# Score Matching Riemannian Diffusion Means

**Frederik Möbius Rygaard**
Technical University of Denmark
fmry@dtu.dk

**Steen Markvorsen**
Technical University of Denmark
stema@dtu.dk

**Søren Hauberg**
Technical University of Denmark
sohau@dtu.dk

**Stefan Sommer**
University of Copenhagen
sommer@di.ku.dk

## ABSTRACT

Estimating means on Riemannian manifolds is generally computationally expensive because the Riemannian distance function is not known in closed-form for most manifolds. To overcome this, we show that Riemannian diffusion means can be efficiently estimated using score matching with the gradient of Brownian motion transition densities using the same principle as in Riemannian diffusion models. Empirically, we show that this is more efficient than Monte Carlo simulation while retaining accuracy and is also applicable to learned manifolds. Our method, furthermore, extends to computing the Fréchet mean and the logarithmic map for general Riemannian manifolds. We illustrate the applicability of the estimation of diffusion mean by efficiently extending Euclidean algorithms to general Riemannian manifolds with a Riemannian $k$-means algorithm and maximum likelihood Riemannian regression.

***Keywords*** Riemannian Manifolds · Diffusion t-means · Intrinsic Statistics · Score Matching

## 1 Introduction

The *average* of a dataset is arguably one of the most essential statistical quantities of interest, and it forms a building block for more elaborate models. For data residing on a Riemannian manifold, the most common definition of an average is the *Fréchet mean*, which minimizes the sum of squared geodesic distances to data [Fréchet, 1948],

$$\mu^{\text{Fréchet}}(x_{1:N}) := \arg\min_{\mu \in \mathcal{M}} \sum_{n=1}^{N} \text{dist}^2(x_n, \mu), \quad (1)$$

where $\text{dist}(\cdot, \cdot)$ is the distance along the manifold, $\mathcal{M}$.



Figure 1: The diffusion mean is estimated by learning the gradients of the log-transition density. (1) Sample points from a Riemannian Brownian motion are (2) passed through a neural network minimizing the score matching loss and (3) outputs the gradients used to estimate the diffusion $t$-mean.

The Fréchet mean is not the only generalization of the Euclidean mean value to non-Euclidean domains, and it has properties, which makes it worthwhile to consider alternatives: *Computationally*, the Fréchet mean is expensive when the manifold distance function is not known in closed-form, which is the case for all but only a selected set of manifolds. *Theoretically*, the Fréchet mean is only a maximum likelihood estimate when the underlying manifold is restricted to being *symmetric* [Fletcher, 2013], which is a very restrictive assumption satisfied by the simplest of manifolds, and the sample Fréchet mean can exhibit slow convergence to the Fréchet mean of the data distribution [Eltzner et al., 2022]. These properties have recently motivated the construction of the *diffusion $t$-mean* [Eltzner et al., 2022]. The *diffusion $t$-mean* is the maximum likelihood estimator for the transition density of a Brownian motion on the manifold [Eltzner et al., 2022],

$$\mu_t^{\text{Diffusion}}(x_{1:N}) := \arg\max_{\mu \in \mathcal{M}} \sum_{n=1}^{N} \log p_t(x_n, \mu), \quad (2)$$

Figure 2: Estimated values of the diffusion mean using score matching for two different manifolds. The two left most plots show the diffusion mean (blue) for a shape space with $30 \times 2$ landmarks on butterfly and cardiac data (black), respectively. The two right most plots shows a learned manifold using a Gaussian process $\mathcal{GP}$ with corresponding expected metric for a rotated MNIST image. The first plot shows the reconstructed rotated MNIST data for half a rotation, while the right most plot shows the estimated diffusion mean for half a rotation.

where $p$ is the transition density with diffusion time $t$. The diffusion mean is per definition a maximum likelihood estimate, and, in many situations, the sample diffusion mean avoids the slow convergence to the distribution mean. However, the construction comes with new computational difficulties. For general Riemannian manifolds, the transition density of a Brownian motion corresponds to the minimal heat kernel for which analytic expressions are rarely available. This implies that evaluating the log-likelihood objective function is intractable. Sommer et al. [2017] suggest using *bridge sampling* to approximate the log-likelihood, but such approximations are computationally expensive.

**In this paper**, we show that *Riemannian score matching* [Bortoli et al., 2022] can be used to efficiently estimate the diffusion t-mean for general manifolds that scale linearly with the manifold dimension. This is a vast improvement over previous approaches with cubic complexity. We further show that the Fréchet mean (1) can be estimated using the limit of the scores. Empirically, we verify that our method is more efficient than Monte Carlo bridge sampling, which is the only available alternative for general manifolds and apply it to learned manifolds. We further illustrate how the efficient estimation of the diffusion mean allows us to perform Riemannian $k$-means clustering and maximum likelihood Riemannian regression. In general, the presented algorithms provide a practical approach to compute statistics on general Riemannian manifolds efficiently.

## 2 Background and Related Work

### 2.1 Riemannian Geometry and Notation

**A Riemannian manifold**, $(\mathcal{M}, g)$, is a differentiable manifold equipped with a Riemannian metric, $g$, which defines a smoothly varying inner product, $\langle \cdot, \cdot \rangle_x$ on each tangent space, $T_x \mathcal{M}$ for $x \in \mathcal{M}$. The tangent space, $T_x \mathcal{M}$, is a vector space consisting of all tangent vectors for some curve $\gamma : (-\epsilon, \epsilon) \to \mathcal{M}$ with $\gamma(0) = x \in \mathcal{M}$ [do Carmo, 1992]. The tangent bundle, $T\mathcal{M}$ is the disjoint union of the tangent spaces, i.e. $T\mathcal{M} = \bigsqcup_{x \mathcal{M}} T_x \mathcal{M}$. Some simple examples of Riemannian manifolds are the $m$-sphere, $\mathbb{S}^m = \{ x \in \mathbb{R}^{m+1} \mid \|x\| = 1 \}$ with metric inherited from the ambient Euclidean space; the torus $\mathbb{T}^m$, which is the direct product of $m$ copies of $\mathbb{S}^1$ equipped with e.g. a flat metric, and also, of course, Euclidean space itself, $\mathbb{R}^m$.

**A geodesic** is a curve, $\gamma : U \subseteq \mathbb{R}_+ \to \mathcal{M}$ that locally minimizes the curve length, $\mathcal{L}(\gamma) = \int_0^1 \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\gamma(t)} \, \mathrm{d}t$. The distance on a Riemannian manifold is defined as the length of the geodesic connecting two points, i.e. $\mathrm{dist}(x, y) = \min_\gamma \mathcal{L}(\gamma)$ for $x, y \in \mathcal{M}$. We will assume that the manifolds is geodesically complete in the sense that between any pair of points, $x, y \in \mathcal{M}$, there exists at least one length-minimizing geodesic, $\gamma$, with $\gamma(0) = x$ and $\gamma(1) = y$. In practice, a geodesic can be found by either minimizing the energy functional, $\mathcal{E}(\gamma) = \frac{1}{2} \int_0^1 \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\gamma(t)} \, \mathrm{d}t$ or by solving the following, usually nonlinear, ODE

$$\frac{\mathrm{d}x^k}{\mathrm{d}t^2} + \Gamma_{ij}^k \frac{\mathrm{d}x^i}{\mathrm{d}t} \frac{\mathrm{d}x^j}{\mathrm{d}t} = 0, \tag{3}$$

