

LDLE: LOW DISTORTION LOCAL EIGENMAPS

Dhruv Kohli & Alexander Cloninger

Department of Mathematics
University of California San Diego
CA 92093, USA
{dhkohli, acloninger}@ucsd.edu

Gal Mishne

Halicioğlu Data Science Institute
University of California San Diego
CA 92093, USA
{gmishne}@ucsd.edu

ABSTRACT

We present Low Distortion Local Eigenmaps (LDLE), a manifold learning technique which constructs a set of low distortion local views of a dataset in lower dimension and registers them to obtain a global embedding. The local views are constructed using the global eigenvectors of the graph Laplacian and are registered using Procrustes analysis. The choice of these eigenvectors may vary across the regions. In contrast to existing techniques, LDLE is more geometric and can embed manifolds without boundary as well as non-orientable manifolds into their intrinsic dimension.

1 INTRODUCTION

Nonlinear dimensionality reduction techniques such as Local Linear Embedding (Roweis & Saul, 2000), Hessian eigenmaps (Donoho & Grimes, 2003), Laplacian eigenmaps (Belkin & Niyogi, 2003), t-SNE (Maaten & Hinton, 2008) and UMAP (McInnes et al., 2018), aim at preserving local distances as they map a manifold embedded in higher dimension into lower (possibly intrinsic) dimension. In particular, UMAP and t-SNE follow a top-down approach where they start with an initial low-dimensional global embedding and then refine it by minimizing a local distortion measure on it. In contrast, similar to the one in (Singer & Wu, 2011), **a bottom-up embedding approach** can be imagined to consist of two steps, first obtaining low distortion local views of the manifold in lower dimension and then registering them to obtain a global embedding of the manifold. These local views can be constructed using the coordinate charts of the manifold. In this work, we take this local perspective to embed a manifold in low dimension.

2 BACKGROUND

Let (\mathcal{M}, g) be a d -dimensional Riemannian manifold with finite volume. By definition, for every x_k in \mathcal{M} , there exists a coordinate chart (\mathcal{U}_k, Φ_k) such that $x_k \in \mathcal{U}_k, \mathcal{U}_k \subset \mathcal{M}$ and Φ_k maps \mathcal{U}_k into \mathbb{R}^d . One can imagine \mathcal{U}_k to be a local view of \mathcal{M} in the ambient space. Using rigid transformations, these local views can be registered to recover \mathcal{M} . Similarly, $\Phi_k(\mathcal{U}_k)$ can be imagined to be a local view of \mathcal{M} in the embedding space \mathbb{R}^d . Again using rigid transformations, these local views can be registered to obtain a d -dimensional embedding of \mathcal{M} . As there may exist multiple mappings which map \mathcal{U}_k into \mathbb{R}^d , a natural strategy would be to choose a mapping with low distortion. Let $d_g(x, y)$ denote the shortest geodesic distance between $x, y \in \mathcal{M}$. Then the distortion of Φ_k on \mathcal{U}_k as defined in (Jones et al., 2007) is given by

$$\text{Distortion}(\Phi_k, \mathcal{U}_k) = \|\Phi_k\|_{Lip} \|\Phi_k^{-1}\|_{Lip} \quad (1)$$

where $\|\Phi_k\|_{Lip}$ (and similarly $\|\Phi_k^{-1}\|_{Lip}$) is the Lipschitz norm of Φ_k given by

$$\|\Phi_k\|_{Lip} = \sup_{\substack{x, y \in \mathcal{U}_k \\ x \neq y}} \frac{\|\Phi_k(x) - \Phi_k(y)\|_2}{d_g(x, y)}. \quad (2)$$

Jones et al. (2007) provide **guarantees on the distortion of the coordinate charts** of the manifold constructed using carefully chosen eigenfunctions of the Laplace-Beltrami operator Δ_g on it with

Dirichlet or Neumann boundary conditions. They proved that for a small enough ball \mathcal{U}_k around a given point x_k on \mathcal{M} , there exist i_1, i_2, \dots, i_d where $i_s \equiv i_s(k)$, such that, if we let ϕ_i to be an eigenfunction of Δ_g then the map $\Phi_k(x) = (\gamma_{ki_1}\phi_{i_1}(x), \dots, \gamma_{ki_d}\phi_{i_d}(x))$, where $x \in \mathcal{U}_k$ and $\gamma_{ki} = (\int_{\mathcal{U}_k} \phi_i^2(y)dy/|\mathcal{U}_k|)^{-1/2}$, has a **distortion bounded from above by κ^2** , a constant that depends on the natural geometric properties of \mathcal{M} . The main idea behind the choice of these eigenfunctions is that they need to constitute a locally orthonormal frame upto scale; the gradients $\nabla\phi_{i_s}(x_k)$ are close to being mutually orthogonal and the local scaling factors $\gamma_{ki_s} \|\nabla\phi_{i_s}(x_k)\|_2$ are close to each other.

