Finding Minimal Parameterizations of Cylindrical Image Manifolds

Michael Dixon, Nathan Jacobs, and Robert Pless
Department of Computer Science and Engineering
Washington University, St. Louis, MO, 63130, USA
{msd2, jacobsn, pless}@cse.wustl.edu

Abstract

Manifold learning has become an important tool to characterize high-dimensional data that vary nonlinearly due to a few parameters. Applications to the analysis of medical imagery and human motion patterns have been successful despite the lack of effective tools to parameterize cyclic data sets. This paper offers an initial approach to this problem, and provides for a minimal parameterization of points that are drawn from cylindrical manifolds—data whose (unknown) generative model includes a cyclic and a non-cyclic parameter. Solving for this special case is important for a number of current, practical applications and provides a start toward a general approach to cyclic manifolds. We offer results on synthetic and real data sets and illustrate an application to de-noising cardiac ultrasound images.

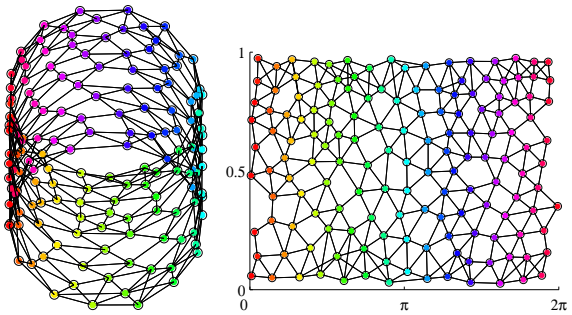

Figure 1. **Minimal cylindrical embedding.** Point sets that are isometric to a cylinder have one cyclic and one linear intrinsic dimension (θ and z). This paper offers an algorithm to embed the original 3D point set (left) into 2D (right) so that the x-axis corresponds to the cyclic dimension and the y-axis corresponds to the linear dimension.

1. Introduction

Many complex image sets can be characterized by a small number of simple variations. For such images, dimensionality reduction techniques are often able to find a low dimensional representation of the data, which can be beneficial to many vision tasks. However, finding cyclic parameterizations poses a problem for current algorithms. This paper addresses this problem, considering the case when one of the underlying variations is cyclic.

In the context of dimensionality reduction, an image is commonly represented as a point in a very high-dimensional space, where each coordinate corresponds to a single pixel value. If a set of images changes by some small number of continuous variations, the images will lie on a low dimensional manifold in this high dimensional space. A number of techniques exist for parameterizing such manifolds. If the underlying manifold is or approximates a linear subspace of the original image space, traditional linear techniques such as PCA [6] or ICA [5] yield effective parameterizations. In the more general case, where the underlying manifold is not linear, this problem is known as nonlinear dimensionality reduction (NLDR). In recent years, numerous NLDR techniques—including Isomap [12], LLE [10], Hessian LLE [2], SDE [14], and Laplacian eigenmaps [1] have been developed and used successfully for a wide range of applications.

However, these algorithms are not designed to handle manifolds with cyclic topologies. Because they embed into Euclidean space, they are unable to find parameterizations that correspond to the underlying degrees of freedom while simultaneously preserving the topology of the manifold. This is a limitation of traditional NLDR techniques, because many natural image sets have cyclic variations, such as a person walking, a heart beating, or an object rotating.

In this work we restrict our focus to cylindrical image manifolds—image sets which vary by one cyclic and one acyclic parameter (Figure 1). We present two novel contributions that exploit properties of cylindrical manifolds in order to find good parameterizations. The first is a direct method for creating a minimal parameterization for data sets sampled from manifolds isometric to right cylindrical manifolds. The second is an algorithm for transforming data with general cylindrical topology to a right cylinder. Combined, these methods allow us to apply our cylindrical

embedding to a broad class of real-world cylindrical manifolds. We characterize our performance versus the alternative approaches to parameterizing cylindrical data and offer parameterization results for a number of synthetic and natural data sets. Finally, we apply this technique to cardiac ultrasound images to reduce image noise.

1.1. Related Work

Recently, a number of algorithms have been developed that parameterize data according to its underlying manifold structure. In general, these algorithms attempt to find global coordinates in a low-dimensional space that preserve some aspect of the local relationships in the original high-dimensional space. Isomap substitutes approximate geodesic distances along the manifold for the Euclidean distance matrix used in multidimensional scaling. LLE attempts to represent the manifold locally by reconstructing each image as a weighted combination of its neighbors. SDE uses semidefinite programming to solve for a kernel matrix, which is then used to find an isometric embedding. These algorithms, among others, have been successfully applied to a number of common problems, including classification, segmentation, and tracking. However, these algorithms can only find Euclidean parameterizations and are not intended to provide cyclic parameterizations. Because cyclic phenomena are common in natural images, a number of methods for extending these algorithms to cyclic manifolds have been proposed.

