

Parallel Transport Unfolding: A Connection-based Manifold Learning Approach

Max Budninskiy¹, Gloria Yin¹, Leman Feng¹, Yiyong Tong², and Mathieu Desbrun¹

¹Caltech

²Michigan State University

Abstract

Manifold learning offers nonlinear dimensionality reduction of high-dimensional datasets. In this paper, we bring geometry processing to bear on manifold learning by introducing a new approach based on metric connection for generating a quasi-isometric, low-dimensional mapping from a sparse and irregular sampling of an arbitrary manifold embedded in a high-dimensional space. Geodesic distances of discrete paths over the input pointset are evaluated through “parallel transport unfolding” (PTU) to offer robustness to poor sampling and arbitrary topology. Our new geometric procedure exhibits the same strong resilience to noise as one of the staples of manifold learning, the Isomap algorithm, as it also exploits all pairwise geodesic distances to compute a low-dimensional embedding. While Isomap is limited to geodesically-convex sampled domains, parallel transport unfolding does not suffer from this crippling limitation, resulting in an improved robustness to irregularity and voids in the sampling. Moreover, it involves only simple linear algebra, significantly improves the accuracy of all pairwise geodesic distance approximations, and has the same computational complexity as Isomap. Finally, we show that our connection-based distance estimation can be used for faster variants of Isomap such as L-Isomap.

Keywords: *Isomap, metric connection, parallel transport, high-dimensional data.*

1 Introduction

With the avalanche of information in our big data era, the demand for numerical tools to analyze and classify high-dimensional data has become paramount. Recognizing and exploiting hidden structures present in the input is particularly desirable when reducing the dimensionality of the data [Vidal et al., 2016]. The manifold assumption, which posits that many high-dimensional datasets actually live on much lower dimensional curved spaces, has been proven surprisingly useful in a variety of contexts. Algorithms working towards extracting reduced dimensional representations of data, also called *manifold learning*, produce a low count of “intrinsic variables” to describe the high-dimensional input—with which markedly faster computations can be performed in the context of machine learning (to address the “curse of dimensionality”), facial animation [Kim, 2007], skeletal animation [Assa et al., 2005], video editing [Pless, 2003], and nonlinear shape editing [Budninskiy et al., 2017] to mention a few.

Perhaps the most popular algorithm for dimensionality reduction is Principal Component Analysis (PCA); yet, it is only useful when data lie on or close to a linear subspace of the high-dimensional space. Manifold learning algorithms can be viewed as nonlinear extensions to PCA: since data points are assumed to be samples from a low-dimensional manifold that is embedded in a high-dimensional space, many nonlinear dimensionality reduction algorithms attempt to uncover a quasi-isometric parametrization of this manifold. More precisely, they seek an as-isometric-as-possible mapping from the input pointset in \mathbb{R}^D into a low-dimensional space \mathbb{R}^d (with $d \ll D$) that represents a low-distortion unfolding of the original data. This is, therefore, the high-dimensional counterpart of *parameterization* in computer graphics, a task that finds a (u, v) parameterization for a two-dimensional patch embedded in 3D (e.g., to texture its surface) and for which a variety of approaches exploiting the three-dimensional nature of the embedding space have been proposed [Hormann et al., 2008, Desbrun et al., 2003, Mullen et al., 2008, Schmidt and Singh, 2009].

In this paper, we bring discrete differential geometry concepts to bear on this ubiquitous and inherently geometric issue, and show how current limitations on geodesic convexity of the input manifold can be lifted elegantly using the Levi-Civita connection, for the same computational complexity.

1.1 Previous Work

There are two distinct flavors of manifold learning in existing work: local methods, which only rely on local measurements and sparse matrix eigenanalysis to recover the low-dimensional manifold, and global methods, that leverage intrinsic measurements between all data points and involve eigenanalysis of a dense matrix.

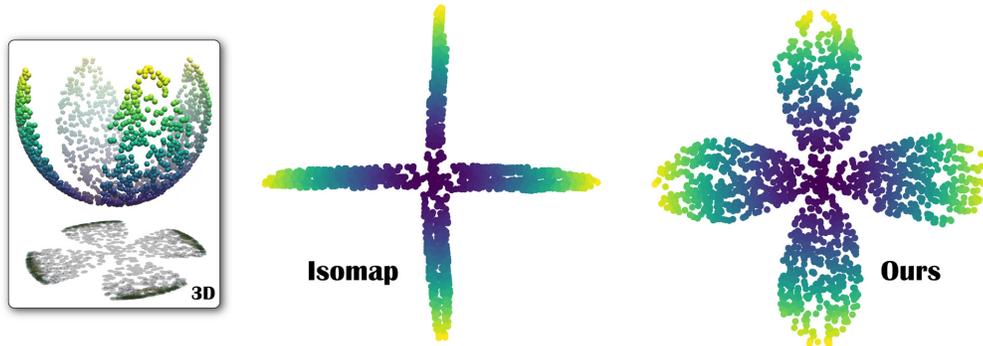


Figure 1: Correctly Unfolding Petals. For a 3D sampling of a 4-petal shaped portion of a sphere (left), Isomap (a staple of manifold learning) fails to find a near isometric 2D parameterization (middle) due to geodesic non-convexity of the intrinsic geometry. Our parallel transport approach, instead, handles this case as expected (right). Lifting the pointset to much higher dimensions and applying random rotations and reflections would not change these results.

Local methods. Local embedding methods (e.g., Laplacian Eigenmaps [Belkin and Niyogi, 2001], Locally Linear Embeddings [Roweis and Saul, 2000] and their many variants) consider local neighborhoods of input points to infer the whole structure of the data through global stitching of these neighborhoods: for each of the n input points, they determine its relative position with respect to its immediate surroundings, and then rely on a sparse matrix eigenanalysis to deduce a global low-dimensional embedding that best preserves local relative positions. While no restrictive assumptions on the geometry or topology of the manifold are required, this approach to nonlinear dimensionality reduction often generates large distortion in the resulting embedding—an issue fortunately addressed by enforcing local linear precision, see [Donoho and Grimes, 2003, Zhang and Wang, 2007, Budninskiy et al., 2017]. The local nature of this family of methods typically implies a computational complexity of $\mathcal{O}(n^{1.5})$ due to the sparsity of the matrix involved in the eigenanalysis. Yet this same locality hampers its robustness: the presence of noise and outliers often leads to degenerate results [Van Der Maaten et al., 2009]. Local embedding methods are thus general and efficient enough to handle arbitrary inputs, but are often not resilient enough to noise to be practical (see Figs. 14 and 6 for instance).

Global methods. The Isomap technique [Tenenbaum et al., 2000] is a variant of MultiDimensional Scaling (MDS) [Torgerson, 1965, Crippen and Havel, 1978] that attempts to reduce distortion in the mapping by preserving pairwise *geodesic* distances between all data points as well as possible. After forming a k -nearest neighbor graph of the data points as a representation of the input manifold, Isomap first solves the all-pairs shortest path problem (using Dijkstra’s or Floyd-Warshall algorithm) to approximate geodesic distances between every pair of points in the graph. Finally, a low-dimensional embedding that best preserves these pairwise distances is then constructed with MDS through eigenanalysis of the corresponding Gram matrix [Borg and Groenen, 2005]. The high computational complexity due to the eigenanalysis of a dense Gram matrix (in $\mathcal{O}(n^3)$) can be further improved through probabilistic linear algebra [Halko et al., 2011], variants such as L-Isomap [De Silva and Tenenbaum, 2002] or by exploiting various numerical improvements of MDS [Harel and Koren, 2002, Brandes and Pich, 2007, Khoury et al., 2012] that approximate its solution in near linear time. Since *all* pairwise geodesic distances are used, Isomap is particularly robust to noise, and for this reason it remains one of the most popular algorithms used for manifold learning.

1.2 Shortcomings of Isomap

Despite its resilience to noise, Isomap lacks severely in generality: it can *only* handle geodesically convex regions correctly. In particular, if the input sampling of the manifold has a large void, Isomap will significantly overestimate geodesic distances since graph-based shortest paths for points on opposite side of the hole will be forced to go around the empty region, thus resulting in a distorted embedding (Fig. 2). This issue is all the more serious as it impacts irregular sampling in general: unlike PCA, Isomap fails to produce a non-distorted embedding for a dataset lying precisely on a linear subspace of the ambient high-dimensional space, since graph-based shortest paths are only approximations of true manifold geodesics, see Fig. 3. There have been attempts to alleviate this limitation, such as the Topologically Constrained Isometric Embedding (TCIE) approach [Rosman et al., 2010] which first tries to detect the boundaries of holes and the geodesic paths that go through these boundaries, before constructing a map that ignores these biased geodesic paths. However, the process of detecting boundaries is highly susceptible to noise, since it can often be unclear whether an empty region delimits a “real” void or whether the local sampling is merely low. Moreover, the

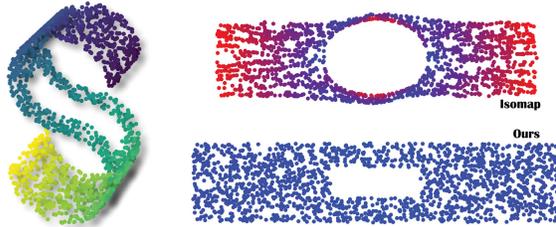


Figure 2: Holey S. *Isomap (top right) fails to find a quasi-isometric embedding of a uniformly sampled developable S-shaped 2-manifold with a sampling void in the middle (left) in 3D. Our approach (bottom right) unfolds it almost perfectly. A color ramp (blue: 0% error, to red: 15% error) is used in the visualization of the two embeddings.*

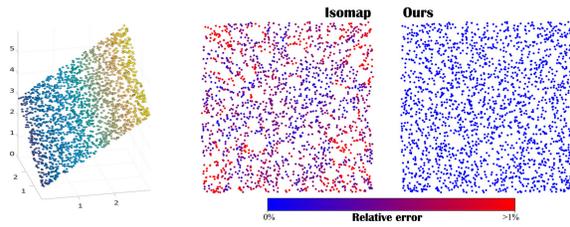


Figure 3: Linear Precision. *Due to its reliance on graph-based paths, Isomap distorts a flat 2D pointset embedded in 3D; instead, our connection-based approach (right) flattens it exactly. A linear color ramp from blue (0% ℓ^2 -error in position relative to bbox size) to red (>1% error) is used in the visualization of the two embeddings.*

subsequent minimization is significantly more difficult to perform as it no longer relies on a simple partial eigendecomposition. Consequently Isomap is often used in practice, even when the topology and geometry of the sampled manifold are not known—because of the lack of another robust learning approach to use instead.

1.3 Shortcomings of Isomap

Despite its resilience to noise, Isomap lacks severely in generality: it can *only* handle geodesically convex regions correctly. In particular, if the input sampling of the manifold has a large void, Isomap will significantly overestimate geodesic distances since graph-based shortest paths for points on the opposite side of the hole will be forced to go around the empty region, thus resulting in a distorted embedding (Fig. 2). This issue is all the more serious as it impacts irregular sampling in general: unlike PCA, Isomap fails to produce a non-distorted embedding for a dataset lying precisely on a linear subspace of the ambient high-dimensional space, since graph-based shortest paths are only approximations of true manifold geodesics, see Fig. 3. There have been attempts to alleviate this limitation, such as the Topologically Constrained Isometric Embedding (TCIE) approach [Rosman et al., 2010] which first tries to detect the boundaries of holes and the geodesic paths that go through these boundaries, before constructing a map that ignores these biased geodesic paths. However, the process of detecting boundaries is highly susceptible to noise, since it can often be unclear whether an empty region delimits a “real” void or whether the local sampling is merely low. Moreover, the subsequent minimization is significantly more difficult to perform as it no longer relies on a simple partial eigendecomposition. Consequently Isomap is often used in practice, even when the topology and geometry of the sampled manifold are not known—because of the lack of another available robust learning approach.

1.4 Overview and contributions

In this paper we introduce Parallel Transport Unfolding (PTU), a manifold learning technique that performs nonlinear dimensionality reduction to produce a quasi-isometric, low-dimensional embedding of an arbitrary set of high-dimensional data points. By simply replacing the Dijkstra path-based geodesic distance estimates with parallel transport based approximations instead, our approach removes the geometric limitations of Isomap: it can reliably handle arbitrary data with strongly irregular sampling while retaining its resilience to noise. We show that this new geometric procedure no longer requires geodesic convexity of the domain sampled by the input data, only involves simple linear algebra, significantly improves the accuracy of all pairwise geodesic distance approximations, and does not change the overall computational complexity of the original Isomap procedure. Moreover, our approach exploits the low dimensionality of the manifold by using a connection on the d -dimensional tangent bundle for efficiency purposes: this markedly differs from previous parallel transport based geodesic computations on 2-manifolds in 3D, where the codimension is 1 and thus, where either the normal field or 2D polar coordinates can be leveraged to derive fast geodesic approximations [Schmidt et al., 2006, Melvær and Reimers, 2012, Schmidt, 2013]. Finally, we also demonstrate that our connection-based distance estimation applies equally well to Landmark-Isomap [De Silva and Tenenbaum, 2002], a variant offering significant improvements in computational time but suffering from the same convexity limitation as Isomap.

