

Integrated Construction of Multimodal Atlases with Structural Connectomes in the Space of Riemannian Metrics

Kristen M. Campbell

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, USA
School of Computing, University of Utah, Salt Lake City, UT, USA

kris@sci.utah.edu

Haocheng Dai

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, USA
School of Computing, University of Utah, Salt Lake City, UT, USA

hdai@sci.utah.edu

Zhe Su

Department of Neurology, University of California Los Angeles, Los Angeles, CA, USA

zsu20@g.ucla.edu

Martin Bauer

Department of Mathematics, Florida State University, Tallahassee, FL, USA

bauer@math.fsu.edu

P. Thomas Fletcher

Department of Electrical & Computer Engineering, University of Virginia, Charlottesville, VA, USA
Department of Computer Science, University of Virginia, Charlottesville, VA, USA

ptf8v@virginia.edu

Sarang C. Joshi

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, USA
Department of Bioengineering, University of Utah, Salt Lake City, UT, USA

sjoshi@sci.utah.edu

Abstract

The structural network of the brain, or structural connectome, can be represented by fiber bundles generated by a variety of tractography methods. While such methods give qualitative insights into brain structure, there is controversy over whether they can provide quantitative information, especially at the population level. In order to enable population-level statistical analysis of the structural connectome, we propose representing a connectome as a Riemannian metric, which is a point on an infinite-dimensional manifold. We equip this manifold with the Ebin metric, a natural metric structure for this space, to get a Riemannian manifold along with its associated geometric properties. We then use this Riemannian framework to apply object-oriented statistical analysis to define an atlas as the Fréchet mean of a population of Riemannian metrics. This formulation ties into the existing framework for diffeomorphic construction of image atlases, allowing us to construct a multimodal atlas by simultaneously integrating complementary white matter structure details from DWMRI and cortical details from T1-weighted MRI. We illustrate our framework with 2D data examples of connectome registration and atlas formation. Finally, we build an example 3D multimodal atlas using T1 images and connectomes derived from diffusion tensors estimated from a subset of subjects from the Human Connectome Project.

1. Introduction

Tractography is one way to represent a structural connectome, or structural network of a brain, which consists of brain regions that are physically connected by a network of neuronal bundles that make up the white matter of that brain. It is not currently possible to image individual neurons in a living brain non-invasively. Instead, we can use diffusion-weighted magnetic resonance images (DWMRI) to infer the pathways of coherent bundles

of neurons that cross through other bundles in white matter and connect with end points in gray matter. There are a variety of tractography algorithms used to infer these white matter pathways, ranging from local methods that integrate local orientation information to form individual streamlines to global methods that estimate all fiber tracts simultaneously. Deterministic (Basser et al., 2000) and probabilistic (Behrens et al., 2003) streamline integration methods are easy to compute, but they suffer from accumulation of local orientation errors leading to tract reconstructions biased toward shorter and straighter tracts. Conversely, global estimation methods (Jbabdi et al., 2007; Christiaens et al., 2015) incorporate prior anatomical knowledge while ensuring that tractography is consistent with the underlying data. However, these global methods have convergence issues, sensitivity to initialization and priors, and a tendency to have estimated tracts end in the middle of white matter regions. Biases in tract reconstruction introduced by either local or global methods affect the accuracy of quantitative measures such as track density and connection strength.

Geodesic tractography algorithms, first introduced by O'Donnell et al. (2002), use a combination of local and global information to determine tracts by formulating white matter pathways as geodesic curves under a Riemannian metric derived from the DWMRI data. In the original work, O'Donnell et al. (2002) use the inverse diffusion tensor as the Riemannian metric. As described by Lenglet et al. (2004), the inverse diffusion tensor metric has a connection to Brownian motion through the Laplace Beltrami operator on the resulting Riemannian manifold. The inverse tensor metric favors paths that follow the principal eigenvector of the diffusion tensor, as this is the locally optimal direction to move. However, geodesic curves for this metric do not consistently follow the principal eigenvector of the diffusion tensor, and tend to be overly straight in regions where the white matter fibers are bending, losing association with the underlying anatomy in these areas. This problem has been addressed by several strategies to improve the adherence of geodesics to the white matter geometry, including "sharpening" the inverse diffusion tensor (Fletcher et al., 2007), using the adjugate of the diffusion tensor (Fuster et al., 2016), and using a conformal metric (Hao et al., 2014). One advantage to the geodesic tractography formulation is that we can do uncertainty and confidence interval analysis of the tractographies following (Sengers et al., 2021).

These tractography techniques help to give insight into the structure of a single brain, but it remains an open challenge to quantitatively measure how these structural connectome pathways vary in a population. To do such a population analysis, we first need a common frame of reference, or atlas space, where spatial differences between subjects' white matter can be measured.

Initially, white matter atlases were constructed by aligning DWMRI to an anatomical template, then transforming and averaging the associated tensor field or distribution function field. For example, Mori et al. (2008) construct a diffusion tensor imaging (DTI) atlas by registering the DWMRI of multiple subjects to a standardized anatomical template. They build the DTI atlas by transforming the diffusion tensors for each subject (Alexander et al., 2001) and then taking the Euclidean average of the transformed diffusion tensors at each voxel. This approach does not use the white matter directionality information encoded in the diffusion images during the registration. It also suffers from the fact that the Euclidean average of diffusion tensors does not take into account the directionality and tends to be fatter (i.e., less anisotropic) than the input tensors (Fletcher and Joshi,

2007). Another approach by Yeh et al. (2018) is to register q -space diffusion images into an anatomical template and estimate the spin distribution function (SDF) at each voxel in the template. Then the SDFs are averaged on a per-voxel basis. While this method does take into account the directionality of the white matter in a local neighborhood, it does not take into account consistency of long-range white matter connections. These methods both rely on low-accuracy registration to the anatomical template image, because at the time high-accuracy diffeomorphic registrations could not be used as they did not preserve the continuity of pathway directions across voxels (Yeh et al., 2018).

The field then focused on constructing atlases directly from tractography. These methods are adept at finding well-known tracts, but they generally require significant expert curation to remove the many false-positive fibers and tracts they create (Zhang et al., 2018). Additionally, variation of population-level tractography characteristics like fiber density of tracts is as likely to be reflective of the tractography process chosen as it is of the underlying physical white matter. These and other biases in tractography quantification are well-characterized by Jeurissen et al. (2019).

In order to get a more complete anatomical atlas, Toga et al. (2006) argue for the integrated derivation of multimodal atlases using techniques such as spatial normalization to produce more comprehensive atlases. Some attempts have been made to create a multimodal population atlas from T1 and DWMRI images. Gupta et al. (2016) rigidly register the DWMRI into the same space as the T1 template image and then apply the same deformations used to create the T1 template to the DWMRI before estimating the diffusion tensors from the transformed DWMRI. This approach puts the DWMRI and T1 images in the same space, but does not use the white matter structure to inform the creation of the T1 atlas. Other previous work on multichannel registration of diffusion and T1 MRI include Avants et al. (2007), who use a Euclidean image match metric on the diffusion tensors, and Uus et al. (2020), who use local angular correlation on the orientation distribution functions (ODFs). Like the white matter-only atlases approaches discussed above, these methods only take local information into consideration and thus do not preserve the consistency of long-range white matter connections. They do demonstrate that registration quality is improved over single channel registration when complementary channels are combined in the objective function.

We are motivated to find an approach that can preserve the best aspects of these atlas and tractography methods, while mitigating their weaknesses. Specifically, we want to create a white matter pathway atlas that preserves local orientations and other anatomical information while maintaining the continuity and integrity of long-range connections. We then want a method to bring subjects into that atlas space to enable statistical quantification of both structural connectivity and geometric variability of white matter structure across a population.