One approach is to embed the points with an additional dimension; however, extra work is then required to extract meaningful parameters. For example, in [15] points from a spherical topology were embedded into three dimensions, where a sphere could then be fit to them. The sphere coordinates were then used to assign a 2D parameterization to the original data.

A large body of literature exists within the graphics community to parameterize points sets and meshes in 3D. These algorithms address different problems; the closest are algorithms for automatically classifying surface topology based on geodesic distances [3], provide conformal mappings of points with a known mesh structure in 3D [4], or generate a mesh of non-zero genus from a 3D point cloud [13]. None of these is immediately adaptable to separate the linear and cyclic parameters of a high-dimensional data set.

One approach designed for high-dimensional data sets is non-flat MDS [9], which explicitly considers the case of embedding cyclic data. This algorithm embeds cyclic and linear dimensions independently, but it requires user input to indicate which parameter causes the most variance, and it fails if the difference in the variances of the dimensions is small.

An alternative to these methods is to modify the distance graph to eliminate non-contractible cycles [8, 7]. This approach, discussed in more detail in Section 3.1, tears the manifold, changing its topology from cyclic to Euclidean. This allows the points to be embedded into Euclidean space; however, it discards all of the cyclic information contained in the graph in the process.

In this work, we consider the case of cylindrical manifolds and present an algorithm that provides an explicit parameterization of the cyclic and linear axes. In addition, we provide a method for correcting cylindrical distortions common in image manifolds.

2. Cylindrical Manifold Embedding

In this work, we restrict our attention to cylindrical manifolds, which can be parameterized by two intrinsic dimensions—one cyclic and one linear. We define the cylindrical embedding objective as follows:

> Given a set of points $X = \{X_1, \ldots, X_n\} \subseteq \mathcal{R}^D$ assume there exists a set of unknown intrinsic cyclic coordinates $L = \{L_1, \ldots, L_n\} \subseteq \mathcal{S}^1 \times \mathcal{R}^1$ and an image generating function $G : \mathcal{S}^1 \times \mathcal{R}^1 \rightarrow \mathcal{R}^D$ such that $X_i = G(L_i)$. The objective is to find an embedding function $E : \mathcal{R}^D \rightarrow \mathcal{S}^1 \times \mathcal{R}$ from the given points X to an embedding $Y = \{Y_1, \ldots, Y_n\} \subseteq \mathcal{S}^1 \times \mathcal{R}^1$ such that there is a one-to-one correspondence between dimensions of Y and dimensions of L. We assume $D \gg 2$, hence the points in X lie on an cylindrical manifold $\mathcal{M}$ in $\mathcal{R}^D$.

In the remainder of this paper we present a method for finding an embedding function if $\mathcal{M}$ is isometric to a right circular cylindrical manifold (which we define as *perfect*), and a set of procedures for transforming general cylindrical manifolds into perfect cylinders, which allows us to parameterize natural image sets.

3. Embedding Perfect Cylinders

Given a set of images drawn from $\mathcal{M}$, which is isometric to a perfect cylinder, our goal is to extract an explicit cylindrical parameterization. Our approach takes advantage of the cylindrical structure of the manifold to distinguish between the cyclic and linear axes.

3.1. Cycle Cutting

One approach to embedding cyclic data is to modify the neighborhood graph to convert its topology from cyclic to Euclidean. The resulting graph can then be embedded using a number of NLDR algorithms. In [8, 7], the authors present a procedure for cutting cyclic manifolds to remove non-contractible cycles.

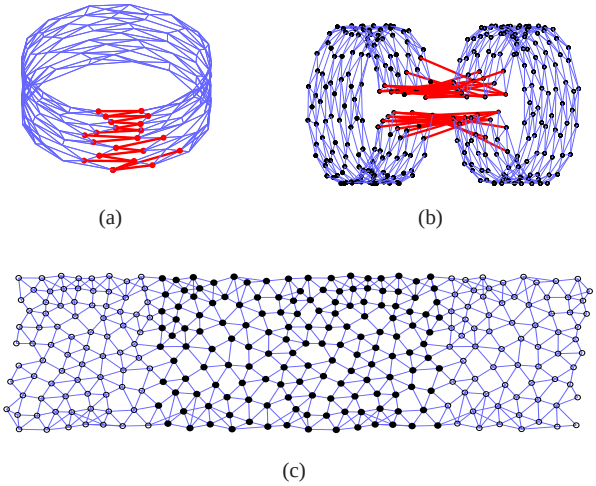

Figure 2. **Constructing the cut-clone-cut graph.** (a) The original graph with cut edges highlighted (red). (b) The combined graph connecting the original graph and its clone using cut edges. Running the cut algorithm on the combined graph creates the cut-combined graph. (c) The cut-clone-cut graph embedded into 2D (original points denoted with filled circles).

