
Learning Distances from Data with Normalizing Flows and Score Matching

Peter Sorrenson*
Heidelberg University

Daniel Behrend-Uriarte
Heidelberg University

Christoph Schnörr
Heidelberg University

Ullrich Köthe
Heidelberg University

Abstract

Density-based distances (DBDs) offer an elegant solution to the problem of metric learning. By defining a Riemannian metric which increases with decreasing probability density, shortest paths naturally follow the data manifold and points are clustered according to the modes of the data. We show that existing methods to estimate Fermat distances, a particular choice of DBD, suffer from poor convergence in both low and high dimensions due to i) inaccurate density estimates and ii) reliance on graph-based paths which are increasingly rough in high dimensions. To address these issues, we propose learning the densities using a normalizing flow, a generative model with tractable density estimation, and employing a smooth relaxation method using a score model initialized from a graph-based proposal. Additionally, we introduce a dimension-adapted Fermat distance that exhibits more intuitive behavior when scaled to high dimensions and offers better numerical properties. Our work paves the way for practical use of density-based distances, especially in high-dimensional spaces.

1 Introduction

Metric learning aims to determine a distance metric that accurately measures the similarity or dissimilarity between data points. Typically, data points are mapped into a space where simple metrics like Euclidean distance can be applied. Although this method is computationally efficient, it is limited by the configurations possible in Euclidean space, where arbitrary distance relationships between sets of points cannot be represented.

A more flexible but computationally expensive approach involves defining a Riemannian metric in the data space and solving for geodesic distances. This allows for arbitrary distance relationships, but it requires choosing an appropriate Riemannian metric. Fermat distances [Groisman et al., 2022], which are a type of density-based distance [Bousquet et al., 2003], offer an elegant solution: the **metric should be conformal (a multiple of the identity matrix)** and inversely proportional to the probability density. In this way, geodesics pass through high-density regions of the data, respecting any inherent manifold structure. This approach is consistent with Fermat’s principle of least time in optics, where the inverse density functions like a refractive index.

To address the computational challenge of computing geodesics with this metric, previous work has approximated trajectories using shortest paths in a graph of nearest neighbors [Bijral et al., 2012]. The edge weights are approximated by raising Euclidean distances to an appropriate power, with theoretical analysis supporting the consistency of this approximation. However, in practice, we

*Correspondence to peter.sorrenson@iwr.uni-heidelberg.de.

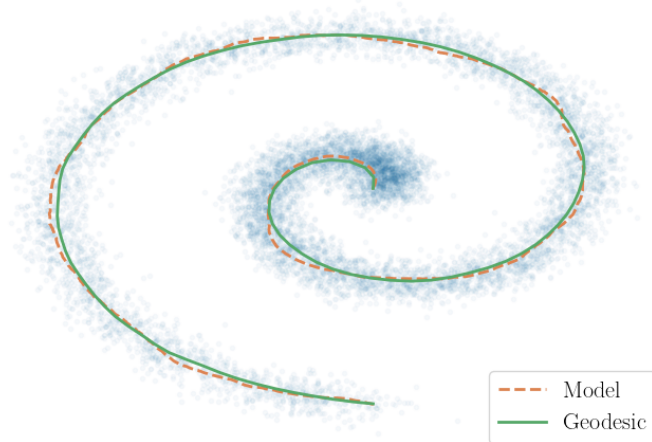


Figure 1: Ground truth and estimated geodesics on a 2-dimensional spiral dataset. The green line shows a ground truth geodesic, solved via relaxation on the true density (Algorithm 1). The dashed orange line is an estimate for the geodesic, found as the shortest path in a nearest neighbor graph with edge weights calculated from a normalizing flow density (Algorithm 2).

observe that convergence quickly plateaus, showing no significant improvement with increasing sample size.

By **employing more accurate edge weights obtained through normalizing flows**, we achieve much faster convergence. Additionally, previous approaches did not account for the limitations of graph-based methods in higher dimensions, where paths become increasingly rough due to the curse of dimensionality. To address this, we use a **relaxation scheme to smooth trajectories** and maintain satisfactory convergence. We find that **training models via score matching is more effective for this purpose than using normalizing flows**.

In summary, our contributions are threefold: First, we introduce the use of normalizing flows to obtain more accurate edge weights for estimating Fermat distances, significantly improving convergence rates. Second, we address the limitations of graph-based methods in high-dimensional spaces by implementing a smooth relaxation scheme using score matching, which better maintains the continuity and smoothness of the trajectories. Finally, we propose a dimension-adapted Fermat distance that scales intuitively with increasing dimensionality and exhibits superior numerical properties, paving the way for practical applications of density-based distances in complex, high-dimensional datasets. Through these innovations, we demonstrate the practicality and effectiveness of our approach, opening up a new possibility for metric learning using density-based distances in high-dimensional spaces.

2 Related work

Bousquet et al. [2003] introduces the idea of changing the geometry of the data space by use of a conformal density-based metric in order to get more informative distance functions. However they do not attempt to calculate geodesics or distances implied by this geometry, instead linking it to a modified SVM algorithm used for semi-supervised learning. Sajama and Orlitsky [2005] also uses a density-based metric for semi-supervised learning. They estimate densities using kernel density estimation and construct a graph of nearest neighbor distances from which they compute geodesics.

Bijral et al. [2012] continues this line of work by approximating the density as inversely proportional to a power of the Euclidean distance between close neighbors. Hwang et al. [2016] gives convergence guarantees for the power-weighted shortest path distances to geodesic distances under the density-dependent metric. Several other papers give similar results, and some practical implementations [Chu et al., 2020, Little et al., 2022, Groisman et al., 2022, Moscovich et al., 2017, Alamgir and Von Luxburg, 2012, Mckenzie and Damelin, 2019, Little et al., 2020]. Terms used are power-weighted

shortest path distances (PWSPD) [Little et al., 2022] and Fermat distances [Groisman et al., 2022, Trillos et al., 2023].

One crucial contribution of our work is to check that proposed methods actually converge in practice, regardless of theoretical guarantees. Previous work does not compare shortest paths to ground truth geodesics, even for a uniform distribution, where geodesics are simply straight lines. We address this major shortcoming by computing the ground truth geodesics for problems with known density and find that existing methods show very poor convergence.

