

L26: Advanced dimensionality reduction

The “snapshot” PCA approach

Oriented Principal Components Analysis

Non-linear dimensionality reduction (manifold learning)

- ISOMAP
- Locally Linear Embedding

The “snapshot” PCA approach

Problem definition

- Imagine that we have collected a small number of samples $(x_i)_{i=1 \dots N}$ each of which has a very high number of features (e.g., high-resolution images or 3D scans)
 - In the conventional PCA approach, we would compute the sample covariance matrix as follows

$$C = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})(x_i - \bar{x})^T$$

- And then try to diagonalize it!

This approach has two problems

- First, the sample covariance matrix will not be full-rank, in which case direct inversion is not possible (we’d need to use SVD)
- Second, C will be very large (e.g., 400MB for 100x100 images with double precision)
- However, we know that at most P eigenvectors will be non-zero.... Is there a better way to find these eigenvectors?

[This material is based upon an unpublished manuscript by Sam Roweis]

7

The “snapshot” trick is based on the fact that the eigenvectors are linear combinations of the data samples

- Note that the eigenvectors capture the directions of variance, and there is no variance in directions normal to the subspace spanned by the data
- Thus, we will seek to express the PCA decomposition in a manner that depends only on the number of samples P

Derivation

- Assume that the data has been centered by subtraction of its mean
- Then, the covariance matrix can be expressed as

$$= \frac{1}{P} \sum_{i=1}^P \mathbf{x}_i \mathbf{x}_i^T$$

- Since the eigenvectors are linear combinations of the data, we can then express them as

$$\mathbf{v}_k = \sum_{i=1}^P \mathbf{x}_i c_{ki}$$

- where \mathbf{v}_k denotes the k -th eigenvector

- With this formulation, our goal becomes finding constants

- Since $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are the eigenvectors, they satisfy the condition

$$\mathbf{A}\mathbf{v}_i = \lambda_i \mathbf{v}_i$$

- Which after derivation becomes

$$\begin{pmatrix} \frac{1}{-1} & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{pmatrix} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$$

- We now define matrix \mathbf{C} , which is the sample inner product of pairs of samples (i.e., the covariance matrix is the sample outer product)

$$C_{12} = \frac{1}{-1} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

- Substituting this matrix into the previous expression yields

$$\sum_{i=1}^2 \sum_{j=1}^2 \mathbf{z}^{(i)} \mathbf{z}^{(j)T} - \sum_{i=1}^3 \mathbf{z}^{(i)} \mathbf{z}^{(i)T}$$

- which, merging subindices 1 and 2 and moving terms to the LHS, yields

$$\sum_{i=1}^2 \sum_{j=1}^2 \mathbf{z}^{(i)} \mathbf{z}^{(j)T} - \sum_{i=1}^2 \mathbf{z}^{(i)} \mathbf{z}^{(i)T} = 0$$

- This condition can be met by finding $\mathbf{z}^{(1)}$ such that*

$$\sum_{i=1}^2 \mathbf{z}^{(i)} \mathbf{z}^{(i)T} - \sum_{i=1}^2 \mathbf{z}^{(i)} \mathbf{z}^{(i)T} = 0, \quad \mathbf{z}^{(1)T} \mathbf{z}^{(1)} = 1$$

- which can be written as

=

- Therefore, the n -dim vectors $\mathbf{z}^{(i)}$ are the eigenvectors of matrix $\mathbf{Z} \mathbf{Z}^T$, which can be found in a conventional manner since $\mathbf{Z} \mathbf{Z}^T$ has size $n \times n$

- Once $\mathbf{z}^{(i)}$ have been found, the actual eigenvectors of the data \mathbf{X} are obtained by a weighted sum of the training samples

$$\mathbf{z}^{(i)} = \sum_{j=1}^n \mathbf{X}^{(j)} \mathbf{z}^{(i)T} \mathbf{X}^{(j)T} \mathbf{z}^{(i)}$$

*This is one out of possibly many solutions (i.e., ways of expressing eigenvectors as linear combinations of the samples), but one that makes the problem very easy to solve, as we see next

Oriented principal components analysis

OPCA is a generalization of PCA that uses a generalized eigenvalue problem with two covariance matrices

- The cost function maximized by OPCA is the signal-to-noise ratio between a pair of high-dimensional signals \mathbf{x} and \mathbf{y}

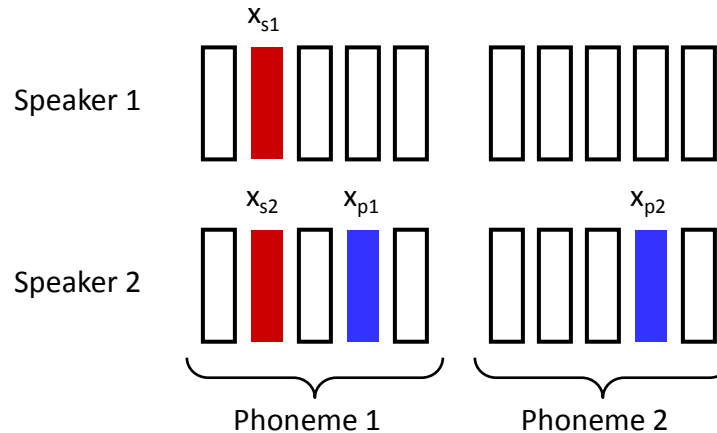
$$\left(\frac{\mathbf{x}^T \mathbf{y}}{\|\mathbf{x}\|} \right)^2 = \frac{[\mathbf{x}^T \mathbf{y}]^2}{\|\mathbf{x}\|^2} = \frac{\mathbf{y}^T \mathbf{S}_x \mathbf{y}}{\mathbf{y}^T \mathbf{S}_y \mathbf{y}} = \mathbf{y}^T \mathbf{S}_x^{-1} \mathbf{S}_y \mathbf{y}$$