written in a local coordinate chart in Einstein notation, where $\Gamma_{ij}^k$ denotes the Christoffel symbols. The exponential map, $\mathrm{Exp}_x : \mathcal{T}_x \mathcal{M} \to \mathcal{M}$, is given by $\mathrm{Exp}_x(v) = \gamma(1)$, where $\gamma$ is a geodesic with $\gamma(0) = x$ and $\dot{\gamma} = v$. It can be shown that the exponential map is a diffeomorphism in a star-shaped neighborhood, $\mathcal{D}(x) \subseteq T_x \mathcal{M}$ [do Carmo, 1992]. Since $\mathcal{M}$ is assumed geodesically complete, the curve $\gamma_x(t) = \mathrm{Exp}_x(tv)$ is either a length-minimizing up to a point $t_0$ or is length minimizing for all $t \in [0, \infty[$ [Pennec, 2006]. In the first case $t_0$ is called a *cut point*, and the set of all cut
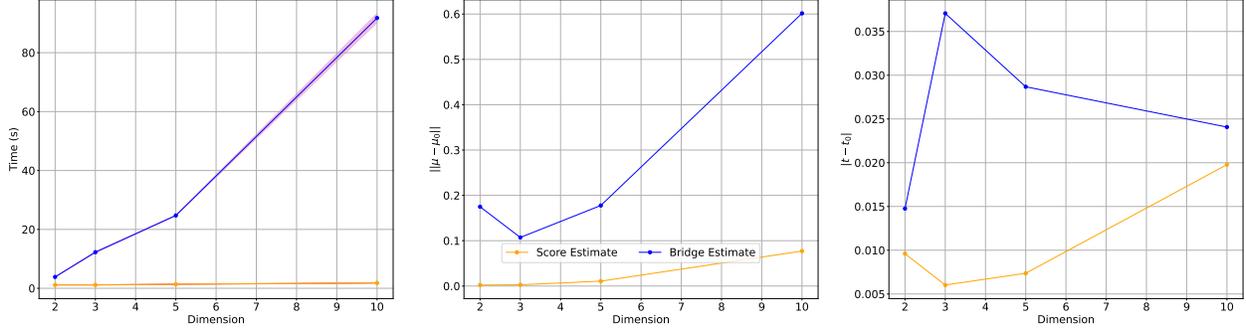
Figure 3: Run-Time and error comparison for $\mathbb{S}^d$ with $d = 2, 3, 5, 10$ between score based gradient descent (yellow) and bridge sampling (blue) to estimate the diffusion mean. The left most plot shows the run-time over dimension, while the center and right plot show the error of the diffusion mean and variance estimate, respectively. The estimates are based on $1,000$ samples of Brownian motion at the north pole for diffusion time $t = 0.5$.

points, $\mathcal{C}(x)$, is called the *cut locus*. Within the cut locus, the inverse of the exponential map is given by the logarithmic map, $\text{Log}_x : \mathcal{M} \setminus \mathcal{C}(x) \to T_x\mathcal{M}$. A measure on a Riemannian manifold, $(\mathcal{M}, g)$, is given by infinitesimal volume on each tangent space $d\mathcal{M}(x) = \sqrt{|g(x)|}\, dx$, where $|g(x)|$ denotes the determinant of the local representation of the metric. The Riemannian divergence, $\text{div}$, is defined in a local coordinates system as $\text{div}(V) = \frac{\partial V^m}{\partial x^m} + V^k \frac{\partial \log \sqrt{\log |g|}}{\partial x^k}$ for a vector field, $V : \mathcal{M} \to T\mathcal{M}$.

**Riemannian Brownian motion** is a stochastic process on $\mathcal{M}$ that, when started at $x \in \mathcal{M}$ has the minimal heat kernel $p_t(x, \cdot)$ as transition density for $t \in \mathbb{R}_+$. The heat kernel is the solution to the PDE

$$\partial_t p_t(x, y) = \frac{1}{2}\Delta_x p_t(x, y), \tag{4}$$

where $\Delta$ denotes the Laplace-Beltrami operator. Note that there are other equivalent ways of defining Riemannian Brownian motion, see e.g. Hsu [2002] for details. We will assume that the manifold $\mathcal{M}$ is stochastic complete such that there exists a solution to eq. 4, and that $\mathcal{M}$ has bounded sectional curvature to ensure the existence of a bounded minimal heat kernel $p$ [Eltzner et al., 2022]. Further, we will assume that $\mathcal{M}$ is orientable and boundaryless without any singularities or boundary points in accordance with Bortoli et al. [2022] in order to sample paths of Brownian motions.

## 2.2 Means on Riemannian Manifolds

**The Fréchet mean** minimizes the expected squared distances with respect to the Riemannian metric [Fréchet, 1948] generalizing the Euclidean mean, which minimizes the expected squared distances with respect to the Euclidean norm.

$$\mu^{\text{Fréchet}}(X) := \underset{y \in \mathcal{M}}{\arg\min}\, \mathbb{E}\left[\text{dist}^2(X, y)\right], \tag{5}$$

for some random variable, $X$, on $\mathcal{M}$ with the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Unlike the Euclidean mean, the Fréchet mean is not necessarily unique. Depending on the curvature of $\mathcal{M}$, $\mu^{\text{Fréchet}}$ may consist of multiple points and thus being a *mean set* [Karcher, 1977]. For a given set of observations $\{x_i\}_{i=1}^n \subset \mathcal{M}$, the sample estimator of the Fréchet mean is given by eq. 1. The Fréchet mean can be numerically estimated using a gradient descent approach with gradient $\nabla_y \mathbb{E}\left[\text{dist}^2(x, y)\right] = \mathbb{E}\left[-2\text{Log}_y(x)\right]$ [Pennec, 2006]. With this approach, computing the sample Fréchet mean requires estimating the logarithmic map, which is numerically expensive, since it involves minimizing either the energy functional or solving the ODE (3) as a boundary value problem. Other approaches to estimate the Fréchet mean without using logarithmic maps have been investigated, where Lou et al. [2021] investigates using direct differentiation of the objective function focusing on Hyperbolic spaces. However, these methods are computationally expensive or rely on closed-form expressions on the manifold. **The diffusion $t$-mean** is inspired by the Euclidean maximum likelihood estimation. Let $p$ denote the minimal heat kernel solving eq. 4 for a given Riemannian manifold, $(\mathcal{M}, g)$, then the diffusion mean is defined as [Eltzner et al., 2022]

$$\mu_t^{\text{Diffusion}}(X) = \underset{y \in \mathcal{M}}{\arg\max}\, \mathbb{E}\left[\log p_t(X, y)\right], \tag{6}$$

Note that $\mu_t^{\text{Diffusion}}$ can in general be a set similar to the Fréchet mean and is denoted as the *diffusion $t$-mean set* [Eltzner et al., 2022]. In Euclidean space, the distribution of the Brownian motion at time $t$ is a normal distribution with variance $t$. The diffusion mean can thus be interpreted as the center of a non-Euclidean normal distribution fitted to the data.

3
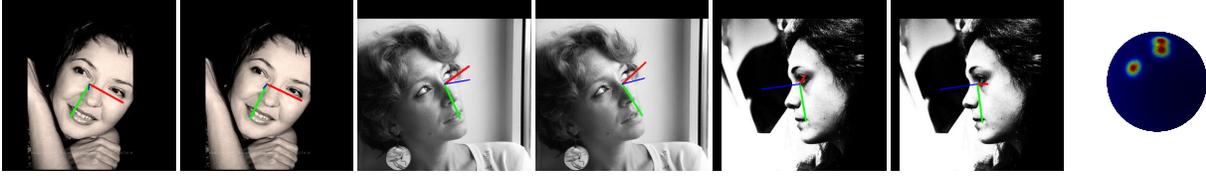
Figure 4: The estimated head-pose position for AFLW-2000 dataset, where the arrows correspond to the head pose position. The left image is the true labeled direction, while the right image is the prediction. Using only the roll and yaw of the Euler angles, we plot the mixture of the maximum likelihood neural regression on $\mathbb{S}^2$ using the estimated un-normalized heat kernel on $\mathbb{S}^2$.

The time $t$ can be fitted to the data as well by minimizing the negative log-likelihood with respect to $t$. This optimal $t$ is called the *diffusion variance*. A sample estimate of the log-likelihood is given in eq. 2. The diffusion mean can be estimated using Monte Carlo simulation [Sommer et al., 2017], which uses samples from Brownian Bridges computed as a guided process [Delyon and Hu, 2006, Papaspiliopoulos and Roberts, 2012]. Further improvements on the Monte Carlo simulation have been proposed by Bui et al. [2022] and Jensen and Sommer [2022]. However, estimating the likelihood using Monte Carlo simulation requires computing the Logarithmic map making it numerically expensive, when there is no closed-form expression available. Unlike this, our method does not rely on closed-form expressions and is applicable to any dataset once the scores have been trained for a fixed manifold independent of the observations.

**Why use diffusion means?** Both the Fréchet mean and diffusion mean are intrinsic means, which take the underlying manifold structure into account rather than using extrinsic methods that can be computationally difficult to project back to the manifold and be less accurate. The diffusion mean and Fréchet mean can be seen as generalizations of the usual Euclidean mean in the sense that both means coincide with the mean in Euclidean space. In terms of the convergence of the estimators for $N$ sample points, $\mu_N^{\text{Fréchet}}$ and $\mu_{t,N}^{\text{Diffusion}}$, Eltzner et al. [2022] shows that the estimator of the diffusion mean for a fixed $t > 0$ is strongly consistent in the sense of Ziezold [1977], i.e. $\cap_{n=1}^{\infty}\overline{\cup_{k=n}^{\infty}\mu_{t,k}^{\text{Diffusion}}} \subset \mu_t^{\text{Diffusion}}$ for almost all $\omega$ and in the sense of Bhattacharya and Patrangenaru [2003], i.e.



