

Computing distances and means on manifolds with a metric-constrained Eikonal approach

Daniel Kelshaw

Department of Aeronautics, Imperial College London, London, UK

djk21@ic.ac.uk

Luca Magri

*Department of Aeronautics, Imperial College London, London, UK
The Alan Turing Institute, British Library, London, UK
Politecnico di Torino, DIMEAS, Torino, Italy*

l.magri@imperial.ac.uk

Abstract

Computing distances on Riemannian manifolds is a challenging problem with numerous applications, from physics, through statistics, to machine learning. In this paper, we introduce the *metric-constrained Eikonal solver* to obtain continuous, differentiable representations of distance functions (geodesics) on manifolds. The differentiable nature of these representations allows for the direct computation of globally length-minimising paths on the manifold. We showcase the use of metric-constrained Eikonal solvers for a range of manifolds and demonstrate the applications. First, we demonstrate that metric-constrained Eikonal solvers can be used to obtain the Fréchet mean on a manifold, employing the definition of a Gaussian mixture model, which has an analytical solution to verify the numerical results. Second, we demonstrate how the obtained distance function can be used to conduct unsupervised clustering on the manifold – a task for which existing approaches are computationally prohibitive. This work opens opportunities for distance computations on manifolds.

1 Introduction

Many high-dimensional systems can be expressed on embedded, low-dimensional manifolds (Fefferman et al., 2016). Standard Euclidean measures of distance are insufficient for providing meaningful measures of similarity, and thus distance measures which take into account the geometry of the underlying manifold are necessary. Distance computations on nonlinear manifolds are crucial in a range of areas, ranging from statistics to machine learning. Computing statistics, or conducting reduced-order modelling on manifolds are notable examples of the importance of distance functions on manifolds (Guigui et al., 2023; Pennec, 2006). Tangent principal component analysis (TPCA) aims to find a mean on the manifold, project points from the manifold to the tangent space at the mean, before finally computing PCA on the tangent space (Huckemann & Ziezold, 2006; Zhang & Zha, 2004). In computing the mean of points on the manifold, the arithmetic mean does not account for the geometry of the manifold, and as a result may not lie on the surface of the manifold. The mean instead must be defined as the point which minimises the square distance to samples on the manifold, known as the Fréchet mean. This notion of distance must adhere to the geometry of the manifold. Applications of distance functions on manifolds extend to the field of machine learning, where embedding methods are prevalent (Wang et al., 2014). Contrastive language-image pre-training (CLIP) make use of a cosine similarity to predict image, text pairings in a zero-shot learning (Radford et al., 2021). The advent of retrieval-augmented generation in large language modelling makes use of vector databases of text embeddings to search for similar representations, which primarily use measures of Euclidean distance, or cosine similarities (Lewis et al., 2021; Gao et al., 2023). The field of physics holds numerous examples of where manifold-based distances are important. The most notable is that of general relativity. General relativity describes how matter can influence spacetime, as evidenced by gravitational lensing. Due to the curvature of spacetime, the distance that light travels cannot be described through standard Euclidean

norms, but rather the notion of geodesic distance is required (Wald, 2010).

Much of the existing work on computing geodesic distances on nonlinear manifolds is framed in a discrete manner. Methods typically fall into one of two approaches: (i) Eikonal-based approaches, and (ii) length-minimising approaches. Distance functions are characterised by having a gradient of unit magnitude everywhere in the domain. The Eikonal equation (Evans, 2010) naturally models this behaviour and, as a result, many methods for computing distance leverage an Eikonal-based approach. The Eikonal equation can be treated as a stationary boundary value problem, or transformed to a time-dependant problem which models arrival time, or wave front propagation. A common means of computing geodesic distances is through the fast marching method, which is a level-set method (Kimmel & Sethian, 1998; Sethian, 1999). The fast marching method treats the problems as a stationary boundary value problem, using a numerical approach to solve a system of equations after discretisation. This fast marching method is similar to Dijkstra’s algorithm (Dijkstra, 1959), which propagates information from a known solution, the boundary values, in a sequential manner – updating a single grid point at a time so that the solution is strictly increasing. The algorithm used to achieve this comprises two steps: (i) a local solver, used to approximate the gradient at a point and estimate distance to surrounding nodes; and (ii) an update step, responsible for selecting the closest point in order to propagate the solution. The quality of the solution depends on the accuracy of the local solver. In a similar vein, there exist fast-sweeping methods for obtaining numerical solutions for the Eikonal equation on a rectangular grid. Fast-sweeping methods use upwind differencing for discretisation, in conjunction with Gauss-Seidel iterations with alternating sweeping ordering (Zhao, 2005; 2007). Ultimately, this allows for parallel updates which reduce the computational complexity of the approach. Another class for computing geodesic distances are heat methods, which leverage a time-dependant solution approach (Crane et al., 2013; Lin et al., 2014; Varadhan, 1967). These methods work by (i) integrating the scalar heat flow for a fixed time; (ii) computing and normalising the gradient of the resulting scalar field; and (iii) solving a Poisson equation to obtain the closest scalar potential, which is the distance. Again, this approach requires discretisation, and relies on a number of assumptions and approximations. The initial condition for the heat flow requires discrete modelling of a point source, and the heat flow assumption states that the computed distance approaches the true distance as the integration time tends to zero. Finally, machine learning approaches to solving the Eikonal equation have been presented by (Gropp et al., 2020), but operate under Euclidean assumptions. These assumptions make them unsuitable for operating on nonlinear manifolds. Though machine learning approaches seek an approximate solution, certain guarantees can be made through the choice of the architecture. The authors of (Grubas et al., 2023) model caustics with solutions to the Eikonal equation and bound solutions based on analytical models.

In contrast with Eikonal-based approaches, the distance between two points can be computed via a length-minimising curve approach. These shortest curves are known as *geodesics* and satisfy the geodesic equation. The geodesic equation, however, describes a class of locally length-minimising curve – with the desired curve being the shortest curve in this class. Rather than solving a boundary value problem to satisfy the equation, methods instead opt to parameterise a curve between two points, optimising parameters of this curve to yield the globally length-minimising curve. Curves used to represent geodesics must be at least two times differentiable to satisfy the geodesic equation; thus, cubic splines are a natural choice. Splines are functions which are defined by piecewise polynomials, constructed in such a manner to guarantee continuity between each constituent spline. In addition, boundary conditions can be enforced through construction. In general, splines are parameterised by the locations of the intersections of the piecewise polynomials, referred to as knots. In obtaining a length-minimising geodesic, the location of these knots is optimised to yield a curve with minimising length (Detlefsen et al., 2021; Ni & Wang, 2023; Rajković, 2023). Instead of enforcing the structure of cubic splines on the geodesic, another approach is to parameterise a geodesic between two points through a neural network (Chen et al., 2018). This neural representation provides a continuous, differentiable representation which can then be integrated to compute the distance. The parameters of the network are optimised to obtain a length-minimising representation of the curve.

In this paper, we provide a method to yield a continuous, differentiable representation of the distance function on a nonlinear manifold. Instead of obtaining a distance field around a fixed point, as in the numerical approaches, or only considering a distinct pair of points, as in the length-minimising curve approaches, our method yields a distance function defined directly on the manifold. We further show how we can leverage the differentiable representation of the distance function to obtain globally length-minimising curves between points without incurring the cost of solving additional sub-optimisation problems.

We first provide a background of material required to specify the Eikonal equation on the manifold in §2. In §3, we introduce the metric-constrained Eikonal equation, which provides a continuous, differentiable distance function on the manifold. Our proposed approach allows us to obtain globally length-minimising paths directly. In §4 we showcase results for three manifolds of interest: Euclidean space, the hypersphere, and a non-convex manifold with multiple local extrema. We demonstrate the efficacy of our approach through comparing results to ground-truth solutions, and the recovery of length-minimising paths. Finally, we showcase the applications of the metric-constrained Eikonal solver to compute the Fréchet mean on the surface of a Gaussian mixture model, and to perform unsupervised clustering on curved manifolds. We end the paper with conclusions in §6. We provide access to our open source library on Github.¹

2 Review on Differential Geometry

Differential geometry provides tools for the study of smooth manifolds, of which Euclidean spaces are a subset. Many real world data lives in non-Euclidean spaces, and so we adopt a non-Euclidean toolbox to deal with this. To provide a motivating example, consider two points $p, q \in \mathbb{S}^2$ on the unit sphere. Using the Euclidean L_2 norm provides a distance between these points, representing the length of a line interior to the surface. Should we wish to compute the length of the shortest line which lies on the surface of the sphere, we need to consider a different form of distance metric. In this section, we provide a brief overview of concepts from differential geometry which form the building blocks of our proposed methodology. For a comprehensive explanation of the subject, we refer the reader to (Lee, 2018; do Carmo, 2013). Notation employed follows the Einstein summation convention.