1.5 Notations and assumptions

Throughout our exposition, we consider an input data set $\mathcal{S} = \{\mathbf{x}_i\}_{i=1..n}$ of n points that irregularly sample (possibly with noise) a connected compact d -manifold \mathcal{M} embedded in a D -dimensional metric space. While the use of arbitrary kernels has been proposed as a unifying view [Ham et al., 2004, Weinberger et al., 2004], we will assume the embedding space to be \mathbb{R}^D for simplicity, with $d \ll D$. We only assume that \mathcal{M} possesses

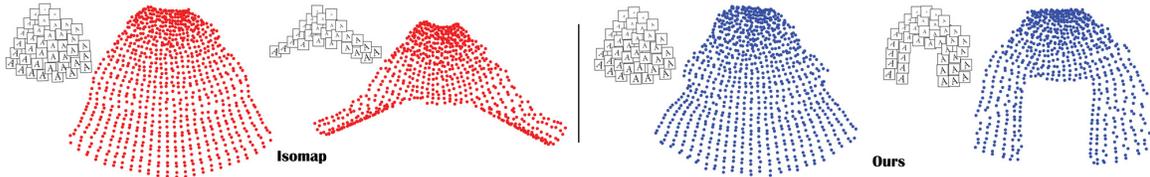


Figure 4: Letter A. From 888 images (120×120 pixels) of rotated and resized letters ‘A’, both Isomap (red) and PTU (blue) produce similar 2D embeddings (left, each dot indicating an image; insets show a subset of the images and their embedding). However, if a part of these input images is removed (rendering the set non-geodesically-convex), Isomap dramatically changes the embedding, while PTU properly reflects the missing images (right).

an atlas with a single chart, so that an injective d -dimensional parametrization of the manifold exists. Our parallel transport based manifold learning algorithm seeks to find an as-isometric-as-possible d -dimensional embedding of the input S as a pointset $\{\mathbf{z}_i\}_{i=1..n} \subset \mathbf{R}^d$. Finally, to simplify further expressions, we assemble a $D \times n$ matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ from the input points and denote the final embedding as a $d \times n$ matrix $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)$.

2 Primer on Isomap

Before describing our method in more detail, we first discuss the original Isomap algorithm.

2.1 Isomap at a glance

The Isomap algorithm [Tenenbaum et al., 2000] for finding a low-dimensional, quasi-isometric embedding of a point set S consists of three steps:

1. Construct a proximity graph G over the point set S ;
2. Evaluate pairwise geodesic distances between elements of S via Dijkstra’s algorithm on G ;
3. Perform MDS on the resulting distances to find a quasi-isometric d -dimensional embedding.

2.2 Computations involved

Each algorithmic step of Isomap involves simple computations that we review next.

Proximity graph. A graph is first constructed by creating undirected edges between neighboring input points. Two simple ways have been proposed to define whether two points should be connected by an edge of the neighborhood graph: the first (k -nearest neighbor, or k -NN) approach declares two points neighbors iff one is among the k nearest neighbors of the other based on Euclidean distances in \mathbf{R}^D ; the second (ϵ -ball) approach declares two points neighbors iff the Euclidean distance between them is smaller than a user-defined threshold ϵ . Both constructions can be done efficiently (typically, in $O(n \log n)$) using a locality sensitive hashing data structure for instance [Datar et al., 2004].

Geodesic distances. After setting edge weights to the Euclidean distances between corresponding pairs of points, the next step is to run Dijkstra’s algorithm on the resulting weighted graph G to compute approximations of all pairwise geodesic distances between points of S in $O(n^2 \log n)$. The squares of these pairwise geodesic distances are then assembled into an $n \times n$ symmetric matrix \mathbf{D} .

Multidimensional scaling. Finally, the (classical) MDS procedure is performed on \mathbf{D} to obtain the low-dimensional embedding that preserves these squared distances as well as possible. Namely, an $n \times n$ Gram matrix \mathbf{G} is constructed from \mathbf{D} via “double-centering” [Borg and Groenen, 2005]; that is, denoting identity matrix by \mathbf{I} and the n -vector of ones by \mathbf{e} , $\mathbf{G} = -\frac{1}{2}(\mathbf{I} - \frac{1}{n}\mathbf{e}\mathbf{e}^t)\mathbf{D}(\mathbf{I} - \frac{1}{n}\mathbf{e}\mathbf{e}^t)$. (This Gram matrix may be further altered to represent a proper kernel, see [Choi and Choi, 2007]). The final embedding \mathbf{Z} is found as the product of square roots of the d largest eigenvalues and corresponding eigenvectors of \mathbf{G} :

$$\mathbf{Z} = \sqrt{\Lambda_d} \mathbf{Q}_d^t$$

where $\mathbf{G} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^t$ is the eigendecomposition of \mathbf{G} , and $(\Lambda_d, \mathbf{Q}_d)$ denote the truncated matrices containing only the d largest eigenvalues and their corresponding eigenvectors respectively. Consequently, the embedding is found by computing the d largest eigenvalues and their associated eigenvectors, a numerical task of expected complexity $O(n^3)$ since the Gram matrix is dense. MDS is, in fact, a variational approach since the solution is the minimizer over all matrices \mathbf{Z} of the squared Frobenius norm $\|\mathbf{Z}'\mathbf{Z} - \mathbf{G}\|_F^2$ subject to $\text{rank}(\mathbf{Z}) = d$.

2.3 Discussion

Despite the dense nature of the Gram matrix which implies a higher computational complexity than local methods, manifold learning through Isomap is one of the most used nonlinear dimensionality reduction methods because of its remarkable robustness to noise. Its entire reliance on graph-based shortest paths has, however, far-reaching calamitous consequences. First, spurious geodesic curvature (i.e., zigzags) in the shortest paths (seen as a piecewise linear curve in \mathbb{R}^D) between two nodes on the graph introduces inaccuracies in the estimation of geodesic distances. In practice, this issue is exacerbated by the sparse sampling that real applications often have to deal with, even if short paths can be locally rectified through a straightening projection [Budninskiy et al., 2017]. Worse, some computational accelerations of Isomap rely on a subsampling of the initial data [De Silva and Tenenbaum, 2004], making this inaccuracy issue all the more limiting. More importantly, Isomap can only provide a quasi-isometric low-dimensional mapping for *geodesically convex sets*: the presence of hole(s) or non-convex boundaries in the domain sampled by the input pointset brings significant overestimations of geodesic distances, thus distorting the results (see Fig. 2 and Fig. 4). Detecting and correcting paths that go around small or large voids is difficult to perform reliably as voids can have a variety of sizes and shapes when dealing with noisy and irregular sampling.

3 Parallel Transport Unfolding

We now present our Parallel Transport Unfolding (PTU) algorithm, which has a very similar structure to Isomap since it comprises the following steps:

- Construct a proximity graph G of the pointset \mathcal{S} and compute local tangent spaces at each point;
- Approximate all pairwise geodesic distances using parallel transport along (shortest) paths on G ;
- Perform MDS to find a d -dimensional embedding that best preserves all the geodesic distances.

We detail each step next, stressing the key differences with Isomap in the first two steps of the algorithm, i.e., the use of approximate tangent spaces and of the Levi-Civita connection to better approximate intrinsic geodesic distances.

3.1 Proximity graph

The construction of a proximity graph G on \mathcal{S} proceeds similarly to the original Isomap algorithm: one can link each point to its neighbors contained in an ϵ -ball, or to its k nearest neighbors—both based on Euclidean distances in \mathbb{R}^D . Variants such as a mix of the two [Wang et al., 2006] or the mutual k nearest neighbors approach [Brito et al., 1997] can also be used to naturally discard outliers. For clarity of presentation, we will use a vanilla k -NN graph in our exposition and all of our tests. The value k should be chosen such that the edges of the resulting graph are good approximations of geodesics between corresponding points; we will typically set k around $4d$ to induce a valence greater than the usual connectivity of a regular grid of the d -dimensional embedding ($k = 2d$), yet less than the number of 1-ring neighbors on that same grid ($k = 3^d - 1$); but knowledge about noise levels in the input pointset can be employed to improve the graph quality by varying the number of neighbors from point to point. Once the graph connectivity is defined, each edge is assigned a weight equal to its Euclidean length in \mathbb{R}^D as an approximation of its geodesic length on the manifold \mathcal{M} .

3.2 Tangent spaces and their orthonormal bases

For each input point \mathbf{x}_i , its K nearest neighbors on the proximity weighted graph (for $K \geq k$) are used to define a *geodesic neighborhood*. The matrix whose columns are the data points (centered by subtracting \mathbf{x}_i) in this neighborhood induces, through its d left singular vectors corresponding to the d largest singular values, an orthonormal basis \mathbb{T}_i of d vectors in \mathbb{R}^D spanning the approximate tangent space:

$$\mathbb{T}_i = (\mathbf{t}_1^i \quad \cdots \quad \mathbf{t}_d^i) \in \mathbb{R}^{D \times d}. \quad (1)$$

While choosing $K = k$ often suffices to provide a good estimate of local tangent spaces, using a value of K distinct from k allows the definition of arbitrarily large geodesic neighborhoods (useful in the presence of strong noise) around each point of \mathcal{S} , while alleviating the traditional issue of “manifold shortcutting” associated with increasing k [Budninskiy et al., 2017]. As a substitute to this frame construction via partial SVD, note that an improved approximation of the local tangent spaces in the presence of strong noise and outliers can also be computed via ℓ_1 -based robust PCA (see, e.g., [Zhang and Lerman, 2014]).

With these orthonormal frames of tangent spaces in place, we can now discuss how to approach discrete parallel transport and how to use it for geodesic length estimation.

3.3 Discrete Parallel transport

We now cover the core of our approach, i.e., exploiting parallel transport in high dimension to better evaluate geodesic distances.

Parallel Transport in Differential Geometry. The notion of parallel transport plays a central role in differential geometry. It induces a way to connect the geometries of nearby points, thus prescribing how a basis of the tangent space at one point of a manifold should be adjusted to produce a “parallel” basis of another tangent space at a nearby point. While this procedure is straightforward for flat spaces (it corresponds to a simple translation), it becomes more involved for manifolds with non-trivial curvature. Its differential geometric treatment involves the definition of a connection on the tangent bundle [Spivak, 1979, Frankel, 2011], which represents an infinitesimal analogue of parallel transport. Most relevant to our work are *metric* connections, i.e., connections such that the parallel transport they define preserves the intrinsic metric of the manifold. The well-known Christoffel symbols are, in fact, the components of a particularly canonical metric connection called the Levi-Civita connection, which we will leverage in our application. Another geometric property we will exploit is the fact that geodesics, usually described through variational analysis as locally shortest curves, can also be defined through parallel transport: a geodesic is a curve that parallel transports its own tangent vector as directly implied by the geodesic equation [Spivak, 1979, Frankel, 2011]. This simple property will guide our evaluation of geodesic distances.

Discrete Parallel Transport. Given points \mathbf{x}_i and \mathbf{x}_j sharing an edge in the proximity graph G , we define the *discrete metric connection* between \mathbf{x}_i and \mathbf{x}_j as the orthogonal $d \times d$ matrix $\mathbf{R}_{j,i}$ in $O(d)$ representing the change of basis that best aligns, in the Frobenius norm, the frames \mathbb{T}_i and \mathbb{T}_j , i.e.,

$$\mathbf{R}_{j,i} = \arg \min_{\mathbf{R} \in O(d)} \|\mathbb{T}_i - \mathbb{T}_j \mathbf{R}\|_F^2 \quad (2)$$

By definition, $\mathbf{R}_{i,j}$ is the inverse of $\mathbf{R}_{j,i}$: $\mathbf{R}_{i,j} = \mathbf{R}_{j,i}^t$. Note that we use the group $O(d)$ of orthogonal matrices because the SVD used in Sec. 3.2 produces arbitrarily-oriented tangent frames. A pre-processing of these tangent bases could be performed (e.g., via a minimum spanning tree) if one wants to ensure that their orientations are consistent, enforcing that the discrete connections are pure rotations (in $SO(d)$) as in the continuous case. The transformation $\mathbf{R}_{j,i}$ can thus be understood as a discrete equivalent to the Levi-Civita connection induced by the metric on \mathcal{S} inherited from the Euclidean space \mathbb{R}^D . It is also a high-dimensional extension of previous discretization of metric connections on triangle meshes [Crane et al., 2010, Liu et al., 2016].

Computing this discrete connection is easily achieved via a singular value decomposition:

Proposition 1: Let $\mathbb{T}_i^t \mathbb{T}_j = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^t$ be the SVD. The discrete connection is expressed as:

$$\mathbf{R}_{j,i} = \mathbf{V} \mathbf{U}^t. \quad (3)$$

Proof: The definition of $\mathbf{R}_{j,i}$ in Eq. (2) maximizes $\text{Tr}[\mathbb{T}_i^t \mathbb{T}_j \mathbf{R}]$ for $\mathbf{R} \in O(d)$ as the other terms do not depend on \mathbf{R} . Since $\mathbb{T}_i^t \mathbb{T}_j \mathbf{R} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^t \mathbf{R}$ and given the invariance of the trace under cyclic permutations, the optimal rotation maximizes $\text{Tr}[(\mathbf{V}^t \mathbf{R} \mathbf{U}) \mathbf{\Sigma}]$ for $\mathbf{V}^t \mathbf{R} \mathbf{U} \in O(d)$. Consequently, one has $\mathbf{V}^t \mathbf{R} \mathbf{U} = \mathbf{Id}$; so $\mathbf{R}_{j,i}$ must be $\mathbf{V} \mathbf{U}^t$. One recognizes the Procrustes superimposition of two nearby frames [Kabsch, 1976], just extended to handle arbitrary choices of frame orientation for tangent spaces. ■

3.4 Local Path Unfolding via Parallel Transport

With a discrete metric connection, we can now “unfold” a polyline path on \mathcal{S} (made out of a series of adjacent proximity graph edges) into \mathbb{R}^d , i.e., we seek to *map a polyline of graph edges into a flat d -dimensional space while best preserving its metric properties such as length, intrinsic curvatures, etc.* This process of “unrolling” a curve onto a tangent space is known as Cartan’s development [Nomizu, 1978, Sharpe, 1997] in differential geometry.

