [0,1]}$. 4- Finally, Φ_t is applied to the object $I(0)$, giving the deformed image $I(t)$ at any time t .

cost function E is a function of I_{ref} and $(\alpha_{0,i})_i$:

$$E(I_{ref}, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N \left(\frac{d(I_i, I_{ref} \circ \Phi_{1,i}^{-1})^2}{\sigma_\epsilon^2} + \|v_{0,i}\|_V^2 \right), \quad (5)$$

The regularity term is the total kinetic energy along the geodesic path related to Φ_i . With the discrete parameterization chosen for the velocity fields (i.e. Equation (3)), the norm $\|\cdot\|_V$ can easily be computed and thus

$$\|v_{0,i}\|_V^2 = \sum_{j=1}^{k_g} \sum_{k=1}^{k_g} \alpha_{j,i}(0) K_g(c_{j,i}, c_{k,i}) \alpha_{k,i}(0). \quad (6)$$

Optimization is performed through gradient descent. Thus, we need to compute the gradient with respect to all the parameters. This is not a trivial task, but an efficient numerical scheme has been proposed in Durrleman et al. [15]. This algorithm relies heavily on the fact that the norm of the vector fields is the one of the RKHS and that the vector fields that are solutions to the problem remain a *finite* sum of kernels at all times (Equation (3)). This enables us to solve a finite-dimensional problem even though the functions we are looking for have infinite dimension. Given a set of momentum vectors α_0 , this strong structural property allows to efficiently deform shapes and compute the cost function by performing the following steps, which are illustrated in Figure 1: integrating Equation (4) gives the evolution of the momenta over time, the velocity field v_t at any time t is computed with Equation (3), the flow of diffeomorphisms $(\Phi_t)_{t \in [0,1]}$ is obtained by solving Equation (2), and the template image I_{ref} is deformed with

the flow: $\Phi_1 \star I_{ref}$. It is then straightforward to infer the distance between the deformed template and the target objects and to compute the regularity term with Equation (6). Finally, one can compute the gradient of the cost function with respect to the template, and use a backward integration along time to compute the gradient with respect to the momentum vectors. These steps correspond to the lines 14-18 in Algorithm 1.

The algorithm is publicly available as part of the open-source software Deformetrica [59].

3.4 Unrealistic local minima issue

There is a dependency between the scale of the kernel and the number of parameters, as a constant vector field has to be well approximated by the finite sum: this imposes that the scale σ_g should be related to the distance between control points. Here, the kernel K_g not only controls the *regularity* of the deformations, but also the *number of parameters* to optimize. This leaves the user to find the balance between a large kernel, (i.e. few control points), which generates smooth but less accurate transformations, and a smaller one, which penalizes large displacements and increases the risk of converging towards unrealistic solutions.

Further, as the optimization problem is not convex, the gradient descent algorithm converges towards a solution that depends on the initialization. The more the numbers of subjects and parameters are high, the more complex the energy landscape of the problem becomes. In the original approach of Durrleman et al. [15], a first step towards the estimation of multiscale deformations was taken by optimizing both the *number* and *position* of the control points. In

areas with low variability, points are inactivated by a L1 sparsity prior on the momenta. Unfortunately, the width of the kernel remains the same and in practice, the vector fields can only be set to zero in image areas of null intensity. Another issue is that the position of the control points cannot change significantly if the distance between points is low. Yet, it is exactly when the number of parameters is high that the risk of converging towards an unrealistic local minimum is higher.

One idea would be to change locally the scale of the kernel so that we can estimate non evenly smooth vector fields, i.e. vector fields with spatially varying scales. Even if such parameterization can be written, one loses the RKHS structure and thereby the ease of computation.

In this paper, we wish to address these two related issues: the dependency of the algorithm on the initialization, which restrains the number of parameters that can be properly optimized, and the difficulty in estimating vector fields that have *locally* varying regularity. In the next section, we will describe a reparameterization of the vector fields which enables us to impose smoothness constraints on the deformations and progressively relax them in a coarse-to-fine fashion. In this way, the algorithm can cope with non evenly smooth transformations while using a small kernel and remaining in the original RKHS setting.

4 Multiscale deformations

In this section, we propose a multiscale optimization procedure based on a Haar-like wavelet representation of the initial velocity fields. This strategy has the advantage of making the algorithm less dependent on the initialization while favoring more multiscale deformations. We rely on the finite parameterization of the velocity fields as a linear combination of RKHS basis elements [15] within the LDDMM setting. Importantly, our strategy enables us to preserve this structural assumption and the efficient numerical scheme that follows. We will show that our algorithm generates more natural template images as well as higher stability regarding the initialization.

4.1 Overview

For the sake of clarity, we provide the reader with an overview of our multiscale strategy. In the *original* algorithm for atlas estimation,

we optimize a cost function E . Two types of parameters are optimized through gradient descent: the template image I_{ref} and N sets of momentum vectors $\alpha_{0,i}$ that parameterize the template-to-subjects velocity fields $v_{0,i}$. The *original* optimization, summarized in Figure 2, iterates between two classical steps (note that the subscripts i have been dropped for clarity):

1. Computation of the gradients ∇E_{α_0} and $\nabla E_{I_{ref}}$ (Section 3.3)
2. Parameters update: $\alpha_0(j+1) \leftarrow \alpha_0(j) - h \times \nabla E_{\alpha_0}$ and $I_{ref}(j+1) \leftarrow I_{ref}(j) - h \times \nabla E_{I_{ref}}$, with h the step size.

In the *multiscale* strategy, summarized in Figure 2, we only modify the parameters update step of the momentum vectors by replacing it with the following steps:

- (i) We use the wavelet transform to obtain a multiscale representation of the gradient of E with regard to the momenta: $\nabla E_{\beta_0} \leftarrow FWT(\nabla E_{\alpha_0})$. (This reparameterization is detailed in Section 4.2)
- (ii) The coordinates of the initial velocity fields v_0 in the wavelet basis are updated: $\beta_0(j+1) \leftarrow \beta_0(j) - h \times \nabla E_{\beta_0}$.
- (iii) The wavelet coefficients in $\beta_0(j+1)$ whose scale is smaller than a current scale S_j are set to zero.
- (iv) The coordinates of v_0 in the RKHS basis are recovered with the Inverse Wavelet Transform: $\alpha_0(j+1) \leftarrow IWT(\beta_0(j+1))$; I_{ref} is updated as in the original step (2). (The coarse-to-fine optimization steps are detailed in Section 4.3.)

4.2 Reparameterization of the initial velocity fields

In this section, we introduce a multiscale reparameterization of the initial velocity fields based on the Haar wavelet transform. The choice of the Haar wavelet is motivated by its ease of implementation and the orthogonality of its transform. We introduce the definition and properties of the continuous Haar Wavelet representation, recall how this construction can be extended to a representation of discrete signals defined on a grid, and we demonstrate how this can be used to obtain a Haar-like representation of the initial velocity fields.

4.2.1 The continuous Haar wavelet transform

Here, we describe the decomposition of a real signal f defined on the d -dimensional space \mathbb{R}^d

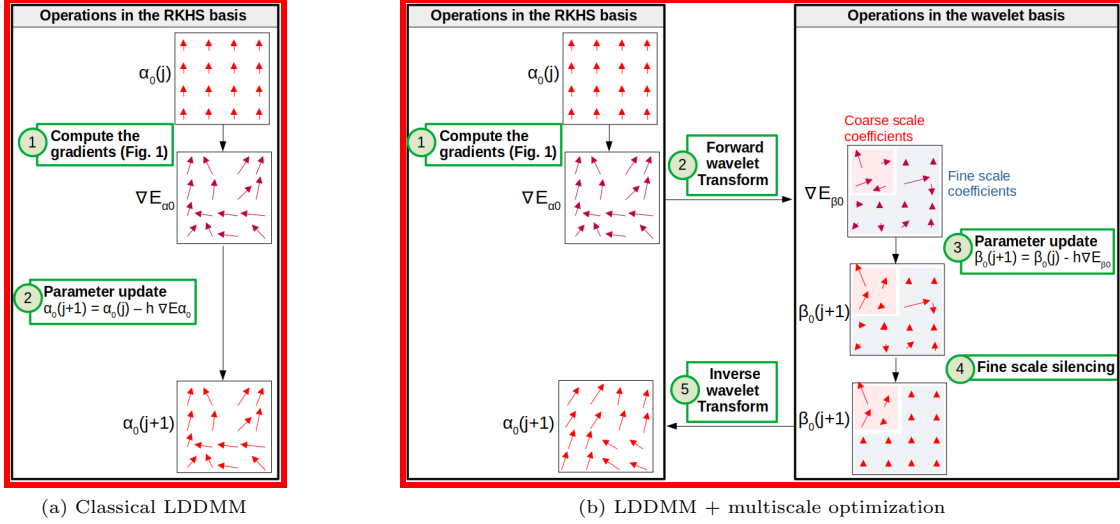


Fig. 2 Overview of the original LDDMM algorithm [15] and the coarse-to-fine optimization strategy in dimension 2. For the sake of clarity, only a **single** gradient descent iteration is presented in each panel. The original algorithm iterates between two classical steps: computing the gradients (1) and updating the parameters (2). In the multiscale strategy, step (1) is preserved and step (2) is replaced by another procedure involving updating the representation of the momentum vectors in the wavelet basis. Red arrows denote the coordinates of the momentum vectors either in the RKHS basis (for $\alpha_0(j)$) either in the wavelet basis (for $\beta_0(j)$). Purple arrows denote the gradients of the cost function ∇E_{α_0} and ∇E_{β_0} with regard to α_0 and β_0 (respectively).




Fig. 3 The 1D Haar wavelet. At scale 0, V_0 is the space of piecewise constant functions of size 1. A basis for V_0 , i.e. $\{\Phi_{0,k}\}_k$, is obtained by translating the scaling function $\Phi_{0,0} = \psi^L$ by factors $k \in \mathbb{Z}$. A basis for W_0 , i.e. $\{\psi_{0,k}\}_k$, is obtained by translating the mother wavelet function $\psi_{0,0} = \psi^H$ by factors $k \in \mathbb{Z}$. At scale s , a basis for the space V_s , i.e. $\{\Phi_{s,k}\}_k$, is obtained by dilating $\Phi_{0,0}$ by 2^s , translating it by $2^s k$ and normalizing it by $2^{-s/2}$. A basis for W_s , i.e. $\{\psi_{s,k}\}_k$, is obtained by performing the same operations on $\psi_{0,0}$.

into a Haar Wavelet basis [28; 60]. The wavelet representation decomposes f into a linear combination of basis functions which have different *resolutions*, *locations* and *orientations*. This representation relies on a collection of embedded spaces V_s that contain functions said of scale s .

In the case of the Haar Wavelet, V_s is the space of piecewise constant functions on a regular grid of size 2^s . Any function f can be approximated in this space by computing a local average: the mean value in each sub-square of the grid.

We define the d -dimensional *scaling function* ϕ by

$$\phi(x) = \prod_{i=1}^d \psi^L(x_i)$$

where $x \in \mathbb{R}^d$ and ψ^L is the 1D piecewise constant function

$$\psi^L(z) = \begin{cases} 1 & \text{for } 0 \leq z < 1 \\ 0 & \text{otherwise} \end{cases}.$$

As illustrated in Figure 3, approximating f at scale s amounts to projecting f onto the space spanned by the orthonormal family

$$\left\{ \phi_{s,k}(x) = 2^{-sd/2} \phi(2^{-s}x - k) \right\}_{k \in \mathbb{Z}}$$

where $\phi_{s,k}$ is the scaling function rescaled by 2^s and then translated by $k2^s$. The factor $2^{-sd/2}$ is a normalization factor that ensures that the $\phi_{s,k}$ have unit energy.

When transitioning from the approximation at scale s to the approximation at the coarser scale $s+1$, some details of f are lost. These details belong to the orthogonal complement W_{s+1} of the space V_{s+1} in V_s . A basis of this space can be obtained by defining the

d -dimensional oriented wavelet functions

$$\psi^o(x) = \prod_{i=1}^d \psi^{o_i}(x_i)$$

where $x \in \mathbb{R}^d$, $o \in \{H, L\}^d$, $\exists i, o_i = H$ and

$$\psi^H(z) = \begin{cases} 1 & \text{for } 0 \leq z < 0.5 \\ -1 & \text{for } 0.5 \leq z < 1 \\ 0 & \text{otherwise} \end{cases}$$

As illustrated in Figure 3, an orthonormal basis of W_s is given by

$$\left\{ \psi_{s,k}^o(x) = 2^{-sd/2} \psi^o(2^{-s}x - k) \right\}_{k \in \mathbb{Z}, \exists i, o_i = H}$$

where $\psi_{s,k}^o$ is the wavelet function of orientation o rescaled by 2^s and then translated by $k2^s$. Note that functions ψ^L and ψ^H act respectively as low and high pass filters. Their combination yields oriented high pass filters, e.g., for $d = 2$, there exist three wavelet functions ψ^{HL} , ψ^{LH} and ψ^{HH} , that express details of the signal along vertical, horizontal and diagonal orientations (respectively).

