



Learning and Using Image Manifolds

Robert Pless


Associate Professor of
Computer Science






What are image manifolds?

- Sets of images that locally have only a few degrees of freedom.
- For example, “8” s from the NIST character database






Example 2



What is manifold learning?

- Given *unorganized* images from low-dimensional manifold, assign each image a low-dimensional coordinate.




Overloading of “Manifold Learning”


1. Given data that comes from a (very) low dimensional manifold, give each data point parameters that reflect relative positions on that manifold.
2. Algorithms for (statistical) learning, when you know your data lies on some (perhaps not low dimensional) manifold in the underlying space.

my talk is about 1, but both are represented in later talks tpdya.

PCA: Learning linear manifolds (on video data)





M \approx **B** **C**




$$M \approx B C$$

Principal Component Analysis



Basis images



Coordinates for each image

Similarity Based Image Analysis

Instead of projection, start with *similarity* measure between images.

Use many images rather than many features in each image.




Image distances to low-dimensional locations

Suppose we had a distance between every pair of images.



Tool: Multi-dimensional scaling

Input: all pairwise distances D .

Output: set of point positions X whose pairwise distances match D .

$$D = \begin{matrix} & 0 & 5 & 3 & 1 & 3 \\ & 5 & 0 & 8 & 2 & 1 \\ & 3 & 8 & 0 & 6 & 4 \\ & 1 & 2 & 6 & 0 & 7 \\ & 3 & 1 & 4 & 7 & 0 \end{matrix}$$

Distance Matrix



Points with that set of pairwise distances.

MDS algorithm:

Squared distance matrix S : $S(i,j) = D(i,j)^2$.

Centering matrix H : $H = I - 1/N$. (identity – uniform matrix of $1/N$)

Dot-product matrix: $t(D) = -HSH/2$

*defined so: $X'X = t(D)$ if for all i,j $(X_i - X_j)'(X_i - X_j) = S(i,j)$,
where X is a matrix whose columns are position vectors in parameter space.*

Consider eigenvalue problem: $X'X = t(D)$

Let λ_p, v_p be the p -th eigenpair of $t(D)$

$$\begin{pmatrix} | & | & | & \dots \\ | & | & | & \dots \\ v_1 & v_2 & v_3 & \dots \\ | & | & | & \dots \\ | & | & | & \dots \end{pmatrix} \begin{pmatrix} \sqrt{\lambda_1} & 0 & 0 & \dots \\ 0 & \sqrt{\lambda_2} & 0 & \dots \\ 0 & 0 & \sqrt{\lambda_3} & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \end{pmatrix} = \begin{pmatrix} X_1 & Y_1 & Z_1 & \dots \\ X_2 & Y_2 & Z_2 & \dots \\ X_3 & Y_3 & Z_3 & \dots \\ \cdot & \cdot & \cdot & \dots \\ \cdot & \cdot & \cdot & \dots \end{pmatrix}$$

Each row is optimal embedding in k -dimensional space, if you use k eigenpairs.

... but image similarity is only meaningful for *small* image distances.

Images which are very similar should be embedded as points which are close to each other.



For images which are not similar, we don't know how close their embedded points should be.



But, we don't believe image similarities for anything but very similar images.

In 2000, two papers presented methods of extending local similarities to give global constraints:

Isomap (Tenenbaum, et al, 2000)

LLE (Roweis and Saul, 2000)

followed by Semi-Definite Embedding, Maximum Variance Unfolding, Diffusion Maps, Laplacian Eigenmaps, Hessian Eigenmaps, Locality Preserving Projections, and others, all of which are techniques for non-linear dimension reduction.

Isomap:

Define $G(V,E)$:


- V is the set of points (in our case, images)
- E is the set of comparable points, (images with differences that are very small)
- $w(e)$ is the image difference.

Algorithm:

- Run all pairs shortest path algorithm on G ,
- Define D , the pairwise distance matrix to be shortest path distance in G .