Figure 5: The estimated Fréchet mean for DTI tensor images corresponding to elements in $\mathcal{P}(3)$ using the time limit of the score. The two leftmost plots show the the Frêchet mean estimated using scores and with the Logarithmic map, respectively, while the right shows the two leading eigenvalues of the observations (black) as well as the approximation using scores (red) and ground truth (blue).

if $\mathcal{M} \neq \emptyset$ and if for almost all $\omega$ and for all $\epsilon > 0$ there exists an $N_{\epsilon,\omega}$ such that $\cup_{k=N}^{\infty}\mu_{t,k}^{\text{Diffusion}}(\omega) \subset B(\mathcal{M},\epsilon)$, where $B(M,\epsilon)$ is the union of all balls with radius $\epsilon$ around all points on $\mathcal{M}$. The estimator for the Fréchet mean is similarly a strongly consistent estimator with respect to the above conditions and satisfies them for the same conditions as the diffusion mean. In the time limit $t \to 0$, $\mu_{t,N}^{\text{Diffusion}} \to \mu_N^{\text{Fréchet}}$ [Eltzner et al., 2022]. The Fréchet mean can be 'smeary' in the sense that the sample Fréchet mean can exhibit slower convergence to the distribution mean than the standard $\sqrt{N}$ rate in the Euclidean central limit theorem. When the diffusion mean and the diffusion variance are estimated simultaneously, the diffusion has less tendency to exhibit smeariness (it can be smeary in some but not all directions) thus requiring fewer samples for a given precision of the sample estimate [Eltzner et al., 2022].

**Computational Cost** The cost of estimating Fréchet means and diffusion means is numerically inexpensive, if the distances or heat kernel, respectively, have closed-form for the given Riemannian manifold. However, this is the case only known for a selected set of Riemannian manifolds. In general, estimating distances on Riemannian manifolds involves either solving the ODE (3) as a boundary value problem or minimizing the energy functional, both of which are computationally expensive. The heat kernel, and thus diffusion means, can be computed using Monte Carlo estimation of the heat kernel with diffusion bridges [Sommer et al., 2017]. The computational complexity of this depends on the precision of the Monte Carlo estimate and is thus not directly comparable to the case for the Fréchet mean. In practice, the two approaches are often similarly expensive. The approach proposed in this paper substantially improves this situation making it computationally feasible to compute statistics on general Riemannian manifolds for both the diffusion and Fréchet mean.

# 3 Estimating of Diffusion t-Means using Score Matching

**Estimating the diffusion $t$-mean.** In practice the diffusion $t$-mean in eq. 6 can be estimated by gradient descent, where the gradient is given by $\nabla_y \mathbb{E}\left[-\log p_t(X, y)\right] = \mathbb{E}\left[-\nabla_y \log p_t(X, y)\right]$. Similarly, we find the diffusion variance by gradient descent using $\partial_t \mathbb{E}\left[-\log p_t(X, y)\right] = \mathbb{E}\left[-\partial_t \log p_t(X, y)\right]$. The optimization scheme is outlined in algorithm 1. However, for general Riemannian manifolds the gradient is not available in closed-form expression. We will instead approximate the gradient using the score of Riemannian Brownian motion.

---
**Algorithm 1:** Estimating Diffusion t-Mean

---
**Input**: $N_{\text{iter}}, x_{1:N}, t^{\text{init}}, \mu^{\text{init}}, \alpha$
**Output**: Diffusion mean and diffusion variance
Learn neural network $s_1$ minimizing eq. 9
Set $\mu = \mu^{\text{init}}, t = t^{\text{init}}$
**for** $i = 0, \ldots, N_{iter}$ **do**
  $\quad \mu \leftarrow \text{Exp}_\mu\left(-\alpha \frac{1}{N} \sum_i s_1(x_i, \mu, t; \theta)\right)$
  $\quad t \leftarrow -\alpha \frac{1}{N} \sum_i \frac{\mathrm{d}}{\mathrm{d}t} f(x_i, \mu, t; \theta)$ with $f$ given by
  $\quad \quad$ eq. 10
**end**
return $\mu, t$

---

**Estimating the score.** Consider a Riemannian Brownian Motion, $W_t^{\mathcal{M}}$, with transition density $p_t(x, y)$. We aim to train a neural network, $s_1 : \mathcal{M}^2 \times \mathbb{R}_+ \to T\mathcal{M}$ with parameters $\theta$ to approximate $\nabla_y \log p_t(x, y)$, by minimizing

$$\mathbb{E}_{\mathbb{P}_t(x,y)}\left[||\nabla_y \log p_t(x, y) - s_1(x, y, t; \theta)||_2^2\right], \tag{7}$$

where $\mathbb{P}_t(x, y)$ denotes the law of Brownian motion on $\mathcal{M}$. Bortoli et al. [2022] shows that $\ell_t$ can be minimized implicitly without evaluating $\log p_t(x, y)$ using the loss function

$$\mathbb{E}_{\mathbb{P}_t(x,y)}\left[\|s_1(x, y, t; \theta)\|_2^2 + 2\text{div}(s_1(x, \cdot, t; \theta))(y)\right], \tag{8}$$

From eq. 8 it can be seen that we can minimize $\ell_t$ without knowing $\log p_t(x, y)$. To avoid evaluating the computational expensive divergence in eq. 8, we will estimate the score by denoising score matching (DSM) [Vincent, 2011]

$$\frac{1}{2}\mathbb{E}_{\mathbb{P}_t(x,y)}\left[\left\|s_1(x, y, t; \theta) + \frac{z}{\sigma}\right\|_2^2\right], \tag{9}$$

where $\sigma^2$ denotes the variance of the infinitesimal time step of Brownian motion and $z := \frac{x-y}{\sigma}$. Note that eq. 9 holds in Euclidean cases. However, since Brownian motion on Riemannian manifolds locally behaves like Euclidean Brownian motion [Hsu, 2002], we can use eq. 9 to estimate the score as long as the time step if sufficiently small. Note that there are other loss functions such as sliced score matching [Song et al., 2019] and variance reducing method [Wang et al., 2020, Meng et al., 2021] that can similarly be applied. Estimating the above loss function requires only sampling paths of Riemannian Brownian motion, which can be done in either local coordinates [Hsu, 2002] or using geodesic random walk in the tangent space [Bradbury et al., 2018]. We refer to appendix A for details on the sampling method.

**Estimating the time derivative** Using the score above we are able to estimate the diffusion mean by gradient descent. To also estimate the diffusion variance, we use the following result in the proof of theorem 4.8 in Eltzner et al. [2022]

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}\left[\log p_t(X, y)\right] = \mathbb{E}\left[\frac{1}{2}\left(\Delta_y \log p_t(X, y) + ||\nabla_y \log p_t(X, y)||^2\right)\right] \tag{10}$$

where $\Delta_y \log p_t(x, y) = g^{jk}\frac{\partial^2 \log p}{\partial y^j \partial y^k} - g^{jk}\Gamma_{jk}^l\frac{\partial \log p}{\partial y^l}$. Thus, the time gradient can be estimated using the score and the trace of second order score. [Meng et al., 2021] proposes methods to estimate higher-order scores, but we found empirically that it was more stable using automatic differentiation of the score to estimate the trace of the Hessian of the heat kernel.

**Complexity in the manifold dimension.** Numerically integrating geodesics using eq. 3 requires inversion of the metric tensor in the Christoffel symbols, which scales cubically in the manifold's dimension. Similarly, sampling Brownian motions requires finding the square root of the inverse metric, again with cubic scaling in the manifold dimension. Thus, finding the Fréchet mean via iterative optimization and diffusion means via bridge sampling has cubic scaling. For the proposed scheme, the cubic scaling persists for the training phase, but, for fixed network layout, inference time complexity scales only linearly in the input dimension.