We can thus decompose any function f in V_s in the two following ways:

$$\begin{aligned} f &= \sum_k a_{s,k} \phi_{s,k} \\ &= \sum_k a_{s+1,k} \phi_{s+1,k} + \sum_{o,k} d_{s+1,k}^o \psi_{s+1,k}^o \end{aligned}$$

As $\sum_k a_{s+1,k} \phi_{s+1,k}$ belongs by construction to the space V_{s+1} , we can further decompose it into a projection onto V_{s+2} and a projection onto W_{s+2} . Repeating this scheme up to scale S leads to the following multiscale decomposition of f :

$$f = \sum_k a_{S,k} \phi_{S,k} + \sum_{s'=s+1}^S \sum_{o,k} d_{s',k}^o \psi_{s',k}^o$$

The classical wavelet construction is concluded by letting s go to $-\infty$, enabling one to decompose any measurable bounded function in such bases.

More importantly, going from the decomposition in V_s to the decomposition in the spaces V_S and $(W_{s'})_{s < s' \leq S}$ corresponds to a change of basis and thus to a discrete operation going from the coefficients $(a_{s,k})_k$ to the coefficients

$(a_{S,k})_k$ and $(d_{s',k}^o)_{s < s' \leq S, k, o}$. This transformation is called the Forward wavelet transform (FWT) and its inverse the Inverse wavelet transform (IWT). Both can be computed directly on the coefficients without relying on the continuous basis functions.

4.2.2 Haar wavelet applied to grids

In our algorithm, rather than using the continuous Haar wavelet decomposition, we have implemented the related discrete Haar wavelet decomposition on a d -dimensional grid $\llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket$, illustrated in Figure 4 in 2D. This transform corresponds to a discrete change of basis related to the continuous Haar transform.

More precisely, to any discrete function $(a_{0,k})_k$ on the grid, one can associate the continuous function f of V_0 defined by

$$f = \sum_{k \in \llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket} a_{0,k} \phi_{0,k},$$

where the $(a_{0,k})_k$ are interpreted as the approximation coefficients at scale 0 of function f .

By construction, f is a piecewise constant function on the related continuous grid. Using the FWT algorithm up to scale S , this function can be decomposed as follows:

$$f = \sum_k a_{S,k} \phi_{S,k} + \sum_{s=1}^S \sum_{o,k} d_{s,k}^o \psi_{s,k}^o$$

One verifies that $a_{S,k} = 0$ when $k_i < 0$ or $k_i 2^S > K_i$ and $d_{s,k}^o = 0$ when $k_i < 0$ or $k_i 2^s > K_i$. Thus, these sums have only a finite number of coefficients. Further, when the K_i are powers of 2, i.e. $K_i = 2^{\zeta_i}$, we impose that the decomposition cannot exceed a maximum scale $S_{max} = \min_i(\zeta_i)$, so that the scaling function support remains within $\prod_i \llbracket 0, K_i \rrbracket$. Thus, when $S \leq S_{max}$, the previous equation reduces to:

$$\begin{aligned} f &= \sum_{0 \leq k_i < 2^{\zeta_i - S}} a_{S,k} \phi_{S,k} \\ &+ \sum_{s=1}^S \sum_{0 \leq k_i < 2^{\zeta_i - s}} d_{s,k}^o \psi_{s,k}^o, \end{aligned} \quad (7)$$

which corresponds exactly to an orthonormal change of basis.

To summarize, from the coefficients $(a_{0,k})_k$, one can compute the coefficients $(a_{S,k})_{0 \leq k < 2^{S_{max} - S}}$ and





Fig. 4 Decomposition of a 4-by-4 grid in two bases. The letters a and d refer to approximation and detail coefficients, respectively. Subscripts indicate the *scale* of the coefficient and the x-y *position* of the related wavelet function. Superscripts indicate the *orientation* of the wavelet function. In the grids, empty cells denote null values.

$(d_{s,k}^o)_{1 \leq s \leq S, 0 < k < 2^{S_{max}-s}, o}$ with the FWT algorithm and perform the inverse operation with the IWT algorithm.

We denote by M_{FWT} and M_{IWT} the $\prod_i K_i \times \prod_i K_i$ matrices associated to the transformations FWT and IWT. M_{FWT} and M_{IWT} are related to one another: we have obviously $M_{FWT} * M_{IWT} = Id$ and, as these are orthonormal transforms, we also verify that $M_{IWT}^T = M_{FWT}$. Furthermore, there exist fast implementations of these transforms that have a linear complexity with respect to the number of coefficients [60].

In this work, we have implemented the FWT using the fast lifting scheme described in [28] (Chapter 7.8). This scheme is strictly equivalent to the previous description when the grid size is a power of 2, with the difference that it handles non-dyadic grids through improved computations at the boundaries and that FWT and IWT remain orthogonal transforms. The pseudocode for the algorithms FWT and IWT can be found in Appendix (Algorithms 5 and 6, respectively), along with a brief explanation and an illustration of how the FWT algorithm operates on a non-dyadic grid (Figure 1).

Note that these algorithms are available in a public Git repository¹.

¹<https://github.com/fleurgaudferneau/Defometrica-multiscale/>

4.2.3 Preservation of the RKHS structure of the velocity fields

The previous subsection transposed the classical Haar description of continuous functions to discrete functions defined on a grid. To implement a coarse-to-fine initialization approach for atlas estimation, we will describe the initial momentum vectors not through their values on the grid but by the decomposition of these values in the discrete Haar basis. The multiscale structure will be used to obtain a smooth initial field by setting fine-scale coefficients to 0. As we will show in this section, this does not change the fact that the initial vector fields $v_{0,i}$, and thus all the vector fields $v_{t,i}$, are by construction finite combinations of the RKHS kernel K_g .

If we consider two spaces V_s and $V_{s'}$ with $s \geq s'$, we observe that any function $\phi_{s,k} \in V_s$ or $\psi_{s,k}^o \in V_s$ is also in the finer space $V_{s'}$. In particular, if we take $s' = 0$, $\phi_{s,k}$ and $\psi_{s,k}^o$ can be decomposed as a linear combination of $\phi_{0,k'}$:

$$\phi_{s,k} = \sum_{k'} \gamma_{s,k,k'} \phi_{0,k'} \quad (8)$$

$$\psi_{s,k}^o = \sum_{k'} \gamma_{s,k,k'}^o \phi_{0,k'} \quad (9)$$

where $\gamma_{s,k,k'}$ and $\gamma_{s,k,k'}^o$ are some fixed real numbers.

Thus, for a function $v = \sum_k a_{0,k} \phi_{0,k}$, its wavelet coefficients $((a_{s,k})_k, (d_{s,k}^o)_{1 \leq s \leq S, o, k})$

can be computed as

$$\begin{aligned} a_{S,k} &= \sum_{k'} \gamma_{S,k,k'} a_{0,k'} \\ d_{s,k}^o &= \sum_{k'} \gamma_{s,k,k'}^o a_{0,k'} \end{aligned}$$

Further, we recall that we optimize shape transformations by using the parameterization of Durrleman et al. [15], i.e. the initial geodesic vector field of a given subject i is defined as a finite linear combination of identical Gaussian kernels that are evaluated at an initial set of points $(c_{k,i}(0))_k$. In our scheme, we set the initial points on a grid so that this initial velocity field writes as

$$v_{0,i}(x) = \sum_{k \in \llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket} K_g(x, c_{k,i}(0)) \alpha_{k,i}(0), \quad (10)$$

where $\alpha_{k,i}(0)$ is a momentum vector attached to $c_{k,i}(0)$.

Instead of optimizing $v_{0,i}$ by optimizing its momentum vectors $\alpha_{k,i}(0)$, we will optimize them in the wavelet domain under the constraints that the finer-scale coefficients are equal to 0.

More precisely, we define

$$((a_{S,k,i})_k, (d_{s,k,i}^o)_{1 \leq s \leq S, o, k}) = FWT((\alpha_{k,i})_k)$$

where the Wavelet Transform has been applied to the momentum vectors component by component. Note that by construction,

$$(\alpha_{k,i})_k = IWT((a_{S,k,i})_k, (d_{s,k,i}^o)_{1 \leq s \leq S, o, k}).$$

Performing optimization with the wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{1 \leq s \leq S, o, k}$ instead of the momentum vectors $(\alpha_{k,i})_k$ amounts to switch from the description of Equation (10) to

$$\begin{aligned} v_{0,i}(x) &= \sum_k a_{S,k,i} \widetilde{\phi}_{S,k}(x) \\ &+ \sum_{s=1}^S \sum_{o,k} d_{s,k,i}^o \widetilde{\psi}_{s,k}^o(x) \end{aligned}$$

where $\widetilde{\phi}_{s,k}$ and $\widetilde{\psi}_{s,k}^o$ are functions defined by replacing $\phi_{0,k'}$ by $K_g(\cdot, c_{k'}(0))$ in Equations 8 and 9:

$$\widetilde{\phi}_{s,k} = \sum_{k'} \gamma_{s,k,k'} K_g(\cdot, c_{k'}(0))$$

$$\widetilde{\psi}_{s,k}^o = \sum_{k'} \gamma_{s,k,k'}^o K_g(\cdot, c_{k'}(0)).$$

Even if $v_{0,i}$ is now defined through the vectorial wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{s,o,k}$ instead of the momentum vectors $(\alpha_{k,i})_k$, it remains a linear combination of the $K_g(x, c_k(0))$ so that we are still in the setting of Durrleman et al. [15] and can rely on

$$v_{t,i}(x) = \sum_k K_g(x, c_{k,i}(t)) \alpha_{k,i}(t).$$

Note that we do not use the Haar parameterization outside the initialization. Indeed, the initial grid is deformed under the action of the diffeomorphism $\Phi_{t,i}$ when the time evolves, so that our Haar coefficients would be hard to interpret for $t > 0$.

4.3 Coarse-to-fine atlas estimation

4.3.1 Reparameterization of the initial velocity fields

The key difference between our scheme and the one of Durrleman et al. [15] is the use of the Haar parameterization in a fixed grid for the initial velocity fields. Thus, we want to optimize the following cost function:

$$\begin{aligned} E(I_{ref}, (\beta_{0,i})_{1 \leq i \leq N}) &= \\ \sum_{i=1}^N \left(\frac{d(I_i, I_{ref} \circ \Phi_{1,i}^{-1})^2}{2\sigma^2} + \|v_{0,i}\|_V^2 \right), \end{aligned} \quad (11)$$


where $\beta_{0,i}$ is a set of wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{s,k,o}$, related to the momentum vectors $\alpha_{0,i}$ by $\alpha_{0,i} = IWT(\beta_{0,i}) = M_{IWT} \beta_{0,i}$.

Since $\alpha_{0,i} = M_{IWT} \beta_{0,i}$, there is a relationship between the gradient of the cost function with respect to the wavelet coefficients $\nabla_{\beta_{0,i}} E$ and the gradient with respect to the momenta $\nabla_{\alpha_{0,i}} E$:


$$\begin{aligned} \nabla_{\beta_{0,i}} E &= M_{IWT}^T \nabla_{\alpha_{0,i}} E \\ &= M_{FWT} \nabla_{\alpha_{0,i}} E \\ &= FWT(\nabla_{\alpha_{0,i}} E) \end{aligned}$$

where we have used the fact that $M_{IWT}^T = M_{FWT}$ because the transform is orthonormal.

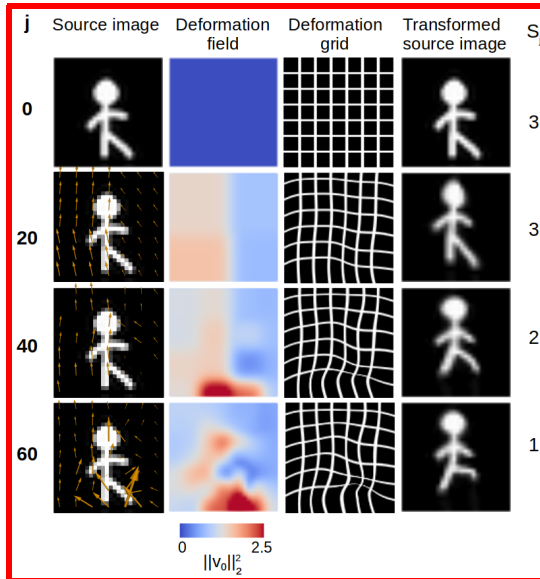
Therefore, if we compute the gradient $\nabla_{\alpha_{0,i}} E$ with Durrleman et al. numerical scheme



(a) Source and target images



(b) Classical LDDMM algorithm



(c) LDDMM + Multiscale algorithm

Fig. 5 Original LDDMM and LDDMM combined with multiscale optimization applied to a registration example. The source and target images (panel (a)) present feature differences at a large scale (character translation) and at a finer scale (orientation of the arms and legs). Registration was performed with $\sigma_g = 4$ ($k_g = 49$ control points). For each algorithm, we display the source-to-target vector fields (orange arrows, scaling factor = 5), the corresponding deformation field L_2 norm in RGB, the deformation grid and the transformed source image every 20 iterations until convergence. S_j is the scale constraining the momentum vectors at iteration j . Notice how the multiscale strategy first estimates coarse displacements to move the character to the top right, then finer transformations to adjust the position of the arms. The original LDDMM algorithm applies fine-scale deformations to the entire image (except the borders), while the multiscale algorithm estimates fine deformations around the character’s right leg, and smoother deformations elsewhere.