2.0.1 Riemannian manifolds and inner product

A Riemannian manifold (Lee, 2018; do Carmo, 2013) is a pair (M, g) , where M is a smooth manifold (Lee, 2012), and g is a Riemannian metric on M . This metric constitutes a symmetric bilinear form $g : T_p M \times T_p M \rightarrow \mathbb{R}$, where $T_p M$ denotes a vector in the tangent space at $p \in M$. Components of the metric are defined as $g_{ij}(p) = \langle \partial_i|_p, \partial_j|_p \rangle$, where $\partial_i = \partial/\partial x^i$. For arbitrary vectors $v, w \in T_p M$, the metric allows for computation of the inner product on the tangent space (Lee, 2018)

$$\langle v, w \rangle_g = g_p(v, w) = g_{ij} v^i w^j. \quad (1)$$

The definition of the inner product allows us to compute distances and angles in the tangent space at a point $p \in M$; ultimately, providing the building blocks for operating on Riemannian manifolds. In Euclidean geometry, the metric is defined as $\tilde{g}_{ij} = \delta_{ij}$, which is the zero-curvature condition. Using this metric in conjunction with the definition of the inner product, we recover the standard definition of the inner product on Euclidean spaces, $\langle v, w \rangle_{\tilde{g}} = v^i w^i$. Every smooth manifold admits a Riemannian metric (Lee, 2018).

2.1 Induced metric

In order to make the problem computationally possible, we work with embedded submanifolds. Suppose (\tilde{M}, \tilde{g}) is a Riemannian manifold, and $M \subseteq \tilde{M}$ is an embedded submanifold. Given a smooth immersion $\iota : M \hookrightarrow \tilde{M}$, where $\dim(M) \leq \dim(\tilde{M})$, the metric $g = \iota^* \tilde{g}$ is referred to as the metric induced by ι , where ι^* is the pullback (Lee, 2018). If (M, g) is a Riemannian submanifold of (\tilde{M}, \tilde{g}) , then for every $p \in M$ and $v, w \in T_p M$, the induced metric is (from Nash embedding theorems)

$$g_p(v, w) = (\iota^* \tilde{g})_p(v, w) = \tilde{g}_{\iota(p)}(d\iota_p(v), d\iota_p(w)). \quad (2)$$

¹All code is available on GitHub: <https://github.com/magrilab/riemax>

This equivalence forms the guiding principle that allows us to compute induced metrics induced by differentiable immersions (do Carmo, 2013). We leverage this idea throughout the paper to numerically define metrics.

2.2 Covariant differentiation and parallel transport

Given the inner product on a manifold, we can introduce the notion of covariant differentiation; a means of specifying derivatives along tangent vectors. For a vector $w \in T_p M$ and a vector field $v \in TM$, the covariant derivative is defined as (do Carmo, 2013)

$$\nabla_v w = v^j \nabla_{\partial_j} w^i \partial_i = (v^j w^k_{;j}) \partial_k = (v^j w^k_{;j} + \Gamma^k_{ij} v^j w^i) \partial_k, \quad (3)$$

where, for brevity, we introduce the notation $w^k_{;j}, w^k_{;j}$ to denote the partial, and covariant derivative respectively – each taken with respect to the j th coordinate vector field. The parenthesised expression in Eq. (3) provides an intuitive explanation for the concept of covariant differentiation. The first term represents the directional derivative, and the second term is responsible for quantifying the twisting of the coordinate system. The term Γ^k_{ij} is referred to as the Christoffel symbols of the second kind, or the affine connection (Lee, 2018), which provide a measure of how the basis changes over the manifold. The Christoffel symbols are defined as

$$\Gamma^k_{ij} = \partial_j(\partial_i) \cdot \partial^k = \frac{1}{2} g^{km} (g_{mi,j} + g_{mj,i} - g_{ij,m}), \quad (4)$$

where $\partial^k = g^{km} \partial_m$ is the index-raised coordinate vector field (contravariant), and g^{km} is the inverse of the metric tensor, such that $\delta^i_k = g^{ij} g_{jk}$. In general, transporting a vector $w \in T_p M$ along the manifold does not guarantee that it remains parallel to itself in the tangent planes along the path. Definition of the covariant derivative provides a sufficient condition to ensure parallelism is conserved along a trajectory. Mathematically this is expressed as

$$\nabla_v w = v^j w^k_{;j} + \Gamma^k_{ij} v^j w^i = 0. \quad (5)$$

This condition for parallelism can be exploited to extend the notion of straight curves to the manifold. Geodesics, which are straight lines in Euclidean space, allow us to define a notion of distance.

2.3 Geodesics

A geodesic $\gamma : [a, b] \subset \mathbb{R} \rightarrow M$ generalises the notion of a straight line on the manifold. These curves are locally length-minimising because they constitute a solution of the Euler-Lagrange equations, which, mechanically, describe the motion of a rigid body devoid of acceleration (do Carmo, 2013). Leveraging the definition of covariant differentiation, we can describe a geodesic as a curve for which tangent vectors remain parallel to themselves as they are transported along the curve, yielding the geodesic equation (Lee, 2018)

$$\nabla_{\dot{\gamma}} \dot{\gamma} = \ddot{\gamma}^k + \Gamma^k_{ij} \dot{\gamma}^i \dot{\gamma}^j = 0, \quad \text{s.t.} \quad \gamma(\lambda = 0) = \gamma_0, \dot{\gamma}(\lambda = 0) = \dot{\gamma}_0, \lambda \in [0, 1] \quad (6)$$

where derivatives $\ddot{\gamma}, \dot{\gamma}$ are taken with respect to the affine parameter λ . This geodesic equation ensures the speed $\langle \dot{\gamma}, \dot{\gamma} \rangle_g^{0.5}$ is constant along the trajectory, allowing us to classify the parameterisations $\lambda \in [0, 1]$ as unit-speed geodesics. Conversely, the geodesic can be parameterised with respect to length to yield a unit-distance geodesic (Lee, 2018).

The definition of the geodesic equation gives rise to two mappings of interest. Given a point $p \in M$, the \exp_p and \log_p maps are, respectively

$$\exp_p : T_p M \rightarrow M, \quad (7)$$

$$\log_p : M \rightarrow T_p M. \quad (8)$$

The \exp_p map defines an initial value problem. Given a vector $v \in T_p M$, there exists a unique geodesic $\gamma : [0, 1] \rightarrow M$ satisfying $\gamma(0) = p, \dot{\gamma}(0) = v$. The exponential map is then defined as $\exp_p(v) = \gamma(1)$. We

can define the \log_p map as the functional inverse of the \exp_p map, that is, $\exp_p \circ \log_p : q \mapsto q \ \forall q \in M$. The log map, however, is not necessarily unique because there may be many $v \in T_p M$ for which $\exp_p v = q$, providing a challenge in defining distance on the manifold. In §3, we overcome this by defining the \log_p^* map, which makes use of the gradient of the distance field to obtain a unique solution.

2.4 Distances on manifolds

The notion of distance on the manifold is defined by geodesics. Although geodesics are length-minimising, this is only in a local sense. Given two points $p, q \in M$, the geodesic distance is defined as the infimum of the length of all valid geodesics (do Carmo, 2013)

$$d_g(p, q) = \inf_{\gamma} \left\{ \int_0^1 \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\gamma(t)}^{1/2} dt : \gamma(0) = p, \gamma(1) = q \right\}. \quad (9)$$

A geodesic γ is said to be minimising *iff* there exist no shorter, valid geodesics with the same endpoints. This forms the basis for approaches which seek to obtain length-minimising curves between two points, as outlined in §1.