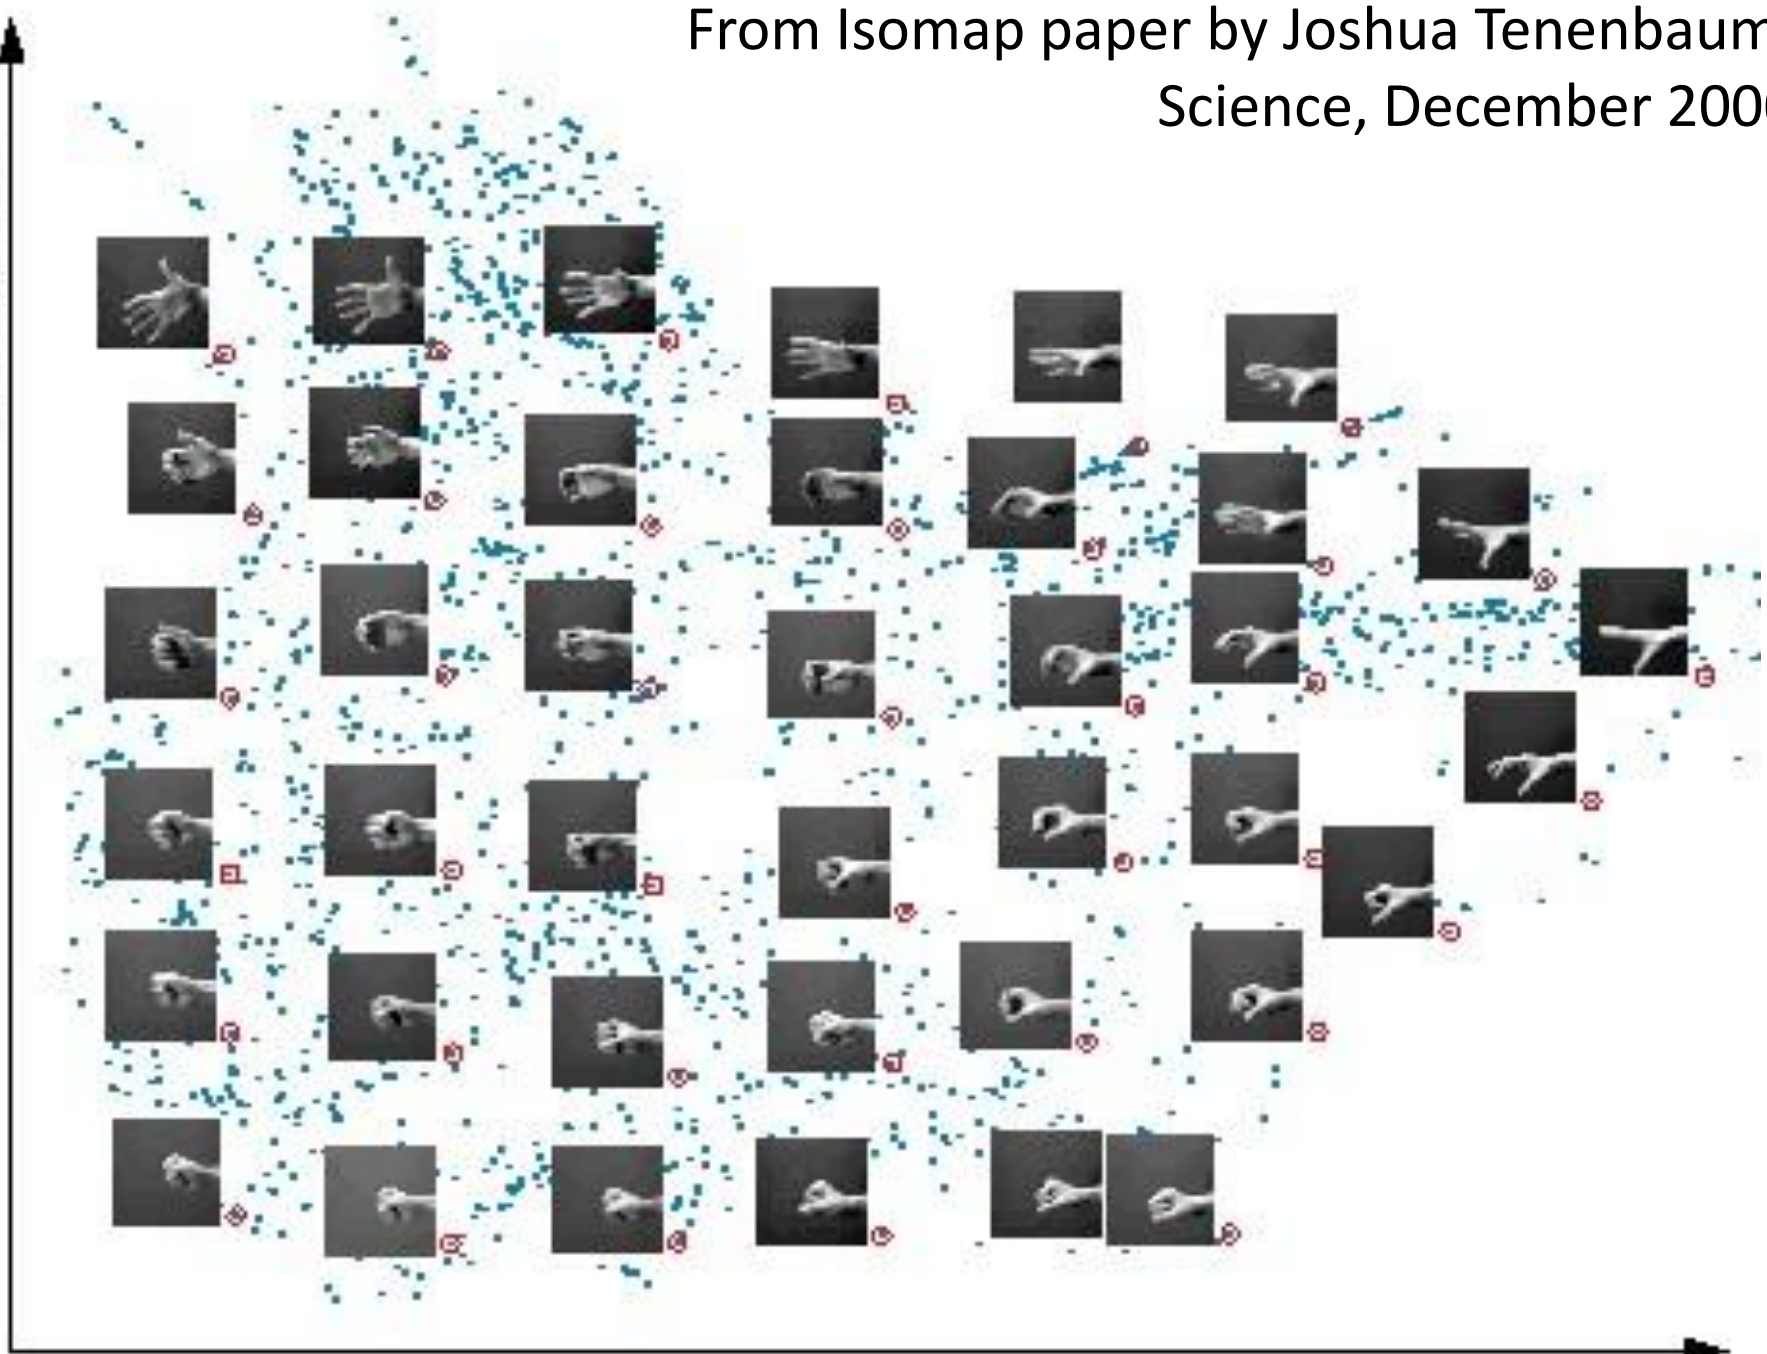
Run MDS, using D as given pairwise distances.

Classic example: Swiss Roll.