[15], we can then easily obtain the values of $\nabla_{\beta_{0,i}} E$ for almost the same cost as the one of this latter gradient. As illustrated in Fig. 2, this means that optimization can now be performed in the wavelet domain: we can compute $\nabla_{\alpha_{0,i}} E$ using the original LDDMM algorithm, and, instead of updating the $\alpha_{0,i}$, compute $\nabla_{\beta_{0,i}} E$ and update the $(\beta_{0,i})_i$. Without further changes, such algorithm would lead to exactly the same results as the original algorithm.

To obtain different and hopefully better results, we enforce some constraints on the wavelet coefficients of the initial velocity fields. Namely, we use a coarse-to-fine initialization strategy by optimizing first the initial velocity fields whose wavelets coefficients are null at the finest scales and adding progressively these fine scale coefficients. In the following, we describe in detail our procedure which is summarized in Algorithm 1 and illustrated on a simple registration example in Figure 5. We chose a source and a target image presenting both large and

small scale differences to illustrate clearly how the original and coarse-to-fine algorithms cope with multiscale deformations.

4.3.2 Coarse-to-fine initialization

The coarse-to-fine optimization can be seen as an initialization of each new scale with the optimal template-to-subject deformations of the previous coarser scale. More precisely, at iteration j , we only optimize the wavelet coefficients of the vector fields whose scales are above or equal to a current decreasing scale S_j . Since the original RKHS setting is preserved, this can easily be done by computing the gradient $\nabla_{\alpha_{0,i}} E$ with the efficient numerical scheme of Durrleman et al. [15] (lines 14-18 in Algorithm 1), applying FWT to $\nabla_{\alpha_{0,i}} E$ to derive the gradient with respect to $\beta_{0,i}$ (line 19) and then setting to 0 the wavelet coefficients whose scale is strictly smaller than S_j (line 22). We then update the coefficients $\beta_{0,i}$ with the modified gradient and

Algorithm 1 Multiscale optimization.

- 1: **Input**
- 2: Set of images $(I_i)_{1 \leq i \leq N}$, template image I_{ref} , kernel width σ_g , step size h , initial scale S_0
- 3: **Initialization**
- 4: $j \leftarrow 0$
- 5: $c_0 \leftarrow$ Grid of control points with spacing σ_g
- 6: Template image $I_{ref}(j) \leftarrow I_{ref}$
- 7: Momentum vectors $\alpha_{0,i}(j) \leftarrow 0 \quad \forall i \in [1, N]$
- 8: Initialize the gradients $\nabla_{\alpha_{0,i}} E \leftarrow 0 \quad \forall i \in [1, N]$ and $\nabla_{I_{ref}} E \leftarrow 0$
- 9: $\beta_{0,i}(j) \leftarrow FWT(\alpha_{0,i}(j)) \quad \forall i \in [1, N]$
- 10: Current scale $S_j \leftarrow S_0$
- 11: **repeat**
- 12: $j \leftarrow j + 1$
- 13: **for** each subject i **do**
- 14: Compute the evolution of $\alpha_i(t)$ using Equation (4) {Gradients computation}
- 15: Compute Φ_i by solving the flow equation (Equation (2))
- 16: Deform the template I_{ref} with Φ_i
- 17: Compute the energy E (Equation (5))
- 18: Compute the gradients $\nabla_{\alpha_{0,i}} E$ & $\nabla_{I_{ref}} E$
- 19: Wavelet transform $\nabla_{\alpha_{0,i}} E$ with Algorithm 5: $\nabla_{\beta_{0,i}} E \leftarrow FWT(\nabla_{\alpha_{0,i}} E) = (a_{S_{max},k}^i)_{k \cup (d_{s,k}^{i,o})_{1 \leq s \leq S_{max},k,o}}$
- 20: **for** each detail coefficient $d_{s,k}^{i,o}$ **do**
- 21: **if** $s < S_j$ **then**
- 22: $d_{s,k}^{i,o} \leftarrow 0$ {Finer scale silencing}
- 23: **end if**
- 24: **end for**
- 25: **end for**
- 26: $\beta_{0,i}(j) \leftarrow \beta_{0,i}(j-1) - h \times \nabla_{\beta_{0,i}} E$ for all subjects i {Parameter update}
- 27: $\alpha_{0,i}(j) \leftarrow IWT(\beta_{0,i}(j)) \forall i \in [1, N]$ for all subjects i {Algorithm 6}
- 28: $I_{ref}(j) \leftarrow I_{ref}(j-1) - h \times \nabla_{I_{ref}} E$
- 29: Compute the mean residuals ∇_j according to Equation (12)
- 30: **if** $\frac{\nabla_{j+1} - \nabla_j}{\nabla_{j-1}} < 0.01$ and $S_j > 1$ **then**
- 31: $S_j \leftarrow S_{j-1} - 1$ {Scale refinement step}
- 32: **end if**
- 33: **until** Convergence
- 34: **return** Template image I_{ref} and momentum vectors $\alpha_{0,i}$

recover the updated $\alpha_{0,i}$ using the IWT function (line 26). If we iterate without modifying the current scale, we optimize the cost function in a subspace of functions that are simpler than in the original algorithm, the wavelet transform scale limitation acting as a regularizer.

As we want to optimize on the full set of functions defined by the momentum vectors, we progressively decrease the current scale S_j . We propose to decrease the scale when we are close to convergence at the current scale (line

31). This is measured by computing the mean residual value over subjects at iteration j :

$$\Delta_j(x) = \frac{1}{N} \sum_i \|I_{j,0} \circ \Phi_{j,i}^{-1} - I_i\|_2^2, \quad (12)$$

where j denotes the current iteration.

If the residual decrease with respect to the previous iteration is below a threshold of 1%, we decide that the algorithm is close to convergence. In the case of Figure 5, the algorithm starts at $S_0 = 3$, performs optimization until (almost) convergence at this scale, goes to scale 2 and performs the subsequent scale transitions in the same manner.

Our optimization procedure ensures that the momenta belonging to the same area are updated with identical values. At a given scale S_j , the velocity fields can vary spatially only at scales coarser than $S_j - 1$. In other words, at scale S_j , the i^{th} initial velocity field implicitly writes as follows:

$$v_{0,i}(x) = \sum_k a_{S_{max},k,i} \widetilde{\phi}_{S_{max},k}(x) + \sum_{s=S_j}^{S_{max}} \sum_{k,o} d_{s,k,i}^o \widetilde{\psi}_{s,k}^o(x)$$

where $\widetilde{\phi}_{S_{max},k}$ and $\widetilde{\psi}_{s,k}^o$ are linear combinations of localized Gaussian kernels $K_g(x, c_k(0))$.

When the algorithm reaches scale 1, the momenta are updated independently of each other. Importantly, unlike previous approaches that represented deformations in a wavelet basis [41; 61; 42; 44; 46], when the algorithm reaches this finest scale, the momentum vectors are free of constraints and the parameterization of the velocity fields is equivalent to its original definition, i.e. a sum of localized small Gaussian kernels. Thus, in theory, our coarse-to-fine algorithm could reach the same solutions as the original one, but as we will see in Section 5, the coarse-to-fine numerical scheme converges to better solutions.

The initial scale S_0 takes values between 1 and S_{max} , i.e. the maximum scale of the wavelet coefficients $\beta_{0,i}(0)$. If $S_0 = 1$, the classical single-scale optimization is performed. By default, $S_0 = S_{max}$ so that coarse-to-fine optimization is performed from scale S_{max} to scale 1.

Our code is available in a public Git repository².

4.3.3 Complexity

Contrary to the previous coarse-to-fine algorithms developed in the LDDMM framework, our strategy does not add any complexity to the mathematical model. The parameterization of the velocity fields remains identical to that of Durrleman et al. [15]. Computation of the gradients and subsequent cost also remains identical. The only additional complexity arises from the algorithms FWT and IWT, which are of linear complexity.

5 Experiments

In this section, we evaluate the benefits of the wavelet reparameterization of the initial velocity fields on different tasks and datasets. In Section 5.1, we use a registration experiment on toy data to illustrate the way our coarse-to-fine algorithm operates and we compare it to the multi-kernel algorithm [50]. Then, we compare the performance of our algorithm to that of the original, single-scale LDDMM version [15] on three atlas estimation tasks of increasing complexity. The training phase consists in atlas estimation and the test phase consists in registering the estimated template image to a set of new images.

Experiments are run on an Ubuntu 18.04.5 machine equipped with a NVIDIA GPU driver with 12 GB memory. The original version of the algorithm is available in Deformetrica Version 4.3.0 [59]. Optimization for the classical LDDMM and LDDMM-multiscale algorithms relies on a gradient descent algorithm in which the step sizes h are first scaled by the squared norm of the gradients and then diminished by a backtracking algorithm to guaranty a descent. Unless stated otherwise, the following parameters are used in all experiments: $\sigma = 0.1$ in the cost criterion; initial step size $h = 0.01$; convergence threshold = 0.0001. The minimum number of iterations between successive coarse-to-fine steps is set to 5. The initial template image for atlas estimation is given by the mean of the intensities of the training images.

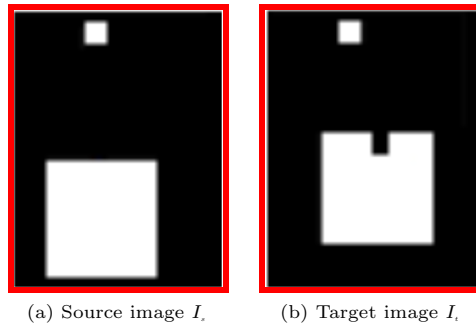


Fig. 6 Registration experiment from Risser et al. [50]

5.1 Toy experiment

In this experiment, we propose to assess the effect of the parameter S_0 on the performance of the multiscale algorithm and compare its behaviour to that of the multi-kernel framework [50]. To this end, we reproduce a toy experiment presented in Risser et al.: the goal is to register a source image I_s onto a target image I_t . Both images, visible in Fig. 6, contain a small (4 by 4 pixels) and a large (20 by 20 pixels) square. From I_s to I_t , the large square is translated to the top-right and a small indentation (4 by 2 pixels) appears on the top of the square. The small square remains at the same location. Thus, the registration task requires a coarse scale deformation (displacement of the large square), a small scale deformation (creation of the indentation), all the while ensuring that the small square is not deformed.

Performance is assessed by using the same metrics as the original experiment:

- **Total residuals** Δ_J , i.e. sum of squared differences in the source image domain at convergence
- **Total residuals in a region of interest covering the indentation**, denoted by $\Delta_{J(\text{ROI})}$
- **Standard deviation of the Jacobian determinant** (SD(J)) [62], quantifying the amount of local volume differences between the template and target images to assess the smoothness of transformations.

In addition, we also compute the algorithms runtimes.

5.1.1 Effect of the initial scale

In the first part of this experiment, we assess the effect of the initial scale S_0 on the performance of the multiscale algorithm. Registration

²<https://github.com/fleurgaudferneau/Deformetrica-multiscale/>




Fig. 7 Registration quality depending on the initial scale S_0 for different kernel widths σ_g .

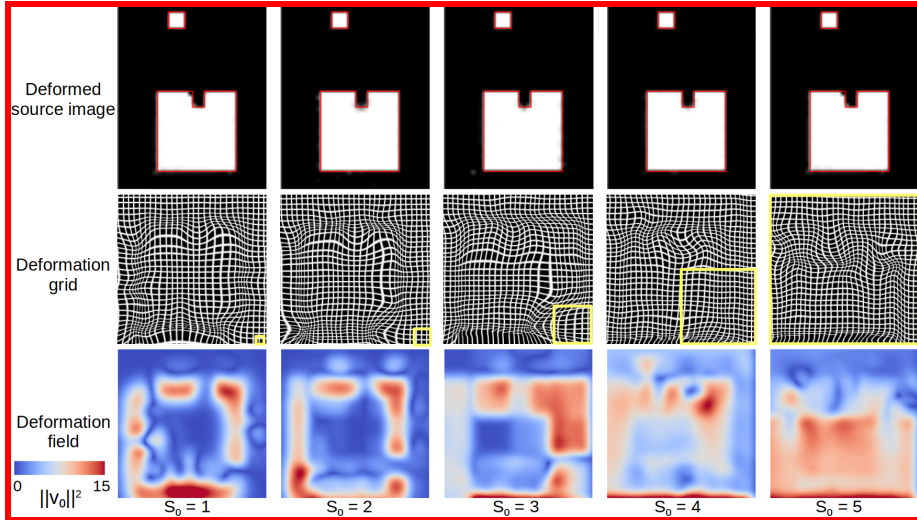


Fig. 8 Registration between I_s and I_t with the LDDMM algorithm combined with multiscale optimization using $\sigma_g = 2$ and different initial scales S_0 . Top row: deformed source images at $t = 1$. The red lines represent the borders of the shapes in I_t . Middle row: corresponding deformation grids, zoomed in on the large square area. The yellow squares represent the support of the wavelet functions of the initial scale S_0 , i.e. the size of the areas in which the vector fields are constant at the beginning of optimization: 32×32 for $S_0 = 5$; 16×16 for $S_0 = 4$, 8×8 for $S_0 = 3$, 4×4 for $S_0 = 2$, and 2×2 pixels for $S_0 = 1$. Bottom row: corresponding deformation fields L^2 norm in RGB, zoomed in on the large square area.

between I_s and I_t is performed with the multiscale algorithm using kernel sizes $\sigma_g \in \{1.7, 2, 2.5, 3\}$, corresponding to a number of control points $k_g \in \{841, 625, 400, 289\}$. For each parameter σ_g , we run the multiscale algorithms with all the possible initial scales $S_0 \in \{1, 2, \dots, S_{max}\}$. Note that $S_0 = 1$ corresponds to the original LDDMM algorithm.