Two-edge Unfolding. Let us define how to unfold in \mathbb{R}^d a three-point $(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k)$ polyline path on \mathcal{M} , i.e., a path made of two adjacent edges of the proximity graph on \mathcal{S} . Its extensions to arbitrary paths will be straightforward. Without loss of generality, we can consider the unfolding to be happening in the tangent space at \mathbf{x}_i , equipped with its orthonormal frame \mathbb{T}_i . The first point \mathbf{y}_i of the unfolded polyline can be chosen to be at the origin of this subspace, or coinciding with \mathbf{x}_i . The first edge, represented by the vector $\mathbf{e}_i = \mathbf{x}_j - \mathbf{x}_i$, is projected (in the ℓ^2 sense) onto the tangent space at \mathbf{x}_i to form a d -dimensional tangent vector \mathbf{v}_i through:

$$\mathbf{v}_i = \mathbb{T}_i^t \mathbf{e}_i.$$

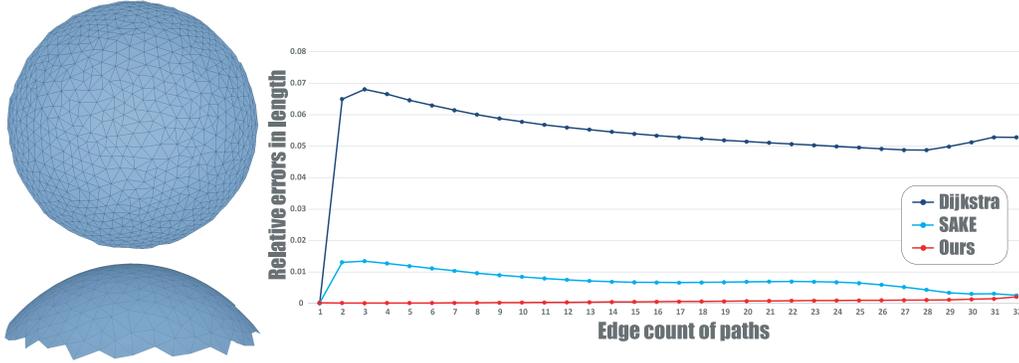


Figure 5: Geodesic Distances. For a graph based on a triangulated spherical cap (left, 2 views), the average relative error of pairwise Dijkstra-based geodesic distances is 5.6%; SAKE [Budninskiy et al., 2017] already reduces the error to 0.8%, but PTU brings it down to 0.046%. Plots (right) show relative errors of pairwise distances as a function of their number of edges.

This vector \mathbf{v}_i is then added to \mathbf{y}_i to form the point \mathbf{y}_j , thus defining the first unfolded edge. The second edge, $\mathbf{e}_j = \mathbf{x}_k - \mathbf{x}_j$ is similarly projected onto the local tangent space defined by frame \mathbb{T}_j at \mathbf{x}_j . The resulting vector can then be parallel transported onto the tangent space at \mathbf{x}_i to become:

$$\mathbf{v}_j = \mathbf{R}_{i,j} \left[\mathbb{T}_j^t \mathbf{e}_j \right],$$

where the term in brackets is the ℓ^2 projection of \mathbf{e}_j onto \mathbb{T}_j . From this vector, now represented in the original tangent frame at \mathbf{x}_i , we construct the final point of the unfolded polyline as: $\mathbf{y}_k = \mathbf{y}_j + \mathbf{v}_j$.

Preservation of geodesic curvature. Due to our use of a discrete metric connection, the unfolding procedure we described has an important property: it nearly preserves the geodesic curvature of the initial curve. Indeed, we used a discretization of the metric-preserving Levi-Civita connection, so the *intrinsic* angle between vectors \mathbf{e}_i and \mathbf{e}_j (meaning, the angle measured on the manifold \mathcal{S}) is preserved by parallel transport, and corresponds to the angle between $(\mathbf{y}_i, \mathbf{y}_j)$ and $(\mathbf{y}_j, \mathbf{y}_k)$ in \mathbb{R}^d —up to discretization errors. This means that if the two-edge polyline were a good approximation of a geodesic on \mathcal{S} , the unfolded polyline would be (nearly) straight in \mathbb{R}^d since a geodesic curve parallel transports its tangent vector. As a corollary, the Euclidean distance $\|\mathbf{v}_i + \mathbf{v}_j\|_2$ between \mathbf{y}_i and \mathbf{y}_k is a good approximation of the geodesic distance between \mathbf{x}_i and \mathbf{x}_k along \mathcal{S} —whether the polyline is an approximate geodesic or not.

Length rescaling. In case of highly-irregular samples with small or no noise, it can be beneficial to preserve the length of each edge \mathbf{e}_i once it is projected onto the corresponding frame, as the projection \mathbf{v}_i on the tangent space can be shortened if the manifold is curved. This adjustment is easily achieved by, for instance, setting

$$\mathbf{v}_i = \mathbb{T}_i^t \mathbf{e}_i \frac{\|\mathbf{e}_i\|_2}{\|\mathbb{T}_i^t \mathbf{e}_i\|_2}. \quad (4)$$

If strong noise and/or outliers are present, this modification should be avoided as it tends to overestimate local distances around noisy points. We only used this length rescaling procedure in Figs. 4 and 13 to slightly improve the results in these extreme cases.

3.5 Distance approximation via Parallel Transport

The unfolding procedure we just introduced can be used to evaluate pairwise geodesic distances without the traditional shortcomings of Dijkstra’s algorithm for graphs over an irregular sampling.

Correcting shortest paths. Consider a Dijkstra shortest polyline $(\mathbf{x}_1, \dots, \mathbf{x}_m)$. Using the unfolding procedure into \mathbb{T}_1 , we iteratively project the edges $\mathbf{e}_s = \mathbf{x}_{s+1} - \mathbf{x}_s$ onto the tangent space spanned by \mathbb{T}_s , before parallel transporting the resulting vector back to the original tangent space, i.e.,

$$\mathbf{v}_s = \left(\prod_{j=1..s-1} \mathbf{R}_{j,j+1} \right) \left[\mathbb{T}_j^t \mathbf{e}_s \right]. \quad (5)$$

After accumulating the results into a single vector $\mathbf{v} = \sum_{s=1..m} \mathbf{v}_s$, we set the approximate geodesic distance between \mathbf{x}_1 and \mathbf{x}_m to be its length $\|\mathbf{v}\|_2$: as discussed above, the span of the unfolded path is a better approximation of the geodesic distance between the end points of the initial path as it ignores the intrinsic twists and turns that the path went through. Since our estimate of the geodesic distance from \mathbf{x}_p to \mathbf{x}_q is in general not the same as the estimate from \mathbf{x}_q to \mathbf{x}_p due to the asymmetry of the discrete parallel transport, we

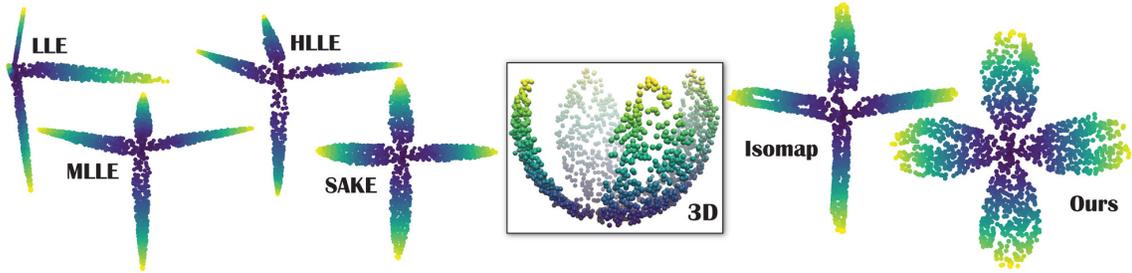


Figure 6: Noisy Petals. Given a 3D sampling of a 4-petal shaped portion of a sphere (see Fig. 1) with added Gaussian noise in the normal direction (σ : 3% of sphere radius, middle), PTU recovers an almost perfect quasi-isometric 2D parametrization, while Isomap still fails (right). Local methods, not exploiting large geodesic distances, fail even worse (left) with the notable exception of SAKE that performs better than Isomap.

average the two spans in a final post-processing step to determine all the pairwise geodesic lengths. Fig. 5 demonstrates an improvement of *over two orders of magnitude* compared to Dijkstra’s approach to estimating geodesic distances.

Leveraging unfolded paths. While we only need the Euclidean distance between the two end points of the unfolded path in \mathbb{R}^d to perform nonlinear manifold learning, the unfolded path can also be useful for gathering additional information. For instance, the largest distance between the straight line between the two end points and the actual unfolded path indicates how far off the path is from being a geodesic. A large value is a typical telltale of the presence of a large void in the data within the manifold, which can be exploited to suggest where to insert new samples in order to improve the results. Similarly, a large difference between the distance estimates computed from one end point to the other and in the opposite direction implies potential issues with sampling density of the data compared to the local curvature of the manifold.

Connection-based Dijkstra’s algorithm. Finally, we point out that our parallel transport approach to estimating geodesic distances could be done along *approximately-shortest* polyline between a given pair of points: fast approximations of shortest paths could thus be used to lower this $\mathcal{O}(n^2 \log n)$ step without significant effect on the results. However, the longest the polyline, the more likely numerical inaccuracies induced by repeated alignment of frame fields will accumulate. Since Dijkstra’s algorithm is not the computational bottleneck in Isomap, we decided to utilize graph-based shortest paths: in fact, our construction can be neatly incorporated within the dynamic programming approach that Dijkstra’s algorithm uses—i.e., storing the predecessor to each point in the shortest path found thus far. Parallel transport unfolding only requires the insertion of six lines in the original Dijkstra’s algorithm in order to reduce the unnecessarily repeated calculations that result from the fact that sub-paths of shortest paths are also shortest paths. Specifically, every time a point is removed from the priority queue (i.e., every time the algorithm finds a shorter path to a point), we can compute the corresponding vector \mathbf{v}_i , before storing the cumulative connection (i.e., the product of connection matrices along the path) to be the cumulative connection of the predecessor multiplied by the local discrete connection. Proceeding in this way guarantees that the calculations are performed without redundancies, adding a negligible amount of computational time to the traditional Dijkstra’s algorithm. Pseudocode is given in Alg. 2. We note finally that our approach of incorporating an intrinsic connection into the classical Dijkstra algorithm can be seen as a variant of the so-called vector-valued Dijkstra methods [Schmidt et al., 2006, Campen et al., 2013].

3.6 MDS

Finally, once we evaluated all pairwise geodesic distances, the results can be processed through MDS: double-centering the matrix of *squared* distances produces the Gram matrix, whose partial eigendecomposition returns the final embedding as in Sec. 2.2. Numerical improvements (see [Harel and Koren, 2002, Brandes and Pich, 2007, Halko et al., 2011, Khoury et al., 2012]) can be applied to reduce the cubic complexity of this stage. Note that robust alternatives to classical MDS could also be used here (e.g., [Cayton and Dasgupta, 2006, Kovnatsky et al., 2016]) to offer resilience to outliers, but at higher computational cost.

3.7 Discussion

We conclude this section with a few properties worth mentioning.