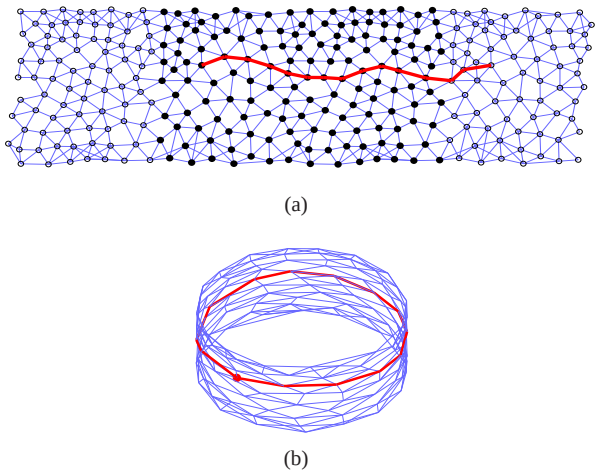

Figure 3. **Geodesic cycles.** (a) The geodesic cycle for a vertex is the shortest path from the point (filled circle) to its clone (open circle) in the cut-clone-cut graph. (b) This corresponds to the shortest non-contractible path around the cylinder from a given vertex.

This procedure, which we call the *cut* algorithm, takes a graph, G, and performs the following:

1. construct a shortest path tree, T, over G,
2. iteratively add to T all edges that do not create chordless cycles longer than a user-specified threshold
3. terminate when no additional edges can be added

The resulting graph is a sub-graph of G in which a set of edges has been removed to eliminate geodesic cycles.

By performing this procedure on a cylindrical manifold, the points can be effectively embedded in 2D; however, interpreting this parameterization is problematic. There is no reliable mechanism for distinguishing cyclic and linear axes in the embedded space because the cutting algorithm discards all cyclic information. The next section presents a method for transforming the neighborhood graph to a planar topology while retaining the cyclic information. This approach is based on the idea that cyclic information can be represented by a repetition of the original points.

3.2. Cut-Clone-Cut Procedure

Given a distance matrix D, we begin by constructing a neighborhood graph $G = (V, E)$ with a vertex for each point and edges between each point and its k nearest neighbors. We then use the cut algorithm, described in Section 3.1, to identify a set of edges, E_r, which, when removed, convert G from a cylindrical to a planar topology (see Figure 2(a)). Next, we clone the original vertices V' and edges E', and makes a graph G_c including both copies, $G_c = (V \cup V', E \cup E')$. We then modify G_c as follows:

For each edge $(v_i, v_j) \in E_r$

Remove (v_i, v_j) and (v'_i, v'_j)

Add (v_i, v'_j) and (v'_i, v_j)

By performing this operation on each edge in E_r, we connect the two graphs in a manner that preserves the all of the geodesic relationships between the points (see Figure 2(b)). This combined graph, G_c, still has a cylindrical topology, so before it can be embedded into 2D we must cut a second time. Because the graph is redundant, no distance information is lost in this operation.

Cutting G_c yields a graph G_{cc} (see Figure 2(c)), which has the following useful properties. First, the graph has a planar topology, which can be embedded in two dimensions. Second, it provides a method for finding geodesic cycles, the shortest non-contractible cycle through a given point. In G_{cc}, the geodesic cycle through a given point is the shortest path between the point and its clone (see Figure 3(a)).

In a right cylindrical manifold, all geodesic cycles are parallel to the cyclic axis of the intrinsic parameterization. We exploit this property in the following embedding step to find an explicit parameterization of the cyclic and linear axes.

3.3. Recovering Minimal Parameterizations

The cut-clone-cut graph, G_{cc}, is then embedded into 2D using Isomap to create an intermediate embedding. This embedding is not necessarily aligned with the intrinsic pa-

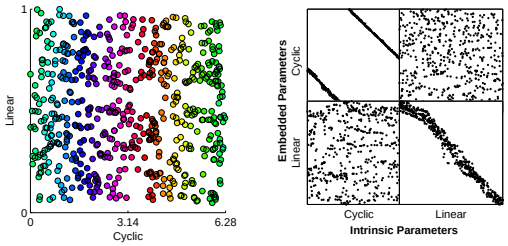

(a) (b)

Figure 4. **Cylindrical embedding of a right cylinder.** (a) The cylindrical embedding generated by the cut-clone-cut procedure and (b) a comparison of the embedding to the known intrinsic coordinates. The linear axes of the embedded and intrinsic coordinates are highly correlated, as are the cyclic coordinates.

rameter axes. However, we can define a cyclic and a linear axis by rotating the embedding to a new basis as follows.

