

**94-775/95-865 Lecture 4:
Manifold learning**


George Chen

PCA reorients data so axes explain variance in “decreasing order”
→ can “flatten” (*project*) data onto a few axes that captures most variance



Image source: http://4.bp.blogspot.com/-USQEgoh1jCU/VfncdNOETcl/AAAAAAAAAGp8/Hea8UtE_1c0/s1600/Blog%2B1%2BIMG_1821.jpg

2D Swiss Roll




PCA would just flatten this thing and
*lose the information that the data actually
lives on a 1D line that has been curved!*




PCA would squash down this Swiss roll (like stepping on it from the top) mixing the red & white parts

Image source: http://4.bp.blogspot.com/-USQEgoh1jCU/VfncdNOETcl/AAAAAAAAAGp8/Hea8UtE_1c0/s1600/Blog%2B1%2BIMG_1821.jpg


2D Swiss Roll




2D Swiss Roll




2D Swiss Roll



2D Swiss Roll



2D Swiss Roll




2D Swiss Roll



This is the desired result

Manifold Learning

- Nonlinear dimensionality reduction (in contrast to PCA which is linear)
- Find low-dimensional “manifold” that the data live on



Basic idea of a manifold:

1. Zoom in on any point (say, x)
2. The points near x look like they're in a lower-dimensional Euclidean space (e.g., a 2D plane in Swiss roll)

Do Data Actually Live on Manifolds?





Image source: <http://www.columbia.edu/~jwp2128/Images/faces.jpeg>

Do Data Actually Live on Manifolds?



Image source: <http://www.adityathakker.com/wp-content/uploads/2017/06/word-embeddings-994x675.png>


Do Data Actually Live on Manifolds?



Mnih, Volodymyr, et al. Human-level control through deep reinforcement learning. Nature 2015.

Manifold Learning with Isomap

Step 1: For each point, find its nearest neighbors, and build a road (“edge”) between them




(e.g., find closest 2 neighbors per point and add edges to them)

Step 2: Compute shortest distance from each point to every other point *where you're only allowed to travel on the roads*

Step 3: It turns out that given all the distances between pairs of points, we can compute what the points should be (the algorithm for this is called *multidimensional scaling*)

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A					
B					
C					
D					
E					

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0				
B		0			
C			0		
D				0	
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5			
B		0	5		
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8		
B		0	5		
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	
B		0	5		
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B		0	5		
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B		0	5	10	
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B		0	5	10	13
C			0	5	
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B		0	5	10	13
C			0	5	8
D				0	5
E					0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D


Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B	5	0	5	10	13
C	8	5	0	5	8
D	13	10	5	0	5
E	16	13	8	5	0

Isomap Calculation Example

In orange: road lengths



2 nearest neighbors of A: B, C

2 nearest neighbors of B: A, C

2 nearest neighbors of C: B, D

2 nearest neighbors of D: C, E

2 nearest neighbors of E: C, D

Build "symmetric 2-NN" graph
(add edges for each point to
its 2 nearest neighbors)

Shortest distances between
every point to every other
point *where we are only
allowed to travel along the
roads*

	A	B	C	D	E
A	0	5	8	13	16
B	5	0	5	10	13
C	8	5	0	5	8
D	13	10	5	0	5
E	16	13	8	5	0


This matrix gets fed into
multidimensional scaling to get
1D version of A, B, C, D, E

The solution is not unique!