To meet these goals, we describe the metric matching framework presented by Campbell et al. (2021) in more detail and then extend it by combining diffeomorphic metric matching with diffeomorphic image matching to enable the construction of both a multimodal white and gray matter atlas simultaneously for the first time. This formulation preserves geodesics transformed by diffeomorphisms, which meets our objective to use local diffusion data and maintain the integrity of long-range connectomics as inferred by tractography (Cheng et al.,

2015). We demonstrate that long-range connections are preserved by performing geodesic tractography on the resulting atlas.

1.1 Contributions of the Article

In this paper, we contribute a mathematical framework for diffeomorphic metric matching that is compatible with existing image matching frameworks to enable the creation of integrated multimodal atlases. We start by using the concept from geodesic tractography to represent connectome fibers as geodesics of a metric, that is, each brain's white matter structure is represented as a point on the infinite-dimensional manifold of Riemannian metrics. This manifold is then equipped with the diffeomorphism-invariant Ebin metric to compute distances and geodesics between connectomes. We explain how this Riemannian manifold is the foundation for the algorithm for diffeomorphic metric registration of structural connectomes and the statistical groupwise metric atlas estimation algorithm. We then extend this model by including an image matching term in those algorithms. Finally, we simultaneously estimate an integrated multimodal white matter pathway and T1 MRI-based image atlas for the first time. This article is an extended version of the Information Processing in Medical Imaging (IPMI) conference paper (Campbell et al., 2021), where we expand the results of the conference proceedings in several major directions. Most importantly, the joint white matter pathway and T1 MRI-based image atlas model is newly introduced and, in contrast to Campbell et al. (2021) which only contained 2D examples, we now present 3D atlas construction examples.

2. Structural Connectomes as Riemannian Metrics

In the white matter of the brain, the diffusion of water is restricted perpendicular to the direction of the axons. Diffusion-weighted MRI measures the microscopic diffusion of water in multiple directions at every voxel in a 3D volume. Thus, the directionality of white matter in the brain can be locally inferred. Traditionally, global connections of the white matter have been estimated by a procedure called *tractography*, which numerically computes integral curves of the vector field formed by the most likely direction of fiber tracts at each point. DTI models anisotropic water diffusion with a 3×3 symmetric positive definite tensor, $D(x)$, at each voxel, $x \in M$, where the manifold $M \subset \mathbb{R}^3$ is the image domain. The principal eigenvector of $D(x)$ is aligned with the direction of the strongest diffusion.

Riemannian metrics that represent connectomics of a subject have been developed in diffusion imaging by O'Donnell et al. (2002) and include the inverse-tensor metric $\tilde{g} = D(x)^{-1}$. However, the geodesics associated with the inverse-tensor metric tend to deviate from the principal eigenvector directions and take straighter paths through areas of high curvature.

In this work we build on the algorithm developed by Hao et al. (2014), which estimates a spatially-varying function, $\alpha(x) : M \rightarrow \mathbb{R}$, that modulates the inverse-tensor metric to create a locally-adaptive Riemannian metric, $g_\alpha = e^{\alpha(x)} \tilde{g}$. We briefly describe the method here for completeness but refer the reader to Hao et al. (2014) for details. This adaptive *connectome metric*, g_α , is conformally equivalent to the inverse-tensor metric and is better at capturing the global connectomics, particularly through regions of high curvature. Figure 1

shows how well the geodesics of each metric match the integral curve of the vector field. The connectome metric geodesics are very closely aligned with the integral curves.

The geodesic between two end-points, $p, q \in M$, associated with the inverse-tensor metric, $\tilde{g}(x) = D(x)^{-1}$, minimizes the energy functional, \tilde{E} , while the geodesic associated with the connectome metric, $g_\alpha(x) = e^{\alpha(x)}D(x)^{-1}$, minimizes the energy functional, E_α .

$$\tilde{E}(\gamma) = \int_0^1 \langle T(t), T(t) \rangle_{\tilde{g}} dt, \quad E_\alpha(\gamma) = \int_0^1 e^{\alpha(x)} \langle T(t), T(t) \rangle_{\tilde{g}} dt, \quad (1)$$

where $\gamma : [0, 1] \rightarrow M$, $\gamma(0) = p$, $\gamma(1) = q$, $T = \frac{d\gamma}{dt}$.

Analyzing the variation of E_α leads to the geodesic equation, $\text{grad } \alpha = 2\nabla_T T$, where the Riemannian gradient of α , $\text{grad } \alpha = \tilde{g}^{-1}(\frac{\partial \alpha}{\partial x^1}, \frac{\partial \alpha}{\partial x^2}, \dots, \frac{\partial \alpha}{\partial x^n})$, and $\nabla_T T$ is the covariant derivative of T along its integral curve.

To enforce the desired condition where the tangent vectors, T , of the geodesic match the vector field, V , of the unit principal eigenvectors of $D(x)$, we minimize the functional, $F(\alpha) = \int_M \|\text{grad } \alpha - 2\nabla_V V\|_{\tilde{g}}^2 dx$. The equation for α that minimizes $F(\alpha)$ is

$$\Delta_{\tilde{g}} \alpha = 2 \text{div}_{\tilde{g}}(\nabla_V V), \quad (2)$$

where div and Δ are the Riemannian divergence and Laplace-Beltrami operator. We discretize the Poisson equation in Equation (2) using a second-order finite difference scheme that satisfies both the Neumann boundary conditions $\frac{\partial \alpha}{\partial \vec{n}} = \langle \text{grad } \alpha, \vec{n} \rangle = \langle 2\nabla_V V, \vec{n} \rangle$ and the governing equation on the boundary. We then solve for α .

Note that we can use this method to match the geodesics of the connectome metric to other vector fields defining the tractogram, e.g., from higher-order diffusion models that can represent multiple fiber crossings in a voxel. In particular, for tractography based on fiber orientation distributions (FODs), we can use the techniques presented in Nie and Shi (2019) to generate the vector field V .




Figure 1: A geodesic of the inverse-tensor metric (blue) and adaptive metric (orange), along with an integral curve (black) associated with the principal eigenvectors for a synthetic tensor field (left) and a subject's connectome metric from the Human Connectome Project (center). Right shows a detailed view of the metric in the corpus callosum.

The fundamental result of the purposed method is the fact that a metric tensor field is estimated such that a given vector field is a geodesic vector field of the estimated Riemannian

metric. It is in this sense that the estimated metric field captures the geometry of the white matter as inferred by tractography. This metric field that captures the tractography will be the primary driving force for the registration algorithms. We will now develop registration and atlas construction algorithms that leverage the rich geometric structure of Manifold of all Riemannian Metrics.

3. The Geometry of the Manifold of all Riemannian Metrics

Once we have estimated a Riemannian metric for a human connectome, it is a point in the infinite-dimensional manifold of all Riemannian metrics, $\text{Met}(M)$, where M is the domain of the image. We will equip this space with a diffeomorphism-invariant Riemannian metric, called the Ebin or DeWitt metric (Ebin, 1970; DeWitt, 1967). The invariance of the infinite-dimensional metric under the group of diffeomorphisms $\text{Diff}(M)$ is a crucial property, as it guarantees the independence of an initial choice of coordinate system on the brain manifold. As we will base our statistical framework on this infinite-dimensional geometric structure, we will now give a detailed overview of its induced geometry on $\text{Met}(M)$.