Fig. 7 presents the total residuals Δ_J and $\Delta_J(\text{ROI})$, the standard deviation of the Jacobian $\text{SD}(J)$ and the runtimes after registration as functions of the initial scale S_0 . Fig. 8 displays the registration results obtained with $\sigma_g = 2$. Compared to our baseline (i.e. $S_0 = 1$), the use of our coarse-to-fine strategy decreases the total residuals value after registration. The highest performance improvement is reached for the smallest kernel $\sigma_g = 1.7$: displacing the large square using fine kernels is very costly,

a problem which the multiscale strategy alleviates by imposing first large transformations. The indentation area also benefits from the multiscale strategy, as indicated by values of $\Delta_J(\text{ROI})$ inferior to 1 pixel when $S_0 > 2$. For all parameters σ_g , the minimum of Δ_J is reached when optimization starts at the second coarsest scale $S_0 = 4$. The lower performance of $S_0 = 5$ indicates that very coarse deformations may not benefit the optimization. At $S_0 = 5$, the algorithm starts by optimizing only wavelet coefficients linked to functions whose support is $[0, 2^4]^d$ control points. With a spacing of 2 pixels between control points, the estimated vector fields are constant on areas of 32×32 pixels, which is very above the size of the large square. This effect is illustrated in Fig. 8: with $S_0 = 5$, the deformation grid shows a transformation affecting a larger part of the image domain than

necessary. We also notice that the final deformation fields estimated with $S_0 = 3$ are strongly influenced by the initial scale, again demonstrating the strong impact of S_0 on the rest of the initialization.

Interestingly, the multiscale strategy also seems to improve the smoothness of the final transformation when the initial scale is $S_0 = 4$ (Fig. 7, panel (b)). In Fig. 8, the deformation grid obtained with $S_0 = 1$ displays some irregularities, whereas the deformation grid and field of $S_0 = 4$ are distinctively smoother.

Finally, we observe in Fig. 7, panel (c) that the algorithm runtime increases linearly with S_0 : for $S_0 = 5$, it is 3 to 4 times higher than for the original LDDMM.

This experiment shows the importance of performing our coarse-to-fine procedure starting from a reasonably coarse scale. It can effectively enhance the outcome of the classical LDDMM when both large and fine scale transformations are needed.

5.1.2 Comparison to the multi-kernel algorithm

In the second part of this experiment, we compare the behaviour and performance of our multiscale optimization vs the multi-kernel algorithm [50]. Direct comparison between the two multiscale algorithms is not possible as 1) our implementation of the classical LDDMM differs from that of Risser et al. 2) experiments in Risser et al. were performed with a fixed (unknown) number of iterations, which is not possible in our setting since our multiscale strategy has an effect on the total number of iterations. It is however possible to compare how the multi-kernel LDDMM improves the classical LDDMM versus how our multiscale optimization improves the classical LDDMM. To this end, we perform registration between I_s and I_t using the same trade-off between regularity and data attachment (i.e. 0.5) and the same kernel widths as in the original experiment, i.e. $\sigma_g \in \{2, 3, 7\}$, plus the intermediary kernels $\sigma_g \in \{1.7, 2.5\}$. Contrary to Risser et al. we do not present results obtained with $\sigma_g > 7$ as the related kernels are too large for our multiscale algorithm to be of use; nor do we present results obtained with $\sigma_g = 1$, which has a low performance on this experiment with our LDDMM implementation. In Table 1, we report the performance of the algorithms for the different values of σ_g : classical LDDMM (Deformetrica version [15]), LDDMM combined with

Table 1 Performance of different algorithms on the toy experiment: original LDDMM (Deformetrica version), LDDMM combined with multiscale optimization (with $S_0 = 3$ for $\sigma_g = 7$ and $S_0 = 4$ for $\sigma_g < 7$), original LDDMM (Risser et al. version and results [50]) and multi-kernel LDDMM (Risser et al. results). The multi-kernel algorithm was tested with different combinations MK_n , with MK_n a sum of n weighted Gaussian kernels with standard deviations linearly sampled between 1 and 10. The reported values are the total residuals values between I_s and I_t in the image domain and a ROI covering the indentation and standard deviation of the Jacobian determinant.

		LDDMM (Deformetrica)			LDDMM + Multiscale		
σ_g	k_g	Δ_J	$\Delta_{J(\text{ROI})}$	SD(J)	Δ_J	$\Delta_{J(\text{ROI})}$	SD(J)
7	64	11.52	6.67	0.09	11.06	6.45	0.10
3	289	1.30	0.31	0.19	0.52	0.03	0.16
2.5	400	1.49	0.11	0.24	0.37	0.05	0.23
2	625	0.37	0.09	0.27	0.20	0.02	0.21
1.7	841	9.36	0.03	0.23	1.21	0.0	0.24

LDDMM (Risser et al.)				Multi-kernel LDDMM			
σ_g	Δ_J	$\Delta_{J(\text{ROI})}$	SD(J)	σ_g	Δ_J	$\Delta_{J(\text{ROI})}$	SD(J)
7	8.3	7.2	0.01	MK ₃	0.88	0.34	0.02
3	3.0	2.5	0.02	MK ₄	1.0	0.35	0.02
1	0.39	0.07	0.07	MK ₅	0.81	0.34	0.02
				MK ₆	0.94	0.37	0.02
				MK ₇	0.94	0.37	0.02

our multiscale algorithm (with $S_0 = S_{max} - 1$), classical LDDMM (Risser et al. version and results [50]), and multi-kernel LDDMM (Risser et al. results). As detailed in Section 2.2, flows in the multi-kernel framework are defined by a weighted sum of Gaussian kernels of different sizes: in Table 1, we report performance values obtained using MK_n , $n \in \{3, 4, 5, 6, 7\}$ [50], where MK_n denotes a sum of weighted Gaussian kernels with standard deviations linearly sampled between 1 and 10.

The classical LDDMM (Risser et al. version) reaches its best performance (in terms of residuals decrease) for the smallest kernel of size $\sigma_g = 1$. Compared to this baseline, the multi-kernel LDDMM has slightly lower performance whichever combination MK_n is used. If we compare our classical LDDMM (Deformetrica version) with our multiscale strategy, the latter yields significantly lower values of Δ_J and $\Delta_{J(\text{ROI})}$ for all values of σ_g . Three cases are discernible:

- When a large kernel is used, i.e. $\sigma_g = 7$, the original LDDMM has poor performance

and the improvement yielded by our coarse-to-fine strategy is only marginal as it is constrained, at core, by the scale σ_g .

- As observed in Section 5.1.1, when a fine kernel is used, i.e. $\sigma_g = 1.7$, the original LDDMM has poor performance. Its outcome is dramatically enhanced by the multiscale algorithm, indicating that it can successfully avoid unrealistic local minima.
- When an adequate kernel is used, e.g. $\sigma_g = 2$, the original LDDMM has good performance (with residual values similar to that of Risser et al. baseline), and the multiscale algorithm is still capable of improving its outcome marginally, reaching the best performance out of all the compared algorithms.

In terms of deformation smoothness, the multi-kernel LDDMM algorithm attains values of $SD(J)$ similar to that of Risser et al. original LDDMM for $\sigma_g = 3$, while reaching better performance. Compared to their baseline ($\sigma_g = 1$), the multi-kernel LDDMM decreases $SD(J)$ by a factor of 3. In contrast, our multiscale algorithm only marginally improves smoothness.

This experiment shows that our multiscale strategy can enhance the outcome of the classical LDDMM in a variety of situations, especially when using fine kernels, and can be competitive with another multiscale algorithm that builds true multiscale flows. In the following experiments, we will compare the performance of the classical, single-scale LDDMM algorithm with that of our multiscale algorithm on three different atlas estimation tasks. We will evaluate their ability to estimate high quality, stable template images as well as natural template-to-subjects deformations in different settings. The multiscale algorithm will be initialized with $S_0 = S_{max} - 1$.

5.2 Handwritten digits

In this section as well as the following experiments, we evaluate the ability of our multiscale scheme to enhance the outcome of the classical LDDMM on an atlas estimation task. We test the original LDDMM algorithm against the coarse-to-fine strategy using an experimental procedure similar to that of Durrleman et al. [15]: atlas estimation is performed using 20 randomly-chosen training images and the estimated template image is registered to 10 randomly chosen test images with the same parameters as those used during training. The

Table 2 Performance (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of handwritten digits. The relative residual errors and SSIM are computed between each train/test image and the deformed template. Bold style indicates statistically significant differences between the algorithms ($p < 0.05$).

		R		SSIM	
σ_g	k_g	LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
3	100	10.2 (3.1)	7.9 (1.9)	0.74 (0.10)	0.79 (0.05)
2	196	10.1 (1.7)	5.3 (0.5)	0.74 (0.05)	0.79 (0.01)
1.5	361	7.3 (1.7)	4.3 (1.4)	0.75 (0.05)	0.79 (0.05)
Registration					
3	100	10.0 (6.1)	7.5 (6.0)	0.74 (0.04)	0.79 (0.03)
2	196	8.5 (5.2)	6.5 (6.3)	0.72 (0.03)	0.78 (0.03)
1.5	361	9.0 (5.1)	5.3 (7.0)	0.71 (0.04)	0.78 (0.03)

experiment is repeated five times with different training and test sets, with no intersection between any of the training and test sets. This procedure is reproduced with different kernel widths σ_g . Since the five experiments are performed on independent datasets, we use paired Student t-tests to compare performance between the two algorithms.

The performance of the algorithms is assessed with the following metrics:

- **Relative residual error:** $R = \frac{\Delta_J}{\Delta_0}$ with Δ_0 and Δ_J and the mean residual value over subjects at iteration 0 and at convergence, respectively.
- **Structural Similarity Index Metric [63] (SSIM)** after training and test: $SSIM(I_1, I_2) = l(I_1, I_2) \times c(I_1, I_2) \times s(I_1, I_2)$ where I_1 and I_2 are the compared images, l is a function comparing the luminance (i.e. the mean pixel/voxel intensity) of the images, c compares the contrast (i.e. the standard deviation of the image intensities), and s quantifies the structural similarity between I_1 and I_2 (i.e. the correlation between luminance- and contrast-normalized intensities). SSIM values range between -1 (dissimilarity) and 1 (near-perfect similarity).
- **Standard deviation of the Jacobian determinant (SD(J))** of the displacement fields linked to the deformations $(\phi_i)_i$ to assess the regularity of the estimated transformations.



Fig. 9 Estimation of the template image by the original LDDMM and the LDDMM combined with multiscale optimization on the dataset of handwritten digits with different values of σ_g . For each experiment, five template images (estimated with non-intersecting training sets) and their standard deviation image are presented on the left, along with the template image from the first training set warped to the first five training images on the right. σ_g : width of the Gaussian kernel; k_g : corresponding number of control points.

• Runtimes

In this section, we use images of the digit 2 extracted from the well-known United States postal database of handwritten digits [64]. The size of the images is 28 by 28 pixels. Table 2 shows the mean relative residual error and similarity yielded by the algorithms over the experiments. Standard deviation of Jacobian values and runtime can be seen in Appendix Table B1. Figure 9 presents the five template images estimated by each algorithm with three different sets of parameters, along with the template image estimated from the first training set warped to five of the training images.

Table 2 shows that the coarse-to-fine algorithm reaches lower residual error and higher similarity than the original LDDMM algorithm during both atlas estimation and registration, with differences that reach significance for $\sigma_g = 2$ and $\sigma_g = 1.5$. Consistent with these results, we observe in Figure 9 that the original algorithm yields highly irregular template images,

a trend which worsens when the number of control points increases, indicating overfitting. The first template image warped towards the training images yields images that are close to the original ones for $\sigma_g = 3$. However, for lower values of σ_g , the reconstructed training images are less realistic. These observations belie the quantitative evaluation, which shows that the performance of the original algorithm increases with the number of parameters. This discrepancy demonstrates that residual error alone is not sufficient to evaluate the accuracy of the algorithms.

Unlike the original LDDMM algorithm, the coarse-to-fine procedure produces realistic template images that are stable across folds and whose quality is preserved when σ_g is decreased. Moreover, all reconstructed images are very close to the original ones. Their quality slightly increases with the number of control points: this is most evident for the third and fifth reconstructed subjects, which become more accurate for lower values of σ_g .

Expectedly, in Appendix Table B1, standard deviations of Jacobian values increase with the number of control points for both algorithms. The multiscale algorithm reaches slightly higher SD(J) values than the original algorithm, indicating more stretching of the deformed source images. In average, the multiscale algorithm has higher runtime, reflecting the time spent performing optimization at each scale (as shown in Section 5.1).