- Since \mathbf{S}_x and \mathbf{S}_y are symmetric, all the generalized eigenvectors are real, and can be sorted by decreasing generalized eigenvalues
- Note that the generalized eigenvectors will NOT be orthogonal but instead will meet the constraint $\mathbf{v}^T \mathbf{S}_y \mathbf{v} = 0$
- The term “oriented” is due to the fact that \mathbf{v} is similar to the ordinary principal direction of \mathbf{S}_y , except that it is oriented towards the least principal direction of \mathbf{S}_x
 - In other words, \mathbf{S}_x “steers” \mathbf{v} away from the directions of high energy in \mathbf{S}_x
 - If \mathbf{S}_x is white noise, then there is no steering, and OPCA is identical to the PCA solution

OPCA: an example

Consider the problem of speech/speaker recognition

- Short-time speech signals are typically represented by a feature vector x , which capture some of its stationary spectral properties (e.g, LPC, cepstrum)
- Assume that we have data from two speakers (s_1, s_2) uttering two different phonemes (p_1, p_2), as illustrated below



- We are interested in finding projection vectors w of the feature vector that maximizes the signal to noise ratio, signal being linguistic information (LI) and noise being speaker information (SI)

= —

- We then extract feature vectors \mathbf{x}_1 and \mathbf{x}_2 from the same speaker but different phonemes; the difference between these two vectors will only contain linguistic information

$$\mathbf{d} = \mathbf{x}_1 - \mathbf{x}_2$$

- We also extract feature vectors \mathbf{x}_1 and \mathbf{x}_2 from the same phoneme but different speakers; the difference between these two vectors will only contain speaker information

$$\mathbf{d} = \mathbf{x}_1 - \mathbf{x}_2$$

- We define the signal and noise covariance matrices as

$$= \left[\begin{matrix} \sigma_s^2 & 0 \\ 0 & \sigma_n^2 \end{matrix} \right]$$

$$= \left[\begin{matrix} \sigma_s^2 & 0 \\ 0 & \sigma_n^2 \end{matrix} \right]$$

- And the SNR becomes

$$= \frac{\sigma_s^2}{\sigma_n^2} = \frac{\left[\begin{matrix} \sigma_s^2 \\ 0 \end{matrix} \right]^2}{\left[\begin{matrix} \sigma_s^2 \\ 0 \end{matrix} \right]^2} = \frac{\sigma_s^2}{\sigma_n^2}$$

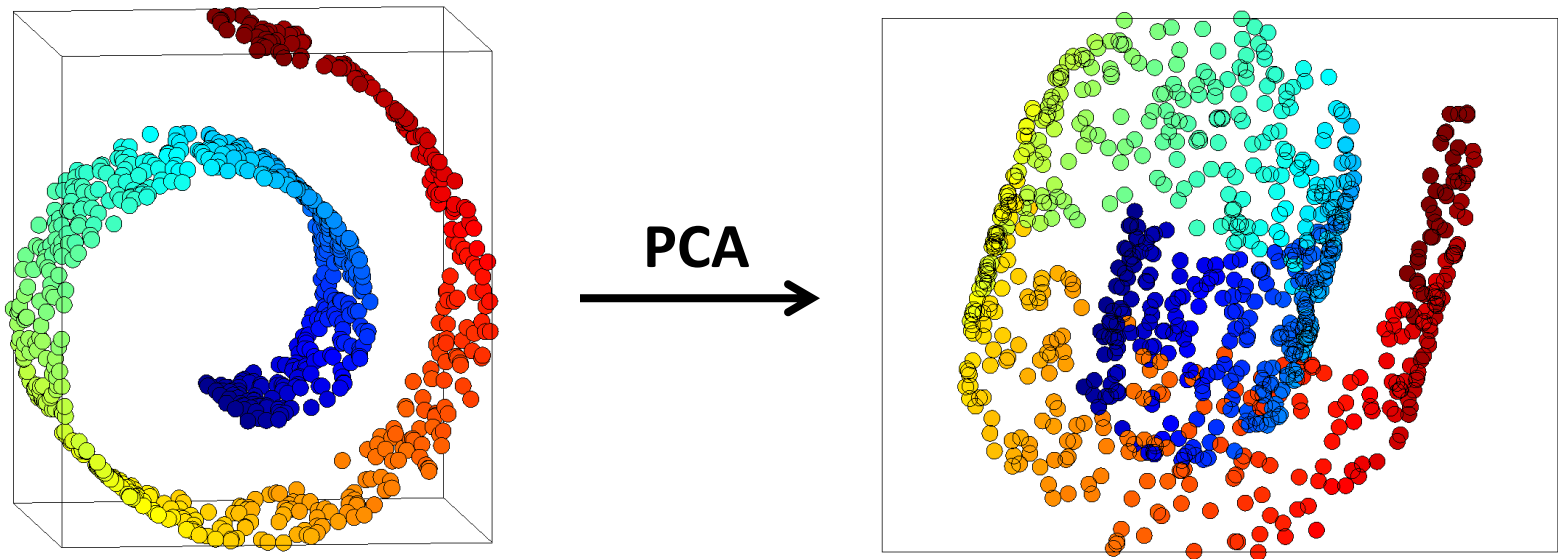
- Thus, the projection vectors that maximize the SNR are derived using an OPCA formulation

