

Nonlinear Spherical Shells for Approximate Principal Curves Skeletonization

Odest Chadwicke Jenkins
Chi-Wei Chu
Maja J Matarić

CJENKINS@USC.EDU
CHUC@USC.EDU
MATARIC@USC.EDU

Abstract

We present Nonlinear Spherical Shells (NSS) as a non-iterative model-free method for constructing approximate principal curves skeletons in volumes of d dimensional data points. NSS leverages existing model-free techniques for nonlinear dimension to remove nonlinear artifacts in data. With nonlinearities removed and topology preserved, data embedded by such procedures are assumed to have properties amenable to simple skeletonization procedures. Given these assumptions, NSS is able extract points in the “middle” of the volume data and hierarchically link them into principal curves, or a set of 1-manifolds connected at junctions.

1. Introduction

Skeletonization is the process of extracting curves or surfaces that describes the shape of an object or connected distribution of points in a d dimensional space. Skeletonization can be performed through a variety of techniques, including morphological thinning, Delaunay triangulation, self-organizing maps, and distance transforms. Additionally, the representations produced through skeletonization take a variety of forms, such as a star-skeleton (Fujiyoshi & Lipton, 1998), a shock scaffold (Leymarie & Kimia, 2003), principal curves (Hastie & Stuetzle, 1989), and the (well-known) medial-axis (Blum, 1967). While skeletonization has applicability to a diverse range of topics, including medical image analysis, robotics, computer graphics, and biometrics, no single method for skeletonization has been demonstrated as the clear choice for all skeletonization problems.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Motivated by recent work in model-free markerless motion capture (Chu et al., 2003), we present *Nonlinear Spherical Shells (NSS)* as a model-free non-iterative skeletonization method that leverages existing nonlinear dimension reduction techniques. Similar to work by Kegl and Krzyzak (Kégl & Krzyzak, 2002), NSS aims to skeletonize d -dimensional distributions of data points as *principal curves* (Hastie & Stuetzle, 1989). Principal curves are a nonlinear generalization of principal components as a set of 1-manifolds that pass through the “middle” of a d -dimensional data distribution and connected by junctions. For data with dimensionality $d > 2$, principal curves provide more compact skeletons than medial-axes, which produces skeletons of higher manifold dimensionality as surfaces.

In the remainder of this paper, we describe the NSS method and its foundation in nonlinear dimension reduction. Results are presented that illustrate the effectiveness of NSS through its application to 2D data of handwritten digits and 3D data of human volumes.

2. Easier Skeletonization through Dimension Reduction

NSS assumes the inherent difficulty in skeletonization is due to shape-specific nonlinearities, obfuscating the topology underlying a data distribution. In this situation, a transformation of great benefit could reinterpret data such that nonlinearities are removed and underlying topology is preserved. Given this transformation, a set of input data points can be reduced to a form amenable to skeletonization. For example, consider a collection of points that comprise the volume of a human arm. In a fully extended state, the arm could plausibly be modeled as a cylinder, whose skeleton could be a segment of the cylinder axis. However, as this arm moves, rotation about the elbow causes this cylinder assumption to be violated. Given a topology-preserving transformation mechanism, any pose of the arm could be reduced a pose-independent configuration (i.e., straightened out) amenable to the cylinder

assumption, allowing for simple skeletonization.

Fortunately for us, recent work in manifold learning has produced a suite of useful techniques for model-free nonlinear dimension reduction, including Isomap (Tenenbaum et al., 2000), Kernel PCA (KPCA) (Schölkopf et al., 1998), and Locally Linear Embedding (LLE) (Roweis & Saul, 2000), based on pairwise relationships between data points. Conceptually, these methods should provide the appropriate topology-preserving nonlinearity-removing we are seeking. Additionally, these techniques do not require *a priori* specification of the topology to be uncovered, unlike a self-organizing map. However, each technique has its own means for characterizing pairwise relationships, leading to significantly different transformations. As shown in Figure 1, KPCA, Isomap, and LLE each produce different embeddings when applied to a set of 3D points comprising a human volume, acquired from multiple calibrated cameras through *voxel carving* (Szeliski, 1993).

