

Geodesic Regression Characterizes 3D Shape Changes in the Female Brain During Menstruation

Adele Myers
UC Santa Barbara
adele@ucsb.edu

Caitlin Taylor
UC Santa Barbara
cmtaylor@ucsb.edu

Emily Jacobs
UC Santa Barbara
emily.jacobs@gmail.com

Nina Miolane
UC Santa Barbara
ninamiolane@ucsb.edu

Abstract

Women are at higher risk of Alzheimer’s and other neurological diseases after menopause, and yet research connecting female brain health to sex hormone fluctuations is limited. We seek to investigate this connection by developing tools that quantify 3D shape changes that occur in the brain during sex hormone fluctuations. Geodesic regression on the space of 3D discrete surfaces offers a principled way to characterize the evolution of a brain’s shape. However, in its current form, this approach is too computationally expensive for practical use. In this paper, we propose approximation schemes that accelerate geodesic regression on shape spaces of 3D discrete surfaces. We also provide rules of thumb for when each approximation can be used. We test our approach on synthetic data to quantify the speed-accuracy trade-off of these approximations and show that practitioners can expect very significant speed-up while only sacrificing little accuracy. Finally, we apply the method to real brain shape data and produce the first characterization of how the female hippocampus changes shape during the menstrual cycle as a function of progesterone: a characterization made (practically) possible by our approximation schemes. Our work paves the way for comprehensive, practical shape analyses in the fields of bio-medicine and computer vision. Our implementation is [publicly available on GitHub](#).

1. Introduction

Women are more likely to experience Alzheimer’s, cognitive decline, and navigational issues after menopause [2, 20]. Yet, topics relevant to female brain health such as menstruation, pregnancy, menopause, and their associated female sex hormone fluctuations only account for 0.3% of the neuroimaging literature between 1995 and 2022 [17]. The hippocampal formation (a brain structure) is an excellent diagnostic tool for investigating the connection between female brain health and sex hormones, as it is the first cor-

tical region to harbor neuropathology in the progression to Alzheimer’s (causing characteristic shape changes visible on magnetic resonance images (MRI)) [5, 13], and it is also very sensitive to sex hormone fluctuations [7]. Sex hormone fluctuations significantly influence brain anatomy and function in healthy subjects [18, 16, 15]. Enhancing our understanding of how hormonal fluctuations affect the healthy brain is crucial to explaining why females are later more at risk for neurological conditions after menopause [4]. We seek to close this knowledge gap by starting with one question: How does the hippocampal formation respond to monthly fluctuations in ovarian hormones during the menstrual cycle?

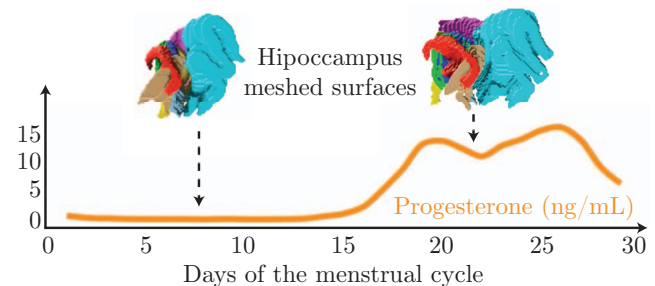


Figure 1. During the menstrual cycle, the ovaries cyclically release female sex hormones such as progesterone into the blood. We propose practical shape analysis tools that quantify 3D shape changes in the brain that occur during progesterone fluctuations, with a focus on the hippocampus: a structure involved in memory and navigation. Visualization created from data by [18, 15].

Recent research has shown that certain substructures of the hippocampal formation change their volume over the course of the menstrual cycle in response to progesterone, but no significant volumetric change was found on the whole-formation level [18]. Our 3D visualizations (Fig. 1) show that the hippocampal formation does change its *shape* on a whole-formation level, but no team has quantified how this effect depends on progesterone levels. This is not surprising because quantifying 3D shape changes is technically challenging, and is in fact is an active research area in itself in mathematics and computer vision. For example,

quantifying surface shape changes as a function of a continuous variable like progesterone can theoretically be performed through geodesic regression on Riemannian manifolds [6, 19]: an extension of linear regression dedicated to shape spaces. However, in its current form, geodesic regression on surface spaces is too slow for practical use. Here, we bridge the gap between computer vision and clinical neuroimaging by presenting a new practical method, a hybrid between geodesic and linear regression, that allows us to quantify how the *shape* of the hippocampal formation changes in response to progesterone.

Contributions We offer several contributions that span the fields of machine learning, differential geometry and clinical neuroimaging. First, in machine learning, we introduce our hybrid geodesic-linear regression method: a faster geodesic regression model which uses linear residuals instead of geodesic residuals in its loss function and also uses a linear regression result to initialize its geodesic regression optimization. Then, we perform extensive synthetic experiments to offer rules of thumb for deciding between linear, geodesic, and geodesic-linear regression. In differential geometry, these results give novel intuition about the curvature of the nonlinear data space of surface shapes. In clinical neuroimaging, these rules of thumb provide practitioners with guidelines to decide on the speed-accuracy trade-off between the regression types, revealing whether a surface mesh sequence can be adequately characterized by the considerably faster linear or geodesic-linear regressions without sacrificing accuracy. Finally, we apply our paradigm to real brain magnetic resonance images (MRIs) of a female brain through the menstrual cycle. We characterize, for the first time, shape changes in the female hippocampal formation as a function of progesterone.

2. Related Works

Consider a series of hippocampal shapes, with the surface of each shape described as a mesh extracted via segmentation from a full-brain MRI (Fig. 1). The *shapes* of these discrete surfaces (meshes) can be described either *extrinsically* or *intrinsically*. The extrinsic approach with the Large Deformation Diffeomorphic Metric Mapping framework [3] represents a surface shape in 3D space by the amount of deformation that one needs to apply on the whole 3D ambient grid containing the surface in order to deform a reference shape into the surface shape of interest. By contrast, the intrinsic approach only deforms the surface itself [1], and hence provides us with two advantages: (i) intuitive deformations that can be discussed with neuroscientists, (ii) higher computational efficiency with up to 10x acceleration [1]. We focus here on the intrinsic approach.

