LDLE: LOW DISTORTION LOCAL EIGENMAPS

Anonymous authors
Paper under double-blind review

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.

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 the local perspective to embed a manifold in low dimension.

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 $(\mathcal{U}_k, \Phi_k)$ such that $x_k \in \mathcal{U}_k, \mathcal{U}_k \subset M$ 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

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

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

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

Jones et al. (2007) provide guarantees on the distortion of the coordinate charts of the manifold. These charts are 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 $i_1, i_2, \ldots, i_d$ where $i_s \equiv i_s(k)$, such that, if we let ϕ_i to be an eigenfunction of Δ_g and $\gamma_{ki} = (\int_{\mathcal{U}_k} \phi_i^2(y) dy / |\mathcal{U}_k|)^{-1/2}$ then the distortion of the map

$$\begin{aligned} \Phi_k : \mathcal{U}_k &\to \mathbb{R}^d \\ x &\to (\gamma_{ki_1}\phi_{i_1}(x), \ldots, \gamma_{ki_d}\phi_{i_d}(x)) \end{aligned} \tag{3}$$

is bounded from above by κ^2 where $\kappa > 1$ 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 orthogonal frame; the gradients $\nabla\phi_{i_s}(x_k)$ are close to being mututally orthogonal and the local

scaling factors $\gamma_{ki_s} \|\nabla \phi_{i_s}(x_k)\|_2$ are close to each other. In this work, we present an approach to realize their work in the discrete setting and obtain low-dimensional low distortion local views of the given dataset. Moreover, we piece together these local views to obtain a global embedding of the manifold.

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 $(U_k)_{k=1}^n$ on the discretized manifold represented by the points $(x_k)_{k=1}^n$. This involves a careful selection of the eigenvectors $\boldsymbol{\phi}_{i_1}, \ldots, \boldsymbol{\phi}_{i_d}$, $i_s \equiv i_s(k)$, of the graph Laplacian to obtain a low distortion map Φ_k of U_k into $\mathbb{R}^d$ where

$$\begin{aligned} \Phi_k &\equiv (\gamma_{ki_1}\boldsymbol{\phi}_{i_1}, \ldots, \gamma_{ki_d}\boldsymbol{\phi}_{i_d}) \text{ where} \\ \Phi_k(x_{k'}) &= (\gamma_{ki_1}\boldsymbol{\phi}_{i_1 k'}, \ldots, \gamma_{ki_d}\boldsymbol{\phi}_{i_d k'}) \text{ and} \\ \Phi_k(U_k) &= (\Phi_k(x_{k'}))_{x_{k'} \in U_k} \text{ (See Figure 1).} \end{aligned} \tag{4}$$

We call U_k and $\Phi_k(U_k)$ the kth local view of the data in the ambient space and the d-dimensonal 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 $\Phi_{k'}$ on U_k as

$$\zeta_{kk'} = \sup_{\substack{x_l, x_{l'} \in U_k \\ x_l \neq x_{l'}}} \frac{\|\Phi_{k'}(x_l) - \Phi_{k'}(x_{l'})\|}{d_e(x_l, x_{l'})} \sup_{\substack{x_l, x_{l'} \in U_k \\ x_l \neq x_{l'}}} \frac{d_e(x_l, x_{l'})}{\|\Phi_{k'}(x_l) - \Phi_{k'}(x_{l'})\|}. \tag{5}$$

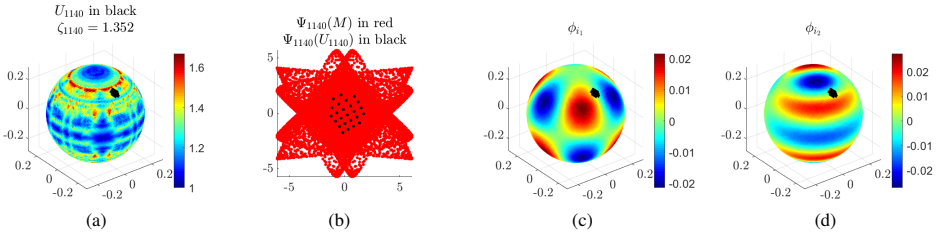

Figure 1: (a) A discretized sphere colored by the distortion ζ_{kk} of Φ_k on U_k as k varies. U_{1140} is a local view in the ambient space. (b) 2d embedding of $(x_k)_{k=1}^n$ and U_{1140} due to Ψ_{1140}. Here $\Psi_{1140}(U_{1140})$ is the corresponding local view in the embedding space. (c, d) Chosen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

In the second step, we develop a clustering algorithm to obtain a small number of intermediate views $\widetilde{\Phi}_m(\widetilde{U}_m)$ 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 register the local views to 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.