From Isomap paper by Joshua Tenenbaum,
Science, December 2000

Anger extension



Wrist rotation



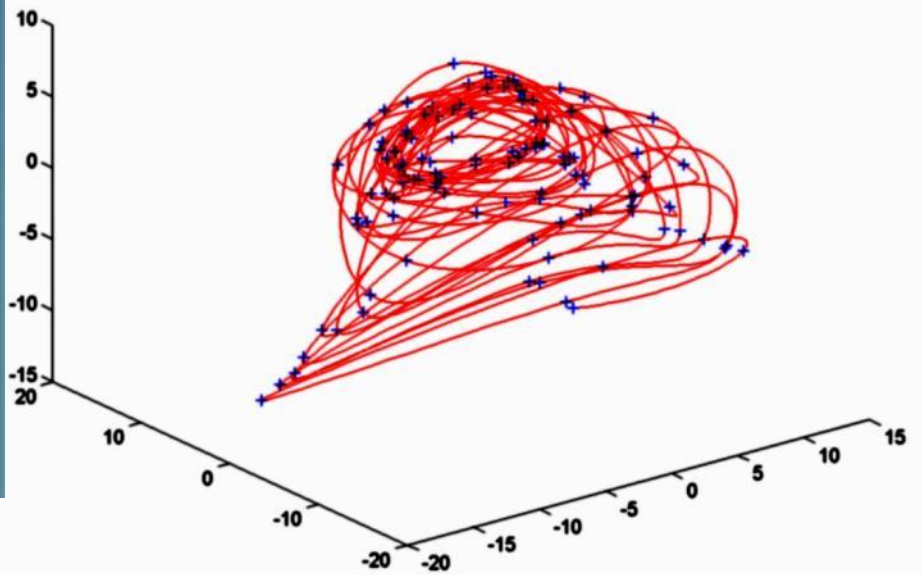

(Isomap – ShortestPath) == PCA

The key to Isomap is the shortest-path distances.


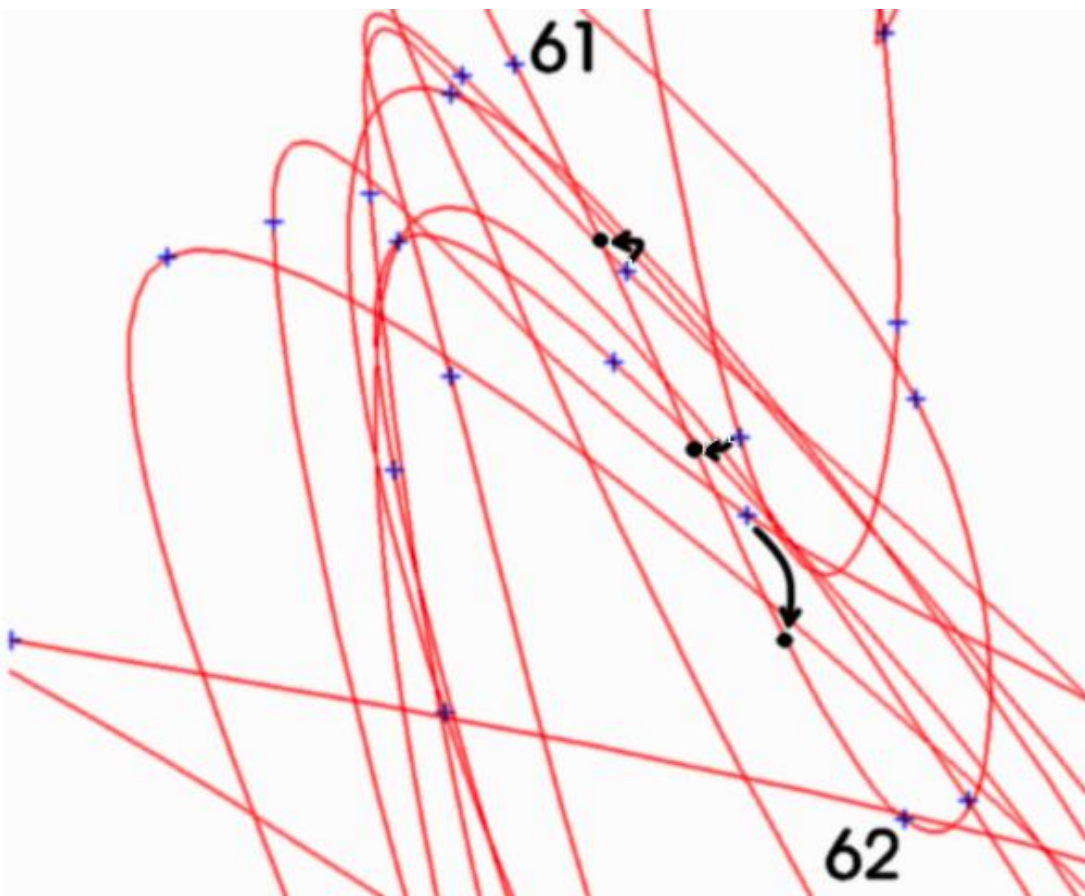
If you run MDS on points with original distances (from high dimensional space), it gives the SAME embedding as PCA, up to Euclidean transformation.

Isomap for video analysis


Example: bird flight




Temporal Super-Resolution




4 x Framerate



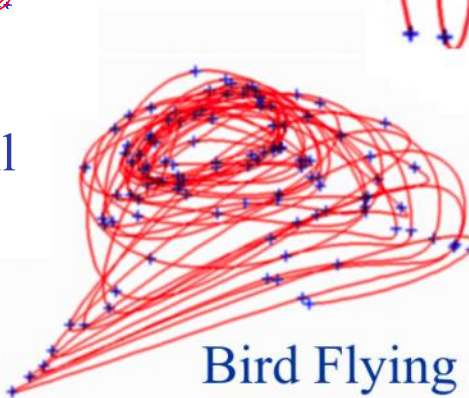
20 x Framerate (different input)



Woman on a treadmill




Bird Flying
then gliding



Bird Flying




Fountain



Video Trajectories

Example: human behaviors




Last example: gait





Applications to medical imagery

- Medical imagery (MRI/CT) is an ideal problem domain for manifold learning.
- Often just a few degrees of freedom:
 - Contrast agent perfusion
 - Viewpoint change
 - Breathing/heartbeat
 - Which may be difficult to measure



Isomap Visualization

Cardiac MRI imagery courtesy of Nikos Tsekos,
Department of Radiology, Washington University Medical School.