3 Background

3.1 Riemannian geometry

Riemannian geometry offers powerful tools for analyzing curved spaces, which can be particularly useful for understanding and optimizing over non-Euclidean spaces such as manifolds. A fundamental object in this field is the metric tensor, where different choices of metric can lead to vastly different geometric structures, such as positive curvature (e.g., a sphere which closes in on itself) or negative curvature (e.g., hyperbolic space which expands exponentially quickly), as well as more complicated mixtures of the two.

More formally, a Riemannian manifold \mathcal{M} is a smooth n -dimensional space that locally resembles \mathbb{R}^n . At each point p on the manifold, there is an associated tangent space \mathcal{T}_p , which is a linear approximation of the manifold at p . The metric tensor g is a smoothly varying positive-definite bilinear form on the tangent space. It defines an inner product in the tangent space, allowing us to measure lengths and angles.

For a tangent vector $v \in \mathcal{T}_p$, the norm is defined as $\|v\| = \sqrt{g(v, v)}$, providing a notion of length in the tangent space. The length of a smooth curve $\gamma(t)$ on the manifold is given by $\int \sqrt{g(\dot{\gamma}, \dot{\gamma})} dt$, where $\dot{\gamma}$ is the derivative of $\gamma(t)$ with respect to t . This integral measures the total length of the curve by summing up the infinitesimal lengths along the curve.

The distance between two points on the manifold is defined as the minimal length of a curve connecting the points. The curves that minimize this length are known as geodesics, which are the generalization of straight lines in Euclidean space to curved spaces.

3.2 Fermat distances

Density-based distances (DBDs) provide a way to define distances between points in a distribution by utilizing the underlying probability density function. This approach can be particularly useful for tasks such as clustering and pathfinding in data analysis and machine learning. If designed well, a DBD can capture desirable properties such as:

1. Shortest paths between points should pass through supported regions of the data. If there is a region with no data between two points, the shortest path should deviate around it to pass through regions with more data.
2. Data should cluster according to its modes. If two data points belong to different modes with little data connecting them, the distance between them should be high.

One way to define a DBD that satisfies these properties is through Fermat distances. Named after Fermat's principle of least time in optics, these distances model the shortest paths between points by considering the inverse probability density (or a monotonic function of it) as analogous to the refractive index.

In the language of Riemannian geometry, we define a metric tensor of the form

$$g(u, v) = \frac{\langle u, v \rangle}{f(p)^2}, \tag{1}$$

where p is the probability density, f is a monotonic function, and $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. This formulation ensures that regions with higher probability density have lower "refractive index," guiding shortest paths through denser regions of the data. In practice, it is convenient to choose




Figure 2: Geodesics of the standard normal distribution, projected into 2d, taken between two orthogonal points at distance \sqrt{D} from the origin. The paths are scaled by \sqrt{D} to enable comparison. (*Left*) Unscaled metric (same temperature for all dimensions) leads to sharply curved geodesics in high dimensions. (*Right*) Scaled metric (temperature equal to dimension) leads to consistent geodesics in all dimensions (note that the lines are overlapping).

$f(p) = p^\beta$, where β is an inverse temperature parameter that controls the influence of the density on the metric. Since our metric enacts a simple scaling of the Euclidean inner product, it is a conformal metric, which locally preserves angles and simplifies key geometric objects such as the geodesic equation and scalar curvature.

The length of a smooth curve $\gamma(t)$ is therefore given by

$$L(\gamma) = \int \frac{\|\dot{\gamma}(t)\|}{p(\gamma(t))^\beta} dt \quad (2)$$

3.2.1 Fermat distances which scale with dimension

If $\beta = 1$, this leads to unintuitive behavior in higher dimensions. Consider the standard normal distribution in D dimensions. Due to rotational symmetry, the geodesic connecting x_1 and x_2 will lie in the plane spanned by these two points and the origin. Without loss of generality, suppose this plane aligns with the first two axis directions. The density restricted to this plane is the same as in the two-dimensional case, up to a constant factor. Therefore, geodesics in the D -dimensional standard normal distribution are the same as in the two-dimensional standard normal distribution, once we account for the rotation into the plane of x_1 and x_2 .

However, despite this similarity, the typical distance from the origin of points sampled from the D -dimensional standard normal distribution is much larger, with a mean of \sqrt{D} . As a result, geodesics will start and end in very low-density regions relative to the two-dimensional standard normal. This means that most of the trajectory is spent moving towards and then outwards from the high-density region near the origin.

To address this issue, if we scale the temperature linearly with dimension ($\beta = 1/D$), the results become consistent with the two-dimensional case (see Fig. 2). Scaling the temperature with dimension also increases numerical stability. This is because the effect of the probability density is reduced, leading to more Euclidean-like geodesics. This effect is clearly visible in Fig. 2, where the geodesics are much closer to straight lines in the scaled case.

3.3 Geodesic equation

From the metric, we can compute the geodesic equation. This is the second-order differential equation which all paths of minimum length (geodesics) must satisfy. In our case, we have the relatively simple solution for a trajectory $\gamma(t)$:

$$\ddot{\gamma} - 2\beta(s(\gamma) \cdot \dot{\gamma})\dot{\gamma} + \beta s(\gamma)\|\dot{\gamma}\|^2 = 0 \quad (3)$$

where $s(x) = \frac{\partial \log p(x)}{\partial x}$ is the score of the distribution and $\|\cdot\|$ is the ordinary Euclidean norm.

Algorithm 1 One step of relaxation. $\varphi = \{\varphi_i\}_{i=0}^n$ is the current state of the trajectory, s is the score function, β the inverse temperature, $\gamma \in [0, 1)$ the smoothing parameter and $\alpha \in (0, 1]$ the step size. This step should be iterated until convergence, which can be measured via the difference between φ and φ' .

```

function RELAXATIONSTEP( $\varphi, s, \beta, \gamma, \alpha$ )
     $v_i \leftarrow \frac{1}{2}(\varphi_{i+1} - \varphi_{i-1})$   $\triangleright i = 1, \dots, n - 1$ 
     $\varphi'_i \leftarrow \frac{1}{2}(\varphi_{i+1} + \varphi_{i-1}) + \frac{\beta}{2} (s(\varphi_i) \|v_i\|^2 - (s(\varphi_i) \cdot v_i)v_i)$ 
     $\varphi'_i \leftarrow (1 - \gamma)\varphi'_i + \frac{\gamma}{2}(\varphi'_{i+1} + \varphi'_{i-1})$ 
     $\varphi'_i \leftarrow \alpha\varphi'_i + (1 - \alpha)\varphi_i$ 
    return  $\varphi'$ 
end function

```

The geodesic equation has the property that the speed is constant with respect to the metric, meaning $\sqrt{g(\dot{\gamma}, \dot{\gamma})}$ is constant. In other words, the geodesic distance traveled in a given time interval is the same at all points of the curve. If γ passes from a high-probability region to a low-probability one, the distances in the low-probability region could be orders of magnitude larger than in the high-probability region. This makes numerically solving the equations challenging due to the vastly different Euclidean velocities along the trajectory.

