CS233: Geometric and Topological Data Analysis

More on Map Networks,

Course Contents Discussion.

6 June 2018

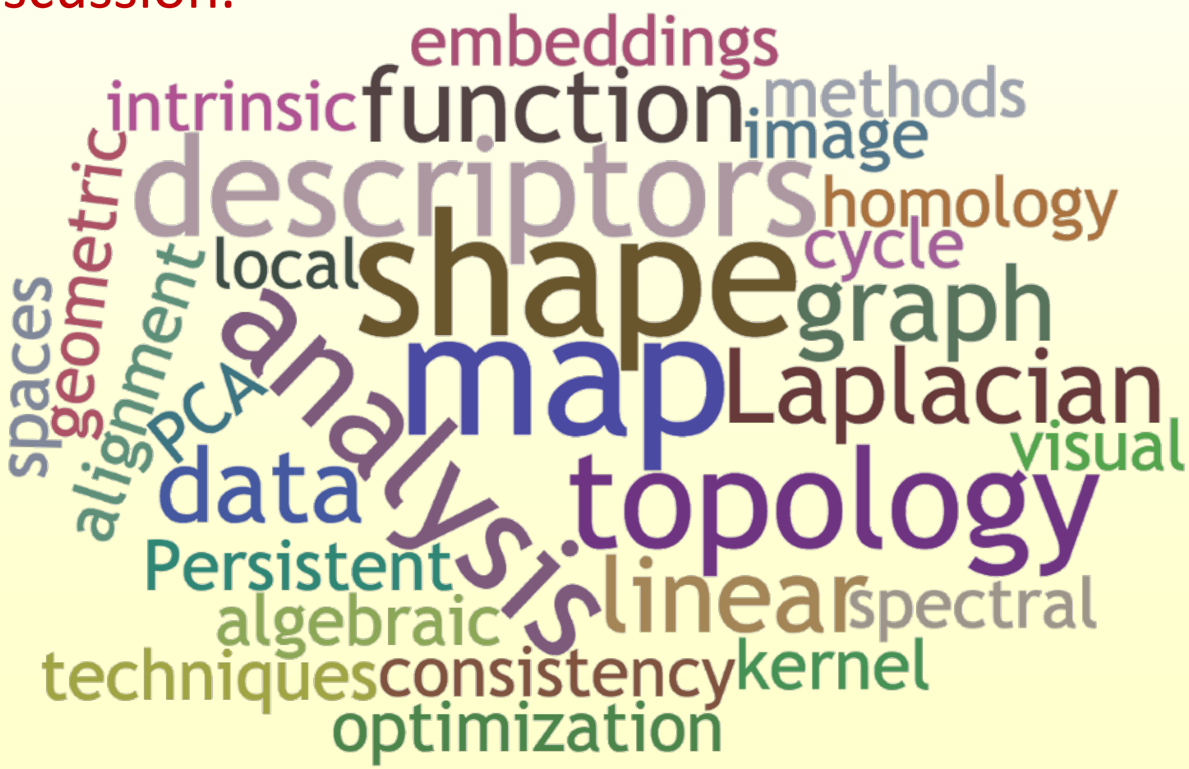

Last Time:
Shape Differences, Map Networks

Area-Based Shape Difference:

