CS233: Geometric and Topological Data Analysis

also CME251

Leonidas J. Guibas

Spring 2017-18

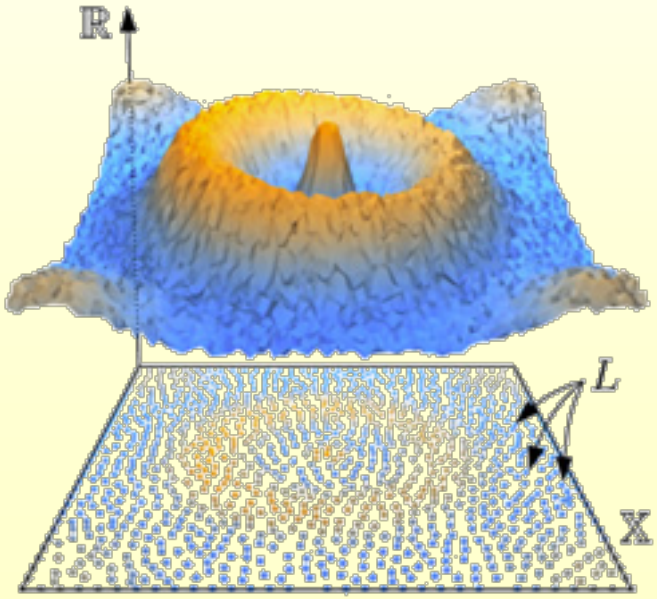

Introduction

Big Data Era

- Data from many kinds of sensors

- Data from simulations

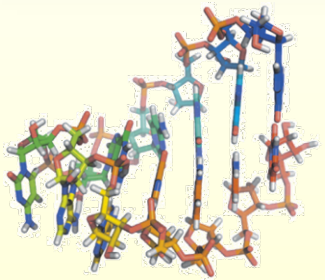

- Data from the activities of individuals on the internet

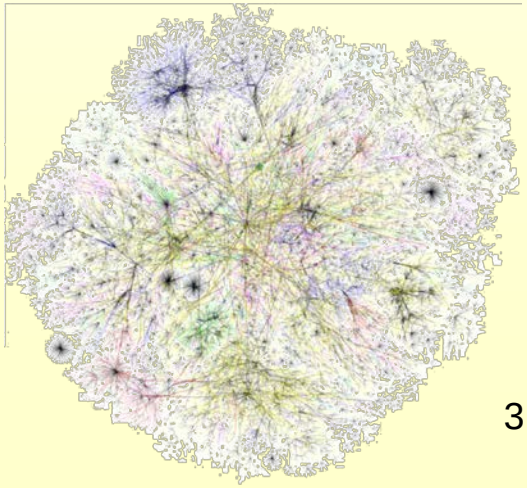

"Data Science"