3 OUR CONTRIBUTION

In this work we present Low Distortion Local Eigenmaps (LDLE), a new manifold learning approach. The main contributions of our work are as follows:

1. We present an algorithmic realization of the construction procedure in (Jones et al., 2007) that applies to the discretized setting. Unlike existing techniques such as UMAP, Laplacian Eigenmaps etc. which construct a global embedding by relying on the globally orthonormal eigenvectors, here, we choose **locally orthonormal eigenvectors** (up to scale) of the graph Laplacian to construct low distortion local embeddings and later piece them together to obtain a **more geometric** global embedding compared to other techniques (see Appendix B).
2. We present an algorithm to obtain a global embedding of the manifold by registering its local views. Unlike existing techniques, the algorithm is designed to embed **manifolds without boundary as well as non-orientable manifolds** into their intrinsic dimension by **tearing** them apart. It also provides **gluing instructions** for the boundary of the embedding by coloring it such that the points on the boundary which are adjacent on the manifold have the same color (see Figure 3).

4 LDLE ALGORITHM

In summary, LDLE consists of three steps. In the first step, we present an algorithmic realization of the result in (Jones et al., 2007) to construct low-dimensional low distortion parameterizations $(\Phi_k)_{k=1}^n$ of the small discrete balls $(U_k)_{k=1}^n$ around the points $(x_k)_{k=1}^n$ on the discretized manifold. This involves a **careful selection of the eigenvectors** $\phi_{i_1}, \dots, \phi_{i_d}$, $i_s \equiv i_s(k)$, of the graph Laplacian to obtain a **low distortion map** Φ_k of U_k into \mathbb{R}^d where

$$\begin{aligned} \Phi_k &\equiv (\gamma_{ki_1}\phi_{i_1}, \dots, \gamma_{ki_d}\phi_{i_d}) \text{ where} \\ \Phi_k(x_{k'}) &= (\gamma_{ki_1}\phi_{i_1k'}, \dots, \gamma_{ki_d}\phi_{i_dk'}) \text{ and} \\ \Phi_k(U_k) &= (\Phi_k(x_{k'}))_{x_{k'} \in U_k} \text{ (See Figure 1)}. \end{aligned} \tag{3}$$

We call U_k and $\Phi_k(U_k)$ the k th local view of the data in the ambient space and the d -dimensional embedding space, respectively. Also, using Eq. 1 and the fact that Euclidean distance d_e approximates the local geodesic distance, we obtain the distortion of Φ_k on U_k as

$$\zeta_k = \sup_{\substack{x_l, x_{l'} \in U_k \\ x_l \neq x_{l'}}} \frac{\|\Phi_k(x_l) - \Phi_k(x_{l'})\|}{d_e(x_l, x_{l'})} \sup_{\substack{x_l, x_{l'} \in U_k \\ x_l \neq x_{l'}}} \frac{d_e(x_l, x_{l'})}{\|\Phi_k(x_l) - \Phi_k(x_{l'})\|}. \tag{4}$$

In the second step, we develop a clustering algorithm to obtain a small number of intermediate views $\tilde{\Phi}_m(\tilde{U}_m)$ with low distortion, from the large number of smaller local views $\Phi_k(U_k)$. In theory, one can use the Generalized Procrustes Analysis (GPA) (Crosilla & Beinat, 2002; Gower, 1975; Ten Berge, 1977) to directly register the local views themselves and recover the global embedding. In practice, too many small local views (high n and small $|U_k|$) result in extremely high computational complexity. Moreover, small overlaps between the local views makes their registration susceptible to errors.

Therefore, we perform clustering to obtain $M \ll n$ intermediate views, \tilde{U}_m and $\tilde{\Phi}_m(\tilde{U}_m)$, of the data in the ambient space and the embedding space, respectively. The algorithm **transforms the notion of a local view per an individual point to an intermediate view per a cluster of points**.

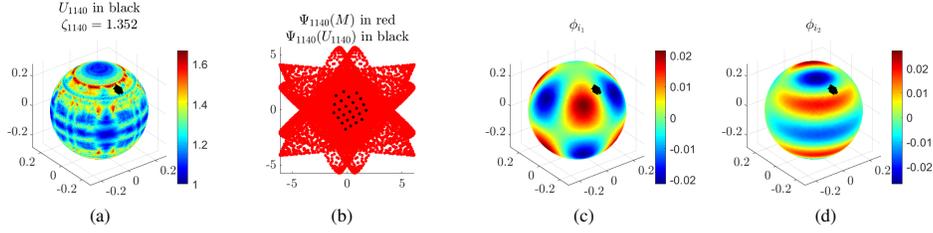


Figure 1: (a) A discretized sphere colored by the distortion ζ_k of Φ_k on U_k as k varies. U_{1140} is a local view in the ambient space. (b) 2d embedding of $(x_k)_{k=1}^n$ and U_{1140} due to Ψ_{1140} . Here $\Psi_{1140}(U_{1140})$ is the corresponding local view in the embedding space. (c, d) Chosen eigenvectors to construct Ψ_{1140} . Notice that the chosen eigenvectors are close to being locally orthogonal.