Consider distance between Gabor Filter Responses at each pixel.

Complex Gabor response separates motion (phase change) from contrast change (magnitude)





1-D embedding by phase shift



Motion axis


1-D embedding from Gabor response magnitude difference.





[Isomap Visualization](#)

4D CT – alignment of ungated images



Acquire data in 16 slice sections (chunks), in cine mode (25 frames).


Reconstruct 3D lung volume for each breathing phase

4D CT Data Acquisition

Patient lung

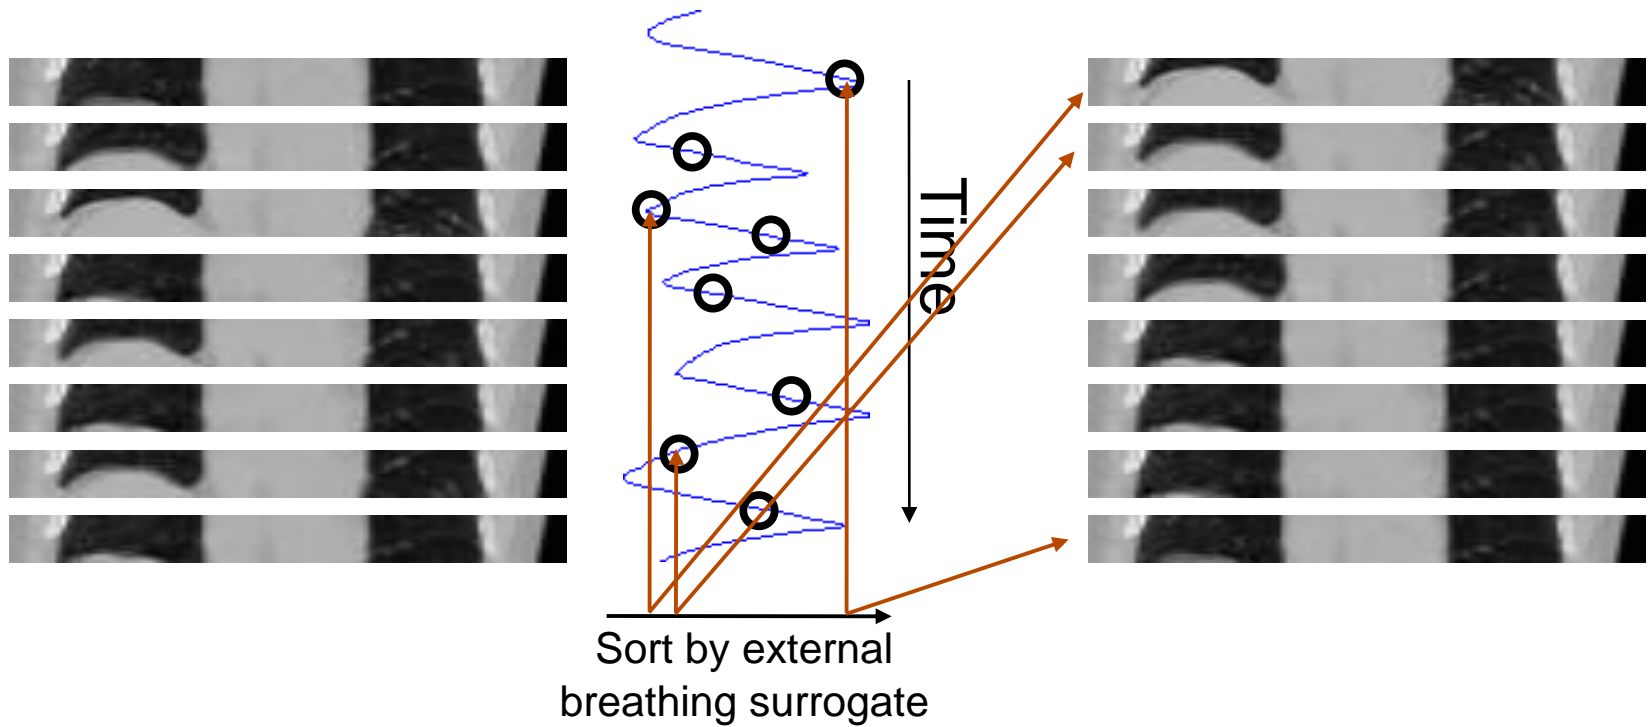


Data acquisition



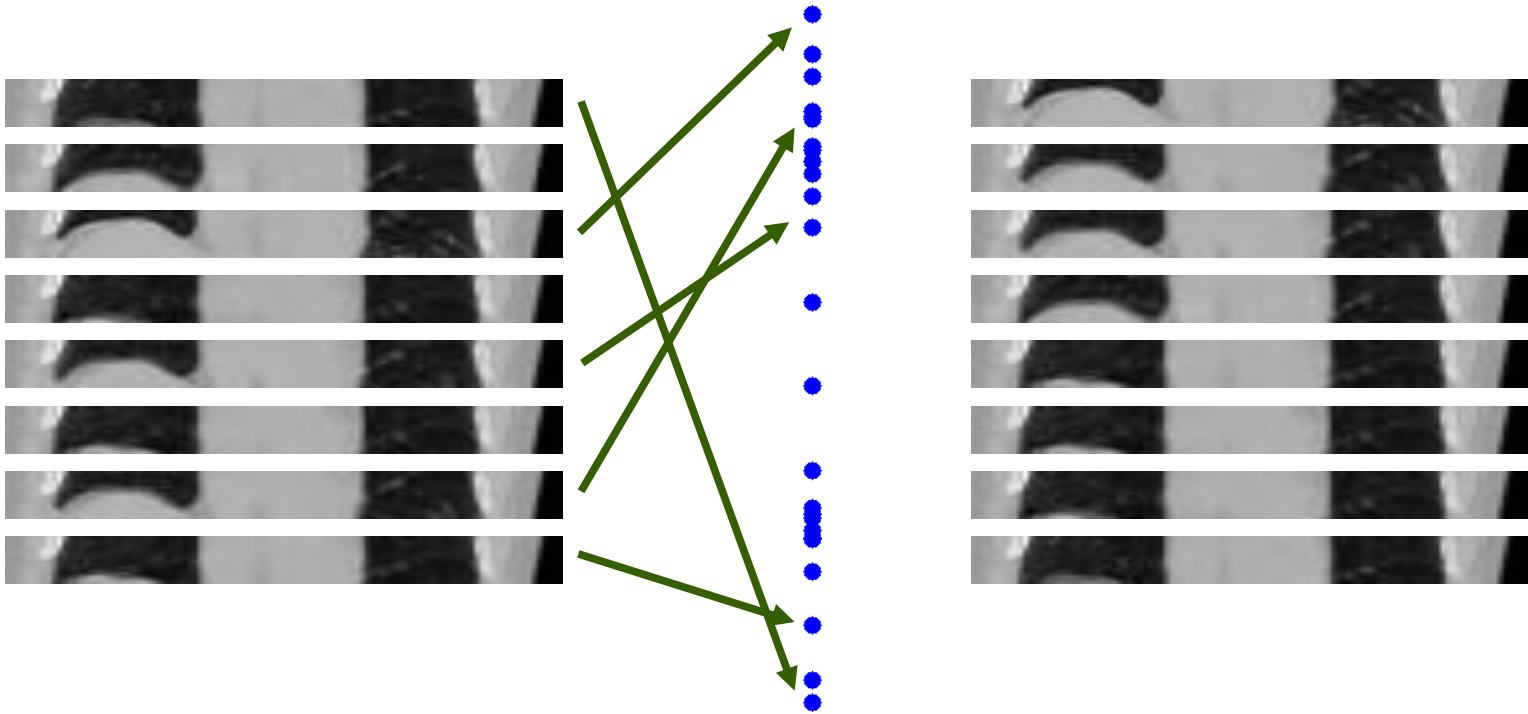
Images essentially unordered