Let M be a smooth n -dimensional manifold; for our targeted applications n will be two or three. We denote by $\text{Met}(M)$ the space of all smooth Riemannian metrics on M , i.e., each element g of the space $\text{Met}(M)$ is a symmetric, positive-definite $\binom{0}{2}$ tensor field on M . It is convenient to think of the elements of M as being point-wise positive-definite sections of the bundle of symmetric two-tensors $S^2 T^* M$, i.e., smooth maps from M with values in $S^2 T^* M$. Thus, the space $\text{Met}(M)$ is an open subset of the linear space $\Gamma(S^2 T^* M)$ of all smooth symmetric $\binom{0}{2}$ tensor fields and hence itself a smooth Fréchet-manifold (Ebin, 1970). Furthermore, let $\text{Diff}(M)$ denote the infinite-dimensional Lie group of all smooth diffeomorphisms of the manifold M . Elements of $\text{Diff}(M)$ act as coordinate changes on the manifold M . This group acts on the space of metrics via pullback

$$\text{Met}(M) \times \text{Diff}(M) \rightarrow \text{Met}(M), \quad (g, \varphi) \mapsto \varphi^* g = g(T\varphi \cdot, T\varphi \cdot). \quad (3)$$

In an analogous way we can define the pushforward action

$$\text{Met}(M) \times \text{Diff}(M) \rightarrow \text{Met}(M), \quad (g, \varphi) \mapsto \varphi_* g = (\varphi^{-1})^* g \quad (4)$$

It is important to note that the geometries of the metrics g and $\varphi^* g$ ($\varphi_* g$ resp.) are also related via φ . In particular, geodesics with respect to g are mapped via φ to geodesics with respect to $\varphi^* g$ (via φ^{-1} for $\varphi_* g$, resp.). On the infinite-dimensional manifold $\text{Met}(M)$, there exists a natural Riemannian metric: the reparameterization-invariant L^2 -metric. To define the metric, we need to first characterize the tangent space of the manifold of all metrics: $\text{Met}(M)$ is an open subset of $\Gamma(S^2 T^* M)$. Thus, every tangent vector h is a smooth bilinear form $h : TM \times_M TM \rightarrow \mathbb{R}$ that can be equivalently interpreted as a map $TM \rightarrow T^* M$. The L^2 -metric is given by

$$G_g^E(h, k) = \int_M \text{Tr}(g^{-1} h g^{-1} k) \text{vol}(g), \quad (5)$$

with $g \in \text{Met}(M)$, $h, k \in T_g \text{Met}(M)$ and $\text{vol}(g)$ the induced volume density of the metric g . This metric, introduced in Ebin (1970), is also known as the Ebin metric. We call

the metric *natural* as it requires no additional background structure and is consequently invariant under the pushforward and pullback actions of the diffeomorphism group, i.e.,

$$G_g(h, k) = G_{\varphi^*g}(\varphi^*h, \varphi^*k) = G_{\varphi_*g}(\varphi_*h, \varphi_*k) \quad (6)$$

for all $\varphi \in \text{Diff}(M)$, $g \in \text{Met}(M)$ and $h, k \in T_g \text{Met}(M)$. Note that the invariance of the metric follows directly from the substitution formula for multi-dimensional integrals.

The Ebin metric induces a particularly simple geometry on the space $\text{Met}(M)$, with explicit formulas for geodesics, geodesic distance and curvature. In the following theorem and corollary, we will present the most important of these formulas, which will be of importance for our proposed metric matching framework.

First we note that a metric $g \in \text{Met}(M)$, in local coordinates, can be represented as a field of symmetric, positive-definite $n \times n$ matrices that vary smoothly over M . Similarly, each tangent vector at g can be represented as a field of symmetric $n \times n$ matrices. By the results of Freed et al. (1989); Gil-Medrano and Michor (1991); Clarke (2013b), one can reduce the investigations of the space of all Riemannian metrics to the study of the geometry of the finite-dimensional space of symmetric, positive-definite $n \times n$ matrices. The pointwise nature of the Ebin metric allows one to solve the geodesic initial and boundary value problem on $\text{Met}(M)$ for each $x \in M$ separately. Consequently the formulas for geodesics, geodesic distance and curvature on the finite-dimensional matrix space can be translated directly to results for the Ebin metric on the infinite-dimensional space of Riemannian metrics.

Note that the space of Riemannian metrics, $\text{Met}(M)$ with the Ebin metric, is not metrically complete and not geodesically convex. Thus the minimal geodesic between two Riemannian metrics may not exist in $\text{Met}(M)$, but only in a larger space; the metric completion $\overline{\text{Met}}(M)$, which consists of all possibly degenerate Riemannian metrics. This construction has been worked out in detail by Clarke (2013a) – including the existence of minimizing paths in $\overline{\text{Met}}(M)$. In the following proof, we will omit these details and refer the interested reader to the article Clarke (2013a) for a more in-depth discussion. In Theorem 1, we present an explicit formula for the minimizing geodesic in $\overline{\text{Met}}(M)$ that connects two given Riemannian metrics.

Theorem 1 (Minimizing geodesics) *For $g_0, g_1 \in \text{Met}(M)$ we define*

$$k(x) = \log(g_0^{-1}(x)g_1(x)), \quad k_0(x) = k(x) - \frac{\text{Tr}(k(x))}{n} \text{Id} \quad (7)$$

$$a(x) = \sqrt[4]{\det(g_0(x))}, \quad b(x) = \sqrt[4]{\det(g_1(x))}, \quad \kappa(x) = \frac{\sqrt{n \text{Tr}(k_0(x)^2)}}{4} \quad (8)$$

$$q(t, x) = 1 + t \left(\frac{b(x) \cos(\kappa(x)) - a(x)}{a(x)} \right), \quad r(t, x) = \frac{tb(x) \sin(\kappa(x))}{a(x)}. \quad (9)$$

Then the minimal path $g(t, x)$ with respect to the Ebin metric in $\overline{\text{Met}}(M)$ that connects g_0 to g_1 is given by

$$g = \begin{cases} (q^2 + r^2)^{\frac{2}{n}} g_0 \exp\left(\frac{\arctan(r/q)}{\kappa} k_0\right) & 0 < \kappa < \pi, \\ q^{\frac{4}{n}} g_0 & \kappa = 0, \\ (1 - \frac{a+b}{a} t)^{\frac{4}{n}} g_0 \mathbb{1}_{[0, \frac{a}{a+b}]} + \left(\frac{a+b}{b} t - \frac{a}{b}\right)^{\frac{4}{n}} g_1 \mathbb{1}_{[\frac{a}{a+b}, 1]} & \kappa \geq \pi, \end{cases} \quad (10)$$




Figure 2: An example of an interpolating geodesic between two metric tensors on the grid with respect to the Ebin metric (5), where the left and the right ellipse fields represent the boundary metrics. One can observe the behavior of the Ebin metric by following the deformations of the ellipses representing the metric tensor.

where $\mathbb{1}$ denotes the indicator function in the variable t . We suppressed the functions' dependence on t and x for better readability.

Proof This theorem is essentially a reformulation of the minimal geodesic formula given in (Clarke, 2013b, Theorem 4.16). In fact, noting that the Ebin metric (5) is point-wise, we can restrict ourselves to each $x \in M$. By (Clarke, 2013b, Theorem 4.5), we know that in the case of $0 \leq \kappa(x) < \pi$, $g_1(x)$ is in the image of the Riemannian exponential map starting at $g_0(x)$, and the Riemannian exponential is a diffeomorphism between $U := S^2 T_x^* M \setminus (-\infty, -4/n] g_0(x)$ and its image. Here $S^2 T_x^* M$ denotes the vector space of all symmetric (0,2) tensors at $x \in M$. Using the formula of the inverse exponential map in (Clarke, 2013b, Theorem 4.5) we calculate the preimage of $g_1(x)$,

$$\text{Exp}_{g_0}^{-1} g_1|_x = \begin{cases} \frac{4}{n} \left(\frac{b(x)}{a(x)} \cos \kappa(x) - 1 \right) g_0(x) + \frac{b(x) \sin \kappa(x)}{\kappa(x) a(x)} g_0(x) k_0(x) & 0 < \kappa(x) < \pi, \\ \frac{4}{n} \left(\frac{b(x)}{a(x)} - 1 \right) g_0(x) & \kappa(x) = 0. \end{cases} \quad (11)$$

The geodesic formula in the case of $0 \leq \kappa(x) < \pi$ then follows immediately from (Clarke, 2013b, Theorem 4.4). For the case of $\kappa(x) \geq \pi$, by (Clarke, 2013b, Theorem 4.14) the minimal geodesic between $g_0(x)$ and $g_1(x)$ is given by the concatenation of the straight segments from $g_0(x)$ to the zero tensor and from the zero tensor to $g_1(x)$, which gives the last statement and finally proves the result. \blacksquare

An example of calculating a geodesic in the space $\overline{\text{Met}}(M)$ using the explicit formula is visualized in Figure 2. The ellipse in the fifth row and fourth column, and the one in the first column and row of this figure are examples of geodesic paths that include possibly degenerate Riemannian metrics from the completion $\overline{\text{Met}}(M)$. Note that handling these degenerate cases are well understood and we did not observe any numerical issues associated with such cases.