Isomap Calculation Example

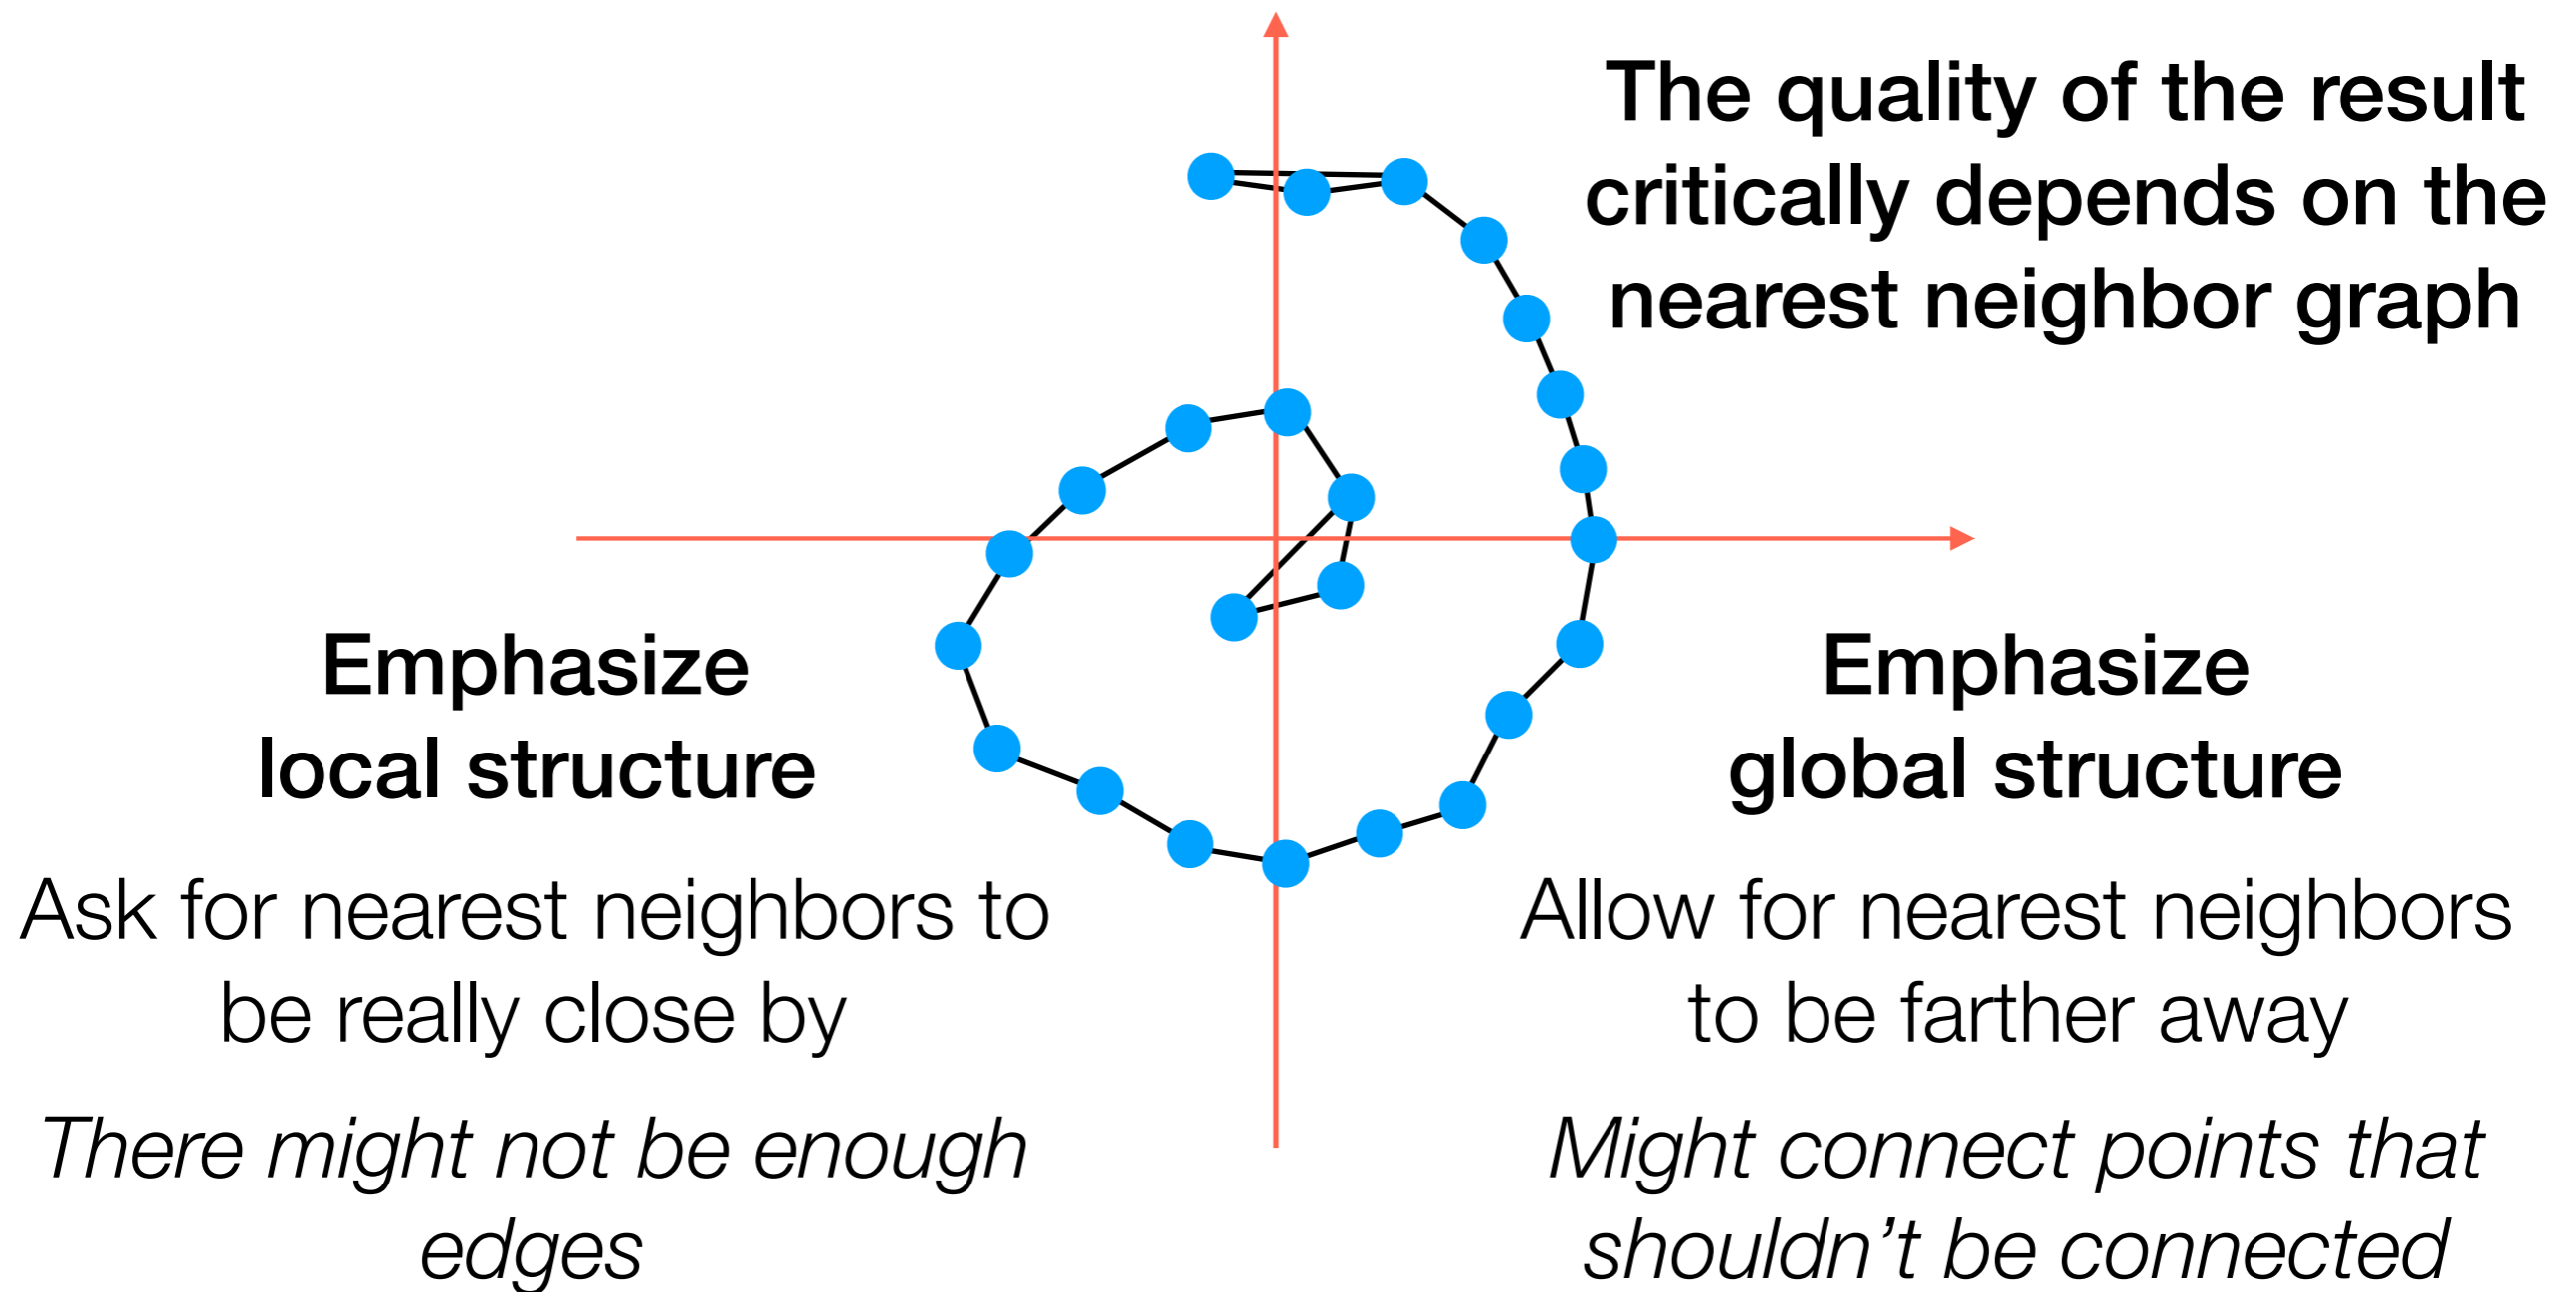
Demo

3D Swiss Roll Example



Joshua B. Tenenbaum, Vin de Silva, John C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science 2000.

Some Observations on Isomap




In general: try different parameters for nearest neighbor graph construction when using Isomap + visualize

t-SNE

**(t-distributed stochastic
neighbor embedding)**


t-SNE High-Level Idea #1

- Don't use deterministic definition of which points are neighbors
- Use probabilistic notation instead




t-SNE High-Level Idea #2

- In low-dim. space (e.g., 1D), suppose we just randomly assigned coordinates as a candidate for a low-dimensional representation for A, B, C, D, E (I'll denote them with primes):





- With any such candidate choice, we can define a probability distribution for these low-dimensional points being similar




t-SNE High-Level Idea #3

- Keep improving low-dimensional representation to make the following two distributions look as closely alike as possible



t-SNE



Also: play with learning rate, # iterations

In practice, often people initialize with PCA

There are some other parameters (less critical)

Manifold Learning with t-SNE

Demo

t-SNE Interpretation

<https://distill.pub/2016/misread-tsne/>

Dimensionality Reduction for Visualization

- There are *many* methods (I've posted a link on the course webpage to a scikit-learn example using ~10 methods)
- PCA is very well-understood; the new axes can be interpreted
- Nonlinear dimensionality reduction: new axes may not really be all that interpretable (you can scale axes, shift all points, etc)
- PCA and t-SNE are good candidates for methods to try first
- If you have good reason to believe that only certain features matter, of course you could restrict your analysis to those!