2.5 Curvature

The Riemann curvature tensor R^l_{ijk} provides a means to quantify curvature on the manifold; a Riemannian manifold with zero curvature is referred to as *flat* and is locally isomorphic to Euclidean space. Given two *infinitesimally close* geodesics, curvature is responsible for geodesic deviation. To obtain a scalar value for the curvature, one can successively contract over indices of the Riemann curvature tensor (Lee, 2018)

$$R^l_{ijk} = \Gamma^l_{ik,j} - \Gamma^l_{ij,k} + \Gamma^l_{jm}\Gamma^m_{ik} - \Gamma^l_{km}\Gamma^m_{ij} \rightarrow R_{ij} = R^m_{imj} \rightarrow R = g^{ij}R_{ij}, \quad (10)$$

where R^l_{ijk} denotes the Riemann curvature tensor; R_{ij} the Ricci curvature tensor, obtained by contracting over the first and third indices; and R , the Ricci scalar – the geometric trace of the Ricci tensor with respect to the metric. The Ricci scalar is a local invariant on the manifold and assigns a single real number to the curvature at a particular point. The Ricci scalar relies solely on the definition of the metric tensor g_{ij} .

In this paper, we provide a computational framework for operating on differentiable manifolds. We leverage the automatic differentiation capabilities of `jax` (Bradbury et al., 2018) to compute the metric induced by smooth immersions as described in Eq. (2), as well as all derivative properties.

3 The metric-constrained Eikonal solver and methodology

Given a smooth immersion $\iota : M \hookrightarrow \mathbb{R}^n$, with $\dim(M) \leq n$, we propose a method to obtain a continuous, differentiable representation of the distance function on M . In this section, we first discuss properties of distance functions and metric spaces, providing criteria that our proposed distance function should satisfy. Second, we build intuition for a distance function constrained around a *single-point* fixed in the domain. Specifically, we (i) introduce the Eikonal equation as a means to express distance functions; (ii) outline a principled means by which to tractably compute \log_p maps; and (iii) deploy neural networks to parameterise the distance function. Third, we extend the approach to compute generic and global distance functions. Finally, we provide a brief outline of network architectures and training techniques.

3.1 Properties of distance functions

A distance function $d : M \times M \rightarrow \mathbb{R}$ is a function with a gradient of unit magnitude at every point in the domain. Defining a distance function on the manifold is equivalent to defining a metric space (M, d) (Lee, 2012), which for all points $p, q, r \in M$ satisfy

$$\begin{aligned} \mathbf{M1:} \quad & d(q, q) = 0 & \mathbf{M3:} \quad & d(p, q) = d(q, p) \\ \mathbf{M2:} \quad & d(p, q) \geq 0 & \mathbf{M4:} \quad & d(p, q) \leq d(p, r) + d(r, q) \end{aligned} \quad (11)$$

In other words, **M1** defines the notion of *zero distance*; **M2** represents *positivity*; **M3** is the notion of *symmetry*; and **M4** states the *triangle inequality* holds. We can leverage the definition of geodesics to impose an additional property

$$\mathbf{M5}: (\iota^* d_E)(p, q) \leq d(p, q), \quad (12)$$

which states that the geodesic distance is not shorter than the Euclidean distance in the image space of the immersion. This follows immediately from acknowledging that geodesics are length-minimising curves on the manifold. Our proposed methodology ensures that $\mathbf{M}\{1, 2, 3, 5\}$ hold by construction.

3.2 Single-point solutions

We first discuss single-point variations of distance functions. We refer to *single-point* solutions as distance functions with a fixed point $p \in M$, such that $d_p : q \mapsto d(p, q)$. This allows us to introduce the methodology, before extending the approach to handle global distance functions in §3.3.

3.2.1 Eikonal equations

We consider a class of nonlinear first-order partial differential equations, the Eikonal equations, for which the solution satisfies the properties of the distance function, i.e., the magnitude of the gradient is unity at all points in the domain. Fixing a point $p \in M$, the distance $\varphi : q \mapsto d(p, q)$ satisfies the equation

$$\|\nabla \varphi\| = 1 \quad \text{s.t.} \quad \varphi|_p = 0. \quad (13)$$

On a Riemannian manifold (M, g) , we augment the formulation with the exterior derivative, which we refer to as *the metric-constrained Eikonal equation*

$$\|\nabla \varphi\| = \langle \text{grad } \varphi, \text{grad } \varphi \rangle_g = \varphi^i \varphi_{,i} = 1 \quad \text{s.t.} \quad \varphi|_p = 0, \quad (14)$$

where $(\text{grad } \varphi)^i = g^{ij} \varphi_{,j}$ is the exterior derivative, and $\varphi^i = g^{ij} \varphi_{,j}$. Taking derivatives in this way ensures that the solution adheres to the geometry of the manifold, as described by the metric induced by the immersion.

3.2.2 Implicit log maps

Consider two points $p, q \in M$, and level-sets of constant distance around the point p . The shortest path connecting p, q must travel orthogonal to the level-sets, that is, geodesics are defined by the gradient of the distance function. Precisely, the exterior derivative of the distance function produces the geodesic flow, which is a section of the tangent bundle. This flow satisfies the geodesic equation (do Carmo, 2013),

$$\nabla_{\text{grad } \varphi} \text{grad } \varphi = 0, \quad (15)$$

This geodesic flow allows us to obtain globally length minimising geodesics. Although the \log_p map is not unique, we formulate a \log_p^* map to obtain the tangent vector which yields the globally shortest geodesic between two points. We exploit the obtained distance function and the corresponding geodesic flow to obtain a vector field describing the dynamics of globally length-minimising paths. We first consider the system

$$\dot{\gamma}(\lambda) = -\text{grad } \varphi(\gamma(\lambda)), \quad \text{with } \gamma : [0, \varphi(q; p)) \rightarrow M. \quad (16)$$

By taking $\gamma(\lambda = 0) = q, \dot{\gamma}(\lambda = 0) = -\text{grad } \varphi(q; p)$ as the initial condition, we can define the \log_p^* map as

$$\log_p^* : q \mapsto \lim_{\lambda \rightarrow \varphi(q; p)^-} \left[\text{grad } \varphi \left(\gamma(\lambda = 0) + \int_0^\lambda \dot{\gamma}(\lambda) d\lambda \right) \right]. \quad (17)$$

Intuitively, given points $p, q \in M$ we compute the \log_p^* map by (i) evaluating the distance, $\varphi(q; p)$; (ii) starting at q , trace the gradient field back to p by integrating the initial value problem along the length of the geodesic; and (iii) evaluate $\text{grad } \varphi(p)$. We cannot directly evaluate $\text{grad } \varphi(p)$ as the solution is not unique because p is

the source of the geodesic flow. Instead, we follow the gradient field in the reverse direction from a point q , tracing out the globally length-minimising geodesic which terminates at p . Standard methods for computing the \log_p^* map rely on parameterising and optimising curves between pairs of points. By employing a metric-constrained Eikonal solver approach, once the distance function has been obtained, the length-minimising geodesics, and thus the \log_p^* map can be computed straightforwardly (§4).

3.2.3 Network parameterisation

We seek a continuous, differentiable solution to the metric-constrained Eikonal equation. To this end, we parameterise the distance function by a neural network $\tilde{\varphi}_\theta : M \rightarrow \mathbb{R}$. Neural networks are universal function approximators of continuous functions (Hornik et al., 1989). The output of the network is augmented to yield the distance function

$$\varphi_\theta(q; p) = (\iota^* d_E)(p, q) [1 + (\sigma \circ \tilde{\varphi}_\theta)(q; p)], \quad (18)$$

where $\sigma \in \mathcal{C}^{\geq 2}(\mathbb{R}, \mathbb{R}_{\geq 0})$ is strictly increasing and at least twice differentiable; \circ denotes function composition; and $d_E : (p, q) \mapsto \|p - q\|_2$ is the Euclidean distance function. Defining the distance function in this manner serves three distinct purposes: (i), the constraint $\varphi|_p = 0$ is handled implicitly (**M1**); the output of the network is guaranteed to be positive (**M2**); and (iii), the inequality $\varphi(q; p) \geq (\iota^* d_E)(p, q)$ is guaranteed (**M5**). We seek a solution $\varphi_{\theta^*} = d_p(q)$ through solving the optimisation problem for the trainable parameters

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta; p) \quad \text{where } \mathcal{L} = \mathbb{E}_q (\|\varphi_{\theta^*, i}(q) \varphi_{\theta, i}(q) - 1\|^2), \quad (19)$$

where the expectation is taken with respect to a distribution $q \sim \mathcal{D}$ in M . This optimisation problem seeks to minimise the residual of the metric-constrained Eikonal equation across the domain.