Next, we recall that the geodesic distance of a Riemannian metric is defined as the infimum of all paths connecting two given points,

$$\text{dist}_{\text{Met}}(g_0, g_1) = \inf \int_0^1 \sqrt{G_g(\partial_t g, \partial_t g)} dt, \quad (12)$$

where the infimum is taken over all paths $g : [0, 1] \rightarrow \text{Met}(M)$ with $g(0) = g_0$ and $g(1) = g_1$. As a direct consequence of Theorem 1, we obtain the explicit formula for the distance function given in Corollary 2.

Corollary 2 (Geodesic distance) *Let $g_0, g_1 \in \text{Met}(M)$ and let k, k_0, a, b and κ be as in Theorem 1. Let $\theta(x) = \min \{\pi, \kappa(x)\}$. Then the squared geodesic distance of the Ebin metric is given by*

$$\text{dist}_{\text{Met}}^2(g_0, g_1) = \frac{16}{n} \int_M (a(x)^2 - 2a(x)b(x) \cos(\theta(x)) + b(x)^2) dx. \quad (13)$$

Proof Using Theorem 1, we obtain the formula of the minimal geodesic that connects g_0 and g_1 . By calculating the time derivative $\partial_t g(t)$, the final statement follows from the definition of the geodesic distance (12) by a direct computation. \blacksquare

Having equipped the space of Riemannian metrics with the distance function (13), we can consider the Fréchet mean, \hat{g} , of a collection of metrics, g_1, \dots, g_N , which is defined as a minimizer of the sum of squared distances.

$$\hat{g} = \operatorname{argmin}_g \sum_{i=1}^N \text{dist}_{\text{Met}}^2(g, g_i). \quad (14)$$

One could directly minimize this functional using a gradient-based optimization procedure. As our distance function is the geodesic distance function of a Riemannian metric and since we have access to an explicit formula for the minimizing geodesics, we will instead use the iterative geodesic marching algorithm, see e.g. Ho et al. (2013), to approximate the Fréchet mean: Given N Riemannian metrics g_i , we approximate the Fréchet mean via $\hat{g} = \hat{g}_N$, where \hat{g}_i is recursively defined as $\hat{g}_0 = g_0$, $\hat{g}_i(x) = g(1/(i+1), x)$ and where $g(t, x)$ is the minimal path, as given in Theorem 1, connecting \hat{g}_{i-1} to the i -th data point g_i . Thus one only has to calculate N geodesics *in total* in the space of Riemannian metrics, whereas a gradient-based algorithm would require one to calculate N geodesic distances *in each step* of the gradient descent.

3.1 The Induced Distance Function on the Diffeomorphism Group

We can use the geodesic distance function of the Ebin metric to induce a right-invariant distance function on the group of diffeomorphisms. As we will be using this distance function as a regularization term in our matching functional, we will briefly describe this construction here. We fix a Riemannian metric $g \in \text{Met}(M)$ and define the “distance” of a diffeomorphism φ to the identity via

$$\text{dist}_{\text{Diff}}^2(\text{id}, \varphi) = \text{dist}_{\text{Met}}^2(g, \varphi^* g) = \text{dist}_{\text{Met}}^2(g, \varphi_* g), \quad (15)$$

where the last equality is due to the invariance of the Ebin metric. To be more precise, this distance can be degenerate on the full diffeomorphism group since the isometries of the Riemannian metric g form the kernel of $\text{dist}_{\text{Diff}}$. For our purposes we will consider the Euclidean metric for the definition of $\text{dist}_{\text{Diff}}$. Thus the only elements in the kernel are

translations and rotations. The right invariance of $\text{dist}_{\text{Diff}}$ follows directly from the $\text{Diff}(M)$ -invariance of the Ebin metric. We note, however, that $\text{dist}_{\text{Diff}}$ is not directly associated with a Riemannian structure on the diffeomorphism group: the orbits of the diffeomorphism group in the space of metrics are not totally geodesic and thus $\text{dist}_{\text{Diff}}$ is not the geodesic distance of the pullback of the Ebin metric to the space of diffeomorphisms. See also Khesin et al. (2013) where this construction has been studied in more detail.

4. Computational Anatomy of the Human Connectome

Fundamental to the precise characterization and comparison of the human connectome of an individual subject or a population as a whole is the ability to map or register two different human connectomes. The framework of Large Deformation Diffeomorphic Metric Mapping (LDDMM) is well developed for registering points (Joshi and Miller, 2000), curves (Glaunès et al., 2008) and surfaces (Vaillant and Glaunès, 2005) all modeled as submanifolds of \mathbb{R}^3 as well as images modeled as an L^2 function (Beg et al., 2005). This framework has also been extended to densities (Bauer et al., 2015) modeled as volume forms. We now extend the diffeomorphic mapping framework to the connectome modeled as Riemannian metrics. The diffeomorphism group deforms the domain, M , and acts naturally on the space of metrics, $\text{Met}(M)$, see Equation (4). With this action and a reparameterization-invariant metric, the problem of registering two connectomes fits naturally into the framework of computational anatomy.

Now the registration of two connectomes is achieved by minimizing the energy

$$E(\varphi) = \text{dist}_{\text{Diff}}^2(\text{id}, \varphi) + \lambda_1 \text{dist}_{\text{Met}}^2(g_0, \varphi_* g_1) \quad (16)$$

over all such diffeomorphisms in $\text{Diff}(M)$. Here $\text{dist}_{\text{Diff}}$ is a right invariant distance on Diff and dist_{Met} is a reparameterization-invariant distance on the space of all Riemannian metrics, e.g., the geodesic distance of the metrics studied above. The first term measures the deformation cost and the second term is a similarity measure between the target and the deformed source connectome. The invariance of the two distances is essential for the minimization problem to be independent of the choice of coordinate system on the brain manifold.

We use the distance function as introduced in Section 3.1 to measure the deformation cost, i.e., $\text{dist}_{\text{Diff}}(\text{id}, \varphi) = \text{dist}_{\text{Met}}(g, \varphi_* g) = \text{dist}_{\text{Met}}(g, \varphi^* g)$ where g is the restriction of the Euclidean metric to the brain domain. This choice greatly increases computational efficiency since we can now use the formulas from Section 3 as explicit formulas for both terms of the energy functional. To minimize the energy functional, we use a gradient flow approach described in Algorithm 1, where the gradient on $\text{Diff}(M)$ is calculated with respect to a right invariant Sobolev metric of order one, which at the identity is given by

$$G_{\text{id}}(u, v) = \int g(\Delta u, v) dx + \lambda \sum_{i=1}^k g(u, \xi_i) g(v, \xi_i) dx \quad (17)$$

where $\lambda > 0$ is a weight parameter, g is the restriction of the Euclidean metric to the brain domain and ξ_1, \dots, ξ_k is an orthonormal basis of the harmonic vector fields on M . This metric, sometimes called the information metric, has been first introduced by Modin




Figure 3: Relations between the information metric on the diffeomorphism group, the Ebin metric on the space of Riemannian metrics, and the Fisher-Rao metric on the space of densities. The mappings I and III are Riemannian submersions and the mapping II is an isometric embedding. Furthermore the diagram is commutative, i.e., $\text{I} = \text{II} \circ \text{III}$. Note, that the Ebin metric G^E and the Fisher-Rao metric G^{FR} are of order zero, while the information metric G^I is a first order Sobolev metric. This discrepancy in the orders of the metric is explained by the fact that the mappings I and II contain a derivative.