Analysis of Parameterized Surfaces In the intrinsic approach, surfaces can be either *parameterized* or *unparameterized*. Each mesh in a dataset of parameterized surfaces is constrained to have a consistent structure: the vertices of meshes in the dataset have one-to-one correspondences. By contrast, datasets of unparameterized surfaces relax this constraint such that statistical analyses can be performed independently of the number and indexation of the mesh vertices. While theoretically grounded, the computational complexity of this approach makes it unpractical for geodesic regression. As an example, statistical analyses within this framework are often limited to population average computations and machine learning algorithms that rely only on distances such as multidimensional scaling or k-means clustering [1, 9, 12].

Therefore, we consider the scenario where each mesh in the dataset is first parameterized to match a reference parameterization, and then statistical analysis such as regression is performed. In this scenario, the first natural choice is to consider linear regression on the set of 3D coordinates of the vertices. Linear regression has the advantage of being conceptually simple while enjoying analytical solutions: it will be our first baseline. However, it has the drawback that it does not enjoy parameterization-invariance. In other words, the distance (or dissimilarity) between two surfaces may change if we choose another reference parameterization, which may change the regression results. In the context of clinical application where the ultimate goal is human health, it is thought that we cannot afford such inconsistency. The alternative is to consider parameterization-invariant distances between parameterized surfaces. In differential geometry, this can be achieved by equipping the space of parameterized surfaces with a Riemannian metric that is invariant under reparameterizations [10, 9]. This process however, turns the data space into a nonlinear manifold, where linear regression needs to be generalized to geodesic regression [6, 19]. Geodesic regression does not enjoy a closed-form estimator and is typically computationally expensive. It will be our second baseline.

Computational Challenges of Geodesic Regression

Geodesic regression [6, 19] for parameterization-invariant Riemannian metrics presents unique computational challenges. In the Riemannian framework, calculating a single geodesic requires a computationally expensive approach. Geodesic regression solves an optimization problem by minimizing a mean square error (MSE) loss function. The MSE requires the computations of $n + 1$ geodesics at each iteration: 1 geodesic representing the generative model, and n geodesics required to compute the residuals in the MSE, where n is the number of surfaces (meshes) in the dataset. We observe that works developing the Riemannian framework limit the number of geodesic computations required

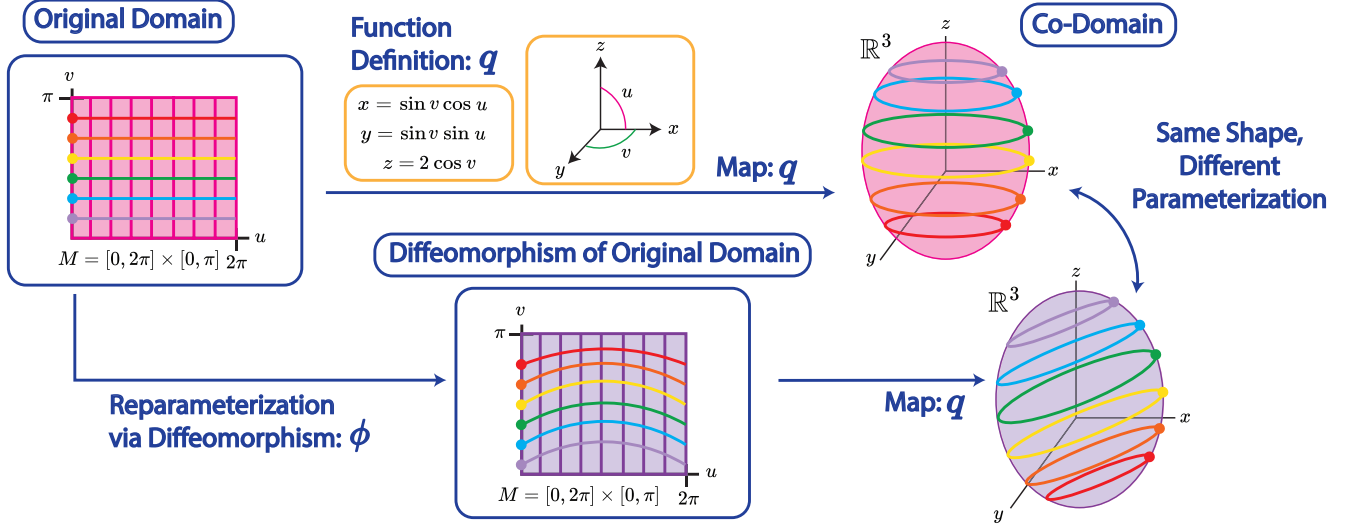


Figure 2. A surface is represented by a function $q : M \rightarrow \mathbb{R}^3$ that maps parameters $(u, v) \in M$ to points in 3D space $q(u, v) \in \mathbb{R}^3$ (top row). Its parameterization can be changed by applying a diffeomorphism ϕ to the domain M before mapping to \mathbb{R}^3 (bottom row).

for their analysis, and do not perform any form of regression. Kilian et al. [11] focus on geodesic interpolation or extrapolation of parameterized surfaces and do not study regression. Kurtek et al. [12] limit their experimental analysis to computing geodesics between pairs of unparameterized surfaces and performing clustering. Jermyn et al. [10] compute geodesics between pairs of parameterized surfaces and provide a classification experiment. Hartman et al [9] estimate population averages and perform dimension reduction with multidimensional scaling (MDS) and tangent PCA for parameterized and unparameterized surfaces. Bauer et al [1] compute geodesics and the population averages of unparameterized surfaces, together with multi-dimensional scaling and k-means clustering. We suspect that the authors did not perform geodesic regression in these works because of their computational costs, which we investigate here.