Data Sets of Geometric Character

Vehicle GPS traces

3D city scans

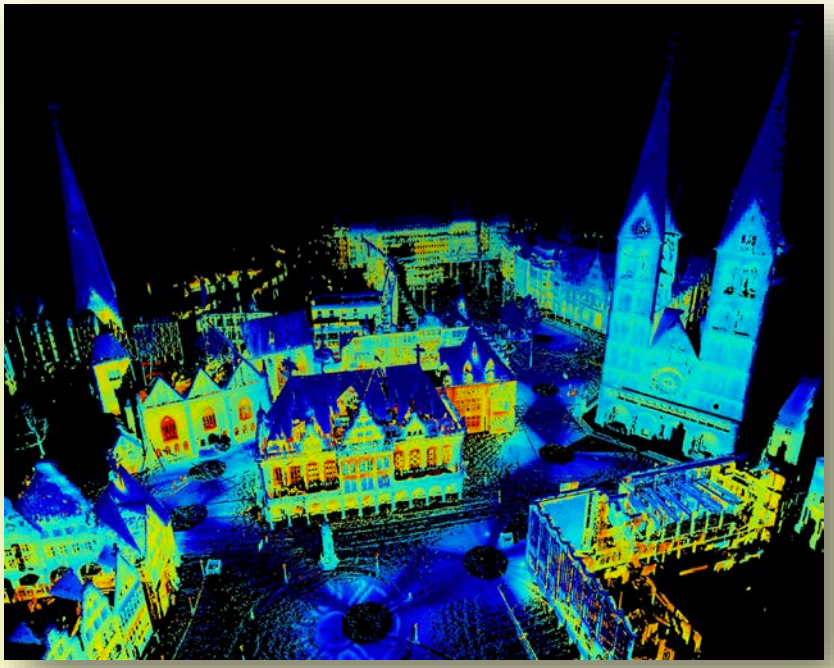

Geometry in Ordinary Images and Videos

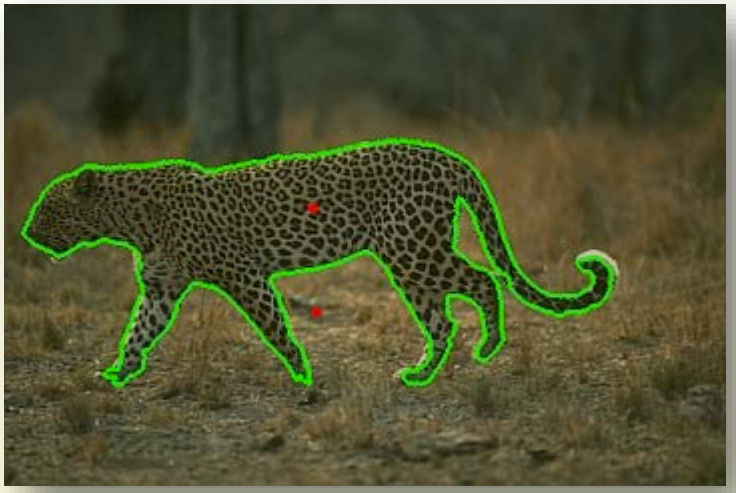

Object recognition

Segmentation involves geometry

Motion involves geometry

Optical flow

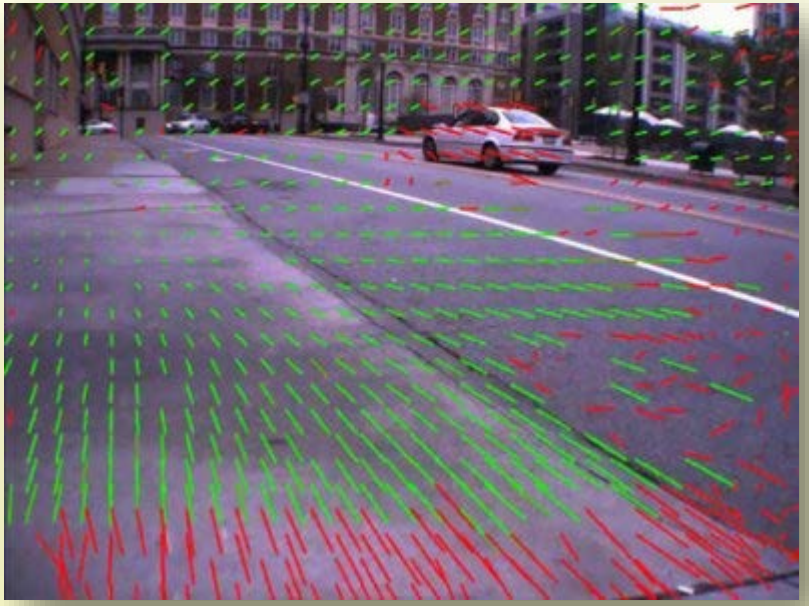

Non-Geometric Data

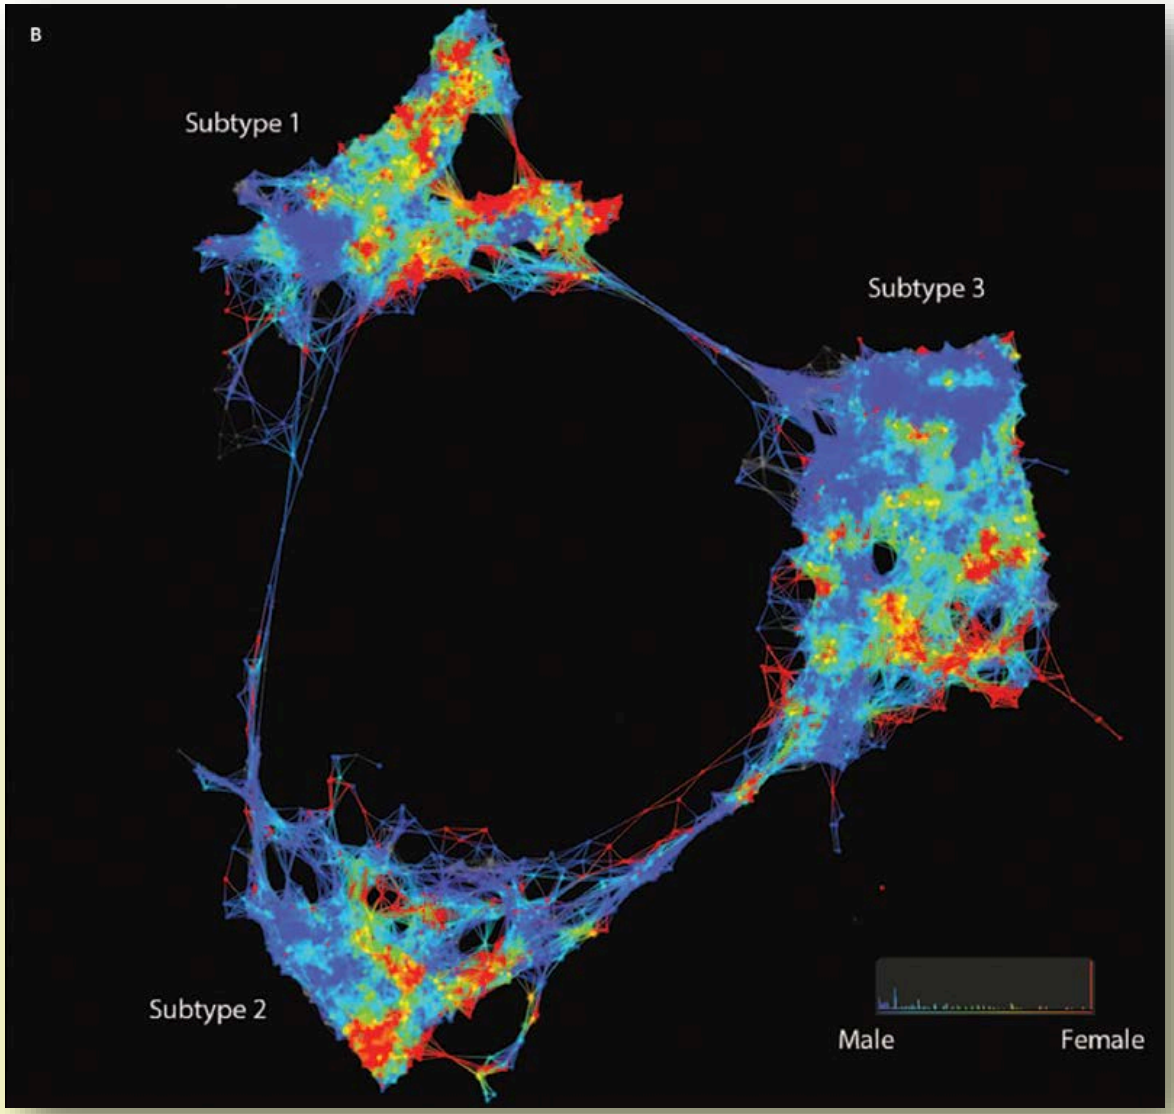

Diabetes II subtypes

Key Course Goals

- Cover basic tools for geometric and topological data analysis, both supervised and unsupervised
- Present mathematical ways to encode and transfer “knowledge” about data
- Present methods for joint data analysis – benefiting from the “wisdom of the collection”

Each Data Set Is Not Alone

- The interpretation of a particular piece of geometric data is deeply influenced by our interpretation of other related data