It is designed so as to ensure low distortion of the parameterizations $\tilde{\Phi}_m$ on \tilde{U}_m . Overall, the **clustering reduces the number of views and increases the overlaps between them, leading to their quick and robust registration**. See Figure 2.

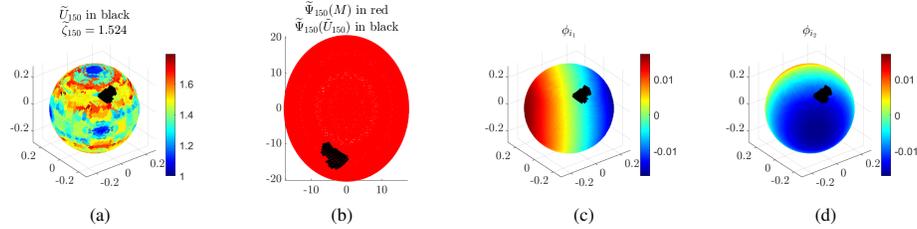


Figure 2: (a) Same discretized sphere colored by the distortion ζ_m of $\tilde{\Phi}_m$ on \tilde{U}_m as m varies (see Eq. (4)). \tilde{U}_{150} is an intermediate view in the ambient space which contains the local view U_{1140} as in Figure 1. (b) 2d embedding of $(x_k)_{k=1}^n$ and \tilde{U}_{150} due to $\tilde{\Psi}_{150}$. Here $\tilde{\Psi}_{150}(U_{150})$ is the corresponding intermediate view in the embedding space. (c, d) Chosen eigenvectors to construct $\tilde{\Psi}_k$ where $\tilde{\Phi}_{150} = \tilde{\Phi}_k$.

In the final step, we present an algorithm based on Global Procrustes analysis (GPA) (Crosilla & Beinat, 2002; Gower, 1975; Ten Berge, 1977) to register the intermediate views and obtain a global embedding. Each view $\tilde{\Phi}_m(\tilde{U}_m)$ is transformed by an orthogonal matrix T_m of size $d \times d$, a d -dimensional translation vector v_m and a positive scalar b_m as a scaling component. The transformed view is given by $\tilde{\Phi}_m^g(\tilde{U}_m)$ such that for all $x_k \in \tilde{U}_m$,

$$\tilde{\Phi}_m^g(x_k) = b_m \tilde{\Phi}_m(x_k) T_m + v_m. \quad (5)$$

In general, the parameters $(T_m, v_m, b_m)_{m=1}^M$ are estimated so that for all m and m' , the two transformed views of the overlap between \tilde{U}_m and $\tilde{U}_{m'}$, obtained using the parameterizations $\tilde{\Phi}_m^g$ and $\tilde{\Phi}_{m'}^g$, align with each other. To be more precise, define the intersection of \tilde{U}_m and $\tilde{U}_{m'}$ as $\tilde{U}_{mm'}$. Then, the parameters are estimated so as to minimize the following alignment error

$$\mathcal{L}((T_m, v_m, b_m)_{m=1}^M) = \sum_{m, m'} \left\| \tilde{\Phi}_m^g(\tilde{U}_{mm'}) - \tilde{\Phi}_{m'}^g(\tilde{U}_{mm'}) \right\|_F^2. \quad (6)$$

In theory, one can start with T_m, v_m and b_m as $I_d, \mathbf{0}$ and 1, and directly use GPA to obtain a local minimum of the above alignment error. However, **naive GPA does not allow tearing apart the manifolds without a boundary**. Therefore, we develop an algorithm by **adapting GPA** so as to tear apart the manifolds without boundary as well as non-orientable manifolds and embed them into their intrinsic dimension. Our algorithm also provides gluing instructions for the boundary of the embedding by coloring it such that the points on the boundary which are adjacent on the manifold have the same color. See Figure 3.

5 EXPERIMENTS

We present experiments to compare LDLE with UMAP (McInnes et al., 2018), t-SNE (Maaten & Hinton, 2008) and Laplacian eigenmaps (Belkin & Niyogi, 2003) on several datasets. These datasets are discretized 2d manifolds embedded in $\mathbb{R}^2, \mathbb{R}^3$ or \mathbb{R}^4 , containing about 10^4 points. These manifolds can be grouped based on the presence of the boundary and their orientability, as shown in Figures 3 and 4 and Figures 5, 8 and 9 provided in the appendix.