- These vectors will span a speaker-independent subspace

Non-linear dimensionality reduction

PCA, LDA and their variants perform a global transformation of the data (rotation/translation/rescaling)

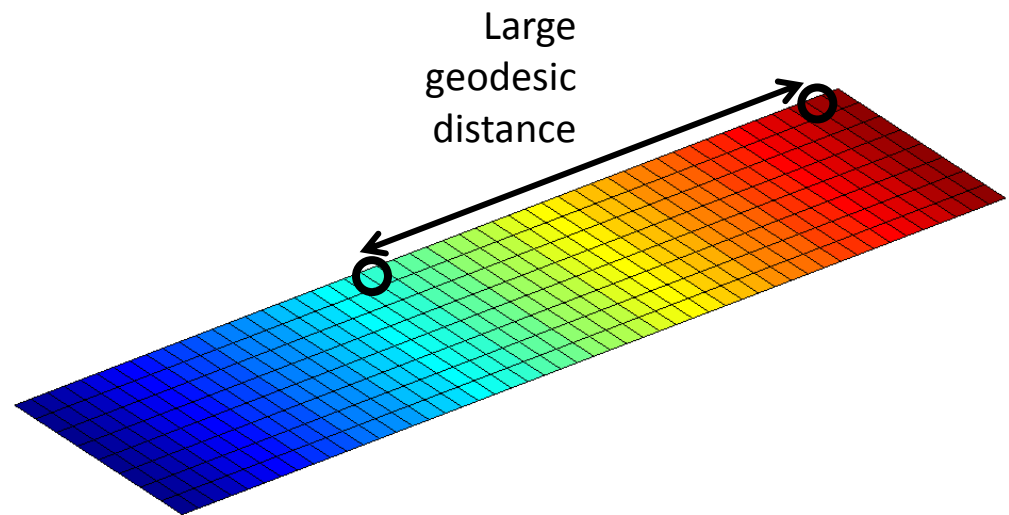
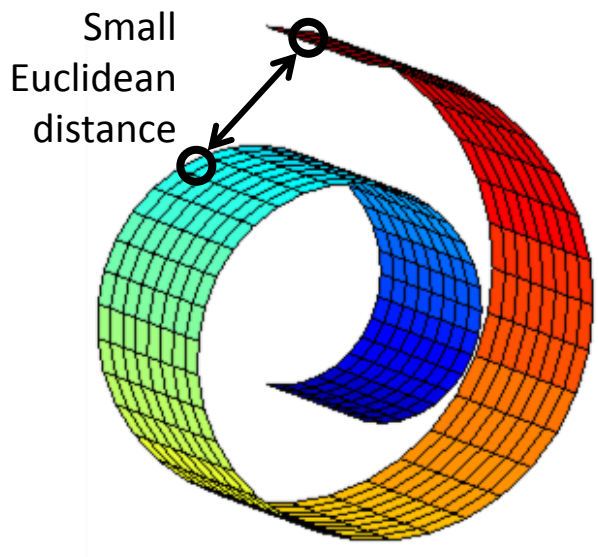
- These techniques assume that most of the information in the data is contained in a linear subspace
- What do we do when the data is actually embedded in a non-linear subspace (i.e., a low-dimensional manifold)?



From <http://www.cs.unc.edu/Courses/comp290-090-s06>

One possible solution (ISOMAP)

- To find the embedding manifold, find a transformation that preserves the geodesic distances between points in the high-dimensional space
 - This approach is related to multidimensional scaling (e.g., Sammon's mapping; L10), except for MDS seeks to preserve the Euclidean distance in the high-dimensional space
- The issue is how to compute geodesic distances from sample data