3. Background

This section presents the mathematical background necessary to formulate our approximation schemes for geodesic regression on the shape space of (hippocampal) surfaces. We refer to [9, 8] for additional details.

A. Riemannian Metrics and Geodesics We first introduce concepts in differential geometry necessary for geodesic regression. A *Riemannian metric* on a smooth manifold \mathcal{N} is a family $(G_p)_{p \in \mathcal{N}}$ of inner products on each tangent space $T_p \mathcal{N}$, such that G_p depends smoothly on the point $p \in \mathcal{N}$. Any Riemannian metric G yields a notion of distance between points q_0, q_1 on \mathcal{N} . Specifically, if $\gamma : [0, 1] \rightarrow \mathcal{N}$ is a smooth trajectory on \mathcal{N} with velocity vector at $t \in [0, 1]$ denoted as $\dot{\gamma}_t \in T_{\gamma(t)} \mathcal{N}$,

its length is defined as $L_\gamma = \int_0^1 \sqrt{G(\dot{\gamma}_t, \dot{\gamma}_t)} dt$ and the distance between any two points $q_0, q_1 \in \mathcal{N}$ is given by $d(q_0, q_1) = \inf_{\gamma: \gamma(0)=q_0, \gamma(1)=q_1} L_\gamma$.

A *geodesic* between two points q_0, q_1 that are “close” in \mathcal{N} is defined as a trajectory γ that locally realizes the shortest distance between q_0 and q_1 . Intuitively, a geodesic is the generalization to manifolds of the concept of a straight line in vector spaces. While some manifolds enjoy analytical expression for their geodesics, this is not case for the manifold of (hippocampal) surface shapes that we will consider here. Thus, geodesics will need to be computed numerically.

To this aim, geodesics are expressed as the solutions of the *geodesic equation*, which is an ordinary differential equation (ODE) which can be written in local coordinates as:

$$\ddot{\gamma}^k(t) + \Gamma_{ij}^k \dot{\gamma}^i(t) \dot{\gamma}^j(t) = 0, \quad (1)$$

for all times $t \in [0, 1]$ where Γ_{ij}^k are the Christoffel symbols associated with the Riemannian metric. Solving this ODE provides numerical solutions for geodesics.

To perform geodesic regression, we will also need two additional operations, called Exp and Log, which we define here. The map $(q, v) \mapsto \gamma_{q,v}(1)$ defined for $(q, v) \in \mathcal{I} \times T_q \mathcal{I}$ is called the *exponential map* (Exp) and essentially computes the point $\gamma_{q,v}(1)$ after following the geodesic of initial point $q \in \mathcal{N}$ and initial velocity $v \in T_q \mathcal{N}$. The inverse of the Exp map on its injectivity domain is called the *logarithm map* (Log).

B. Surfaces and Their Parameterizations A *continuous surface* can be described by a function $q : M \rightarrow \mathbb{R}^3$, where M is a two-dimensional space of parameters $(u, v) \in M$ that parameterize the 3D points $q(u, v) \in \mathbb{R}^3$ on the surface.

Intuitively, the function q deforms the space of parameters M to give the surface its distinct shape, e.g., the ellipsoid shown in the top row of Fig. 2. Mathematically, q is required to be an oriented smooth mapping in $C^\infty(M, \mathbb{R}^3)$ that is also regular in the sense that its differential dq is injective everywhere on M .

The *parameterization of a surface* refers to the placement of points on the surface. If we define one surface as $q : M \rightarrow \mathbb{R}^3$, then we can describe the same surface with a different parameterization by $q \circ \phi : M \rightarrow \mathbb{R}^3$, where ϕ is an orientation-preserving diffeomorphism of M . Intuitively, ϕ smoothly deforms the placement of parameters on the domain M , which in turn smoothly changes the placement of points in the co-domain, as shown with rainbow colors in the bottom row of Fig. 2. The change of parameterization ϕ does not change the shape of the surface, which is an ellipsoid in both rows of Fig. 2.

C. Space of Surfaces The space of surfaces is denoted $\mathcal{I} \subset C^\infty(M, \mathbb{R}^3)$. The space \mathcal{I} is an infinite dimensional manifold immersed in the infinite dimensional vector space $C^\infty(M, \mathbb{R}^3)$. The Riemannian metric we choose to equip the manifold \mathcal{I} with defines the distance between its points $q_0, q_1 \in \mathcal{I}$ and thus the notion of dissimilarity between the two surfaces q_0, q_1 . We consider the second-order Sobolev metric [9]:

$$G_q(h, k) = \int_M (a_0 \langle h, k \rangle + a_1 g_q^{-1} (dh_m, dk_m) + b_1 g_q^{-1} (dh_+, dk_+) + c_1 g_q^{-1} (dh_\perp, dk_\perp) + d_1 g_q^{-1} (dh_0, dk_0) + a_2 \langle \Delta_q h, \Delta_q k \rangle) \text{vol}_q, \quad (2)$$

where h, k are tangent vectors at point $q \in \mathcal{N}$; g_q^{-1} is the pullback metric from \mathbb{R}^3 that defines distances on the surface q itself; Δ_q is the Laplacian induced by q ; $dh_m, dh_+, dh_\perp, dh_0$ are orthogonal vector-valued one-forms and vol_q is the surface area measure of q . The scalars $a_0, a_1, a_2, b_1, c_1, d_1$ are weighting parameters that define distance between two surfaces based on how they are sheared, scaled, bent, or parameterized with respect to each other.