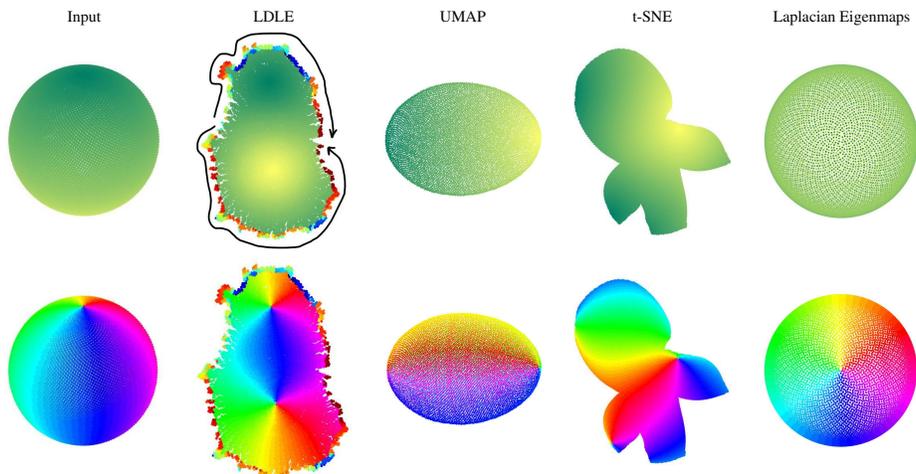


Figure 3: Embeddings of the sphere in \mathbb{R}^3 into \mathbb{R}^2 . The top and bottom row contain the same plots colored by the height and the azimuthal angle of the sphere ($0 - 2\pi$), respectively. LDLE automatically colors the boundary so that the points on the boundary which are adjacent on the sphere have the same color. The arrows are manually drawn to help the reader identify the two pieces of the boundary which are to be stitched together to recover the original sphere. UMAP and Laplacian eigenmaps squeezed the sphere into two different viewpoints of \mathbb{R}^2 (side or top view of the sphere). t-SNE also tore apart the sphere but the embedding lacks interpretability as it is “unaware” of the boundary.

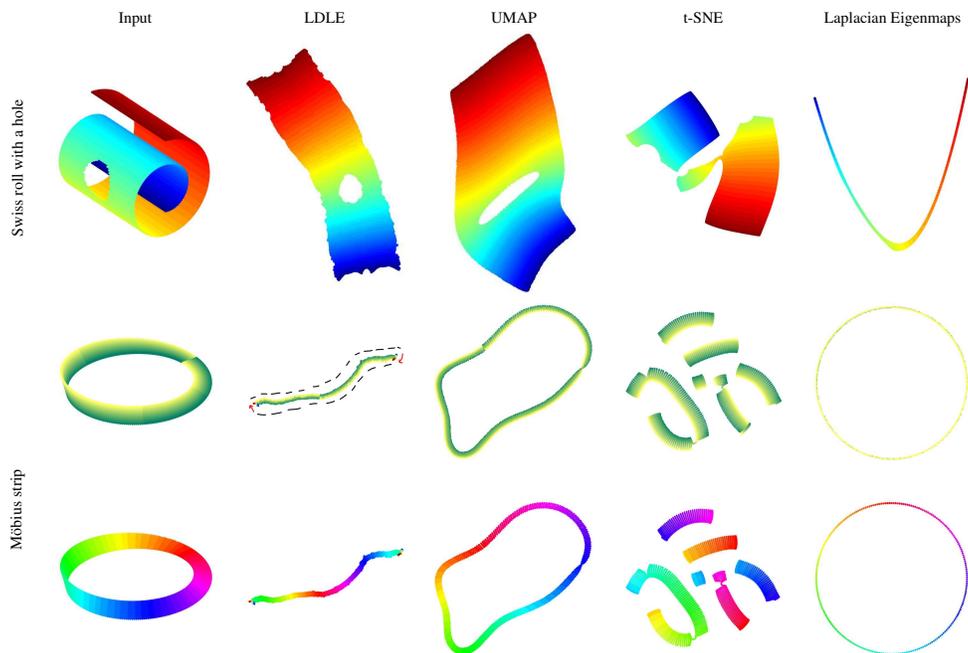


Figure 4: Embedding of swissroll with a hole and a Möbius strip in \mathbb{R}^3 into \mathbb{R}^2 . The last two rows contain the same plots colored by the two parameters used to construct the Möbius strip. The red arrows are manually drawn to help the reader identify the two pieces of the boundary which are to be stitched together to recover the original strip. The pieces of the boundary along the dashed lines are not to be stitched together.

6 CONCLUSION

In summary, LDLE competes with the other methods in terms of visualization quality. In particular, the embeddings produced by LDLE are geometrically more accurate than those produced by UMAP, t-SNE and Laplacian Eigenmaps (see Appendix B). We also demonstrated that LDLE can embed manifolds without boundary as well as non-orientable manifolds into their intrinsic dimension, a feature that is missing from the existing techniques.