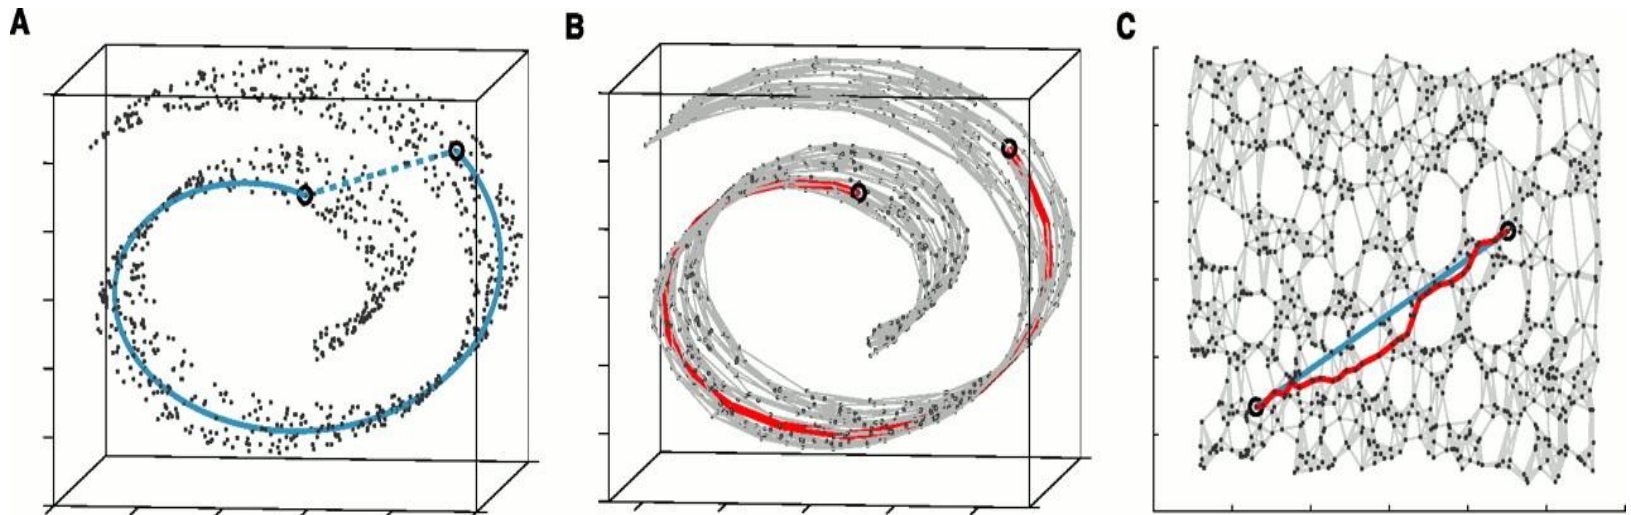
ISOMAP

ISOMAP [Tenenbaum et al., 2000] is based on two simple ideas

- For neighboring samples, Euclidean distance provides a good approximation to geodesic distance
- For distant points, geodesic distance can be approximated with a sequence of steps between clusters of neighboring points

ISOMAP operates in three steps

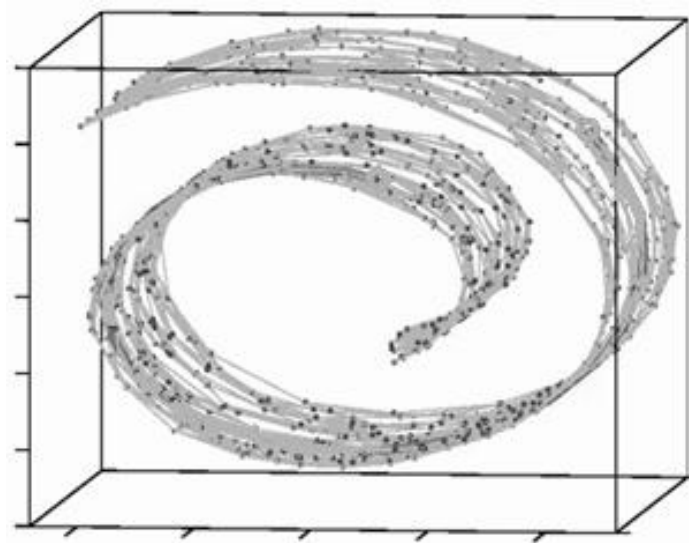
- Find nearest neighbors to each sample
- Find shortest paths (e.g., Dijkstra)
- Apply MDS



[Tenenbaum et al., 1997]

STEP 1

- Determine which points are neighbors in the manifold, based on the distances (x, y) in the input space
- This can be performed in two different ways
 - Connect each point to all points within a fixed radius
 - Connect each point to all of its K nearest neighbors
- These neighborhood relations are represented as a weighted graph G , each edge of weight (x, y) between neighboring points
- Result:



STEP 2

- Estimate the geodesic distances d_{ij} between all pair of points on the manifold by computing their shortest path distances d_{ij}^G in the graph G
- This can be performed, e.g., using Dijkstra's algorithm

STEP 3

- Find d -dim embedding \mathbf{X} that best preserves the manifold's estimated distances
 - In other words, apply classical MDS to the matrix of graph distances $D = \{d_{ij}^G\}$
- The coordinate vectors \mathbf{x}_i are chosen to minimize the following cost function
$$J(\mathbf{X}) = \sum_{i,j} \| \mathbf{x}_i - \mathbf{x}_j \|^2$$
- where D denotes the matrix of Euclidean distances $\{d_{ij} = \| \mathbf{x}_i - \mathbf{x}_j \|\}$, and operator B converts distances to inner products
$$B_{ij} = -\frac{d_{ij}^2}{2}$$
- where D^2 is the matrix of squared distances $\{d_{ij}^2 = d_{ij}^2\}$, and C is a centering matrix, defined as
$$C_{ij} = \frac{1}{n} \delta_{ij} \quad ; \quad C = \frac{1}{n} \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & 1 & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & 1 \end{bmatrix}$$
- It can be shown that the global minimum of $J(\mathbf{X})$ is obtained by setting the coordinates \mathbf{x}_i to the top d eigenvectors of the matrix B

(1) Construct neighborhood graph

(a) Define graph by connecting points i and j if they are [as measured by $d(i, j)$] closer than epsilon (*epsilon-Isomap*), or if i is one of the K nearest neighbors of j (*K-Isomap*).

(b) Set edge lengths equal to $d(i, j)$

(2) Compute shortest paths

(a) Initialize

$d(i, j) = d(i, j)$ if i, j are linked by an edge;
 $d(i, j) = \infty$ otherwise.