**Implementation.** In algorithm 1, the step on the Riemannian manifold is done using the exponential map. In practice, the step can also be done without the exponential map in the chart of the Riemannian manifold. This can be more efficient, especially for learned Riemannian manifolds. Further, it should be noted that the estimation of the diffusion $t$-mean using gradient descent will have local convergence, potentially not finding the entire set of diffusion $t$-means. We will restrict the diffusion variance $t \in (0, 1)$ corresponding to the values of $t$ used for learning the score and time derivative. We note also that since the score is an approximation of the gradient of the log heat kernel, then convergence is not necessarily guaranteed. However, there exists well-known bounds on the score [Bortoli et al., 2022, Song et al., 2021], we refer



Figure 6: Estimates of the diffusion t-mean for $\mathbb{S}^2$ (first row) and $2 \times 2$ landmark shape space (second row) for synthetic data corresponding to $1,000$ samples of Brownian motion at time $t = 0.5$. The first column shows the diffusion mean, while the second column shows the estimated diffusion variance.

to these for details. We estimate the bridge estimate of the diffusion mean in accordance with Jensen and Sommer [2023], which requires evaluating the logarithmic map. We note that there other method such as Corstanje et al. [2024], which can be used if there is a closed-form expression of the heat kernel on a comparison manifold.

## 4 Experiments and Applications

The following section contains experiments and results regarding the estimation of the diffusion t-mean. Further we show its applicability to estimating the Fréchet mean as well as statistical algorithms on Riemannian manifolds. A full description of the manifolds used in the section can be found in appendix B. Details on implementation and training is described in appendix E.

**Estimating the Diffusion t-Mean.** To illustrate the applicability of using scores we compute the diffusion mean of a high-dimensional $30 \times 2$ landmarks shape space [Arnaudon et al., 2023]. We compute the diffusion mean on landmarks of butterfly wings as well on cardiac data, where bridge sampling is too expensive to compute. We also compute the diffusion mean on the expected metric for a learned manifold using a Gaussian process ($\mathcal{GP}$) [Tosi et al., 2014] applied to a rotated MNIST image [Deng, 2012], where no closed form expressions are available. The estimated diffusion means can be seen in Fig. 2, where we see that the estimated diffusion means seem reasonable given the observed data. Table 1 shows the estimated diffusion mean using scores compared to bridge sampling for different manifolds and varying dimension, where we plot the result for $\mathbb{S}^d$ in Fig. 3. Each manifold has a sample of $1,000$ end points of Brownian motion with a fixed starting point and fixed diffusion time, $t = 0.5$. Fig. 6 shows the estimates and sampled data for $\mathbb{S}^2$ and $2 \times 2$ landmarks. When the logarithmic map is available in closed-form, we compute the diffusion mean using bridge sampling. If the heat kernel is not known, we compare the estimates to the starting point and diffusion variance of the sampled data. From table 1 we see that using the score the estimates of the diffusion t-mean is comparable to bridge sampling, while having significantly lower computational time.

**Estimating the Logarithmic map and Fréchet mean** As a side-effect of computing the score, we estimate the Logarithmic map as the time-limit of the score. Hsu [2002] shows that $\lim_{t \searrow 0} t \log p(x, y, t) = -\frac{\text{dist}^2(x,y)}{2}$ uniformly, which implies that

$$-\lim_{t \searrow 0} t \nabla_y \log p(x, y, t) = \text{Log}_y(x). \quad (11)$$

| Manifold | $\|\mu^{\text{score}} - \mu_0\|_2$ | $\|\mu^{\text{bridge}} - \mu_0\|_2$ | $\|t^{\text{score}} - t_0\|$ | $\|t^{\text{bridge}} - t_0\|$ | Score Time (s) | Bridge Time (s) |
|---|---|---|---|---|---|---|
| $\mathbb{R}^2$ | 0.0172 | **0.0168** | **0.0120** | 0.0406 | **0.3324 ± 0.0127** | 0.3867 ± 0.0036 |
| $\mathbb{R}^3$ | **0.0269** | 0.1027 | **0.0092** | 0.0603 | **0.3449 ± 0.0063** | 0.4505 ± 0.0064 |
| $\mathbb{R}^5$ | **0.0400** | 0.0779 | **0.0001** | 0.0585 | **0.3668 ± 0.0084** | 0.6140 ± 0.0086 |
| $\mathbb{R}^{10}$ | **0.0395** | 0.1055 | **0.0079** | 0.0564 | **0.4341 ± 0.0241** | 1.0747 ± 0.0190 |
| $\mathbb{S}^2$ | **0.0022** | 0.1749 | **0.0096** | 0.0148 | **1.1519 ± 0.0130** | 3.8635 ± 0.0203 |
| $\mathbb{S}^3$ | **0.0027** | 0.1072 | **0.0060** | 0.0371 | **1.1534 ± 0.0205** | 12.2643 ± 0.0878 |
| $\mathbb{S}^5$ | **0.0107** | 0.1777 | **0.0074** | 0.0287 | **1.3828 ± 0.0178** | 24.7303 ± 0.0462 |
| $\mathbb{S}^{10}$ | **0.0771** | 0.6015 | **0.0198** | 0.0241 | **1.8075 ± 0.0240** | 91.7757 ± 0.8074 |
| Sym(2) | **0.0584** | 1.6426 | **0.0337** | 0.0813 | 0.8898 ± 0.0111 | **0.6773 ± 0.0061** |
| Sym(3) | **0.0336** | 2.4352 | 0.0977 | **0.0876** | 1.0214 ± 0.0097 | 1.1938 ± 0.0424 |
| Sym(5) | **0.0653** | 3.8941 | 0.1202 | **0.0905** | 1.4501 ± 0.0102 | 2.6307 ± 0.1142 |
| LM(2 × 2) | 0.0473 | – | 0.0263 | – | 1.3179 ± 0.0119 | – |
| LM(3 × 2) | 0.0384 | – | 0.0118 | – | 1.4223 ± 0.0159 | – |
| LM(5 × 2) | 0.0704 | – | 0.0212 | – | 1.6141 ± 0.0163 | – |

Table 1: The estimated values and runtime of the diffusion $t$-mean using score and bridge sampling for different manifolds and dimension. The data consists of $1,000$ paths of Riemannian Brownian motion at time $T = 0.5$. For $\mathbb{S}^{10}$ we only use 100 samples. If the heat kernel is not known in closed-form we compare the estimates using scores and bridge sampling to the staring point and diffusion variance $T = 0.5$. For details on the heat kernels, see appendix C. We only compute the diffusion mean with bridge sampling when the logarithmic map is available in closed-form. The estimated run-time is based on five repeats after an initial compile of the code and show the time for 5 iterations. For details on hardware we refer to appendix E. LM referes to Landmarks.

Thus, choosing a sufficiently small $T > 0$, we can approximate the Logarithmic map as the limit of the learned gradients. Further, the Fréchet mean is given by $\mu^{\text{Fréchet}}(x_{1:N}) = \arg\min_{\mu \in \mathcal{M}} \mathbb{E}\left[\text{dist}^2(X, \mu)\right]$ [Fréchet, 1948], which can be estimated by gradient descent with $\nabla_\mu \mathbb{E}\left[\text{dist}^2(X, \mu)\right] = \mathbb{E}\left[2\text{Log}_\mu(X)\right]$ [Pennec, 2006]. Choosing a sufficiently small $T > 0$, we can approximate the Logarithmic map as the limit of the learned gradients using eq. 11 with gradient-descent. Since we only have to evaluate the scores in a fixed, sufficiently small value $T > 0$, we train specifically the scores for this approximation by only evaluating at a fixed time to increase performance. We set $t = 0.01$ for these experiments. Similar to Crane et al. [2013], we normalize in the gradient steps the estimated Logarithm map in eq. 11, since only the direction of the Logarithm map is important in the gradient descent method. We illustrate the estimate on DTI tensor data [Glasser et al., 2013, Sotiropoulos et al., 2013, Van Essen et al., 2013] in Fig. 5. The DTI tensor data consists of coronal slice of an HCP subject, which corresponds to $3 \times 3$ positive definite symmetric matrices. We plot the two leading eigenvalues of the data as well as the estimate using the closed-form logarithmic map compared to the logarithmic map estimated by the scores. Table 2 illustrates the estimates of the Fréchet mean using scores compared to using the closed-form expression of the logarithmic map. From Table 2 we see that the difference between the estimates is low, while the computational time is also reasonable. Note that the computational time of ground truth is significantly faster, since we use closed form expressions of the logarithmic map rather than solving the ODE (3) or minimizing the energy functional.