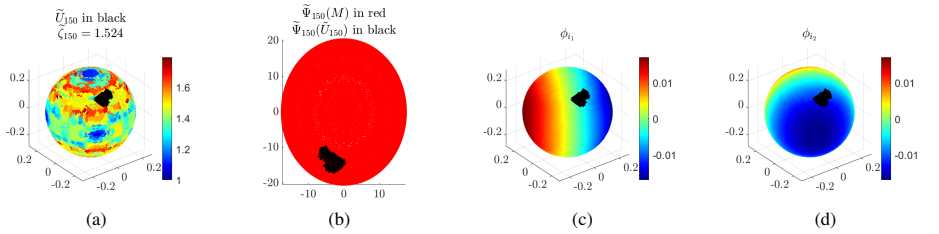

Figure 2: (a) Same discretized sphere colored by the distortion $\widetilde{\zeta}_{mm}$ of $\widetilde{\Phi}_m$ on $\widetilde{U}_m$ as m varies. $\widetilde{U}_{150}$ is an intermediate view in the ambient space which contains the local view U_{1140} as in Figure 1. (b) 2d embedding of $(x_k)_{k=1}^n$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

Therefore, we perform clustering to obtain $M \ll n$ intermediate views, $\widetilde{U}_m$ and $\widetilde{\Phi}_m(\widetilde{U}_m)$, 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. It is designed so as to ensure low distortion of the parameterizations $\widetilde{\Phi}_m$ on $\widetilde{U}_m$. Overall, the clustering

reduces the number of views and increases the overlaps between them, leading to their quick and robust registration. See Figure 2.

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 $\widetilde{\Phi}_m(\widetilde{U}_m)$ 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 $\widetilde{\Phi}^g_m(\widetilde{U}_m)$ such that for all $x_k \in \widetilde{U}_m$,

$$\widetilde{\Phi}^g_m(x_k) = b_m \widetilde{\Phi}_m(x_k) T_m + v_m. \tag{6}$$

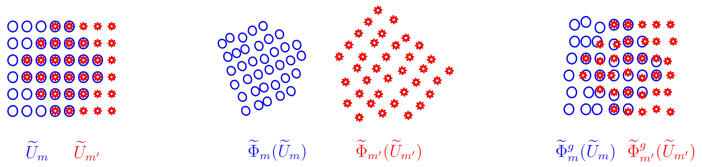

Figure 3: (left) The intermediate views $\widetilde{U}_m$ and $\widetilde{U}_{m'}$ of a 2d manifold in a possibly high dimensional ambient space. These views trivially align with each other. The red star in blue circles represent their overlap $\widetilde{U}_{mm'}$. (middle) The mth and m'th intermediate views in the 2d embedding space. (right) The views after aligning the $\widetilde{\Phi}_m(\widetilde{U}_{mm'})$ with $\widetilde{\Phi}_{m'}(\widetilde{U}_{mm'})$.

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 $\widetilde{U}_m$ and $\widetilde{U}_{m'}$, obtained using the parameterizations $\widetilde{\Phi}^g_m$ and $\widetilde{\Phi}^g_{m'}$, align with each other. To be more precise, define the overlap between $\widetilde{U}_m$ and $\widetilde{U}_{m'}$ as $\widetilde{U}_{mm'}$. It is natural to ask for a low distortion global embedding of the data. Therefore, we must ensure that the embeddings of $\widetilde{U}_{mm'}$ due to the mth and the m'th view in the embedding space, also align with each other. Thus, the parameters $(T_m, v_m, b_m)_{m=1}^M$ are estimated so that $\widetilde{\Phi}^g_m(\widetilde{U}_{mm'})$ aligns with $\widetilde{\Phi}^g_{m'}(\widetilde{U}_{mm'})$ for all m and m' (see Figure 3). However, due to the distortion of the parameterizations it is usually not possible to perfectly align the two embeddings. We can represent both embeddings of the overlap as matrices with $|\widetilde{U}_{mm'}|$ rows and d columns. We then measure the error in the the alignment as the squared Frobenius norm of the difference of the two matrices. The error is trivially zero if $\widetilde{U}_{mm'}$ is empty. Overall, the estimated parameters minimize the following alignment error