Proposition 2: *Unlike Isomap, PTU is linearly precise as long as one uses a proximity graph G where each sample point has enough neighbors to span a d dimensional subspace: assuming the pointset S samples a linear d -dimensional subspace \mathcal{R} of \mathbb{R}^D , the PTU embedding \mathbf{Z} is isometric to \mathbf{X} .*

Proof: When the sampled data lie on a linear subspace of dimension d , each orthonormal frame \mathbb{T}_i , computed using SVD as described in Sec. 3.5, forms a basis for \mathcal{R} . Moreover, every pair \mathbb{T}_i and \mathbb{T}_j are perfectly aligned by the discrete connection $\mathbf{R}_{j,i}$ (Eq. (2)). As a result, unfolding a polyline $(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$ reduces to rewriting it in the basis \mathbb{T}_{i_1} and our geodesic length estimation recovers the exact (Euclidean) distance between \mathbf{x}_{i_1} and \mathbf{x}_{i_m} , independent of the sampling irregularities or of the geodesic convexity of the domain. PTU thus becomes equivalent to classical MDS, producing a d -dimensional embedding \mathbf{Z} that is isometric to \mathbf{X} by construction. ■

Proposition 3: *Under mild assumptions on the regularity of a manifold \mathcal{M} and assuming that the input pointset \mathcal{S} samples \mathcal{M} finely enough with potential sampling voids over regions of small (sectional) curvature of the manifold, the PTU estimate $d_{\text{PTU}}(\mathbf{x}_i, \mathbf{x}_j)$ of the geodesic distance between points \mathbf{x}_i and \mathbf{x}_j based on a Dijkstra (shortest) polyline $(\mathbf{x}_i, \dots, \mathbf{x}_j)$ computed on the proximity graph G of \mathcal{S} approximates the real geodesic distance $d_g(\mathbf{x}_i, \mathbf{x}_j)$.*

Discussion: We quantify this proposition more rigorously in App. A by providing a concrete error bound (and its proof) between d_{PTU} and d_g . Our bound relies on three key components: 1) if the pointset \mathcal{S} is dense enough, then for appropriate choice of parameters k and K the discrete tangent spaces approximate their continuous equivalents, and the resulting discrete connection converges to the Levi-Civita connection in the sampling limit as proved in [Singer and Wu, 2012], where the authors used an equivalent notion of parallel transport to define and study a discrete connection Laplacian; 2) sampling voids in \mathcal{S} are allowed as long as the integral of the intrinsic sectional curvature of \mathcal{M} over the regions of the underlying manifold corresponding to these voids is small—in other words, we will assume that a Dijkstra polyline lies within a tubular neighborhood of its corresponding geodesic, for a tubular diameter less than $O(1/\sqrt{\kappa_s})$ where κ_s is the local maximum absolute value of the sectional curvature of \mathcal{M} ; and 3) for a dense enough sampling of \mathcal{M} , a straight \mathbb{R}^n vector between two nearby points \mathbf{x}_i and \mathbf{x}_j on \mathcal{M} has approximately the same length as the Cartan development of the geodesic between them on the tangent space $T_i\mathcal{M}$ at \mathbf{x}_i . Note that our statement involves no geodesic convexity requirement.

The benefit of using parallel transport over regular Dijkstra geodesic length approximation is thus clear: our construction eliminates spurious geodesic curvature that graph approximations inevitably suffer from, bringing significant improvement even for well sampled domains (see Fig. 5). It also allows for almost perfect recovery of pairwise geodesic distances for developable manifolds ($\kappa_s = 0$) with *arbitrary* topology in the sampling limit, even if the sampled manifold is not geodesically convex. When voids are present in pointsets that sample non-developable manifolds, PTU computes approximate geodesic distances without having to explicitly fill in the voids. Instead, it extracts geometric information from paths surrounding each hole to recover high accuracy geodesic estimates on the manifold, provided that the voids were not over regions with large curvature.

Proposition 4: *Just like Isomap, the complexity of PTU is $O(n^3)$.*

Proof: Our approach does not change the computational complexity of the various steps compared to Isomap: the proximity graph construction is still $O(n \log n)$, the construction of tangent planes takes $O(nDK^2)$, our Parallel Transport Unfolding using Dijkstra shortest paths is in $O(n^2 \log n)$ (the unfolding process part itself requires $O(n^2(Dd^2 + d^3))$ to perform matrix multiplications and SVDs), while the partial eigendecomposition of a dense matrix is expected to take $O(n^3)$ operations. ■

Additional control. Compared to Isomap, our parallel transport based approach has a few extra parameters that can be exploited to offer more control over the manifold learning process:

- Since we compute local tangent spaces at each input point, the neighborhood size K can be adapted (either globally or locally) to the input. While using $K = k$ is sufficient in most cases (see our examples), raising this value can help deal with very noisy inputs as mentioned in Sec. 3.2.
- Similarly, local PCA of these neighborhoods may not lead to the best estimations of tangent spaces in extreme cases: robust PCA, or even local averaging of the PCA estimates¹ can help manage large amounts of noise and outliers. Again, we did not have recourse to these variants to prove the robustness of PTU in its default form; but they can be easily incorporated in a practical implementation of our approach to add flexibility.

¹As a side note, having a discrete connection makes the local averaging of vectors or frames particularly simple, as neighboring values can be parallel transported to a common point, where a pointwise average is computed.

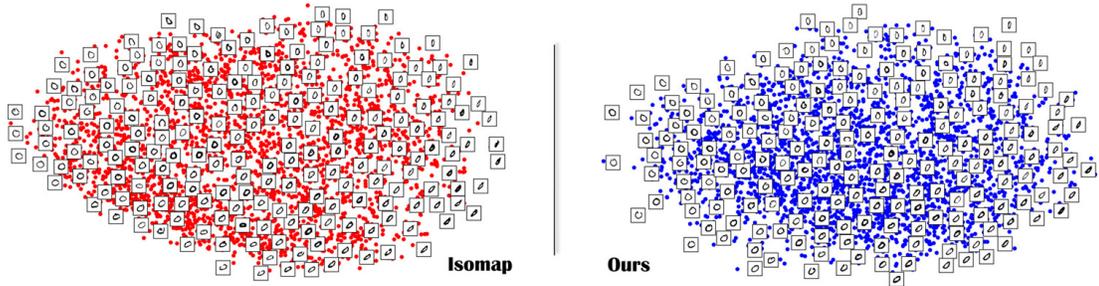


Figure 7: Digit zero. The 3000 images of handwritten zeros (with a resolution of 28x28 pixels) from the MNIST dataset are mapped quite similarly by both Isomap and PTU for the first two most significant coordinates (here, $d=4$ was used). A few images are displayed next to their corresponding 2D point for visualization purposes.

- Finally, our parallel transport estimation of geodesic lengths assumes the knowledge of the dimension d of the data, just like Isomap requires as well—and many approaches have been proposed to estimate this dimension directly from the data, see, e.g., [Pettis et al., 1979, Levina and Bickel, 2004]. Yet, this intrinsic dimension can, in fact, differ from the dimensionality of the visualization one wishes to produce. Fig. 7 illustrates this point: dimensionality reduction approaches applied on the MNIST dataset of digits often use a 2D illustration of their results for easy visual display; however, local analysis of the dimensionality of the zero digit image set indicates an intrinsic dimension of $d=4$, reflecting the variety of ways to pen a zero (slant, thickness, smoothness, ...). While this information cannot be exploited in Isomap if a 2D depiction is desired, PTU can exploit this estimate of d for its parallel transport procedure, but use only the first two eigenvectors of the Gram matrix—essentially showing a 2D projection of a 4D parameterization of the dataset. Fig. 11 shows another example of a 2D visualization of an intrinsic parameterization for $d=3$.

Related geometric methods. Note finally that two related works have proposed using parallel transport for data analysis, albeit for different purposes. Vector diffusion maps (VDM [Singer and Wu, 2012]) also exploit parallel transport, but focus instead on computing a low-dimensional embedding that preserves *vector diffusion distances* derived from a connection Laplacian, not true geodesics (see their Figs. 6.2 & 6.3). Parallel vector field embedding (PVF [Lin et al., 2013]) proposes a different discretization of the connection-Laplacian relying on an extrinsic definition of the covariant derivative. Coordinates of an embedding are constructed via Poisson solves, so that the coordinate lines in \mathbb{R}^d are mapped to parallel vector fields on the original data. Just like our approach, PVF perfectly recovers an isometric parametrization if the manifold is developable, since it is then isometric to a subset of \mathbb{R}^d . However, for non-developable manifolds, coordinate lines of the resulting parameterization do *not* correspond to geodesic curves on the original manifold. While they provide a valid definition of *an* embedding, PVF look for a low-dimensional “quasi-parallel” embedding. PTU, instead, targets the same goal as Isomap, i.e., a quasi-isometric mapping for arbitrary sampled manifolds.

3.8 Acceleration via landmarks

A particularly simple way to accelerate Isomap is through the use of “landmarks” [De Silva and Tenenbaum, 2002], a small fraction of samples of the original pointset: the MDS procedure is applied just to the landmarks to find their quasi-isometric embedding \mathbb{Z} in \mathbb{R}^d , before positioning all other points *relative to those landmarks* in linear time, significantly reducing computational complexity. This approach, however, often fails in practice: the sensitivity of the original Isomap method to poor sampling quality makes the low-dimensional embedding of a few landmarks very brittle.

Given its much greater robustness to irregular sampling, PTU is particularly amenable to this landmark-based acceleration without any other alteration than replacing distance estimation by our parallel transport approach. If ℓ landmarks are used, the amount of computations can decrease quite dramatically: the MDS complexity (which was the bottleneck for Isomap and PTU) changes from $O(n^3)$ to $O(\ell^3)$. Let us briefly discuss the implementation of such an L-PTU variant.

Setup. From the n original points, we extract ℓ landmarks with $\ell \ll n$, used as a coarse approximation of the input geometry. Landmark selection is not a sensitive part of our approach as long as the landmarks provide a good spatial coverage of the initial pointset. Note also that we compute our discrete connection for the entire pointset, since 1) it is not the computational bottleneck of the original PTU treatment, and 2) we need the distance from every data point to every landmark in order to compute the final embedding anyway. The

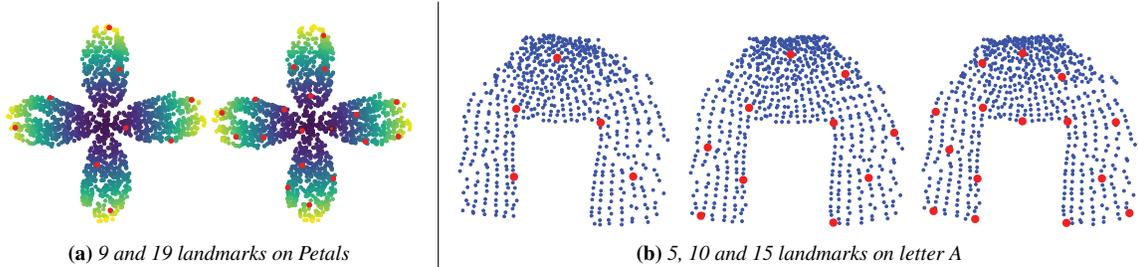


Figure 8: Landmark-PTU. Combining parallel transport unfolding with the landmark-based approach of L-Isomap [De Silva and Tenenbaum, 2002] drastically reduces computational times, with only small differences in the embedding: (top) for 9 (left) and 19 (right) landmarks, results on the Noisy Petals dataset are visually indistinguishable from the full treatment (see Fig. 6); (bottom) for the Letter A dataset, 5 landmarks (left) already capture the proper embedding, but 10 (middle) and 15 (right) landmarks result in a better approximation of the full treatment found in Fig. 4 (landmarks are in red).

choice of computing a “full resolution” connection guarantees accurate estimation of geodesic distances even if very few landmarks are used.

From landmark embedding to pointset embedding. After computing a low-dimensional embedding of the ℓ landmarks using PTU as described in Sec. 2.2 through

$$\mathbb{Z} = \sqrt{\Lambda_d} \mathbb{Q}_d^t,$$

(where now Λ_d and \mathbb{Q}_d are the d largest eigenvalues stored in a diagonal matrix and corresponding eigenvectors of the Gram matrix derived from the $\ell \times \ell$ matrix \mathbb{D} of the *squared* geodesic distances between landmarks only), the embedding of the remaining points can be performed using its pseudoinverse \mathbb{Z}^\dagger and the knowledge of geodesic distances between input points and landmarks. The pseudoinverse of \mathbb{Z} has an explicit formulation due to its basic form, requiring no additional computations:

$$\mathbb{Z}^\dagger = \mathbb{Q}_d \sqrt{\Lambda_d}^{-1}$$

Denote by $\bar{\mathbf{d}}$ the columnwise mean of \mathbb{D} (i.e., $\bar{\mathbf{d}}_p = \frac{1}{\ell} \sum_q \mathbb{D}_{pq}$). Then as described in L-Isomap [De Silva and Tenenbaum, 2002], the position of a non-landmark point $\mathbf{x}_i \in \mathcal{S}$ in the low-dimensional embedding can be directly computed using a vector \mathbf{d}_i of squared geodesic distances from \mathbf{x}_i to the ℓ landmarks through:

$$\mathbf{z}_i = (\mathbb{Z}^\dagger)^t (\bar{\mathbf{d}} - \mathbf{d}_i)$$

That is, knowing how distant a point is from the landmarks originally, we deduce its final position based on the low-dimensional embedding of the landmarks in $\mathcal{O}(\ell^2)$. Consequently, the proposed L-PTU procedure (consisting of computing all pairwise geodesic distances between ℓ landmarks and all input points, constructing embedding of the landmarks via MDS, and enriching it with non-landmarks through matrix-vector multiplications) brings the complexity down to $\mathcal{O}(\ell n \log(n) + \ell n(Dd^2 + d^3) + \ell^3 + \ell^2 n)$. As Fig. 8 demonstrates, this simple variant returns nearly the same embeddings as full PTU with as little as 0.1% to 0.5% of the points used as landmarks. More results using L-PTU are available in the Supplementary Material.

4 Results

We now discuss implementation details and provide a series of tests, on synthetic and real datasets, to compare our parallel transport approach to the original Isomap and other nonlinear dimensionality reduction methods.

4.1 Implementation details

Our implementation follows the steps described in the previous sections as summarized in Alg. 1. We use a modified Dijkstra’s algorithm to compute shortest paths and find geodesic distances concurrently, as detailed in Alg. 2 (the only new lines of code to handle connections are in blue). The final partial eigensolve was implemented using the Spectra C++ library [Qiu et al., 2015]. If the dimension D is high (i.e., larger than 100), vectorizing matrix-matrix and matrix-vector multiplications is crucial for efficiency.