$$V = A_M^{-1} F^\top A_N F$$

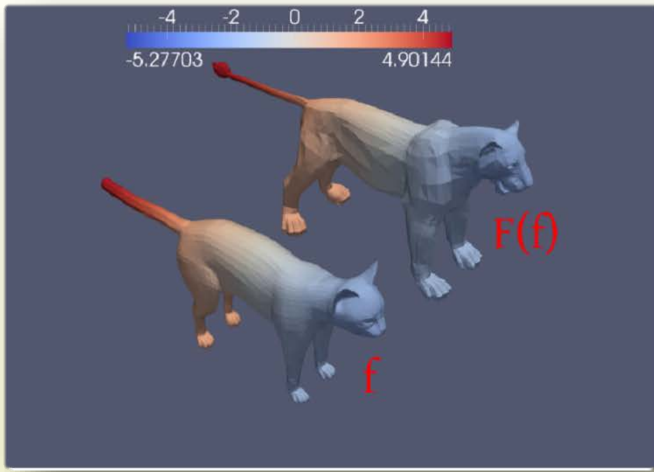

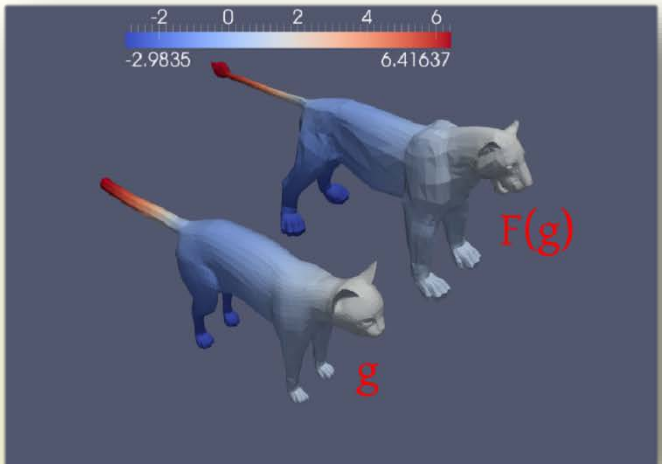

$$\int_{lion} F(f)F(g) \neq \int_{cat} fg$$

$$\downarrow$$

$$\int_{lion} F(f)F(g) = \int_{cat} f V(g)$$

Intrinsic Shape Space

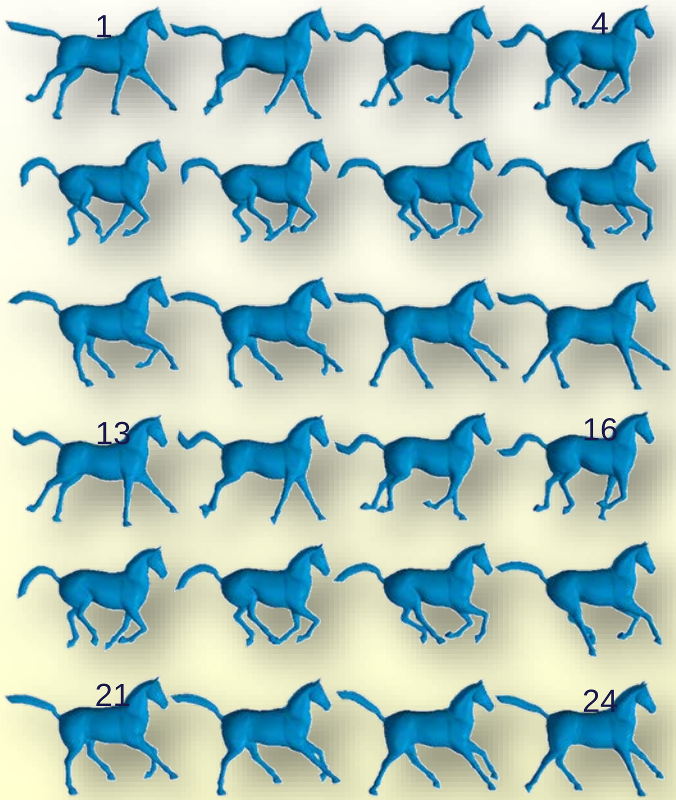

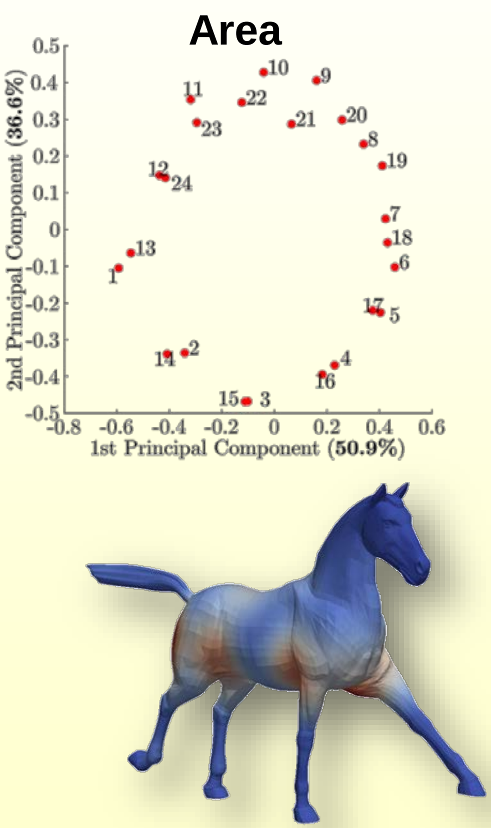

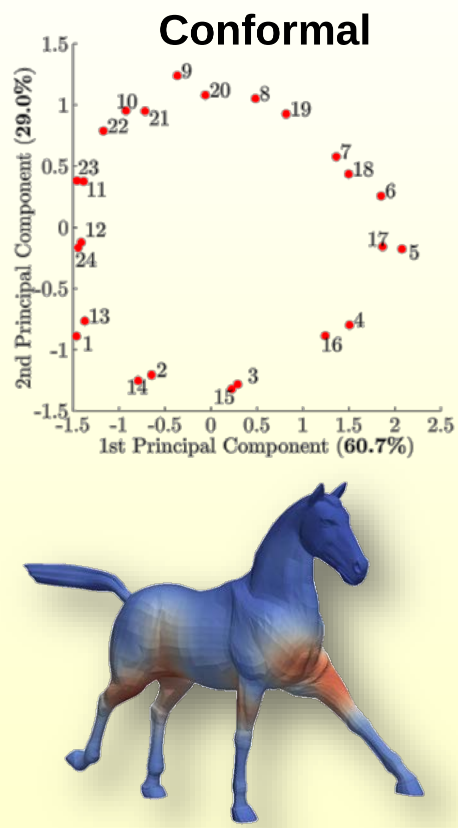

Shape Analogies

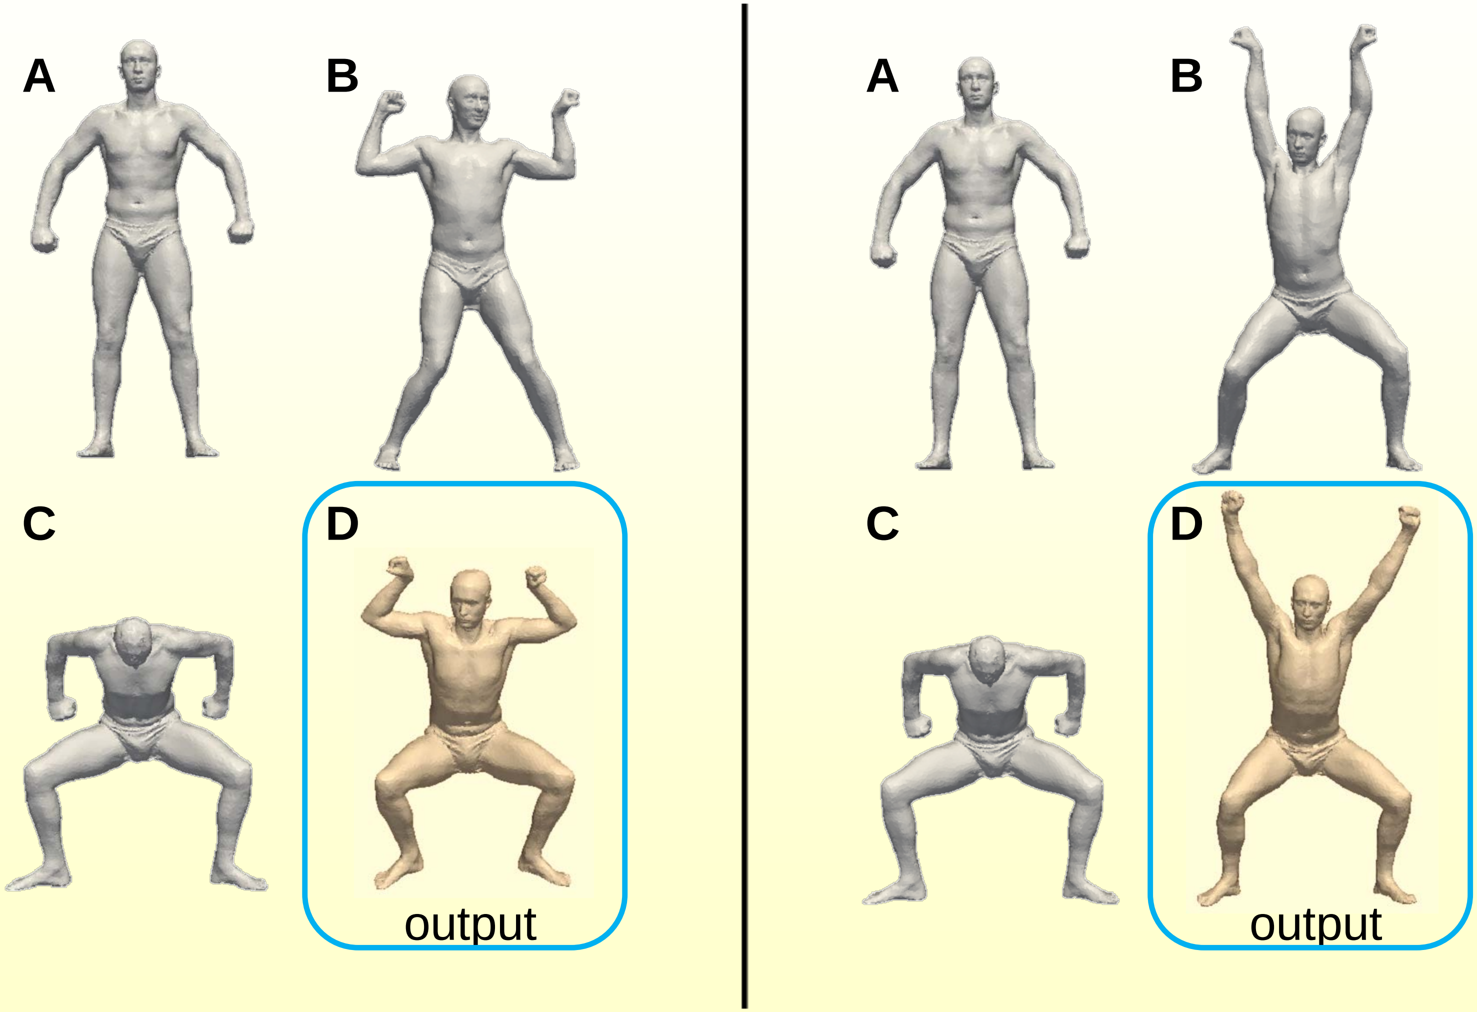

Map Networks for Related Data

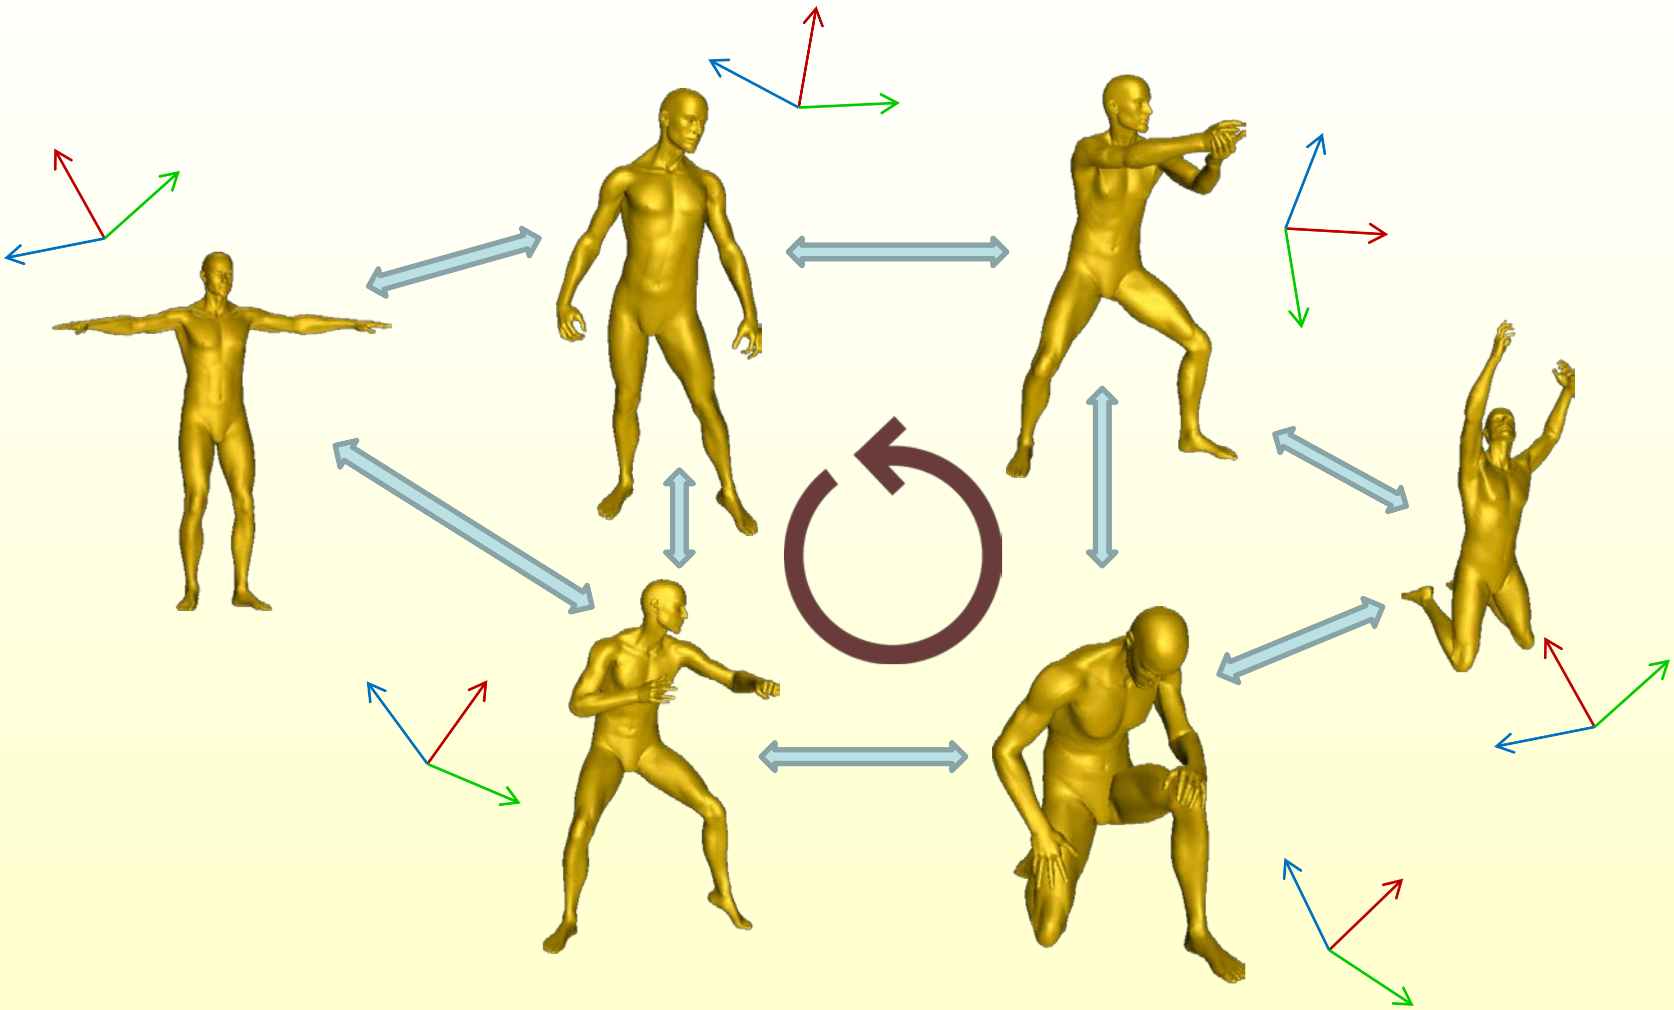

Cycle consistency

Cycle-Consistency ≡ Low-Rank

- In a map network, commutativity, path-invariance, or cycle-consistency are equivalent to a low rank or semidefiniteness condition on a big mapping matrix

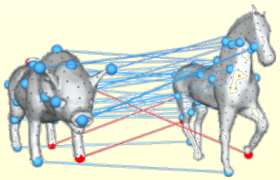

$$X = \begin{pmatrix} I_m & X_{1,2} & \cdots & X_{1,n} \\ X_{1,2} & I_m & \cdots & \cdots \\ \vdots & \vdots & I_m & X_{(n-1),n} \\ X_{n,1} & \vdots & X_{n,(n-1)} & I_m \end{pmatrix}.$$

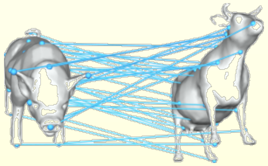

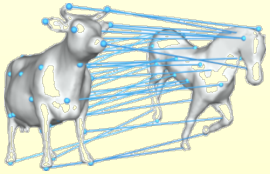

- Conversely, such a low-rank condition can be used to
  - regularize and clean up functional maps
  - extract shared structure

Map processing!

Entity Extraction in Images

[F. Wang, Q. Huang, L. G., ICCV '13]

- Task: jointly segment a set of related images
  - same object, different viewpoints/scales:

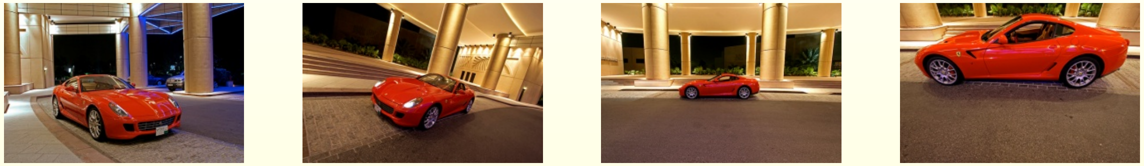

  - similar objects of the same class:

- Benefits and challenges:
  - Images can provide weak supervision for each other
  - But exactly how should they help each other? How to deal with clutter and irrelevant content?

PASCAL: 10 images per class are shown

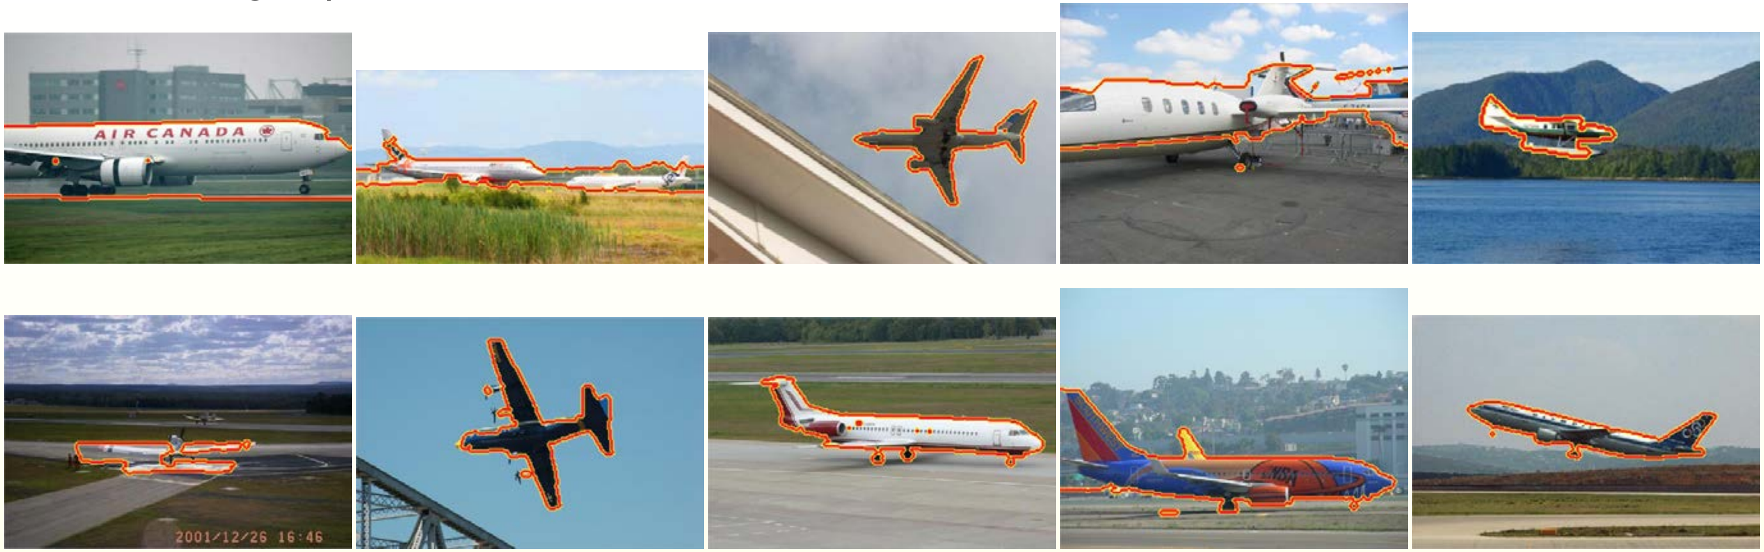

Multiple Latent Spaces

Multi-Class Co-Segmentation

[F. Wang, Q. Huang, M. Ovsjanikov, L. G., CVPR'14]