Of these dimension reduction mechanisms, Isomap intuitively appears the best suited for skeletonization purposes. When applied to human volumes (Figure 1), Isomap consistently reduces the data to a configuration similar to a “*Da Vinci*” pose in 3 dimensions. Given distinguishable limbs and volumes of genus 0, Isomap-reduced volumes have desirable *reduction properties* of: *i*) the center of mass of the embedded volume is located at the origin, *ii*) the distance between the origin and other points in the volume monotonically increase while moving along the body, and *iii*) articulations of the volume are separated as far apart as possible. The applicability of Isomap to this problem is not surprising considering the use of shortest-path geodesic distance in several skeletonization methods (Zhou & Toga, 1999; Fricout et al., 2002). These methods, however, are sensitive to seed points selected to serve as root points for shortest-path computation. In contrast, Isomap performs multidimensional scaling on a matrix of all-pairs shortest-paths, allowing use of feature space centering to avoid explicit seed point selection.

3. Nonlinear Spherical Shells

NSS is a five-step procedure, described as follows:

1. embedding of input data into topology-preserving nonlinearity-removing coordinates (e.g., Isomap)
2. partitioning the embedded data into concentric spherical shells
3. clustering data points in each shell partition

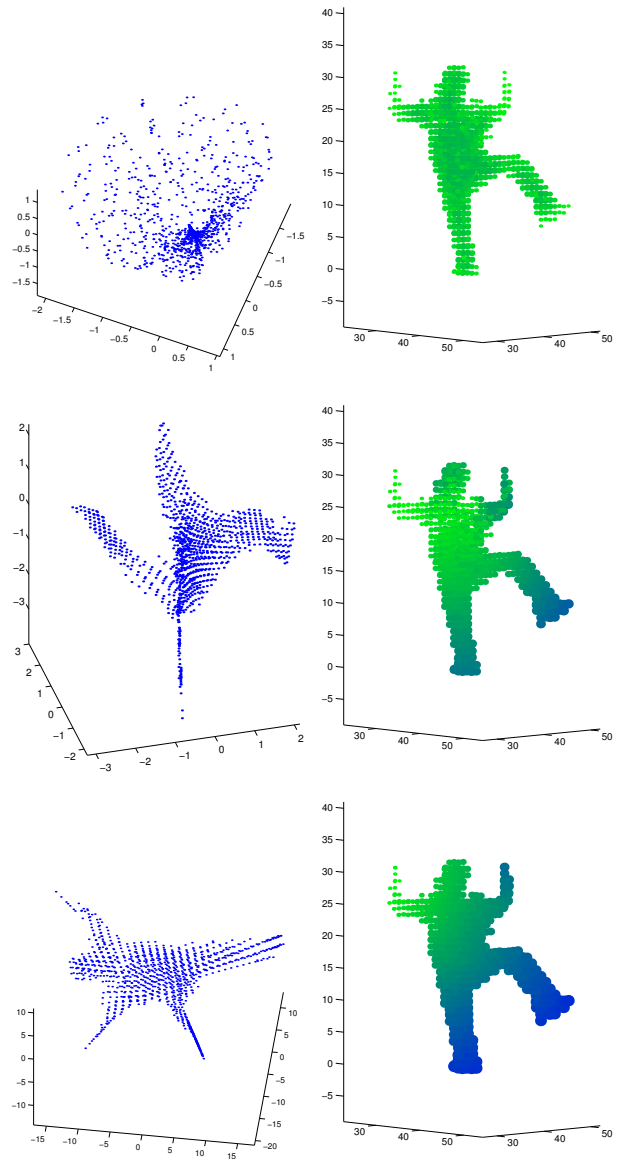


Figure 1. (Left column) A 3D human volume consisting of 1073 points embedded through KPCA (top, Gaussian RBF neighborhood = 3), Isomap (middle, neighborhood radius = 2), and LLE (bottom, 7 nearest neighbors). (Right column) Plot input volume points weighted by embedding-space distances with respect to the most extremal point on the right hand. The embedding-space distance for each point is illustrated as small and green for small distances and large and blue for large distances.