External breath surrogate



1-D manifold

- Sort breath using Isomap of images






Color coded by breath surrogate (belt) measurement.





Solve for affine parameters of one couch position that maximizes smoothness over volume segment boundary to next affine parameter.






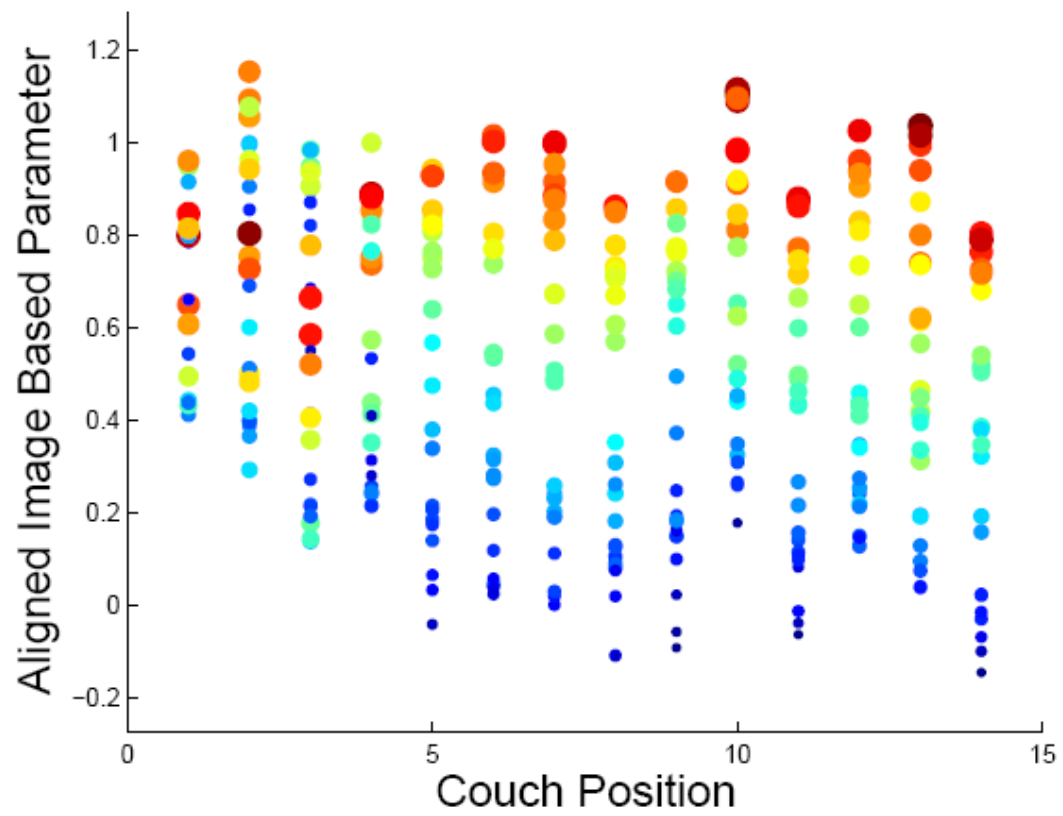





































Limitations





Each chunk needs data in each part of breath


Breath must be present in all images

Top of lung difficult to order


Lots of data needed for secondary variations (heart beat, hysteresis)

“is there more to manifold learning than re-sampling, re-ordering, and de-noising?”





