10

Topologically Constrained Isometric Embedding

Guy Rosman, Alexander M. Bronstein, and Michael M. Bronstein, and Ron Kimmel

Technion – Israel Institute of Technology, Department of Computer Science, Haifa 32000, Israel

Summary. Presented is an algorithm for nonlinear dimensionality reduction that uses both local(short) and global(long) distances in order to learn the intrinsic geometry of manifolds with complicated topology. Since our algorithm matches nonlocal structures, it is robust even to strong noise. We show experimental results on both synthetic and real data demonstrating the advantages of our approach over state-of-the-art manifold learning methods.

10.1 Introduction

Analysis of high-dimensional data is encountered in numerous pattern recognition applications. It appears that often a small number of dimensions is needed to explain the high-dimensional data. Dimensionality reduction methods such as *principal components analysis* (PCA) [15] and *multidimensional scaling* (MDS) [5] are often used to obtain a low-dimensional representation of the data. This is a commonly used pre-processing stage in pattern recognition.

While methods like PCA assume the existence of a linear map between the data points and the parametrization space, such a map often does not exist. Applying linear dimensionality reduction methods to nonlinear data may result in a distorted representation. *Nonlinear dimensionality reduction* (NLDR) methods attempt to describe a given high-dimensional data set of points as a low-dimensional manifold, by a nonlinear map preserving certain properties of the data. This kind of analysis has applications in numerous fields, such as color perception, pathology tissue analysis [11], enhancement of MRI images [13], face recognition [7], and biochemistry [24], to mention a few. In the field of motion understanding, several methods use NLDR techniques [22, 29], including works presented in this book, see Chapter 2.

As the input data, we assume to be given N points in the M-dimensional Euclidean space, $\{\mathbf{z}_i\}_{i=1}^N \subset \mathbb{R}^M$. The points constitute vertices of a proximity graph with the set of edges E; the points $\mathbf{z}_i, \mathbf{z}_j$ are neighbors if $(i,j) \in E$. The data points are samples of an m-dimensional manifold $\mathcal{M} \subset \mathbb{R}^M$, where

B. Rosenhahn et al. (eds.), Human Motion – Understanding, Modelling, Capture, and Animation, 243–262.
© 2008 *Springer.*

$M \gg m$. The manifold is assumed to have a parametrization, represented by the smooth bijective map $\varphi : \mathcal{C} \subset \mathbb{R}^m \to \mathcal{M}$.

The goal is to find a set of points $\{\mathbf{x}_i\}_{i=1}^N \subset \mathbb{R}^m$ representing the parametrization. We will use the $N \times m$ matrix $\mathbf{X}$ representing the coordinates of the points in the parametrization space. Many NLDR techniques attempt finding an m-dimensional representation for the data, while preserving some *local* invariant. Another class of algorithms preserves *global* invariants, like the *geodesic distances* d_{ij}, approximated as shortest paths on the proximity graph.

The *locally linear embedding* (LLE) algorithm [30] expresses each point as a linear combination of its neighbors,

$$\mathbf{z}_i = \sum_{j:(i,j)\in E} w_{ij}\mathbf{z}_j,$$

and minimizes the least-squares deviation from this local relation. This gives rise to a minimal eigenvalue problem

$$\mathbf{X}^* = \underset{\mathbf{X}^T\mathbf{X}=\mathbf{I}}{\operatorname{argmin}} \operatorname{trace}(\mathbf{X}^T\mathbf{M}\mathbf{X}),$$

where $\mathbf{M}$ is an $N \times N$ sparse matrix with elements

$$m_{ij} = \begin{cases} 1 - w_{ij} - w_{ji} + \sum_k w_{ki}w_{kj} & i = j, \\ -w_{ij} - w_{ji} + \sum_k w_{ki}w_{kj} & i \neq j. \end{cases}$$

A similar approach is the *Laplacian eigenmaps* algorithm proposed by Belkin and Niyogi in [1]. It solves the following minimum eigenvalue problem,

$$\mathbf{X}^* = \underset{\substack{\mathbf{X}^T\mathbf{D}\mathbf{X}=\mathbf{I} \\ \mathbf{X}^T\mathbf{D}\mathbf{1}=\mathbf{0}}}{\operatorname{argmin}} \operatorname{trace}(\mathbf{X}^T\mathbf{L}\mathbf{X}),$$

where $\mathbf{1}$ is an all-ones vector, and $\mathbf{L}$ is the Laplacian of the graph, an $N \times N$ sparse matrix with elements

$$l_{ij} = \begin{cases} w_{ij} & (i,j) \in E \text{ and } i = j, \\ -\sum_{k\neq i} w_{ki} & i = j, \\ 0 & \text{else}, \end{cases}$$

and $\mathbf{D}$ is a diagonal matrix with elements $d_{ii} = -l_{ii}$. The weights w_{ij} can be selected, for example, as $w_{ij} = \exp\{-\frac{1}{\sigma^2}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2\}$. The Laplacian $\mathbf{L}$ is a discrete approximation of the Laplace–Beltrami operator and its eigenfunctions are locally flat. This geometric interpretation is valid if the sampling of the manifold is uniform.

The *diffusion maps* framework [11] tries to generalize the work of Belkin and Niyogi and other local methods. Writing the quadratic expression in $\mathbf{X}$ in a slightly different way,