| Manifold | $\|\mu^{\text{score}} - \mu_0\|_2$ | $\|\mu^{\text{Fréchet}} - \mu_0\|_2$ | Score Time (s) | Fréchet Time (s) |
|---|---|---|---|---|
| $\mathbb{R}^2$ | **0.0206** | 0.0348 | 0.0869 ± 0.0020 | **0.0407 ± 0.0407** |
| $\mathbb{R}^3$ | **0.0199** | 0.0664 | 0.0910 ± 0.0011 | **0.0399 ± 0.0399** |
| $\mathbb{R}^5$ | **0.0214** | 0.0638 | 0.0877 ± 0.0017 | **0.0420 ± 0.0420** |
| $\mathbb{R}^{10}$ | **0.0161** | 0.0600 | 0.0952 ± 0.0025 | **0.0487 ± 0.0487** |
| $\mathcal{S}^2$ | **0.0661** | 0.1352 | 0.5629 ± 0.0225 | **0.3021 ± 0.3021** |
| $\mathbb{S}^3$ | **0.0477** | 0.3562 | 0.4425 ± 0.0055 | **0.2356 ± 0.2356** |
| $\mathbb{S}^5$ | **0.0536** | 1.0873 | 0.4873 ± 0.0268 | **0.2623 ± 0.2623** |
| $\mathbb{S}^{10}$ | **0.5727** | 1.5465 | 0.4993 ± 0.0108 | **0.2768 ± 0.2768** |
| $\mathcal{P}(2)$ | 0.0134 | **0.0000** | 0.6206 ± 0.0102 | **0.6399 ± 0.6399** |
| $\mathcal{P}(3)$ | 0.0085 | **0.0000** | 0.6534 ± 0.0097 | 0.7171 ± 0.7171 |
| $\mathcal{P}(5)$ | 0.0204 | **0.0000** | 0.7284 ± 0.0089 | 0.7767 ± 0.7767 |
| LM(2 × 2) | 0.0532 | – | 0.3305 ± 0.0060 | – |
| LM(3 × 2) | 0.0303 | – | 0.3511 ± 0.0032 | – |
| LM(5 × 2) | 0.0762 | – | 0.3671 ± 0.0042 | – |

Table 2: The estimated values and runtime of the Frêchet mean using the score and bridge sampling for different manifolds and dimension. The data is the same as in table 1. We only compute the Frêchet mean with the logarithmic map when the logarithmic map is available in closed-form. Expect the Eucldiean case we compare it to the starting point of the data sampled. The estimated run-time is based on five repeats after an initial compile of the code and show the time for 5 iterations. For details on hardware we refer to appendix E.

**Riemannian $k$-means** The $k$-means algorithm is an unsupervised learning [Bishop, 2006], which has been generalized into a Riemannian setting [Arvanitidis et al., 2018]. It consists of initializing K centroids and then assign each data point to the cluster of the closest centroid, and then iteratively update the centroid of each cluster as the mean value of data points in the cluster to minimize $\sum_{i=1}^{N} \sum_{k=1}^{K} \delta_{ik} ||x_i - \mu_k||^2$ with $x_{1:N}$ denoting the observations and $\delta_{ik} \in \{0, 1\}$ for all $i, k$ with $\sum_k \delta_{ik} = 1$. Similarly to the Euclidean method we can update the centroids computing the Fréchet mean using the gradient descent using eq. 11. In order to assign observations to their respective cluster, we have to use to estimate the Riemannian distance. By Varadhan's formulas [Varadhan, 1967] we have that $\text{dist}(x, y)^2 = -4 \lim_{t \searrow 0} t \log p_t(x, y)$. It follows from l'Hôspitals rule that

$$\lim_{t \searrow 0} t \log p_t(x, y) = \lim_{t \searrow 0} \frac{\log p_t(x, y)}{1/t}$$
$$= \lim_{t \searrow 0} \frac{\partial_t \log p_t(x, y)}{\frac{d}{dt} \frac{1}{t}} \qquad (12)$$
$$= \lim_{t \searrow 0} -\partial_t \log p_t(x, y) t^2.$$



Figure 7: Each row shows the three images from the same cluster using Riemannian $K$-means algorithm on the AFLW-2000 dataset. The Riemannian $K$-means algorithm is estimated for the corresponding points of the head position in $\mathbb{S}^3$.

As argued in [Crane et al., 2013] Varadhan's formula is sensitive to approximation errors. However, since we in the $k$-means algorithm only use the distance for cluster assignment, we are not required to have an accurate approximation of the distance. The algorithm can be found in appendix E. To illustrate the method we consider the AFLW-2000 dataset [Zhu et al., 2019], which consists of 1667 images with landmarks annotation as well Euler angles corresponding to the head pose of face images. The Euler angles can be converted into points on $\mathbb{S}^3$. We compute the $K$-means estimate on $\mathbb{S}^3$ and show the corresponding images for $K = 3$ in Fig. 7. We see that the method seem to accurately distinguish the images between looking straight ahead, looking right and looking to the left.

**Maximum Likelihood Riemannian Regression (MLRR)** Geodesic regression (GR) aims to find a relationship between observations $(x_1, y_1), \ldots, (x_N, y_N)$ by minimizing $\min_{\mu,v} \sum_{i=1}^{N} \text{dist} (y_i, \text{Exp}(\text{Exp}(\mu, x_i v), \epsilon_i))^2$, where $\epsilon \in T\mathcal{M}$ [Fletcher, 2013] such that the observations $y_i$ is described by a geodesic at $\mu$ with initial velocity $x_i v \in T_\mu \mathcal{M}$. Instead of having the observation noise $\epsilon$ in the tangent bundle we will consider the measurement noise as independent dispersions of Riemannian Brownian motion similar to [Sommer, 2018]. Rather than restricting the regression to a geodesic we will consider a general function, $f_\theta : \mathbb{R}^d \to \mathcal{M}$, with parameters $\theta$, and a diffusion variance function, $\sigma_\Phi : \mathbb{R}^d \to \mathbb{R}$, with parameters $\Phi$ such that

$$y_i \sim W^{\mathcal{M}}_{\sigma_\Phi(x)^2} (f_\theta(x)),$$



Figure 8: Maximum likelihood geodesic regression (orange) on $\mathbb{S}^2$ for synthetic data corresponding to noisy measurements (black) around a geodesic.

such that the observations can be seen as samples around a diffusion mean along $f$. Unlike the original formulation in Fletcher [2013] the measurement noise is not restricted as geodesics from the tangent space in direction $\epsilon_i$, but rather as measurement noise directly on the manifold by Brownian motion. Assuming independence between the observations we have that the log likelihood is given by

$$\mathcal{L} = \sum_{i=1}^{N} \log p_{\sigma_\Phi(x)^2} (f_\theta(x), y_i),$$

where $\sigma^2$ corresponds to the diffusion variance and $p$ the heat kernel. By symmetry of heat kernel, i.e. $\nabla_x p_t(x, y) = \nabla_y p_t(x, y)$, then the gradient of the log likelihood is given by

$$\nabla_\theta \mathcal{L} = \sum_{i=1}^{N} \nabla_y \log p_{\sigma_\Phi(x)^2}(f_\theta(x), y_i) \nabla_\theta f_\theta(x),$$
$$\nabla_\Phi \mathcal{L} = 2 \sum_{i=1}^{N} \partial_t \log p_{\sigma_\Phi(x)^2}(f_\theta(x), y_i) \nabla_\Phi \sigma_\Phi(x), \qquad (13)$$

8

We therefore propose to estimate the parameters by gradient descent. If $\sigma$ is constant and $f_\theta = \text{Exp}$, then the above formulation corresponds to maximum likelihood geodesic regression. Note further that if $\sigma \to 0$ if fixed, then the objective function above is equivalent to minimizing the distance by Hsu [2002]. In this case the original objective function from Fletcher [2013] appears. We show in Fig. 8 the result of maximum likelihood geodesic regression on $\mathbb{S}^2$. Further using the AFLW-2000 dataset we estimate the head position given landmark observations on the images. We show the result in Fig. 4. In order to compute the uncertainty estimate on $\mathbb{S}^3$ we train a neural network $g_t(x, y)$ such that $\nabla_y g_t(x, y)$ approximates the score as described in Salimans and Ho [2021].