3.3 Global solutions

In this subsection, instead of constraining the solution to the Eikonal equation to a single point, we consider learning a global solution $\varphi : (p, q) \mapsto d(p, q)$. We first introduce an additional equation to enforce the Eikonal constraint on both variables

$$\|\nabla_p \varphi(p, q)\| = \|\nabla_q \varphi(p, q)\| = 1 \quad \text{s.t.} \quad \varphi|_{p=q} = 0, \quad (20)$$

where ∇_p, ∇_q denotes partial derivatives with respect to the respective function arguments. As before, we parameterise the solution with a neural network, this time taking points $p, q \in M$ as inputs. As well as the constraints imposed for the single-point case, we also wish to impose symmetricity (**M3**). To achieve this, we augment the output of the network to define the global distance function

$$\varphi_\theta(p, q) = (\iota^* d_E)(p, q) [1 + \sigma \circ (\tfrac{1}{2}(\tilde{\varphi}_\theta(p, q) + \tilde{\varphi}_\theta(q, p)))] . \quad (21)$$

As a result of this imposed reciprocity, we know the equality $\|\nabla_p \varphi(p, q)\| = \|\nabla_q \varphi(p, q)\|$ holds. The symmetricity constraint ensures that we still need only a single equation to specify the global solution to the metric-constrained Eikonal equation (Grubas et al., 2023), and so we follow a similar optimisation approach as in the *single-point* case described in §3.2. Thus, we seek a solution $\varphi_{\theta^*} = d(p, q)$ through solving

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta) \quad \text{where } \mathcal{L} = \mathbb{E}_q (\|\varphi_{\theta^*, i}(p, q) \varphi_{\theta, i}(p, q) - 1\|^2), \quad (22)$$

where the expectation is taken with respect to a distribution $q \sim \mathcal{D}$ in M . (The network outputs are augmented in such a manner to provide a lower-bound on the distance based on the pullback of the Euclidean distance, $\iota^* d_E$ (**M5**). (We provide an upper-bound of the distance in Appendix A. This method is not showcased in the results of this paper as it increased computational complexity with no significant impact on the quality of the results.)

3.4 Network architecture and training

We parameterise the distance functions φ_θ with a neural network; in particular, we employ the modified multi-layer perceptron (MLP) architecture as described by (Wang et al., 2021). This choice of architecture has

two main advantages: (i) the introduction of positional embeddings; and (ii) the use of residual connections aid in mitigating numerical gradient instability. Given the reliance on our proposed methodology on the use of automatic differentiation for computing the loss, the modified MLP is a natural choice. For an extensive overview on modified MLPs, we refer the reader to (Wang et al., 2023).

The optimisation of the network $\tilde{\varphi}_\theta$ is performed through a gradient-descent approach, which employs the **adam** optimiser (Kingma & Ba, 2017) to minimise the expectation of the loss. The initial learning rate for the optimiser is prescribed in each case, and then follows an exponential decay schedule to aid convergence. The expectation of the loss should provide a meaningful estimate of the residual error, and is taken over a distribution $q \sim \mathcal{D} \in M$. A standard approach would be to sample points in a uniform distribution, selecting all points in the domain with equal probability. However, we can exploit properties of the geometry to provide improved estimates.

3.4.1 Curvature-based sampling

The solution φ to the metric-constrained Eikonal equation is more sensitive in regions of high curvature, which is a consequence of induced geodesic deviation (do Carmo, 2013). We mitigate the impact of geodesic deviation by employing a curvature-based sampling strategy, which ensures that we can capture the solution appropriately in regions of high curvature. Employing the Metropolis-Hastings (Metropolis et al., 1953; Hastings, 1970) sampling scheme, we draw samples in regions of the manifold that exhibit higher degrees of scalar curvature, which is characterised by the magnitude of the Ricci scalar (Lee, 2018). An outline of this curvature-based sampling is provided in algorithm 1.

Algorithm 1 Metropolis-Hastings curvature-based sampling

Require: $q_0 \in M, \delta \in \mathbb{R}$ *// q_0 is the initial condition, δ controls the step-size*
for $t \leftarrow 1, \dots, n$ **do**
 $q' \leftarrow q_{t-1} + \delta \mathcal{N}(0, 1)$ *// propose the location of a new sample*
 $u \leftarrow \mathcal{U}(0, 1)$
 $\alpha \leftarrow \log(|R(q')|) - \log(|R(q_{t-1})|)$ *// compute acceptance ratio*
 if $\log(u) \leq \alpha$ **then** *// accept the proposed point*
 $q_t \leftarrow q'$
 else *// reject the proposed point*
 $q_t \leftarrow q_{t-1}$

4 Results

In order to demonstrate the efficacy of our proposed methodology, we show results for three manifolds of interest: (i) Euclidean space, (ii) the unit hypersphere, and (iii) the peaks manifold. We first consider the Euclidean case, providing a baseline for which we can verify the method in the absence of curvature. Second, we introduce constant curvature by obtaining distance functions on the unit hypersphere. In each of these cases, there exists an analytical form of the distance function; parameterising the manifolds by the number of dimensions allows us to investigate how the proposed methodology scales with increasing dimensionality. Finally, we investigate the peaks manifold, which is a non-convex manifold with multiple local extrema. There exists no analytical form for the distance function on the peaks manifold, and as a result we show comparisons against the length-minimising approach. An overview of the manifolds used throughout the paper, with the exception of Euclidean space, is provided in Figure 1 where visualisations are shown in \mathbb{R}^3 . This section focuses on demonstrating the proposed methodology, and §5 highlights the benefits that using a continuous-differentiable distance function can bring.

4.1 Euclidean Space

We first consider the limiting case of the Euclidean manifold, for which there is zero curvature across the entire domain. As discussed in §2, the metric tensor is the identity matrix. Despite the apparent simplicity




Figure 1: Overview of the manifolds used throughout the paper, visualised in \mathbb{R}^3 . Panel (a) depicts a patch of the two-sphere, panel (b) shows a model defined by the probability density function of a Gaussian mixture model, and panel (c) shows a multi-modal manifold.

of the manifold, this provides a foundation for analysis. The Euclidean metric is induced by the smooth immersion $\iota : x \mapsto x$ with $x \in \mathbb{R}^n$ (the identity function). The corresponding distance function is the ℓ^2 norm, $d_g : (p, q) \mapsto \|p - q\|_2$. Pairing the manifold with the corresponding distance function yields a metric space.

4.1.1 Performance

In order to assess the performance of the proposed methodology, we seek to recover the distance function for dimensions $n \in \{2, 5, 10, 15, 20\}$. In each case, we seek a solution $\varphi_{\theta^*} = d_g$. Performance is measured using the relative ℓ^2 error between the predicted, and ground truth distance. We define the error

$$\varepsilon = \sqrt{\frac{\mathbb{E}[(\varphi_{\theta}(p, q) - d_g(p, q))^2]}{\mathbb{E}[d_g(p, q)^2]}} \quad \text{where } p, q \sim \mathcal{U}_{[M^+, M^-]}, \quad (23)$$

where M^+, M^- denote the bounds imposed for training on the manifold. For the purposes of this work, we consider a bounded subset of the manifold, namely $M \in [-3, 3] \subset \mathbb{R}^n$, so we can standardise the inputs to the network.