5.3 Artificial characters

In the previous experiment, one can remark that the performance of the algorithms diverge most when a high number of parameters is used. Therefore, one might simply be tempted to employ the original LDDMM algorithm with a lower number of parameters, as in Durrleman et al. [15] whose experiments were performed with 36 control points. However, datasets that present a higher amount of details and inter-subject variability may benefit from our coarse-to-fine strategy even when a lower number of control points is used. To confront our algorithm with a more difficult task, we manually designed a dataset of 30 characters. The size of the images is 28 by 28 pixels.

We compare our algorithm to the original LDDMM using cross-validation: the dataset is randomly split into a training set (24 images) and a test set (6 images). Each algorithm independently estimates a template image from

the training set, and then registers the template to each image in the test set with the same parameters as those used during training. This procedure is repeated five times, and reproduced with different values of σ_g . No statistical tests are performed because of the overlap between the training sets and between the test sets. Performance is assessed using the same metrics as in Section 5.2. Table 3 displays the mean relative residual error and SSIM after atlas estimation and registration and Figure 10 shows the five template images estimated by each algorithm with three different sets of parameters, along with five reconstructed training images.

Table 3 shows that the coarse-to-fine algorithm reaches lower residual error than the original LDDMM algorithm. The performance of the coarse-to-fine strategy increases with the number of control points during training and test, while the original LDDMM algorithm demonstrates the opposite trend. These results are supported by the qualitative evaluation of the template images. In Figure 10, for $\sigma_g = 5$, the two algorithms generate template images that present discrete but noticeable differences. With the original LDDMM version, the arms and legs of the characters appear slightly fuzzier, and the second template image is noisy. The reconstructed images yielded by the original algorithm are blurry (and even erroneous in case of the third subject), while the template and reconstructed images yielded by the coarse-to-fine algorithm seem sharp and accurate.

As in the previous experiment, the quality of the template images estimated by the original LDDMM algorithm deteriorates when the number of control points increases: with $\sigma_g = 4$ and $\sigma_g = 3$, images become fuzzier and display erroneous features inherited from the initial templates (mean intensity images), indicating high dependency on the initialization. The morphology of all but one reconstructed characters is also completely erroneous. In contrast, the multiscale algorithm is able to produce stable, sharp and correct template images for all parameters. Similar observations can be made regarding the transformation of the template image towards the five training images: the multiscale strategy succeeds in generating images that are nearly identical to the original ones.

These differences have a simple explanation: the original version simultaneously estimates the overall shape of the characters and details

Table 3 Performance (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of artificial characters. The relative residual errors and SSIM are computed between each train/test image and the deformed template.

		R		SSIM	
σ_g	k_g	LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
5	36	14.0 (4.5)	5.5 (0.8)	0.76 (0.04)	0.82 (0.01)
4	49	16.4 (3.8)	4.4 (0.4)	0.74 (0.03)	0.84 (0.00)
3	100	18.4 (4.5)	3.3 (0.5)	0.68 (0.03)	0.83 (0.00)
Registration					
5	36	9.8 (6.2)	6.0 (6.3)	0.75 (0.02)	0.78 (0.03)
4	49	10.7 (6.0)	4.6 (4.2)	0.73 (0.02)	0.82 (0.01)
3	100	13.4 (7.5)	3.6 (3.8)	0.64 (0.05)	0.80 (0.00)

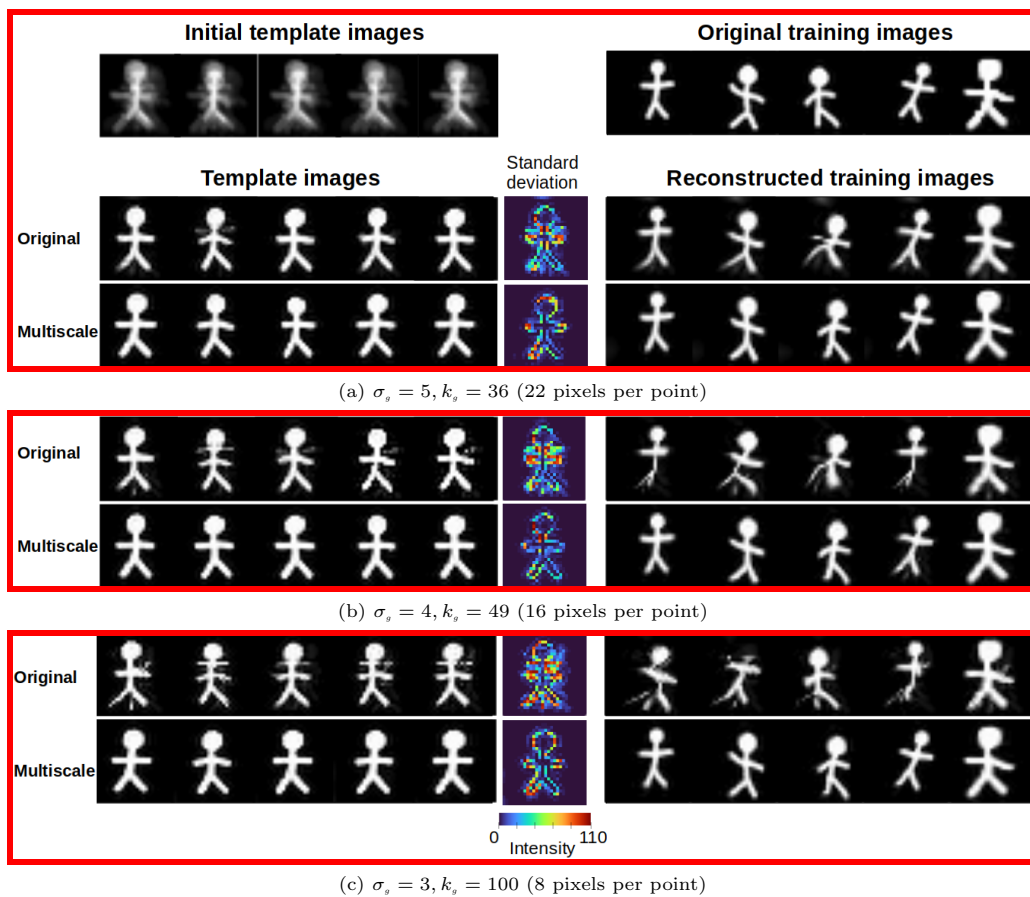


Fig. 10 Estimation of the template image by the original LDDMM and the LDDMM combined with multiscale optimization on the toy dataset with different parameters σ_g . For each experiment, five estimated template images (for each fold of cross-validation) and their standard deviation image are presented on the left, along with the template image estimated from the first training set warped to the first five training images on the right.

such as the location and orientation of the arms and legs, making it more dependent on the initial template image and leading to the selection of erroneous features, while the coarse-to-fine strategy first focuses on estimating the characters main features, which are then refined when the finer scales are optimized. This phenomenon is illustrated by movies showing template optimization across iterations, available at the first author’s webpage³.

In Appendix Table B2, it can be seen that the multiscale algorithm has slightly lower $SD(J)$ values during registration with $\sigma_g = 3$ compared to the original LDDMM algorithm. This suggests a better ability of our multiscale strategy to regularize deformations with a high number of parameters.

This experiment shows that on a dataset with high variability, the original LDDMM algorithm is unable to estimate templates that

are satisfying with respect to quantitative and qualitative criteria. The coarse-to-fine algorithm outperforms the former in both criteria, regardless of the number of parameters.

5.4 Fetal brain images

To evaluate the performance of our coarse-to-fine approach on a dataset of clinical images, we use 30 fetal brain MRIs with agenesis of the corpus callosum acquired in Hopital Trousseau, France [11]. Gestational ages are comprised between 32 and 34 weeks of gestation (mean = 32.9 ± 0.6).

Agenesis of the corpus callosum is a developmental anomaly characterized by the total or partial absence of the corpus callosum. It is often associated to anatomical features such as widening of the lateral ventricles. Atlas estimation can help better understand congenital anomalies by providing an insight into how these anatomical characteristics vary together [11]. However, as abnormal fetal brains may present a wide range of defects, this makes atlas

³https://fleurgaudfernau.github.io/Multiscale_atlas_estimation

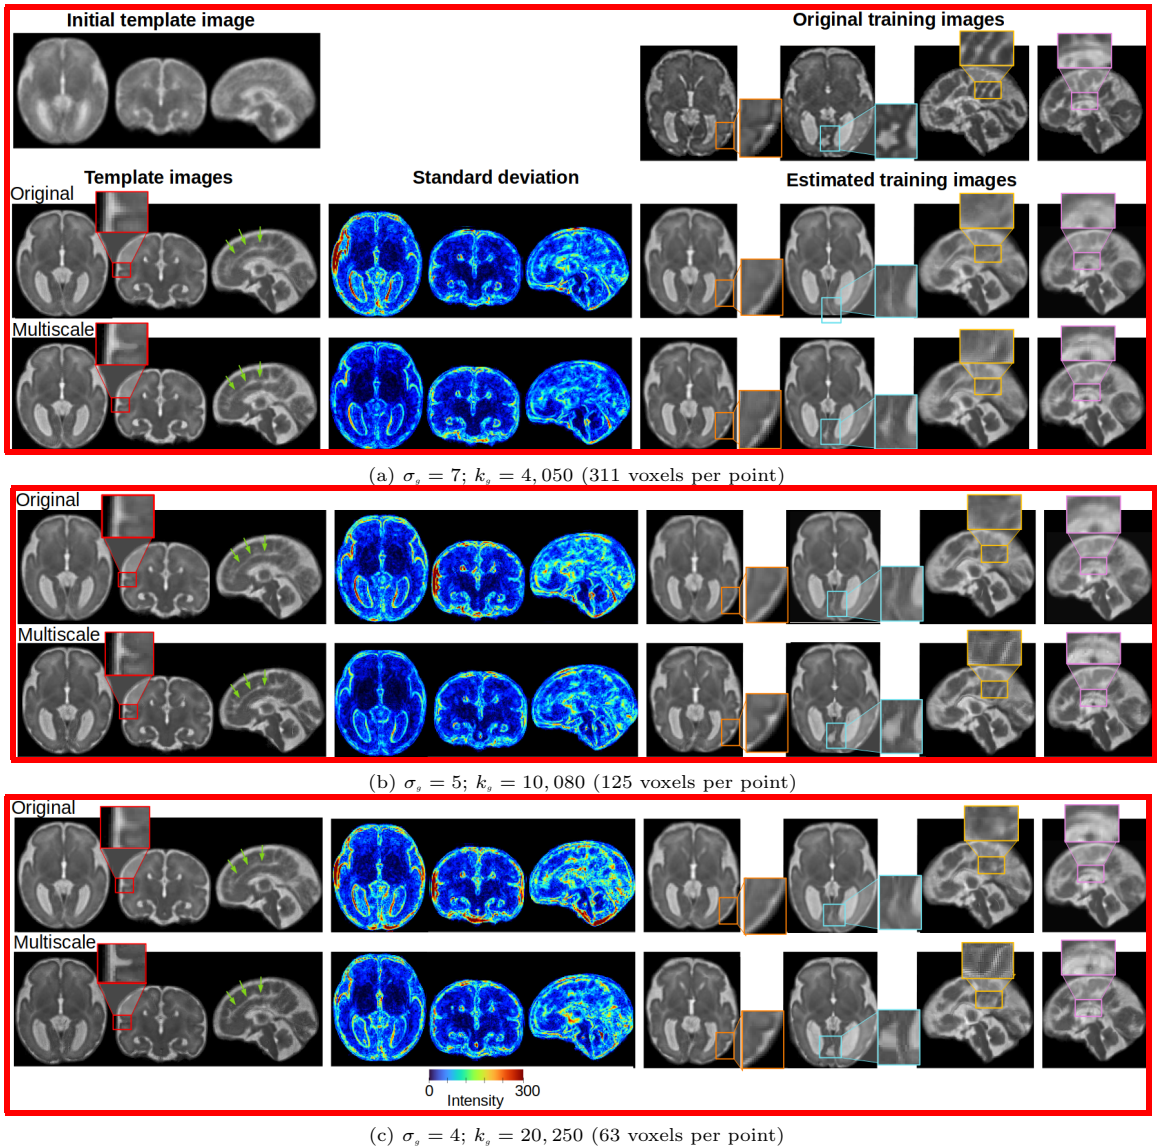


Fig. 11 Atlas estimation by the original LDDMM and the LDDMM combined with multiscale optimization on the dataset of fetal brain images. For each experiment, the estimated template volume from the first fold of cross-validation is presented in the left column. The first column displays the standard deviation of the template intensities over five folds of cross-validation. The first row displays four different training images in axial or sagittal view; the four rightmost columns display the corresponding reconstructed images (i.e. the template image warped to the training images). Salient differences between images are zoomed in with a factor 2. Red square: superior temporal sulcus. Green arrows: cingulate sulcus. Orange square: temporal cortex. Blue square: interhemispheric fissure. Yellow square: interhemispheric sulci. Pink square: pericallosal area.

estimation more difficult and prone to errors than with datasets of healthy fetuses. Thus, it is crucial to develop algorithms that are able to estimate realistic templates on both healthy and abnormal subjects.