(b) For $i = 1, 2, \dots, n$, replace all entries $d(j, k)$ by $\min\{d(j, i), d(i, k) + d(i, j)\}$

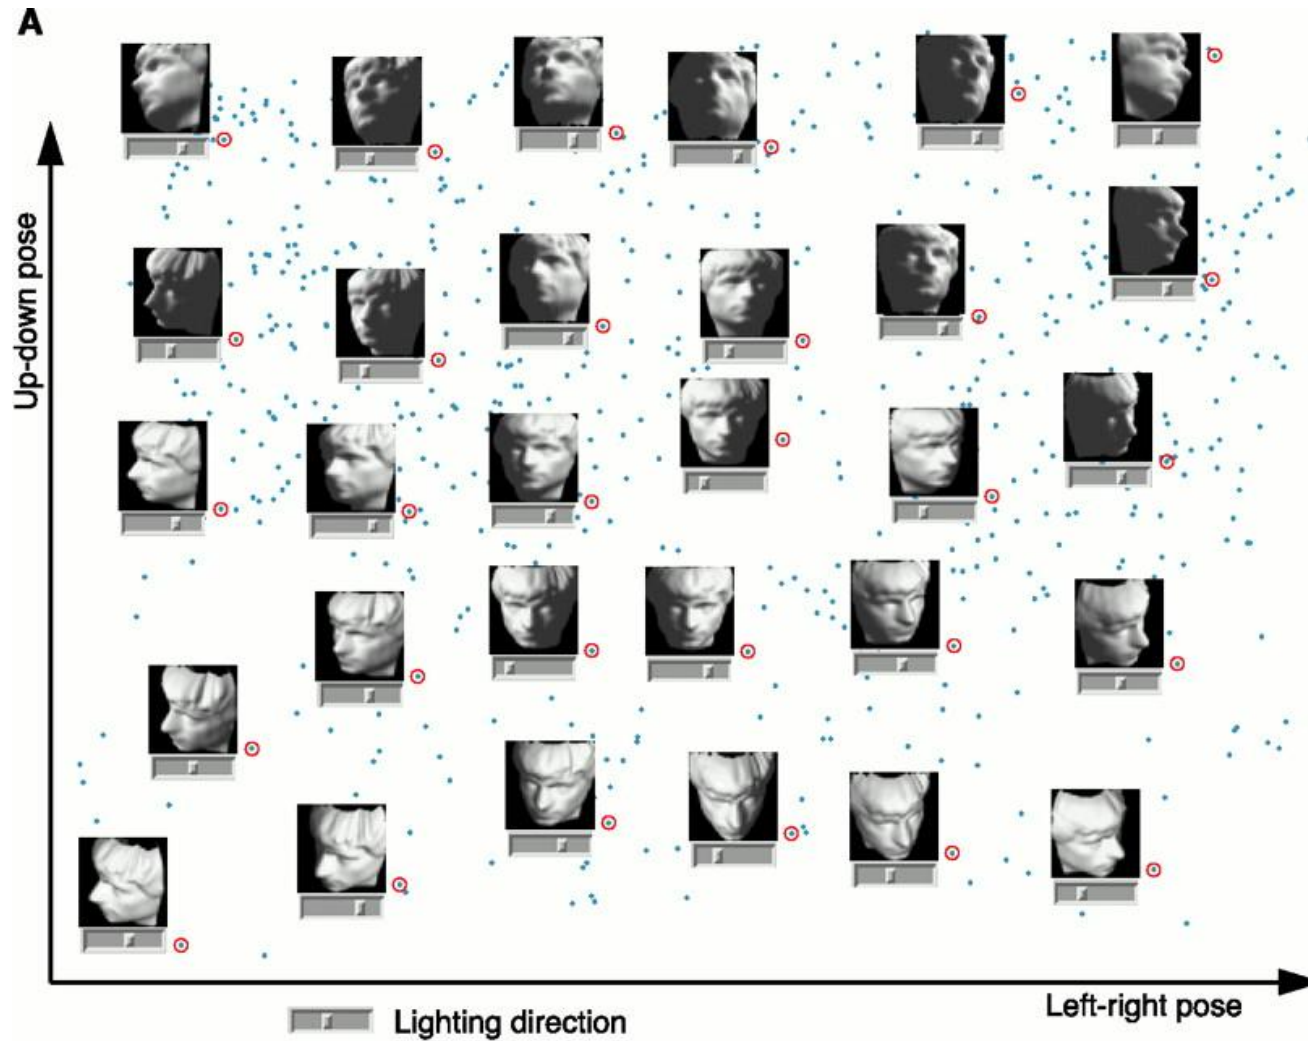
(c) Matrix $D = \{d(i, j)\}$ will contain the shortest path distances between all pairs of points in

(3) Construct d-dimensional embedding

(a) Let λ_i be the i th eigenvalue (in decreasing order) of matrix D , and v_i be the i th component of the i th eigenvector

(b) Set the i th component of the d -dimensional coordinate vector equal to $\sqrt{\lambda_i} v_i$

ISOMAP results



[Tenenbaum et al., 1997]

ISOMAP discussion

ISOMAP is guaranteed asymptotically to recover the true dimensionality and geometric structure of a certain class of Euclidean manifolds

- These are manifolds whose intrinsic geometry is that of a convex region of Euclidean space, but whose geometry in the high-dimensional space may be highly folded, twisted or curved
 - Intuitively, in 2D, these include any physical transformations one can perform on a sheet of paper without introducing tears, holes, or self-intersections
- For non-Euclidean manifolds (hemispheres or tori), ISOMAP will still provide a globally-optimal low-dimensional representation

These guarantees are based on the fact that, as the number of samples increases, the graph distances (d_{ij}) provide increasingly better approximations of the intrinsic geodesic distances (g_{ij})

- However, this proof has limited application [Carreira-Perpiñán, in press] because
 - Data in high-dimensions is scarce (!) and
 - Computational complexity would preclude the use of large datasets anyway

Mapping

- Note that ISOMAP does not provide a mapping function $f = (\cdot)$
- One could however be learned from the pairs $\{ \cdot, \cdot \}$ in a supervised fashion