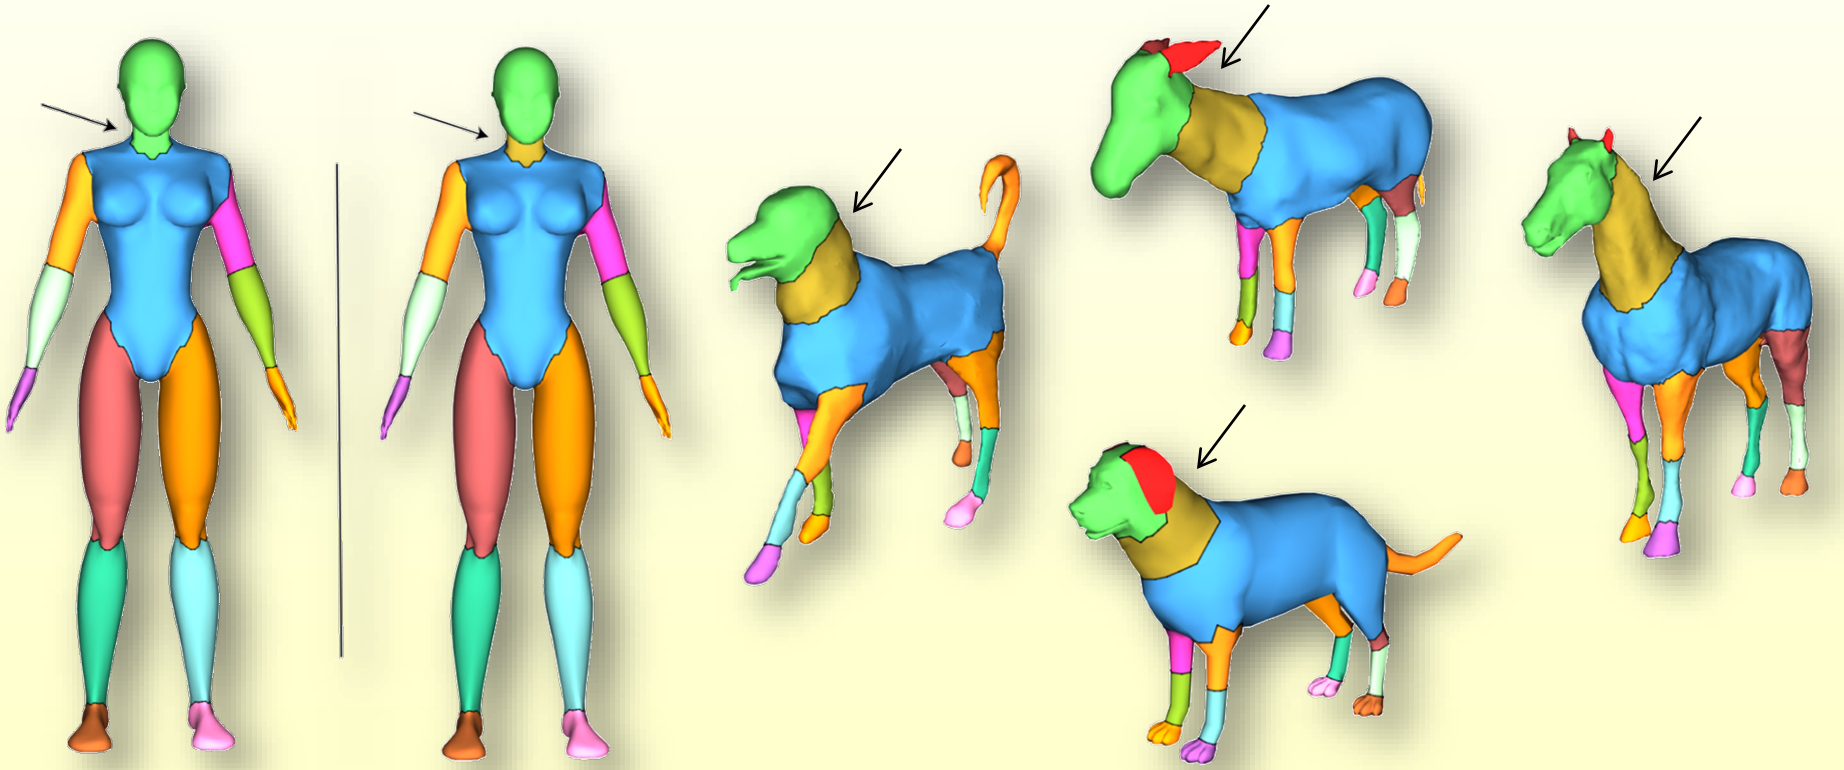

3D Segmentation

Look at Data from Many Views

Caveats

- 2nd time course is offered – still some rough edges
- Will cover a wide variety of techniques: LA, ML, Stat, optimization, geometry processing, …
- Data sets will be primarily visual/geometric (images, 2D/3D point clouds, meshes, or CAD models)