We define a *cycle vector* to be the vector from a vertex to its clone in the intermediate embedding. These vectors approximate continuous geodesic cycles around the cylinder and should therefore be parallel to the cyclic axis. Based on this, we define the cyclic axis to be the mean cycle vector and the linear axis to be perpendicular to the cyclic axis. Projecting the intermediate embedding onto these new axes creates the aligned embedding.

The final coordinates are found by subtracting the mean from the aligned embedding, and, for each point, using the coordinates of its clone or original (choosing the closer to the y-axis). Note that this allows our algorithm to handle cases in which the initial cut is not straight. The coordinates are then rescaled so that the cyclic period (defined by the mean cyclic vector magnitude) is 2π and the linear coordinates range from $[0, 1]$.

3.4. Cut-Clone-Cut Results

Given a set of points isometric to a perfect cylinder, the results in Figure 4 show that there is a one-to-one correlation between intrinsic and the embedded coordinates output by the cut-clone-cut algorithm. However, given a set of points sampled from a non-perfect cylindrical manifold (a truncated cone shown in Figure 6(b)), the results in Figure 5(a) show that the intrinsic and embedded coordinates do not correlate one-to-one. Distorted cylinders such as this are common in natural image sets; we present a solution to this problem in the following section.

4. Embedding Distorted Cylinders

In this section we define two common types of distortion found in cylindrical image manifolds, describe an iterative procedure that modifies image distances to reduce this distortion and use this procedure to recover a cylindrical em-

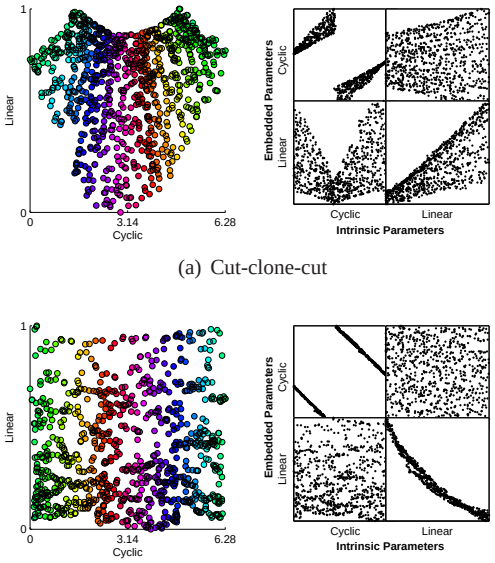

(a) Cut-clone-cut

(b) Radial-distortion reduced

Figure 5. **Cylindrical embedding of a cone.** (left) The cylindrical embedding generated by (a) the cut-clone-cut procedure and (b) the cut-clone-cut procedure with radial-distortion reduced. (right) Coordinates from the embeddings compared to known intrinsic coordinates. In (a) there is correlation between all coordinates. After distortion is reduced in (b), there is a one-to-one correlation between the intrinsic and embedded coordinates.

bedding from non-perfect (distorted) cylindrical manifolds.

The first type of distortion, *radial distortion*, can be visualized in 3D as changing the radius of the circle as it sweeps along the central axis of the cylinder. The second type, *axial distortion*, can be visualized as bending of the central axis. Examples of simple image sets that illustrate these distortions are given in Figure 6; examples of natural image sets with both types of distortion are given in Section 5.

These methods for reducing distortion operate by rescaling edge distances in the G_{cc} graph computed in the cut-clone-cut procedure. The final embedding and axis-alignment discussed in Section 3.3 is performed after these modifications to G_{cc}.

Properties of geodesic cycles on perfect cylinders are used to guide updates to edge lengths. The goal is to modify G so that geodesic cycles in G have the same properties as geodesic cycles on a perfect cylinder. To help preserve relative neighborhood structure we do not rescale edge lengths independently. Instead we use a vertex-scaling framework in which each vertex has a scale factor. Each edge is rescaled by the average scale factor of its vertices.