It is convenient to **reparameterize the equation such that the Euclidean velocity is constant instead**, as this is much easier to handle numerically. This is especially useful for solving via relaxation, where having equal-length segments leads to much more stable behavior. Making this change results in a very similar equation (removing only the factor of 2 in the second term):

$$\ddot{\varphi} - \beta(s(\varphi) \cdot \dot{\varphi})\dot{\varphi} + \beta s(\varphi) \|\dot{\varphi}\|^2 = 0 \quad (4)$$

See Appendix A.1 for derivations of these results.

4 Methods

4.1 Ground truth distances through relaxation

A relaxation scheme allows one to solve a differential equation constrained by boundary conditions by iteratively updating the solution such that it better obeys the differential equation. We divide the interval $[0, 1]$ into equal time steps of size $h = 1/n$ and write $\varphi_i = \varphi(ih)$. Initialize the values of the curve φ at $t = ih$ for $i = 0, \dots, n$, with the constraint that φ_0 and φ_n are the chosen endpoints. The initial curve could be the straight line connecting the end points, or a more informative guess if additional information is available.

Equation (4) leads to the following relaxation scheme. Define

$$v_i = \frac{\varphi_{i+1} - \varphi_{i-1}}{2} \quad (5)$$

then update with

$$\varphi'_i = \frac{\varphi_{i+1} + \varphi_{i-1}}{2} + \frac{\beta}{2} (s(\varphi_i) \|v_i\|^2 - (s(\varphi_i) \cdot v_i)v_i) \quad (6)$$

See Appendix A.2 for a derivation.

4.1.1 Implementation

In practice, some additional features are needed to ensure stable convergence of the relaxation scheme, especially when the score function is not particularly smooth. Firstly, we found that oscillatory features commonly arise. To address this, we use smoothing after the update, returning a value which is interpolated between the value at the node and the average of its neighbors. Secondly, for rougher score functions we use a step size between 0 and 1, where 0 corresponds to no update and 1 corresponds to the ordinary relaxation step. See Algorithm 1.

4.2 Weighted graph algorithm based on Euclidean distance

If we have a large number of samples from $p(x)$, we can approximate the density between close neighbors x_1 and x_2 as constant and hence the shortest path between them as a straight line. The density is roughly proportional to the inverse of the Euclidean distance, raised to the power d , where d is the intrinsic dimension of the distribution [Bijral et al., 2012]:

$$p(\gamma) \propto \frac{1}{\|x_1 - x_2\|^d} \quad (7)$$

Hence the distance between x_1 and x_2 according to the metric is approximately proportional to a power of the Euclidean distance

$$\text{dist}(x_1, x_2) \propto \|x_1 - x_2\|^{\beta d + 1} \quad (8)$$

Bijral et al. [2012] therefore proposes to approximate shortest paths between more distant points by constructing a k -nearest neighbors graph of the samples with edge weights given by $\|x_1 - x_2\|^{\beta d + 1}$ and applying Dijkstra’s algorithm.

Hwang et al. [2016] gives consistency guarantees for this estimate, proving that a scaled version of it converges to the ground truth distance with a rate of $\exp(-\theta_0 n^{1/(3d+2)})$, with n the sample size, d the intrinsic dimension, and θ_0 a positive constant. Groisman et al. [2022] additionally show that the shortest paths converge to geodesics in a large sample size limit.

While this approach is appealing in its simplicity, can be easily implemented using standard libraries, and comes with theoretical guarantees, the density estimate is too poor to provide accurate estimates of Fermat distances in practice. In the experimental section 5.1.1 we show that paths computed with this approach do not converge to the ground truth geodesics as sample size increases, even in very simple 2-dimensional distributions such as the standard normal.

We are unsure why we observe such a large deviation from the theoretical prediction. It is possible that the unknown constant θ_0 is very small, leading to extremely slow convergence, despite the exponential function. We leave investigation into this point to future work.

4.3 Normalizing flows

Normalizing flows [Rezende and Mohamed, 2015, Kobyzev et al., 2020, Papamakarios et al., 2021] are a set of techniques to learn a probability density via a parameterized diffeomorphism f_θ which maps from data to latent space. The probability density is obtained through the change of variables formula:

$$p_\theta(x) = \pi(f_\theta(x)) \left| \det \left(\frac{\partial f_\theta}{\partial x} \right) \right|^{-1} \quad (9)$$

where π is a simple latent distribution. A common strategy to design normalizing flows is through coupling blocks [Dinh et al., 2014, 2016, Durkan et al., 2019], where f_θ is the composition of a series of blocks, each with a triangular Jacobian. This makes the determinant easy to calculate, and hence the above formula is tractable.

Normalizing flows are trained by minimizing the negative log-likelihood on a training set:

$$\mathcal{L} = \mathbb{E}_{p_{\text{data}}(x)} [-\log p_\theta(x)] \quad (10)$$

4.4 Density-weighted graph algorithm

In order to better estimate probability densities, we use a normalizing flow trained on the data and directly use the density predicted by the model. In order to approximate the geodesic distance between close neighbors, we assume that the geodesic connecting them is approximately a straight line, and estimate its length by a discrete approximation:

$$\text{dist}(x_1, x_2) \approx \sum_{i=1}^S \frac{\|y_i - y_{i-1}\|}{p_\theta(y_{i-1/2})} \quad (11)$$

Here the y_i are points which interpolate x_1 and x_2 with $y_0 = x_1$ and $y_S = x_2$, and $y_{i-1/2}$ is the midpoint of y_{i-1} and y_i .