4. linking of data clusters in adjacent shell partitions into a hierarchy
5. computing skeleton points as centroids of data clusters

Nonlinear dimension reduction is used in the first step of the NSS procedure to produce an embedding that realizes a shape-independent topology-preserving configuration of the input data. The resulting configuration is expected to have the reduction properties mentioned in the previous section. For our NSS implementation, we applied the Isomap¹ directly to the input data. Isomap requires user specification specify of only the dimensionality of the embedding and a function for constructing local neighborhoods about each data point (e.g., K-nearest neighbors or an epsilon radius). Dimension reduction is not necessarily the aim of the step. Thus, the embedded space can potentially be of equal or greater dimensionality as the input space.

Assuming adherence to the reduction properties, the embedded data can be partitioned into concentric spherical shells S :

$$\frac{(i-1) \|x_{max}\|}{|S|} \leq \|x\| < \frac{(i) \|x_{max}\|}{|S|} \Rightarrow x \in S_i \quad (1)$$

where S_i is the set of points in the i^{th} shell partition, $|S|$ is the number of shell partitions, and x_{max} is the furthest point from the origin. The motivation for this partitioning is that separable clusters corresponding to individual articulations of the volume will form as we move away from the origin. Additionally, the boundary of a spherical shell in the embedding space corresponds to a potentially non-spherical and nonlinear boundary in the input space.

A clustering mechanism is applied to individual shell partitions to identify clusters $C_{i,j}$, where i and j are shell and cluster identifiers, respectively. Many clustering mechanisms (Jain & Dubes, 1988) can be used for this purpose. We chose to cluster using the one-dimensional “sweep-and-prune” technique (Cohen et al., 1995) for bounding clusters by axis-aligned boxes. Sweep-and-prune clustering requires specification of a partitioning distance threshold data projected onto a single axis. Clustering is performed by partitioning data with respect to its projection on a single axis and iterating over every axis. Clustering in this fashion avoids the need to estimate the expected number of clusters in a partition.

Once established, partitioned clusters in adjacent shell partitions are linked together to define the topology of the skeleton. For each partitioned cluster $C_{i,j}$, parent

¹Isomap code was provided by the authors of Isomap (Thank you!) and is available at <http://isomap.stanford.edu/>.

clusters $P_{i,j}$ in the parent shell partition S_{i-1} can be found through the minimum distance between members of different clusters:

$$\min_{a \in C_{i,j}, b \in C_{i-1,k}} \|a - b\| < \epsilon_m \Rightarrow C_{i-1,k} \in P_{i,j} \quad (2)$$

or (for lighter computation) through the minimum distance between the cluster centroid $m'_{i,j}$ and the members of the potential parent cluster:

$$\min_{b \in C_{i-1,k}} \|m'_{i,j} - b\| < \epsilon_c \Rightarrow C_{i-1,k} \in P_{i,j} \quad (3)$$

where ϵ_m and ϵ_c are user-selected thresholds (potentially based on values for $\|x_{max}\|$ and $|S|$).

Given the partitioned clusters $C_{i,j}$ and cluster linkages $P_{i,j}$, the approximate principal curve for the input data can be computed by one of two different means. Both of these mechanisms rely on the fact that 1-to-1 correspondences exist between input and embedded data points to estimate principal points $m_{i,j}$. The simplest of these mechanisms computes $m_{i,j}$ as the centroid of $C_{i,j}$ in the input space. The other mechanism uses an interpolation mechanism (e.g., Shepard’s interpolation (Shepard, 1968)) to map embedding space centroids $m'_{i,j}$ into input space centroids $m_{i,j}$. The principal curve is finally formed by linking $m_{i,j}$ to its parents specified by $P_{i,j}$. Chu et al. (Chu et al., 2003) discuss additional refinement steps for the procedure to produce more aesthetic principal curves.

4. Experiments

In this section, we describe two experiments with our MATLAB NSS implementation applied to data of 2D handwritten digits and 3D human volumes. No modifications to the NSS code were made in order to accommodate the differing dimensionality between these data sets.

4.1. 2D Handwritten Digit Skeletonization

Figure 2 shows results from applying NSS to 10 handwritten digits, from the Yonsei University Digit Database, comprised of points in 2D. Skeletonization of the digits was performed using no embedding, Isomap embedding, and LLE embedding. While the Isomap embeddings typically produced the most visually appropriate skeletons, we use these results to discuss issues in the NSS method.