7 ACKNOWLEDGEMENT

Dhruv Kohli and Gal Mishne were supported by the National Institutes of Health (R01EB026936).

REFERENCES

- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Fabio Crosilla and Alberto Beinat. Use of generalised Procrustes analysis for the photogrammetric block adjustment by independent models. *ISPRS Journal of Photogrammetry and Remote Sensing*, 56:195–209, 04 2002. doi: 10.1016/S0924-2716(02)00043-6.
- David L Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.
- John C Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
- Peter W Jones, Mauro Maggioni, and Raanan Schul. Universal local parametrizations via heat kernels and eigenfunctions of the Laplacian. *arXiv preprint arXiv:0709.1975*, 2007.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- Amit Singer and Hau-tieng Wu. Orientability and diffusion maps. *Applied and computational harmonic analysis*, 31(1):44–58, 2011.
- Jos MF Ten Berge. Orthogonal Procrustes rotation for two or more matrices. *Psychometrika*, 42(2): 267–276, 1977.

A MANIFOLDS WITH BOUNDARY

Figure 5 compares the visualization quality of LDLE with other techniques. It shows the 2d embeddings of 2d manifolds with boundary in \mathbb{R}^2 or \mathbb{R}^3 . For the square with two holes, unlike UMAP and Laplacian Eigenmaps, LDLE nearly preserved the shape of the holes. For the sphere with a hole which is a curved 2d manifold with boundary in \mathbb{R}^3 , both UMAP and Laplacian eigenmaps flattened it into \mathbb{R}^2 while LDLE and t-SNE tore it apart. Note that the boundaries of the LDLE embeddings are usually distorted. It turns out that when the points in the input which lie on the boundary are known apriori then the distortion near the boundary can be reduced by using the doubled manifold. The resulting embeddings are shown in Figure 6.

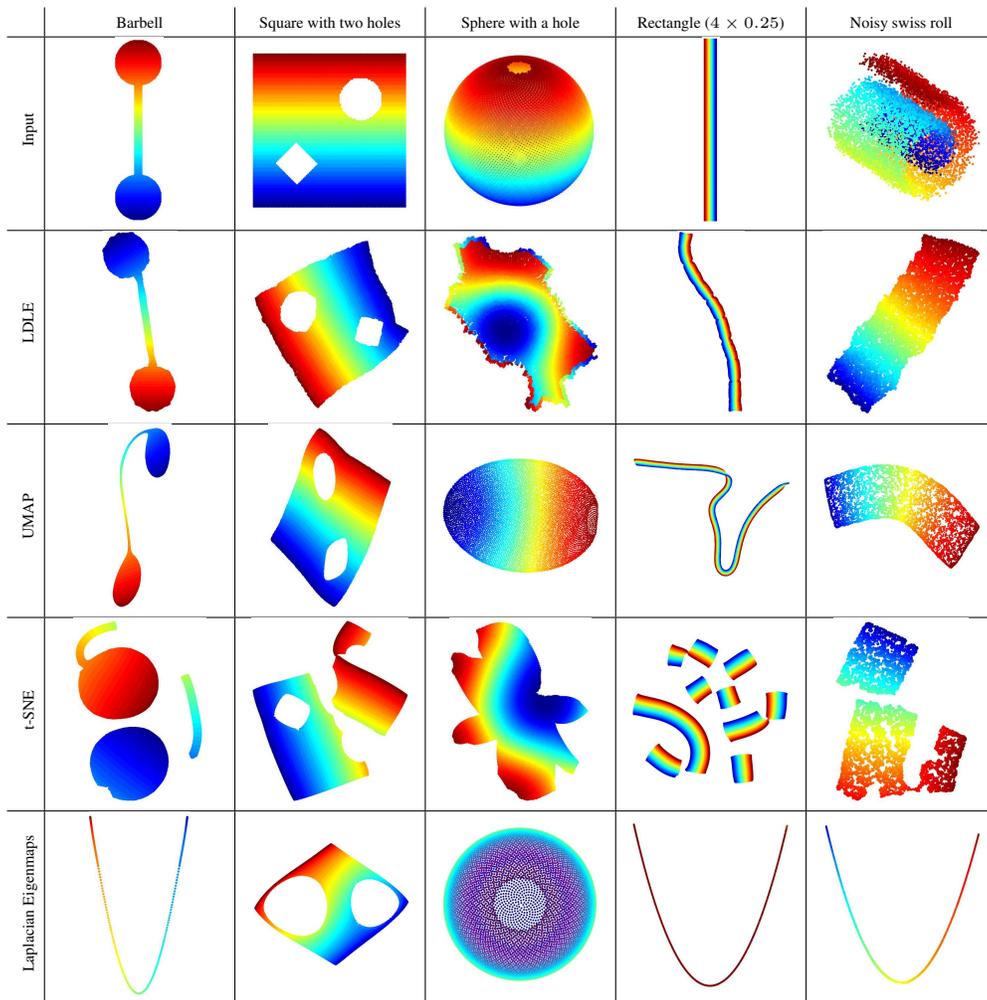


Figure 5: Embeddings of 2d manifolds with boundary into \mathbb{R}^2 . The noisy swiss roll is constructed by adding uniform noise in all three dimensions, with support on $[0, 0.05]$.