(2015); see also Bauer et al. (2015). We choose this specific gradient because of the relation of the information metric to both the Ebin metric on the space of metrics and the Fisher-Rao metric on the space of probability densities. We summarize the relations between these geometries in Figure 3; for a precise description of the underlying geometric picture we refer to Khesin et al. (2013) and Bauer et al. (2015).

Note, that our framework allows for the immediate inclusion of points, curves, surfaces and images in the registration problem. Image intensity information, for example, can be easily incorporated in the registration problem by simply adding an appropriate similarity measure for the image term (e.g. the standard L^2 metric between the deformed moving image and the fixed image) to the energy functional. The minimization problem incorporating image intensities naturally becomes:

$$E(\varphi) = \text{dist}_{\text{Diff}}^2(\text{id}, \varphi) + \lambda_1 \text{dist}_{\text{Met}}^2(g_0, \varphi_* g_1) + \lambda_2 \text{dist}_{L^2}^2(I_0, \varphi_* I_1), \quad (18)$$

where λ_1, λ_2 are the relative weights and $\varphi_* I_1 = I_1 \circ \varphi^{-1}$ is the natural left action of Diff on L^2 images. In this convention I_1 is called the moving image (g_1 the moving metric, resp.). It is important to note that the energy functional only depends on φ^{-1} and not on φ . To see this we use the invariance of the geodesic distance on Diff and the definition of the pushforward to rewrite (18) as

$$E(\varphi) = \text{dist}_{\text{Diff}}^2(\text{id}, \varphi^{-1}) + \lambda_1 \text{dist}_{\text{Met}}^2(g_0, g_1(T\varphi^{-1}, T\varphi^{-1}\cdot)) + \lambda_2 \text{dist}_{L^2}^2(I_0, I_1 \circ \varphi^{-1}). \quad (19)$$

Consequently, we will use φ^{-1} rather than φ as our optimization variable, c.f. Algorithm 1 below.

4.1 Estimating the Atlas for a Population of Connectomes

Given a collection of connectomes modeled as points on an abstract Riemannian manifold, we can directly apply least-squares estimation to define the average connectome. Thus the

Algorithm 1 Inexact Metric Matching Algorithm**Inputs:**Fixed and moving metrics g_0, g_1 ; Fixed and moving images I_0, I_1 **Initialize:**

```

step size  $\epsilon$ ; weight parameters  $\lambda_1, \lambda_2$ ; max iteration times MaxIter
 $\varphi^{-1}, E \leftarrow \text{id}, 0$ 
for iteration = 0 : MaxIter do
   $\varphi_* g_1 \leftarrow (d\varphi^{-1})^T (g_1 \circ \varphi^{-1}) (d\varphi^{-1})$             $\triangleright$  Pushforward of  $g_1$  by  $\varphi$ 
   $\varphi_* I_1 \leftarrow I_1 \circ \varphi^{-1}$             $\triangleright$  Pushforward of  $I_1$  by  $\varphi$ 
   $E \leftarrow \text{EbinEnergy}(\varphi_* g_1, g_0, \varphi_* I_1, I_0, \lambda_1, \lambda_2)$             $\triangleright$  Calculate energy by Equation (18)
   $v \leftarrow \Delta^{-1}(E. \text{grad})$             $\triangleright$  Transfer gradient w.r.t. information metric to  $L^2$ 
   $\psi \leftarrow \text{id} - \epsilon v$             $\triangleright$  Construct the approximation
   $\varphi^{-1} \leftarrow \psi \circ \varphi^{-1}$             $\triangleright$  Update the diffeomorphism
end for
return  $\varphi^{-1}$ 

```

template estimation problem can be formulated as a joint minimization problem:

$$\hat{g} = \operatorname{argmin}_{g, \varphi_i} \sum_{i=1}^N \text{dist}_{\text{Diff}}^2(\text{id}, \varphi_i) + \lambda_1 \text{dist}_{\text{Met}}^2(g, (\varphi_i)_* g_i) \quad (20)$$

We use the iterative alternating algorithm proposed in Joshi et al. (2004) for solving the above optimization problem: we alternate gradient steps between optimizing with respect to each diffeomorphism, $\varphi_i, i = 1, \dots, N$, and minimizing with respect to the metric average \hat{g} . In the metric optimization step we use the Fréchet mean algorithm described in Section 3. The above procedure for estimating a connectome atlas can again be trivially extended to jointly estimate an image atlas consistent with the connectome atlas. To estimate the atlases jointly we use the following extended joint minimization problem:

$$\hat{g}, \hat{I} = \operatorname{argmin}_{g, I, \varphi_i} \sum_{i=1}^N \text{dist}_{\text{Diff}}^2(\text{id}, \varphi_i) + \lambda_1 \text{dist}_{\text{Met}}^2(g, (\varphi_i)_* g_i) + \lambda_2 \text{dist}_{L^2}^2(I, (\varphi_i)_* I_i) . \quad (21)$$

For the extended image-connectome atlas alternating algorithm at each iteration of the algorithm, the atlas image is updated by the simple average of the deformed individual subject images. This is a consequence of using the simple L^2 metric for the images. See Algorithm 2 for details of this process.

4.2 Implementation Details

In this section, we describe some implementation details that are important for reliable and efficient algorithm performance, particularly when applied to 3D brain data.

4.2.1 METRIC ESTIMATION

As done in Hao et al. (2014), we apply a mask to both the connectome metric estimation process and the atlas-building algorithm for two reasons. First, it is important that we constrain the problem to biologically realistic white matter tracts by not allowing tractography

Algorithm 2 Atlas Building Algorithm**Inputs:**metric fields $G = \{g_1, \dots, g_N\}$, images $I = \{I_1, \dots, I_N\}$ **Initialize:**

```

max iteration times MaxIter
for iteration = 0 : MaxIter do
     $g_{\text{mean}} \leftarrow \text{FrechetMean}(G)$             $\triangleright$  Section 3, compute with random order each time
     $I_{\text{mean}} \leftarrow \text{EuclideanMean}(I)$ 
    for  $i = 1 : N$  do
         $\varphi_i^{-1} \leftarrow \text{MetricMatching}(g_{\text{mean}}, g_i, I_{\text{mean}}, I_i)$             $\triangleright$  Algorithm 1
         $g_i \leftarrow (\varphi_i)_* g_i$             $\triangleright$  Update  $g_i$  by pushforward of  $\varphi_i$ 
         $I_i \leftarrow (\varphi_i)_* I_i$             $\triangleright$  Update  $I_i$  by pushforward of  $\varphi_i$ 
    end for
end for
return  $g_{\text{mean}}$ 

```

to flow through regions of CSF. Second, we avoid numeric issues associated with processing air and other noisy regions outside the skull. This also speeds up computation, as we need to look only at voxels inside the masked region instead of the entire image volume. Care must be taken with both first and second derivatives to use an appropriate and accurate finite difference stencil near the boundaries of the mask to ensure only points inside the mask are used. When matching the geodesics to a vector field consisting of the principal eigenvector directions, it is important to ensure that the eigenvector signs are consistent prior to computing derivatives of the vector field.

4.2.2 ATLAS BUILDING

For the atlas building algorithm, we deform each individual mask into atlas space at each outer iteration, and then apply the union of these deformed masks when computing the current atlas estimate. For each iteration of the atlas building algorithm, we perform only two iterations inside the metric matching function to avoid overfitting the individual metrics to early estimates of the Fréchet mean. Though we have not observed any subject order-related problems with the Fréchet mean calculation in practice, we randomize the order of the subjects each time we compute the mean to avoid the possibility of introducing problems related to subject order. In practice, we find the algorithm behaves well when we update ϵ in Algorithm 1 such that $1/\epsilon$ is approximately equal to the energy from (18).