4.1.2 Optimisation

We seek a solution to the global Eikonal equation (global optimisation problem Eq. (22)). Following the initialisation of the network, optimisation is performed with the `adam` optimiser. Further details on the network and optimisation parameters can be found in Appendix B. For each case $n \in \{2, 5, 10, 15, 20\}$, we train a total of five models, each initialising weights of the network with a different random seed. Each model is optimised for a total of 5×10^4 parameter updates, empirically determined to provide sufficient convergence. The loss is computed using using 2^{14} points uniformly sampled for each p, q (we do not employ curvature-based sampling as the curvature is uniformly zero across the domain).

To provide an illustrative example, we first consider the case \mathbb{R}^2 for which we can visualise the resulting distance fields. Figures 2(a, b) show the predicted distance, and true distance around the point $p = 0$. The distance field is characterised by level-sets forming concentric circles around the origin. Qualitatively, we observe that the predicted and ground-truth fields are identical, demonstrating that the model is able to recover the distance function appropriately. Figure 2c provides quantitative results for each of the cases explored, quoting the mean relative ℓ^2 error achieved across the five runs. The standard deviation for each n did not exceed 6.239×10^{-8} , and as such error bars are not reported. We observe that, with the exception of

$n = 20$, the error increases with increasing dimension n . This behaviour is expected, as the distance function becomes more challenging to model in higher dimensions.




Figure 2: Results for the Euclidean manifold. Panel (a) shows the predicted distance around a point, panel (b) shows the corresponding true distance from the point, and panel (c) shows the mean relative ℓ^2 error across five repeats for increasing dimensions $n \in \{2, 5, 10, 15, 20\}$.

4.2 Hypersphere

We continue our investigation by considering a patch of the hypersphere \mathbb{S}^n , another manifold for which there exists an analytical distance function. A visualisation is provided in Figure 1a. While our previous example of Euclidean space exhibits zero curvature, the hypersphere is characterised by constant curvature, with Ricci scalar $R = n(n-1)r^{-2}$ where r is the radius (do Carmo, 2013). Without loss of generality, we consider the unit hypersphere, $r = 1$. The introduction of curvature marks a notable challenge for which Euclidean-based methods are not positioned to handle. An induced metric of the hypersphere manifold is provided by the immersion $\iota : \mathbb{S}^n \rightarrow \mathbb{R}^{n+1}$. A definition of the immersion is provided in Appendix C. The distance function on the unit hypersphere is provided by the cosine rule, $d_g : (p, q) \mapsto \arccos(\iota(p) \cdot \iota(q))$. This manifold, in conjunction with the associated distance function, provides the metric space.

4.2.1 Performance

Similarly to the experiments run in the Euclidean case, we seek to recover the distance function for each $n \in \{2, 5, 10, 15, 20\}$. The performance is measured by the relative ℓ^2 error between the predicted, and ground-truth distances; employing the definition in Eq. (23). Imposing a coordinate system on the hypersphere inevitably leads to a singularity which can cause numerical issues. As a result of this singularity, and the periodicity of the manifold, we impose bounds $x \in [\pi/6, 5\pi/6] \subset \mathbb{S}^n$. These bounds are used to standardise inputs to the network.

4.2.2 Optimisation

The optimisation scheme used for obtaining a solution to the metric-constrained Eikonal equation in the case of the hypersphere is identical to that used in the Euclidean case, described in §4.1. In this case, we also do not employ the curvature-based sampling method, as we know *a-priori* that the manifold exhibits constant scalar curvature.

For the case of \mathbb{S}^2 , we provide a visualisation of the distance function around the point $p = \pi/2$. Results in Figures 3(a, b) show the predicted distance, and true distance fields respectively. The presence of curvature, though subtle, is apparent in these examples, with the level-sets of distance appearing warped when compared with the original Euclidean case. Despite the presence of curvature, the predicted and true distance fields show good agreement. Figure 3c shows the results for recovering the distance function on the hypersphere for each of the cases, quoting the mean relative ℓ^2 error achieved across the five runs. The standard deviation

for each case did not exceed 3.049×10^{-3} , and as such error bars are not reported as they are too small. We observe error increasing as a function of manifold dimension, which is due to accumulation of errors.




Figure 3: Results for the hypersphere manifold. Panel (a) shows the predicted distance around a point, panel (b) shows the corresponding true distance from the point, and panel (c) shows the mean relative ℓ^2 error across five repeats for dimensions $n \in \{2, 5, 10, 15, 20\}$.

4.3 Peaks Manifold

The manifolds we have examined thus far have analytical distance functions; however, in many scenarios, we do not have access to the analytical distance functions. Therefore, we further test the metric-constrained Eikonal solver by considering a non-convex manifold with multiple local extrema (Figure 1c), exhibiting non-negligible scalar curvature; namely the peaks function (?)

$$\iota : [x; y] \mapsto [x; y; 3(1-x)^2 e^{-x^2-(y+1)^2} - 10\left(\frac{x}{5} - x^3 - y^5\right) e^{-x^2-y^2} - \frac{1}{3} e^{-(x+1)^2-y^2}]. \quad (24)$$

We first consider single-point solutions to the distance function, as described in §3.2, obtaining a solution around the point $p = 0$. Next, we extend this to examine the global solution, as described in §3.3, comparing results with a cubic-spline-based length-minimising approach (Detlefsen et al., 2021).

4.3.1 Single-point solutions

We first consider single-point solutions to the metric-constrained Eikonal equation. For optimisation, we use the `adam` optimiser with exponential weight decay; details of which can be found in the Appendix B. We train for a total of 2×10^5 parameter updates, drawing 2×10^{14} random samples from the domain for each parameter update. In this case, we employ the curvature-based sampling method described in §3.4, drawing random samples from the average of a uniform distribution and the distribution defined by the probability density function (PDF) of the Ricci scalar across the manifold. The curvature-based sampling is shown in Figure 4a, with the colour map corresponding to the probability density function, and points representing random samples drawn from the distribution. Results in Figure 4b show the predicted distance field from the centre of the domain p , with arrows representing the geodesic flow, $\text{grad } \varphi_\theta$, overlaid. We observe that the geodesic distance adheres to the geometry. Using the predicted geodesic flow, we obtain length-minimising curves (yellow); these are compared with valid geodesics (red), generated through integration of the geodesic equation (5) with $\gamma(0) = q, \dot{\gamma}(0) = -\text{grad } \varphi_\theta(q)$. The overlap of these two trajectories demonstrates the validity of the obtained geodesic flow. Finally, to assess how the choice of $p \in M$ affects the performance, we investigate the symmetricity of the distance function across the domain. We consider a 7×7 grid of origin points on the manifold, distributed equally in the Euclidean sense, forming the set \mathcal{X} . For each $p, q \in \mathcal{X}$, we obtain a distance function and compute pair-wise distances, $\varphi_\theta(q; p), \varphi_\theta(p; q)$, as shown in Figure 4c. We observe that as the distance increases, the variability in predicted distance increases correspondingly as distance predictions accumulate error with increasing distance.




Figure 4: Results for the single-point solution on the peaks manifold. Panel (a) shows the probability density function generated from the curvature-based sampling (Algorithm 1), with samples from the distribution overlaid. Panel (b) shows the distance field around a point in the centre of the domain with arrows depicting the geodesic flow overlaid. The yellow curves denote length-minimising geodesics obtained from the geodesic flow, while the red curves denotes the corresponding valid geodesics. Panel (c) shows results for the symmetry test.

4.3.2 Global Solutions

Obtaining a solution constrained to a single point has limited uses. In general, we wish to obtain a distance function defined over the entire manifold. Analysis for the single-point case demonstrates that, although we obtain accurate results, we cannot guarantee symmetry. Moving to a *global* solution, we can ensure that we satisfy the condition **M3**. In absence of an analytical solution for arbitrary manifolds, including the peaks manifold, we compare the metric-constrained Eikonal approach with the standard approach found in the literature (optimising parameterised curves (Detlefsen et al., 2021; Chen et al., 2018)). For the standard approach, we represent geodesics as cubic-splines. These cubic splines are defined in such a manner that the boundary conditions are enforced automatically. The curves are parameterised in such a manner that any parameterisation yields a valid curve. We optimise the parameters of the cubic splines through a gradient-descent based approach, seeking to minimise the energy of the resulting curve. Optimisation is performed with a gradient-descent approach, using the `adam` optimiser with a learning rate of 10^{-3} for a total of 50×10^3 iterations. Convergence is achieved when the absolute magnitude of the gradient updates do not exceed 10^{-4} .