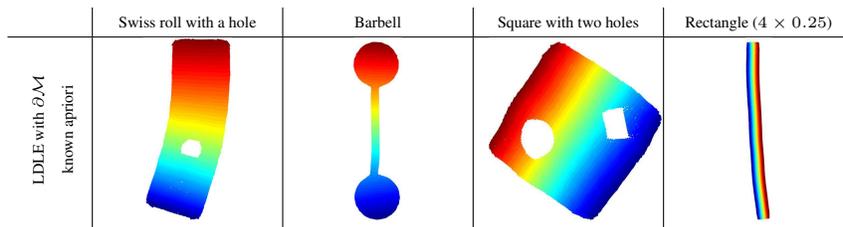


Figure 6: LDLE embeddings when the points on the boundary are known apriori.

B QUANTITATIVE COMPARISON

To compare LDLE with other techniques in a quantitative manner, we compute the distortion \mathcal{D}_k of the embeddings of the geodesics originating from x_k and then plot the distribution of \mathcal{D}_k (see Figure 7). The procedure to compute \mathcal{D}_k follows. In the discrete setting, we first define the geodesic between two given points as the shortest path between them which in turn is computed by running Dijkstra algorithm on the graph of 5 nearest neighbors. Here, the distances are measured using the Euclidean metric d_e . Let us denote the number of nodes on the geodesic between x_k and $x_{k'}$ by $n_{kk'}$ and the sequence of nodes by $(x_s)_{s=1}^{n_{kk'}}$ where $x_1 = x_k$ and $x_{n_{kk'}} = x_{k'}$. Then the length of the geodesic between x_k and $x_{k'}$ is given by

$$L_{kk'} = \sum_{s=2}^{n_{kk'}} d_e(x_s, x_{s-1}). \quad (7)$$

Denote the embedding of x_k by y_k . Then the length of the embedding of the geodesic between x_k and $x_{k'}$ is given by

$$L_{kk'}^g = \sum_{s=2}^{n_{kk'}} d_e(y_s, y_{s-1}). \quad (8)$$

Finally, the distortion \mathcal{D}_k of the embeddings of the geodesics originating from x_k is given by the ratio of maximum expansion and minimum contraction, that is,

$$\mathcal{D}_k = \sup_{k'} \frac{L_{kk'}^g}{L_{kk'}} / \inf_{k'} \frac{L_{kk'}^g}{L_{kk'}} = \sup_{k'} \frac{L_{kk'}^g}{L_{kk'}} \sup_{k'} \frac{L_{kk'}}{L_{kk'}^g}. \quad (9)$$

A value of 1 for \mathcal{D}_k means the geodesics originating from x_k have same length in the input and in the embedding space. If $\mathcal{D}_k = 1$ for all k then the embedding is geometrically (and therefore topologically too) the same as the input. Figure 7 shows the distribution of \mathcal{D}_k due to LDLE and other algorithms for various examples. Clearly, LDLE least distorts the geodesics in the input space, especially when the points on the boundary are known a priori. In this sense, LDLE is more geometric than other techniques.

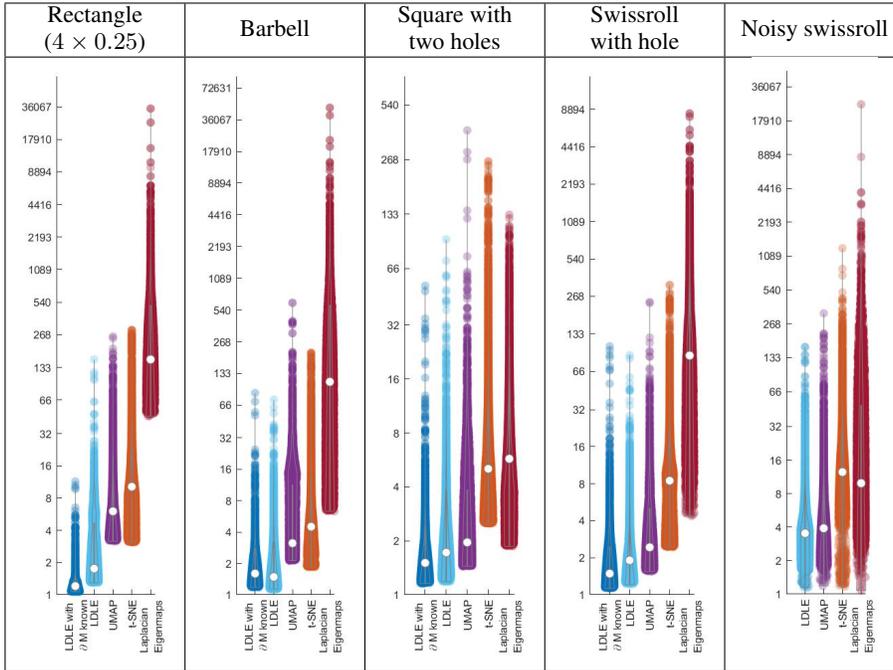


Figure 7: Violin plots for the distribution of \mathcal{D}_k (See Eq. (9)). The white point inside the violin represents the median.