Prerequisites / Overlaps

- Presumes some knowledge of linear algebra, optimization, algorithmic thinking, basic geometry …
- As compared to CS229, STATS216, much greater emphasis on diverse data types (e.g., point clouds, graphs), in addition to the geometry / topology angle

The Linear Space View of Data

Unsupervised Methods

Embedding Data into a Euclidean Space

- Attributes or features can be used to map data to a Euclidean space

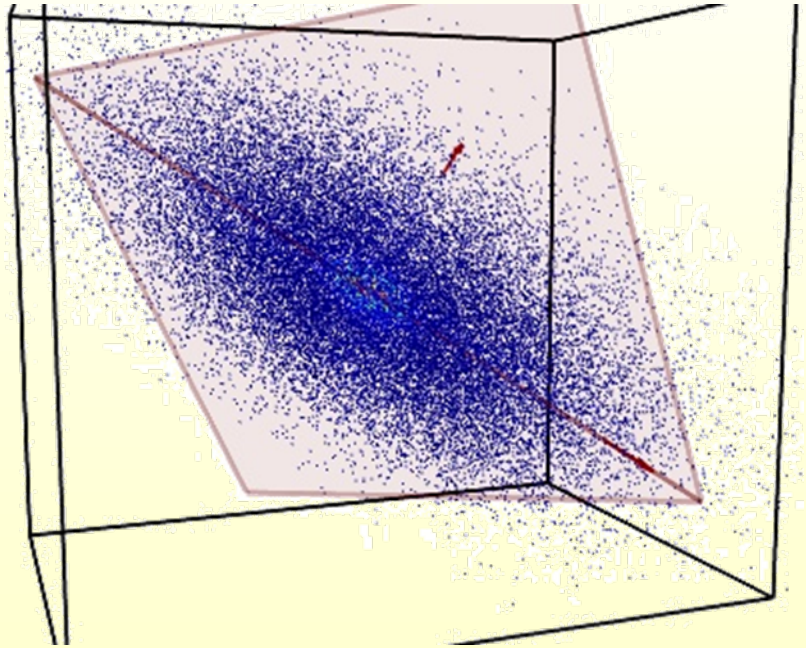

Direct Embeddings

- Input

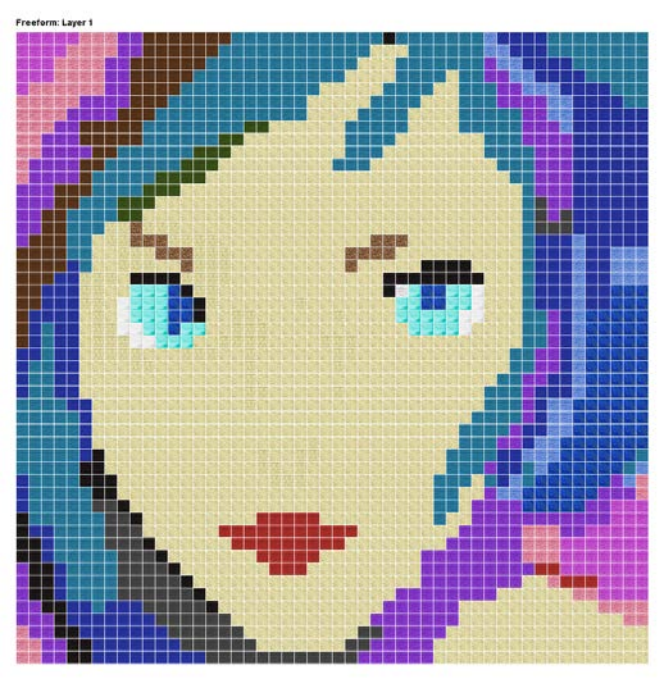

- Latent code

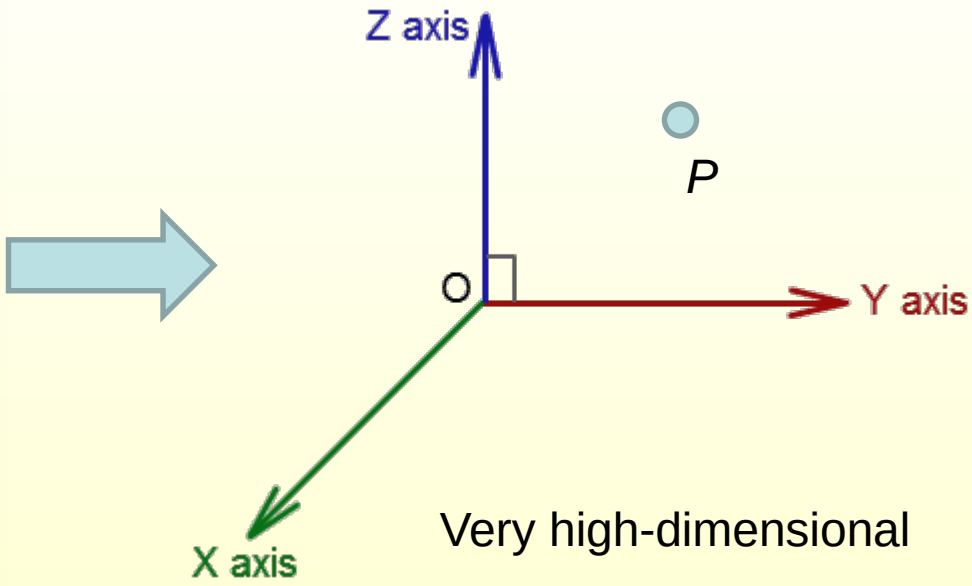

Very high-dimensional

Linear Space Methods

- Principal components analysis (PCA)
- Canonical correlation analysis (CCA)