The brain MRIs are preprocessed, volume reconstructed and rigidly aligned according to the procedure described in Gaudfernau et al. [11]. The final images have size 105x100x120 voxels. Cross-validation is performed in the same manner as in Section 5.3, with 24 images used for atlas estimation and 6 images used

for testing. Figure 11 presents an example of estimated template image during cross validation along with 4 reconstructed training images for different parameters σ_g , and Table 4 displays the mean relative residual error and SSIM after atlas estimation (training) and registration (test). Visual examination of the templates is performed by an expert radiologist.

Table 4 shows that the coarse-to-fine algorithm achieves lower residual error and higher SSIM values than the original algorithm during atlas estimation and registration for all

Table 4 Performance (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of fetal brains images. The relative residual error (R) and the SSIM are computed between each train/test image and the deformed template.

σ_g	k_g	R		SSIM	
		LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
7	4,050	43.6 (3.7)	37.8 (0.9)	0.73 (0.01)	0.74 (0.02)
5	10,080	38.5 (4.0)	32.5 (1.6)	0.74 (0.02)	0.77 (0.0)
4	20,250	34.5 (3.8)	31.9 (3.8)	0.76 (0.02)	0.78 (0.02)
Registration					
7	4,050	35.0 (6.2)	31.1 (5.7)	0.71 (0.01)	0.72 (0.0)
5	10,080	31.4 (5.0)	25.6 (5.2)	0.74 (0.01)	0.76 (0.01)
4	20,250	26.7 (1.4)	22.8 (1.1)	0.75 (0.01)	0.78 (0.01)

values of σ_g . In Appendix Table B3, both algorithms show similar values of $SD(J)$, though slightly higher for the multiscale algorithm, suggesting a better ability to capture some of the volume changes characterizing abnormal subjects, e.g. shrunk corpus callosum. The template images in Figure 11 present subtle but noticeable differences: the brain volumes estimated with the multiscale optimization display sharper features and enhanced contrast between structures, especially for higher values of σ_g . The template images estimated by the original LDDMM algorithm display more fuzzy areas, especially at the boundary between cortical gray matter and white matter (see for example the superior temporal sulcus, red squares) and the medial surface of the brain (e.g. the cingulate sulcus, green arrows). Compared to the original algorithm, the multiscale strategy yields template images that are more stable across folds, especially in regions with high inter-subjects variability, e.g. cortical folds, lateral ventricles and corpus callosum area.

The multiscale template-to-subject transformations also build more accurate anatomical structures: see for example the more pronounced gyration patterns (orange squares) and clearly delineated interhemispheric fissure and lateral ventricle (blue squares). Interestingly, the multiscale template image warped to the training subjects reveals more abnormal features associated to corpus callosum agenesis, such as the typical radiating sulci (yellow

squares) and missing corpus callosum (pink squares).

As in the previous experiments, the multiscale algorithm has higher runtimes than the original LDDMM algorithm in Appendix Table B3.

Altogether, these results indicate that our coarse-to-fine strategy can successfully enhance the results of atlas estimation applied to real-world, complex clinical data.

6 Discussion

In this paper, we took advantage of the hierarchical property of the wavelet decomposition to develop a coarse-to-fine optimization procedure in the LDDMM framework. Specifically, we proposed a Haar-like wavelet representation of the initial velocity fields to enhance the outcomes of the classical RKHS-based LDDMM algorithm. The transfer of information from coarse to fine scales ensures smarter initialization of the deformations at each level, leading the algorithm to favor more accurate solutions and avoid unrealistic local minima. Contrary to previous coarse-to-fine algorithms introduced in the LDDMM framework [19; 52; 54; 55], our approach adds no complexity to the mathematical model. Specifically, the reparameterization of the velocity fields can be seen as an additional layer of spatial regularization, which preserves the RKHS structure of the vector fields and the efficient numerical scheme used to compute the gradients. This reparameterization can easily be translated to other mathematical frameworks which model deformations using vector fields. For example, our multiscale strategy could be combined with models that express diffeomorphisms as a composition of deformations of increasingly fine scales [54; 55].

We first performed a registration experiment to assess the influence of the initial scale parameter and compare our algorithm to the multi-kernel strategy [50]. Results suggest that our multiscale algorithm can be competitive with the multi-kernel framework and lead to more accurate matching. While the focus of the multi-kernel framework is the design of an explicitly multiscale deformation model, our strategy’s primary goal is efficiency, in the spirit of classical coarse-to-fine strategies. A clear advantage of our multiscale optimization is its convenience: we only need to identify *one* relatively well-performing kernel and apply our multiscale scheme, while in the multi-kernel

framework, one needs to identify *several* relevant kernels, combine them, and optimize their weights, with no guarantee to converge to a better local minima [50]. However, our method does not improve deformation smoothness, which reflects the different bases upon which the two algorithms are built: our strategy is multiscale in the sense that it solves the optimization problem in spaces of increasing resolution; the multi-kernel framework is multiscale in the way deformations are defined as coexisting flows. Thus, it would be interesting to adapt our multiscale scheme to enhance optimization in the multi-kernel framework in order to estimate true multiscale flows and favor deformation smoothness.

We then evaluated the performance of our multiscale algorithm on three atlas estimation tasks of increasing difficulty. Compared to the original LDDMM algorithm, the coarse-to-fine algorithm yields higher quality templates with better stability (estimated by the variability of the template intensities across cross-validation folds), that are able to generalize to unseen images. Not only does our strategy produce images that have a realistic anatomy, but it leads to enhanced preservation of anatomical details, including unusual or abnormal ones. This makes it particularly appropriate for tasks involving high inter-subject variability, specifically clinical images. Our results suggest that the multiscale LDDMM algorithm can estimate a more diverse range of transformations all the while preserving reasonable smoothness.

Some limitations of this algorithm have to be highlighted. Unlike approaches based on a mixture of kernels, our deformation model relies on a single Gaussian kernel. While this provides the advantage of introducing no additional parameters, the results of our algorithm depend on the choice of the kernel width σ_g - but to a lesser extent than the original version. As shown in Section 5.2, the multiscale LDDMM algorithm fails to avoid irregular deformations when using a very small kernel. In such scenarios, a solution could reside in local adaptation, i.e. maintain smoothness constraints in areas where the matching is near-perfect and constrain the velocity fields to be unevenly smooth, in line with strategies based on simultaneously coexisting flows [50]. As illustrated in Section 5.1, the runtime of our multiscale optimization increases with the number of scales of the coarse-to-fine procedure. This can be prohibitive when working with a

high number of parameters, which increases the number of scales, or, as shown in Section 5.4, when performing atlas estimation on clinical images. Methodological improvements could be made to our multiscale algorithm. The Haar wavelet, which produces sharp transitions between nearby areas of the vector fields, could be replaced with smoother wavelet functions such as the Daubechies wavelet to favor more regular vector fields.

The properties of the wavelet transform offer interesting avenues to explore. Notably, the wavelet functions $\phi_{s,k}$ and $\psi_{s,k}^o$ are normalized depending on their scale s . By modifying the weights attributed to the fine and coarse scales in this normalization, one can change the relative importance attributed to high and low frequency coefficients during optimization. It would be interesting to refine our wavelet-based spatial regularizer in this way and observe how it impacts the results.

The simplicity of our algorithm makes it easy to implement, which opens up interesting perspectives. The coarse-to-fine strategy could be applied to other types of atlases such as spatio-temporal ones [12], or to other statistical frameworks such as the Bayesian Mixed Effect Model [9]. In regard to the latter point, a significant advance would be to integrate our reparameterization of the velocity fields into this Bayesian framework by introducing priors on the wavelet coefficients, in the spirit of Downie et al. [61] who decomposed deformations into a Haar wavelet basis and modelled the coefficients as independent random variables with a mixture distribution.

In addition, we will focus on developing a dual coarse-to-fine strategy, by applying a hierarchical representation to the images, as already attempted in other mathematical models, e.g. with B-spline deformations [30; 65; 66] and in the hyperelasticity framework [49]. Alternating both coarse-to-fine strategies would very likely provide template images of even higher quality. Ultimately, we will also apply our multiscale strategies to other models available in the LDDMM framework, namely geodesic regression [8] and its variant piecewise geodesic regression [67] and evaluate their usefulness on complex clinical challenges such as the modelling of the fetal brain growth in a continuous manner [68], a well-knowingly challenging task in Computational Anatomy.

Declarations

Acknowledgments. This work was partly funded by the third author’s chair in the PRAIRIE institute funded by the French national agency ANR as part of the “Investissements d’avenir” programme under the reference ANR-19- P3IA-0001.

Conflict of interest. The authors have no conflicts of interest to declare that are relevant to the content of this article.

Ethics approval. This work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of human subjects. The experiments conducted on the dataset of fetal brain MRI has been granted ethics approval by the Institutional Review Board of the Comité d’éthique de la recherche en imagerie médicale (CERIM) under the reference CRM-2112-215.

Data availability The dataset of handwritten digits is publicly available at: <https://www.kaggle.com/datasets/bistaumanga/usps-dataset>. The dataset of artificial characters is available at the first author’s webpage: https://fleurgaudferneau.github.io/Multiscale_atlas_estimation/. The dataset of fetal brain MRI is not available due to ethical considerations.

Author’s contributions. Conceptualization: F.G.; Methodology: F. G., E.L.P.; Validation: E.B.; Writing - original draft preparation: F.G.; Writing - review and editing: E.L.P., E.B., S.A.; Funding acquisition: S.A.; Resources: E.B.; Supervision: E.L.P., S.A.

Appendix A Algorithms FWT and IWT

To compute the Forward and Inverse Wavelet Transforms, explicit computation of the related matrices \mathbf{M}_{FWT} and \mathbf{M}_{IWT} is not necessary since the algorithms rely on local operations on basis coefficients. To compute the wavelet coefficients of an array X of dimension d , the FWT algorithm (Algorithm 5) relies on a 1D Haar Forward algorithm (Algorithm 2), which computes local means and differences along one axis. Uneven numbers of rows/columns are handled by computing *weighted* averages and differences so that the boundaries are given the same importance as the rest of the array. Fig. 1 illustrates how the FWT algorithm operates on a 2-dimensional array, corresponding to lines 4-12 in Algorithm 5.

Finally, the output wavelet coefficients have to be normalized to preserve the input signal energy: this is done by computing the matrix \mathbf{M}_{FWT} in Algorithm 4 and computing the renormalization matrix \mathbf{R} from \mathbf{M}_{FWT} :

$$\mathbf{R}[i] = \frac{1}{\|\mathbf{M}_{\text{FWT}}[i, :]\|_2}$$

where $\|\mathbf{M}_{\text{FWT}}[i, :]\|_2$ is the L_2 norm of the i_{th} row of \mathbf{M}_{FWT} .

The IWT algorithm (Algorithm 6) runs in a manner that is symmetrical to the FWT algorithm, by using the 1D Haar Backward algorithm (Algorithm 3) to compute finer-scale coefficients from coarse scale coefficients one axis at a time.

Algorithm 2 1D Haar Forward step

```

1: Input
2:  $\beta$ : array of shape  $(k_1, \dots, k_D)$ ,  $d$ : axis along which to compute the transform,  $K_d$ : size of the original array along axis  $d$ ,  $s$ : current scale,  $w_s$ : list of scales for each axis  $[1, \dots, D]$  in ascending order
3: Initialization
4: Swap axes in  $\beta$  to put axis  $d$  in position 0
5: if  $k_d > 1$  then
6:   if  $k_d$  is even then
7:      $\beta_a \leftarrow (\beta[0 :: 2] + \beta[1 :: 2])/2$  {Average the consecutive rows of  $\beta$ }
8:     if  $K_d \neq 2 \times k_d^s$  then
9:        $\delta \leftarrow K_d/2^s - (k_d - 1)$  {Weighting of the border}
10:       $\beta_a[-1] \leftarrow (\beta[-2] + \delta * \beta[-1])/(1 + \delta)$ 
11:     end if
12:      $\beta_d \leftarrow \beta[0 :: 2] - \beta_a$  {Difference between the consecutive rows of  $\beta$ }
13:   else
14:      $\beta_a \leftarrow (\beta[0 : -1 : 2] + \beta[1 :: 2])/2$  {Average the consecutive rows of  $\beta$ }
15:      $\beta_d \leftarrow \beta[0 : -1 : 2] - \beta_a$  {Difference between the consecutive rows of  $\beta$ }
16:      $\beta_a \leftarrow \text{concatenate}([\beta_a, \beta[-1]])$  {Add the last unpaired row}
17:   end if
18:    $\beta \leftarrow \text{concatenate}([\beta_a, \beta_d])$ 
19: end if
20:  $w_s[d] \leftarrow \lceil k_d/2 \rceil + w_s[d]$ 
21: return  $\beta, w_s$ 

```

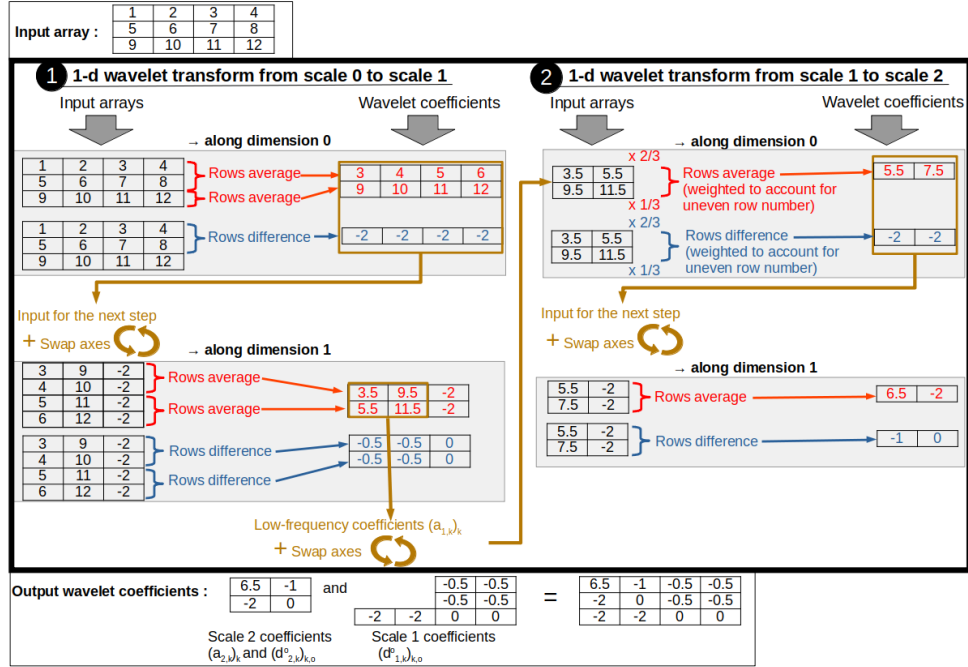


Fig. 1 The FWT Algorithm (Algorithm 5) applied to a 3-by-4 array. Steps in grey squares illustrate the 1D Haar Forward step (Algorithm 2). Running the IWT Algorithm (Algorithm 6) amounts to running the illustrated steps backward. For the sake of simplicity, the renormalization step (line 16 in Algorithm 5) is not featured.