The iterative vertex-scaling framework used by each distortion reduction procedure consists of iterations of the following:

1. **Determine prospective edge lengths.** For each edge (u, v) with length $d(u, v)$ in the graph $G = (V, E)$, de-

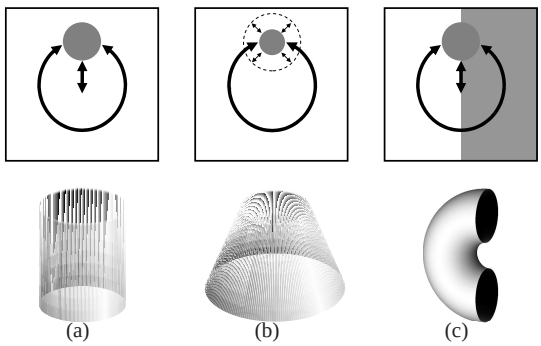

Figure 6. **Simple cylindrical image manifolds.** The top figures represent image sets with one cyclic and one linear intrinsic parameter—each dimension designated by an arrow. Differences in the image generating function result in different manifolds in image space. The revolving and translating circle in (a) generates an image set that is isometric to a perfect cylinder. The rotating circle with variable radius (b) generates a cone because image distances are greater when the circle is larger. The circle moving as in example (a) against a variable background (c) generates a bent cylinder because image distances are greater when the circle is over lighter regions.

termine a prospective length $\hat{d}(u,v)$. The prospective lengths depend on the type of distortion being reduced and will be detailed later.

2. **Solve for vertex scale factors.** Each edge contributes a constraint of the form:

$$(s_u + s_v)\frac{d(u,v)}{2} = \hat{d}(u,v), \tag{1}$$

where s_u, s_v are the unknown *vertex scale factors*.

   (a) Using Equation 1, form the $|E| \times |V|$ sparse matrix A such that the i-th row, which corresponds to edge (u,v), has two non-zero entries:

$$a_{i,u} = a_{i,v} = \frac{d(u,v)}{2}.$$

   Let the column vector R be defined as:

$$\forall i \in \{1,\ldots,|E|\},\ R_i = \hat{d}(u,v)$$

   (b) Solve for the set of vertex scale factors $S = (s_1,\ldots,s_n)^\mathsf{T}$, that minimizes the squared residual error of the following linear system:

$$AS = R.$$

3. **Update edge lengths.** Update G_c using the the vertex scale factors S to calculate new edge lengths. The new edge lengths are the entries of AS.

Iteration continues until vertex scale factors stabilize; in our experiments, this is typically four or five iterations.

We now describe the method for determining prospective edge lengths for both types of distortion.

4.1. Reducing Radial Distortion

On a perfect cylinder, all geodesic cycles are equal length. We approximate this property by iteratively modifying edge lengths until all geodesic cycles in the nearest-neighbor graph G are approximately equal length.

The method for determining prospective edge lengths $\hat{d}(u,v)$ for use in the iterative vertex-scaling framework is as follows:

1. **Extract geodesic cycles from the cut-clone-cut graph.** Let C be the set of geodesic cycles in G and let $|c^*|$ be the length of the longest cycle $c^* \in C$.

2. **Determine prospective edge lengths.** Observe that geodesic cycle c can be made as long as c^* by setting edge lengths as follows:

$$\forall (u,v) \in c,\ \hat{d}_c(u,v) = \frac{|c^*|}{|c|} d(u,v).$$

   Then, the prospective length for each edge (u,v) is:

$$\hat{d}(u,v) = \max_{c \in C} \hat{d}_c(u,v). \tag{2}$$

If we ignore vertex scaling, choosing $\hat{d}(u,v)$ by Equation 2 causes the shortest geodesic cycle to become as long as the previous longest geodesic cycle. However, at each iteration many geodesic cycles lengthen—even the longest—because they often share edges with shorter geodesic cycles.

We apply radial-distortion reduction to the truncated-cone point set from Section 3.4. This results in an explicit cylindrical embedding, with good correlation between the intrinsic and embedded coordinates, as illustrated in Figure 5(b).

4.2. Reducing Axial Distortion

On a perfect cylinder all geodesic cycles are parallel—the length of the shortest path from each point on one cycle to the nearest on another is the same. In the nearest-neighbor graph G, we approximate this property by modifying lengths of edges on shortest paths between each the geodesic cycles.

The method for determining prospective edge lengths, based on the parallel geodesic-cycle objective, is as follows:

1. **Extract geodesic nearest neighbors.** Let C be the set of geodesic cycles in G. Define the geodesic nearest neighbors for a pair of geodesic cycles $U, V \in C; U \neq V$ as a set

$$Q_{U,V} = \{(u_i, v_i) \mid u_i \in U; v_i \in V\}$$

   such that for each pair (u_i, v_i) u_i is the closest vertex to v_i and vice-versa.

2. **Define UV prospective edge lengths.** Define $\mu_{U,V}$ to be the average length of the shortest paths between the pairs of vertices in $Q_{U,V}$. Using this notation, if two geodesics cycles U, V are parallel then $\forall (u,v) \in Q_{U,V},\ d(u,v) = \mu_{U,V}$.

   Observe that we can make geodesic cycles U, V parallel by setting edge lengths so that all shortest paths are equal length. For a pair of vertices $(u_i, v_i) \in Q_{U,V}$, let $(p_1, p_2), \ldots, (p_{r-1}, p_r)$ be the set of edges on the shortest path between u_i, v_i, where p_1 is u_i and p_r is v_i.

