

Manifold-augmented Eikonal Equations: Geodesic Distances and Flows on Differentiable Manifolds.

Daniel Kelshaw

Department of Aeronautics, Imperial College London

DJK21@IC.AC.UK

Luca Magri

*Department of Aeronautics, Imperial College London
The Alan Turing Institute, British Library*

L.MAGRI@IMPERIAL.AC.UK

Editors: Sophia Sanborn, Christian Shewmake, Simone Azeglio, Nina Miolane

Abstract

Manifolds discovered by machine learning models provide a compact representation of the underlying data. Geodesics on these manifolds define locally length-minimising curves and provide a notion of distance, which are key for reduced-order modelling, statistical inference, and interpolation. In this work, we propose a model-based parameterisation for distance fields and geodesic flows on manifolds, exploiting solutions of a manifold-augmented Eikonal equation. We demonstrate how the geometry of the manifold impacts the distance field, and exploit the geodesic flow to obtain globally length-minimising curves directly. This work opens opportunities for statistics and reduced-order modelling on differentiable manifolds.

Keywords: Riemannian manifolds, Geodesic, Distance

1. Introduction

The manifold hypothesis states that many high-dimensional systems are expressible 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. These distance measures are of particular importance for computing statistics on manifolds, or for reduced-order modelling (Guigui et al., 2023; Pennec, 2006).

Related Work. Standard methods compute geodesic distances on the manifold by parameterising and minimising the energy of paths between points (Chen et al., 2018; Kalatzis et al., 2020); however, these methods yield pairwise distances rather than a continuous distance function. Solutions to the Eikonal equation can provide such functions, but prior applications on manifolds are limited. Euclidean methods presented by the likes of Gropp et al. (2020) do not handle nonlinear manifolds. Works by Crane et al. (2013); Spira and Kimmel (2004) develop discrete solutions, which prove intractable in high-dimensions and do not offer differentiable representations. To our knowledge, no existing work solves a manifold-augmented Eikonal equation to obtain a continuous, differentiable representation of the distance function. Differentiable representations are crucial for obtaining geodesic flows, which characterise globally shortest-paths on the manifold, thus providing an alternative method for computation of geodesics (do Carmo, 2013; Lee, 2018).

Contributions. We first provide a background of material required to specify the Eikonal equation on the manifold in §2. Second, in §3 we describe a method for obtaining solutions to a manifold-augmented Eikonal equation, providing a continuous, differentiable distance function on the manifold. Our proposed approach allows us to obtain length-minimising paths directly. Third, in §4 we showcase results on a complex manifold with many local extrema. We demonstrate the efficacy of our approach through the successful recovery of globally length-minimising paths. Finally, we discuss limitations and future work in §5.

2. Background

In this section, we provide a brief overview of concepts from differential geometry which form the building blocks of our proposed methodology. Note that, throughout this paper, we employ the Einstein summation convention (do Carmo, 2018; Lee, 2012).

Riemannian manifolds and the 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 choice of Riemannian metric on M . This metric constitutes a symmetric bilinear form $g : T_p M \times T_p M \rightarrow \mathbb{R}$, allowing for computation of the inner product on the tangent space,

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

Defining the induced metric. We consider the case of embedded submanifolds. Suppose (\tilde{M}, \tilde{g}) is a Riemannian manifold, and $M \subseteq \tilde{M}$ is an embedded submanifold. Given a smooth immersion $\iota : M \hookrightarrow \tilde{M}$, the metric $g = \iota^* \tilde{g}$ is referred to as the metric induced by ι , where ι^* is the pullback. 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 defined as

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

Defining distance on the manifold. While the induced metric g provides the means to compute distances on the tangent space, the notion of distance on the manifold arises from considering the existence of shortest-paths, known as geodesics (Lee, 2018, 2012; do Carmo, 2018, 2013). These geodesic curves, γ , satisfy the second-order ordinary differential equation

$$\nabla_{\dot{\gamma}} \dot{\gamma} = \ddot{\gamma}^k + \Gamma^k_{ij} \dot{\gamma}^i \dot{\gamma}^j = 0, \quad (3)$$

where ∇ denotes the covariant derivative, and Γ^k_{ij} are the Christoffel symbols. While 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; precisely

$$d_g(p, q) = \inf_{\gamma} \left\{ \int_0^1 \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle_{\dot{\gamma}(t)}^{1/2} dt : \gamma(0) = p, \gamma(1) = q \right\}. \quad (4)$$

Conversely, this implies that the shortest curve between two points is itself a geodesic.