$$\mathcal{L}((T_m, v_m, b_m)_{m=1}^M) = \sum_{m,m'} \left\| \widetilde{\Phi}^g_m(\widetilde{U}_{mm'}) - \widetilde{\Phi}^g_{m'}(\widetilde{U}_{mm'}) \right\|_F^2. \tag{7}$$

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, using naive GPA to align a view with all of its adjacent views would prevent us from 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 4.

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 4 and 5 and Figures 6, 9 and 10 provided in the supplementary material.

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 Figure 8). 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.

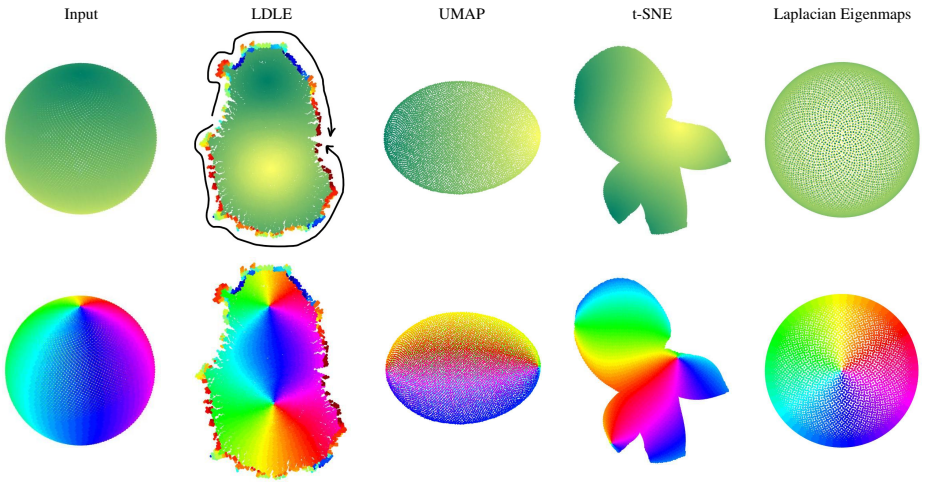

Figure 4: Embeddings of the sphere in $\mathbb{R}^3$ into $\mathbb{R}^2$. The top and bottom row contain the same plots colored by the height and the azimuthal angle of the sphere ($0 - 2\pi$), respectively. LDLE automatically colors the boundary so that the points on the boundary which are adjacent 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 spehere). t-SNE also tore apart the sphere but the embedding lacks interpretability as it is "unaware" of the boundary.