   The UV prospective length is then

$$\hat{d}_{U,V}(p_i, p_{i+1}) = d(p_i, p_{i+1}) \frac{\mu_{U,V}}{d(u_i, v_i)}$$

   for edges on shortest paths between vertices in $Q_{U,V}$; $\hat{d}_{U,V}$ is zero elsewhere.

3. **Combine UV prospective edge lengths.** For a given edge (u, v), the prospective length is the average of all non-zero UV prospective length:

$$\hat{d}(u,v) = \frac{1}{s} \sum_{U,V \in C; U \neq V} \hat{d}_{U,V}(u,v) \qquad (3)$$

   where s is the number of non-zero values of $\hat{d}_{U,V}(u,v)$ for all pairs of geodesic cycles $U, V \in C$ such that $U \neq V$.

Results of applying this algorithm to a cylinder with axial distortion (a half-torus) are shown in Figure 7.

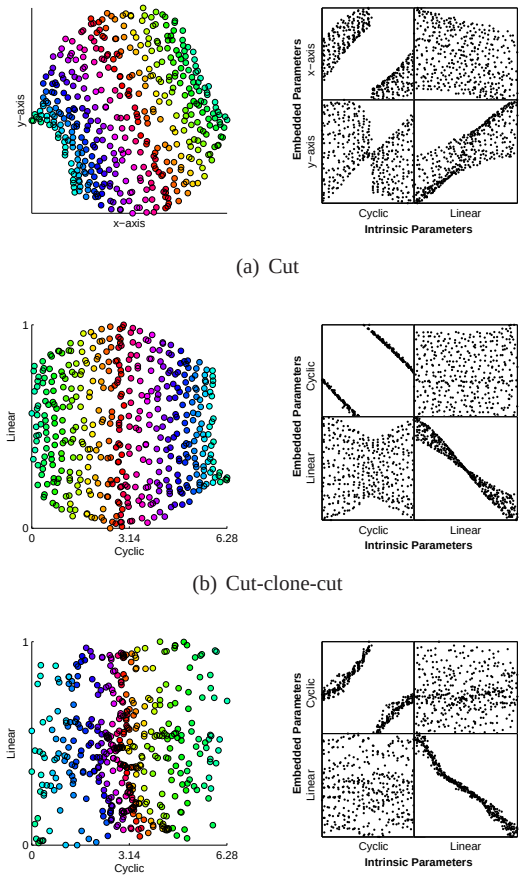

(a) Cut

(b) Cut-clone-cut

(c) Axial-distortion reduced

Figure 7. **Cut-clone-cut embedding of a half-torus.** (left) Scatter plots of the cylindrical embeddings generated by (a) cutting the cylinder, (b) the cut-clone-cut procedure, and (c) axial distortion reduction. (right) Coordinates from the embeddings compared to known intrinsic coordinates.

5. Results

In this section, we show the results of our algorithm, compared against the cyclic cut described in Section 3.1, the 2D Isomap embedding and the known intrinsic parameters (when available). We evaluated our algorithm on the following image sets.

'Rice' image set The 'rice' image set consists of 10,000 synthetic 16×16 intensity images with two intrinsic parameters (rotation and translation of the grain), and was previously used to test algorithms for cyclic data sets [8]. Vector quantization is applied in image space to reduce the image set to 300 images. This image set has significant axial distortion because the distance between two images due to translation depends upon the angle of the grain—there is a greater distance when the long axis of the grain is perpendicular to the translation.

Statue image set This data set consists of multiple views of a rigid object, used previously to illustrate the "non-flat" embedding procedure [9]. The object was placed on a turntable and rotated 360 degrees, and the camera was mounted to a rotating arm which covered a 48 degree range of elevations. Samples were taken uniformly in 6 degree increments, resulting in a total of 480 images. The scene has two underlying degrees of freedom; the table rotation is cyclic and the camera elevation is linear, thus the underlying image manifold is cylindrical.

Despite regular sampling, distortions in the graph exist that cause the embedding to not be perfectly cylindrical. Distortions result because the image differences caused by rotation depend on the camera elevation.

Drill image set This data set consists of 76 images of a rotating drill. As the drill turns, wood shavings gradually accumulate. The phase of the rotation and the amount of

accumulated debris result in cyclic and linear variations, respectively.

The results of our method are shown in Figure 8. For comparison we also show the results of cutting without alignment or correction for distortions and traditional Isomap. In each example, the cut algorithm gives interpretable results, but the embedded coordinates do not correspond to the intrinsic parameters. These results demonstrate our algorithm's ability to extract meaningful coordinates from natural image sets, even when the underlying cylindrical manifold includes axial and radial distortion. In the following section, we apply our algorithm to cardiac ultrasound images, and use the extracted coordinates to eliminate noise.