C MANIFOLDS WITHOUT BOUNDARY

In Figure 8, we show the 2d embeddings of 2d manifolds without boundary, a curved torus in \mathbb{R}^3 and a flat torus in \mathbb{R}^4 . A flat torus is a parallelogram whose opposite sides are identified. In our case, we construct a discrete flat torus using a rectangle with sides 2 and 0.5 and embed it in four dimensions as follows,

$$X(\theta_i, \phi_j) = \frac{1}{4\pi}(4\cos(\theta_i), 4\sin(\theta_i), \cos(\phi_j), \sin(\phi_j)) \tag{10}$$

where $\theta_i = 0.01i\pi$, $\phi_j = 0.04j\pi$, $i \in \{0, \dots, 199\}$ and $j \in \{0, \dots, 49\}$.

LDLE produced similar representation for both the inputs. None of the other methods do that. The main difference in the LDLE embedding of the two inputs is based on the length of the same colored pieces of the boundary. Note that these pieces are adjacent in the input space. For the flat torus, the two pieces with same colors are almost equal in length while for the curved torus, they usually have different lengths. This is because of the difference in the curvature of the two inputs, zero everywhere for the flat torus and non-zero almost everywhere on the curved torus. The mathematical correctness of the LDLE embeddings using the cut and paste argument is shown in Figure 10.