The 7 digit is a case illustrating the purpose of nonlinear dimension reduction. The non-embedded 7 skele-

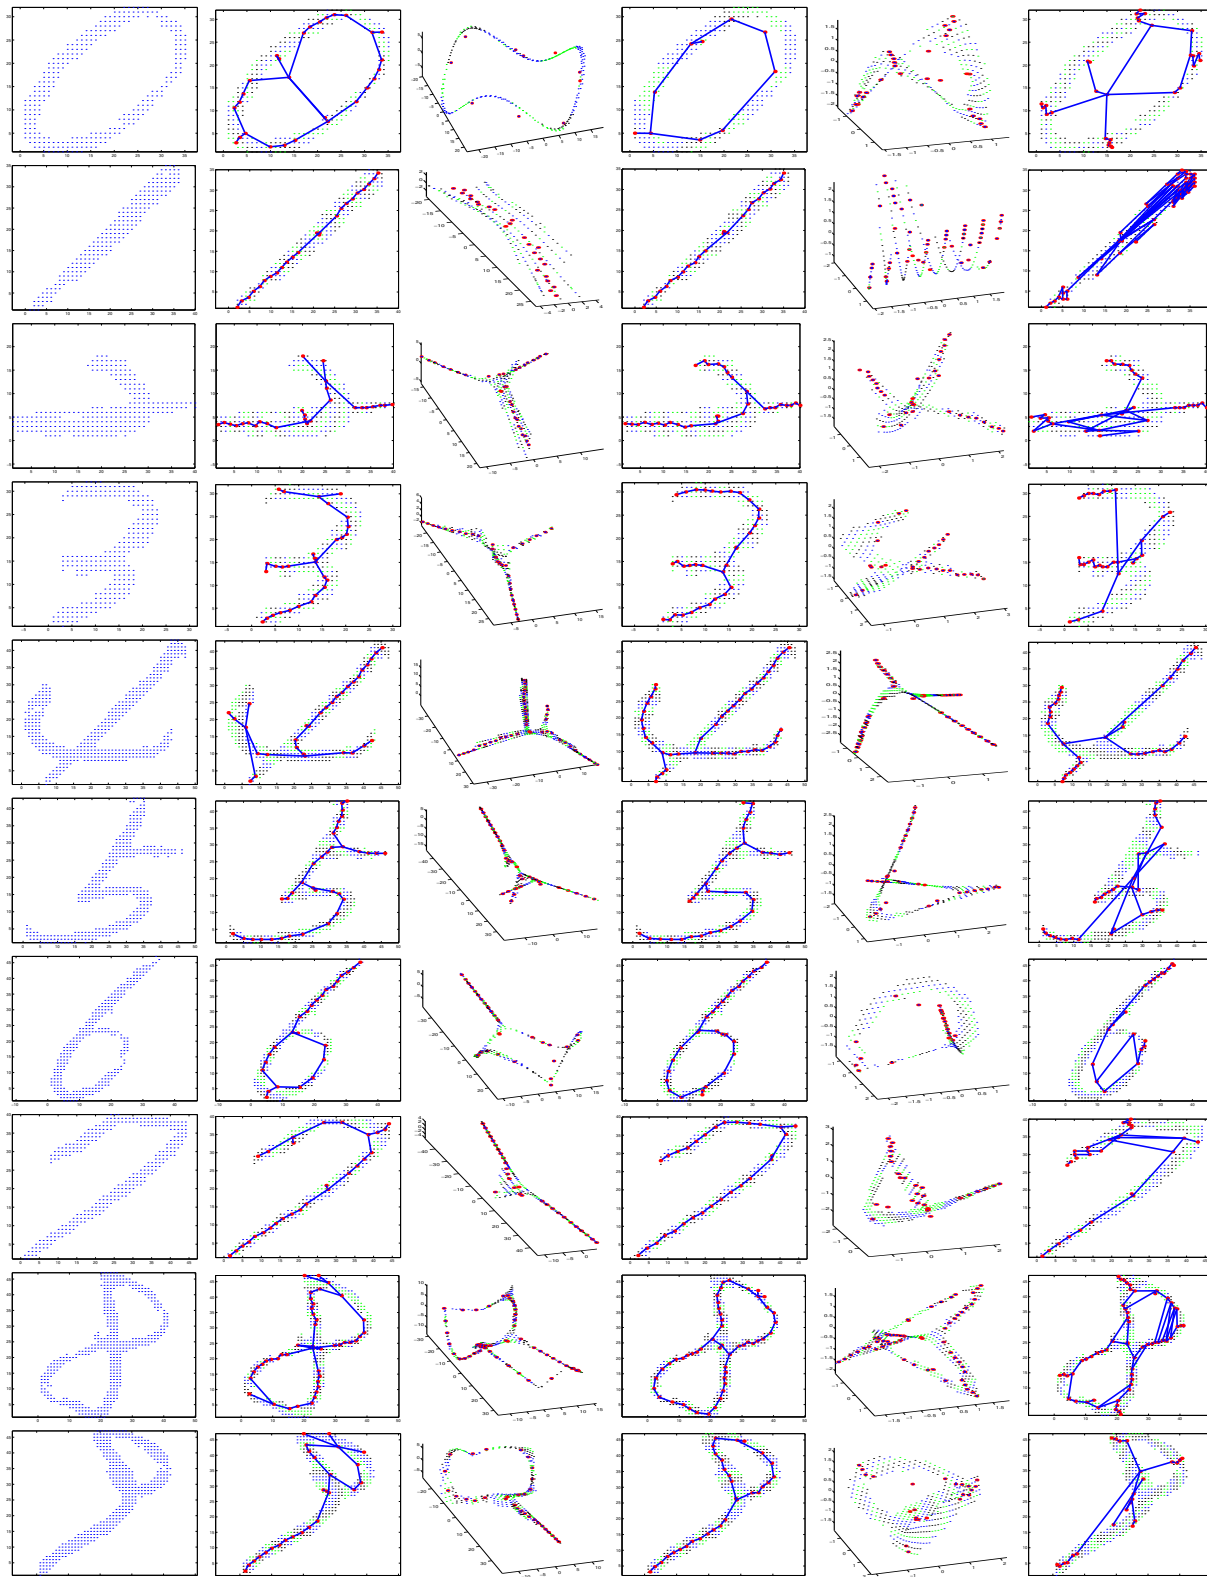


Figure 2. Results from applying NSS to 10 sets of 2D point comprising handwritten digits with 15 shell partitions. (First column) Handwritten digit images from the Yonsei University Digit Database. (Second column) NSS skeletonization in the input space without embedding. (Third column) NSS skeletonization and (fourth column) Isomap embedding with an epsilon neighborhood of 2. (Fifth column) NSS skeletonization and (sixth column) LLE embedding with 7 nearest neighbors.

ton is rooted in the middle of the long branch of the 7. Spherical partitioning results in the short branch of the 7 having no parents but instead being connected to the root through its children at the top of the 7. In contrast, the Isomap 7 skelton produces partitions that geodesically move along the curvature of the 7. Thus, the skeletonization of the 7 is reduced to the skeletonization of the 1 digit. Additionally, the left branch of the 4 digit has a significant nonlinearity that is appropriately handled by the Isomap skeleton and not handled by the non-embedded skeleton. However, junctions in the volume, such as the upper junction of the 5 digit, present problems for tuning cluster partitioning thresholds.

The 6, 8, and 9 digits illustrates how NSS works for data of genus greater than 0. NSS provides visually and topologically appropriate skeletons for these volumes. However, these skeletons are not guaranteed to be topologically consistent with the input data. For instance, the use of centroids for parent linking (Figure ??), produces that are not guaranteed to close around the loops of digit. Thus, the skeleton is unable to match the genus of the input data. While minimum inter-member distance works for these digits, we currently have no general guarantee for appropriate genus matching. However, NSS is able to handle some sparsity in the input, such as in the rightmost side of the loop in the 6 digit.

The 0 digit is a case where NSS, in its current form, is not suited for skeletonization, but provides a rough approximation. Even though the embedding of the 0 digit does not violate the reduction properties assumed by NSS, degenerate conditions are encountered that cause errors in the procedure. First, no points are located with the first several shell partitions. Thus, root cluster determination must be handled as a singular case. Our heuristic for handling this case is to consider the first populated shell partition as the first partition. The root shell partition in this case is likely to consist of a single cluster. Consequently, the centroid of the root cluster will not be among the points of this cluster. In general, the placement of cluster centers is not quite accurate for this digit, but the topology of the digit is captured and skeleton is a roughly located in an appropriate position.