5.1. Application to Cardiac Ultrasound

Ultrasound is a low-cost imaging modality that is fast, radiation free, and adaptable to many tasks. The chief drawback is low image quality. When ultrasound images are taken of repetitive motions within the body, the images often lie along a low-dimensional image manifold. When those motions include the heart, the image manifold often has a cyclic structure.

We applied our techniques for cylindrical manifold learning to a sequence of 105 cardiac ultrasound images, and we focus on the valve region. The two degrees of freedom within this video are the changing appearance of the valve (indexed by the cyclic phase of the heartbeat), and a rigid transformation of the scene due to transducer motion.

Manifold embedding of the images offers an opportunity for noise reduction [11], by registering and then combining images taken during approximately the same part of the heartbeat cycle. Averaging multiple images offers the potential for noise reduction, but the average of images taken at different parts of the heartbeat leads to motion blur. Figure 9 shows promising initial results when using cylindrical embedding for image de-noising. Compared to other simple techniques, averaging images with similar cyclic parameters provides sharp images with less noise and noticeably less motion blur, especially for images taken when the valve is moved significantly.

6. Conclusion

To our knowledge, this paper is the first to provide a minimal parameterization directly from a distance matrix derived by measuring geodesic distances between points on a cylinder. This extends the benefits of NLDR for perceptual organization of many biological data sets which have both cyclic and linear parameters, and is a first step toward learning natural image manifolds with complex topology.

References

[1] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, 2002.

[2] D. L. Donoho and C. E. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *PNAS*, 100(10):5591–5596, 2003.

[3] X. Gu and S. Yau. Surface classification using conformal structures, 2003.

[4] X. Gu and S.-T. Yau. Global conformal surface parameterization. In *SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 127–137, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[5] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley and Sons, 2001.

[6] I. T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 1986.

[7] J. A. Lee and M. Verleysen. How to project 'circular' manifolds using geodesic distances? In *Proceedings of European Symposium on Artificial Neural Networks*, pages 223–230, April 2004.

[8] J. A. Lee and M. Verleysen. Nonlinear dimensionality reduction of data manifolds with essential loops. *Neurocomputing*, 67:29–53, August 2005.

[9] R. Pless and I. Simon. Embedding images in non-flat spaces. In *Proc. of the International Conference on Imaging Science, Systems, and Technology*, 2002.

[10] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[11] R. Souvenir and R. Pless. Isomap and non-parametric models of image deformation. In *IEEE Workshop on Motion and Video Computing (WACV/MOTION'05)*, volume 2, pages 195–200, 2005.

[12] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[13] G. Tewari, C. Gotsman, and S. J. Gortler. Meshing genus-1 point clouds using discrete one-forms. Technical Report CS TR-14-05, Harvard University, June 2005.

[14] K. Q. Weinberger and L. K. Saul. Unsupervised learning of image manifolds by semidefinite programming. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-04)*, volume II, pages 988–995, Washington D.C., 2004.

[15] H. Winnemöeller, A. Mohan, J. Tumblin, and B. Gooch. Light waving: Estimating light positions from photographs alone. *Computer Graphics Forum*, 24(3):433–438, 2005.