B



breathing



A



heartbeat



Snake
 $C(s)$
→
Single image



$C(s,t)$
→
Time sequence


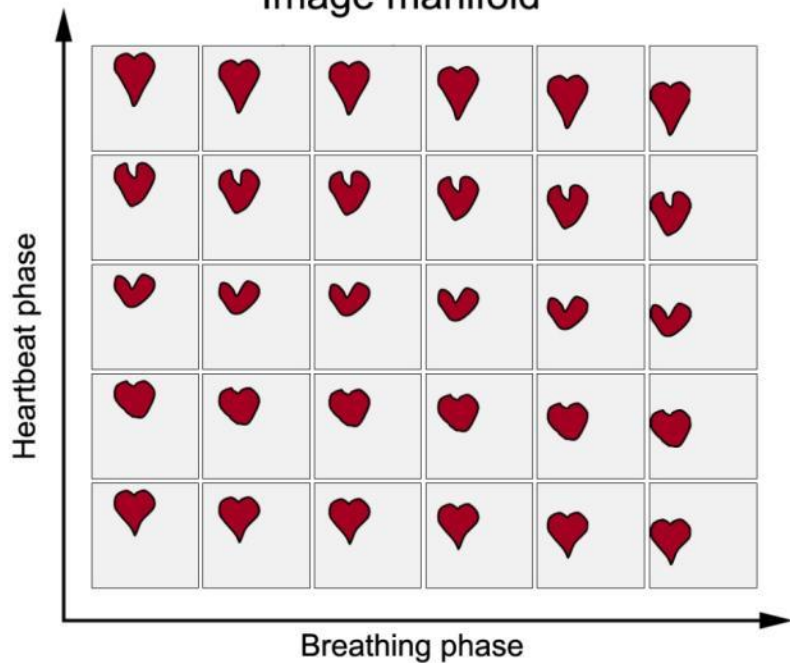
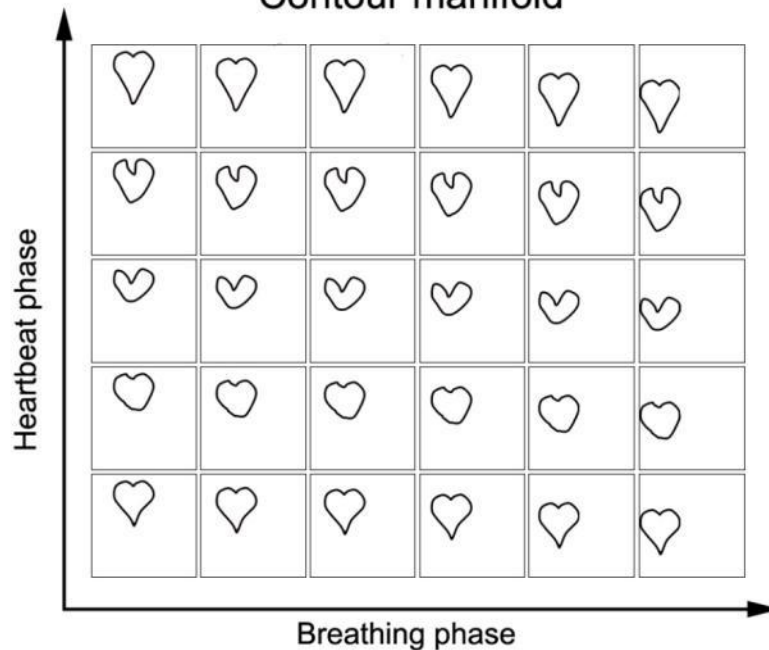