4.2 Simulated Datasets

We first test the performance of PTU on artificial datasets (embedded in 3D or higher) to quantifiably evaluate its behavior.

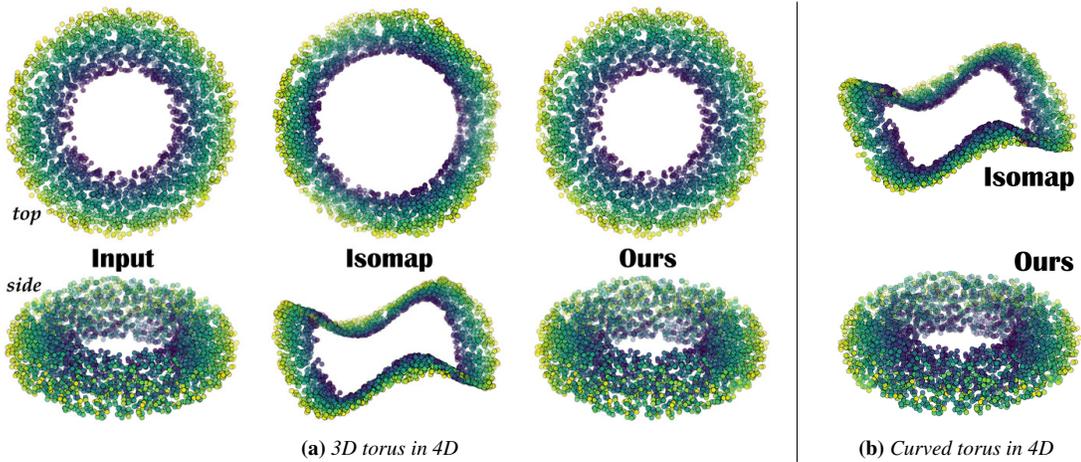


Figure 9: Tori embedded in 4D. (a) From a 3D pointset filling up a toroidal domain that we trivially embedded in 4D by adding an extra constant coordinate, the 3D embedding computed via Isomap leads to an unexpected global distortion due to the manifold being not convex. Instead, PTU is nearly perfect: Isomap has a normalized average relative error of 40.6%, while PTU is 0.06%. (b) A mildly curved torus in 4D is obtained by mapping a set of points $(x, y, z) \in T^2$ to $(x, y, z, (x^2 + y^2)/2)$. Now it has non-zero curvature as a manifold in 4D. 3D Isomap embedding still suffers from similar global distortion (left); PTU recovers the torus very well (right) despite the non-trivial curvature.

Linear Precision. A nonlinear extension to PCA should, at the very least, be linear-precise, i.e., input data lying on a flat d -manifold in \mathbb{R}^D should be isometrically mapped to \mathbb{R}^d . From an input pointset densely sampling a flat square shape embedded in 3D space, we compare the performance of Isomap and PTU (with $k = K = 10$) in Fig. 3: we visualize the error for each point based on the normalized distance between ground truth and its embedding location. As expected PTU recovers the data exactly (to numerical precision), while Isomap introduces distortions due to its use of graph-based distances. We also use a 4D pointset that randomly samples a 3D torus in Fig. 9a, with the embedding space being 3-dimensional. Unlike Isomap, which introduces severe distortion because of the non-convexity of the sampled domain, PTU ($k = K = 10$) perfectly reproduces the torus. This last example also proves that Isomap can significantly distort data (in unexpected ways) in higher dimensions.

Non-Geodesically-Convex Domains. An S-shaped manifold with a rectangular void (see Fig. 2) is well parameterized by PTU, while Isomap introduces spurious distortion associated with biased geodesic distance estimations around the void. Since this S-shaped manifold is isometrically developable, we can compare errors in parameterization on a per-point basis: Isomap reaches 15% of relative error, while PTU (still using $k = K = 10$) stays below 0.2%. The effect of non-geodesically-convex domains is even more pronounced in higher dimensions: Fig. 1 shows a 3D pointset sampling four petals from a surface of a sphere in 3D; this pointset is then lifted to $D = 100$ dimensions and rotated by a random orthogonal transformation. The resulting non-developable and non-convex 100D dataset is then embedded in 2D by Isomap, clearly demonstrating that graph-based distances bias results dramatically. Our algorithm, run with $k = K = 10$, correctly unfurls the petal-like set.

Non-developable manifolds. We also demonstrate our results on highly non-developable manifolds. Fig. 10 shows a dense, noise-free pointset that is sampling a 2D height field with two Gaussian bumps. The differences between Isomap and PTU ($k = 8$) are small with such a dense dataset. This example also demonstrates that the size of the geodesic neighborhood used for tangent space estimation does not dramatically affect the quality of results: PTU outputs for $K = 8$ and $K = 24$ are visually indistinguishable. We also test an input pointset sampling a slightly curved 3D torus in 4D: from the 3D torus in Fig. 9a, each point (x, y, z) is mapped in 4D to $(x, y, z, (x^2 + y^2)/2)$. While Isomap gets even more distorted, PTU ($k = K = 10$) still produces a toroidal 3D embedding, see Fig. 9b.

Geodesic distance estimation. We also compare the accuracy of Dijkstra, PTU, and SAKE-corrected [Budninskiy et al., 2017] estimations of geodesic distances, by applying these 3 methods to a low-density, regular sampling of a spherical cap in Fig. 5 so that geodesic distances are known analytically. PTU geodesics are over 120 times more accurate than Dijkstra’s, and more than 20 times better than the local geodesic correction method used in SAKE [Budninskiy et al., 2017] (where the whole cap is treated as a single neighborhood).

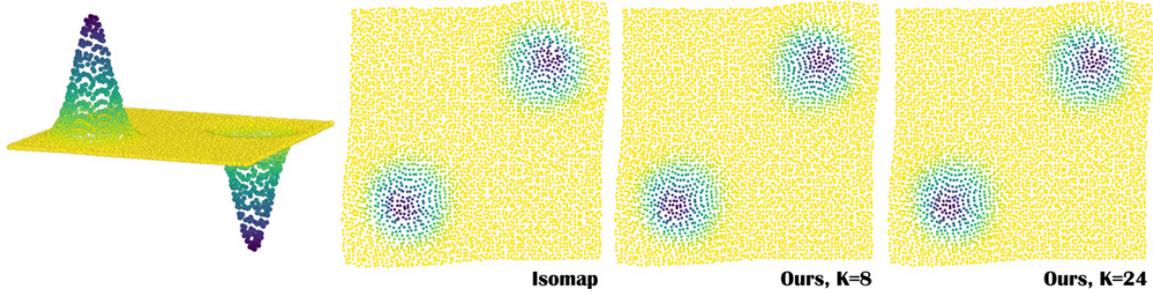


Figure 10: Gaussian landscape. For a noise-free dense sampling of a non-developable height field (left), Isomap and PTU return comparable quasi-isometric 2D parameterizations. Using $K=8$ or $K=24$ neighbors (right) does not visually affect the result of the PTU embedding.

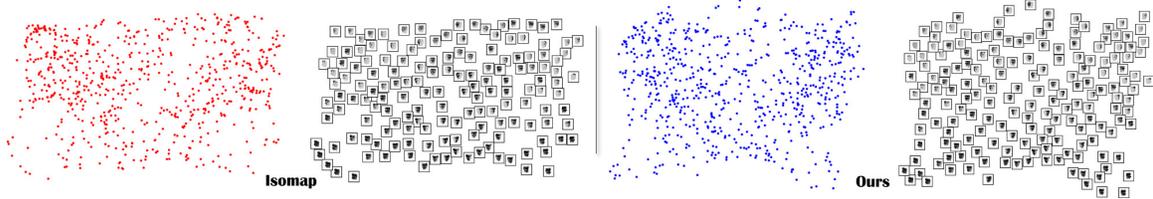


Figure 11: Faces Dataset. For the classical Faces dataset (known to be, by construction, of intrinsic dimension 3), the first two coordinates of both Isomap and PTU are quite similar. Left: 2D parameterization; right: a few of the face images are shown at their actual positions to better understand the parameterization.

Sensitivity to Irregular Sampling. For a highly irregular sampling of a simple S-shaped dataset, the robustness of PTU is shown in Fig. 13. Observe that Isomap completely collapses a section of the data in the right bottom corner, while PTU correctly captures the features of the input data even for a non-adapted choice of neighborhood parameters ($k=K=10$), introducing only small distortion throughout the domain. These two behaviors are visualized by the error plots using a linear color ramp from blue (0% error) to red (10% error).

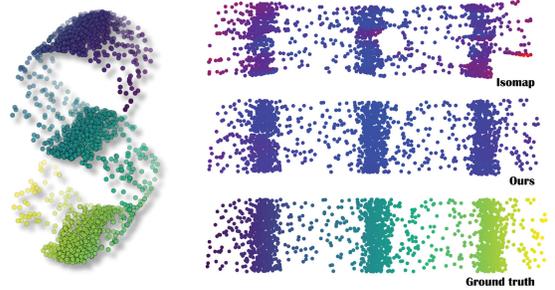


Figure 13: Varying Density S. For a widely varying density of points (left), Isomap (top right) introduces large spurious distortions, unlike PTU (middle right).

Sensitivity to Noise. Figs. 12a & 12b verify that PTU (with $k = 10$) is as resilient to strong noise and outliers as Isomap even without locally adapting neighborhood sizes ($K = 25$ was used in the Gaussian noise case, and $K = 10$ in the case of sparse noise). Local methods such as LLE [Roweis and Saul, 2000], Hessian LLE [Donoho and Grimes, 2003], LTSA [Zhang and Zha, 2004], MLE [Zhang and Wang, 2007] or SAKE [Budninskiy et al., 2017] (with $k = 10$ for fairness of comparison) fail to unwrap noisy datasets (see Fig. 14) as they do not exploit the geodesic distance estimates between pairs of points that are far apart. Fig. 6 shows that adding noise to the petals dataset from Fig. 1 does not alter the result of our approach significantly (we used $k = 10, K = 30$); yet, Isomap remain unable to unfurl the data properly, and on this seemingly simple data, all local methods fail, at times spectacularly. Please refer to the Supplemental Material for more systematic testing of the effect of noise levels on both local and global methods; as expected, PTU is systematically as good or better than all other methods.

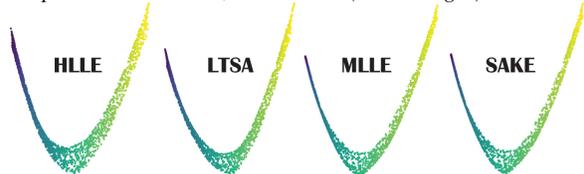


Figure 14: Failure of Local Methods. For the very noisy Swiss Roll example of Fig. 12b, none of the local manifold learning methods returns a decent parameterization, as no large intrinsic distances are exploited.

Timing. A performance analysis of our algorithm shows perfect agreement with the expected time complexity orders: the eigensolve dominates the computational time and scales as $O(n^3)$, parallel transport Dijkstra scales as $O(n^2 \log n)$, and graph construction as $O(n \log n)$, while tangent estimates are linear in n . Examples of timings on an Intel i7 2GHz, 8 GB RAM laptop are 6.4s for the petals parametrization in Fig. 1, 10.9s for the irregular sampling of S-shaped manifold in Fig. 13, and 9.8s for the noisy Swiss roll in Fig. 12b (with $n = 2000$). Note that these timings are for the full-blown MDS procedure, with roughly half of time spent on

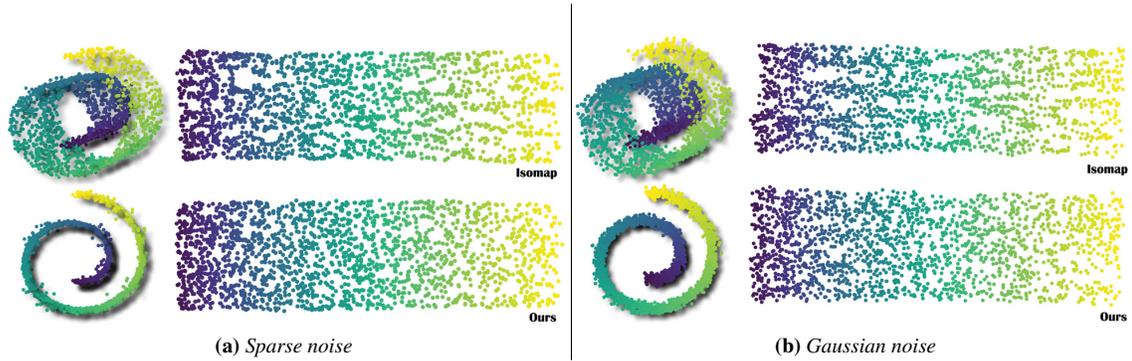


Figure 12: Noisy Swiss rolls. (a) If one adds noise to the Swiss Roll dataset in the normal direction by displacing 10% of points with a uniform distribution of amplitude equal to 8% of the max bounding box size and adding a Gaussian noise to the other points with standard deviation equal to 0.4% of the bounding box (see two views of the 3D dataset on the left), Isomap accentuates a few low sampled regions compared to PTU. (b) For a strong Gaussian noise (standard deviation equal to 2% of the max bounding box size), Isomap suffers from clear visual artifacts while PTU returns a good parameterization without the need for robust estimations.

the other steps. Using the L-PTU variant with 0.05%-0.1% of the points as landmarks improved the efficiency of the MDS step by 500 to 5000 times on tested datasets, with virtually no visual difference compared to the full treatment, see Fig. 8 and App. B.