We employ the network φ_θ to represent the global distance function. This network is trained via the augmentation defined in Eq.(21), and results compared with those obtained from the parameterised curve approach. Results for this comparison are shown in Figure 5. In many cases, the parameterised curve approach fails to converge, and so comparisons are shown only in cases where convergence was achieved. We observe that the Eikonal-based distances tend to be equal to, or less than, the lengths of the parameterised curves. This shows that the parameterised curves provide an upper-bound of the distance.

Given a point $p \in M$, we can visualise the resulting distance field around the point. Examples of this are provided in Figure 6 where distance fields are shown for five distinct points. In addition to the distance fields, the geodesic flow is denoted by the red vector fields in the figure. The geodesic flow is defined across the entire domain, and can be employed to obtain globally length-minimising geodesics. In the bottom right panel of Figure 6, we take four of the points, labelled P1 – P4, and obtain the geodesics connecting each point to its neighbours. We observe that we obtain a complete cycle with each geodesic connecting the points successfully.




Figure 5: Comparison of results from standard parameterised curve approach against results from the metric-constrained Eikonal solver. Results are shown only in the cases where the parameterised curve approach converges. We observe good agreement between the proposed metric-constrained Eikonal approach, and the standard length-minimising curve approach.




Figure 6: Distance fields and geodesic flows on the peaks manifold. The first five panels show a colour map of the geodesic distance with the vector field from the geodesic flow overlaid. The final panel in the lower right shows four geodesics generated from these vector fields, completing a cycle: P1 – P4.

5 Applications

There exist a wide range of applications for which manifold-based distance computations are useful. Given a differentiable representation of the distance function on the manifold, we are able to perform optimisation with respect to the distance. To this end, we examine two applications of differentiable distance functions on manifolds: (i) computing the Fréchet mean (Guigui et al., 2023); and (ii) unsupervised clustering (Jin & Han, 2010).

5.1 Computing the Fréchet mean

The arithmetic mean is broadly incompatible with arbitrary manifolds, being suitable for operating in Euclidean spaces only (Pennec, 2006; Guigui et al., 2023). To provide an illustrative example, consider two distinct points on the surface of \mathbb{S}^2 . Taking the arithmetic mean of these points yields a mean interior to the sphere, rather than on the surface. The Fréchet mean accounts for this and is defined as the point on the manifold which minimises the squared distance to samples on the manifold, where the distance function is defined on the manifold (Pennec, 2006). The Fréchet mean is

$$\mu^* = \arg \min_{\mu} \sum_{q \in \mathcal{X}} d_g(\mu, q)^2, \quad (25)$$

where \mathcal{X} denotes the set of points on the manifold for which we wish to compute the mean. Standard approaches for computing the Fréchet mean rely on computing the log map between pairs of points, ultimately requiring solving sub-optimisation problems, which becomes computationally expensive for large numbers of points. This standard approach is outlined in algorithm 2. Instead, we propose to first obtain a differentiable representation of the distance function using a metric-constrained Eikonal solver, outlined in algorithm 3. The obtained distance function provides fast inference, and allows us to optimise directly for the Fréchet mean. To provide a quantitative example, we must first define a manifold for which we can analytically compute the mean. For the purposes of this work, we employ a Gaussian mixture model.

Algorithm 2 Standard algorithm for computing Fréchet mean

Require: $q \in \mathcal{X} \subset M$

```

 $\mu_0 \leftarrow \frac{1}{|\mathcal{X}|} \sum_{p \in \mathcal{X}} p$ 
for  $i \leftarrow 1, \dots, n$  do
  for  $p \in \mathcal{X}$  do
     $\gamma_p \leftarrow \text{compute\_geodesic}(\mu_{i-1}, p)$ 
     $L_p \leftarrow \frac{\dot{\gamma}_p(0)}{\|\dot{\gamma}_p(0)\|} L(\gamma_p)$ 
   $\hat{\mu}_i \leftarrow \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{L}} x$  // compute tangent space mean
   $\mu_i \leftarrow \exp_{\mu_{i-1}}(\hat{\mu}_i)$  // map tangent space mean back to the feature space

```

Algorithm 3 Computing the Fréchet mean with the metric-constrained Eikonal solver

Require: $q \in \mathcal{X} \subset M$

```

 $\varphi_\theta \leftarrow \text{solve\_metric\_constrained\_eikonal}()$ 
 $\mu_0 \leftarrow \frac{1}{|\mathcal{X}|} \sum_{p \in \mathcal{X}} p$ 
for  $i \leftarrow 1, \dots, n$  do
   $\tau \leftarrow \sum \varphi_\theta(\mu_{i-1}, p)^2$  // compute action to minimise
   $\mu_i \leftarrow \text{gradient\_step}(\mu_{i-1}, \partial \tau / \partial \mu_{i-1})$  // conduct gradient-based update

```

5.1.1 Gaussian mixture models as a manifold

The probability density function (PDF) of a Gaussian mixture model is a weighted sum of the PDFs of the component Gaussians

$$p(x; \mu, \Sigma, \alpha) = \sum_{i=1}^{N_G} \alpha_i f_{\mathcal{N}}(x; \mu_i, \Sigma_i) \quad \text{s.t.} \quad \sum_{i=1}^{N_G} \alpha_i = 1 \quad \text{with} \quad x, \mu_i \in \mathbb{R}^n, \Sigma_i \in \mathbb{R}^{n \times n} \quad (26)$$

where for N_G component Gaussians, α_i is a weighting term, μ_i is the mean of the i th Gaussian, Σ_i is the covariance matrix of the i th Gaussian, and $f_{\mathcal{N}}$ is the density function for a Gaussian distribution. Given the definition of a Gaussian mixture model, the mean is defined as the weighted sum of the constituent means, that is $\mu = \sum_{i=1}^{N_G} \alpha_i \mu_i$. By using the Gaussian mixture model, we define a manifold as the surface of the resulting probability density function. We define an immersion $\iota : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1}$ as

$$\iota : x \mapsto [x; p(x; \mu, \Sigma, \alpha)]. \quad (27)$$

For the applications explored in this paper, we consider a manifold of dimension $n = 2$. In combination, the manifold and the corresponding, yet to be found, distance function define our metric space. The manifold employed for the Fréchet mean case can be seen in Figure 1b.

5.1.2 Fréchet mean of the Gaussian Mixture Model

We first define the manifold as the surface of a Gaussian mixture model with four constituent Gaussians. The domain is then bounded such that $p \in [-3, 3] \subset \mathbb{R}^2$, which allows us to standardise inputs to the network. Parameters of the network are trained for a total of 5×10^4 parameter updates, using the `adam` optimiser to update the weights. The optimiser has an initial learning rate $\eta = 10^{-3}$, which decays exponentially. Further information on the network and optimiser can be found in Appendix B. Using the same Gaussian mixture model as that used in training, we sample points on the surface of the manifold, discarding points that lie outside of the prescribed domain. The mean of these points is by definition the expected value of the given Gaussian mixture model. Given these points, and the distance function, we seek to recover the mean of the distribution by way of computing the Fréchet mean. We compute the Fréchet mean through gradient-based optimisation, which is possible as a result of the differentiable nature of the distance function. We provide an initial guess μ_0 , and use the `adam` optimiser with learning rate $\eta = 10^{-2}$ to update the candidate mean. In practice, given the fast inference afforded by the distance function, we can run multiple optimisation chains in parallel. We provide sixteen initial guesses and optimise in parallel. Figure 7 shows the manifold, the samples, and the optimisation trajectories from each of the initial conditions. We observe that each of the optimisation trajectories terminates at the same point. For each these optimisation chains, we compute the relative ℓ^2 error between the true mean, and the predicted mean. The mean of these errors is $\bar{\varepsilon} = 4.873 \times 10^{-3}$, showing we are able to accurately recover the Fréchet mean on the manifold defined by a Gaussian mixture model.