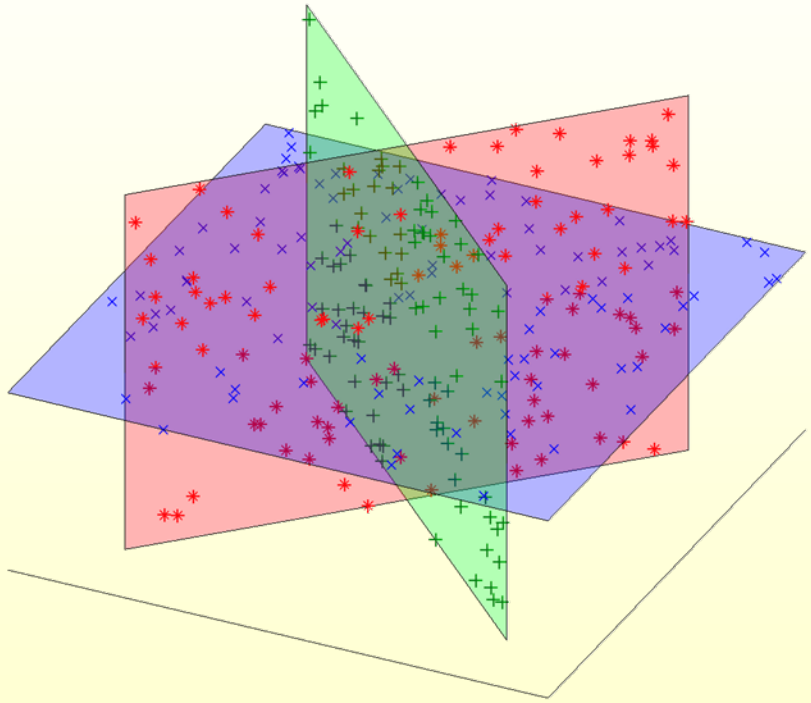

Dimensionality reduction

Data as Points on a Manifold

- Non-linear dimensionality reduction
- Low-d inside high-d

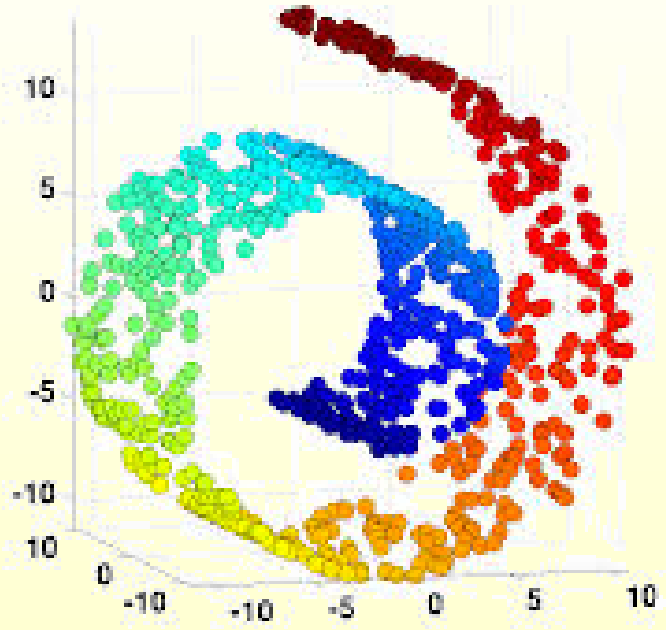