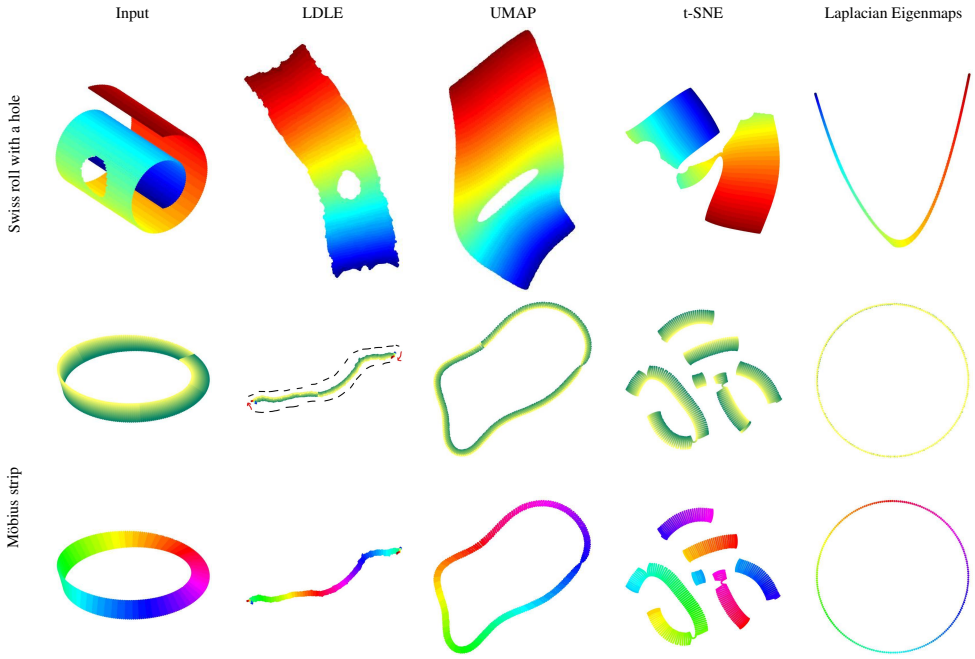

Figure 5: More examples of global embeddings. We embedded a Möbius strip in $\mathbb{R}^3$ into $\mathbb{R}^2$. The second and third row contain the same plots colored by the two parameters of the manifold. The red arrows are manually drawn to help the reader identify the two pieces of the boundary which are to be stitched together to recover the original strip. The pieces of the boundary corresponding to the dashed lines are not to be stitched together.

A SUPPLEMENTARY MATERIAL

In Figure 6, we show the 2d embeddings of 2d manifolds with boundary, in $\mathbb{R}^2$ or $\mathbb{R}^3$. To a large extent, LDLE preserved the shape of the holes. UMAP and Laplacian eigenmaps distorted the shape of the holes and the region around them, while t-SNE produced dissected embeddings. For the sphere with a hole which is a curved 2d manifold with boundary, embedded in $\mathbb{R}^3$, both UMAP and Laplacian eigenmaps squeezed it into $\mathbb{R}^2$ while LDLE and t-SNE tore it apart. The correctness of the LDLE embedding is proved in Figure 11.