Algorithm 2 Construction of a density-weighted graph which can subsequently be used to compute shortest paths and distances between points. X is the data matrix, p the probability density, β the inverse temperature, k the number of neighbors, and S the number of segments with which to approximate the integral. The computations can also be performed in the log space, returning log edge weights, which is more numerically stable in higher dimensions.

```

function DENSITYGRAPH( $X, p, \beta, k, S$ )
   $\mathcal{G} \leftarrow \text{KNN}(X, k)$                                 ▷ Construct  $k$ -nearest neighbor graph
  for  $(l, m) \in \mathcal{G}$  do                                    ▷ Iterate over edges
     $y_i \leftarrow (1 - \frac{i}{S})X_l + \frac{i}{S}X_m$                 ▷ For  $i = 0, \dots, S$ 
     $y_{i-1/2} \leftarrow \frac{1}{2}(y_{i-1} + y_i)$             ▷ For  $i = 1, \dots, S$ 
     $\mathcal{G}_{lm} \leftarrow \sum_{i=1}^S \|y_i - y_{i-1}\|/p(y_{i-1/2})^\beta$ 
  end for
  return  $\mathcal{G}$ 
end function

```

This allows us to take a k -nearest neighbors graph of the data and replace the edge weights by the approximate distance. This results in a sparse graph, to which we can apply Dijkstra’s algorithm to estimate shortest paths and distances between any two points. See Algorithm 2.

4.5 Relaxation with a learned score model

The score function is the gradient of the log-likelihood with respect to spatial inputs, i.e., $s(x) = \nabla_x \log p(x)$. Note the difference from the definition commonly used in statistic, where the derivative is with respect to model parameters.

Score matching [Hyvärinen, 2005] is a family of methods to estimate this function from data, including variants such as sliced score matching [Song et al., 2020a], denoising score matching [Vincent, 2011] and diffusion models [Song et al., 2020b].

Sliced score matching We can learn the score of $p(x)$ by minimizing the following objective (known as the score matching objective):

$$\mathbb{E}_{p(x)}[\text{tr}(\nabla_x s_\theta(x)) + \frac{1}{2}\|s_\theta(x)\|^2] \tag{12}$$

In high dimensions the trace term is expensive to evaluate exactly, so sliced score matching uses a stochastic approximation (the Hutchinson trace estimator [Hutchinson, 1989]):

$$\mathbb{E}_{p(x)p(v)}[v^T \nabla_x s_\theta(x)v + \frac{1}{2}\|s_\theta(x)\|^2] \tag{13}$$

where v is sampled from an appropriate distribution, typically standard normal.

Since only the score is needed for the relaxation algorithm described in Section 4.1, we can use a trained score model to perform relaxation to approximate geodesics. In practice, we initialize the path from a graph-based method to speed up convergence and avoid local minima.

5 Experiments

Performance metric After solving for a shortest path φ between two points x_1 and x_2 , we compute its length $L(\varphi)$ by using the ground truth density function, then compare this value to the ground truth distance between the two points:

$$M(\varphi) = \log \frac{L(\varphi)}{\text{dist}(x_1, x_2)} \tag{14}$$

Provided we have access to the ground truth density, this performance metric can be calculated independently of how φ is obtained, and has a value of zero only if φ is the ground truth geodesic. The metric will otherwise take positive values.

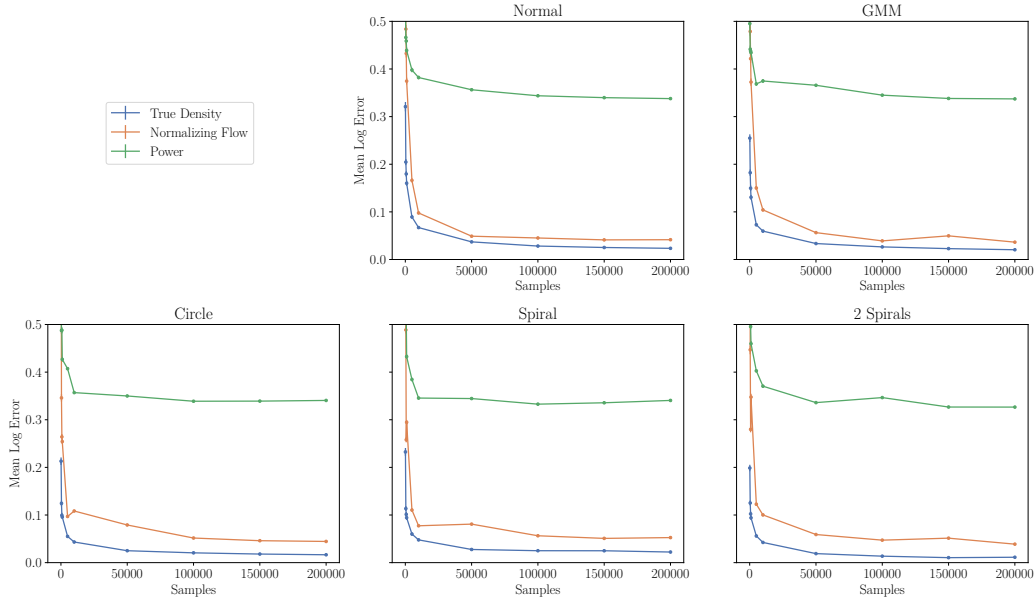


Figure 3: The length of shortest paths in power-weighted distance graphs do not converge to the ground truth with increasing sample size (green). In contrast, learned and ground-truth density weighted graphs (orange and blue) do converge, at similar rates.

5.1 Scaling the dataset

In this section we perform experiments on a set of 5 two-dimensional datasets for which we know the ground truth density (see Fig. 3). Details on how the datasets were created can be found in Appendix B.

5.1.1 Power-weighted graph paths do not converge with increasing sample size

We first estimate paths using the power-weighted graph method [Bijral et al., 2012], applied to a k-nearest neighbor graph. We find that in all 5 datasets, increasing the sample size leads to saturating performance (see Fig. 3, green lines). In Fig. 6 in the appendix we show that increasing the number of neighbors does not improve performance.

5.1.2 Poor convergence is due to poor density estimation

In order to investigate the root cause of the poor convergence, we designed other graph methods which use the nearest neighbor estimator for the density, but do not use power-weighted graph edges. We use 4 variants for the density estimate (see Appendix C.1). We find that all 4 have very similar performance to the power-weighted graph, whereas the equivalent estimators using the ground truth density quickly converge. See Appendix C.1 for plots supporting this claim.