Locally Linear Embedding

LLE uses a different strategy than ISOMAP to recover the global non-linear structure

- ISOMAP estimates pairwise geodesic distances for all points in the manifold
- Instead, LLE uses only distances within locally linear neighborhoods

Intuition

- Assume that the dataset consists of n vectors x_i , each having d dimensions, sampled from an underlying manifold
- Provided that there is sufficient data (i.e., the manifold is well sampled) one can expect that each point and its neighbors will lie on a linear patch of the manifold

Approach

- LLE solves the problem in two stages
 - First, compute a linear approximation of each sample in the original space
 - Second, find the coordinates in the manifold that are consistent with the previous linear approximation

This discussion is borrowed from [Roweis and Saul, 2000]

Algorithm (part 1)

- The local geometry of these patches is modeled by linear weights that reconstruct each data point as a linear combination of its neighbors
- Reconstruction errors are measured with the L2 cost function

$$(\text{error}) = \sum_{i=1} \left| \text{data point} - \sum_{j=1}^n \text{weights}_j \cdot \text{neighbors}_j \right|^2$$

- where weights measure the contribution of the neighbors to the reconstruction of the data point
- Weights are minimized subject to two constraints
 - Each data point is reconstructed only from its neighbors
 - Rows of the weight matrix sum up to one: $\sum_{j=1}^n \text{weights}_j = 1$
- These constraints ensure that, for any particular sample, the weights are invariant to translation, rotation or scaling
- This problem can be solved using least-squares (details next)

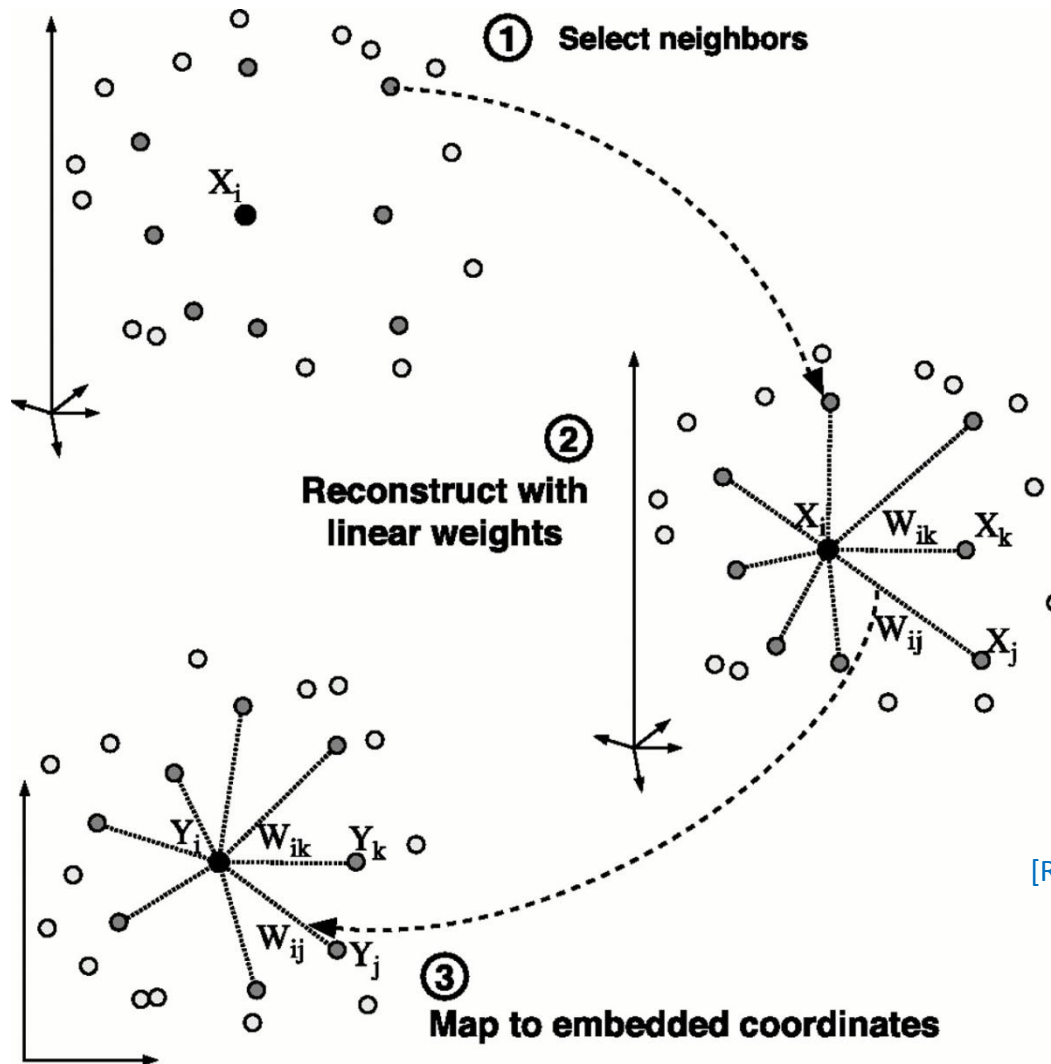
Algorithm (part 2)

- Given that these weights reflect intrinsic properties of the local geometry, we expect that they will also be valid for local patches in the manifold
- Therefore, we seek to find d -dimensional coordinates \mathbf{z} that minimize the following cost function

$$J(\mathbf{z}) = \left\| \mathbf{z} - \sum_i w_i \mathbf{x}_i \right\|^2$$

- Note that this function is similar to the previous one, except that here \mathbf{x}_i are fixed and we solve for \mathbf{z}
- This minimization problem can be solved as a sparse $n \times n$ eigenvalue problem, whose bottom d non-zero eigenvectors are the orthogonal coordinates of the data in the manifold (details follow)