Image manifold

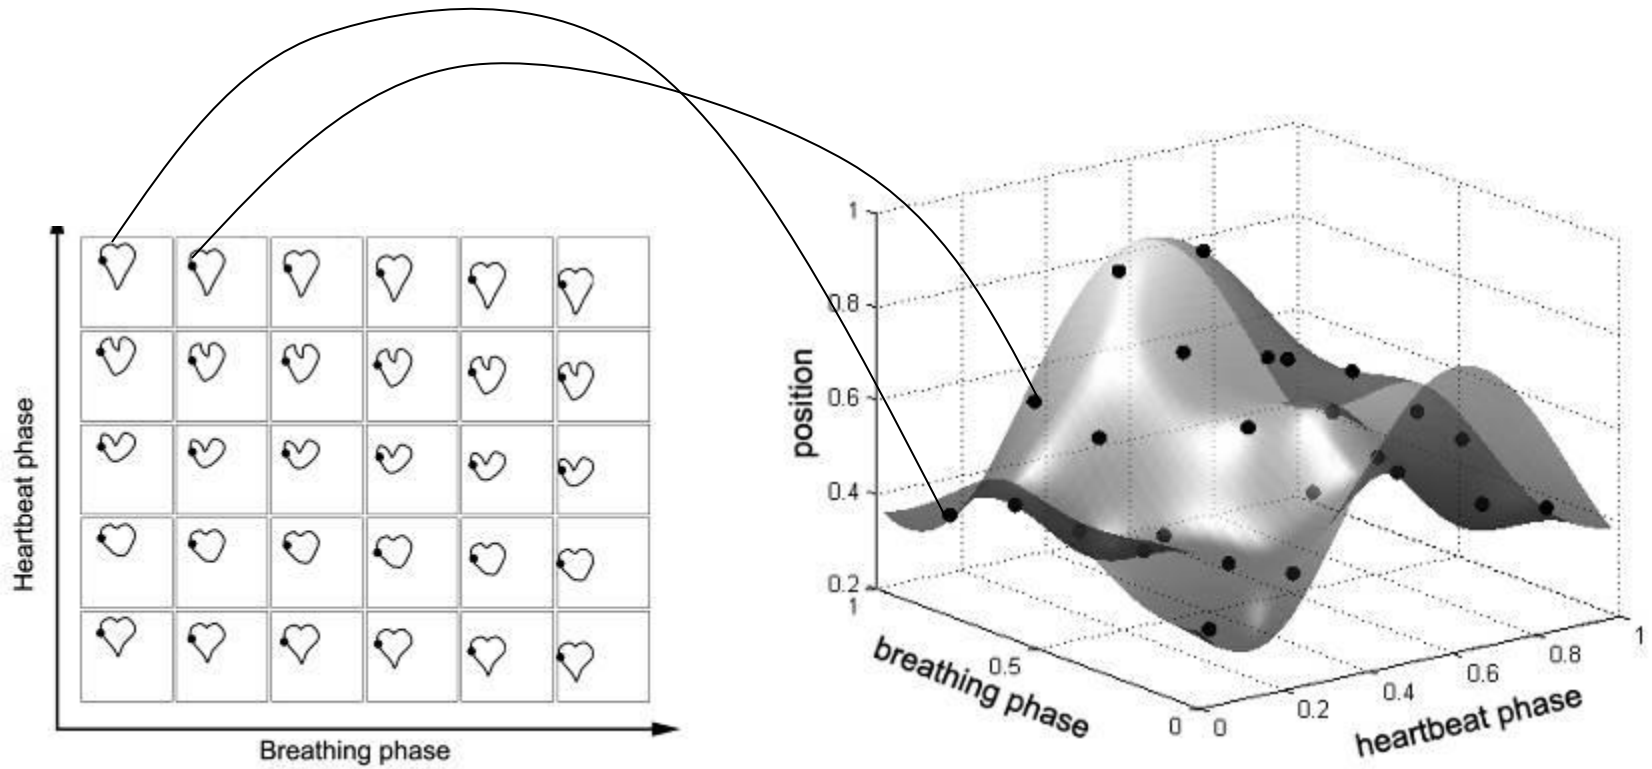


Contour manifold




$C(s,\phi,\theta)$
→

- Use a cubic B-spline surface to specify how each control point varies with ϕ and θ .



$$C(s, \phi, \theta)$$



- Heartbeat phase ϕ , breathing phase θ .
- contour $C(s, \phi, \theta)$

$$E = \oint E_{int}(C(s, \phi, \theta)) + E_{img}(C(s, \phi, \theta)) + E_{con}(C(s, \phi, \theta)) ds d\phi d\theta$$