- Input:
  - A collection of N images sharing M objects
  - Each image contains a subset of the objects





- Output
  - Discovery of what objects appear in each image
  - Their pixel-level segmentation

Framework

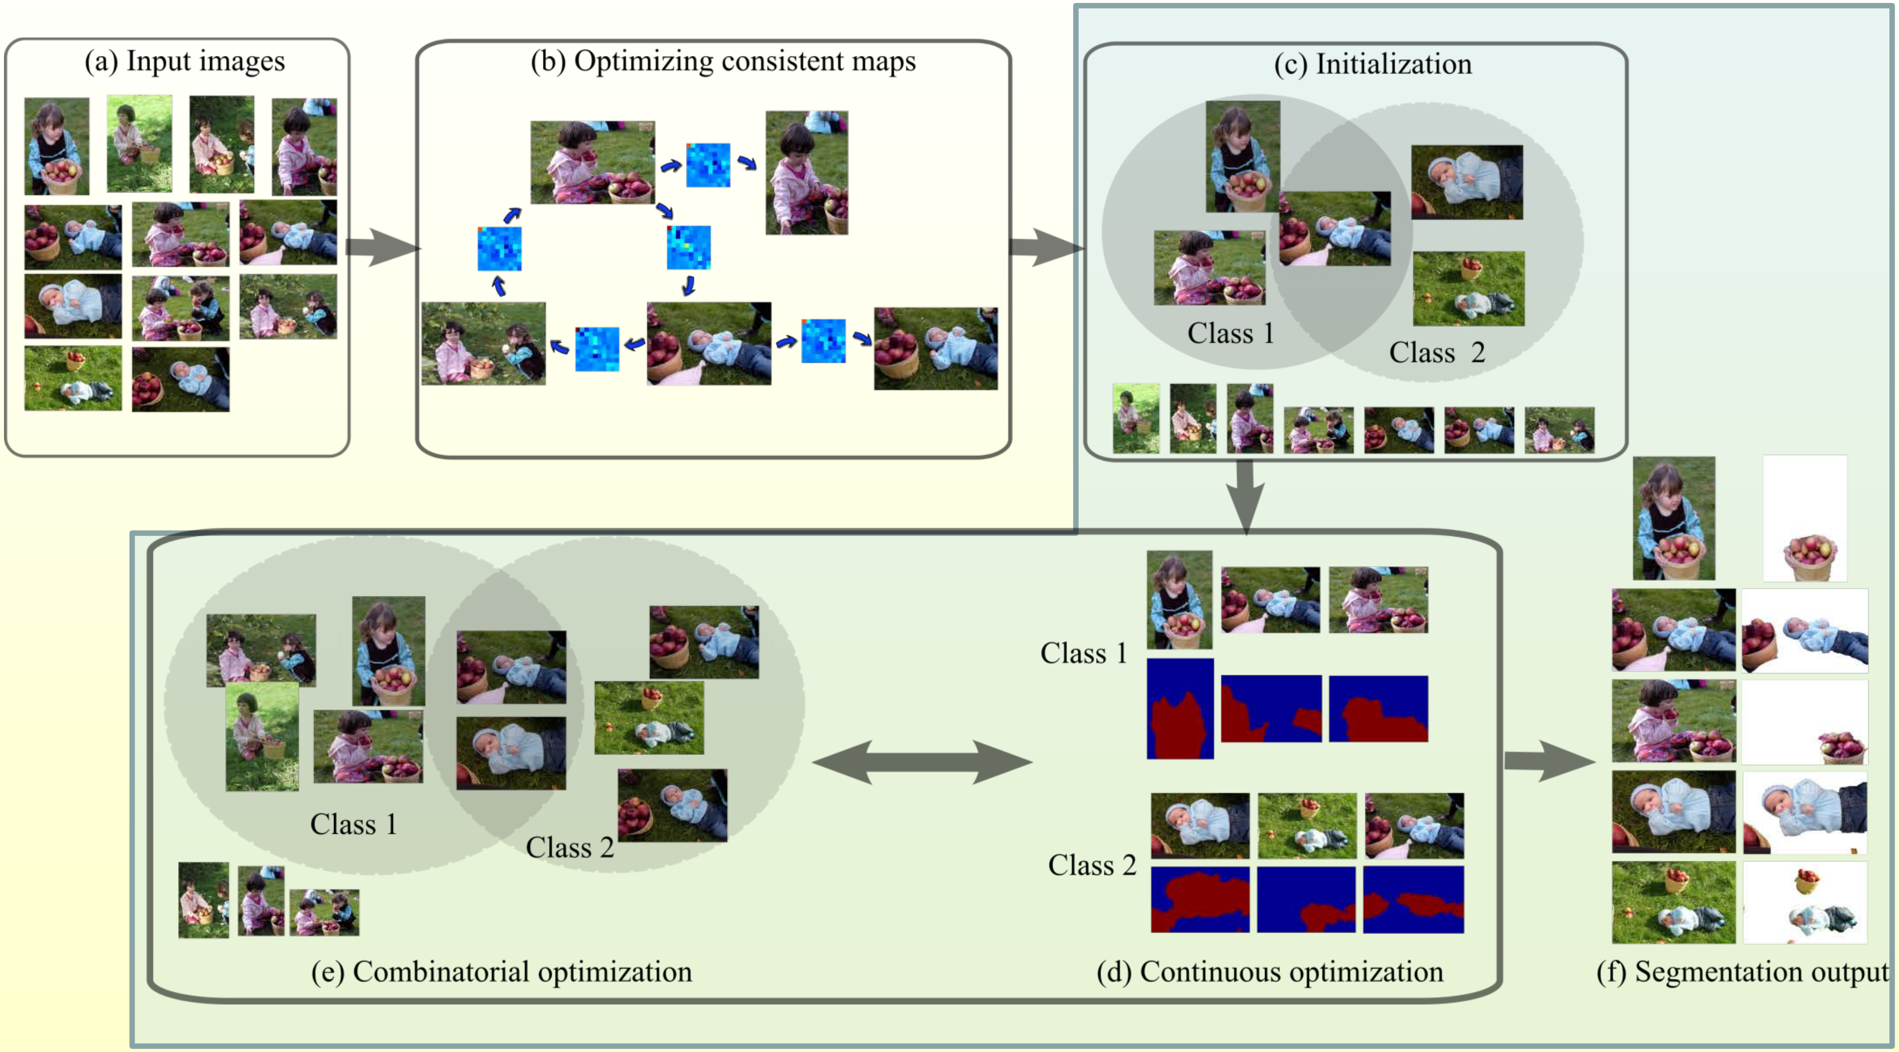
(a) Input images
(b) Optimizing consistent maps
(c) Initialization
Class 1
Class 2
(e) Combinatorial optimization
Class 1
Class 2
(d) Continuous optimization
Class 1
Class 2
(f) Segmentation output

Consistent Functional Maps

- Partial cycle consistency:

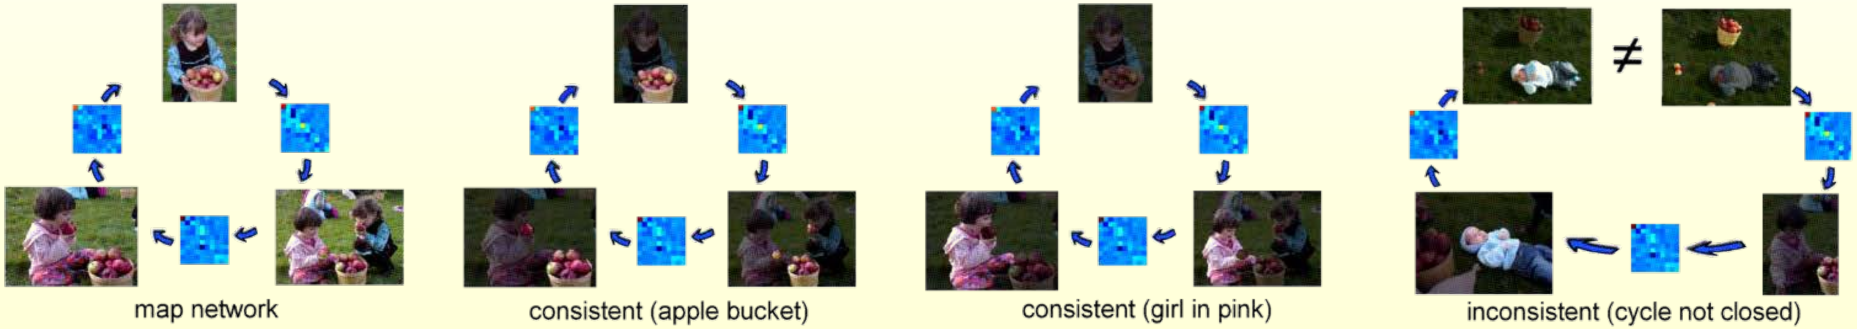

Must deal with non-total maps

Related to topological persistence / persistent homology

Consistent Functional Maps

- Latent functions: $\boldsymbol{Y}_i = (\boldsymbol{y}_{i1}, \cdots, \boldsymbol{y}_{iL})$
- Discrete variables: $\boldsymbol{z}_i = \{z_{il} \in \{0, 1\}, 1 \leq l \leq L\}$
- Relationship: $\mathbf{Y}_i \mathrm{Diag}(\mathbf{z}_i) = \mathbf{Y}_i$
- Consistency:

$$\boldsymbol{X}_{ij}\boldsymbol{Y}_i = \boldsymbol{Y}_j \mathrm{Diag}(\boldsymbol{z}_i), \quad (i, j) \in \mathcal{E}.$$

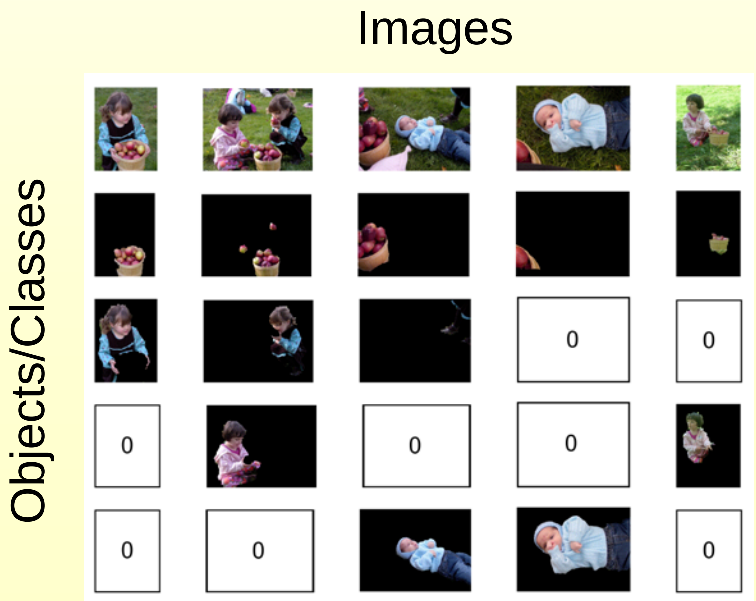

Optimizing Segmentation Functions

- Alternating between:
  - Continuous optimization:
    - Optimal segmentation functions in each class
  - Combinatorial optimization:
    - Class assignment by propagating segmentation functions