4.3 High-dimensional datasets

We also present results on a number of real and/or high-dimensional datasets. While no ground truth is available for these examples, they allow us to compare PTU and Isomap on inputs that may not even satisfy the (single-chart) manifold assumption.

Faces Dataset. On this classic set of 698 images of 64x64 pixels, PTU and Isomap recover the same two characteristic features of the data: Fig. 11 shows that both arrange the images based on the azimuth and elevation of the camera, with a fairly similar global structure (2D visualization, $d=3$, $k=6$, and $K=18$).

Digit zero. When applied to 3000 digit zero images (28x28 pixels) from the MNIST dataset, both PTU and Isomap create a parameterization of the different ways people write a zero, separating left-leaning from right-leaning and circular from oval zeros as shown in Fig. 7 (2D visualization, $d=4$, $k=K=10$).

Knights. We tried our approach on the dataset from [Budninskiy et al., 2017] of reflectance fields captured using the Light Stage apparatus [LightScape, 2016]. A static character (in a knight costume) was captured in 608×456 RGB images under 221 individual lighting directions covering a large sphere of illumination. From this pointset in \mathbb{R}^{831744} , Isomap and PTU learn a flat 2D manifold that best fits this high-dimensional dataset. The result for $d=2$ and $k=K=8$, shown in Fig. 15, recovers positions related to light angles without any knowledge of the setup; the highly distorted results of local approaches on this dataset can be found in [Budninskiy et al., 2017]. The black background of each image was removed for clarity.

Letter A. From an input set of 888 images with 120x120 pixels of a rotated and resized letter ‘A’, the structure of this intrinsically 2-dimensional manifold in 14400 dimensional space revealed in PTU ($k=8$, $K=32$) and Isomap embeddings in Fig. 4 (left) is very intuitive. However, when a section of the data is removed (right), Isomap suffers from its characteristic global distortion, caused by the presence of the hole; instead, the structure of PTU embedding remains basically the same.

Algorithm 1 Dimensionality reduction via parallel unfolding

Input: Pointset $\mathcal{S} = \{\mathbf{x}_i \in \mathbb{R}^D\}_{i=1..n}$

- 1: Construct proximity graph G of \mathcal{S} , and pairwise Dijkstra’s shortest paths (Sec. 3.1)
- 2: Compute tangent frames $\{\mathbb{T}_i\}_{i=1..n}$ (Sec. 3.2)
- 3: Compute geodesic distances \mathbf{D} via parallel transport (Sec. 3.5)
- 4: Perform MDS on \mathbf{D} to obtain \mathbf{Z} (Sec. 2.2)

Output: Low-dimensional embedding $\mathbf{Z} = \{\mathbf{z}_i \in \mathbb{R}^d\}_{i=1..n}$

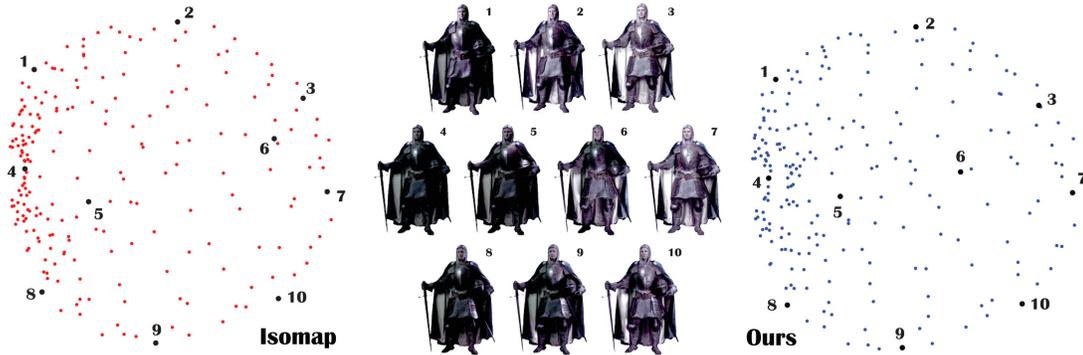


Figure 15: Knights. From 221 RGB images of 608×456 pixels (10 examples shown on the right) of an actor captured in full costume from different lighting directions, a 2D embedding is computed solely based on local pixel differences using Isomap and PTU (left). Both methods find a parameterization of the images corresponding to lighting direction and intensity (the knight images correspond to black dots).

5 Conclusions

The use of parallel transport on high-dimensional datasets remedies an important limitation of the Isomap approach: unfolding paths between pairs of points based on the Levi-Civita connection significantly improves the estimation of geodesic distances and removes the restriction for geodesic convexity of the input data. We demonstrated on a series of examples that our approach does indeed recover similar unfolding to Isomap for geodesically convex inputs of low- and high-dimensional data, but neither overestimates geodesic distances if large voids are present, nor suffers from large deformation in the case of non-geodesically convex inputs. This property is particularly crucial to the success of our landmark-based variant, L-PTU, which can efficiently approximate the low-dimensional embedding of a large datasets in $O(n^2 \log n)$, even in the presence of noise.

Our work offers multiple opportunities for future research. First, connections and parallel transport are rarely used in data analysis (a notable exception being in unsupervised domain adaptation [Shrivastava et al., 2014]), so other applications of this common concept of geometry processing may be found valuable as well, including the use of Levi-Civita connections for arbitrary, non-Euclidean metrics. Additionally, we have assumed that the input data lie on a manifold of a given dimension d ; but real data is sometimes better captured by a CW complex, i.e., a set made out of regions of different dimensionality glued together. Being able to handle geodesic distances in this case would be a nice and useful extension. Exploiting current algorithmic work in approximating shortest paths on graphs could also be valuable to replace Dijkstra’s algorithm; maybe the geodesics in heat approach [Crane et al., 2013] could, similarly, be extended to high-dimensional datasets for the same purpose. Finally, manifold learning concepts can also be extended to achieve clustering and classification of unlabeled data; using our improved geodesic distances (and potentially, its extension to CW complexes) may help provide better algorithms for these tasks, pointing to the relevance of connections in other fields of applications.

Acknowledgments

We gratefully acknowledge the help of Mark Gillespie and Prof. Venkat Chandrasekaran who provided early feedback on our work. This work was partially supported by a generous grant from Pixar Animation Studios.

A Formal Statement of Proposition 3

Before formulating a precise error bound on our geodesic approximation to quantify Proposition 3, we first review existing results and state a few reasonable assumptions on the sampling \mathcal{S} of the manifold \mathcal{M} .

A similar notion of discrete parallel transport was considered by the authors of [Singer and Wu, 2012] in order to define and prove convergence of a connection Laplacian operator. We will make use of one of the theorems they proved towards their goal, with a proof (omitted here) relying on geometric properties of parallel transport and probabilistic guarantees (obtained through Bernstein’s inequality) on the quality of tangent bases approximation. Following their notations, let $\varphi: \mathcal{M} \hookrightarrow \mathbb{R}^D$ be the embedding of a smooth and compact Riemannian d -manifold \mathcal{M} with its metric induced from the embedding space \mathbb{R}^D . Points from \mathcal{S} are considered to be sampled from \mathcal{M} according to a probability density function $p \in \mathcal{C}^3(\mathcal{M})$. Denote the tangent bundle of \mathcal{M} by $T\mathcal{M}$, the tangent space at point \mathbf{x}_i by $T_i\mathcal{M}$, the differential of the embedding at \mathbf{x}_i

by $d\varphi_i : T_i\mathcal{M} \mapsto \mathbb{R}^D$ and the parallel transport operator from \mathbf{x}_j to \mathbf{x}_i along a geodesic connecting them by $\mathbf{P}_{i,j} : T_j\mathcal{M} \mapsto T_i\mathcal{M}$. Let ε be the maximum radius of the geodesic K -neighborhoods used to construct our approximate tangent bases $\{\mathbb{T}_i\}_{i=1}^n$ (see Eq. (1) and its description) and finally, let $\varepsilon_g > 0$.

Theorem (from [Singer and Wu, 2012]). Assume $\varepsilon = O(n^{-\frac{1}{d+2}})$ and consider two points \mathbf{x}_i and \mathbf{x}_j of \mathcal{M} , such that the geodesic distance between them is $O(\varepsilon_g)$. Then for any vector $\mathbf{u}_j \in T_j\mathcal{M}$, with high probability:

$$\mathbf{R}_{i,j}\mathbb{T}_j^t d\varphi_j[\mathbf{u}_j] = \mathbb{T}_i^t d\varphi_i[\mathbf{P}_{i,j}\mathbf{u}_j] + O(\gamma_{i,j}) \quad (6)$$

where $\gamma_{i,j} = \varepsilon^3 + \varepsilon_g^3$ if \mathbf{x}_i and \mathbf{x}_j are at least ε -away from the boundary of \mathcal{M} and $\gamma_{i,j} = \varepsilon + \varepsilon_g^3$ otherwise.

In addition to this theorem, we will use two more assumptions:

- I. Let $\mathbf{c}_p = (\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$ be the shortest polyline connecting \mathbf{x}_{i_1} and \mathbf{x}_{i_m} in the proximity graph G . First, we assume that the polyline is included in an ε_d -thickening of the real geodesic \mathbf{c}_g connecting the same endpoints, where ε_d is a positive constant such that $\varepsilon_d^2 \kappa_s \ll 1$, with κ_s denoting the maximum absolute value of intrinsic sectional curvature of \mathcal{M} . Note that this condition is less stringent than the one used to prove convergence of Dijkstra polylines to geodesic curves [Tenenbaum et al., 2000]: sampling voids of maximum geodesic diameter ε_d are allowed in the input pointset \mathcal{S} .
- II. Let \mathbf{g}_{i_s} be a tangent vector in $T_{i_s}\mathcal{M}$ that connects the endpoints of the Cartan development of the geodesic curve between \mathbf{x}_{i_s} and $\mathbf{x}_{i_{s+1}}$ onto $T_{i_s}\mathcal{M}$. Denoting $\mathbf{e}_{i_s} = (\mathbf{x}_{i_{s+1}} - \mathbf{x}_{i_s})$, we will also assume:

$$\mathbb{T}_{i_s}^t d\varphi_{i_s} \mathbf{g}_{i_s} = \mathbb{T}_{i_s}^t \mathbf{e}_{i_s} + O(\gamma_{i_s, i_{s+1}}) \quad (7)$$

This condition links local curvature and sampling density: the projection of \mathbf{e}_{i_s} onto the approximate tangent basis \mathbb{T}_{i_s} must be close to the unwrapped geodesic in the same basis. Note that the length rescaling step (Eq. (4)) can help in practice to tighten the error bound $O(\gamma_{i_s, i_{s+1}})$.

Proposition 3 revisited: Under assumptions I, II and the assumptions of the theorem, the PTU estimate d_{PTU} of the geodesic distance between points \mathbf{x}_{i_1} and \mathbf{x}_{i_m} , based on a Dijkstra shortest polyline $\mathbf{c}_p = (\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$ on G , provides an approximation of the length d_g of the geodesic curve \mathbf{c}_g with the same endpoints in the following sense:

$$d_{PTU}(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) = d_g(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) + O(\delta),$$

where $O(\delta)$ is between $O(m^2(\varepsilon + \varepsilon_g^3 + \varepsilon_g \varepsilon_d^2 \kappa_s))$ and $O(m^2(\varepsilon^3 + \varepsilon_g^3 + \varepsilon_g \varepsilon_d^2 \kappa_s))$ depending on how many polyline segments are close to the manifold boundary.