5.1.3 Learning the density leads to convergent paths

Since the problem lies in poor density estimation, we hope to fix the problem by better density estimation. To this end, we use a normalizing flow to learn the density from samples, then use Algorithm 2 to construct a graph over which we compute shortest paths. We find that the paths obtained converge at close to the same rate as paths obtained by an equivalent method using the ground truth densities (Fig. 3 orange and blue lines). Full experimental details can be found in Appendix B.




Figure 4: Mean log error with a sample size of 200,000. Graph-based paths (the first three methods) converge very slowly in higher dimensions with increasingly worse performance as dimension increases. When using the scaled metric ($\beta = 1/D$), convergence improves over the unscaled metric, but all graph-based methods perform similarly in higher dimensions. Doing relaxation with a learned score function (red line) leads to much faster convergence, avoiding the curse of dimensionality.

5.2 Scaling the dimension

Machine learning datasets often lie in high-dimensional spaces, so it is crucial that any method applied to them scales well, avoiding the curse of dimensionality as much as possible. We test this by taking one of our two-dimensional datasets, namely the standard normal, and scaling it up to 25 dimensions. We perform experiments both for the scaled and unscaled version of the metric.

5.2.1 Graph-based paths do not converge well in higher dimensions

Our first finding is that all graph-based methods show declining performance as the dimension increases, for both unscaled and scaled metrics (Fig. 4). This can be straightforwardly explained by the curse of dimensionality: with increasing dimension but fixed sample size, it is increasingly unlikely that points in the dataset lie close to the true trajectory. The shortest path in the graph becomes increasingly rough in contrast to the smooth ground truth trajectory.

5.2.2 Relaxation using a learned score function improves convergence rates

In order to obtain smoother paths, we turn to the same relaxation algorithm we use to solve for ground truth geodesics. Instead of using the ground truth score function, we use a learned score function. An obvious choice would be to use automatic differentiation to get the score from a trained normalizing flow, but we find the resulting scores are too noisy to be useful, inhibiting reliable convergence of the relaxation scheme. Instead, we learn a score model with sliced score matching, which results in a smoother function suitable for performing relaxation. Full experimental details can be found in Appendix B.

We find that this fixes the convergence problems, both for the unscaled and scaled metrics, though convergence is more reliable in the scaled case (Fig. 4). Please see Fig. 8 and Fig. 9 in the appendix for more detailed evidence.

6 Conclusion

In conclusion, our study focuses on understanding the convergence properties of various algorithms to ground-truth geodesics by employing toy distributions with known densities. We investigate the standard normal distribution in dimensions higher than 2, highlighting the challenges faced by graph-based methods even in this simple scenario. Our findings demonstrate significant improvements through the use of normalizing flows and score models, yet further research is needed to address more

complex distributions and real-world data where the underlying density is unknown. Future work should aim to unify the roles of normalizing flows and score models into a single cohesive model, enhancing the method’s applicability and efficiency. Additionally, a deeper theoretical analysis is required to understand why our proposed methods converge effectively, while previous methods with consistency guarantees do not. This work lays the groundwork for practical applications of density-based distances, setting the stage for more robust and scalable solutions in metric learning.

References

- Morteza Alamgir and Ulrike Von Luxburg. Shortest path distance in random k-nearest neighbor graphs. *arXiv preprint arXiv:1206.6381*, 2012.
- Lynton Ardizzone, Till Bungert, Felix Draxler, Ullrich Köthe, Jakob Kruse, Robert Schmier, and Peter Sorrenson. Framework for Easily Invertible Architectures (FrEIA), 2018-2022. URL <https://github.com/vislearn/FrEIA>.
- Avleen S Bijral, Nathan Ratliff, and Nathan Srebro. Semi-supervised learning with density based distances. *arXiv preprint arXiv:1202.3702*, 2012.
- Olivier Bousquet, Olivier Chapelle, and Matthias Hein. Measure based regularization. *Advances in Neural Information Processing Systems*, 16, 2003.
- Timothy Chu, Gary L Miller, and Donald R Sheehy. Exact computation of a manifold metric, via Lipschitz embeddings and shortest paths on a graph. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 411–425. SIAM, 2020.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv preprint arXiv:1605.08803*, 2016.
- Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in neural information processing systems*, 32, 2019.
- Pablo Groisman, Matthieu Jonckheere, and Facundo Sapienza. Nonhomogeneous Euclidean first-passage percolation and distance learning. *Bernoulli*, 28(1):255–276, 2022.
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
- John D Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(03): 90–95, 2007.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Sung Jin Hwang, Steven B Damelin, and Alfred O Hero III. Shortest path through random points. 2016.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979, 2020.
- Anna Little, Mauro Maggioni, and James M Murphy. Path-based spectral clustering: Guarantees, robustness to outliers, and fast algorithms. *Journal of machine learning research*, 21(6):1–66, 2020.

- Anna Little, Daniel McKenzie, and James M Murphy. Balancing geometry and density: Path distances on high-dimensional data. *SIAM Journal on Mathematics of Data Science*, 4(1):72–99, 2022.
- Daniel Mckenzie and Steven Damelin. Power weighted shortest paths for clustering Euclidean data. *arXiv preprint arXiv:1905.13345*, 2019.
- Amit Moscovich, Ariel Jaffe, and Nadler Boaz. Minimax-optimal semi-supervised regression on unknown manifolds. In *Artificial Intelligence and Statistics*, pages 933–942. PMLR, 2017.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- Sajama and Alon Orlitsky. Estimating and computing density based distance metrics. In *Proceedings of the 22nd international conference on Machine learning*, pages 760–767, 2005.
- Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pages 574–584. PMLR, 2020a.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020b.
- Nicolás García Trillos, Anna Little, Daniel McKenzie, and James M Murphy. Fermat distances: Metric approximation, spectral convergence, and clustering algorithms. *arXiv preprint arXiv:2307.05750*, 2023.
- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, 17(3):261–272, 2020.

A Derivations

A.1 Geodesic equation

The geodesic equation which characterizes paths of zero acceleration γ is

$$\ddot{\gamma}^i + \Gamma_{jk}^i \dot{\gamma}^j \dot{\gamma}^k = 0 \quad (15)$$

where Γ_{jk}^i are the Christoffel symbols of the second kind

$$\Gamma_{jk}^i = \frac{1}{2} g^{il} \left(\frac{\partial g_{lj}}{\partial x^k} + \frac{\partial g_{lk}}{\partial x^j} - \frac{\partial g_{jk}}{\partial x^l} \right) \quad (16)$$

and g^{il} is the inverse of g , i.e. such that $g^{il} g_{lj} = \delta_j^i$.