The choice of second-order Sobolev metric is motivated by the following facts. First, the zero-order and first-order Sobolev metrics yield less stable results in geodesic interpolation between complex 3D shapes [9]. Second, the weighting parameters $a_0, a_1, a_2, b_1, c_1, d_1$ defining the second-order Sobolev metric in Eq. (2) can be linked to observable physical deformations (shearing, bending, etc) which helps with intuitively comparing physical objects. Last, the metric in Eq. (2) yields a distance that is rotation and reparameterization invariant [9]. In other words, if all the surfaces in the dataset are rotated and reparameterized in the same way, i.e., using the same rotation matrix and reparameterization diffeomorphism ϕ , then their pairwise distances are

unchanged. We note that this property is practical only if we first assume that all the surfaces (are oriented and) have valid point-to-point correspondences.

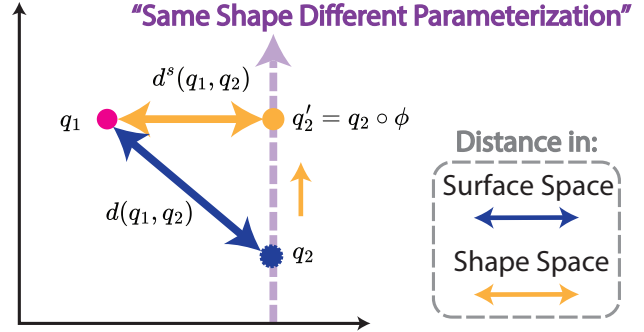


Figure 3. Distances in surface space vs shape space. q_1 and q_2 are two surfaces with different parameterization and different shape. The distance given by the second-order Sobolev metric 2 measures both the parameterization and shape differences. The shape space distance 3 only measures difference in shape.

D. Space of Surface Shapes In the space of surfaces \mathcal{I} , if two surfaces have the same shape but different orientations or parameterizations, they correspond to different points. By contrast, we introduce the space of surface shapes where two surfaces with the same shape correspond to the same point, regardless of differences in their orientation or parameterization. Mathematically, the space of surface shapes is defined as the quotient space: $\mathcal{I}/(\text{Rot}(\mathbb{R}^3) \times \text{Diff}(M))$ —see [9] for details. For simplicity, we consider the case of parameterizations with the shape space $\mathcal{S} = \mathcal{I}/\text{Diff}(M)$ while the case of orientations can be treated similarly.

In the shape space \mathcal{S} , the distance between two surface shapes q_1 and q_2 is given by:

$$d^{\mathcal{S}}(q_1, q_2) = \inf_{\phi} d(q_1, q_2 \circ \phi) = d(q_1, q_2'), \quad (3)$$

where ϕ represents a choice in parameterization. In Eq. (3), the parameterization of q_2 is varied until the second-order Sobolev distance d in Eq. (2) between q_1 and q_2 reaches an infimum as shown in Fig. 3. This operation matches the parameterization of q_2 to the parameterization of q_1 so that any remaining discrepancy between them is due to difference in shape, rather than difference in parameterization. Ideally, we would perform our geodesic regression methods directly in the shape space \mathcal{S} . However, the high computational cost of this approach leads us to instead compute in the surface space \mathcal{I} after choosing a reference parameterization that corresponds to the first hippocampal surface of our dataset.

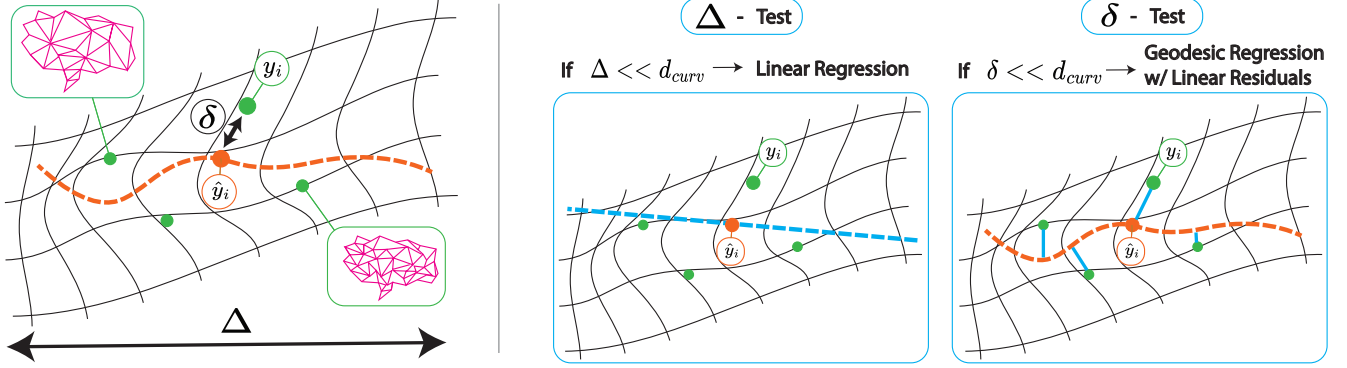


Figure 4. Overview: Approximation schemes for geodesic regression approximation. δ -test: if the residual magnitudes are small compared to the curvature of the manifold, we can use geodesic regression with linear residuals (GRLR). Δ -test: if the distance covered by the data set is small compared to the curvature, we use linear regression (LR). Reducing geodesic regression (GR) to either GRLR or LR provides up to four orders of magnitude speed-up, while sacrificing little accuracy.

4. Methods

We seek to quantify the *anatomical changes* in the hippocampal formation that emerge from progesterone variations during the menstrual cycle. To achieve this, we propose approximations to geodesic regression on the space of 3D brain shapes that make it computationally fast enough for practical use. We further propose rules of thumb for determining when each approximation can be used, as summarized in Fig. 4.

4.1. Linear Regression

Model *Linear regression* (LR) models the relationship between an independent variable $X \in \mathbb{R}$ and the dependent variable Y taking values in \mathbb{R}^D as:

$$Y = \alpha + X\beta + \epsilon, \quad (4)$$

where $\alpha \in \mathbb{R}^D$ is the intercept, $\beta \in \mathbb{R}^D$ is the slope, and ϵ represents the noise.