3. Methodology

We propose a new method for computing geodesic distances, theoretically yielding the ability to compute globally length-minimising geodesics directly. We consider a function $\varphi : q \mapsto d_g(p, q)$, which computes the geodesic distance between points $q, p \in M$. This distance function φ satisfies the manifold-augmented Eikonal equation

$$\|\nabla\varphi\| = \langle \text{grad } \varphi, \text{grad } \varphi \rangle_g = \varphi^i \varphi_{,i} = 1 \quad \text{s.t.} \quad \varphi|_p = 0. \quad (5)$$

where $(\text{grad } \varphi)^i = g^{ij} \varphi_{,j}$ is the exterior derivative; g^{ij} is the inverse of the metric tensor g_{ij} ; $\varphi_{,i}$ denotes the derivative along the i th coordinate vector field; and $\varphi^i = g^{ij} \varphi_{,j}$. Notably, the exterior derivative of the distance field provides the geodesic flow, a section of the tangent bundle (Lee, 2018, Thm. 6.31, 6.32). This flow satisfies the geodesic equation (3),

$$\nabla_{\text{grad } \varphi} \text{grad } \varphi = 0, \quad (6)$$

where integral paths constitute unit-distance geodesics travelling orthogonal to the level-sets imposed by the distance field. Obtaining the distance function φ_θ yields a method to compute globally length-minimising geodesics directly, alleviating the need to parameterise length-minimising curves between pairs of points.

Neural Network Parameterisation. We parameterise the distance function by a neural network $\tilde{\varphi}_\theta : M \rightarrow \mathbb{R}$, providing a continuous, differentiable representation of the solution. The output of the network is augmented to yield the distance function

$$\varphi_\theta(q; p) = (\iota^* d_E)(p, q) [1 + (\sigma \circ \tilde{\varphi}_\theta)(q; p)], \quad (7)$$

where $\sigma \in \mathcal{C}^{\geq 2}(\mathbb{R}, \mathbb{R}_{\geq 0})$ is strictly increasing; and $d_E : (p, q) \mapsto \|p - q\|_2$ is the Euclidean distance function. Defining the distance function in this manner serves two distinct purposes: (i) the constraint $\varphi|_p = 0$ is handled implicitly; and (ii) the inequality $\varphi(q; p) \geq (\iota^* d_E)(p, q)$ is guaranteed (do Carmo, 2013). We seek a solution $\varphi_{\theta^*} = d_g(q, p)$ through solving

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta; p) \quad \text{where } \mathcal{L} = \mathbb{E}_q (\|\varphi_{\theta^i}(q) \varphi_{\theta,i}(q) - 1\|^2), \quad (8)$$

where the expectation is taken with respect to some distribution $q \sim \mathcal{D}$ in M .

Curvature-based Sampling. The solution φ to the Eikonal equation is more sensitive in regions of high curvature, a consequence of induced geodesic deviation (do Carmo, 2013). We mitigate the impact of geodesic deviation through employing a curvature-based sampling strategy, which ensures 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).¹

1. All code is publically available: <https://github.com/danielkelshaw/riemax>

4. Results

In order to provide an illustrative example, we demonstrate results for a function $\iota : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, where intrinsic coordinates are interpretable in two-dimensional space. We choose a non-convex manifold with multiple local extrema, exhibiting non-negligible scalar curvature; namely, the peaks function (MathWorks). Parameterising the distance function φ_θ with a neural network, we employ the `adam` optimiser (Kingma and Ba, 2017), followed by L-BFGS (Liu and Nocedal, 1989), to seek a solution to the optimisation problem (8). Further details on the network architecture and optimisation scheme can be found in Appendix A.

We first use our proposed sampling method, in conjunction with kernel density estimation (Rosenblatt, 1956), to construct a probability density function (PDF) correlated to the magnitude of the scalar curvature; shown in Figure 1(a). During optimization, our sample distribution \mathcal{D} is defined by the mean of the curvature PDF and a uniform distribution.

Results in Figure 1(b) 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, showing a tendency to circumvent local extrema where possible. 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 (3) 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, 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 depicted in Figure 1(c). We observe that as the distance increases, the variability in predicted distance increases correspondingly; effectively, compounding the error. This result is not unexpected, as distance predictions are naturally reliant on predictions closer to the fixed origin.

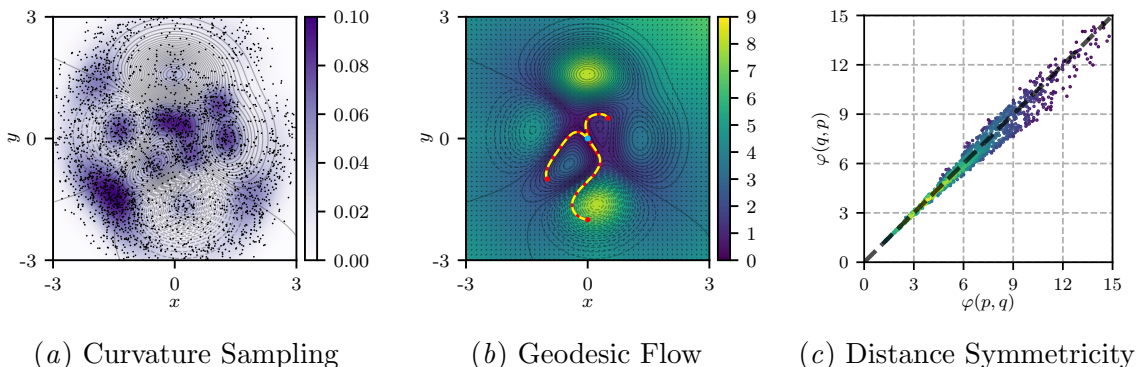


Figure 1: Results for the distance function, and geodesic flow on the peaks manifold.