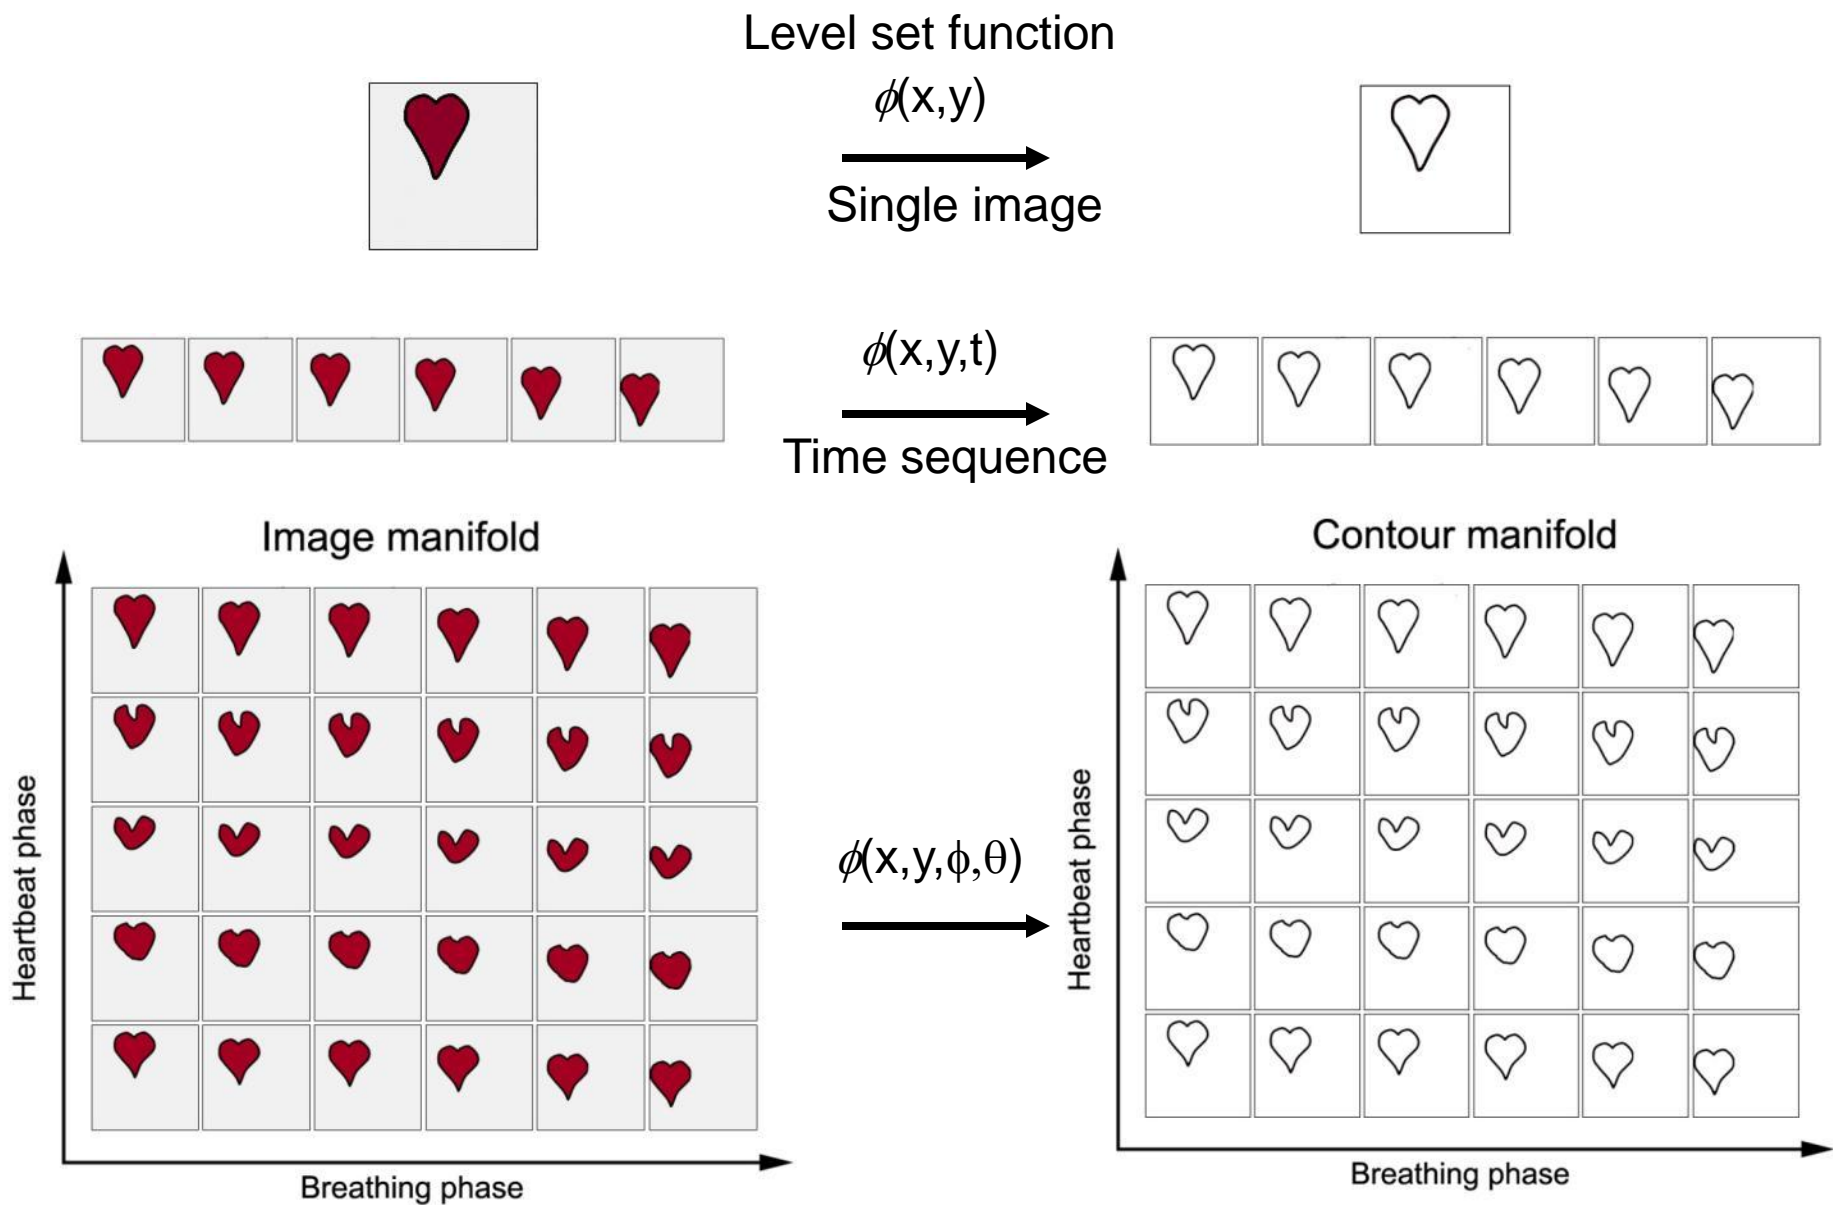


$$E_{int} = \frac{1}{2} \left[\alpha \left| \frac{\partial C}{\partial s} \right|^2 + \beta \left| \frac{\partial^2 C}{\partial s^2} \right|^2 + \mu \left(\left| \frac{\partial C}{\partial \phi} \right|^2 + \left| \frac{\partial C}{\partial \theta} \right|^2 \right) + \gamma \left(\left| \frac{\partial^2 C}{\partial \phi^2} \right|^2 + \left| \frac{\partial^2 C}{\partial \theta^2} \right|^2 \right) \right]$$

$$E = \int E_{int} + E_{img} + \boxed{E_{con}} ds d\phi d\theta$$

$$E_{con} = \begin{array}{l} \text{Term to penalize non-} \\ \text{translational motion} \end{array} + \begin{array}{l} \text{Term to penalize} \\ \text{deformation other than} \\ \text{expansion/contraction} \end{array}$$








4D Level set function



Worst single image
results from best
parameters



Manifold Learning for
Segmentation, Q.
Zhang, R. Souvenir,
R. Pless, EMMCVPR,
2005




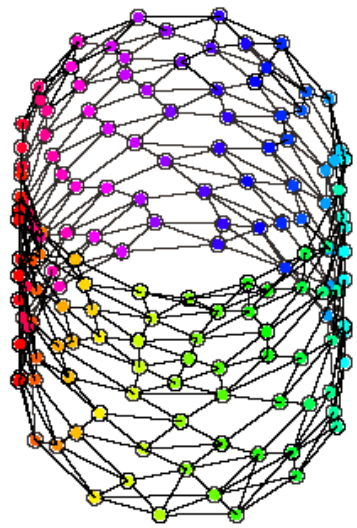







Cardio-Pulmonary Level
Sets, Q. Zhang, R. Souvenir,
R. Pless, CVBIA 2005

Cine-MR segmentation summary

- Manifold learning re-arranges original data frames in order to provide *additional constraints* which improve segmentation
- Resulting optimization problem remains very similar to standard Snakes or Level-Sets.

Conclusions

- Manifold learning is an important tool for the analysis of images that vary due to motion and deformation.
- Especially useful in medical --- images of one patient often form very clean manifolds of just a few dimensions.



Looking forward?

Acknowledgements / Other Interests



Richard Souvenir
UNC-Charlotte




Nathan Jacobs,
Kentucky



Manfred Georg,
Google/ Youtube




Michael Dixon,
Willow Garage



Qilong Zhang,
Nomura


Omni-directional Vision




Surveillance and Tracking



(a) Intersection



(b) \mathcal{N}_1



(c) \mathcal{N}_2

Remote Sensing