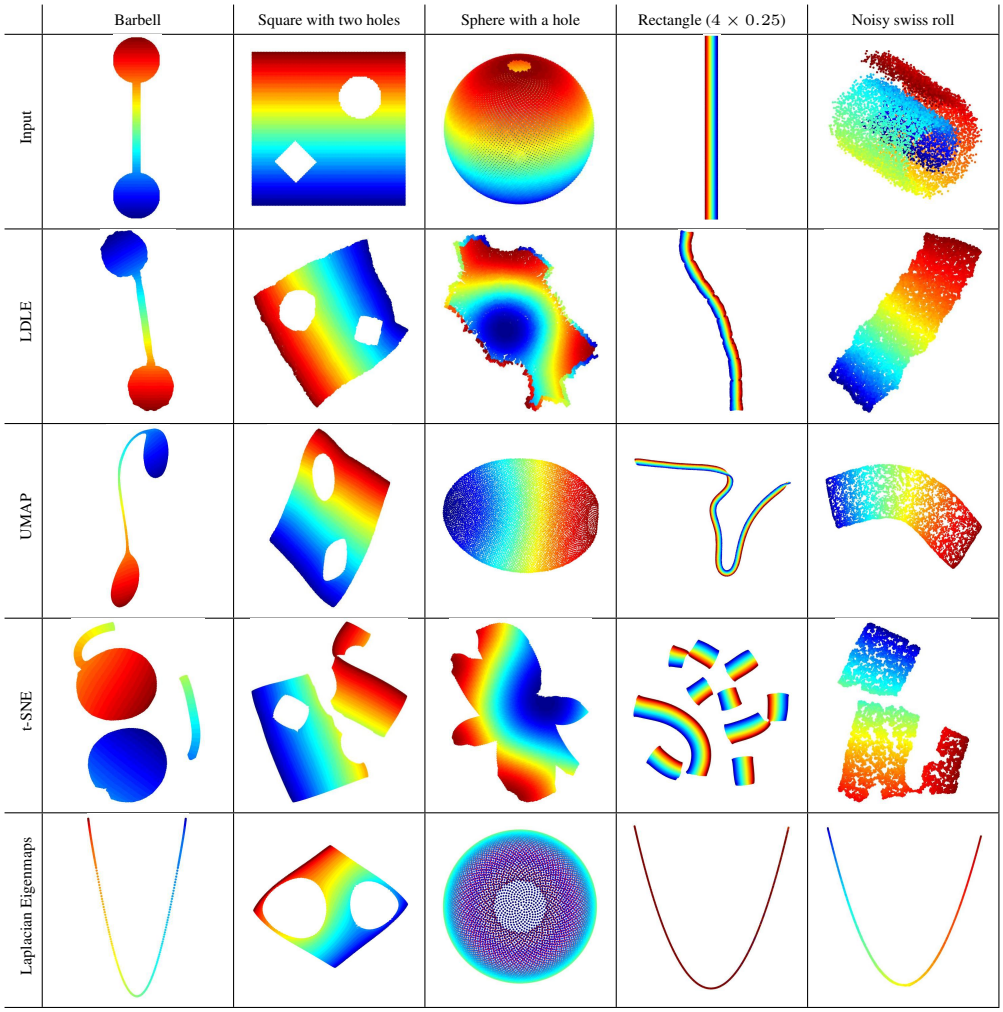

Figure 6: Embeddings of 2d manifolds with boundary into $\mathbb{R}^2$. The noisy swiss roll is constructed by adding uniform noise in all three dimensions, with support on $[0, 0.05]$.

We note that the boundaries of the LDLE embeddings in Figure 6 are usually distorted. This is because the injectivity radius at the points close to the boundary is low and precisely zero at the points on the boundary. As a result, the size of the balls around these points exceeds the limit beyond which Jones et al. (2007) result is applicable. However, they point 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 using the doubled manifold. Therefore, we obtained LDLE embeddings when the points on the boundary are known, which are shown in Figure 7.

To compare the embeddings in a quantitative manner, we plot the distortion of the geodesics originating from each point in the input in Figure 8. In the discrete setting, we take the geodesic between two given points as the shortest path between them. Using Dijkstra algorithm over a graph of 5 nearest neighbors with d_e as the underlying metric, we compute the sequence of nodes on the shortest path between x_k and $x_{k'}$. Denote the number of nodes on this path by $n_{kk'}$, the sequence of nodes by $(x_{kk'_s})_{s=1}^{n_{kk'}}$ and the length of the path in the

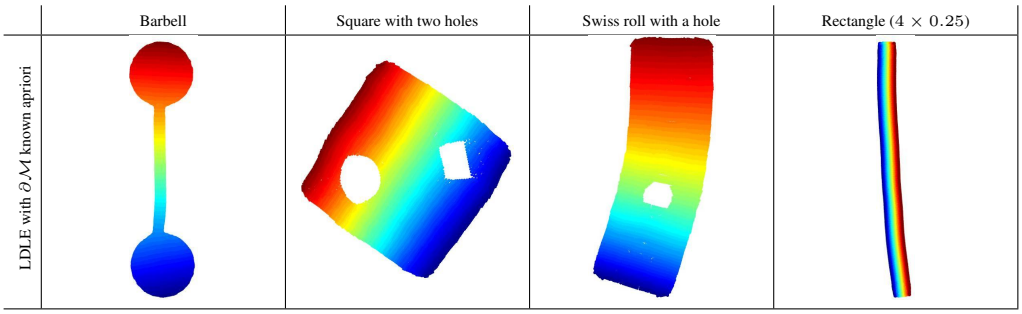

Figure 7: LDLE embeddings when the points on the boundary are known apriori.

input space by

$$L_{kk'} = \sum_{s=2}^{n_{kk'}} d_e(x_{kk'_s}, x_{kk'_{s-1}}). \tag{8}$$

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

$$L^g_{kk'} = \sum_{s=2}^{n_{kk'}} d_e(y_{kk'_s}, y_{kk'_{s-1}}). \tag{9}$$

Finally, the distortion of the paths originating from x_k is given by

$$\mathcal{D}_k = \sup_{k'} \frac{L^g_{kk'}}{L_{kk'}} \sup_{k'} \frac{L_{kk'}}{L^g_{kk'}}. \tag{10}$$

Compared to the other algorithms, the global embedding produced by LDLE least distorts the shortest geodesics in the input space, especially when the points on the boundary are known apriori.