Isomap, locally linear embeddings, Laplacian eigenmaps

Deep Learning

Supervised Methods

Deep Learning Networks

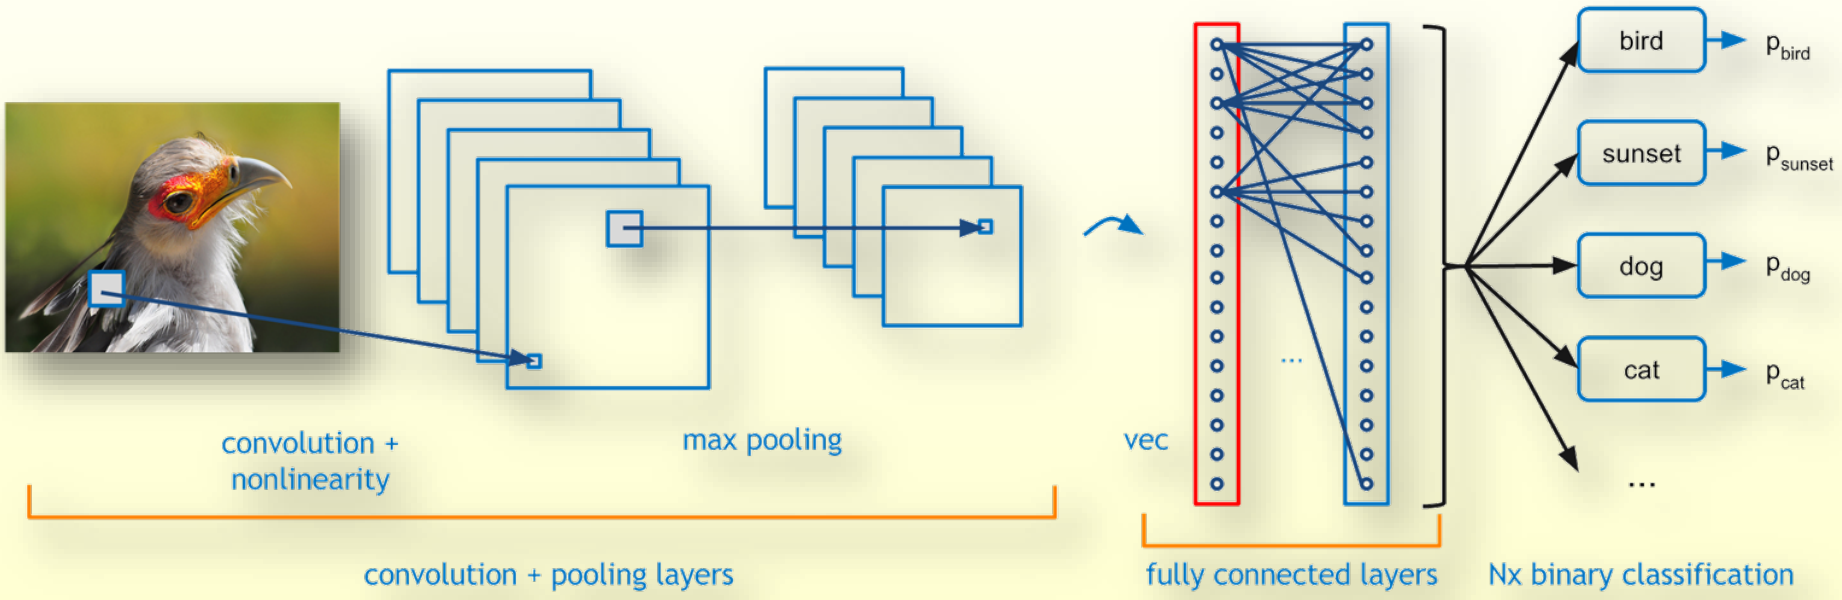

Classification

Success Made Possible By ....