Loss Given data $(x_i, y_i) \in \mathbb{R} \times \mathbb{R}^D$, for $i = 1, \dots, n$, we fit the linear regression model through least squares, i.e., we compute the estimates for the intercept and slope $\hat{\alpha}, \hat{\beta}$ as:

$$(\hat{\alpha}, \hat{\beta}) = \arg \min_{(\alpha, \beta)} \frac{1}{2} \sum_{i=1}^n \|y_i - \hat{y}_i\|^2 \text{ for } \hat{y}_i = \alpha + x_i\beta, \quad (5)$$

which minimizes the summed squared magnitude of the (linear) *residuals*: $y_i - \alpha - x_i\beta$, for $i = 1, \dots, n$.

Learning Importantly for computational purposes, this minimization has an analytical solution given by the normal equations $\hat{\beta} = \frac{\frac{1}{n} \sum x_i y_i - \bar{x}\bar{y}}{\frac{1}{n} \sum x_i^2 - \bar{x}^2}$ and $\hat{\alpha} = \bar{y} - \bar{x}\hat{\beta}$, where \bar{x} and \bar{y} are the sample means of the x_i and y_i , respectively. We will use linear regression as our first baseline, where X is the level of progesterone, and Y is the hippocampal surface

discretized as a mesh, which takes values in $\mathbb{R}^{N \times 3}$ where N is the number of mesh vertices.

4.2. Geodesic Regression

Model *Geodesic regression* (GR) [6, 19] models the relationship between an independent variable $X \in \mathbb{R}$ and the dependent variable Y , whose values lie on a manifold \mathcal{N} , as:

$$Y = \text{Exp}(\text{Exp}(p, Xv), \epsilon), \quad (6)$$

where ϵ is noise in the tangent space at $\text{Exp}(p, Xv)$, and Exp is the operation defined in the previous section. Note that when the manifold of interest is $\mathcal{N} = \mathbb{R}^D$, the exponential operator simplifies to addition: $\text{Exp}(p, v) = p + v$. Consequently, the geodesic regression generative model simplifies to the linear regression generative model of Eq. (4) with $p = \alpha$ and $v = \beta$. We also note that the exponential operation appears twice: to model the geodesic itself, and to model the noise ϵ . In what follows, we consider geodesic regression on the manifold $\mathcal{N} = \mathcal{I}$ equipped with a second-order Sobolev metric from Eq. (2).

Loss Given data $(x_i, y_i) \in \mathbb{R} \times \mathcal{I}$, for $i = 1, \dots, n$, we seek to learn estimates of the intercept and slope $(p, v) \in \mathcal{I} \times T_p\mathcal{I}$. In the manifold setting, the loss function associated with the geodesic given by (p, v) is minimized as:

$$(\hat{p}, \hat{v}) = \arg \min_{(p, v)} \frac{1}{2} \sum_{i=1}^n d(y_i, \hat{y}_i)^2, \quad (7)$$

$$= \arg \min_{(p, v)} \frac{1}{2} \sum_{i=1}^n \|\text{Log}(\hat{y}_i, y_i)\|_{\hat{y}_i}^2, \quad (8)$$

$$\text{for } \hat{y}_i = \text{Exp}(p, x_i v). \quad (9)$$

We compute estimates for intercept and slope (\hat{p}, \hat{v}) which minimize the summed squared magnitude of the (geodesic) *residuals* $\text{Log}(\text{Exp}(p, x_i v), y_i)$ for $i = 1, \dots, n$. The

geodesic residuals differ from the linear residuals as they are calculated with exponentials and logarithms instead of additions and subtractions.

Learning In contrast to linear regression, the least squares problem of Eq. (7) above does not have an analytical solution for general manifolds \mathcal{I} . Instead, we need to compute the estimates of the intercept and slope with gradient descent, which is typically computationally expensive. Gradient descent comes in two flavors depending on the strategy used to compute the gradient, which can be either a Riemannian gradient as originally proposed in [6] or an extrinsic gradient. The Riemannian gradient writes [6]:

$$\begin{aligned}\nabla_p l &= - \sum_{i=1}^N d_p \text{Exp}(p, x_i v)^\dagger \epsilon_i, \\ \nabla_v l &= - \sum_{i=1}^N x_i d_v \text{Exp}(p, x_i v)^\dagger \epsilon_i,\end{aligned}$$

where l is the loss function, $\epsilon_i = \text{Log}(\text{Exp}(p, x_i v), y_i)$ are the residuals, d_v and d_p are derivatives, and \dagger denotes the adjoint. In the general case, the expression of these derivatives and their respective adjoint operators are not known although they can be derived analytically for some manifolds as in [19]. However, to the best of our knowledge, no such formula exists for shape spaces of parameterized surfaces, so we use a numerical approach.

4.3. Why is Geodesic Regression Slow?

The geodesic regression optimization is slow due to the Exp and Log maps in Eq. (7). Computation of the exponential and logarithm maps do not enjoy an analytical expression for the manifold that we are interested in, and neither do their differentials. Consequently, we compute them only numerically, as implemented in Geomstats [14] as follows.

For the *exponential map*, we consider the geodesic equation as a coupled system of first-order ODEs:

$$\begin{cases} \dot{v}(t) = \dot{\gamma}(t) \\ \dot{v}(t) = f(v(t), \gamma(t), t) \end{cases}$$

where f is a smooth function given by Eq. (1) and the state variable is $(\gamma(t), \dot{\gamma}(t))$. Given initial conditions, we use a first-order forward Euler scheme to integrate this system. For a given step dt we compute:

$$v(t+dt) = v(t) + \dot{v}(t)dt = v(t) + f(v(t), \gamma(t), t)dt. \quad (10)$$

Introducing the parameter n_{steps} , if we integrate this geodesic equation between $t = 0$ and $t = 1$ in n_{steps} then we use $dt = (n_{steps})^{-1}$. Consequently, the parameter n_{steps} controls the numerical precision of the computation of the exponential map. The computation of Exp is slow due to this numerical integration.