5.2 Unsupervised clustering on the manifold with the metric-constrained Eikonal solver

Clustering-based methods follow a similar principle to that of computing the Fréchet mean. Given a set of points \mathcal{X} , we seek to identify centroids and their associated clusters. For k -means clustering (Lloyd, 1982), we define the objective

$$(\mu^*, \mathcal{S}^*) = \arg \min_{(\mu, \mathcal{S})} \sum_{i=1}^k \sum_{x \in \mathcal{S}_i} d_g(x, \mu)^2, \quad (28)$$

where for k clusters, μ_i denotes the centroid of each cluster \mathcal{S}_i . To solve the optimisation problem, we make use of Lloyd’s algorithm (Lloyd, 1982), which alternates between two steps: (i) assignment, which assigns each observation to the cluster with the nearest centroid; and (ii) update, which recomputes the centroids of each cluster based on the observations assigned to the cluster.

A large number of distance computations are required for convergence. In the assignment stage a total of $k|\mathcal{X}|$ distance computations are required, and the update stage relies on computing the Fréchet mean for

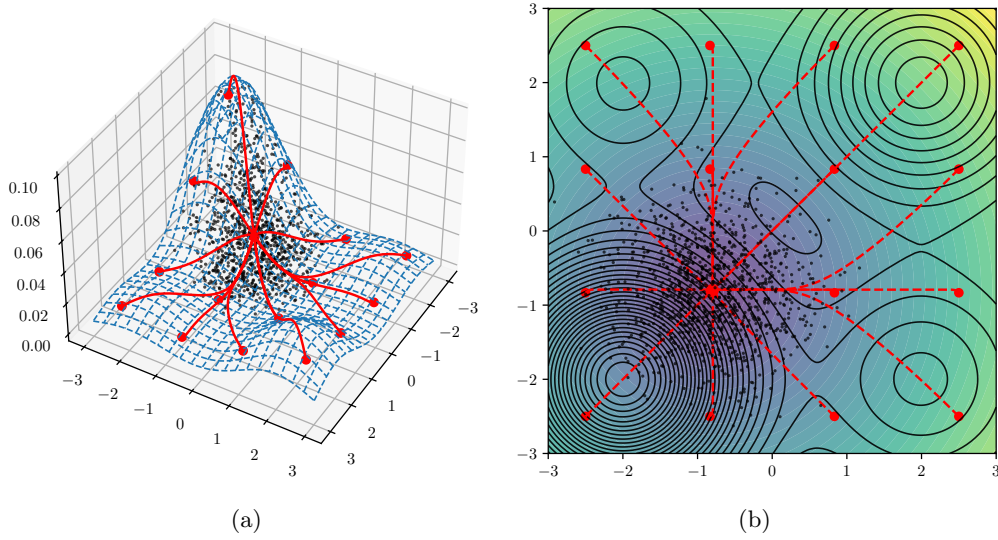


Figure 7: Computing the Fréchet mean. Panel (a) provides an extrinsic view of the manifold, the samples, and the optimisation trajectories – each converging to the mean of the underlying distribution. Panel (b) provides an intrinsic view of the same result. The arithmetic mean does not lie on the surface of the manifold, but it is interior to the surface.

each cluster as described in §5.1. Existing methods are ill-equipped to solve such a problem: (i) the large number of distance computations required at each iteration is prohibitive; and (ii) the nested optimisations required for obtaining Fréchet means at each iteration is infeasible.

Algorithm 4 Unsupervised clustering on the manifold

Require: $(\mu^{(0)}, S^{(0)})$ *// assign initial guesses for clusters and their centroids*
 $i \leftarrow 0$
 $\varphi_\theta \leftarrow \text{solve_metric_constrained_eikonal}()$
while $S^{(i)} \neq S^{(i-1)}$ **do**
 $d_{\mu \rightarrow p} \leftarrow \varphi_\theta(\mu^{(i-1)}, p)$ *// compute distance between each centroid and each data point*
 $S^{(i)} \leftarrow \text{assign_clusters}(d_{\mu \rightarrow p})$ *// assign points to centroids that minimise the distance*
 $\mu^{(i)} \leftarrow \text{frechet_mean}(S^{(i)})$ *// recompute centroids of each cluster, as per algorithm 3*
 $i \leftarrow i + 1$

5.2.1 Clustering on the Peaks manifold

We demonstrate the feasibility of this approach for points randomly sampled on the surface of the peaks manifold. First, samples \mathcal{X} are generated on the surface of the manifold. We then use the previously obtained global distance function from §4.3.2 in conjunction with Lloyd’s algorithm to obtain a total of five distinct clusters. Results for this clustering are shown in Figure 8. We first provide a baseline by conducting standard k -means clustering with a Euclidean distance function on the manifold, results for which are shown in Figure 8a. As a result of using the ℓ^2 norm as the distance measure, the clusters form circles around each of the centroids. In Figure 8b, we extend the Euclidean-based clustering, using the pullback of the ℓ^2 norm as the distance metric used for clustering, i.e, clustering in three-dimensional space. We observe that clusters are disjointed and overlapping; a characteristic that is unbecoming of clustering methods. Finally, the proposed clustering method is shown in Figure 8c. We observe that points are clustered based on geodesic distance, conforming to the geometry of the manifold. This is particularly noticeable for the dark-blue cluster in the lower right which extends into the centre of the domain. The deployment of the metric-augmented Eikonal solver in Figure 8c exemplifies the improvements afforded by the proposed methodology. Thanks to the differential nature of the obtained distance function, the Fréchet mean, as shown in §7, can be obtained

with ease and for relatively low computational expense. Notably, this example seeks to obtain five clusters over 2500 samples. A total of 20 iterations were required to reach convergence for the clustering, comprising 100 evaluations of the Fréchet mean. This took just 113.084s of wall-clock time on a single NVIDIA GeForce RTX 4090.




Figure 8: Clustering on the peaks manifold. Panel (a) shows euclidean-based clustering of the samples directly on the manifold. Panel (b) shows the equivalent clustering, where distance is computed with the pullback of the Euclidean ℓ^2 norm. Panel (c) shows the clustering results for the metric-constrained Eikonal approach. Red points, in each case, depict the centroids of the clusters.

6 Conclusion

In this work, we address the problem of computing geodesic distances on manifolds. First, we introduce the metric-constrained Eikonal equation as a means for representing distance. Second, we propose a model-based parameterisation of the distance function providing a continuous, differentiable representation with theoretical guarantees. Obtaining the distance function in this manner is the main contribution of the manuscript; ultimately, alleviating issues encountered with all methods currently described in the literature. Third, we demonstrate the efficacy of our method for obtaining distance functions on a range of manifolds. A comparison is made to traditional methods in both cases, with our manifold-based Eikonal solver providing faster inference, and upholding theoretic guarantees in all cases. Finally, we tackle two key applications, which are the computation of the Fréchet mean and k -means clustering on manifolds. These applications have been computationally infeasible until now as a result of the comparatively expensive methods currently available in the literature. This work opens opportunities for reduced-order modelling, statistical modelling on manifolds, and distance-based computations in general.

Acknowledgments

D.K. is supported by a PhD studentship sponsored by the Aeronautics department at Imperial College London. L.M. acknowledges support from the ERC Starting Grant No. PhyCo 949388 and EU-PNRR YoungResearcher TWIN ERC-PI 0000005. The authors thank Prof. Marika Taylor (Pro-Vice Chancellor and Head of College of Engineering and Physical Sciences at the University of Birmingham) for fruitful discussions on manifolds and geodesics.

References

- 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.
- Nutan Chen, Alexej Klushyn, Richard Kurlle, Xueyan Jiang, Justin Bayer, and Patrick van der Smagt. Metrics for Deep Generative Models, February 2018.