In Figure 9, 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,

$$X(\theta_i, \phi_j) = \frac{1}{4\pi}(4cos(\theta_i), 4\sin(\theta_i), \cos(\phi_j), \sin(\phi_j)) \tag{11}$$

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

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 11.

In Figure 10, we show the 2d embedding of another non-orientatble manifold, a Klein bottle which is a non-orientable two dimensional manifold without boundary. We construct a discrete Klein bottle using its 4D Möbius tube representation as follows,

$$X(\theta_i, \phi_j) = (R(\phi_j)\cos\theta_i, R(\phi_j)\sin\theta_i, r\sin\phi_j\cos\frac{\theta_i}{2}, r\sin\phi_j\sin\frac{\theta_i}{2}) \tag{12}$$

$$R(\phi_j) = R + r\cos\phi_j \tag{13}$$

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

Laplacian eigenmaps produced incorrect embeddings, t-SNE produced dissected and non-interpretable embeddings and UMAP squeezed the inputs into $\mathbb{R}^2$. On the other hand, LDLE produced mathematically correct embeddings by tearing apart both inputs to embed them into $\mathbb{R}^2$.

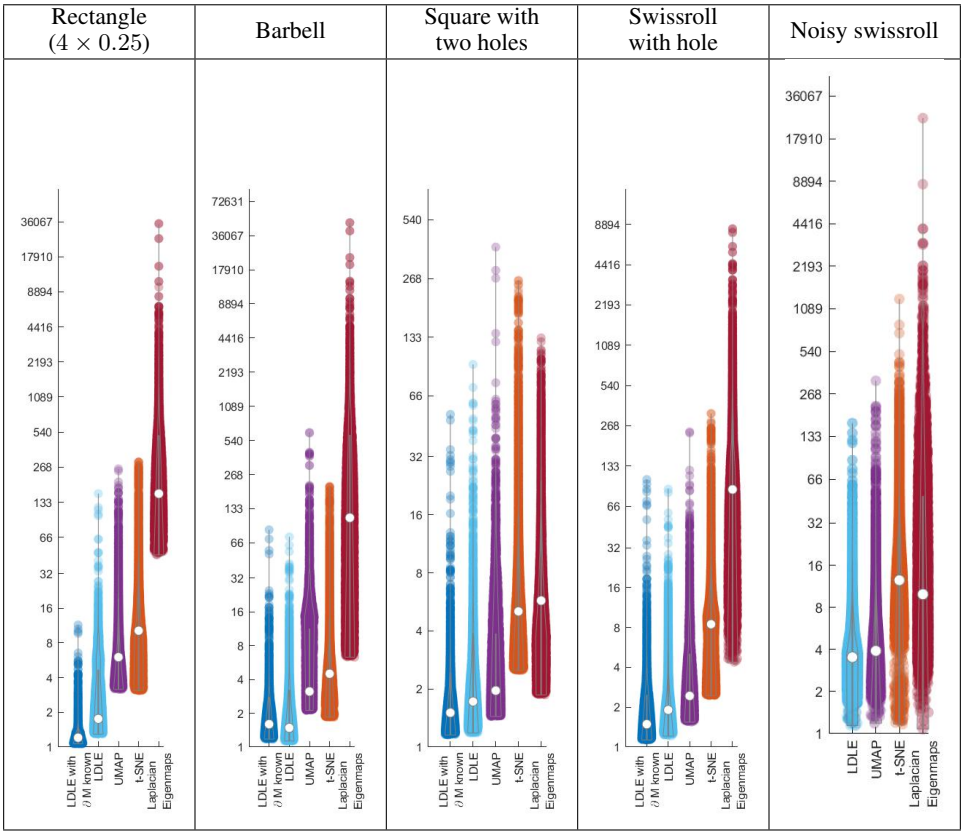

Figure 8: Distortion of the geodesics originating from each point in the input (See Eq. (10)).