For the *logarithm map*, we solve the optimization problem in v :

$$\min d^2(\text{Exp}(p, v), q),$$

that represents the fact that Log is the inverse map of Exp. This minimization is solved by gradient descent (GD) until a convergence tolerance is reached. It uses scipy for minimization method and computes the gradient of the exponential map with automatic differentiation. The computation of Log is slow due to this optimization process.

4.4. Approximations Schemes with Rules of Thumb

Curved spaces are locally linear. Thus, if a data set falls on a “small” portion the shape space, addition and subtraction will offer excellent approximations of exponentials and logarithms and avoid costly computations. To speed up geodesic regression, we propose two approximation schemes shown in Fig. 4: (i) linear regression, and (ii) geodesic regression with linear residuals, where (ii) represents a novel approach for geometric machine learning—which we describe in the next section.

We also propose rules of thumb to determine when each of these methods will yield sufficiently accurate approximations of geodesic regression. For (i), we propose the Δ -Test, which explores when the magnitude of the geodesic length of the data set (Δ) is small at the scale of the curvature of the manifold (Fig. 4 middle). For (ii), we propose the δ -Test, which explores when the magnitude of the noise (δ) is small at the scale of the curvature of the manifold (Fig. 4 right). In the experiments section, we explore the curvature of the space of 3D discrete surfaces to give numerical values to these guidelines.

4.5. Geodesic Regression with Linear Residuals

Model We propose *geodesic regression with linear residuals* (GRLR) to model the relationship between an independent variable $X \in \mathbb{R}$, the noise-free dependent variable taking values in a manifold, and the (noisy) dependent variable Y taking values in \mathbb{R}^D . In other words, we propose the following generative model:

$$Y = \text{Exp}(p, Xv) + \epsilon, \quad (11)$$

where $\text{Exp}(p, Xv)$ is the noise-free dependent variable and ϵ is the noise. The noise-free dependent variable is constrained to be a surface in \mathcal{I} , and thus Y 's dependency on X is modelled using the Exp operation. However, in practical applications, the data's noise may push the data off of \mathcal{I} . Thus, in addition to its computational gain, this generative model acknowledges the fact that there is no reason for the noise to be constrained on the manifold.

Loss Given data $(x_i, y_i) \in \mathbb{R} \times \mathbb{R}^D$, for $i = 1, \dots, n$, we fit this regression model through least squares, i.e., we

compute the estimates for the intercept and slope as:

$$(\hat{p}, \hat{v}) = \arg \min_{(p,v)} \frac{1}{2} \sum_{i=1}^n \|\hat{y}_i - y_i\|^2 \text{ for } \hat{y}_i = \text{Exp}(p, x_i v), \quad (12)$$

where the squared *geodesic distance* of Eq. 7 has been replaced by the *squared Euclidean distance*, but we keep the exponential map defining the geodesic.

Learning Like geodesic regression, this least squares problem still requires gradient descent. The gradient can be computed as a Riemannian gradient or as an extrinsic gradient. The Riemannian gradient is given by:

$$\begin{aligned} \nabla_p l &= - \sum_{i=1}^N d_p \text{Exp}(p, x_i v)^\dagger \epsilon_i, \\ \nabla_v l &= - \sum_{i=1}^N x_i d_v \text{Exp}(p, x_i v)^\dagger \epsilon_i, \end{aligned}$$

where l is the loss function, $\epsilon_i = y_i - \text{Exp}(p, x_i v)$ are the residuals, d_v and d_p are derivatives, and \dagger is the adjoint. By avoiding the computation of n logarithms, the learning process enjoys a significant speed-up. Even within the use of the extrinsic gradient, the loss function avoids the computations of these logarithms and is thus accelerated. We note that we still need to compute the Exp by numerical integrations and their derivatives which we do by automatic differentiation. Our implementation is [publicly available on GitHub](#).

5. Experiments

We investigate the curvature of the space of surfaces to quantify which approximation scheme should be used on which dataset. Guided by this analysis, we approximate geodesic regression on 3D hippocampal surfaces, giving the first characterization of hippocampal formation’s shape change as a function of progesterone.

5.1. Curvature Estimation with δ - and Δ - Tests

Simulations We perform experiments on synthetic meshes to provide rules-of-thumb that help the practitioner decide when linear regression or geodesic regression with linear residuals can be used with little loss in accuracy on the space of discrete surfaces.

Specifically, our experiments explore cases when lines can be used to approximate geodesics. First, we compute both a line and a geodesic between two meshes q^{start} and q^{end} . Then, we compare the meshes along the line in $\mathbb{R}^{N \times 3}$ to the meshes along the geodesic in \mathcal{I} , where N is the number of vertices in the 3D meshes. We have either $n = 5$ or $n = 10$ meshes along each sequence. The start mesh q^{start}

is an ellipsoid whose principal axes have length 2, 2, and 3. The end mesh is a deformed version q^{end} of the reference q^{start} , where the amount of deformation is controlled by a factor that we vary in $\{1\%, 10\%, 50\%, 100\%\}$. This deformation factor indicates by how much each vertex in q^{end} has been moved compared to the vertex’s position in q^{start} and is given as a percentage of the diameter of q^{start} . q^{end} is generated by adding isotropic Gaussian noise to each vertex in q^{start} . In other words, the standard deviation of the Gaussian noise is: $\sigma = \text{deformation} \times D$ where D is the diameter of the mesh.

We determine how much the geodesic and the line between q^{start} and q^{end} differ by computing the root mean square deviation (RMSD) between the geodesic mesh sequence and the line mesh sequence, which we then normalize by the diameter D of the mesh:

$$\text{RMSD} = \frac{1}{D} \sqrt{\frac{1}{TN} \sum_{t=1}^T \sum_{j=1}^N \|v_{tj}^{\text{line}} - v_{tj}^{\text{geodesic}}\|^2}, \quad (13)$$

where N is the number of vertices, T the number of meshes in the sequence (5 or 10) and D the diameter. We also time the computation of the geodesic and the line and report their ratio.