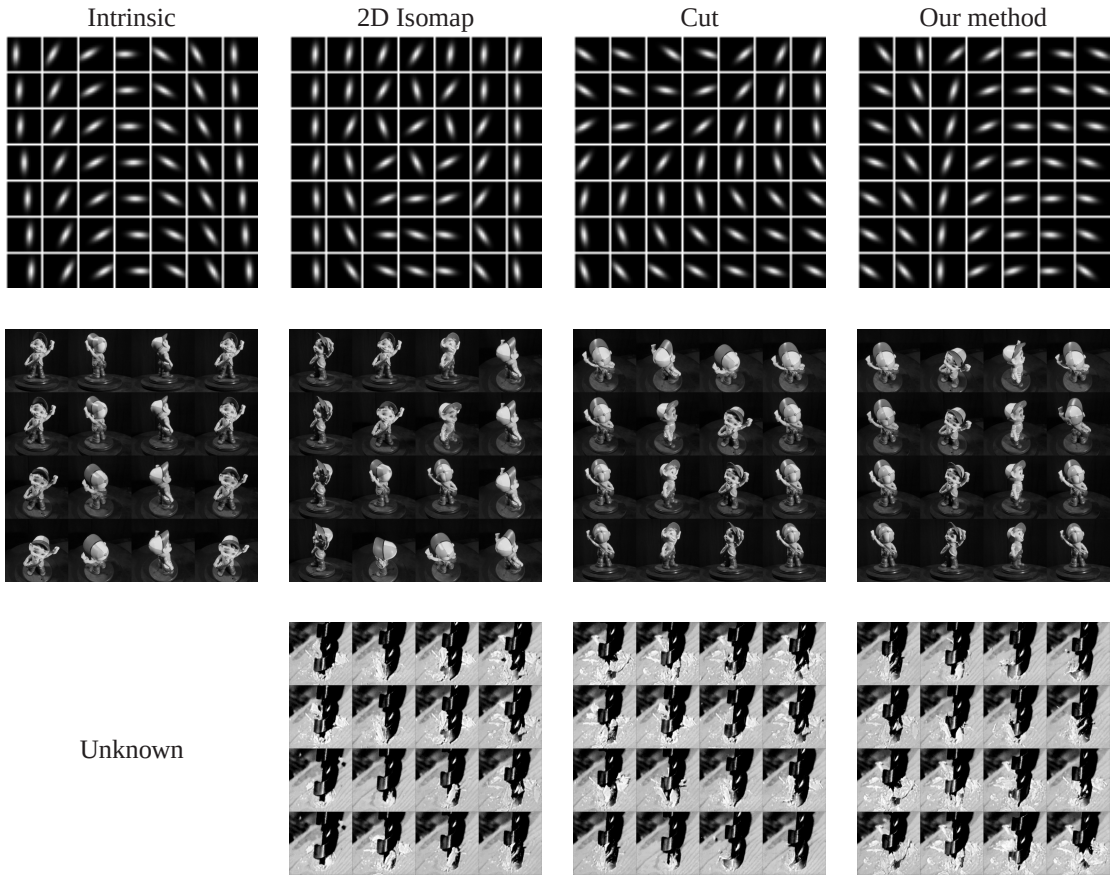

Figure 8. **Cylindrical embeddings of natural image sets.** In these plots, the underlying coordinates are divided into an evenly spaced grid, and the image closest to the center of each grid cell is shown. When an explicit cylindrical parameterization is known, the x-axis represents the cyclic coordinate and the y-axis represents the linear. The first row shows the 'rice' image set. The second row shows the statue image set. The third row shows the drill image set. In each example, our method successfully recovers a parameterization that distinguishes between the cyclic and non-cyclic variations in the scene, whereas the other parameterizations often combine the two.

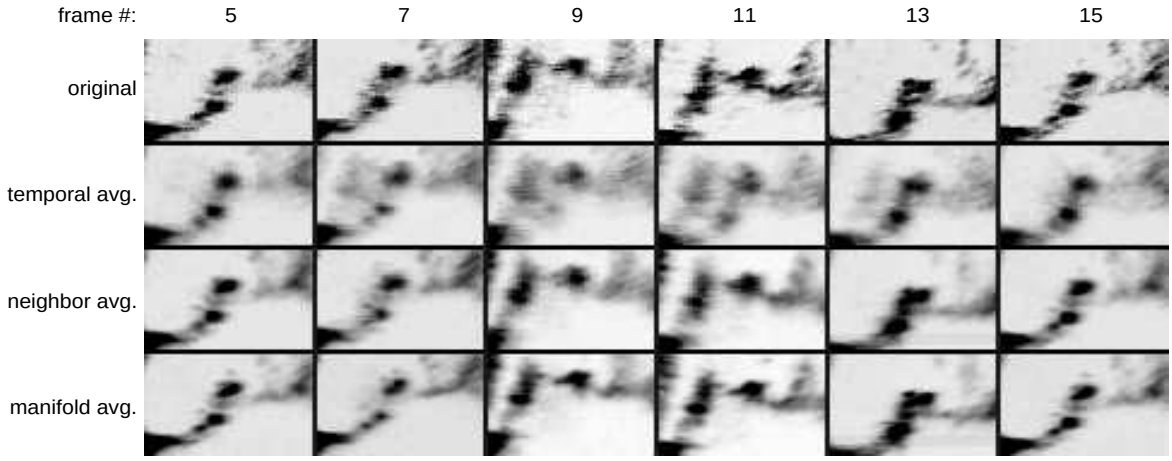

Figure 9. **Noise reduction in cardiac ultrasound.** This figure shows six frames from a cardiac ultrasound video of the valve of a beating heart. The top row shows six original images, the second row shows the results of temporal averaging. The third row shows the results of finding the 10 most similar images, performing translational alignment then averaging, and the final column shows the results of finding the 10 images with the most similar cyclic parameter, performing translational alignment then averaging. This shows the most improvement during frames 7, 9, and 11, where there is significant motion. Using manifold parameters to average images that do not introduce motion blur results in improved image sharpness over naive averaging of nearby images.