Continuous Optimization

- Optimize segmentations in each object class
  - Consistent with functional maps
  - Align with segmentation cues
  - Mutually exclusive

$$
\begin{aligned}
\min_{s_{ik}, i \in \mathcal{C}_k} \quad & \sum_{k=1}^{M} \sum_{(i,j) \in \mathcal{E} \cap (\mathcal{C}_k \times \mathcal{C}_k)} \|X_{ij} s_{ik} - s_{jk}\|^2 \\
& + \gamma \sum_{l \neq k} \sum_{i \in \mathcal{C}_k \cap \mathcal{C}_l} (s_{il}^T s_{ik})^2 + \mu \sum_{k=1}^{M} \sum_{i \in \mathcal{C}_k} s_{ik}^T L_i s_{ik} \\
\text{subject to} \quad & \sum_{i \in \mathcal{C}_k} \|s_{ik}\|^2 = |\mathcal{C}_k|, \quad 1 \leq k \leq K.
\end{aligned}
$$

Combinatorial Optimization

- Expand each object class by propagating segmentations to other images

$$\max_{s_{ik}} \quad \frac{1}{|\mathcal{N}(i) \cap \mathcal{C}_k|} \sum_{j \in \mathcal{N}(i) \cap \mathcal{C}_k} (s_{ik}^T X_{ji} s_{jk})^2 - \gamma \sum_{l \neq k, i \in \mathcal{C}_l} (s_{ik}^T s_{il})^2 - \mu s_{ik}^T L_i s_{ik}$$

$$\text{subject to} \quad \|s_{ik}\|^2 = 1$$

Optimizing Segmentation Functions

- More images will be included in each object class

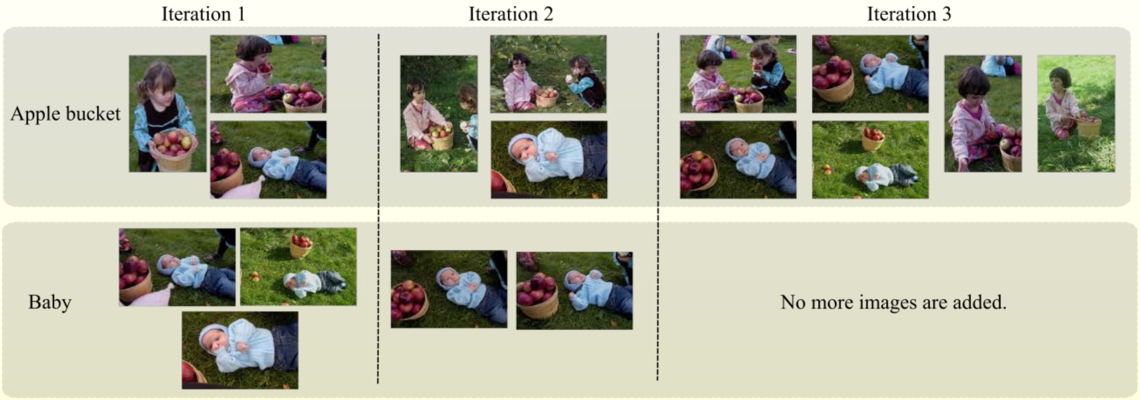

- Segmentation functions are improved during iterations

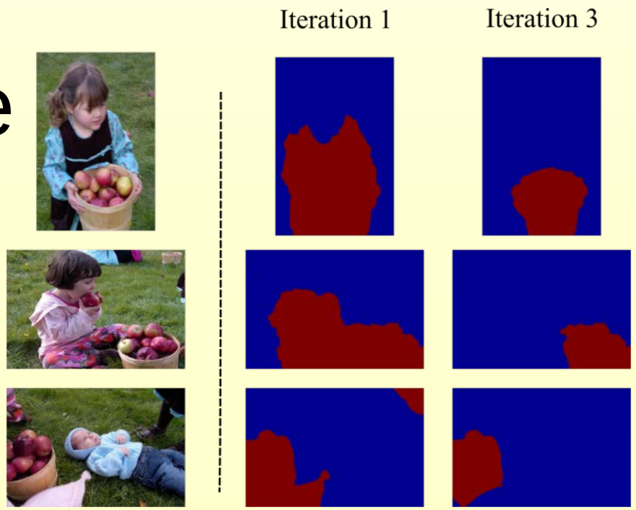

Experimental Results

- Accuracy
  - Intersection-over-union
  - Find the best one-to-one matching between each cluster and each ground-truth object.
- Benchmark datasets
  - MSRC: 30 images, 1 class (degenerated case);
  - FlickrMFC data set: 20 images, 3~6 classes
  - PASCAL VOC: 100~200 images, 2 classes

Experimental Results

class	N	M	Kim'12	Kim'11	Joulin '10	Mukherjee '11	Ours
Apple	20	6	40.9	32.6	24.8	25.6	**46.6**
Baseball	18	5	31.0	31.3	19.2	16.1	**50.3**
butterfly	18	8	29.8	32.4	29.5	10.7	**54.7**
Cheetah	20	5	32.1	40.1	50.9	41.9	**62.1**
Cow	20	5	35.6	43.8	25.0	27.2	**38.5**
Dog	20	4	34.5	35.0	32.0	30.6	**53.8**
Dolphin	18	3	34.0	47.4	37.2	30.1	**61.2**
Fishing	18	5	20.3	27.2	19.8	18.3	**46.8**
Gorilla	18	4	41.0	38.8	41.1	28.1	**47.8**
Liberty	18	4	31.5	41.2	44.6	32.1	**58.2**
Parrot	18	5	29.9	36.5	35.0	26.6	**54.1**
Stonehenge	20	5	35.3	49.3	47.0	32.6	**54.6**
Swan	20	3	17.1	18.4	14.3	16.3	**46.5**
Thinker	17	4	25.6	34.4	27.6	15.7	**68.6**
Average	-	-	31.3	36.3	32.0	25.1	**53.1**

Performance comparison on the MFCFlickr dataset

class	N	Joulin'10	Kim'11	Mukherjee'11	Ours
Bike	30	43.3	29.9	42.8	**51.2**
Bird	30	47.7	29.9	-	**55.7**
Car	30	59.7	37.1	52.5	**72.9**
Cat	24	31.9	24.4	5.6	**65.9**
Chair	30	39.6	28.7	39.4	**46.5**
Cow	30	52.7	33.5	26.1	**68.4**
Dog	30	41.8	33.0	-	**55.8**
Face	30	70.0	33.2	40.8	60.9
Flower	30	51.9	40.2	-	**67.2**
House	30	51.0	32.2	66.4	56.6
Plane	30	21.6	25.1	33.4	**52.2**
Sheep	30	66.3	60.8	45.7	**72.2**
Sign	30	58.9	43.2	-	**59.1**
Tree	30	67.0	61.2	55.9	**62.0**

Performance comparison on the MSRC dataset

class	N	NCut	MNcut	Ours
Bike + person	248	27.3	30.5	**40.1**
Boat + person	260	29.3	32.6	**44.6**
Bottle + dining table	90	37.8	39.5	**47.6**
Bus + car	195	36.3	39.4	**49.2**
bus + person	243	38.9	41.3	**55.5**
Chair + dining table	134	32.3	30.8	**40.3**
Chair + potted plant	115	19.7	19.7	**22.3**
Cow + person	263	30.5	33.5	**45.0**
Dog + sofa	217	44.6	42.2	**49.6**
Horse + person	276	27.3	30.8	**42.1**
Potted plant + sofa	119	37.4	37.5	**40.7**

Performance comparison on the PASCAL-multi dataset

Apple + picking

Baseball + kids

Butterfly + blossom

Apple + picking (red: apple bucket; magenta: girl in red; yellow: girl in blue; green: baby; cyan: pum

Baseball + kids (green: boy in black; blue: boy in grey; yellow: coach.)

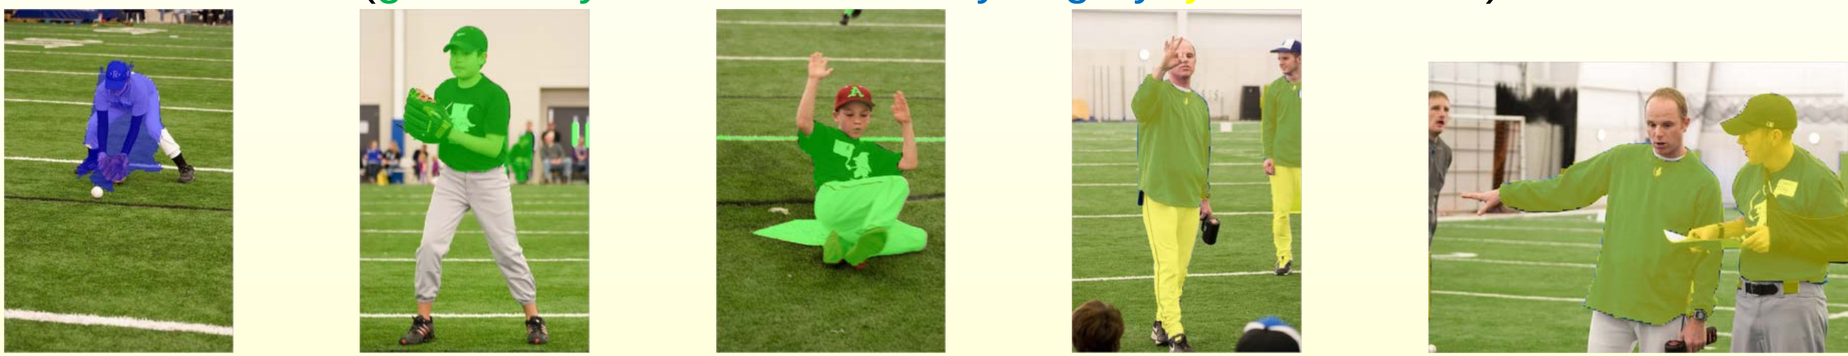

Butterfly + blossom (green: butterfly in orange; yellow: butterfly in yellow; cyan: red flowe

Cheetah + Safari

Cow + pasture

Dog + park

Dolphin + aquarium

Cheetah + Safari (red: cheetah; yellow: lion; magenta: monkey.)

Cow + pasture (red: black cow; green: brown cow; blue: man in blue.)

Dog + park (red: black dog; green: brown dog; blue: white dog.)

Dolphin + aquarium (red: killer whale; green: dolphin.)

Fishing + Alaska





Gorilla + zoo





Liberty + statue



Parrot + zoo





Fishing + Alaska (blue: man in white; green: man in gray; magenta: woman in gray; yellow: salmon.

Gorilla + zoo (blue: gorilla; yellow: brown orangutan)

Liberty + statue (blue: empire state building; green: red boat; yellow: liberty statue.)

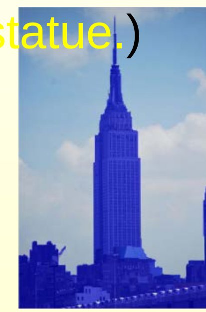

Parrot + zoo (red: hand; green: parrot in green; blue: parrot in red.)

Stonehenge

Swan + zoo

Thinker + Rodin

Stonehenge (blue: cow in white; yellow: person; magenta: stonehenge.)

Swan + zoo (blue: gray swan; green: black swan.)

Thinker + Rodin (red: sculpture Thinker; green: sculpture Venus; blue: Van Gogh.)

Apple + picking (red: apple bucket; magenta: girl in red; yellow: girl in blue; green: baby; cyan: pum

Baseball + kids (green: boy in black; blue: boy in grey; yellow: coach.)