Algorithm 3 1D Haar Backward step

- 1: **Input**
- 2: β : array of wavelet coefficients of shape (k_1, \dots, k_D) , d : axis of β along which to apply the backward transform, K_d : size of the original array along axis d , s : current scale, w_s : list of scales for each axis $[1, \dots, D]$ in ascending order
- 3: **Initialization**
- 4: $\beta \leftarrow$ Swap axes in β to put axis d in position 0
- 5: $X \leftarrow$ zero array of the shape of β
- 6: $n_{low} \leftarrow w_s[d][0]$
- 7: $\beta_a \leftarrow \beta[:n_{low}]$ {Low-frequency coefficients}
- 8: $\beta_d \leftarrow \beta[n_{low} :]$ {High-frequency coefficients}
- 9: **if** $w_s[d][1] > 1$ **then**
- 10: **if** $w_s[d][1]$ is even **then**
- 11: $X[0 :: 2] \leftarrow \beta_a + \beta_d$
- 12: $X[1 :: 2] \leftarrow 2 \times \beta_a - X[0 :: 2]$
- 13: **if** $K_d \neq 2 \times k_d^s$ **then**
- 14: $\delta \leftarrow K_d/2^s - (k_d - 1)$
- 15: $X[-1] \leftarrow ((1 + \delta) \times \beta_a[-1] - X[-2])/\delta$
- 16: **end if**
- 17: **else**
- 18: $X[0 : -1 : 2] \leftarrow \beta_a[: -1] + \beta_d$
- 19: $X[-1] \leftarrow \beta_a[-1]$
- 20: $X[1 :: 2] \leftarrow 2 \times \beta_a[: -1] - X[0 : -1 : 2]$
- 21: **end if**
- 22: **end if**
- 23: $w_s[d] \leftarrow w_s[d][1 :]$
- 24: **return** X, w_s

Algorithm 4 Renormalization

- 1: {Compute the Haar Forward Matrix M_{FWT} and the renormalization matrix R }
- 2: **Input**
- 3: (K_1, \dots, K_D) : shape of the array to wavelet transform
- 4: **Initialization**
- 5: $n \leftarrow \prod_{i=1}^D K_i$
- 6: $M_{FWT} \leftarrow$ zero array of shape (n, n)
- 7: **for** i in range(n) **do**
- 8: $Z \leftarrow$ zero array of shape (K_1, \dots, K_D)
- 9: i^{th} element of $Z \leftarrow 1$
- 10: $\beta, w_s, R \leftarrow$ **Haar_Forward**($Z, \rho = 0$) {Algorithm 5}
- 11: $M_{FWT}[:, i] \leftarrow$ flatten(β)
- 12: **end for**
- 13: $R \leftarrow$ zero array of shape (K_1, \dots, K_D)
- 14: **for** i in range(n) **do**
- 15: i^{th} element of $R \leftarrow \frac{1}{\|M_{FWT}[:, i]\|_2}$
- 16: **end for**
- 17: **return** R

Algorithm 5 Forward Haar Wavelet Transform (FWT algorithm)

```

1: {Compute the wavelet coefficients  $\beta$  of an array  $X$ }
2: Input
3:  $X$ : array of shape  $(K_1, \dots, K_D)$ ,  $\rho$ : renormalization factor (default: 1)
4: Initialization
5:  $\beta \leftarrow X$  {Array to store the wavelet coefficients}
6:  $w_s \leftarrow \lceil [K_1], \dots, [K_D] \rceil$  {For each axis, store the scales for which the wavelet coefficients have been computed in ascending order}
7:  $S_{max} \leftarrow \lceil \log_2(\max(K_1, \dots, K_D)) \rceil$ 
8: for  $s$  in range( $S_{max}$ ) do
9:    $\beta_{current} \leftarrow \beta[: w_s[0][0], \dots, : w_s[D][0], ]$ 
   {Low-frequency coefficients at scale  $s - 1$ }
10:  for  $d$  in range( $D$ ) do
11:     $\beta_{current}, w_s, R \leftarrow$ 
    Haar_forward_1d_step( $\beta_{current}, d, K_d,$ 
     $s, w_s$ ) {Algorithm 2}
12:     $\beta[: w_s[0][0], \dots, : w_s[D][0], ] \leftarrow \beta_{current}$ 
13:  end for
14: end for
15: if  $\rho \neq 0$  then
16:    $R \leftarrow$  renormalize( $(K_1, \dots, K_D)$ )
   {Algorithm 4}
17:    $\beta \leftarrow \beta \times R^{\rho}$ 
18: end if
19: return  $\beta, w_s, R$ 

```

Algorithm 6 Inverse Haar Wavelet Transform (IWT algorithm)

```

1: {Given its wavelet coefficients  $\beta$ , compute array  $X$  (i.e. coefficients of scale 0)}
2: Input
3:  $\beta$ : array of wavelet coefficients of shape  $(K_1, \dots, K_D)$ ,  $\rho$ : renormalization factor (default: 1),  $w_s$ : list of scales for each axis  $[1, \dots, D]$ , stored in ascending order
4: Initialization
5:  $X \leftarrow \beta$ 
6:  $S_{max} \leftarrow \lceil \log_2(\max(K_1, \dots, K_D)) \rceil$ 
7: if  $\rho \neq 0$  then
8:    $R \leftarrow$  renormalize( $(K_1, \dots, K_D)$ )
   {Algorithm 4}
9:    $X \leftarrow X/R^{\rho}$ 
10: end if
11: for  $s$  in range( $S_{max} - 1, -1, -1$ ) do
12:    $X_{current} \leftarrow X[: w_s[0][1], \dots, : w_s[D][1]]$ 
13:   for  $d$  in range( $D$ ) do
14:      $X_{current} \leftarrow$ 
     Haar_Backward_1d_step( $X_{current}, d,$ 
      $K_d, s, w_s$ ) {Algorithm 3}
15:      $X[: w_s[0][1], \dots, : w_s[D][1]] \leftarrow X_{current}$ 
16:   end for
17: end for
18: return  $X$ 

```

Appendix B Algorithms smoothness and runtime

Table B1 Standard deviation of the Jacobian determinant (SD(J)) and runtimes (RT) (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of handwritten digits. Bold style indicate statistically significant differences between the algorithms ($p < 0.05$).

σ_g	k_g	SD(J)		RT (min.)	
		LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
3	100	0.15 (0.02)	0.19 (0.01)	8.50 (3.26)	13.22 (2.88)
2	196	0.18 (0.01)	0.21 (0.01)	8.61 (1.29)	12.29 (1.13)
1.5	361	0.20 (0.01)	0.22 (0.01)	7.0 (1.44)	12.34 (1.12)
Registration					
3	100	0.20 (0.01)	0.20 (0.01)	0.31 (0.05)	0.52 (0.03)
2	196	0.22 (0.01)	0.24 (0.02)	0.38 (0.05)	0.56 (0.02)
1.5	361	0.23 (0.02)	0.26 (0.03)	0.32 (0.03)	0.58 (0.08)

Table B2 Standard deviation of the Jacobian determinant (SD(J)) and runtimes (RT) (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of artificial characters.

		SD(J)		RT (min.)	
σ_g	k_g	LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
5	36	0.10 (0.01)	0.11 (0.01)	9.51 (3.04)	13.65 (2.52)
4	49	0.12 (0.01)	0.13 (0.01)	9.43 (2.19)	13.78 (2.62)
3	100	0.13 (0.01)	0.12 (0.01)	11.10 (3.77)	13.21 (1.77)
Registration					
5	36	0.13 (0.02)	0.12 (0.01)	0.37 (0.04)	0.56 (0.05)
4	49	0.14 (0.01)	0.14 (0.01)	0.38 (0.03)	0.51 (0.05)
3	100	0.17 (0.02)	0.14 (0.01)	0.39 (0.03)	0.47 (0.02)

Table B3 Standard deviation of the Jacobian determinant (SD(J)) and runtimes (RT) (mean and standard deviation) of the classical LDDMM and the LDDMM combined with multiscale optimization over five folds of cross-validation after atlas estimation and registration on the dataset of fetal brains images.

		SD(J)		RT (min.)	
σ_g	k_g	LDDMM	LDDMM+ Multiscale	LDDMM	LDDMM+ Multiscale
Atlas estimation					
7	4,050	0.02 (0.01)	0.03 (0.01)	134.40 (20.67)	189.97 (67.39)
5	10,080	0.02 (0.01)	0.03 (0.01)	95.80 (34.61)	191.37 (89.82)
4	20,250	0.03 (0.01)	0.05 (0.01)	111.38 (37.24)	189.74 (77.0)
Registration					
7	4,050	0.03 (0.01)	0.03 (0.01)	4.12 (0.53)	6.66 (0.62)
5	10,080	0.03 (0.01)	0.03 (0.01)	4.27 (0.48)	7.22 (0.42)
4	20,250	0.04 (0.01)	0.05 (0.01)	3.68 (0.70)	7.84 (0.14)

References

- [1] Thompson, D.W.: On Growth and Form. Canto. Cambridge University Press, ??? (1992). <https://doi.org/10.1017/CBO9781107325852>
- [2] Miller, M.I., Arguillère, S., Tward, D.J., Younes, L.: Computational anatomy and diffeomorphometry: A dynamical systems model of neuroanatomy in the soft condensed matter continuum. *WIREs Systems Biology and Medicine* **10**(6), 1425 (2018) <https://doi.org/10.1002/wsbm.1425>
- [3] Oliveira, F., Tavares, J.: Medical image registration: A review. *Computer methods in biomechanics and biomedical engineering* **17**, 73–93 (2014) <https://doi.org/10.1080/10255842.2012.670855>
- [4] Trouvé, A.: Diffeomorphisms groups and pattern matching in image analysis. *Int. J. Comput. Vision* **28**(3), 213–221 (1998) <https://doi.org/10.1023/A:1008001603737>
- [5] Miller, M., Trouvé, A., Younes, L.: On the metrics and euler-lagrange equations of computational anatomy. *Annual review of biomedical engineering* **4**, 375–405 (2002) <https://doi.org/10.1146/annurev.bioeng.4.092101.125733>
- [6] Christensen, G., Rabbitt, R., Miller, M.: Deformable template using large deformation kinematics. *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society* **5**, 1435–47 (1996) <https://doi.org/10.1109/83.536892>
- [7] Vaillant, M., Miller, M., Younes, L., Trouvé, A.: Statistics on diffeomorphisms via tangent space representations. *NeuroImage* **23 Suppl 1**, 161–9 (2004) <https://doi.org/10.1016/j.neuroimage.2004.07.023>
- [8] Fletcher, T.: Geodesic regression on riemannian manifolds. In: Pennec, X., Joshi, S., Nielsen, M. (eds.) *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy: Geometrical and Statistical Methods for Modelling Biological Shape Variability*, Toronto, Canada, pp. 75–86 (2011)
- [9] Allasonnière, S., Durrleman, S., Kuhn, E.: Bayesian mixed effect atlas estimation with a diffeomorphic deformation model. *SIAM Journal on Imaging Sciences* **8**, 1367–1395 (2015) <https://doi.org/10.1137/140971762>
- [10] Grenander, U., Miller, M.: Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics* **LVI**, 617–694 (1998) <https://doi.org/10.1090/qam/1668732>
- [11] Gaudfernau, F., Blondiaux, E., Allasonnière, S.: Analysis of the Anatomical Variability of Fetal Brains with Corpus Callosum Agenesis. In: *MICCAI 2021 - Perinatal, Preterm and Paediatric Image Analysis Workshop*. LNCS - Lecture Notes in Computer Science, vol. 12959. Springer, Strasbourg, France (2021). https://doi.org/10.1007/978-3-030-87735-4_26
- [12] Debavelaere, V., Durrleman, S., Allasonnière, S.: Learning the clustering of longitudinal shape data sets into a mixture of independent or branching trajectories. *International Journal of Computer Vision* **128** (2020) <https://doi.org/10.1007/s11263-020-01337-8>
- [13] Miller, M., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. *Journal of mathematical imaging and vision* **24**, 209–228 (2006) <https://doi.org/10.1007/s10851-005-3624-0>
- [14] Younes, L.: Jacobi fields in groups of diffeomorphisms and applications. *Quarterly of Applied Mathematics* **65** (2007) <https://doi.org/10.1090/S0033-569X-07-01027-5>
- [15] Durrleman, S., Allasonnière, S., Joshi, S.: Sparse adaptive parameterization of variability in image ensembles. *Int. Journal of Computer Vision* **101**, 1–23 (2012) <https://doi.org/10.1007/s11263-012-0556-1>
- [16] Durrleman, S., Prastawa, M., Gerig, G., Joshi, S.: Optimal data-driven sparse parameterization of diffeomorphisms for population analysis, vol. 22, pp. 123–34 (2011). https://doi.org/10.1007/978-3-642-22092-0_11
- [17] Sommer, S., Nielsen, M., Darkner, S., Pennec, X.: Higher-order momentum distributions and locally affine lddmm registration. *SIAM Journal on Imaging Sciences* **6**(1), 341–367 (2013) <https://doi.org/10.1137/110859002>
- [18] Zhang, M., Fletcher, P.: Finite-dimensional lie algebras for fast diffeomorphic image registration, vol. 24 (2015). https://doi.org/10.1007/978-3-319-19992-4_19
- [19] Sommer, S., Lauze, F., Nielsen, M., Pennec, X.: Sparse multi-scale diffeomorphic registration: The kernel bundle framework. *J. of Mathematical Imaging and Vision* **46** (2012) <https://doi.org/10.1007/>