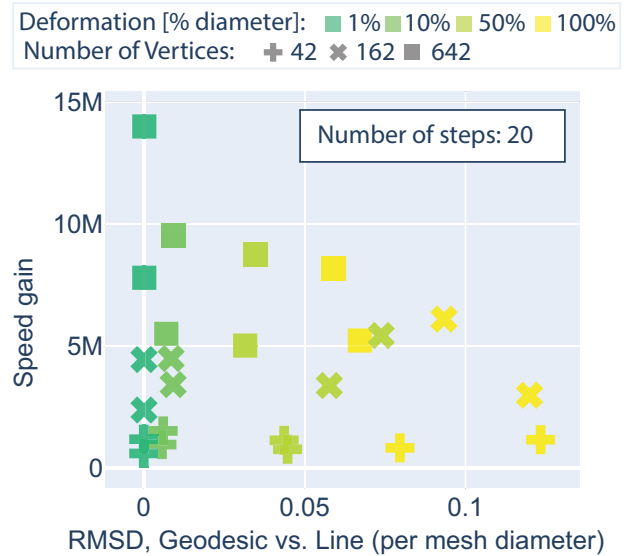


Figure 5. Accuracy-speed trade-off between a geodesic and its linear approximation in joining two meshes q^{start} and q^{end} . The x-axis (error) shows how their meshes differ by computing distances between their vertices. The y-axis (speed) shows the ratio of their computational times. The color represents the deformation factor, i.e. how deformed the mesh q^{end} is from the reference mesh q^{start} . The symbols represent the number of vertices in the meshes. The number of steps is a parameter controlling the numerical integration computing the geodesic. The results are similar when $n_{\text{steps}} = 5$.




Figure 6. Linear regression reveals hippocampal deformation associated with an increase in progesterone during the menstrual cycle. The two rows show two views of each 3D mesh along the sequence. For visualization purposes, a coloring of the mesh is used to represent depth. While volumetric analyses did not capture a volumetric change, our shape analysis reveals that an increase in progesterone corresponds to a shear deformation of the hippocampal formation.

δ -Test Fig. 5 shows that deformation factors of 1% yield errors below 0.05% of the diameter of the shape (below 0.0005 on the figure). This is true across two values of the number of steps used for the numerical computation of the exponential map: $n_{steps} = 20$, and $n_{steps} = 5$ (see Section). Consequently, for our δ -Test: when the measurement noise on the vertices of the meshed shapes is expected to be less than 1% of the total mesh diameter, we recommend using linear residuals instead of geodesic residuals. In this case, we assess that the shape manifold can be approximated as linear at the scale of residual length: the curvature is low compared to the magnitude of the noise. By using linear residuals, we enjoy a considerable speed-up, up to $14M \times$ (for a number of steps of 20), as shown in Fig. 5, and up to $1.5M \times$ (for a number of steps of 5).

As a real-world example, consider q^{start} as the mesh corresponding to a true hippocampal shape at a given level of progesterone, and q^{end} as the mesh that we observe in practice after segmenting and extracting the mesh from the MRI data. In this case, the “measurement noise” is the MRI noise and the segmentation error. We expect measurement noise to displace each vertex by around 1% of the total mesh diameter, since MRI images have a good resolution and both segmentation and meshing algorithms are reasonably accurate. Thus, brain MRIs are in the regime where geodesic regression can enjoy considerable speed-ups by utilizing linear residuals.

Δ -Test Additionally, Fig. 5 shows that deformation factors of 10-50% yield an error that is less than 10% of the diameter (below 0.1 on the figure). Consequently, for our Δ -Test: if the data set’s largest deformation between two meshes is 10-50% of the diameter of the mesh, and if practitioners can tolerate a maximum loss of accuracy of 10% in their results, then using linear regression instead of geodesic regression allows them to significantly speed up their pipeline. In this case, we assess that the shape manifold can be approximated as linear on the scale of the data

set: the curvature is low compared to the spread of the data. Even more strikingly, the error decreases from 10% to only 3% if practitioners consider meshes with hundreds of vertices, as shown in Fig. 5.

5.2. Hippocampal Shape Change Characterization

Data set We use a time-series of 3D brain images recorded from magnetic resonance imaging (MRI): 11 images from 11 consecutive days, capturing the progesterone peak of a single female subject’s natural menstrual cycle [15], as analyzed by volumetric analyses in [18]. We choose to focus on the progesterone peak (11 days), as opposed to the full menstrual cycle (30 days) for simplicity. The female also measured hormone levels in her blood in conjunction with each MRI session.

Pre-processing We align the 3D images to correct for the position and orientation of the subject’s head in the MRI scanner, and to extract the surface of each substructure of the hippocampal formation. The results of this pre-processing are shown in Figure 1 where each sub-structure surface is color-coded and shown at two different levels of progesterone (low and high). Here, we can visually observe the hippocampal formation’s shape evolution, which our analyses will seek to characterize with the slope and intercept learned from the proposed approximation of geodesic regression. We then use Eq. (3) to give every hippocampal mesh the same parameterization.

Characterization Consider q^{start} , q^{end} the meshes corresponding to hippocampus shapes at the lowest and highest progesterone levels respectively. Fig. 1 shows that there is a deformation factor of around 10%: each vertex is displaced by around 10% of the total mesh diameter between the two meshes shown. Thus, following Δ -Test, we use linear regression to provide a characterization of 3D shape changes in the hippocampal formation during the menstrual cycle. Fig. 6 reveals for the first time that the hippocampal forma-

tion shears in response to an increase in progesterone. Additionally, this computation provides an important educational tool. Clinical neuroscientists can use the result of our regression model to query, for a given progesterone level, what is the associated hippocampal shape.