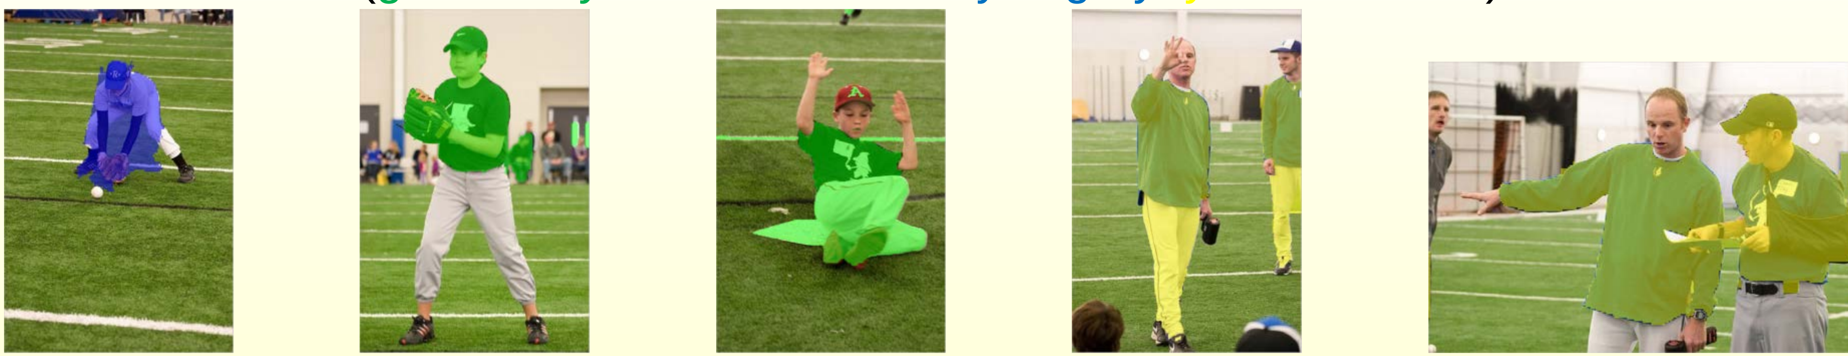

Butterfly + blossom (green: butterfly in orange; yellow: butterfly in yellow; cyan: red flowe

3D Co-Segmentation

Consistent Shape Segmentation

[Q. Huang, F. Wang, L. Guibas, ’14]

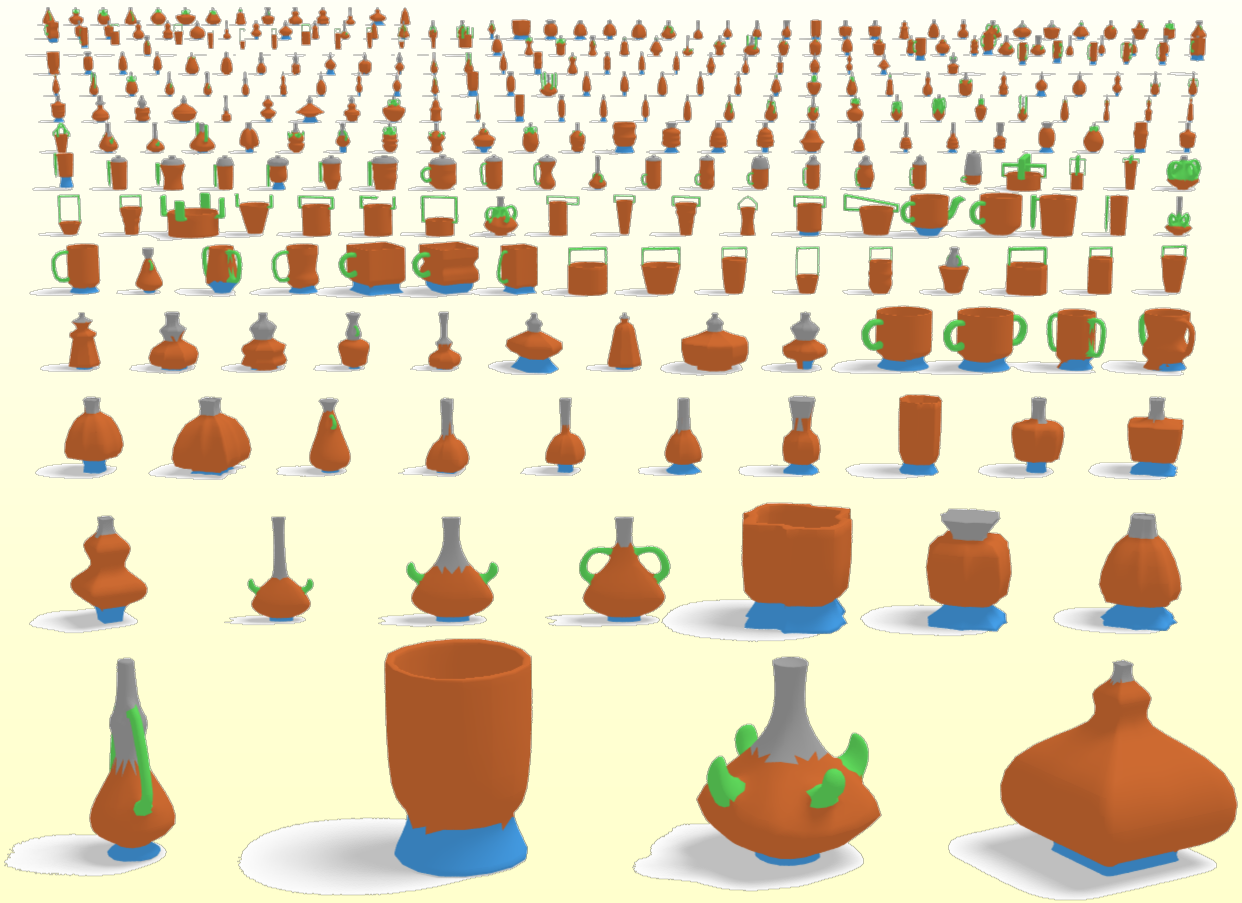

First Build a Network

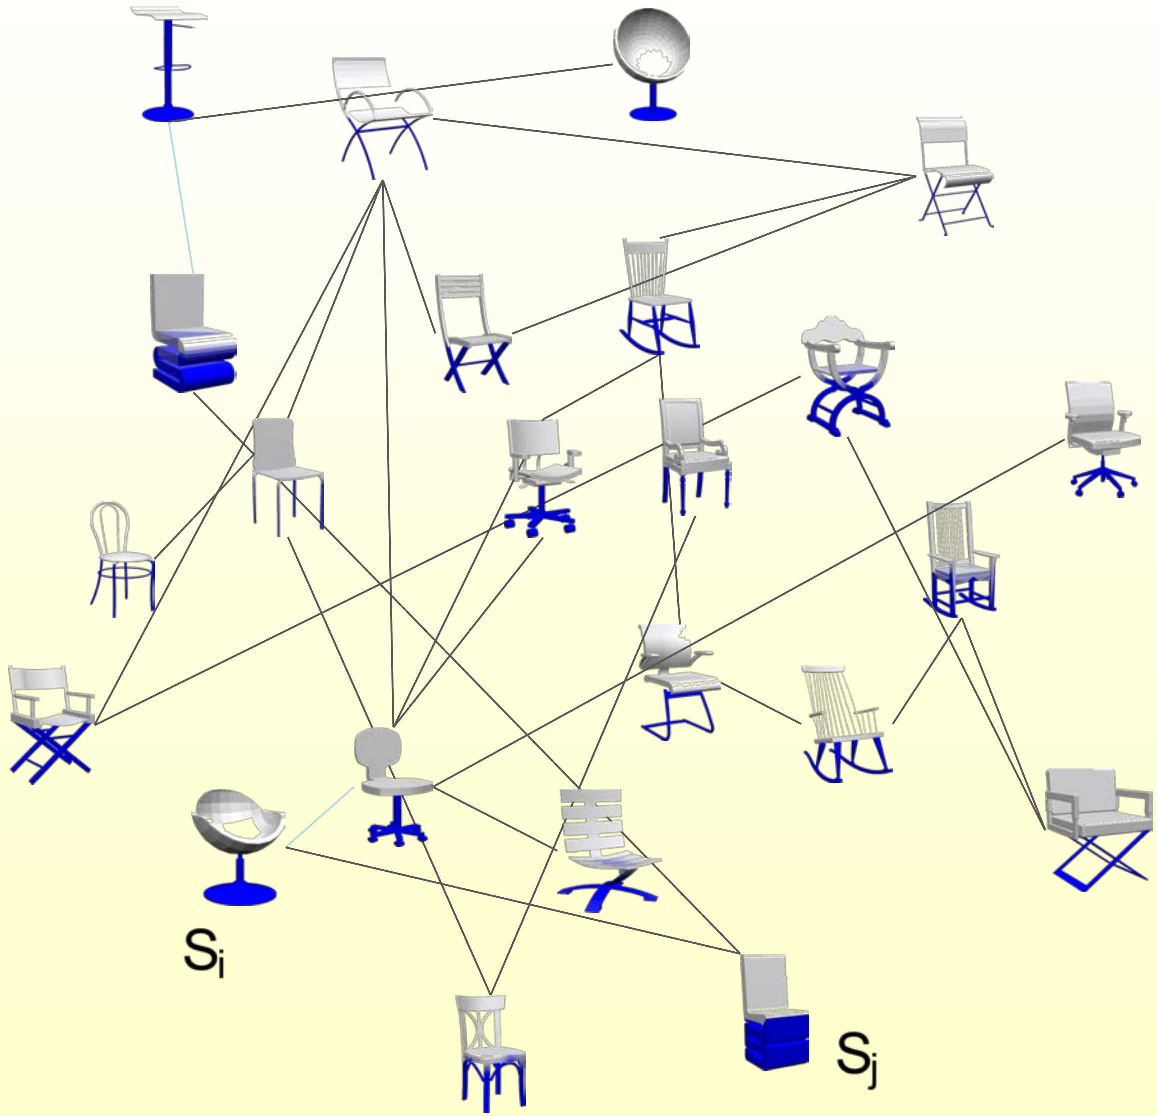

distance histogram

Use the D2 shape descriptor and connect each shape to its nearest neighbors

$$\mathcal{G} = (\mathcal{F}, \mathcal{E})$$

Start From Noisy Shape Descriptor Correspondences

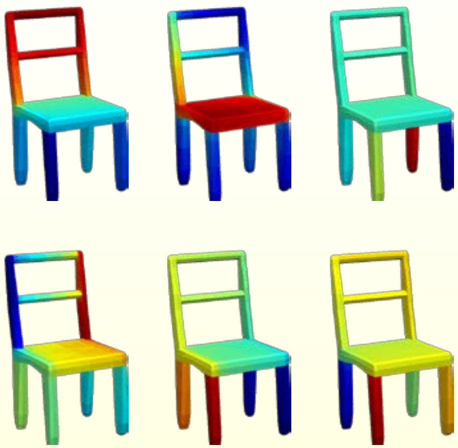

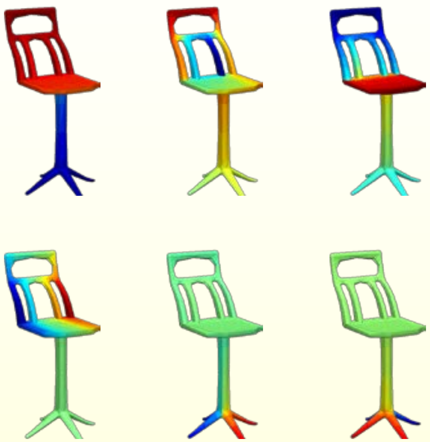

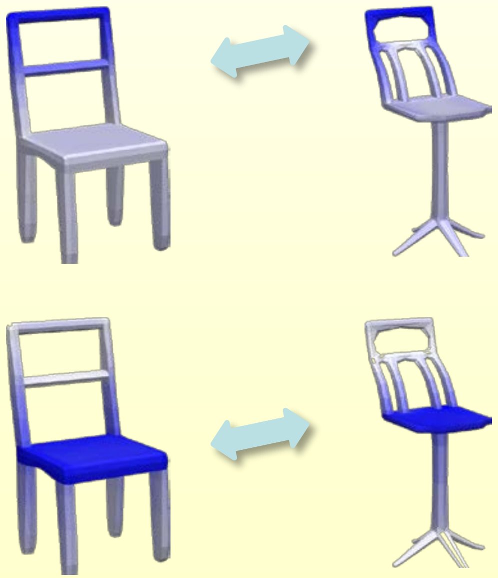

$$C_i X_{ij} \approx D_j$$

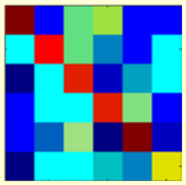

Lift to
functional form

$C_i \bullet \quad \bullet \quad \bullet \; D_i$

The Pipeline

Step 1 Step 2

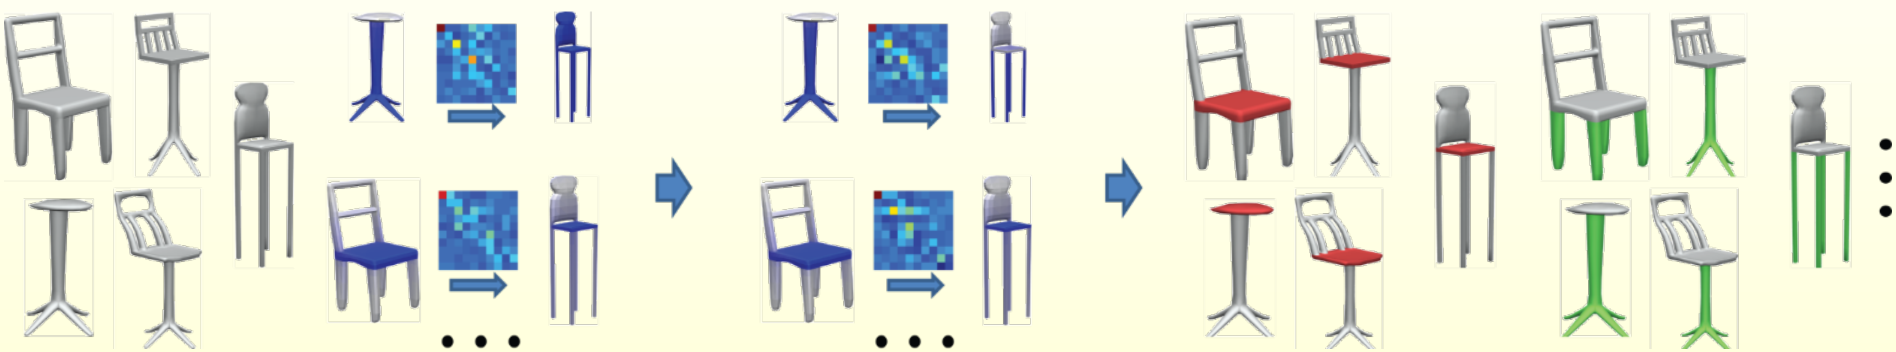

Original shapes with noisy maps

Cleaned up maps

Consistent basis functions extracted

Joint Map Optimization

- Step 1: Convex low-rank recovery using robust PCA – we minimize over all *X*

convex!