LLE summary



[Roweis and Saul, 2000]

LLE: solving for W_{ij}

The linear weights can be solved as a constrained LS problem

- Consider a particular sample x with K nearest neighbors $\{x_i\}_{i=1}^K$ and reconstruction weights w_i (that sum up to one)
- These weights can be found in three steps:
- **STEP 1:** Evaluate inner product between neighbors to compute the neighborhood correlation matrix C and its inverse C^{-1}

$$C = \begin{bmatrix} x_1 \cdot x_1 & \dots & x_1 \cdot x_K \\ \vdots & \ddots & \vdots \\ x_K \cdot x_1 & \dots & x_K \cdot x_K \end{bmatrix}$$

- **STEP 2:** Compute Lagrange multiplier λ that enforces the constraint $\sum w_i = 1$

$$\lambda = \frac{1 - C^{-1} \mathbf{1}}{-\mathbf{1}^T C^{-1} \mathbf{1}}$$

- **STEP 3:** Compute the reconstruction weights as

$$w_i = C^{-1} (x_i \cdot x + \lambda)$$

- NOTE: If C is singular, regularize by adding a small multiple of the identity matrix
- Since this 3-step process requires matrix inversion, a more efficient solution is to
 - First, solve the linear system of equations $C w = x - x$, and
 - Then, rescale the weights so they sum up to one (which yields the same results)

[Roweis and Saul, 2000; Saul and Roweis 2001]

LLE: solving for Y_i

The embedding vectors are found by minimizing the cost function

$$J(Y) = \sum_i \|x_i - Y y_i\|^2$$

- To make the optimization problem well-posed we introduce two constraints
 - Since coordinates can be translated without affecting the cost function, we remove this degree of freedom by imposing that they are centered

$$\sum_i y_i = 0$$

- To avoid degenerate solutions, we constraint the embedding vectors to have unit covariance matrix

$$\frac{1}{n} \sum_i y_i y_i^T = I$$

- This allows the cost function to be expressed in a quadratic form involving inner products of the embedding vectors