## 5  Conclusion

In this paper we have shown that the Riemannian diffusion mean can be estimated using scores from diffusion models. The estimation of the diffusion mean was empirically compared to bridge simulation to estimate the diffusion mean, where score matching performed well in terms of both the estimations error and computational time. We further showed that the method extends to computing the Frêchet mean efficiently. We showed the applicability of the method by introducing a Riemannian $K$-means algorithm as well as maximum likelihood Riemannian regression model, where distances or the density are not known in closed-form expression. This work shows how the score of Riemannian Brownian motion can be used to build efficient algorithms on general Riemannian manifolds.

**Limitations and further research.**  In the paper we have assumed that the manifolds are geodesically and stochastically complete, have bounded sectional curvature, orientable and boundary less. Even though this covers most manifolds some will not be able fall into this category. Further research can be done to generalize the diffusion mean and sampling of Brownian motion to other cases that can generalize the above approach. The training of the score function can be numerically expensive if no closed-form expressions exist, which can restrict learning the score for high-dimensional data. In this paper we only estimate the log gradient of the score, however, to improve the model higher order derivatives could also be investigated to make the optimization scheme for the diffusion mean more efficient. We also note that we learn the score by simulating paths on the manifold. However, for high dimensional spaces there might be regions, where the score has not seen sampled paths before, which would make the estimate of the score less accurate.

# References

Alexis Arnaudon, Darryl Holm, and Stefan Sommer. *Stochastic Shape Analysis*, pages 1325–1348. Springer, Switzerland, 2023. ISBN 9783030986605. doi:10.1007/978-3-030-98661-2_86. Publisher Copyright: © Springer Nature Switzerland AG 2023.

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=SJzRZ-WCZ`.

Rabi Bhattacharya and Vic Patrangenaru. Large sample theory of intrinsic and extrinsic sample means on manifolds. *The Annals of Statistics*, 31(1):1 – 29, 2003. doi:10.1214/aos/1046294456. URL `https://doi.org/10.1214/aos/1046294456`.

Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. ISBN 0387310738.

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling, 2022.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

Mai Ngoc Bui, Yvo Pokern, and Petros Dellaportas. Inference for partially observed Riemannian Ornstein-Uhlenbeck diffusions of covariance matrices, November 2022. URL `http://arxiv.org/abs/2104.03193`. arXiv:2104.03193 [stat].

Marc Corstanje, Frank van der Meulen, Moritz Schauer, and Stefan Sommer. Simulating conditioned diffusions on manifolds, 2024.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics*, 32(5):1–11, September 2013. ISSN 1557-7368. doi:10.1145/2516971.2516977. URL `http://dx.doi.org/10.1145/2516971.2516977`.

Bernard Delyon and Ying Hu. Simulation of conditioned diffusion and application to parameter estimation. *Stochastic Processes and their Applications*, 116(11):1660–1675, 2006. ISSN 0304-4149. doi:https://doi.org/10.1016/j.spa.2006.04.004. URL `https://www.sciencedirect.com/science/article/pii/S0304414906000469`.

Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6): 141–142, 2012.

M.P. do Carmo. *Riemannian Geometry*. Mathematics (Boston, Mass.). Birkhäuser, 1992. ISBN 9783764334901. URL `https://books.google.dk/books?id=uXJQQgAACAAJ`.

Benjamin Eltzner, Pernille Hansen, Stephan F. Huckemann, and Stefan Sommer. Diffusion means in geometric spaces. *Bernoulli (in print)*, 2022.

P. Thomas Fletcher. Geodesic regression and the theory of least squares on riemannian manifolds. *International Journal of Computer Vision*, 105, 11 2013. doi:10.1007/s11263-012-0591-y.

Maurice Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l'institut Henri Poincaré*, 10(4): 215–310, 1948. URL `http://www.numdam.org/item/AIHP_1948__10_4_215_0/`.

Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the human connectome project. *NeuroImage*, 80:105–124, 2013. ISSN 1053-8119. doi:https://doi.org/10.1016/j.neuroimage.2013.04.127. URL `https://www.sciencedirect.com/science/article/pii/S1053811913005053`. Mapping the Connectome.

Søren Hauberg, Michael Schober, Matthew Liptrot, Philipp Hennig, and Aasa Feragen. A random riemannian metric for probabilistic shortest-path tractography. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 597–604, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24553-9.

E.P. Hsu. *Stochastic Analysis on Manifolds*. Contemporary Mathematics. American Mathematical Soc., 2002. ISBN 9780821883884. URL `https://books.google.dk/books?id=2NMOZ7svRmEC`.

Mathias Højgaard Jensen and Stefan Sommer. Mean estimation on the diagonal of product manifolds. *Algorithms*, 15(3), 2022. ISSN 1999-4893. doi:10.3390/a15030092. URL `https://www.mdpi.com/1999-4893/15/3/92`.

Mathias Højgaard Jensen and Stefan Sommer. Simulation of conditioned semimartingales on riemannian manifolds, 2023.

Erik Jørgensen. The central limit problem for geodesic random walks. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 32:1–64, 1975.

Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30: 509–541, 1977.

Aaron Lou, Isay Katsman, Qingxuan Jiang, Serge Belongie, Ser-Nam Lim, and Christopher De Sa. Differentiating through the fréchet mean, 2021.

Chenlin Meng, Yang Song, Wenzhe Li, and Stefano Ermon. Estimating high order gradients of the data distribution by denoising, 2021.

Omiros Papaspiliopoulos and G. Roberts. *Importance sampling techniques for estimation of diffusion models*, pages 311–340. 01 2012.

Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25:127–154, 07 2006. doi:10.1007/s10851-006-6228-4.

Tim Salimans and Jonathan Ho. Should EBMs model the energy or the score? In *Energy Based Models Workshop - ICLR 2021*, 2021. URL `https://openreview.net/forum?id=9AS-TF2jRNb`.

Stefan Sommer. An infinitesimal probabilistic model for principal component analysis of manifold valued data, 2018.

Stefan Sommer, Alexis Arnaudon, Line Kuhnel, and Sarang Joshi. Bridge simulation and metric estimation on landmark manifolds, 2017.

Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation, 2019.

Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models, 2021.

S. N. Sotiropoulos, S. Moeller, S. Jbabdi, J. Xu, J. L. Andersson, E. J. Auerbach, E. Yacoub, D. Feinberg, K. Setsompop, L. L. Wald, T. E. J. Behrens, K. Ugurbil, and C. Lenglet. Effects of image reconstruction on fiber orientation mapping from multichannel diffusion mri: Reducing the noise floor using sense. *Magnetic Resonance in Medicine*, 70(6):1682–1689, 2013. doi:https://doi.org/10.1002/mrm.24623. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24623`.

Alessandra Tosi, Søren Hauberg, Alfredo Vellido, and Neil D. Lawrence. Metrics for probabilistic geometries, 2014.

David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E.J. Behrens, Essa Yacoub, and Kamil Ugurbil. The wu-minn human connectome project: An overview. *NeuroImage*, 80:62–79, 2013. ISSN 1053-8119. doi:https://doi.org/10.1016/j.neuroimage.2013.05.041. URL `https://www.sciencedirect.com/science/article/pii/S1053811913005351`. Mapping the Connectome.

S. R. S. Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2):431–455, 1967. doi:https://doi.org/10.1002/cpa.3160200210. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160200210`.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. doi:10.1162/NECO_a_00142.

Ziyu Wang, Shuyu Cheng, Yueru Li, Jun Zhu, and Bo Zhang. A wasserstein minimum velocity approach to learning unnormalized models, 2020.

Chenchao Zhao and Jun S. Song. Exact heat kernel on a hypersphere and its applications in kernel svm. *Frontiers in Applied Mathematics and Statistics*, 4, 2018. ISSN 2297-4687. doi:10.3389/fams.2018.00001. URL `https://www.frontiersin.org/articles/10.3389/fams.2018.00001`.

Xiangyu Zhu, Xiaoming Liu, Zhen Lei, and Stan Z. Li. Face alignment in full pose range: A 3d total solution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1):78–92, January 2019. ISSN 1939-3539. doi:10.1109/tpami.2017.2778152. URL `http://dx.doi.org/10.1109/TPAMI.2017.2778152`.

Herbert Ziezold. *On Expected Figures and a Strong Law of Large Numbers for Random Elements in Quasi-Metric Spaces*, pages 591–602. Springer Netherlands, Dordrecht, 1977. ISBN 978-94-010-9910-3. doi:10.1007/978-94-010-9910-3_63. URL `https://doi.org/10.1007/978-94-010-9910-3_63`.