4.2. 3D Human Volume Skeletonization

Figure 2 shows results from applying NSS to 10 human volumes, acquired using methods described in (Chu et al., 2003), comprised of points in 3D. Each of the volumes were of genus 0. NSS skeletonization of the volumes was performed using Isomap embed-




Figure 3. A case in which NSS inappropriately skeletonizes data of a 9 digit of genus 1 (left) to a genus 0 skeleton (right) using centroid parent linking.

ding. Visually appropriate skeletons were produced for each volume. The skeleton in the eighth frame with a visually errant skeleton link is actually appropriate due to outlier points inappropriately embedded by Isomap. Although not as aesthetic as skeletons produced by other methods, NSS relied on no postprocessing mechanism. Postprocessing could greatly improve the resulting skeleton aesthetic. Skeletons can also be improved by using and increased number of shell partitions. Higher resolution partitioning, however, would require data of higher resolution.

5. Conclusion

We have presented Nonlinear Spherical Shells (NSS) as a non-iterative model-free method for constructing approximate principal curves skeletons in volumes of d dimensional data points. We examine the use of nonlinear dimension reduction to reduce data distributions into configurations amenable to simple principal curve skeletonization. While NSS is not as matured as existing skeletonization methods, its simplicity and flexibility exhibits promise for compact skeltonization of various types of volumetric shapes without significant *a priori* assumptions and manual tuning. In addition, NSS procedures will greatly benefit and improve with progress of techniques in manifold learning.

References

- Blum, H. (1967). A transformation for extracting new descriptors of shapes. *Models for the Perception of Speech and Visual Form* (pp. 362–380). Cambridge, MA, USA: MIT Press.
- Chu, C.-W., Jenkins, O. C., & Matarić, M. J. (2003). Markerless kinematic model and motion capture from volume sequences. *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR 2003)* (pp. II: 475–482). Madison, Wisconsin, USA.
- Cohen, J. D., Lin, M. C., Manocha, D., & Ponamgi,

- M. K. (1995). I-COLLIDE: An interactive and exact collision detection system for large-scale environments. *Symposium on Interactive 3D Graphics* (pp. 189–196, 218). Monterey, California, USA.
- Fricout, G., Jeulin, D., Cullen-McEwen, L., Harper, I., & Bertram, J. (2002). 3d-skeletonization of ureteric trees. *International Symposium on Mathematical Morphology* (pp. 157–164). Sydney, New South Wales, Australia.
- Fujiyoshi, H., & Lipton, A. (1998). Real-time human motion analysis by image skeletonization. *Proc. of the Workshop on Application of Computer Vision*.
- Hastie, T., & Stuetzle, W. (1989). Principal curves. *Journal of the American Statistical Association*, *84*, 502–516.
- Jain, A., & Dubes, R. (1988). *Algorithms for clustering data*. Englewood Cliffs, NJ, USA: Prentice Hall.
- Kégl, B., & Krzyzak, A. (2002). Piecewise linear skeletonization using principal curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *24*, 59–74.
- Leymarie, F. F., & Kimia, B. B. (2003). Computation of the shock scaffold for unorganized point clouds in 3d. *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR 2003)* (pp. I: 821–827). Madison, Wisconsin, USA.
- Roweis, S. T., & Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, *290*, 2323–2326.
- Schölkopf, B., Smola, A. J., & Müller, K.-R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, *10*, 1299–1319.
- Shepard, D. (1968). A two-dimensional interpolation function for irregularly-spaced data. *Proceedings of the ACM National Conference* (pp. 517–524). ACM Press.
- Szeliski, R. (1993). Rapid octree construction from image sequences. *Computer Vision, Graphics, and Image Processing: Image Understanding*, *58*, 23–32.
- Tenenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, *290*, 2319–2323.
- Zhou, Y., & Toga, A. W. (1999). Efficient skeletonization of volumetric objects. *IEEE Transactions on Visualization and Computer Graphics*, *5*, 196–209.




Figure 4. Results from applying NSS to 10 sets of 3D points comprising human volumes. (Top rows) Human volumes obtained from a moving human captured from multiple calibrated cameras. (Middle rows) Skeletons produced for each volume. (Bottom rows) Isomap embedding (epsilon=2), clusters, and color coded partitions for each volume.