trace norm

$\|X\|_\star = \sum_i \sigma_i(X)$

$$X^\star = \lambda \|X\|_\star + \min_X \sum_{(i,j)\in\mathcal{G}} \|X_{ij}C_{ij} - D_{ij}\|_{2,1}$$

$\|A\|_{2,1} = \sum_i \|\vec{a}_i\|$

Dual ADMM

- Step 2: Perturb the above *X* to force the factorization

$$\sum_{1\leq i,j\leq N} \|X^\star_{ij} - Y_j^{+} Y_i\|_F^2 + \mu \sum_{i=1}^{N} \sum_{1\leq k<l\leq L} (\mathbf{y}_{ik}^T \mathbf{y}_{il})^2$$

Non-linear least squares
Gauss-Newton descent

The Y_i give us the desired latent spaces

Consistent Shape Segmentation

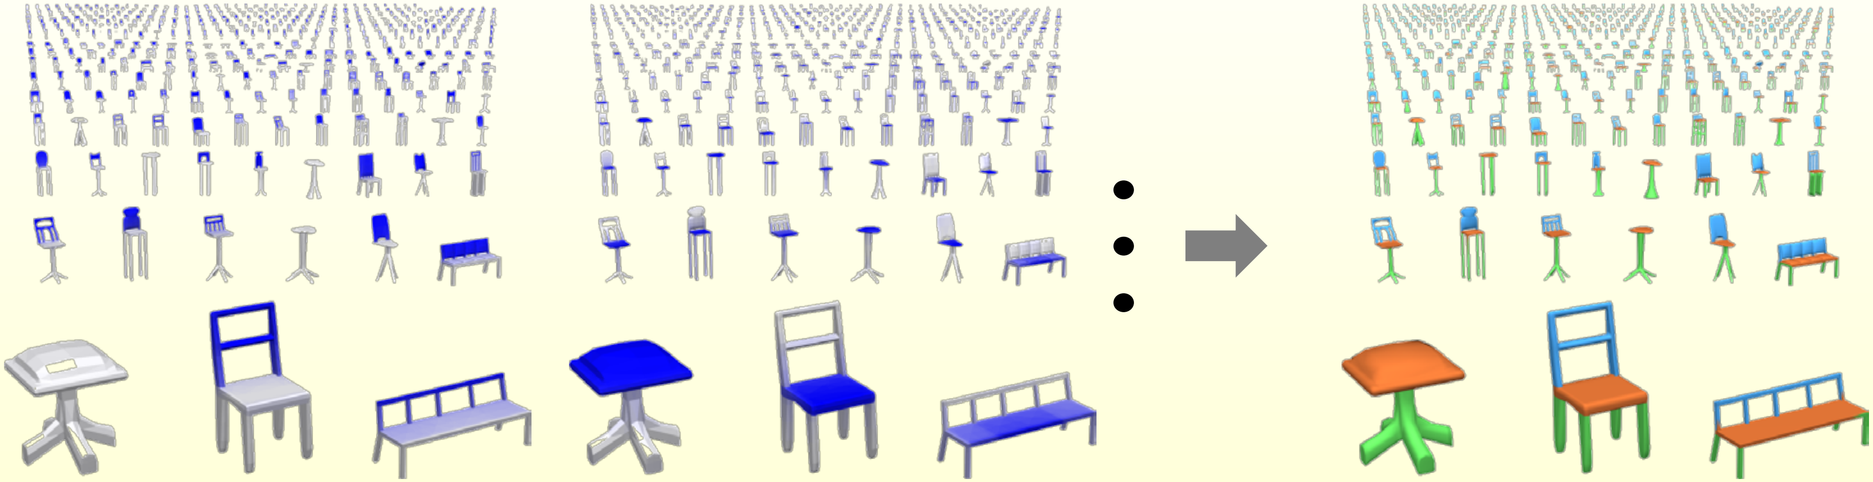

Via 2nd order MRF on each shape independently

Hierarchical Scaling

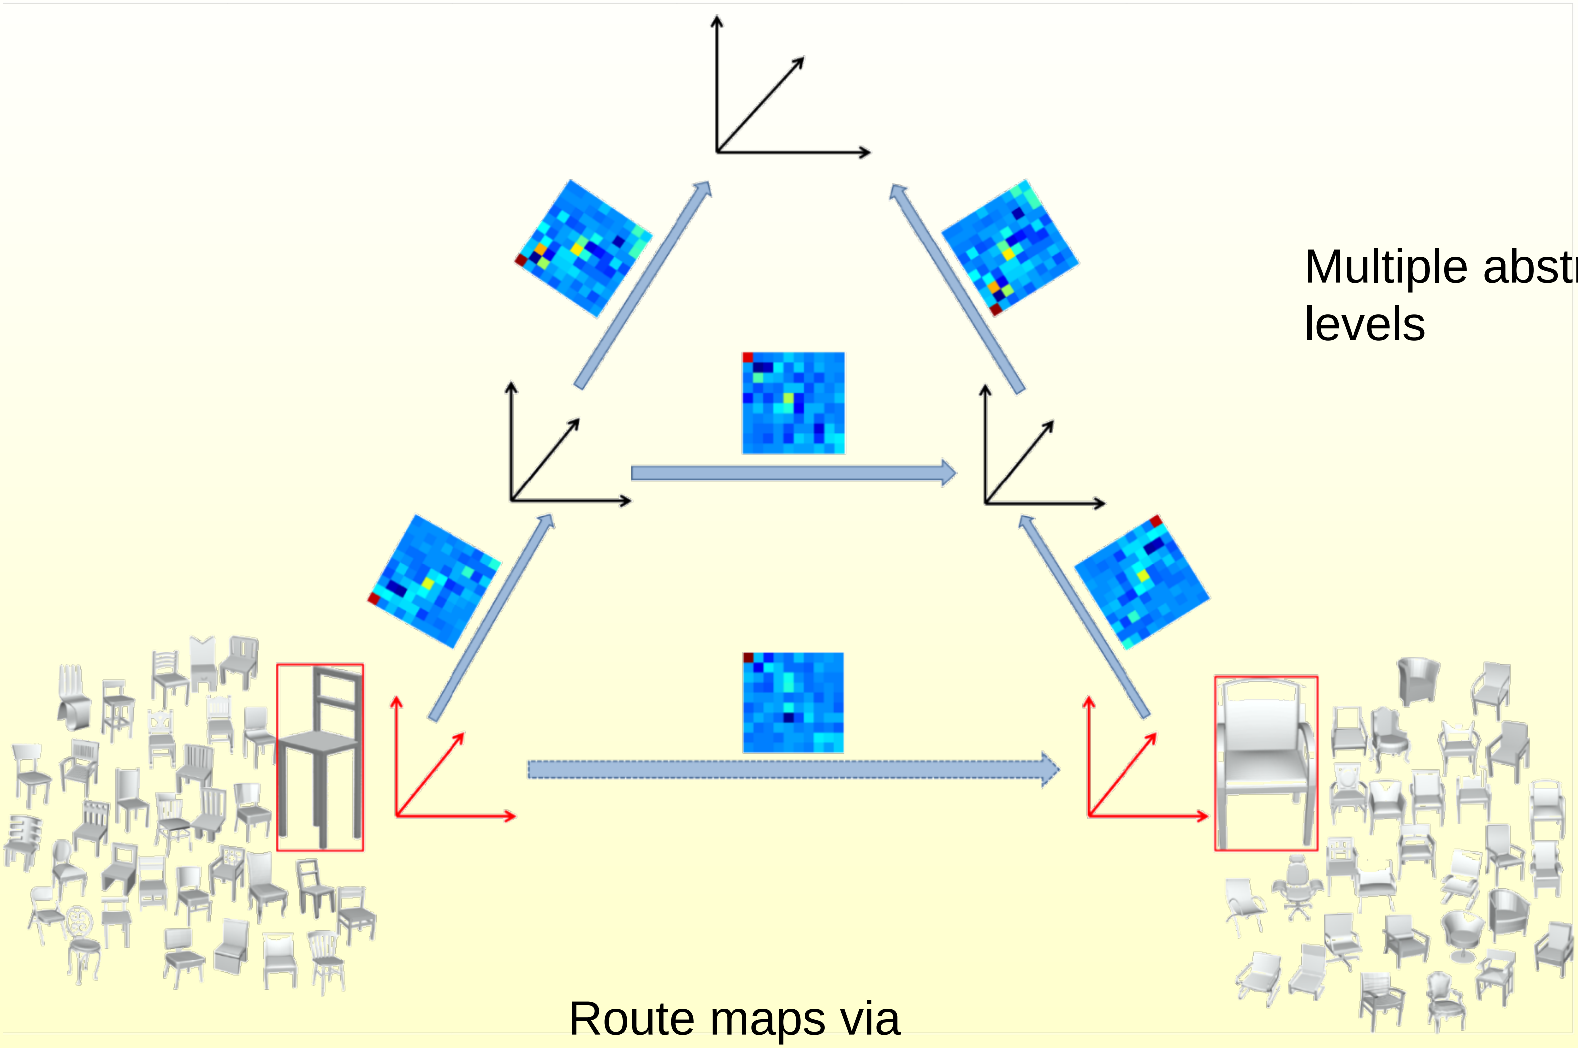

Multiple abstraction levels

Route maps via the abstractions

Co-Limits:
The Network is the Abstraction

Mosaicking or SLAM at the Level of Functions

http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15463-f08/www/proj4/www/gme/