For computational efficiency, all our algorithms are implemented in **PyTorch**, which allows us to take advantage of the built-in GPU acceleration and automatic differentiation. The computationally most expensive part in the geodesic distance of the Ebin metric is the calculation of the matrix logarithm, i.e., the calculation of $k = \log(g_0^{-1} g_1)$. Currently the matrix logarithm function is not implemented in **PyTorch**, and other alternatives such as the function provided by **Scipy** do not support automatic differentiation. Therefore we calculate the Cholesky factorization of $g_0 = GG^T$. We aim to use this factorization to reduce the eigendecomposition of the nonsymmetric matrix $g_0^{-1} g_1$ to the eigendecomposition of the symmetric matrix $W = G^{-1} g_1 (G^T)^{-1}$: writing $W = Q\Lambda Q^T$ for the eigendecomposition of

W , we directly obtain the eigendecomposition of $g_0^{-1}g_1 = V\Lambda V^{-1}$, where $V = (G^T)^{-1}Q$. This in turn allows us to calculate $k = \log(g_0^{-1}g_1) = V\log(\Lambda)V^{-1}$.

Consequently, the bottleneck of our atlas building algorithm is the large amount of eigendecomposition problems that have to be computed – in each iteration step we have to solve $N \times \text{Res}$ eigendecompositions of $n \times n$ matrices, where Res is the resolution of the image and n is the dimension of the domain.

It turns out that the GPU implementation of `torch.linalg.eig` is not well-designed for solving eigendecomposition of numerous small matrices. The extremely low speed made our algorithms unusable for any experiments in 3D. To speed up these calculations, we reimplement the eigendecomposition based on the work of Lenssen et al. (2019), which leads to an order of magnitude increase in performance.

Powered by the `torch.autograd` module, we can now easily solve the gradient of the Ebin energy in Algorithm 1 without a closed-form gradient solution. Our 3D atlas building code, including both metric-only matching and joint matching, will become publicly available at <https://github.com/aarentai/Atlas-Building-3D>.

4.2.3 GEODESIC TRACTOGRAPHY

Several computational strategies have been proposed to compute white matter tractography as geodesic curves. O’Donnell et al. (2002) develop a level set approach to solve the Eikonal-type equation of the geodesic distance transform from a seed point. Fletcher et al. (2007) extend this approach to simultaneously solve for the entire set of minimal geodesics between two brain regions by solving two Hamilton-Jacobi PDEs. These Hamilton-Jacobi PDEs can be solved quickly on GPUs using the method of Jeong et al. (2007). While these strategies are better choices for a production pipeline, for this paper we chose to directly integrate the geodesic equation from seed points using the principal eigenvector at each seed as the initial direction for shooting.

Given a Riemannian metric g , we compute the corresponding Christoffel symbol Γ via

$$\Gamma_{ij}^k = \frac{1}{2} \sum_l^n g^{kl} \left(\frac{\partial g_{jl}}{\partial x^i} + \frac{\partial g_{il}}{\partial x^j} - \frac{\partial g_{ij}}{\partial x^l} \right), \quad (22)$$

where g^{ij} denotes the components of the inverse metric tensor. Together with the position γ and velocity $\dot{\gamma}$, the Christoffel symbols enable us to find the acceleration $\ddot{\gamma}$ at time t by solving the geodesic equation $\ddot{\gamma}^{(k)} + \Gamma_{ij}^k \dot{\gamma}^{(i)} \dot{\gamma}^{(j)} = 0$, which gives

$$\ddot{\gamma}^{(k)}(t) = - \begin{bmatrix} \dot{\gamma}^{(1)}(t) & \dot{\gamma}^{(2)}(t) & \dot{\gamma}^{(3)}(t) \end{bmatrix} \begin{bmatrix} \Gamma_{11}^k(\gamma(t)) & \Gamma_{12}^k(\gamma(t)) & \Gamma_{13}^k(\gamma(t)) \\ \Gamma_{21}^k(\gamma(t)) & \Gamma_{22}^k(\gamma(t)) & \Gamma_{23}^k(\gamma(t)) \\ \Gamma_{31}^k(\gamma(t)) & \Gamma_{32}^k(\gamma(t)) & \Gamma_{33}^k(\gamma(t)) \end{bmatrix} \begin{bmatrix} \dot{\gamma}^{(1)}(t) \\ \dot{\gamma}^{(2)}(t) \\ \dot{\gamma}^{(3)}(t) \end{bmatrix} \quad (23)$$

where $\ddot{\gamma}^{(k)}, \dot{\gamma}^{(k)}$ are the components of the acceleration vector $\ddot{\gamma}$ and velocity vector $\dot{\gamma}$.

After the initial conditions for the position and velocity are given and $\ddot{\gamma}^{(k)}(0)$ is computed, we update the acceleration at subsequent time steps using a fourth-order Runge-Kutta scheme (Press et al., 1986). See also the function `algo.geodesic.geodesicpath()` in our open access repository.

5. Results

We first demonstrate our framework in 2D using synthetic data and 2D data extracted from brain images. Next, we construct a 3D connectome atlas from DWMRI for a subset of subjects from the Human Connectome Project. We show that we can use the complementary information from T1-weighted MRI for those same subjects to build an integrated multimodal atlas. Finally, we demonstrate that the multimodal atlas preserves both local and long-range connectivity information by computing both whole-brain and seed region-based geodesic tractography of the atlas.

5.1 2D Simulated Data

We verified our method by generating vector fields whose central integral curves are a family of parameterized cubic functions. We used the method of parallel curves to add vectors for additional integral curves parallel to the central curve with a distance $k \in [-0.2, 0.2]$ from the central curve. We then constructed tensors whose principal eigenvectors align with the generated vector fields and that have a specified major axis to minor axis ratio of 6:1.

We first estimated the adaptive metric conformal to the inverse-tensor metric such that the geodesics of the adaptive metrics align with the integral curves of the simulated vector fields. After finding the connectome metric for each subject, we ran 400 iterations of the atlas building Algorithm 2 using only the metric distance as shown in (20) to estimate the atlas in Figure 4. To help the diffeomorphisms update smoothly, we set $\lambda_1 = 100$ in (16) and the step size $\epsilon = 5$ in Algorithm 1.

We compared a geodesic of the atlas starting from a particular seed point with geodesics of the four connectome metrics starting from the atlas seed point mapped into individual space. Figure 4 shows these individual geodesics in atlas space before and after applying the diffeomorphisms. We see that the atlas geodesic is nicely centered in the middle of the undeformed individual geodesics as expected. Also, the deformed individual geodesics align well with the atlas geodesic. The upper right panel shows the mean distance between the atlas geodesic and deformed subject geodesic in atlas space in different stages. After 75 iterations of optimization, we can see the distance of the four pairs of geodesics almost converged.

5.2 2D Real Data

To illustrate how our method works with real data, we used two subjects from the Human Connectome Project Young Adult (HCP) dataset Glasser et al. (2013). For each subject, we fit a diffusion tensor model to the images with a b -value of 1000 using `dtifit` from FSL Basser et al. (1994) and generated a white-matter mask based on fractional anisotropy values thresholded between 0.25 and 1. To process in 2D, we extracted a single axial slice from each image along with the x and y components of the associated tensors. We estimated the adaptive connectome metric from the inverse-tensor metric associated with the 2D diffusion tensors. We would like to emphasize that the 2D atlas construction is only for illustrative purposes. In real human data no tract is expected to live in a single plane and 2D processing is not appropriate.




Figure 4: Left and center: geodesics of four synthetic metrics starting from the atlas seed point (X) mapped into each metric’s space. Upper right: mean distance trend between atlas geodesic and every deformed subject geodesic in atlas space. Lower right: estimated atlas with geodesic (orange) overlaid on geodesics from the four metrics deformed into atlas space.

To generate the atlas shown in Figure 5, we ran atlas building with only the metric distance terms from (20) for 5000 iterations with $\lambda_1 = 100$, $\epsilon = 1$, which took 50 minutes on an Intel Xeon Silver 4108 CPU. The regularization term, λ_1 , balances the magnitudes of the diffeomorphisms from each subject’s connectome metric to the atlas. To ensure that the final geodesics in the atlas also follow the major eigenvectors of the atlas tensors, we solve for the α conformal factor for the atlas as described in Section 2.