-
- 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 0730-0301, 1557-7368. doi: 10.1145/2516971.2516977.
- Nicki S. Detlefsen, Alison Pouplin, Cilie W. Feldager, Cong Geng, Dimitris Kalatzis, Helene Hauschultz, Miguel González-Duque, Frederik Warburg, Marco Miani, and Søren Hauberg. StochMan. *GitHub. Note*: <https://github.com/MachineLearningLifeScience/stochman/>, 2021.
- E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, December 1959. ISSN 0945-3245. doi: 10.1007/BF01386390.
- Manfredo Perdigão do Carmo. *Riemannian Geometry*. Mathematics: Theory & Applications. Birkhäuser, Boston Basel Berlin, 14., corrected print edition, 2013. ISBN 978-0-8176-3490-2 978-3-7643-3490-1.
- Lawrence C. Evans. *Partial Differential Equations*. American Mathematical Soc., 2010. ISBN 978-0-8218-4974-3.
- Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, October 2016. ISSN 0894-0347, 1088-6834. doi: 10.1090/jams/852.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2023.
- Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit Geometric Regularization for Learning Shapes, July 2020.
- Serafim Grubas, Anton Duchkov, and Georgy Loginov. Neural Eikonal solver: Improving accuracy of physics-informed neural networks for solving eikonal equation in case of caustics. *Journal of Computational Physics*, 474:111789, February 2023. ISSN 0021-9991. doi: 10.1016/j.jcp.2022.111789.
- Nicolas Guigui, Nina Miolane, and Xavier Pennec. 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.
- W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):97–109, 1970. ISSN 0006-3444. doi: 10.2307/2334940.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, January 1989. ISSN 0893-6080. doi: 10.1016/0893-6080(89)90020-8.
- Stephan Huckemann and Herbert Ziezold. Principal component analysis for Riemannian manifolds, with an application to triangular shape spaces. *Advances in Applied Probability*, 38(2):299–319, June 2006. ISSN 0001-8678, 1475-6064. doi: 10.1239/aap/1151337073.
- Xin Jin and Jiawei Han. K-Means Clustering. In Claude Sammut and Geoffrey I. Webb (eds.), *Encyclopedia of Machine Learning*, pp. 563–564. Springer US, Boston, MA, 2010. ISBN 978-0-387-30164-8. doi: 10.1007/978-0-387-30164-8_425.
- R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, July 1998. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.95.15.8431.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017.
- John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 2012. ISBN 978-1-4419-9981-8 978-1-4419-9982-5. doi: 10.1007/978-1-4419-9982-5.

-
- John M. Lee. *Introduction to Riemannian Manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer International Publishing, Cham, 2018. ISBN 978-3-319-91754-2 978-3-319-91755-9. doi: 10.1007/978-3-319-91755-9.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, April 2021.
- Binbin Lin, Ji Yang, Xiaofei He, and Jieping Ye. Geodesic Distance Function Learning via Heat Flow on Vector Fields, May 2014.
- Stuart Lloyd. Least squares quantization in PCM. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- Qian Ni and Xuhui Wang. Shape Analysis by Computing Geodesics on a Manifold via Cubic B-splines. *Communications in Mathematics and Statistics*, October 2023. ISSN 2194-671X. doi: 10.1007/s40304-023-00373-3.
- Xavier Pennec. Intrinsic Statistics on Riemannian Manifolds: Basic Tools for Geometric Measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, July 2006. ISSN 0924-9907, 1573-7683. doi: 10.1007/s10851-006-6228-4.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763. PMLR, July 2021.
- Marko Rajković. *Geodesics, Splines, and Embeddings in Spaces of Images*. Thesis, Universitäts- und Landesbibliothek Bonn, August 2023.
- J. A. Sethian. Fast Marching Methods. *SIAM Review*, 41(2):199–235, 1999. ISSN 0036-1445.
- 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. ISSN 1097-0312. doi: 10.1002/cpa.3160200210.
- Robert M Wald. *General Relativity*. University of Chicago press, 2010.
- Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and Mitigating Gradient Flow Pathologies in Physics-Informed Neural Networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, January 2021. ISSN 1064-8275. doi: 10.1137/20M1318043.
- Sifan Wang, Shyam Sankaran, Hanwen Wang, and Paris Perdikaris. An Expert’s Guide to Training Physics-informed Neural Networks, August 2023.
- Wei Wang, Yan Huang, Yizhou Wang, and Liang Wang. Generalized autoencoder: A neural network framework for dimensionality reduction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 490–497, 2014.
- Zhenyue Zhang and Hongyuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM journal on scientific computing*, 26(1):313–338, 2004.
- Hongkai Zhao. A fast sweeping method for eikonal equations. *Mathematics of computation*, 74(250):603–627, 2005.
- Hongkai Zhao. Parallel implementations of the fast sweeping method. *Journal of Computational Mathematics*, pp. 421–429, 2007.

A Imposing upper bounds in network augmentation

Considering points $p, q \in M$, there must exist no shorter curve connecting the points than the length-minimising geodesic, as described in Eq. (9). Under this assumption, the length of an arbitrary curve γ_L that connects the points must be greater than, or equal to, the length of the minimising geodesic. We can prescribe the bounds

$$(\iota^* d_E)(p, q) \leq d(p, q) \leq \int_0^1 \langle \dot{\gamma}_L(\lambda), \dot{\gamma}_L(\lambda) \rangle_{\gamma_L(\lambda)} d\lambda. \quad (29)$$

We can use these bounds to alter the way in which we augment the network outputs. Empirically, prescribing an upper-bound incurs additional computational cost with little improvement in performance. For this reason, we do not show results for upper-bounded augmentations in the paper. Choosing an arbitrary curve between two points prescribes this upper bound. In order to reliably obtain curves, we can linearly interpolate between points in the intrinsic coordinates, that is:

$$\gamma_L(t) = q + (p - q)t \quad \text{where, } t \in [0, 1], \quad (30)$$

allowing us to compute the upper-bounded length as

$$L_\gamma^+ = \int_0^1 \langle \dot{\gamma}_L(t), \dot{\gamma}_L(t) \rangle_{\gamma(t)}^{1/2} dt \quad (31)$$

or,

$$L_\gamma^+ = \int_0^1 \langle p - q, p - q \rangle_{q+(p-q)t}^{1/2} dt \quad (32)$$

Using this definition, we can alter how we augment the network output to provide a bounded solution. For the single-point solution, we would write

$$\varphi_\theta(q; p) = (\iota^* d_E)(p, q) \left[1 + \left(\frac{L_\gamma^+}{(\iota^* d_E)(p, q)} - 1 \right) (\sigma \circ \tilde{\varphi})(q; p) \right], \quad (33)$$

where σ now represents the **sigmoid** function, and L_γ^+ is the length of the linearly interpolated curve in intrinsic coordinates. We can extend this approach to the two-point solution

$$\varphi_\theta(p, q) = (\iota^* d_E)(p, q) \left[1 + \left(\frac{L_\gamma^+}{(\iota^* d_E)(p, q)} - 1 \right) \sigma \left(\frac{1}{2} (\tilde{\varphi}_\theta(p, q) + \tilde{\varphi}_\theta(q, p)) \right) \right], \quad (34)$$

With these approaches, we have retained the same guarantees as the network augmentations described in the main text, as well as provided an upper-bound for the distance.

B Training parameters

In each of the cases, we employ a modified MLP architecture. The architecture has a depth of 8 layers, each with a width of 256 neurons. Optimisation was conducted with the **adam** optimiser. The optimiser uses an exponentially decaying learning rate schedule, reducing by a factor 0.9 every 2000 parameter updates.

C Hypersphere

For the hypersphere, we define a coordinate system with $n - 1$ angular coordinates x_1, \dots, x_{n-1} , with angles $x_1, \dots, x_{n-2} \in [0, \pi]$, and $x_{n-1} \in [0, 2\pi)$ radians. The immersion is the mapping $\iota : x \mapsto y$, where

$$y_1 = \cos(x_1) \quad (35)$$

$$y_2 = \sin(x_1) \cos(x_2) \quad (36)$$

$$y_3 = \sin(x_1) \sin(x_2) \cos(x_3) \quad (37)$$

$$\vdots \quad (38)$$

$$y_{n-1} = \sin(x_1) \dots \sin(x_{n-2}) \cos(x_{n-1}) \quad (39)$$

$$y_n = \sin(x_1) \dots \sin(x_{n-2}) \sin(x_{n-1}). \quad (40)$$