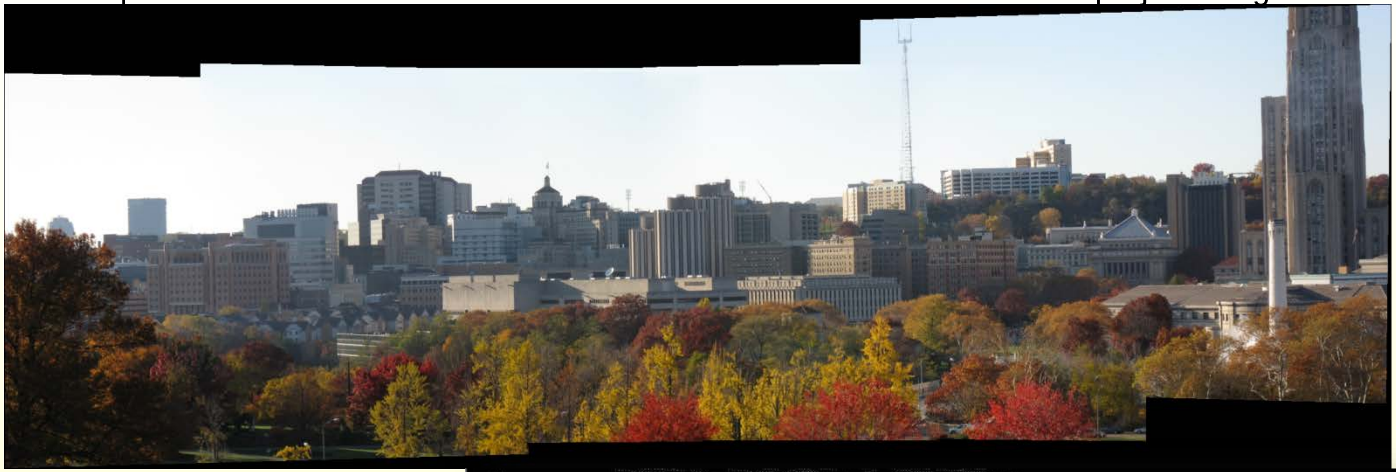

robotics.ait.kyushu-u.ac.jp

The Network is the Abstraction

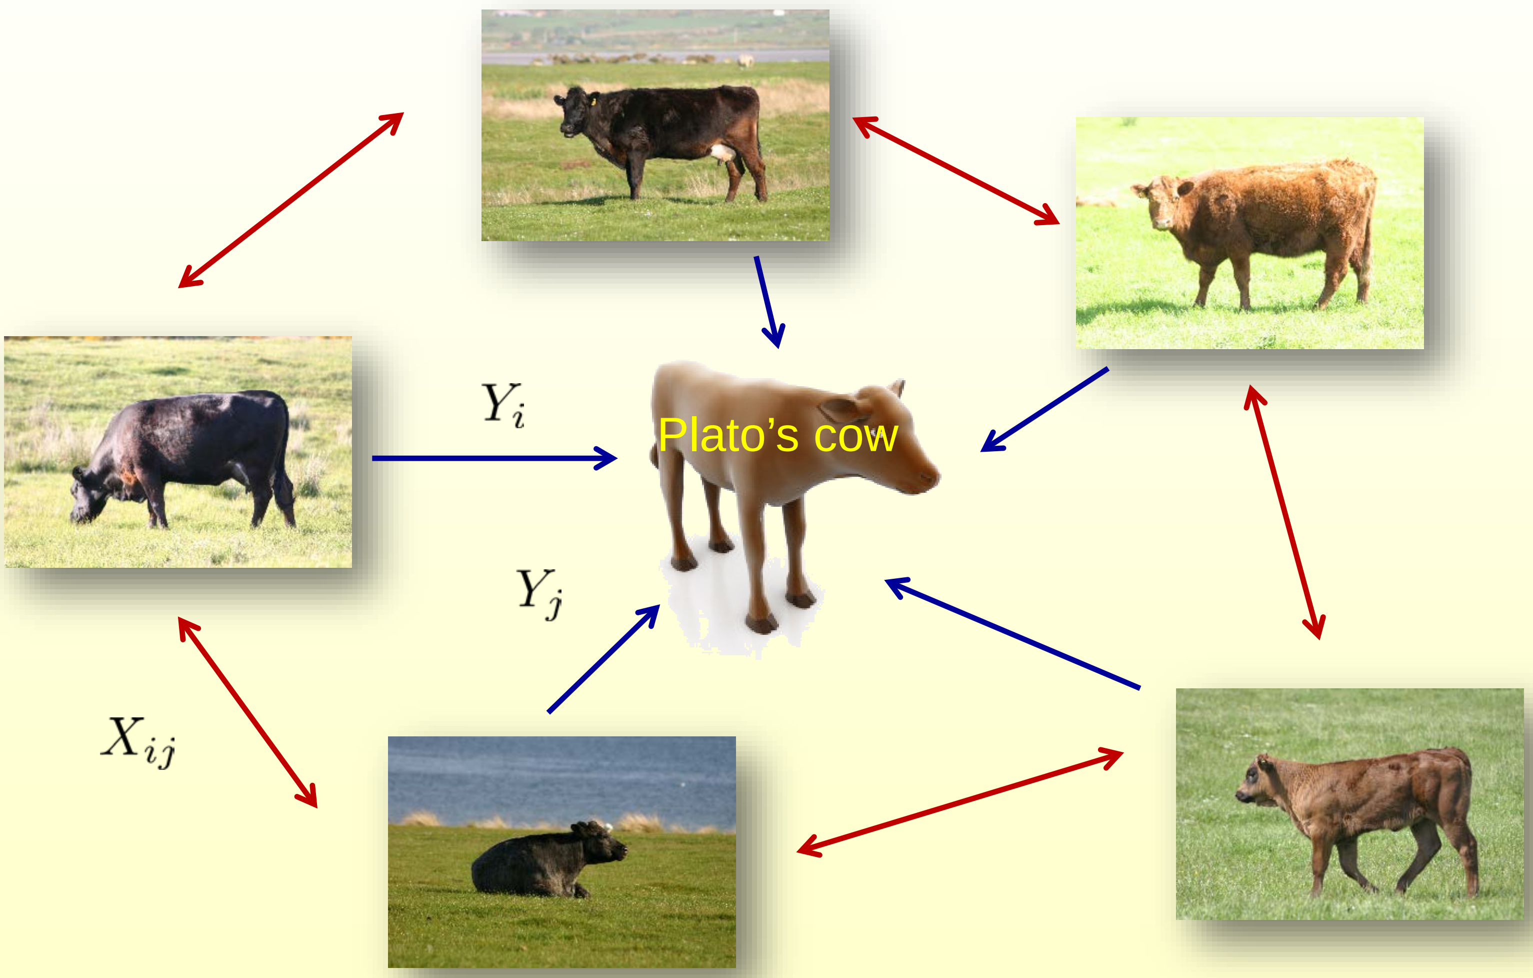

Abstractions Emerge from he Network

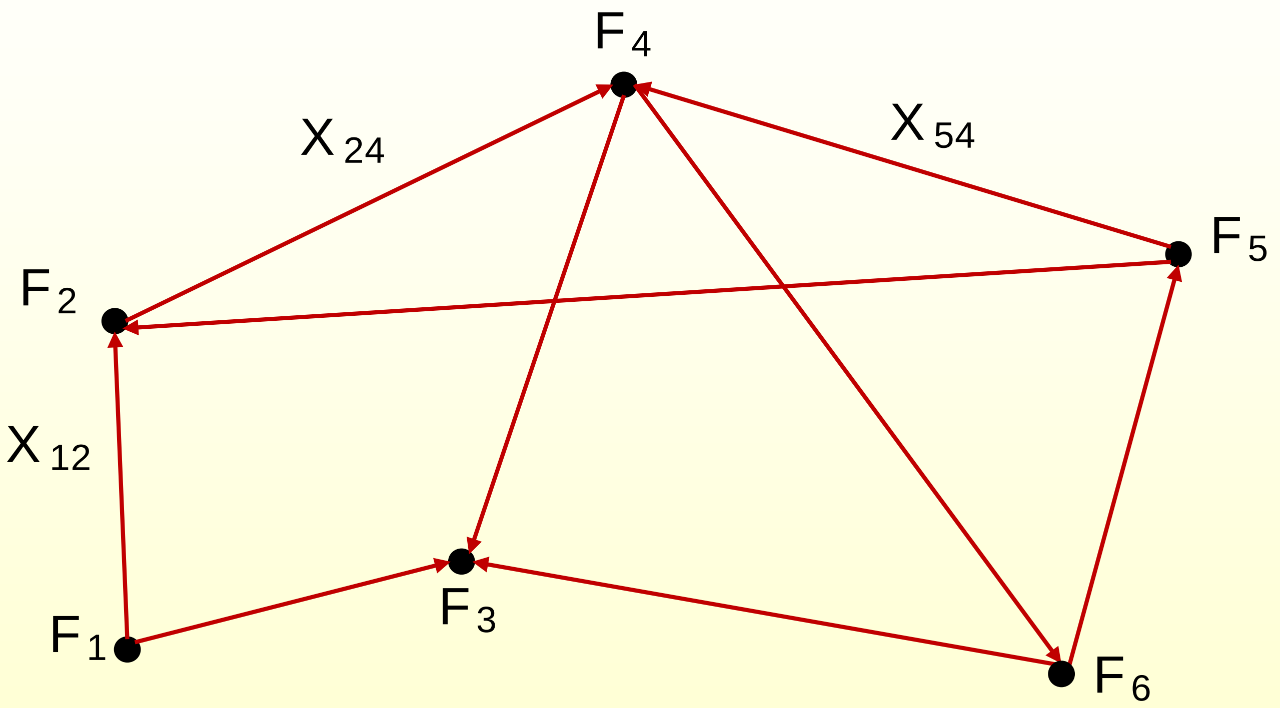

(Approximately) Cycle-Consistent Diagram

Abstraction – Colimit

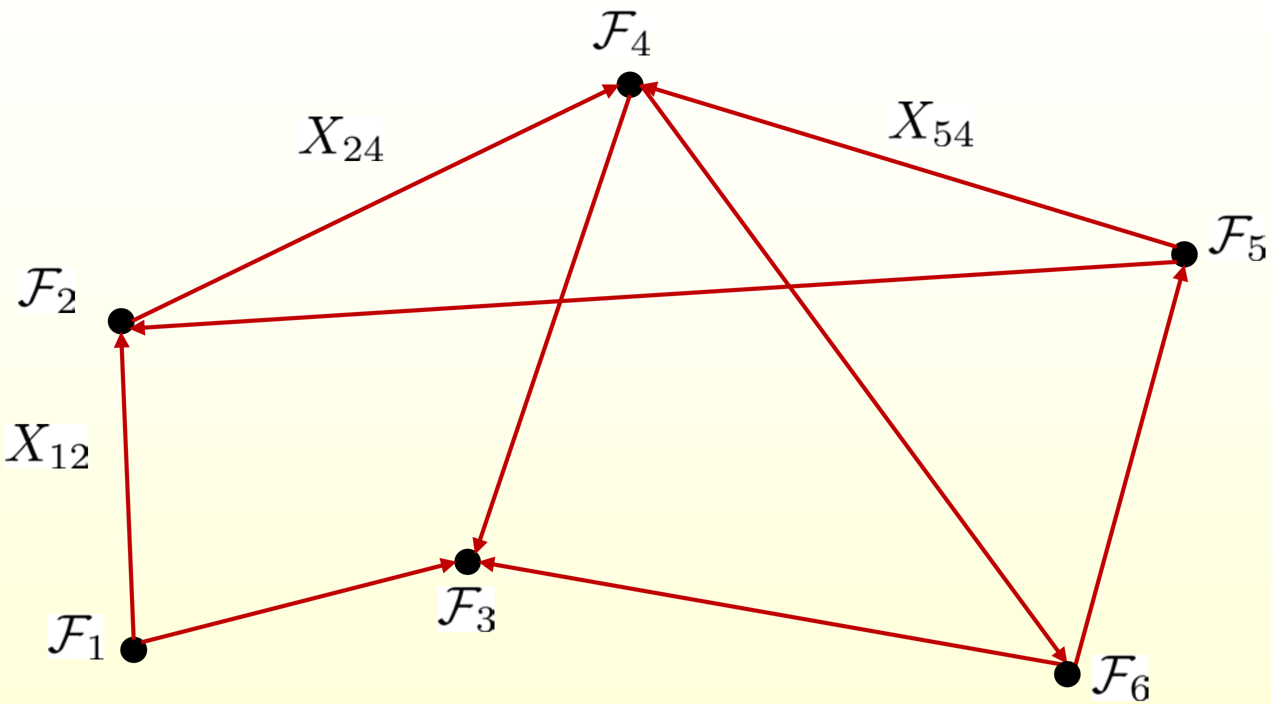

Colimits glue parts together to make a whole

$$\varinjlim \mathcal{F}_i = \coprod_i \mathcal{F}_i \Big/ \sim$$

Abstraction – Approximate Colimit

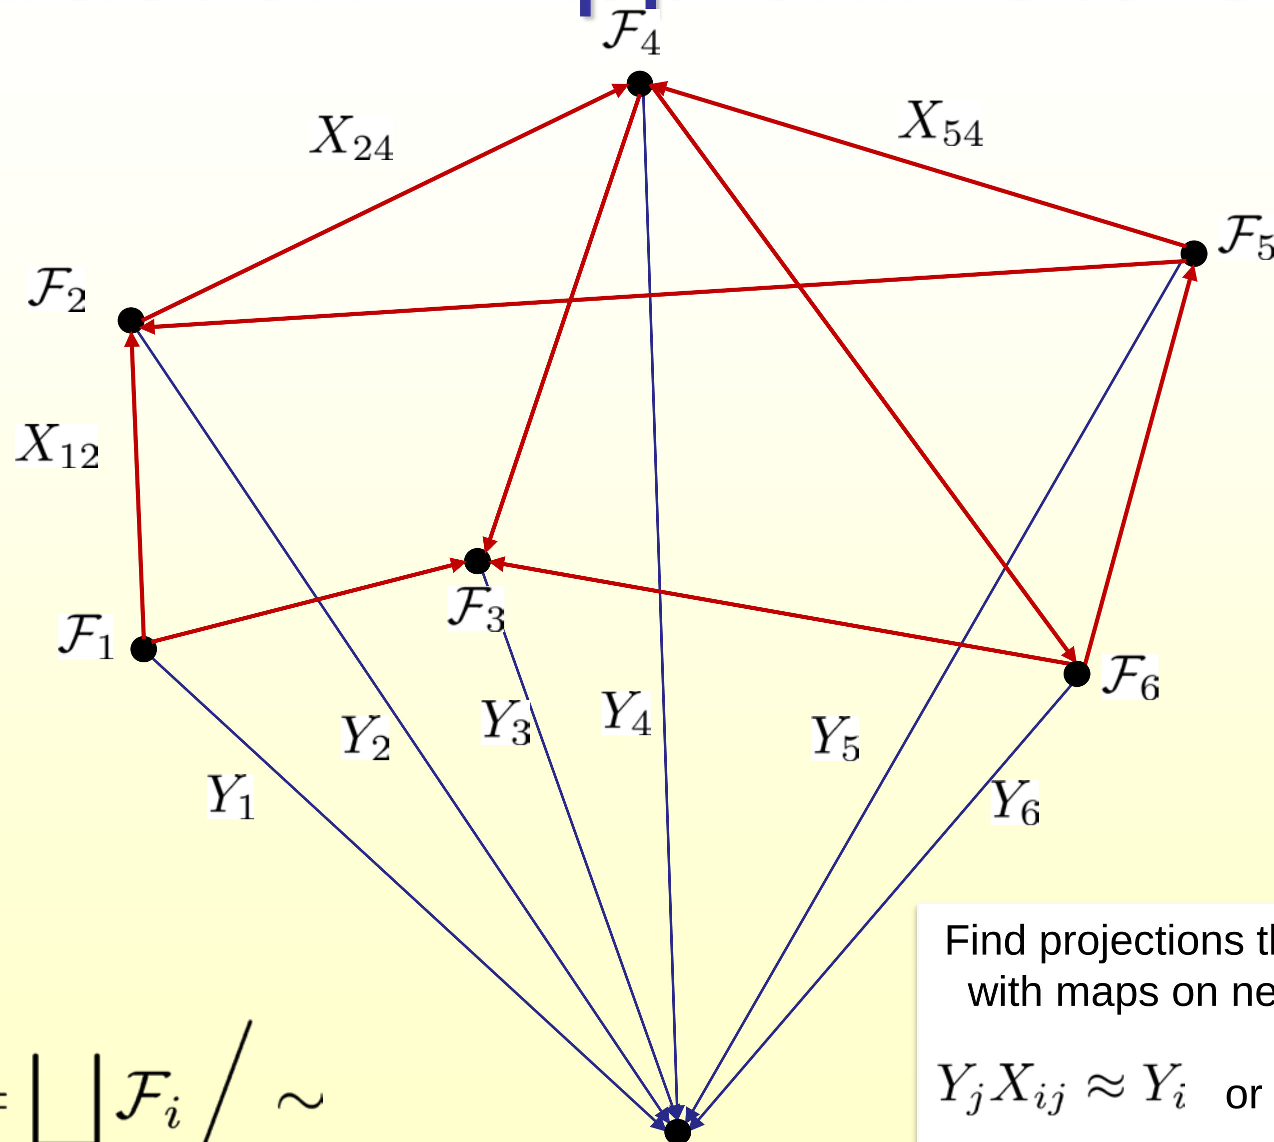

$$\varinjlim \mathcal{F}_i = \bigsqcup_i \mathcal{F}_i \Big/ \sim$$

Find projections that "play well" with maps on network edges,

$Y_j X_{ij} \approx Y_i$ or $X_{ij} \approx Y_j^+ Y_i$

"Colimit" = Latent space = Abstraction

The Network is the Abstraction

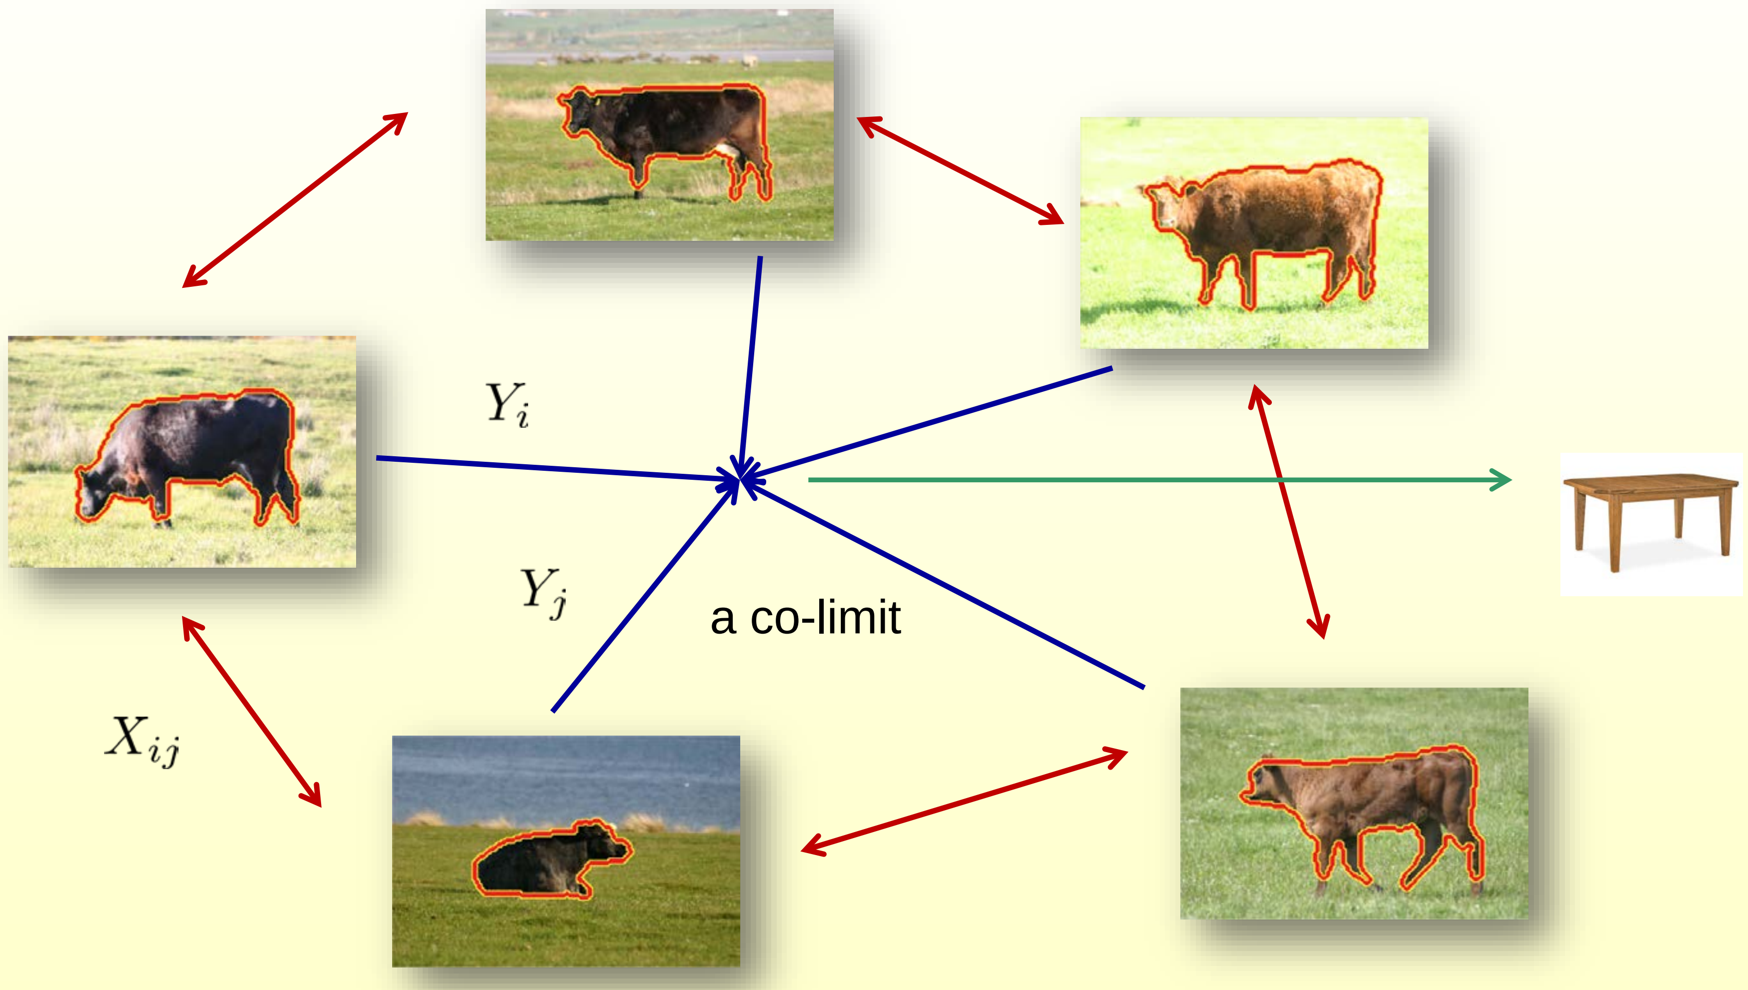

The Course

Data Analysis

- (Geometric and Topological Data) (Analysis)
- (Geometric and Topological) (Data Analysis)

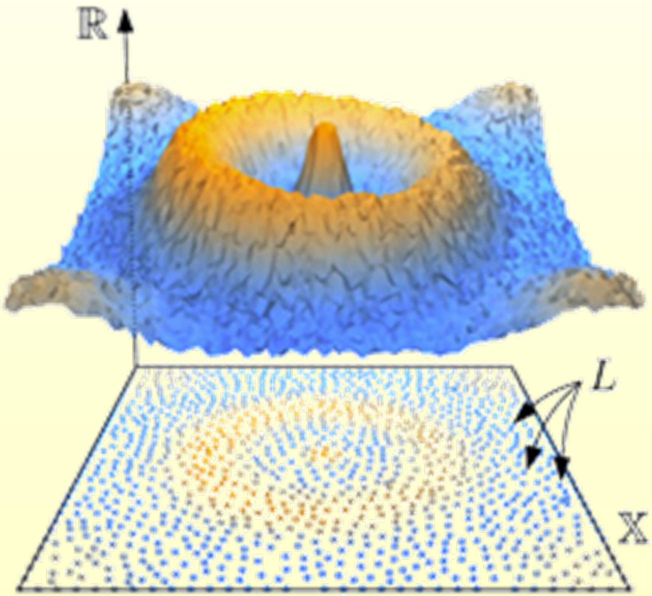

Topics Covered

- Visual data sets
- PCA
- CCA
- Spectral graph methods
- MDS; NLDR
- Intro to computational topology
- Homology and persistent homology
- Mapper
- 3D geometry representations
- Shape descriptors (shape context, spin images, HKS, WKS)
- Rigid alignments (ICP, RANSAC, geometric hashing)
- Non-rigid alignments (isometric, conformal)
- Shape correspondences
- Volumetric and Multiview CNNs
- Deep learning on point clouds
- Graph and mesh CNNs
- Functional Maps
- Shape differences
- Map networks and cycle consistency

Topics Not Covered

- Factor analysis
- Independent components analysis
- Nearest neighbor search
- Locality sensitive hashing
- Clustering
- Topic modeling (LDA, etc)
- Tensor decompositions
- Dictionary learning
- Sparse recovery / compressive sensing
- Mixture models
- Non-negative matrix factorization
- Matrix completion
- Zig-zag persistence
- Reeb graphs
- Random graphs and network models


optimization
geometric
spectral
Persistent
homology
cycle
methods
embeddings
descriptors
shape
image
algebraic
function
linear
Laplacian
map
data
PCA
graph
analysis
alignment
kernel
intrinsic
local
topology
spaces
techniques
visual
consistency

Thank You