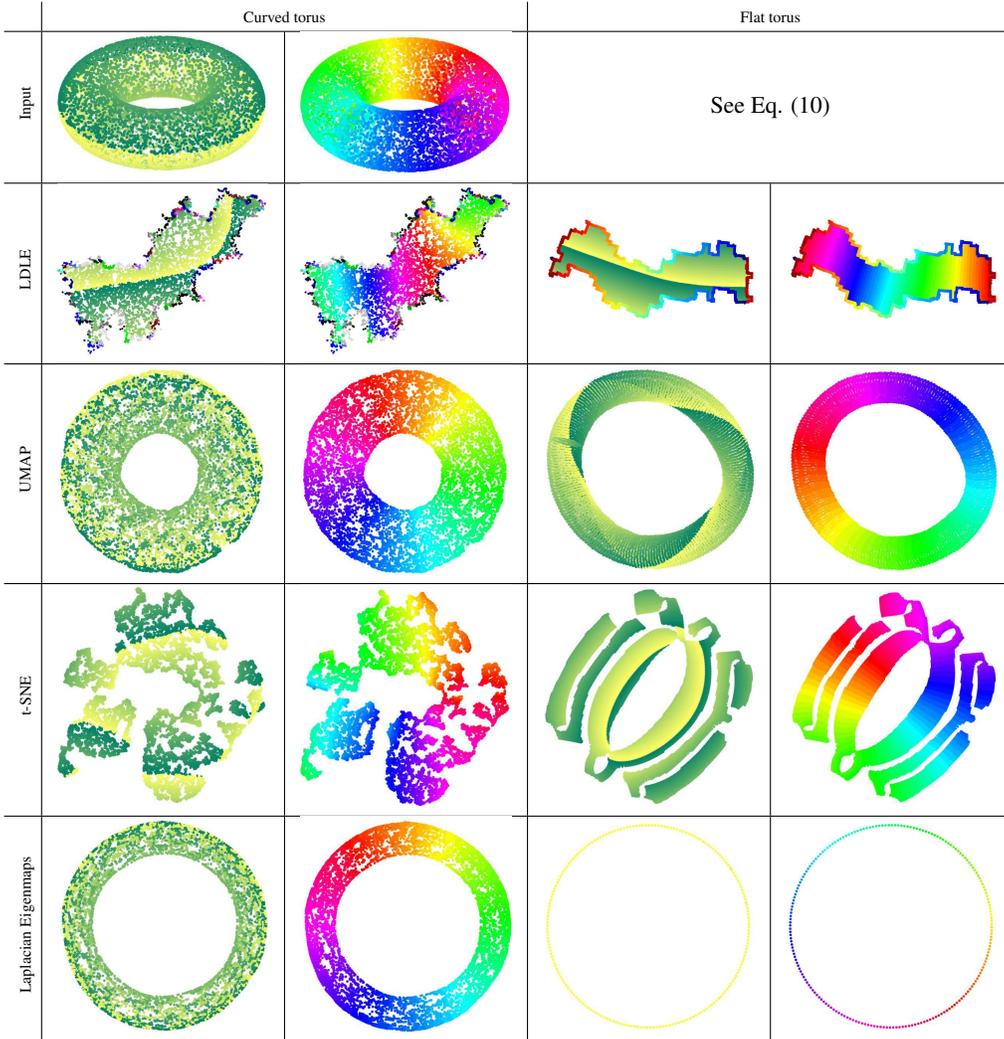


Figure 8: Embeddings of 2d manifolds without boundary into \mathbb{R}^2 . For each manifold, the left and right columns contain the same plots colored by the two parameters of the manifold. A proof of the mathematical correctness of the LDLE embeddings is provided in Figure 10.

D NON-ORIENTABLE MANIFOLDS

In Figure 9, we show the 2d embedding of a Klein bottle which is a non-orientable 2d manifold without boundary. We construct a discrete Klein bottle using its 4d Möbius tube representation as follows,

$$\begin{aligned} X(\theta_i, \phi_j) &= (R(\phi_j) \cos \theta_i, R(\phi_j) \sin \theta_i, r \sin \phi_j \cos \frac{\theta_i}{2}, r \sin \phi_j \sin \frac{\theta_i}{2}) \\ R(\phi_j) &= R + r \cos \phi_j \end{aligned} \quad (11)$$

where $\theta_i = i\pi/100$, $\phi_j = j\pi/25$, $i \in \{0, \dots, 199\}$ and $j \in \{0, \dots, 49\}$.

Laplacian eigenmaps produced incorrect embeddings, t-SNE produced dissected and non-interpretatable embeddings and UMAP squeezed the input into \mathbb{R}^2 . On the other hand, LDLE produced mathematically correct embedding by tearing apart the Klein bottle to embed it into \mathbb{R}^2 (see Figure 10).

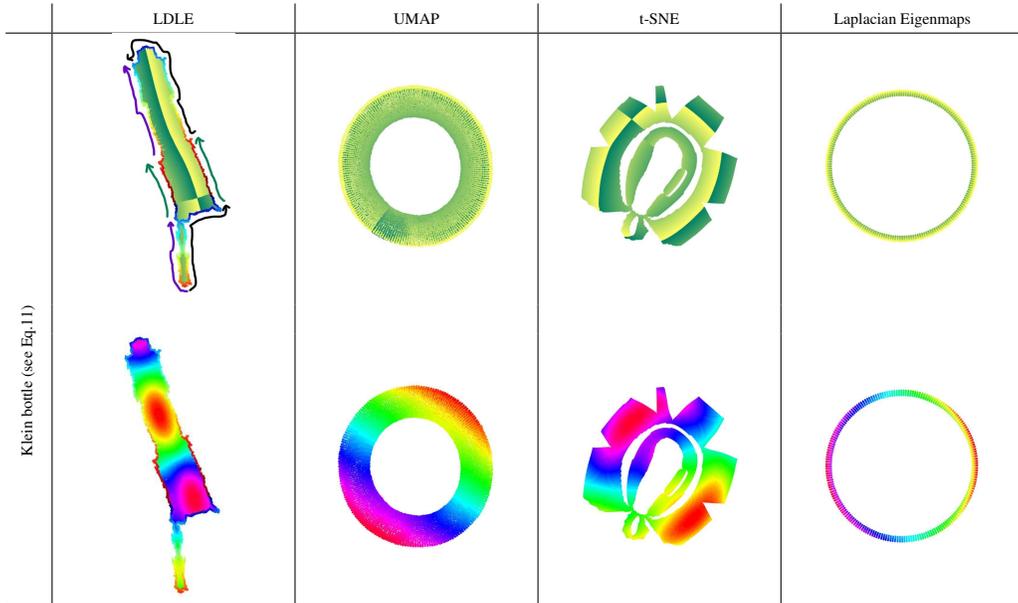


Figure 9: Embedding of a 2d non-orientable Klein bottle in \mathbb{R}^4 into \mathbb{R}^2 . The top and bottom row contain the same plots colored by the two parameters of the manifold. A proof of the mathematical correctness of the LDLE embeddings is provided in Figure 10.

E MATHEMATICAL CORRECTNESS OF LDLE EMBEDDINGS

	LDLE with arrows	Derived cut and paste diagrams
Sphere		
Sphere with a hole		
Curved torus		
Flat torus		
Möbius strip		
Klein bottle		

Figure 10: (Left) LDLE embedding with arrows drawn by tracing the colored boundary. (Right) Derived cut and paste diagrams to prove the correctness of the embedding. Filled arrows represent pieces of the boundary which are to be stitched along the arrows of the same color. Pieces of the boundary represented by dashed lines are not to be stitched. Dotted lines and shallow arrows represent cut and paste instructions.