# Appendix

## A Sampling Riemannian Brownian Motion

Brownian motion, $W_t^{\mathcal{M}}$ on a Riemannian manifold, $\mathcal{M}$, can be simulated as a random walk in the tangent space mapped onto the Riemannian manifold by the exponential map [Jørgensen, 1975].

---

**Algorithm 2:** Simulating Brownian Motion in Tangent Space

---

**Input**: $T$, $x_0 \in \mathcal{M}$, $N_{\text{steps}}$, $D$ (dimension of manifold)
**Output**: Samples of Brownian Motion $W_t$ with $W_0 = x_0$
Set $\delta = \frac{T}{N_{\text{steps}}}$
Set $W_0 = x_0$
**for** $k = 1, \ldots, N_{steps}$ **do**
    Sample $v_k \sim \mathcal{N}(0, I_D)$
    $W_{k\delta} = \text{Exp}_{\sqrt{\delta}v}\left(W_{(k-1)\delta}\right)$
**end**
return $\{W_t\}_{t \in [0, \delta, 2\delta, \ldots, T]}$

---

Brownian Motion on a manifold can equivalently also be simulated in a local coordinate system, $\{x^i\}_{i=1}^D \subset \mathcal{M}$ with $\sigma_j^i = g^{ji}$ [Hsu, 2002]

---

**Algorithm 3:** Simulating Brownian Motion in Local Coordinates

---

**Input**: $T$, $x_0 \in \mathcal{M}$, $N_{\text{steps}}$, $D$ (dimension of manifold)
**Output**: Samples of Brownian Motion $W_t$ with $W_0 = x_0$
Set $\delta = \frac{T}{N_{\text{steps}}}$
Set $W_0 = x_0$
**for** $k = 1, \ldots, N_{steps}$ **do**
    Sample $v_k \sim \mathcal{N}(0, I_D)$
    $W_{k\delta}^i = \frac{1}{2}g^{jk}\left(W_{(k-1)\delta}\right)\Gamma_{jk}^i\left(W_{(k-1)\delta}\right)dt + \sigma_j^i\left(W_{(k-1)\delta}\right)v_k^j, \quad \forall i \in \{1, \ldots, D\}$
**end**
return $\{W_t\}_{t \in [0, \delta, 2\delta, \ldots, T]}$

---

## B Riemannian Manifolds

Table 3 describes the various manifolds used in the papers as well as their embeddings.

To avoid singularities in the chart for the embedded manifolds the basis of the chart is updated, when the chart deviates from the initial center. This is for instance the case for stereographic projection, where the antipodal point is undefined.

## C Heat Kernel

The following section will briefly state the heat kernels that have closed form expressions for the Riemannian manifolds studied in this paper.

### C.1 Euclidean Space

In $\mathbb{R}^m$ the heat kernel is given by [Eltzner et al., 2022]

$$p(x, y, t) = \frac{1}{(2\pi t)^{m/2}} e^{-\frac{-||x-y||^2}{2t}}, \tag{14}$$

with $x, y \in \mathbb{R}^m$. This implies that

$$\ln p(x, y, t) = -\frac{||x-y||^2}{2t} - \frac{m}{2}\ln(2\pi t). \tag{15}$$

| Manifold | Description | Parameters | Local Coordinates | Embedding |
|---|---|---|---|---|
| $\mathbb{R}^n$ | The Euclidean $n$-dimensional vector space | - | Standard basis of $\mathbb{R}^n$ | - |
| $\mathbb{S}^n$ | The n-sphere, $\{x \in \mathbb{R}^{n+1} \mid ||x||_2 = 1\}$ | - | Stereographic coordinates | $15800 \pm 239$ |
| $\mathcal{P}(n)$ | The space of positive definite matrices | - | Local coordinates are given respect to the standard basis in $\mathbb{R}^{n(n+1)/2}$ | Embedded into $\mathbb{R}^{n^2}$ by the mapping $f(x) = l(x)l(x)^T$, where $l : \mathbb{R}^{n(n+1)/2} \to \mathbb{R}^{n \times n}$ maps $x$ into a lower triangle matrix consisting of the elements in $x$. |
| $\mathrm{Sym}(n)$ | The space of symmetric matrices | - | Local coordinates are given respect to the standard basis in $\mathbb{R}^{n(n+1)/2}$ | Embedded into $\mathbb{R}^{n^2}$ by converting the local coordinates into a symmetric triangular matrix. |
| $\mathrm{Landmarks}(n \times d)$ | The space of landmarks, $\{x_k\}_{k=1}^n$ is a shape space that consist of $n$ distinct points $x_k \in \mathbb{R}^d$. The landmark space can be equipped with a right-invariant Riemannian metric [Arnaudon et al., 2023]. | The Gaussian Kernel with parameter $\alpha = 1$ | All landmarks are given with respect to the standard basis in $\mathbb{R}^{nd}$ | - |

Table 3: Description of the manifolds used in the paper.

Thus

$$\nabla_y \ln p(x, y, t) = \frac{x - y}{t},$$
$$\frac{d}{dt} \ln p(x, y, t) = \frac{||x - y||^2}{2t^2} - \frac{m}{2t} \tag{16}$$

It can clearly be seen that the optimal diffusion mean is given by

$$\frac{1}{N} \sum_{i=1}^N \nabla_\mu \ln p(x_i, \mu^*, t) \Rightarrow \mu^* = \frac{1}{N} \sum_{i=1}^N x_i,$$

which is independent of the diffusion time. The optimal diffusion time is then given by

$$\frac{1}{N} \sum_{i=1}^N \frac{d}{dt} \ln p(x_i, \mu^*, t^*) = \frac{\frac{1}{N} \sum_{i=1}^N ||x_i - \mu^*||^2}{m}.$$

## C.2 Circle

The heat kernel on $\mathbb{S}^1$ is given by [Eltzner et al., 2022]

$$p(x, y, t) = \frac{1}{\sqrt{2\pi t}} \sum_{k \in \mathbb{Z}} e^{-\frac{-(x - y + 2\pi k)^2}{2t}}, \tag{17}$$

with $x, y \in \mathbb{R}/2\pi\mathbb{Z}$, which implies that

$$\ln p(x, y, t) = -\frac{1}{2} \ln 2\pi t + \ln \sum_{k \in \mathbb{Z}} e^{-\frac{-(x - y + 2\pi k)^2}{2t}}. \tag{18}$$

Assuming that the infinite sum is uniformly and absolute convergent, then

| Manifold | $\mathbf{x}_0$ | Architecture | Activation Function | Embedded |
|---|---|---|---|---|
| $\mathbb{R}^n$ | $\mathbf{0}_n$ | Linear$(128, 128, 128)$ | TANH | - |
| $\mathbb{S}^n$ | North pole | Linear$(512, 512, 512, 512, 512)$ | TANH | $\times$ |
| $\mathcal{P}(n)$ | $10 \cdot \text{diag}(n)$ | Linear$(512, 512, 512)$ | TANH | $-$ |
| Sym$(n)$ | $\text{diag}(n)$ | Linear$(512, 512, 512)$ | TANH | $-$ |
| Landmarks$(n \times 2)$ | See description above | Linear$(512, 512, 512)$ | TANH | $-$ |

Table 4: The Architecture for learning the score

$$\nabla_y \ln p(x,y,t) = \frac{\sum_{k \in \mathbb{Z}} (2\pi k + x - y) e^{-\frac{(x-y+2\pi k)^2}{2t}}}{\sqrt{2\pi t} p(x,y,t)},$$

$$\frac{d}{dt} \ln p(x,y,t) = \frac{\frac{1}{2\sqrt{2\pi t}} \sum_{k \in \mathbb{Z}} \frac{-(x-y+2\pi k)^2}{2t^{-2}} e^{-\frac{(x-y+2\pi k)^2}{2t}}}{p(x,y,t)} - \frac{\frac{1}{(2t)^{3/2}\sqrt{\pi}} \sum_{k \in \mathbb{Z}} e^{-\frac{(x-y+2\pi k)^2}{2t}}}{p(x,y,t)} \tag{19}$$