$$J(Y) = \sum_i (x_i - Y y_i)^T (x_i - Y y_i)$$

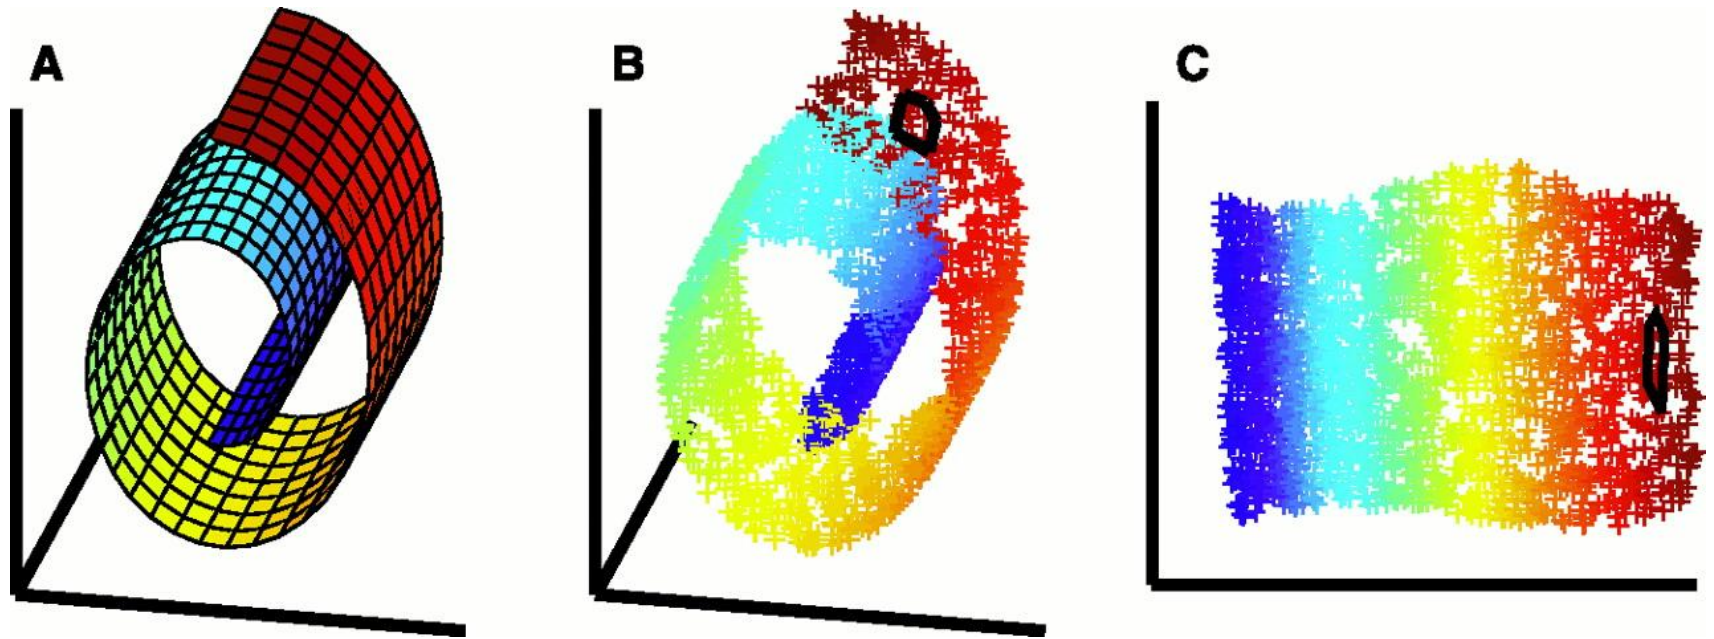
- where

$$= \sum_i x_i^T x_i - 2 \sum_i x_i^T Y y_i + \sum_i y_i^T Y^T Y y_i$$

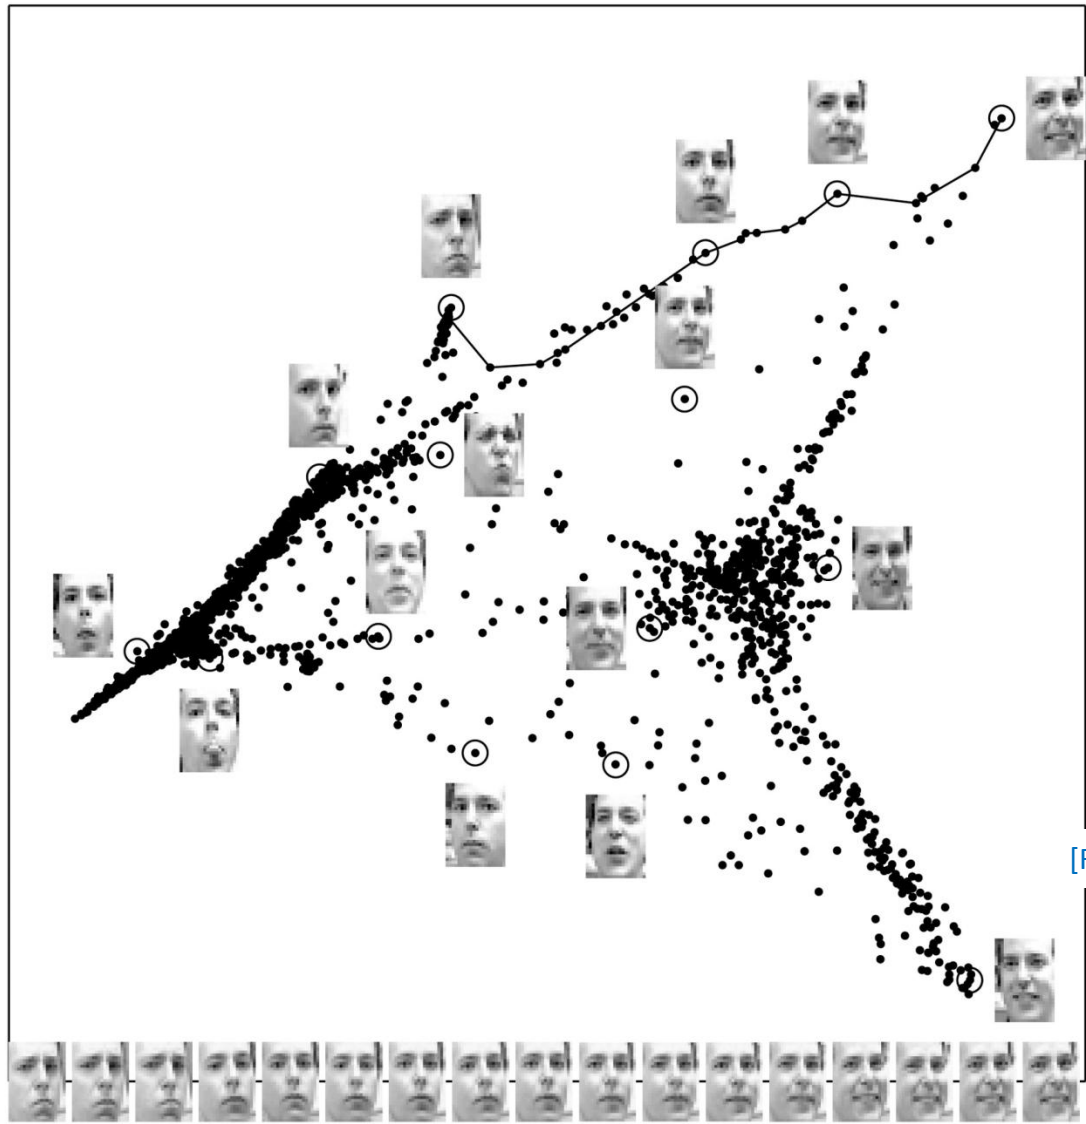
It can be shown that the optimal embedding is found by computing the bottom $k - 1$ eigenvectors of matrix

- We discard the bottom eigenvector, which is the unit vector
 - This eigenvector represents a free translation mode with eigenvalue zero
 - Discarding this eigenvector enforces the constraint that the embedding coordinates have zero mean
- The remaining $k - 1$ eigenvectors form the $k - 1$ embedding coordinates found by LLE

LLE examples



[Roweis and Saul, 2000]



[Roweis and Saul, 2000]

LLE discussion

LLE is simple and attractive, but shares some of the limitations of MDS methods

- Sensitivity to noise
- Sensitivity to non-uniform sampling of the manifold
- Does not provide a mapping (though one can be learned in a supervised fashion from the pairs $\{x_i, y_i\}$)
- Quadratic complexity on the training set size
- Unlike ISOMAP, no robust method to compute the intrinsic dimensionality, and
- No robust method to define the neighborhood size K