5.3 3D Metric Atlas

We used six subjects from the Human Connectome Project Young Adult (HCP) dataset Glasser et al. (2013) in this experiment as shown in Figure 6. For each subject, we fit a diffusion tensor model to the images using b -values of 1000, 2000, and 3000 using `dtifit` from FSL Basser et al. (1994) and generated a white-matter mask by keeping voxels with fractional anisotropy between 0.25 and 1. We rigidly registered the T1 images for each subject using `brainsfit` from Johnson et al. (2007), and applied that rigid registration to the subject’s white matter mask and diffusion tensors, being sure to reorient the individual tensors. We estimated α for the adaptive connectome metric from the inverse-tensor metric




Figure 5: Left: diffeomorphism from HCP subjects (103818, 111312) to the atlas. Center: each subject’s connectome metric and a geodesic (blue, red) starting from the atlas seed (X) mapped to subject space. Right: atlas and a geodesic (orange) starting at the seed (X). Subject geodesics are mapped to atlas space (blue, red). Bottom: detailed view of corpus callosum.

associated with the diffusion tensors after smoothing them with a Gaussian filter, $\sigma = 1.5$ and cleaning any tensors that were not positive semi-definite. The estimated α was clipped to the range $[-2, 2]$ before applying it to the inverse-tensor metric associated with the unsmoothed diffusion tensors to create the adaptive connectome metric. In the atlas building process, we used the metric distance terms from (20) and set the step size $\epsilon = 0.005$ in Algorithm 1 and $\lambda_1 = 1.0$. For 800 iterations, the algorithm took about 12 hours on an Nvidia Titan RTX GPU, which features 24GB VRAM.

To generate the T1 atlas for this experiment, we computed the mean of each subject’s T1 image after deforming it by the diffeomorphism found from the metric atlas construction. As shown in Figure 7, the white matter atlas looks reasonable, but the T1 atlas is blurry, especially in gray matter regions because the atlas construction process is not using any gray matter information to guide the diffeomorphisms in these regions.

5.4 3D Joint Atlas

We used the same six subjects to construct a joint T1 and metric atlas by adding in the image matching term as shown in (21). We set the weight of the image term, λ_2 in (18), to 1.0×10^{-8} , which balances the image term loss to approximately the same magnitude as the metric term loss. All other parameters were kept the same as in the metric-only atlas building process. The resulting metric and T1 atlases are shown in Figure 7. Note how much more well-defined the gray matter regions are when the T1 information is included




Figure 6: Axial, coronal, and sagittal views of the 6 HCP subjects' T1 images and their metrics colored by the direction of the principal eigenvector of the inverse of the metric. Red is oriented from left to right, green from anterior to posterior and blue from inferior to superior directions.

in the optimization process. At the same time, the metric atlas is not degraded by adding this term in, and may even be slightly better than before.

The whole-brain tractography shown in Figure 8 was performed by seeding eight geodesics per voxel in the white matter mask, stopping when each geodesic left the averaged white matter mask. Individual tracts were created by drawing seed regions in the atlas space and transforming the regions back to subject space. Each voxel in the seed region had 27 seeds. Geodesics terminated when fractional anisotropy dropped below 0.2 for the subject tractography, or when the geodesic left the averaged white matter mask for the atlas tractography. Other than the seed region and FA thresholding, we did not apply any stopping or rejection criteria based on angle, tract length, anatomical priors or any of the other techniques used to clean up false positives. We performed tractography visualization using 3D Slicer (www.slicer.org) via the SlicerDMRI project (<http://dmri.slicer.org>) (Norton et al., 2017; Zhang et al., 2020). Note that the whole-brain tractography can also be registered to




Figure 7: Axial, coronal and sagittal views of the T1 and metric atlas. Left columns: the atlas produced using only the metric matching term. Right columns: the atlas produced using metric and image matching terms jointly.

the O'Donnell atlas using the `WhiteMatterAnalysis` toolkit (O'Donnell and Westin, 2007; O'Donnell et al., 2012; Zhang et al., 2018) in order to do clustering based on that atlas if desired. Example tracts for the genu of the corpus callosum, corticospinal tract, and cingulum computed by integrating geodesics of the joint atlas are shown in Figure 9.

6. Conclusions

In this paper, we introduce a novel framework for statistically analyzing structural connectomes by representing them as a point on the manifold of Riemannian metrics, enabling us to perform geometric statistics. Using this representation, we build a framework for connectome atlas construction based on the action of the diffeomorphism group and the natural Ebin metric on the space of all Riemannian metrics. Because this framework is compatible with existing image atlas construction frameworks, we are then able to construct an integrated multimodal atlas using complementary white matter and cortical information from DWMRI and T1-weighted MRI simultaneously. In future work we plan to investigate in more detail the convergence properties of the proposed algorithms and quantitatively compare our approach to other existing methods. We expect this new methodology to open up opportunities for a deeper understanding of structural connectomes, their variabilities and their relationships to cortical structure.




Figure 8: Top row: metric colored by orientation overlaid on T1 for joint atlas of six HCP individuals. Bottom row: whole-brain geodesic tractography for joint atlas, with tract segments colored by their orientation.




Figure 9: Examples of joint atlas tractography generated from seeds placed in atlas seed regions, with tract segments colored by orientation. Left: axial view of tractography seeded in the genu of the corpus callosum. Center: coronal view of tractography seeded in the corticospinal tract. Right: axial view of tractography seeded in the cingulum.

Acknowledgments

M. Bauer was supported by NSF grants DMS-1912037, DMS-1953244. K. Campbell, H. Dai, S. Joshi and P. Fletcher were supported by NSF grant DMS-1912030. Z. Su was supported by NSF grant DMS-1912037, NIH/NIAAA award R01-AA026834. Data were provided in

part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

Conflicts of Interest

The authors have no conflicts of interest.

References

Daniel C Alexander, Carlo Pierpaoli, Peter J Basser, and James C Gee. Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging*, 20(11):1131–1139, 2001.

Brian Avants, Jeffrey T Duda, Hui Zhang, and James C Gee. Multivariate normalization with symmetric diffeomorphisms for multivariate studies. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 359–366. Springer, 2007.

Peter J Basser, James Mattiello, and Denis LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance, Series B*, 103(3):247–254, 1994.

Peter J Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic resonance in medicine*, 44(4):625–632, 2000.

Martin Bauer, Sarang Joshi, and Klas Modin. Diffeomorphic density matching by optimal information transport. *SIAM Journal on Imaging Sciences*, 8(3):1718–1751, 2015.

M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, 2005.

Timothy EJ Behrens, Mark W Woolrich, Mark Jenkinson, Heidi Johansen-Berg, Rita G Nunes, Stuart Clare, Paul M Matthews, J Michael Brady, and Stephen M Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 50(5):1077–1088, 2003.

Kristen M Campbell, Haocheng Dai, Zhe Su, Martin Bauer, P Thomas Fletcher, and Sarang C Joshi. Structural connectome atlas construction in the space of Riemannian metrics. In *International Conference on Information Processing in Medical Imaging*, pages 291–303. Springer, 2021.

Guang Cheng, Hesamoddin Salehian, John R Forder, and Baba C Vemuri. Tractography from HARDI using an intrinsic unscented Kalman filter. *IEEE transactions on medical imaging*, 34(1):298–305, 2015.

Daan Christiaens, Marco Reisert, Thijs Dhollander, Stefan Sunaert, Paul Suetens, and Frederik Maes. Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model. *Neuroimage*, 123:89–101, 2015.

Brian Clarke. The completion of the manifold of Riemannian metrics. *Journal of Differential Geometry*, 93(2):203–268, 2013a.

Brian Clarke. Geodesics, distance, and the CAT(0) property for the manifold of Riemannian metrics. *Mathematische Zeitschrift*, 273(1-2):55–93, 2013b.