Our metric has the form $g_{ij} = \lambda^2 \delta_{ij}$, meaning that

$$\Gamma_{jk}^i = \frac{1}{2\lambda^2} \delta^{il} \left(\frac{\partial \lambda^2}{\partial x^k} \delta_{lj} + \frac{\partial \lambda^2}{\partial x^j} \delta_{lk} - \frac{\partial \lambda^2}{\partial x^l} \delta_{jk} \right) \quad (17)$$

$$= \delta^{il} \left(\frac{\partial \log \lambda}{\partial x^k} \delta_{lj} + \frac{\partial \log \lambda}{\partial x^j} \delta_{lk} - \frac{\partial \log \lambda}{\partial x^l} \delta_{jk} \right) \quad (18)$$

$$= \frac{\partial \log \lambda}{\partial x^k} \delta_j^i + \frac{\partial \log \lambda}{\partial x^j} \delta_k^i - \frac{\partial \log \lambda}{\partial x^l} \delta^{il} \delta_{jk} \quad (19)$$

In this case $\lambda = p^{-\beta}$ and hence $\frac{\partial \log \lambda}{\partial x} = -\beta \frac{\partial \log p}{\partial x} = -\beta s$, where $s = \frac{\partial \log p}{\partial x}$ is the score function of the probability density. As a result

$$\Gamma_{jk}^i = -\beta s_k \delta_j^i - \beta s_j \delta_k^i + \beta s_l \delta^{il} \delta_{jk} \quad (20)$$

Substituting into the geodesic equation leads to

$$\ddot{\gamma}^i - \beta s_k \dot{\gamma}^i \dot{\gamma}^k - \beta s_j \dot{\gamma}^j \dot{\gamma}^i + \beta s_l \delta^{il} \delta_{jk} \dot{\gamma}^j \dot{\gamma}^k = 0 \quad (21)$$

which can be written in vector notation as

$$\ddot{\gamma} - 2\beta(s(\gamma) \cdot \dot{\gamma})\dot{\gamma} + \beta s(\gamma) \|\dot{\gamma}\|^2 = 0 \quad (22)$$

Let $f : [0, 1] \rightarrow [0, 1]$ be a strictly increasing diffeomorphism. Given a curve $\gamma : [0, 1] \rightarrow \mathbb{R}^n$, consider the curve $\varphi : [0, 1] \rightarrow \mathbb{R}^n$ defined such that $\varphi(u) = \gamma(f^{-1}(u))$ or equivalently $\varphi(f(t)) = \gamma(t)$. The derivatives of the two curves are related by

$$\dot{\varphi} f = \dot{\gamma} \quad (23)$$

and

$$\ddot{\varphi} f^2 + \dot{\varphi} \ddot{f} = \ddot{\gamma} \quad (24)$$

where derivatives of φ are with respect to $u = f(t)$. By substituting these expressions into the geodesic equation, we arrive at the equivalent equation for φ :

$$\ddot{\varphi} f^2 + \dot{\varphi} \ddot{f} - 2\beta(s(\gamma) \cdot \dot{\varphi} f) \dot{\varphi} f + \beta s(\gamma) \|\dot{\varphi} f\|^2 = 0 \quad (25)$$

We would like φ to have a constant Euclidean speed (the magnitude but not necessarily the direction of the velocity is constant), meaning

$$\frac{d}{du} \left(\frac{1}{2} \|\dot{\varphi}\|^2 \right) = \dot{\varphi} \cdot \ddot{\varphi} = 0 \quad (26)$$

Multiplying the geodesic equation by $\dot{\varphi}$, this is equivalent to requiring

$$\|\dot{\varphi}\|^2 \ddot{f} - 2\beta(s(\gamma) \cdot \dot{\varphi}) \|\dot{\varphi}\|^2 \dot{f}^2 + \beta(s(\gamma) \cdot \dot{\varphi}) \|\dot{\varphi}\|^2 \dot{f}^2 = 0 \quad (27)$$

implying that

$$\ddot{f} = \beta(s(\gamma) \cdot \dot{\varphi}) \dot{f}^2 \quad (28)$$

Resubstituting into the geodesic equation and dividing by \dot{f}^2 leads to the result:

$$\ddot{\varphi} - \beta(s(\varphi) \cdot \dot{\varphi}) \dot{\varphi} + \beta s(\varphi) \|\dot{\varphi}\|^2 = 0 \quad (29)$$

A.2 Relaxation scheme

The central finite difference approximations for the first and second derivatives are

$$\dot{\varphi}_i \approx \frac{\varphi_{i+1} - \varphi_{i-1}}{2h} \quad (30)$$

and

$$\ddot{\varphi}_i \approx \frac{\varphi_{i+1} - 2\varphi_i + \varphi_{i-1}}{h^2} \quad (31)$$

By substituting this into the differential equation we have

$$\frac{\varphi_{i+1} - 2\varphi_i + \varphi_{i-1}}{h^2} - \beta(s(\varphi_i) \cdot \dot{\varphi}_i) \dot{\varphi}_i + \beta s(\varphi_i) \|\dot{\varphi}_i\|^2 \approx 0 \quad (32)$$




Figure 5: The 5 two-dimensional datasets.

where $\dot{\varphi}_i$ is estimated using finite differences. We can rearrange this to

$$\varphi_i \approx \frac{\varphi_{i+1} + \varphi_{i-1}}{2} + \frac{h^2 \beta}{2} (s(\varphi_i) \|\dot{\varphi}_i\|^2 - (s(\varphi_i) \cdot \dot{\varphi}_i) \dot{\varphi}_i) \quad (33)$$

and use the equation as an update rule, updating each position of the curve except the endpoints at each iteration.