5. Conclusion

In this work, we address the problem of computing geodesic distance on manifolds. First, we introduce the manifold-augmented Eikonal equation as a means of representing distance functions on the manifold. Second, we propose a model-based parameterisation of this distance function, providing a continuous, differentiable representation with theoretic guarantees. Third, we demonstrate the efficacy of our proposed methodology on a complex manifold, showcasing the ability to recover globally length-minimising paths through leveraging the geodesic flow. Future work will focus on improving robustness as geodesic distance increases, and removing the fixed origin limitation. This work opens opportunities for enhanced distance computation in statistics and reduced-order modelling.

Acknowledgments

D. Kelshaw. and L. Magri. acknowledge support from the UK EPSRC. L. Magri gratefully acknowledges financial support from the ERC Starting Grant PhyCo 949388.

References

- Nutan Chen, Alexej Klushyn, Richard Kurle, 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.
- 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.
- Manfredo Perdigão do Carmo. *Differential Geometry of Curves & Surfaces*. Dover Publications, INC, Mineola, New York, revised & updated second edition edition, 2018. ISBN 978-0-486-80699-0.
- 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.
- Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit Geometric Regularization for Learning Shapes, July 2020.
- 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 Keith Hastings. Monte Carlo sampling methods using Markov chains and their applications. 1970.

Dimitrios Kalatzis, David Eklund, Georgios Arvanitidis, and Soren Hauberg. Variational Autoencoders with Riemannian Brownian Motion Priors. In *Proceedings of the 37th International Conference on Machine Learning*, pages 5053–5066. PMLR, November 2020.

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.

Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, August 1989. ISSN 1436-4646. doi: 10.1007/BF01589116.

MathWorks. Peaks function - MATLAB peaks. <https://www.mathworks.com/help/matlab/ref/peaks.html#>.

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.

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 1573-7683. doi: 10.1007/s10851-006-6228-4.

Murray Rosenblatt. Remarks on Some Nonparametric Estimates of a Density Function. *The Annals of Mathematical Statistics*, 27(3):832–837, September 1956. ISSN 0003-4851, 2168-8990. doi: 10.1214/aoms/1177728190.

Alon Spira and Ron Kimmel. An efficient solution to the eikonal equation on parametric manifolds. *Interfaces and Free Boundaries*, 6(3):315–327, September 2004. ISSN 1463-9963. doi: 10.4171/ifb/102.

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.

Appendix A. Network Architecture and Optimisation

Here, we provide a brief overview of the network architecture and optimisation scheme used throughout this paper.

Network Architecture We parameterise the distance function $\varphi_\theta : q \rightarrow d_g(p, q)$ with a neural network $\tilde{\varphi}_\theta$. The architecture of the neural network is characterised by a simple modification of a standard multi-layer perceptron (MLP), as proposed by Wang et al. (2021, 2023). For the purposes of this paper, we train the modified MLP with four hidden layers, each with 128 neurons. The `tanh` activation is used on all but the output layer. Inputs to the network are standardised according to the domain of interest.

Optimisation Scheme We employ a two-stage optimisation scheme: first using the gradient-descent optimiser, `adam` (Kingma and Ba, 2017); before finetuning the result with L-BFGS (Liu and Nocedal, 1989). When using the `adam` optimiser, we initialise the parameters with standard defaults: $\beta_1 = 0.9, \beta_2 = 0.999, \varepsilon = 10^{-8}, \bar{\varepsilon} = 0$. We define an initial learning rate of $\eta(0) = 10^{-3}$, exponentially decaying by a factor 0.9 every 2000 steps. This exponential decay only takes effect after a warm-up period of 5000 parameter updates, accelerating convergence. We train for a total of 10^5 parameter updates, randomly sampling a total of 2^{13} values of $q \in M$ each time. For the second stage of training, we use the L-BFGS optimiser for a further 1000 parameter updates. This approach stabilises training and improves convergence, avoiding local minima. All other training parameters remain the same. Training takes approximately 15mins on a single NVIDIA RTX 6000 GPU.

Curvature-based Sampling We use the Metropolis-Hastings sampling method to generate random samples based on the magnitude of the scalar curvature. In order to ensure the Markov chains are statistically representative, we discard the initial 5×10^4 samples. We also draw samples from $N = 5$ chains, each with a random initialisation to provide improved coverage of the manifold. In order to generate a probability density function, a kernel density estimate was obtained from the given samples.