6. Conclusion

We have proposed a shape analysis technique to reveal what volumetric analyses could not: that the overall shape of hippocampal formation changes during progesterone level fluctuation. The implications for women's health are profound. Because each structure of the brain is dedicated to a specific function, and the hippocampal formation is directly related to functions that deteriorate in women after menopause, characterizing how the hippocampal formation changes in response to sex hormones changes is critical. Not only does it provide a diagnostic for disease prediction, but it also offers a method to probe relationships between hormone level, hippocampal shape, and brain health.

Here, we provide a practical method to characterize such changes with slopes and intercepts learned through approximation schemes for geodesic regression on the space of 3D discrete surfaces. This work aims to open research avenues for automated, fast, and statistically sound diagnostics of female brain health.

References

- [1] Martin Bauer, Nicolas Charon, Philipp Harms, and Hsi-Wei Hsieh. A numerical framework for elastic surface matching, comparison, and interpolation. *International Journal of Computer Vision*, 129(8):2425–2444, 2021. [2](#), [3](#)
- [2] Christopher R Beam, Cody Kaneshiro, Jung Yun Jang, Chandra A Reynolds, Nancy L Pedersen, and Margaret Gatz. Differences between women and men in incidence rates of dementia and alzheimer's disease. *Journal of Alzheimer's disease*, 64(4):1077–1083, 2018. [1](#)
- [3] 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*, 61:139–157, 2005. [2](#)
- [4] Janine A Clayton and Francis S Collins. Policy: Nih to balance sex in cell and animal studies. *Nature*, 509(7500):282–283, 2014. [1](#)
- [5] AT Du, Nea Schuff, D Amend, MP Laakso, YY Hsu, WJ Jagust, K Yaffe, JH Kramer, B Reed, D Norman, et al. Magnetic resonance imaging of the entorhinal cortex and hippocampus in mild cognitive impairment and alzheimer's disease. *Journal of Neurology, Neurosurgery & Psychiatry*, 71(4):441–447, 2001. [1](#)
- [6] Thomas Fletcher. Geodesic regression on riemannian manifolds. In *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy-Geometrical and Statistical Methods for Modelling Biological Shape Variability*, pages 75–86, 2011. [2](#), [5](#), [6](#)
- [7] Liisa AM Galea, Steven R Wainwright, MM Roes, P Duarte-Guterman, C Chow, and DK Hamson. Sex, hormones and neurogenesis in the hippocampus: hormonal modulation of neurogenesis and potential functional implications. *Journal of neuroendocrinology*, 25(11):1039–1061, 2013. [1](#)
- [8] Nicolas Guigui, Nina Miolane, Xavier Pennec, et al. Introduction to riemannian geometry and geometric statistics: from basic theory to implementation with geomstats. *Foundations and Trends® in Machine Learning*, 16(3):329–493, 2023. [3](#)
- [9] Emmanuel Hartman, Yashil Sukurdeep, Eric Klassen, Nicolas Charon, and Martin Bauer. Elastic shape analysis of surfaces with second-order sobolev metrics: a comprehensive numerical framework. *International Journal of Computer Vision*, 131(5):1183–1209, 2023. [2](#), [3](#), [4](#)
- [10] Ian H Jermyn, Sebastian Kurtek, Hamid Laga, Anuj Srivastava, Gerard Medioni, and Sven Dickinson. *Elastic shape analysis of three-dimensional objects*. Springer, 2017. [2](#), [3](#)
- [11] Martin Kilian, Niloy J Mitra, and Helmut Pottmann. Geometric modeling in shape space. In *ACM SIGGRAPH 2007 papers*, pages 64–es. 2007. [3](#)
- [12] Sebastian Kurtek, Eric Klassen, Zhaohua Ding, and Anuj Srivastava. A novel riemannian framework for shape analysis of 3d objects. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 1625–1632. IEEE, 2010. [2](#), [3](#)
- [13] Nasim Maleki, Lino Becerra, Jennifer Brawn, Bruce McEwen, Rami Burstein, and David Borsook. Common hippocampal structural and functional changes in migraine. *Brain Structure and Function*, 218:903–912, 2013. [1](#)
- [14] Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, Hatem Hajri, Yann Cabanes, Thomas Gerald, Paul Chauchat, Christian Shewmake, Daniel Brooks, Bernhard Kainz, Claire Donnat, Susan Holmes, and Xavier Pennec. Geomstats: A Python package for Riemannian geometry in machine learning. *Journal of Machine Learning Research*, 21:1–9, 2020. [6](#)
- [15] Laura Pritschet, Tyler Santander, Caitlin M Taylor, Evan Layher, Shuying Yu, Michael B Miller, Scott T Grafton, and Emily G Jacobs. Functional reorganization of brain networks across the human menstrual cycle. *Neuroimage*, 220:117091, 2020. [1](#), [8](#)
- [16] Lisa R Taxier, Kellie S Gross, and Karyn M Frick. Oestradiol as a neuromodulator of learning and memory. *Nature Reviews Neuroscience*, 21(10):535–550, 2020. [1](#)
- [17] Caitlin M Taylor, Laura Pritschet, and Emily G Jacobs. The scientific body of knowledge—whose body does it serve? a spotlight on oral contraceptives and women's health factors in neuroimaging. *Frontiers in neuroendocrinology*, 60:100874, 2021. [1](#)
- [18] Caitlin M Taylor, Laura Pritschet, Rosanna K Olsen, Evan Layher, Tyler Santander, Scott T Grafton, and Emily G Jacobs. Progesterone shapes medial temporal lobe volume across the human menstrual cycle. *NeuroImage*, 220:117125, 2020. [1](#), [8](#)

- [19] P Thomas Fletcher. Geodesic regression and the theory of least squares on riemannian manifolds. *International journal of computer vision*, 105:171–185, 2013. [2](#), [5](#), [6](#)
- [20] Kjersti Grøtta Vetvik and E Anne MacGregor. Sex differences in the epidemiology, clinical features, and pathophysiology of migraine. *The Lancet Neurology*, 16(1):76–87, 2017. [1](#)