- s10851-012-0409-0
- [20] Modersitzki, J.: FAIR: Flexible Algorithms for Image Registration. Society for Industrial and Applied Mathematics, ??? (2008). <https://doi.org/10.1137/1.9780898718843>
- [21] Witkin, A.P.: Scale-space filtering. In: Fischler, M.A., Firschein, O. (eds.) Readings in Computer Vision, pp. 329–332. Morgan Kaufmann, San Francisco (CA) (1987). <https://doi.org/10.1016/B978-0-08-051581-6.50036-2>
- [22] Koenderink, J.J.: The structure of images. *Biological Cybernetics* **50**, 363–370 (2004)
- [23] Lucas, B., Kanade, T.: An iterative image registration technique with an application to stereo vision (ijcai), vol. 81 (1981)
- [24] Anandan, P.: A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision* **2**, 283–310 (1989) <https://doi.org/10.1007/BF00158167>
- [25] Enkelmann, W.: Investigations of multigrid algorithms for the estimation of optical flow fields in image sequences. *Computer Vision, Graphics, and Image Processing* **43**(2), 150–177 (1988) [https://doi.org/10.1016/0734-189X\(88\)90059-X](https://doi.org/10.1016/0734-189X(88)90059-X)
- [26] Terzopoulos, D.: Image analysis using multigrid relaxation methods. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **8**, 129–139 (1986) <https://doi.org/10.1109/TPAMI.1986.4767767>
- [27] Unser, M., Aldroubi, A., Gerfen, C.: Multiresolution image registration procedure using spline pyramids. *Proc SPIE* **2034** (1993) <https://doi.org/10.1117/12.162061>
- [28] Mallat, S.: Theory for multi-resolution signal decomposition - the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence - PAMI* **11** (1989)
- [29] Thirion, J.-P.: Image matching as a diffusion process: An analogy with maxwell’s demons. *Medical Image Analysis* **2** (1998)
- [30] Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Transactions on Medical Imaging* **18**(8), 712–721 (1999) <https://doi.org/10.1109/42.796284>
- [31] Schnabel, J., Rueckert, D., Quist, M., Blackall, J., Castellano-Smith, A., Hartkens, T., Penney, G., Hall, W., Liu, H., Truwit, C., Gerritsen, F., Hill, D., Hawkes, D.: A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations, vol. 1, pp. 573–581 (2001). https://doi.org/10.1007/3-540-45468-3_69
- [32] Hellier, P., Micheli-Tzanakou, E.: Inter-Subject Non-Rigid Registration: An Overview with Classification and the Romeo Algorithm, pp. 273–337 (2007). https://doi.org/10.1007/0-306-48608-3_8
- [33] Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95–113 (2007) <https://doi.org/10.1016/j.neuroimage.2007.07.007>
- [34] Avants, B.B., Epstein, C., Grossman, M., Gee, J.: Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* **12**, 26–41 (2008) <https://doi.org/10.1016/j.media.2007.06.004>
- [35] Tustison, N., Avants, B.: Explicit b-spline regularization in diffeomorphic image registration. *Frontiers in Neuroinformatics* **7** (2013) <https://doi.org/10.3389/fninf.2013.00039>
- [36] Ashburner, J., Friston, K.J.: Nonlinear spatial normalization using basis functions. *Human Brain Mapping* **7**(4), 254–266 (1999) [https://doi.org/10.1002/\(SICI\)1097-0193\(1999\)7:4<254::AID-HBM4>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1097-0193(1999)7:4<254::AID-HBM4>3.0.CO;2-G)
- [37] Zhang, M., Liao, R., Dalca, A.V., Turk, E.A., Luo, J., Grant, P.E., Golland, P.: Frequency diffeomorphisms for efficient image registration. In: Niethammer, M., Styner, M., Aylward, S., Zhu, H., Oguz, I., Yap, P.-T., Shen, D. (eds.) *Information Processing in Medical Imaging*, pp. 559–570. Springer, Cham (2017)
- [38] Mang, A., Biros, G.: An inexact newton-krylov algorithm for constrained diffeomorphic image registration. *SIAM Journal on Imaging Sciences* **8**, 1030–1069 (2015) <https://doi.org/10.1137/140984002>
- [39] Zhang, M., Fletcher, P.: Fast diffeomorphic image registration via fourier-approximated lie algebras. *International Journal of Computer Vision* **127** (2019) <https://doi.org/10.1007/s11263-018-1099-x>
- [40] Christensen, G.E., Johnson, H.J.: Consistent image registration. *IEEE Transactions on Medical Imaging* **20**(7), 568–582 (2001) <https://doi.org/10.1109/42.932742>

- [41] Amit, Y.: A nonlinear variational problem for image matching. *Siam Journal on Scientific Computing* **15** (1994) <https://doi.org/10.1137/0915014>
- [42] Wu, Y.-T., Kanade, T., Li, C., Cohn, J.: Image registration using wavelet-based motion model. *Int. Journal of Computer Vision* **38**, 129–152 (2000) <https://doi.org/10.1023/A:1008101718719>
- [43] Cathier, P.: Iconic feature registration with sparse wavelet coefficients, vol. 9, pp. 694–701 (2006). https://doi.org/10.1007/11866763_85
- [44] Sun, W., Niessen, W.J., Klein, S.: Wavelet based free-form deformations for nonrigid registration. In: Ourselin, S., Styner, M.A. (eds.) *Medical Imaging 2014: Image Processing*, vol. 9034, pp. 934–939. SPIE, ??? (2014). <https://doi.org/10.1117/12.2043846>. International Society for Optics and Photonics
- [45] Cai, W., Wang, J.: Adaptive multiresolution collocation methods for initial-boundary value problems of nonlinear pde s. *SIAM Journal on Numerical Analysis* **33**, 937–970 (1996) <https://doi.org/10.1137/0733047>
- [46] Gefen, S., Tretiak, O., Bertrand, L., Rosen, G., Nissanov, J.: Surface alignment of an elastic body using a multiresolution wavelet representation. *IEEE transactions on bio-medical engineering* **51**, 1230–41 (2004) <https://doi.org/10.1109/TBME.2004.827258>
- [47] Musse, O., Heitz, F., Armspach, J.: Topology preserving deformable image matching using constrained hierarchical parametric models. *Image Processing, IEEE Transactions on* **10**, 1081–1093 (2001) <https://doi.org/10.1109/83.931102>
- [48] Noblet, V., Heinrich, C., Heitz, F., Armspach, J.: 3-d deformable image registration: A topology preservation scheme based on hierarchical deformation models and interval analysis optimization. *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society* **14**, 553–66 (2005) <https://doi.org/10.1109/TIP.2005.846026>
- [49] Debroux, N., Le Guyader, C., Vese, L.A.: Multiscale Registration. In: *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 115–127. Springer International Publishing, ??? (2021). https://doi.org/10.1007/978-3-030-75549-2_10
- [50] Risser, L., Vialard, F.-X., Wolz, R., Deprez, M., Holm, D., Rueckert, D.: Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping. *IEEE transactions on medical imaging* **30**, 1746–59 (2011) <https://doi.org/10.1109/TMI.2011.2146787>
- [51] Bruveris, M., Risser, L., Vialard, F.-X.: Mixture of kernels and iterated semidirect product of diffeomorphisms groups. *SIAM Journal on Multiscale Modeling and Simulation* **10** (2011) <https://doi.org/10.1137/110846324>
- [52] Gris, B., Durrleman, S., Trouvé, A.: A sub-riemannian modular framework for diffeomorphism-based analysis of shape ensembles. *SIAM Journal on Imaging Sciences* **11** (2016) <https://doi.org/10.1137/16M1076733>
- [53] Tan, M., Qiu, A.: Large deformation multiresolution diffeomorphic metric mapping for multiresolution cortical surfaces: A coarse-to-fine approach. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* **25** (2016) <https://doi.org/10.1109/TIP.2016.2574982>
- [54] Modin, K., Nachman, A., Rondi, L.: A multiscale theory for image registration and nonlinear inverse problems. *Advances in Mathematics* **346** (2018) <https://doi.org/10.1016/j.aim.2019.02.014>
- [55] Miller, M.I., Tward, D.J., Trouvé, A.: Coarse-to-fine hamiltonian dynamics of hierarchical flows in computational anatomy. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 3760–3765 (2020). <https://doi.org/10.1109/CVPRW50498.2020.00438>
- [56] Niethammer, M., Kwitt, R., Vialard, F.-X.: Metric learning for image registration. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8455–8464 (2019) <https://doi.org/10.1109/CVPR.2019.00866>
- [57] Shen, Z., Vialard, F.-X., Niethammer, M.: Region-specific diffeomorphic metric mapping. In: Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., ??? (2019). <https://proceedings>.

- neurips.cc/paper_files/paper/2019/file/291597a100aadd814d197af4f4bab3a7-Paper.pdf
- [58] Berlinet, A., Thomas-Agnan, C.: Reproducing Kernel Hilbert Space in Probability and Statistics, pp. 1–54. Springer, ??? (2004). <https://doi.org/10.1007/978-1-4419-9096-9>
- [59] Bône, A., et al.: Deformetrica Version 4.3.0 (2018). <https://gitlab.com/icm-institute/aramislab/deformetrica/-/wikis/home>
- [60] Mallat, S.: A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way, 3rd edn. Academic Press, Inc., USA (2008)
- [61] Downie, T., Silverman, B.: A wavelet mixture approach to the estimation of image deformation functions (2000) <https://doi.org/10.2307/25053170>
- [62] Leow, A., *et al.*: Statistical properties of jacobian maps and the realization of unbiased large-deformation nonlinear image registration. *IEEE transactions on medical imaging* **26**, 822–32 (2007) <https://doi.org/10.1109/TMI.2007.892646>
- [63] Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: From error visibility to structural similarity. *Image Processing, IEEE Transactions on* **13**, 600–612 (2004) <https://doi.org/10.1109/TIP.2003.819861>
- [64] Hastie, T., Tibshirani, R., Friedman, J., Franklin, J.: The elements of statistical learning: Data mining, inference, and prediction. *Math. Intell.* **27**, 83–85 (2004) <https://doi.org/10.1007/BF02985802>
- [65] Loeckx, D., Slagmolen, P., Maes, F., Vandermeulen, D., Suetens, P.: Non-rigid image registration using conditional mutual information. *Information processing in Medical Imaging Conference* **20**, 725–37 (2007) <https://doi.org/10.1109/TMI.2009.2021843>
- [66] Staring, M., Heide, U., Klein, S., Viergever, M., Pluim, J.: Registration of cervical mri using multifeature mutual information. *IEEE transactions on medical imaging* **28**, 1412–21 (2009) <https://doi.org/10.1109/TMI.2009.2016560>
- [67] Chevallier, J., Oudard, S., Allasonnière, S.: Learning spatiotemporal piecewise-geodesic trajectories from longitudinal manifold-valued data. In: 31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, United States (2017)
- [68] Licandro, R., Langs, G., Kasprian, G., Sablatnig, R., Prayer, D., Schwartz, E.: A longitudinal diffeomorphic atlas-based tissue labeling framework for fetal brains using geodesic regression. (2016)