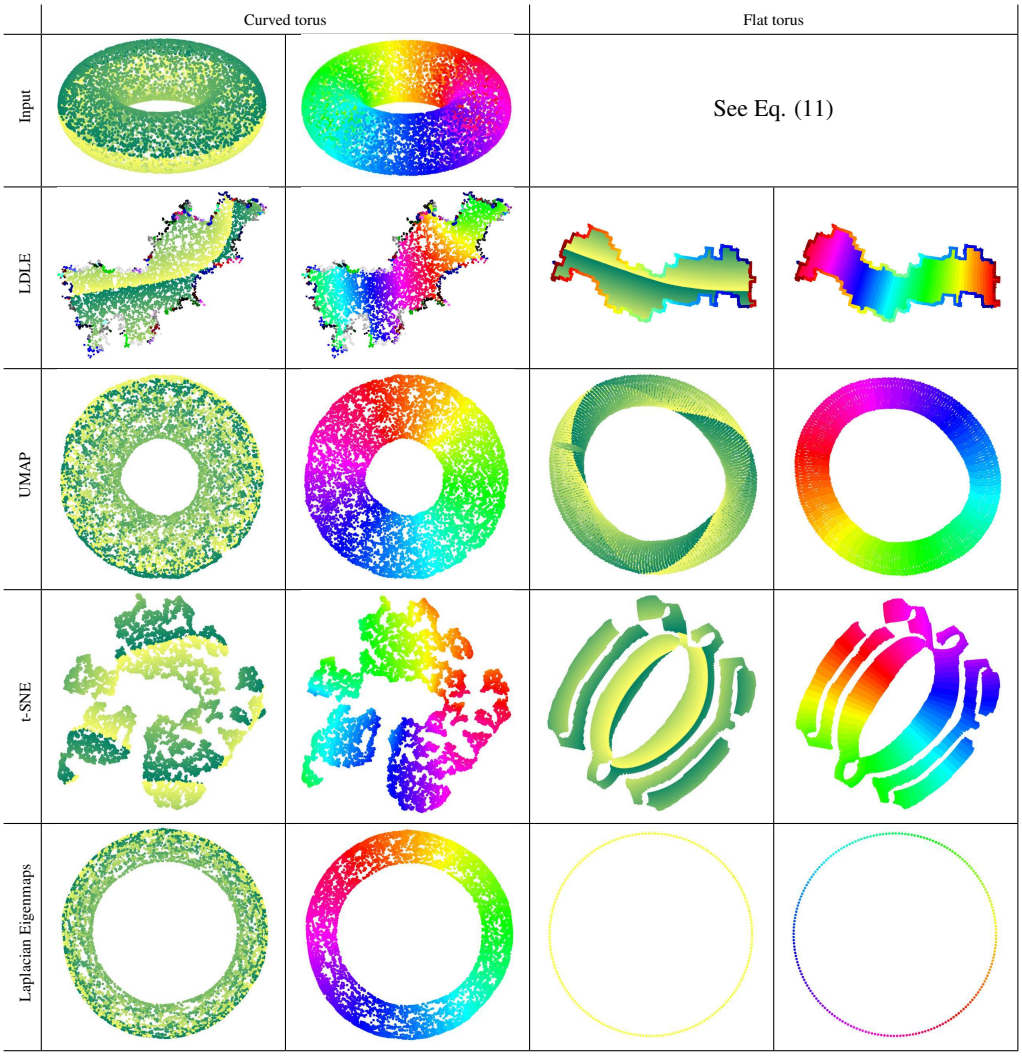

Figure 9: Embeddings of 2d manifolds without boundary into $\mathbb{R}^2$. For each manifold, the left and right columns contain the same plots colored by the two parameters of the manifold. A proof of the mathematical correctness of the LDLE embeddings is provided in Figure 11.

Input	LDLE	UMAP	t-SNE	Laplacian Eigenmaps
See Eq. (13)		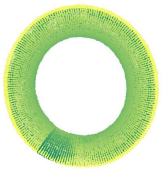		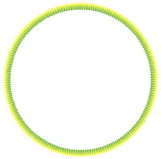
See Eq. (13)	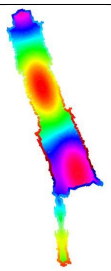	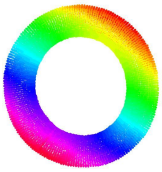	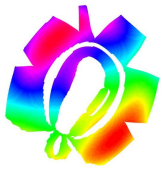	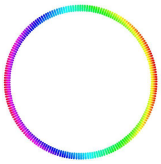

Figure 10: Embeddings of a 2d non-orientable Klein bottle into $\mathbb{R}^2$. The top and bottom row contain the same plots colored by the two parameters of the manifold. A proof of the mathematical correctness of the LDLE embeddings is provided in Figure 11.