## C.3 m-Sphere

For the m-sphere, $\mathbb{S}^m$, with $m \geq 2$ the heat kernel is given by uniformly and absolutely convergent infinite sum in theorem 1 in Zhao and Song [2018]

$$p(x,y,t) = \sum_{l=0}^{\infty} e^{-l(l+m-1)\frac{t}{2}} \frac{2l+m-1}{m-1} \frac{1}{A_S^m} C_l^{(m-1)/2} \left( \langle x, y \rangle_{\mathbb{R}^{m+1}} \right), \tag{20}$$

where $x, y \in \mathbb{R}^{m+1}$, $C_l^\alpha$ denotes the Gegenbauer polynomials and

$$A_S^m = \frac{2\pi^{(m+1)/2}}{\Gamma\left( (m+1)/2 \right)}, \tag{21}$$

which gives the gradients

$$\nabla_y \ln p(x,y,t) = \frac{\sum_{l=0}^{\infty} e^{-l(l+m-1)\frac{t}{2}} (2l+m-1) \frac{1}{A_S^m} C_{l-1}^{(m+1)/2} \left( \langle x, y \rangle_{\mathbb{R}^{m+1}} \right) x}{p(x,y,t)},$$

$$\frac{d}{dt} \ln p(x,y,t) = \frac{-\sum_{l=0}^{\infty} \frac{l(l+m-1)}{2} e^{-l(l+m-1)\frac{t}{2}} \frac{2l+m-1}{m-1} \frac{1}{A_S^m} C_l^{(m-1)/2} \left( \langle x, y \rangle_{\mathbb{R}^{m+1}} \right)}{p(x,y,t)}. \tag{22}$$

The above is implemented using the recursion

$$l C_l^{(\alpha)}(x) = 2(l-1+\alpha) x C_{l-1}^{(\alpha)}(x) - (n+2\alpha-2) C_{l-2}^{(\alpha)}(x). \tag{23}$$

# D  Data

The synthetic data consists of the end points of $1,000$ sampled Brownian motions using algorithm 3 with starting point, $x_0$, described in table 4. The cardiac data consist of annotated images of cardiac, while the butterfly data consists of annotated images of butterflies. We have scaled the cardiac data, so it has similar numerical values to the butterfly data to be in area, where the estimated score has seen sufficiently many sample paths. The DTI tensor data consists of pre-processed diffusion data of 20 subjects from the Q3 release of the Human Connectome Project (HCP) [Hauberg et al., 2015]. The data is processed in accordance with Hauberg et al. [2015], which gives corresponding values in $\mathcal{P}(3)$ for four segmentation. We use 100 samples of segmentation 1 in examples. The MNIST data is described in Deng [2012].

# E  Implementation

The implementation is based on JAXGEOMETRY. JAXGEOMETRY is open-source and can be found at `https://github.com/ComputationalEvolutionaryMorphometry/jaxgeometry`. The initial seed value is set to 2712 when generating random numbers.

## E.1 Hardware

The score has all been trained for 24 hours on the system with 4 nodes, where for each node is configured with

- 2x Intel Xeon Processor 2650v4 (12 core, 2.20GHz)

- 256 GB memory

- FDR-Infiniband

- 480 GB-SSD disk

All other computations have been performed on a *HP* computer with Intel Core i9-11950H 2.6 GHz 8C, 15.6" FHD, 720P CAM, 32 GB (2×16GB) DDR4 3200 So-Dimm, Nvidia Quadro TI2004GB Descrete Graphics, 1TB PCle NVMe SSD, backlit Keyboard, 150W PSU, 8cell, W11Home 64 Advanced, 4YR Onsite NBD.

## E.2 Sampling

We sample paths of Brownian motion using algorithm 3 with $N_{\text{steps}} = 1000$ and $T = 1.0$. Initially, a batch Brownian motion paths are simulated from a fixed starting point, $x_0 \in \mathcal{M}$. For each $x_0$ we sample a fixed number of paths, and similarly for each $t$ we sample a fixed number of times. For each batch of sample paths we choose SAMPLES PER $x_0$ end points of the samples as new starting points for the next batch of Brownian motion. Table 5 shows the hyper-parameters from sampling.

| Parameter | Value |
|---|---|
| Samples per $x_0$ | 1 |
| Samples per $t$ | 100 |
| Number of $x_0$ | 1024 |
| Time Steps for sampling Brownian Motion | 100 |
| Sampling time for Brownian Motion | 1.0 |

Table 5: Hyper-Parameters for Sampling

## E.3 Score Matching

Table 4 shows the architecture for each Riemannian manifold. The column, 'embedded', indicates whether the score was learned in the embedded space. If there is no mark, the score is learned in local coordinates. The column, 'Activation Function', indicates the activation function applied after each hidden layer. All scores are trained for 24 hours or up to a maximum of $50,000$ epochs.

The initial starting point, $x_0$, in table 4 for the landmark spaces consisting of 2 and 5 landmarks are 2d points, where the first coordinate is equally spaces between $[-5, 5]$, and the second coordinate is zero. For 10-landmarks and 20-landmarks $x_0$ is set as the first observation of the butterfly dataset with sub-sampling landmarks accordingly. In table 4 the initial coordinate for SPDN is written in local coordinates.

All manifolds are trained using the ADAM optimizer with learning rate 0.001 with linear annealing for the first 1000 epcohs and the cosine annealing similar to Bortoli et al. [2022]. For the embedded scores, the score is projected the tangent space in embedded space.

## E.4 Optimization for Diffusion $t$-Mean and Fréchet Mean

The estimation of the diffusion $t$-mean is done using the gradient descent algorithm to jointly estimate $t$ and $\mu$. The hyper-parameters for estimating the diffusion t-mean is shown in table 6.

| Manifold | Learning rate | $t_0$ | Iterations | Method |
|---|---|---|---|---|
| $\mathbb{R}^n$ | 0.1 | 0.2 | 1000 | ADAM |
| $\mathbb{S}^n$ | 0.1 | 0.2 | 1000 | Gradient-Descent |
| $\text{Sym}(n)$ | 0.1 | 0.2 | 1000 | ADAM |
| $\text{Landmarks}(n)$ | 0.1 | 0.2 | 1000 | ADAM |

Table 6: The Hyper-parameters for estimating the diffusion t-mean.

The hyper-parameters for estimating the Frêchet mean is shown in table 7

| Manifold | Learning rate | $t_0$ | Iterations | Method |
|---|---|---|---|---|
| $\mathbb{R}^n$ | 0.1 | 0.2 | 1000 | ADAM |
| $\mathbb{S}^n$ | 0.01 | 0.2 | 1000 | Gradient-Descent |
| $\mathrm{Sym}(n)$ | 0.1 | 0.2 | 1000 | ADAM |
| $\mathcal{P}(n)$ | 0.1 | 0.2 | 1000 | ADAM |
| $\mathrm{Landmarks}(n)$ | 0.1 | 0.2 | 1000 | ADAM |

Table 7: The Hyper-parameters for estimating the Frêchet mean.

## E.5  Riemannian K-Means Estimation

We estimate the Logarithmic map with a fixed time of 0.1. The Fréchet mean is estimated using 100, a step size of 0.1 and a fixed time of 1.0. The $K$-means algorithm is estimated using 10 iterations with $K = 4$ for both $\mathbb{S}^2$ and 20-landmarks.

---

**Algorithm 4:** Riemannian $k$-Means

---

**Input**: K, $N_{\text{iter}}$, $\{\mu_i^{\text{init}}\}_{i=1}^K$, $\{x_n\}_{i=1}^N$
**Output**: K Centroids and Clusters
Learn neural network $s_\theta$ minimizing eq. 8
Set $\{\mu_i\}_{i=1}^K = \{\mu_{(i)}^{\text{init}}\}_{i=1}^K$
**for** $i=1,\ldots,N_{iter}$ **do**
    **for** $j=1,\ldots,K$ **do**
        **for** $n=1,\ldots,N$ **do**
            Set $d_n = \sqrt{4t^2 \partial_t \log p_t(x_n, \mu_j)}$ for small $t$
        **end**
        Set $z_{n,\arg\min_n d_n} = 1, z_{nk} = 0, \forall k \neq \arg\min_n d_n$
    **end**
    Set $\mu_k$ as the Fréchet mean using eq. 11 for all $\{x_n\}_{n=1}^N$ in cluster $k$.
**end**
return $(\{\mu_k\}_{k=1}^K, \{z_{nk}\}_{k=1}^K, \forall n \in \{1, \ldots, N\})$

---

## E.6  Gaussian Process

We train Gaussian process for a MNIST image rotated $2\pi$, where we sample 200 images along the rotation. We fix the latent variables as the 200 corresponding points on the circle and compute the posterior distribution and expected Riemannian metric as described in Tosi et al. [2014].