Dividing by very small h could lead to numerical instability. We can avoid dividing and multiplying by h by the following update rule. First define

$$v_i = \frac{\varphi_{i+1} - \varphi_{i-1}}{2} \quad (34)$$

then update with

$$\varphi_i = \frac{\varphi_{i+1} + \varphi_{i-1}}{2} + \frac{\beta}{2} (s(\varphi_i) \|v_i\|^2 - (s(\varphi_i) \cdot v_i) v_i) \quad (35)$$

B Experimental Details

Software libraries We used the `numpy` [Harris et al., 2020] and `scipy` [Virtanen et al., 2020] Python libraries for basic numerical and graph operations and `matplotlib` [Hunter, 2007] for plotting. We used `scikit-learn` [Pedregosa et al., 2011] to compute nearest neighbor graphs and `pytorch` [Paszke et al., 2019] for training neural networks. We used `FrEIA` [Ardizzone et al., 2018-2022] for constructing normalizing flows.

Two-dimensional datasets We use the following two-dimensional datasets. All are implemented as Gaussian mixture models. The first two are explicitly Gaussian, the final three are GMMs with 50 components fitted using `scikit-learn` to data generated by adding noise to certain geometric structures.

- Standard normal

- Gaussian mixture with 3 components
- Circle
- 1 spiral
- 2 spirals

Normalizing flow training We used the FrEIA library [Ardizzone et al., 2018-2022] to design normalizing flows using spline coupling blocks [Durkan et al., 2019]. We used the following hyperparameters:

- Blocks: 12 for two-dimensional data, 5 for standard normal in D dimensions
- Spline bins: 10
- Domain clamping in the splines: 5
- ActNorm between blocks
- Noise of 10^{-5} added to the training data
- LR = 5×10^{-4}
- Weight decay = 10^{-6}
- Fully-connected coupling block subnets with 3 hidden layers of 64 dimensions each, with ReLU activation and BatchNorm
- Training iterations (as a function of training set size n): $2500 \times 2^{n//1000}$
- Batch size = 256 if $n > 400$, else $\lfloor n/3 \rfloor$
- Adam optimizer

Score model training We use fully connected networks to predict the score, with 6 hidden layers with the width depending on the dimension of the input. We use the following hyperparameters:

- Hidden layer width (as a function of input dimension D)
 - $3 \leq D \leq 5$: 128
 - $6 \leq D \leq 8$: 170
 - $9 \leq D \leq 25$: 200
- LR = 10^{-3}
- Weight decay = 10^{-6}
- Softplus activation function
- Training iterations (as a function of training set size n): $2500 \times 2^{n//1000}$
- Batch size = 256 if $n > 400$, else $\lfloor n/3 \rfloor$
- Adam optimizer

In all cases, the chosen model checkpoint was that which achieved the lowest training loss.

Compute resources For the two-dimensional datasets, we performed all experiments on a machine with a NVIDIA GeForce RTX 2070 GPU and 32 GB of RAM. It took approximately 24 hours to run all experiments.

For the higher dimensional standard normal distributions, we performed all experiments on a machine with two GPUs: a NVIDIA GeForce RTX 2070 and a NVIDIA GeForce RTX 2080 Ti Rev. A. The machine has 32 GB of RAM. It took approximately 24 hours to run all experiments.

In addition we did some preliminary experiments on the first machine, totalling approximately 48 hours of compute time.




Figure 6: Increasing the number of nearest neighbors does not significantly improve the convergence of the power-weighted graph method on two-dimensional datasets.

C Additional experimental results

C.1 Convergence of alternative graph edge weightings

We try out other edge weights, as described in the main text. Given the nearest-neighbor density estimates $\hat{p}(X_l) \propto (\min_m \|X_l - X_m\|)^{-d}$ for data points X_l , we approximate the graph edges by $\mathcal{G}_{lm} = \|X_l - X_m\| / \tilde{p}((X_l + X_m)/2)$ where \tilde{p} is a density estimate for the midpoint. We use 4 variants:

1. **Inverse of mean:** $\tilde{p}((X_l + X_m)/2) = (\hat{p}(X_l) + \hat{p}(X_m))/2$
2. **Mean of inverse:** $1/\tilde{p}((X_l + X_m)/2) = (1/\hat{p}(X_l) + 1/\hat{p}(X_m))/2$
3. **Max:** $\tilde{p}((X_l + X_m)/2) = \max(\hat{p}(X_l), \hat{p}(X_m))$
4. **Min:** $\tilde{p}((X_l + X_m)/2) = \min(\hat{p}(X_l), \hat{p}(X_m))$

We find that all 4 have very similar performance to the power-weighted graph, whereas the equivalent estimators using the ground truth density quickly converge. See Fig. 7.




Figure 7: Alternative edge weightings based on the nearest neighbor density estimator do not converge. The same weightings based on the ground truth density do converge. Experiments performed on the GMM two-dimensional dataset.