Proof: First, observe that condition (7) implies:

$$\mathbf{v}_{i_s} = \left(\prod_{j=1}^{r-1} \mathbf{R}_{i_j, i_{j+1}} \right) \left[\mathbb{T}_{i_r}^t \mathbf{e}_{i_s} \right] = \left(\prod_{j=1}^{r-1} \mathbf{R}_{i_j, i_{j+1}} \right) \left[\mathbb{T}_{i_r}^t \mathbf{g}_{i_s} \right] + O(\varepsilon_g^3)$$

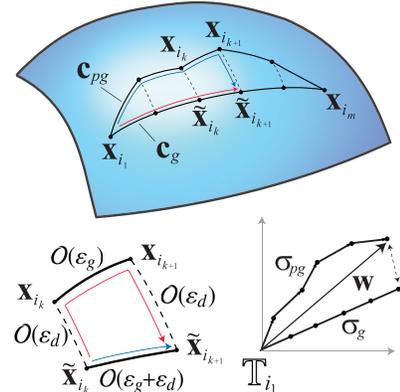
Using Eq. (6) repeatedly, we obtain that \mathbf{v}_{i_s} is approximately equal to \mathbf{g}_{i_s} parallel-transported to \mathbb{T}_{i_1} along a piecewise-geodesic curve \mathbf{c}_{pg} passing through the points $(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$, i.e.:

$$\mathbf{v}_{i_s} = \mathbb{T}_{i_1}^t d\varphi_{i_1} \left(\prod_{j=1}^{s-1} \mathbf{P}_{i_j, i_{j+1}} \right) \left[\mathbf{g}_{i_s} \right] + O \left(\sum_{j=1}^{s-1} \gamma_{i_j, i_{j+1}} \right) + O(\varepsilon_g^3)$$

Denoting the first term of the right-hand side by \mathbf{w}_{i_s} and calling $\mathbf{w} = \sum_{s=1}^m \mathbf{w}_{i_s}$, we obtain:

$$\sum_{s=1}^m \mathbf{v}_{i_s} = \sum_{s=1}^m \mathbf{w}_{i_s} + O(\gamma), \text{ i.e., } \mathbf{v} = \mathbf{w} + O(\gamma)$$

where $O(\gamma)$ takes value between $O(m^2(\varepsilon + \varepsilon_g^3))$ and $O(m^2(\varepsilon^3 + \varepsilon_g^3))$ depending on how many polyline segments are close to the boundary; note that \mathbf{w} is the vector connecting the endpoints of Cartan's development $\sigma_{pg} \in T_{i_1}\mathcal{M}$ of the piecewise-geodesic curve \mathbf{c}_{pg} interpolating the polyline vertices (see inset, bottom right). To show that the curve σ_{pg} stays close to Cartan's development σ_g of the geodesic \mathbf{c}_g , recall that for a patch of \mathcal{M} with diameter $O(\varepsilon_d)$, the change in direction of a vector parallel-transported along different paths within the patch is bounded by $O(\kappa_s \varepsilon_d^2)$ (this result relies on a bounded curvature transformation, which in turn follows from our bound on absolute value of sectional curvature from Assumption I). Using Assumption II, we can cut the geodesic \mathbf{c}_g connecting \mathbf{x}_{i_1} and \mathbf{x}_{i_m} into curved segments with vertices $\tilde{\mathbf{x}}_k$, such that $d_g(\tilde{\mathbf{x}}_k, \mathbf{x}_{i_k}) = O(\varepsilon_d)$ for



$k = 2, \dots, m-1$ (we set $\tilde{\mathbf{x}}_{i_1} = \mathbf{x}_{i_1}$ and $\tilde{\mathbf{x}}_{i_m} = \mathbf{x}_{i_m}$ for notational consistency). Because the geodesic length of $\mathbf{x}_{i_{k+1}}\mathbf{x}_{i_k}$ is $O(\varepsilon_g)$, the diameter of a patch enclosed by four geodesic segments $\mathbf{x}_{i_k}\mathbf{x}_{i_{k+1}}$, $\tilde{\mathbf{x}}_{i_{k+1}}\tilde{\mathbf{x}}_{i_k}$, $\tilde{\mathbf{x}}_{i_k}\mathbf{x}_{i_k}$ and $\mathbf{x}_{i_{k+1}}\tilde{\mathbf{x}}_{i_{k+1}}$ is $O(\varepsilon_g + \varepsilon_d) = O(\varepsilon_d)$ assuming $\varepsilon_g \lesssim \varepsilon_d$ (see inset). Thus the result of parallel transporting a vector along 3 geodesic segments $\tilde{\mathbf{x}}_{i_k}\mathbf{x}_{i_k}$, $\mathbf{x}_{i_k}\mathbf{x}_{i_{k+1}}$ and $\mathbf{x}_{i_k}\tilde{\mathbf{x}}_{i_k}$ deviates in direction from parallel transporting the same vector along geodesic segment $\tilde{\mathbf{x}}_{i_k}\tilde{\mathbf{x}}_{i_{k+1}}$ by $O(\kappa_s \varepsilon_d^2)$. Summing up contributions from all the patches for $k = 1, \dots, \ell$, and taking into account cancellations from segments $\mathbf{x}_{i_k}\tilde{\mathbf{x}}_{i_k}$ and $\tilde{\mathbf{x}}_{i_k}\mathbf{x}_{i_k}$ for $k = 2, \dots, \ell-1$, the directional error in parallel transporting a vector along the piecewise geodesic curve $\mathbf{x}_{i_1}\mathbf{x}_{i_2} \dots \mathbf{x}_{i_\ell}\tilde{\mathbf{x}}_{i_\ell}$ compared to its transport along the true geodesic segment $\mathbf{x}_{i_1}\tilde{\mathbf{x}}_{i_\ell}$ becomes $O(\ell \kappa_s \varepsilon_d^2)$. As a result, performing Cartan’s development patch by patch, the final discrepancy between the endpoint positions of σ_{pg} and σ_g incurred by developing them using parallel transport along piecewise-geodesic curve \mathbf{c}_{pg} vs. the true geodesic \mathbf{c}_g is $O(m^2 \kappa_s \varepsilon_g \varepsilon_d^2)$, as it combines the cumulative directional errors of parallel transported tangent vectors and the lengths of the corresponding segments.

Given that Cartan’s development of a geodesic curve is a straight segment, we conclude that $\|\mathbf{w}\|_2 = d_g(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) + O(m^2 \kappa_s \varepsilon_g \varepsilon_d^2)$. By construction, we have $d_{PTU}(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) = \|\mathbf{v}\|_2$, implying that:

$$d_{PTU}(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) = d_g(\mathbf{x}_{i_1}, \mathbf{x}_{i_m}) + O(\delta),$$

where $O(\delta) = O(m^2 \varepsilon_g \varepsilon_d^2 \kappa_s + \gamma)$. Note that our discrete unfolding of *any* polyline converges to the corresponding Cartan’s development; however, in general it has to be a (nearly-)shortest polyline for its extremities to be at a distance approximating the proper geodesic length, as geodesic curves are only locally (and not globally) shortest. ■

Finally, we note that while our assumptions are weaker than the ones used to prove convergence of Dijkstra polylines to geodesic curves [Tenenbaum et al., 2000], the error for graph-based distance approximations has linear dependence on the number of polyline segments, while our bound depends quadratically on the number of segments. However, in practice our parallel transport based method consistently outperforms Dijkstra path approximation (see Fig. 5), potentially pointing to the existence of tighter bounds.

B Additional Results

This last appendix contains additional results to provide further tests of the Parallel Transport Unfolding (PTU) approach and comparisons to other existing Manifold Learning algorithms.

B.1 PTU vs. Local Methods

One of the key advantages of Parallel Transport Unfolding is its resilience to noise: like Isomap [Tenenbaum et al., 2000], PTU uses *all* geodesic distances to embed a dataset into a low-dimensional space, which allows for a much increased robustness to noise and outliers. In this section, we provide further numerical tests to confirm this statement.

Petals dataset. First off, the noiseless Petals dataset makes clear that Isomap fails due to the obvious geodesic non-convexity while PTU has no problem finding the proper four petals as demonstrated in Fig. 1. In this noiseless case, it turns out that most local methods, such as Modified LLE [Zhang and Wang, 2007], Hessian LLE [Donoho and Grimes, 2003], and SAKE [Budninskiy et al., 2017], do also remarkably well (see Fig. 16), as convexity (or lack thereof) plays basically no role in their embeddings.



Figure 16: Local Methods for Noiseless Petals. From a 3D sampling of 4-petal shaped portion of a sphere (see Fig. 1), local methods such as Modified LLE [Zhang and Wang, 2007], Hessian LLE [Donoho and Grimes, 2003], or SAKE [Budninskiy et al., 2017] have no issue with the non-convexity of the intrinsic geometry (unlike Isomap), and give results nearly equivalent to PTU. A notable exception is LLE [Roweis and Saul, 2000], which returns a near degenerate solution.

However, if a bit of noise is added to this example, local methods fail (at times spectacularly): for a moderate Gaussian noise with a standard variation of 3% of the radius of the sphere on which the petals

lie, Fig. 6 shows that local methods all fail while PTU keeps a very similar embedding since it relies on all pairwise geodesic distances—Isomap too to a certain extent, even if it is clearly more deformed than in the noiseless case.

Study of noise effects on local and global methods. In order to better demonstrate the robustness of global methods to noise, we use the simple (and very widely used) Swiss Roll dataset, and use both local and global methods on this dataset with an increasing amount of Gaussian noise along the normal of this roll (varying the standard deviation from 0 to 2.8% of the bounding box size). The number of neighbors is set to 10 for all methods to offer a fair comparison. As Fig. 17 clearly shows, the best local method (SAKE on this example) starts failing at half of the maximum standard deviation that global methods can handle. When the standard deviation of the noise reaches 2.8%, even global methods fail as the proximity graph starts having numerous connections across branches: pairwise geodesic distances will have too many incorrect values to be able to recover a decent embedding. Note that at this level of noise, the dataset is far from the manifold assumption we are making about input data.

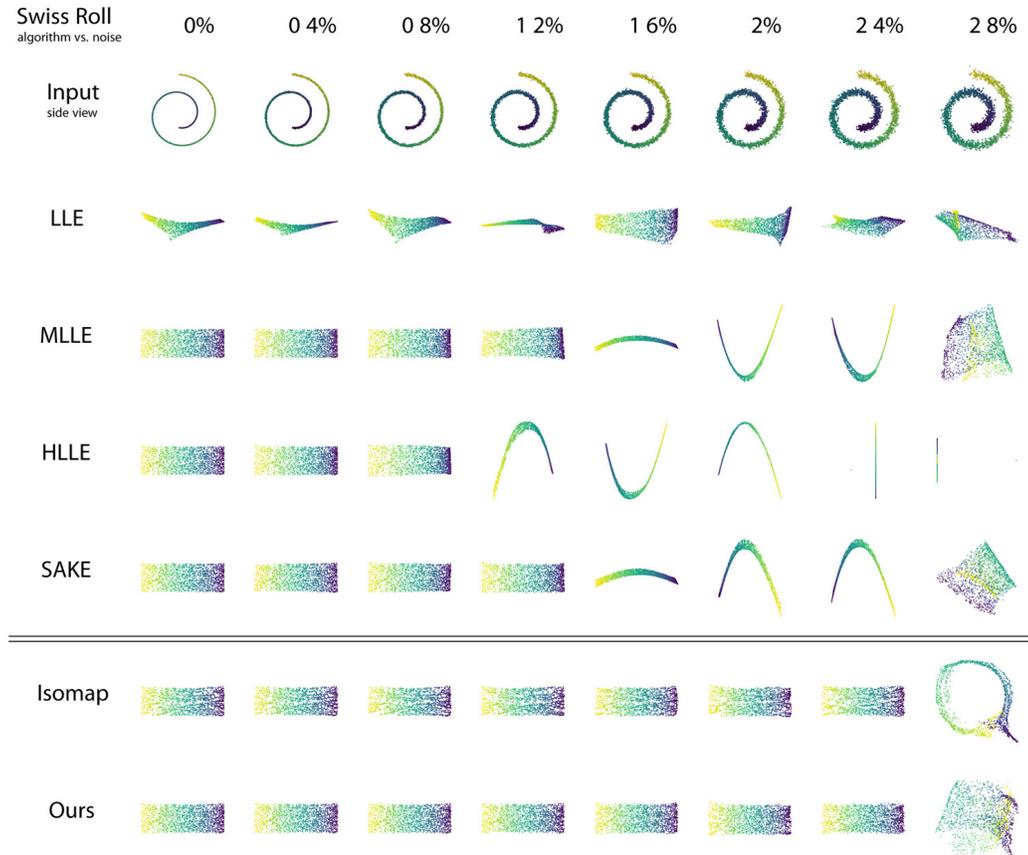


Figure 17: Effects of noise on local and global methods. Using the Swiss Roll dataset, Gaussian noise with standard deviation given as a percentage of the bounding box of the original noiseless swiss roll is added along the normal. We use the same number of neighbors (10) for local methods to provide a fair comparison (it prevents shortcutting as much as possible; using larger values would make the local methods fail earlier). Local methods all failed around $\sigma = 1.3\%$, while global methods (Isomap and our approach) fare well until 2.7%. At 2.8%, the neighbors of a datapoint may belong to several different branches of the roll, which makes it impossible even for global methods to handle.

B.2 Landmark-PTU

Fig. 18 provides more results for L-PTU: in order to complement Fig. 8, we also provide the results for landmarks on the *noiseless* Petals dataset, as well as on the highly-irregular S-shaped dataset from Fig.13. Here again, less than 1% of landmarks is enough to capture the shape almost perfectly.

References

[Assa et al., 2005] Assa, J., Caspi, Y., and Cohen-Or, D. (2005). Action synopsis: Pose selection and illustration. *ACM Trans. Graph.*, 24(3):667–676.