B. S. DeWitt. Quantum theory of gravity. I. The canonical theory. *Phys. Rev.*, 160 (5): 1113–1148, 1967.

David G. Ebin. The manifold of Riemannian metrics. In *Global Analysis (Proc. Sympos. Pure Math., Vol. XV, Berkeley, Calif., 1968)*, pages 11–40. Amer. Math. Soc., Providence, R.I., 1970.

P Thomas Fletcher and Sarang Joshi. Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, 87(2):250–262, 2007.

P Thomas Fletcher, Ran Tao, Won-Ki Jeong, and Ross T Whitaker. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In *Information Processing in Medical Imaging (IPMI)*, pages 346–358. Springer, 2007.

Daniel S Freed, David Groisser, et al. The basic geometry of the manifold of Riemannian metrics and of its quotient by the diffeomorphism group. *The Michigan Mathematical Journal*, 36(3):323–344, 1989.

Andrea Fuster, Tom Dela Haije, Antonio Tristán-Vega, Birgit Plantinga, Carl-Fredrik Westin, and Luc Florack. Adjugate diffusion tensors for geodesic tractography in white matter. *Journal of Mathematical Imaging and Vision*, 54(1):1–14, 2016.

Olga Gil-Medrano and Peter W Michor. The Riemannian manifold of all Riemannian metrics. *Quarterly Journal of Mathematics (Oxford)*, 42:183–202, 1991.

Matthew F Glasser, Stamatios N Sotropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, and Jonathan R Polimeni. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

Joan Glaunès, Anqi Qiu, Michael I Miller, and Laurent Younes. Large deformation diffeomorphic metric curve mapping. *International journal of computer vision*, 80(3):317, 2008.

Vikash Gupta, Grégoire Malandain, Nicholas Ayache, and Xavier Pennec. A framework for creating population specific multimodal brain atlas using clinical T1 and diffusion tensor images. In *Computational Diffusion MRI*, pages 99–108. Springer, 2016.

Xiang Hao, Kristen Zygmunt, Ross T Whitaker, and P Thomas Fletcher. Improved segmentation of white matter tracts with adaptive Riemannian metrics. *Medical Image Analysis*, 18(1):161–175, 2014.

Jeffrey Ho, Guang Cheng, Hesamoddin Salehian, and Baba Vemuri. Recursive Karcher expectation estimators and geometric law of large numbers. In *Artificial Intelligence and Statistics*, pages 325–332, 2013.

Saad Jbabdi, Mark W Woolrich, Jesper LR Andersson, and TEJ Behrens. A Bayesian framework for global tractography. *Neuroimage*, 37(1):116–129, 2007.

Won-Ki Jeong, P Thomas Fletcher, Ran Tao, and Ross Whitaker. Interactive visualization of volumetric white matter connectivity in DT-MRI using a parallel-hardware Hamilton-Jacobi solver. *IEEE transactions on visualization and computer graphics*, 13(6):1480–1487, 2007.

Ben Jeurissen, Maxime Descoteaux, Susumu Mori, and Alexander Leemans. Diffusion mri fiber tractography of the brain. *NMR in Biomedicine*, 32(4):e3785, 2019.

Hans Johnson, Greg Harris, Kent Williams, et al. BRAINSFit: mutual information rigid registrations of whole-brain 3D images, using the insight toolkit. *Insight J*, 57(1):1–10, 2007.

Sarang Joshi, Brad Davis, Matthieu Jomier, and Guido Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23:S151–S160, 2004.

Sarang C Joshi and Michael I Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8):1357–1370, 2000.

B. Khesin, J. Lenells, G. Misiołek, and S. C. Preston. Geometry of diffeomorphism groups, complete integrability and geometric statistics. *Geom. Funct. Anal.*, 23(1):334–366, 2013. ISSN 1016-443X.

Christophe Lenglet, Rachid Deriche, and Olivier Faugeras. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In *European Conference on Computer Vision*, pages 127–140. Springer, 2004.

Jan Eric Lenssen, Christian Osendorfer, and Jonathan Masci. Deep iterative surface normal estimation, 2019.

Klas Modin. Generalized Hunter-Saxton equations, optimal information transport, and factorization of diffeomorphisms. *J. Geom. Anal.*, 25(2):1306–1334, 2015. ISSN 1050-6926. doi: 10.1007/s12220-014-9469-2. URL <http://dx.doi.org/10.1007/s12220-014-9469-2>.

Susumu Mori, Kenichi Oishi, Hangyi Jiang, Li Jiang, Xin Li, Kazi Akhter, Kegang Hua, Andreia V. Faria, Asif Mahmood, Roger Woods, Arthur Toga, Bruce Pike, Pedro Rosa Neto, Alan Evans, Jiangyang Zhang, Hao Huang, Michael I. Miller, Peter van Zij, and John Mazziotta. Stereotaxic white matter atlas based on diffusion tensor imaging in an ICBM template. *Neuroimage*, 40(2):570–582, 2008.

Xinyu Nie and Yonggang Shi. Topographic filtering of tractograms as vector field flows. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 564–572. Springer, 2019.

Isaiah Norton, Walid Ibn Essayed, Fan Zhang, Sonia Pujol, Alex Yarmarkovich, Alexandra J Golby, Gordon Kindlmann, Demian Wassermann, Raul San Jose Estepar, Yogesh Rathi, et al. SlicerDMRI: open source diffusion MRI software for brain cancer research. *Cancer research*, 77(21):e101–e103, 2017.

Lauren O’Donnell, Steven Haker, and Carl-Fredrik Westin. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 459–466, 2002.

Lauren J O’Donnell and Carl-Fredrik Westin. Automatic tractography segmentation using a high-dimensional white matter atlas. *IEEE transactions on medical imaging*, 26(11):1562–1575, 2007.

Lauren J O’Donnell, William M Wells, Alexandra J Golby, and Carl-Fredrik Westin. Unbiased groupwise registration of white matter tractography. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 123–130. Springer, 2012.

William H Press, William T Vetterling, Saul A Teukolsky, and Brian P Flannery. *Numerical recipes*, volume 818. Cambridge university press Cambridge, 1986.

Rick Sengers, Luc Florack, and Andrea Fuster. Geodesic tubes for uncertainty quantification in diffusion MR. In *Information Processing in Medical Imaging (IPMI)*, pages 279–290. Springer, 2021.

Arthur W Toga, Paul M Thompson, Susumu Mori, Katrin Amunts, and Karl Zilles. Towards multimodal atlases of the human brain. *Nature Reviews Neuroscience*, 7(12):952–966, 2006.

Alena Uus, Maximilian Pietsch, Irina Grigorescu, Daan Christiaens, Jacques-Donald Tournier, Lucilio Cordero Grande, Jana Hutter, David Edwards, Joseph Hajnal, and Maria Deprez. Multi-channel registration for diffusion MRI: Longitudinal analysis for

the neonatal brain. In *International Workshop on Biomedical Image Registration*, pages 111–121. Springer, 2020.

Marc Vaillant and Joan Glaunès. Surface matching via currents. In *International Conference on Information Processing in Medical Imaging*, pages 381–392. Springer, 2005.

Fang-Cheng Yeh, Sandip Panesar, David Fernandes, Antonio Meola, Masanori Yoshino, Juan C Fernandez-Miranda, Jean M Vettel, and Timothy Verstynen. Population-averaged atlas of the macroscale human structural connectome and its network topology. *NeuroImage*, 178:57–68, 2018.

Fan Zhang, Ye Wu, Isaiah Norton, Laura Rigolo, Yogesh Rathi, Nikos Makris, and Lauren J O’Donnell. An anatomically curated fiber clustering white matter atlas for consistent white matter tract parcellation across the lifespan. *NeuroImage*, 179:429–447, 2018.

Fan Zhang, Thomas Noh, Parikshit Juvekar, Sarah F Frisken, Laura Rigolo, Isaiah Norton, Tina Kapur, Sonia Pujol, William Wells III, Alex Yarmarkovich, et al. SlicerDMRI: Diffusion MRI and tractography research software for brain cancer surgery planning and visualization. *JCO clinical cancer informatics*, 4:299–309, 2020.