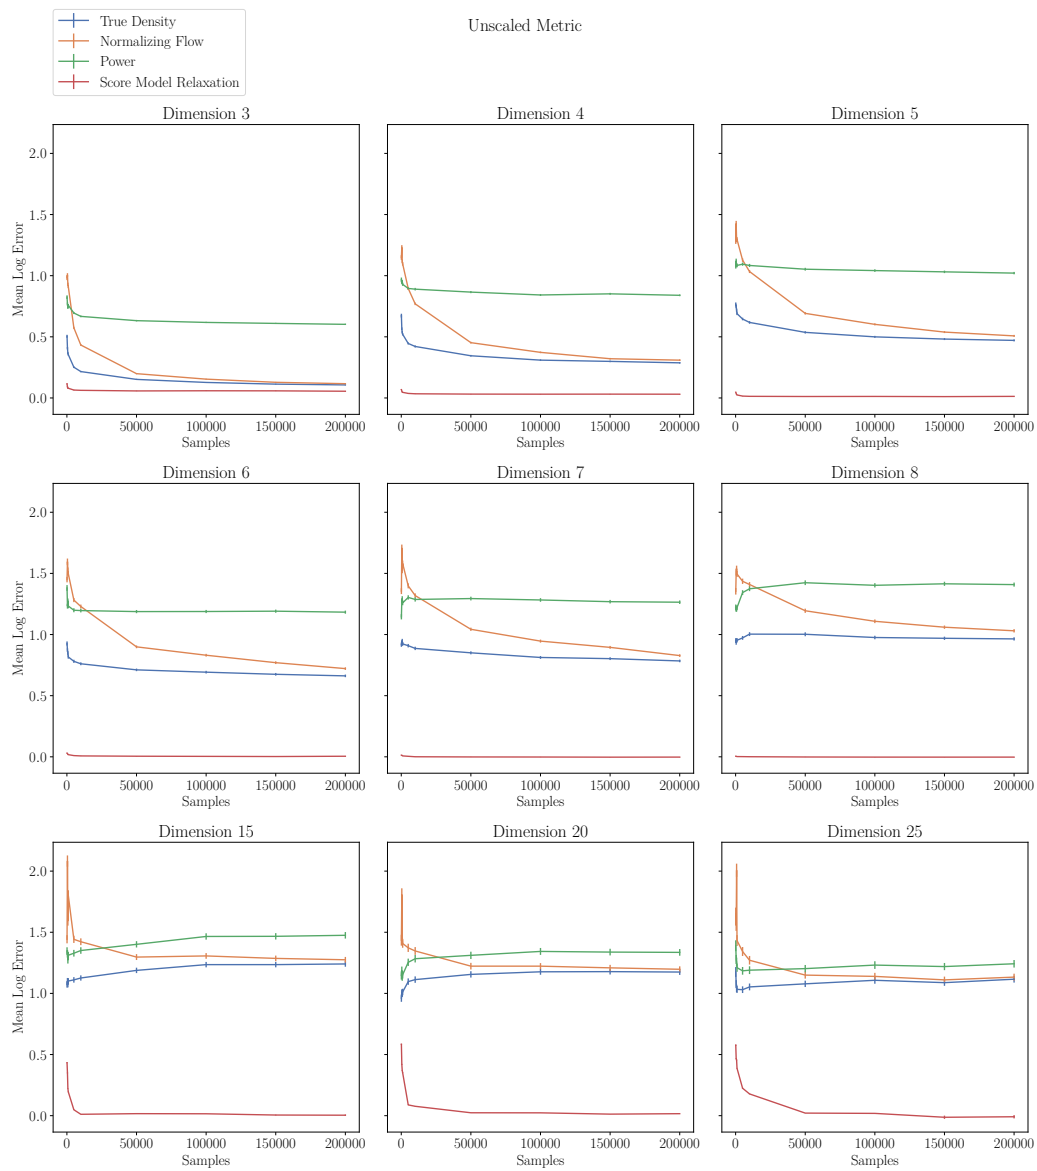


Figure 8: Detailed plot of convergence behavior in higher dimensions, unscaled metric. We see that all graph-based methods plateau in performance, whereas the relaxation-based method continues to show good convergence.

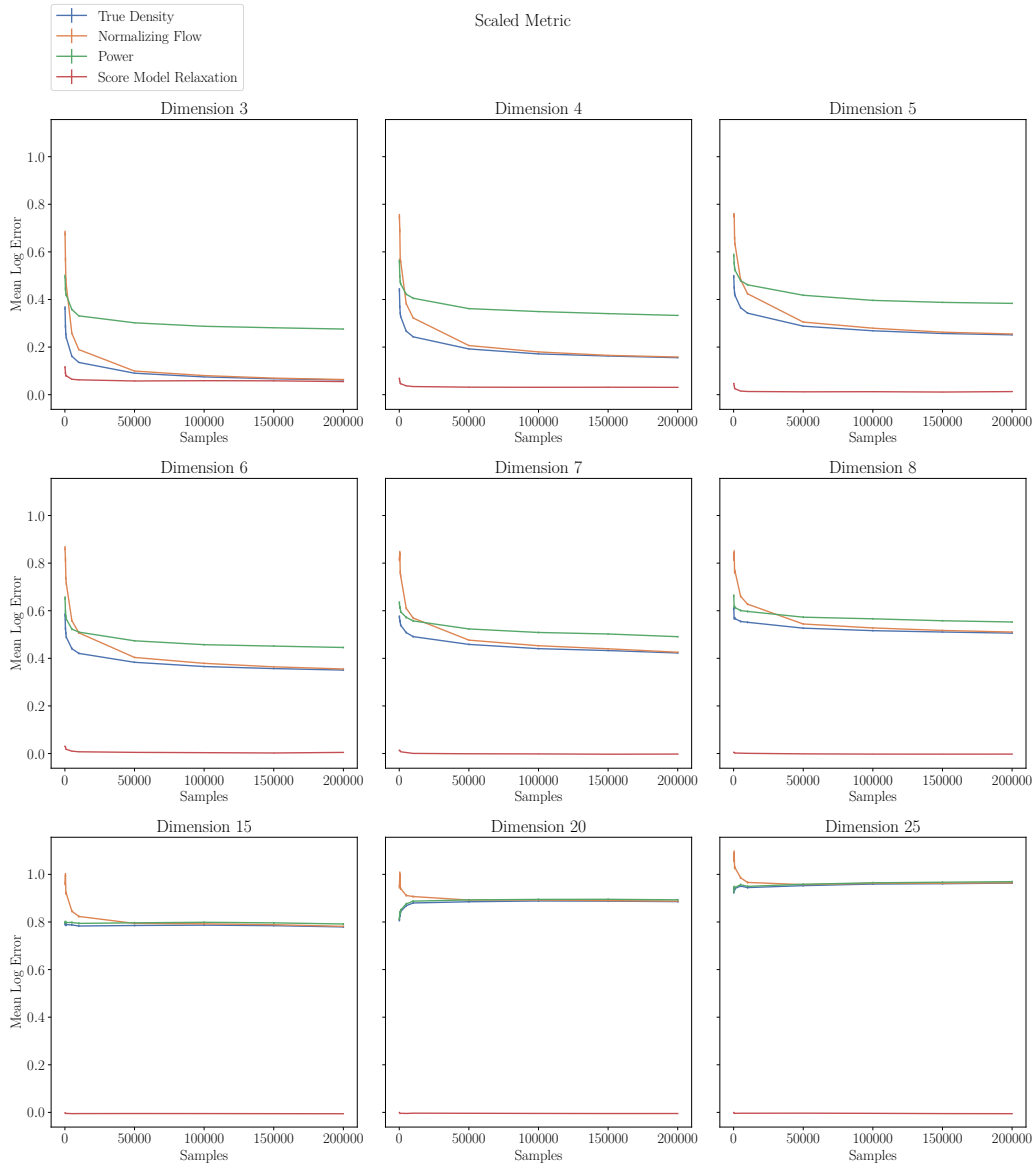


Figure 9: Detailed plot of convergence behavior in higher dimensions, scaled metric. We see that all graph-based methods plateau in performance, whereas the relaxation-based method continues to show good convergence.