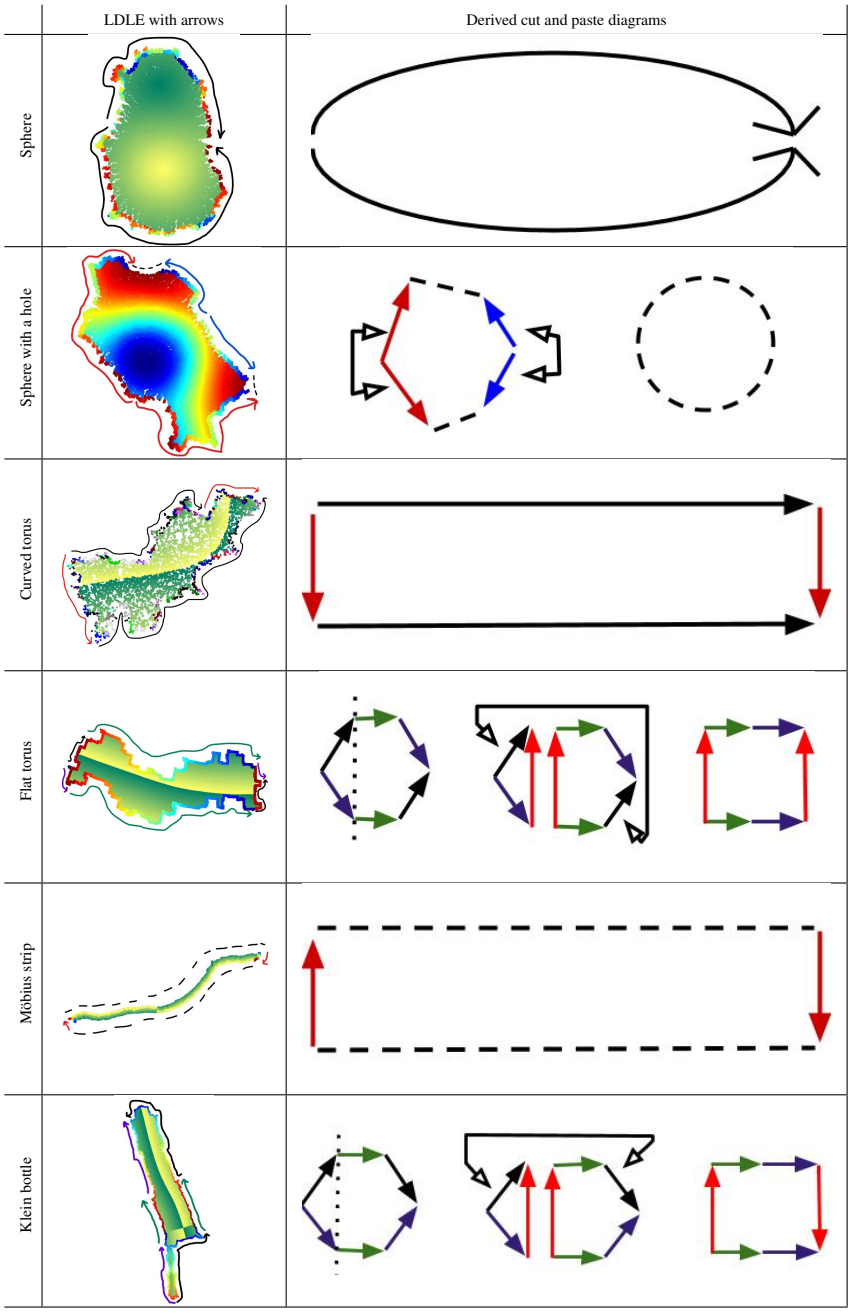

Figure 11: (Left) LDLE embedding with arrows drawn by tracing the colored boundary. (Right) Derived cut and paste diagrams to prove the correctness of the embedding. Filled arrows represent pieces of the boundary which are to be stitched along the arrows of the same color. Pieces of the boundary represented by dashed lines are not to be stitched. Dotted lines and shallow arrows represent cut and paste instructions.

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.