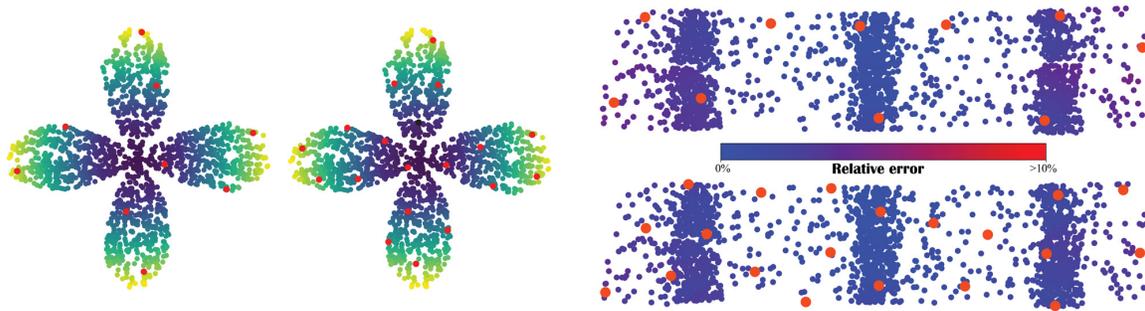


Figure 18: Landmark-PTU. Landmarks are colored red for clarity. (Left) the use of 9 landmarks (left) or 19 landmarks (right) is enough to reconstruct the petals in the noiseless example of Fig. 1. (Right) using 10 landmarks (top) vs. 20 landmarks (bottom) is visually very similar on this 2000-point datasets, although one can notice a slight distortion as indicated by the color of the mapped points using the color ramp based on distortion error compared to the expected perfect embedding.

- [Belkin and Niyogi, 2001] Belkin, M. and Niyogi, P. (2001). Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Neural Information Processing Systems*, pages 585–591.
- [Borg and Groenen, 2005] Borg, I. and Groenen, P. J. F. (2005). *Modern multidimensional scaling: Theory and applications*. Springer, 3rd edition.
- [Brandes and Pich, 2007] Brandes, U. and Pich, C. (2007). Eigensolver methods for progressive multidimensional scaling of large data. In Kaufmann, M. and Wagner, D., editors, *International Symposium on Graph Drawing*, pages 42–53.
- [Brito et al., 1997] Brito, M. R., Chávez, E. L., Quiroz, A. J., and Yukich, J. E. (1997). Connectivity of the mutual k-nearest-neighbor graph in clustering and outlier detection. *Statistics & Probability Letters*, 35(1):33–42.
- [Budninskiy et al., 2017] Budninskiy, M., Liu, B., Tong, Y., and Desbrun, M. (2017). Spectral affine-kernel embeddings. *Computer Graphics Forum*, 36(5):117–129.
- [Campen et al., 2013] Campen, M., Heistermann, M., and Kobbelt, L. (2013). Practical anisotropic geodesy. In *Symposium on Geometry Processing*, pages 63–71.
- [Cayton and Dasgupta, 2006] Cayton, L. and Dasgupta, S. (2006). Robust Euclidean embedding. In *International Conference on Machine Learning*, pages 169–176.
- [Choi and Choi, 2007] Choi, H. and Choi, S. (2007). Robust Kernel Isomap. *Pattern Recognition*, 40(3):853–862.
- [Crane et al., 2010] Crane, K., Desbrun, M., and Schröder, P. (2010). Trivial connections on discrete surfaces. *Comp. Graph. Forum*, 29(5):1525–1533.
- [Crane et al., 2013] Crane, K., Weischedel, C., and Wardetzky, M. (2013). Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (TOG)*, 32(5):152.
- [Crippen and Havel, 1978] Crippen, G. M. and Havel, T. F. (1978). Stable calculation of coordinates from distance information. *Acta Crystallographica Section A*, 34(2):282–284.
- [Datar et al., 2004] Datar, M., Immorlica, N., Indyk, P., and Mirrokni, V. S. (2004). Locality-sensitive hashing scheme based on p-stable distributions. In *Symposium on Computational Geometry*, pages 253–262.
- [De Silva and Tenenbaum, 2002] De Silva, V. and Tenenbaum, J. B. (2002). Global versus local methods in nonlinear dimensionality reduction. In Becker, S., Thrun, S., and Obermayer, K., editors, *Neural Information Processing*, volume 15 of *Advances in Neural Information Processing Systems*, pages 705–712. MIT Press.
- [De Silva and Tenenbaum, 2004] De Silva, V. and Tenenbaum, J. B. (2004). Sparse multidimensional scaling using landmark points. Technical report, Stanford University.
- [Desbrun et al., 2003] Desbrun, M., Meyer, M., and Alliez, P. (2003). Intrinsic parameterizations of surface meshes. *Computer Graphics Forum*, 21(3):209–218.
- [Donoho and Grimes, 2003] Donoho, D. L. and Grimes, C. (2003). Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596.
- [Frankel, 2011] Frankel, T. (2011). *The Geometry of Physics: An Introduction*. Cambridge University Press.
- [Halko et al., 2011] Halko, N., Martinsson, P.-G., and Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288.
- [Ham et al., 2004] Ham, J., Lee, D. D., Mika, S., and Schölkopf, B. (2004). A kernel view of the dimensionality reduction of manifolds. In *International Conference on Machine Learning*, pages 369–376.

- [Harel and Koren, 2002] Harel, D. and Koren, Y. (2002). Graph drawing by high-dimensional embedding. In *International Symposium on Graph Drawing*, pages 207–219.
- [Hormann et al., 2008] Hormann, K., Polthier, K., and Sheffer, A. (2008). Mesh parameterization: Theory and practice. In *ACM SIGGRAPH ASIA Course Notes*, pages 12:1–12:87.
- [Kabsch, 1976] Kabsch, W. (1976). A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A*, 32(5):922–923.
- [Khoury et al., 2012] Khoury, M., Hu, Y., Krishnan, S., and Scheidegger, C. (2012). Drawing large graphs by low-rank stress majorization. *Computer Graphics Forum*, 31(3.1):975–984.
- [Kim, 2007] Kim, S.-H. (2007). Generation of expression space for realtime facial expression control of 3D avatar. In Gagalowicz, A. and Philips, W., editors, *Proceedings of MIRAGE*, pages 318–329.
- [Kovnatsky et al., 2016] Kovnatsky, A., Glashoff, K., and Bronstein, M. M. (2016). MADMM: A generic algorithm for non-smooth optimization on manifolds. In *European Conference on Computer Vision*, pages 680–696.
- [Levina and Bickel, 2004] Levina, E. and Bickel, P. J. (2004). Maximum likelihood estimation of intrinsic dimension. In *Neural Information Processing Systems*, pages 777–784.
- [LightScape, 2016] LightScape (2016). Light Stage Data Gallery, University of Southern California’s Institute for Creative Technologies; <http://gl.ict.usc.edu/Data/LightStage/>.
- [Lin et al., 2013] Lin, B., He, X., Zhang, C., and Ji, M. (2013). Parallel vector field embedding. *Journal of Machine Learning Research*, 14:2945–2977.
- [Liu et al., 2016] Liu, B., Tong, Y., Goes, F. D., and Desbrun, M. (2016). Discrete connection and covariant derivative for vector field analysis and design. *ACM Trans. Graph.*, 35(3):23:1–23:17.
- [Melvær and Reimers, 2012] Melvær, E. L. and Reimers, M. (2012). Geodesic polar coordinates on polygonal meshes. *Computer Graphics Forum*, 31(8):2423–2435.
- [Mullen et al., 2008] Mullen, P., Tong, Y., Alliez, P., and Desbrun, M. (2008). Spectral conformal parameterization. In *Symposium on Geometry Processing*, pages 1487–1494.
- [Nomizu, 1978] Nomizu, K. (1978). Kinematics and differential geometry of submanifolds. *Tohoku Math. J.*, 30(4):623–637.
- [Pettis et al., 1979] Pettis, K. W., Bailey, T. A., Jain, A. K., and Dubes, R. C. (1979). An intrinsic dimensionality estimator from near-neighbor information. *IEEE Trans. Pattern Anal. Mach. Intell*, 1(1):25–37.
- [Pless, 2003] Pless, R. (2003). Image spaces and video trajectories: using isomap to explore video sequences. In *IEEE International Conference on Computer Vision*, pages 1433–1440.
- [Qiu et al., 2015] Qiu, Y., Guennebaud, G., and Niesen, J. (2015). Spectra: C++ library for large scale eigenvalue problems. <http://yixuan.cos.name/spectra>.
- [Rosman et al., 2010] Rosman, G., Bronstein, M. M., Bronstein, A. M., and Kimmel, R. (2010). Nonlinear dimensionality reduction by topologically constrained isometric embedding. *International Journal of Computer Vision*, 89:56–68.
- [Roweis and Saul, 2000] Roweis, S. T. and Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326.
- [Schmidt, 2013] Schmidt, R. (2013). Stroke parameterization. *Computer Graphics Forum*, 32(2):255–263.
- [Schmidt et al., 2006] Schmidt, R., Grimm, C., and Wyvill, B. (2006). Interactive decal compositing with discrete exponential maps. *ACM Trans. Graph.*, 25(3):605–613.
- [Schmidt and Singh, 2009] Schmidt, R. and Singh, K. (2009). Approximate conformal parameterization of point-sampled surfaces. Report, Technical report CSRG-605, Computer Science Department, University of Toronto.
- [Sharpe, 1997] Sharpe, R. (1997). *Differential Geometry: Cartan’s Generalization of Klein’s Erlangen Program*, volume 166 of *Graduate Texts in Mathematics*. Springer-Verlag New York.
- [Shrivastava et al., 2014] Shrivastava, A., Shekhar, S., and Patel, V. M. (2014). Unsupervised domain adaptation using parallel transport on Grassmann manifold. In *IEEE Winter Conference on Applications of Computer Vision*, pages 277–284.
- [Singer and Wu, 2012] Singer, A. and Wu, H.-T. (2012). Vector diffusion maps and the connection Laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144.
- [Spivak, 1979] Spivak, M. (1979). *A Comprehensive Introduction to Differential Geometry*, volume 1. Publish of Perish, Inc., second edition.
- [Tenenbaum et al., 2000] Tenenbaum, J. B., Silva, V. d., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323.
- [Torgerson, 1965] Torgerson, W. S. (1965). Multidimensional scaling of similarity. *Psychometrika*, 30(4):379–393.

- [Van Der Maaten et al., 2009] Van Der Maaten, L., Postma, E., and Van den Herik, J. (2009). Dimensionality reduction: a comparative review. *J. Mach. Learn. Res.*, 10:66–71.
- [Vidal et al., 2016] Vidal, R., Ma, Y., and Sastry, S. S. (2016). *Generalized Principal Component Analysis*. Springer.
- [Wang et al., 2006] Wang, J., Tong, X., Lin, S., Pan, M., Wang, C., Bao, H., Guo, B., and Shum, H.-Y. (2006). Appearance manifolds for modeling time-variant appearance of materials. *ACM Trans. Graph.*, 25(3):754–761.
- [Weinberger et al., 2004] Weinberger, K. Q., Sha, F., and Saul, L. K. (2004). Learning a kernel matrix for nonlinear dimensionality reduction. In *International Conference on Machine Learning*, pages 106–114.
- [Zhang and Lerman, 2014] Zhang, T. and Lerman, G. (2014). A novel M-estimator for robust PCA. *J. Mach. Learn. Res.*, 15(1):749–808.
- [Zhang and Wang, 2007] Zhang, Z. and Wang, J. (2007). MLLE: Modified locally linear embedding using multiple weights. In *Neural Information Processing Systems*, pages 1593–1600.
- [Zhang and Zha, 2004] Zhang, Z. and Zha, H. (2004). Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM J. Sci. Comput.*, 26(1):313–338.

Algorithm 2 Parallel Transport Dijkstra

Input: Pointset \mathcal{S} , proximity graph G , tangent frames $\{\mathbb{T}_i\}_{i=1..n}$

```
1: Create minimum priority queue  $P$ 
2: for  $i \in [1, n]$  do
3:    $\mathbf{R}[i] \leftarrow$  identity matrix,  $\mathbf{v}[i] \leftarrow 0$ 
4:   for all  $\mathbf{x}_j$  adjacent to  $\mathbf{x}_i$  in  $G$  do
5:      $Pred[j] \leftarrow i$ 
6:      $dist[j] \leftarrow |\mathbf{x}_j - \mathbf{x}_i|$ 
7:      $P.push(dist[j], j)$ 
8:   while not  $P.isEmpty()$  do
9:      $\mathbf{x}_r \leftarrow P.pop\_min()$ 
10:     $q \leftarrow Pred[r]$ 
11:     $\mathbf{U}\mathbf{S}\mathbf{V}^t \leftarrow \text{SVD}(\mathbb{T}_q^t \mathbb{T}_r)$ 
12:     $\mathbf{R}[r] = \mathbf{R}[q] \cdot \mathbf{U}\mathbf{V}^t$ 
13:     $\mathbf{v}[r] = \mathbf{v}[q] + \mathbf{R}[q] \cdot \mathbb{T}_q^t(\mathbf{x}_r - \mathbf{x}_q)$  (Eq. (5))
14:     $geo\_dist[r] = |\mathbf{v}[r]|$ 
15:    for all  $\mathbf{x}_j$  adjacent to  $\mathbf{x}_r$  do
16:       $temp\_dist \leftarrow dist[r] + |\mathbf{x}_j - \mathbf{x}_r|$ 
17:      if  $temp\_dist < dist[j]$  then
18:         $dist[j] = temp\_dist$ 
19:         $Pred[j] = r$ 
20:         $P.update(dist[j], j)$  ▷  $update()$  inserts  $j \notin P$ 
21:     $\mathbf{D}[i, 1 : n] \leftarrow geo\_dist[1 : n]$ 
22: Symmetrize distance matrix  $\mathbf{D} \leftarrow (\mathbf{D} + \mathbf{D}')/2$ 
Output: Pairwise geodesic distance matrix  $\mathbf{D} \in \mathbb{R}^{n \times n}$ 
```