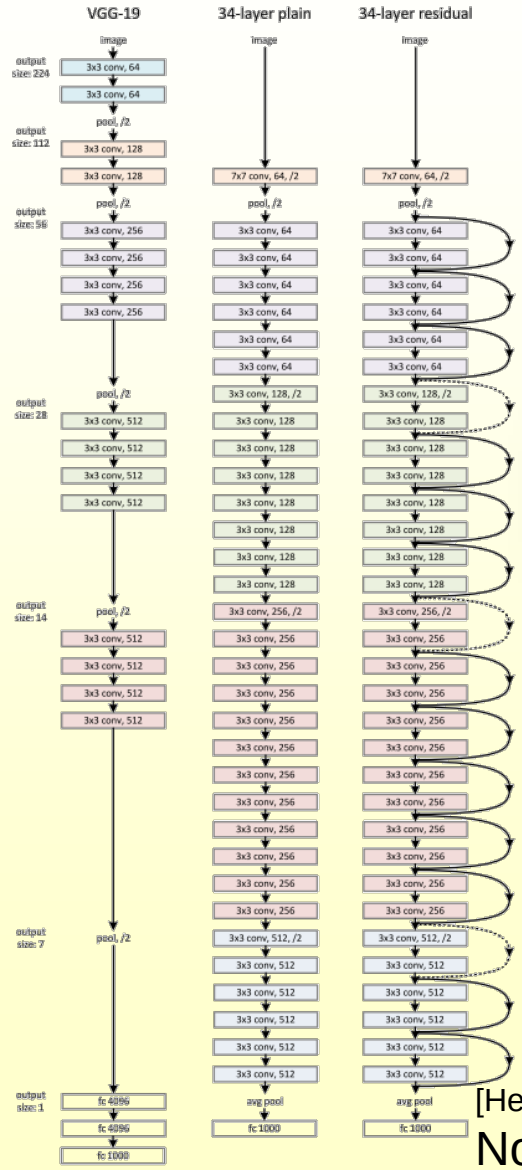

[He et al., 2015]

Novel deep architectures

Plenty of annotated data

Lots of computing power

Data Abstraction Across Layers

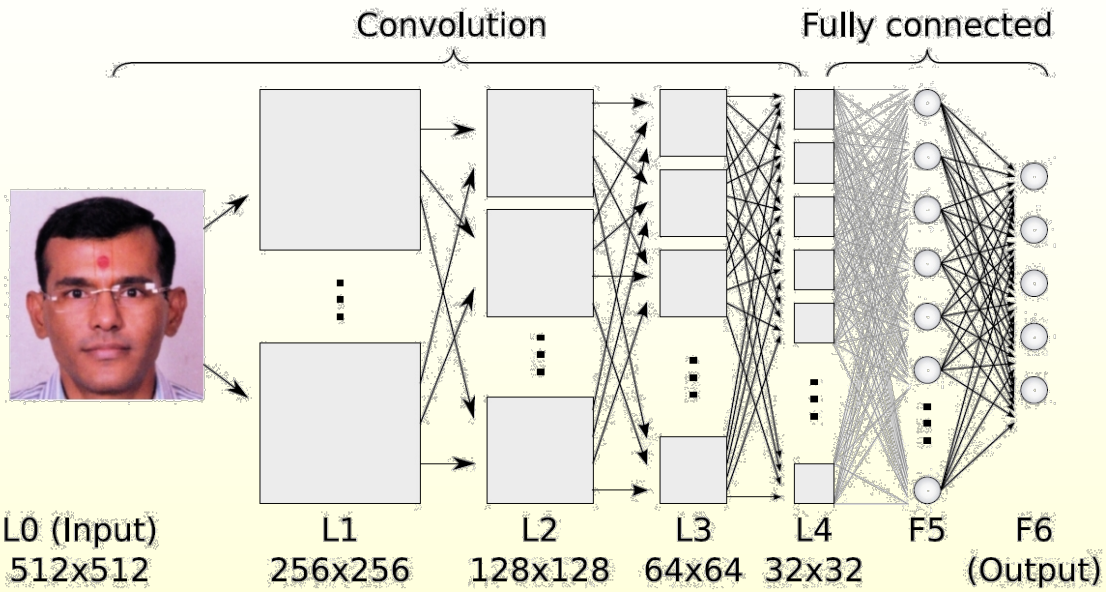

[Makwana, 2016]

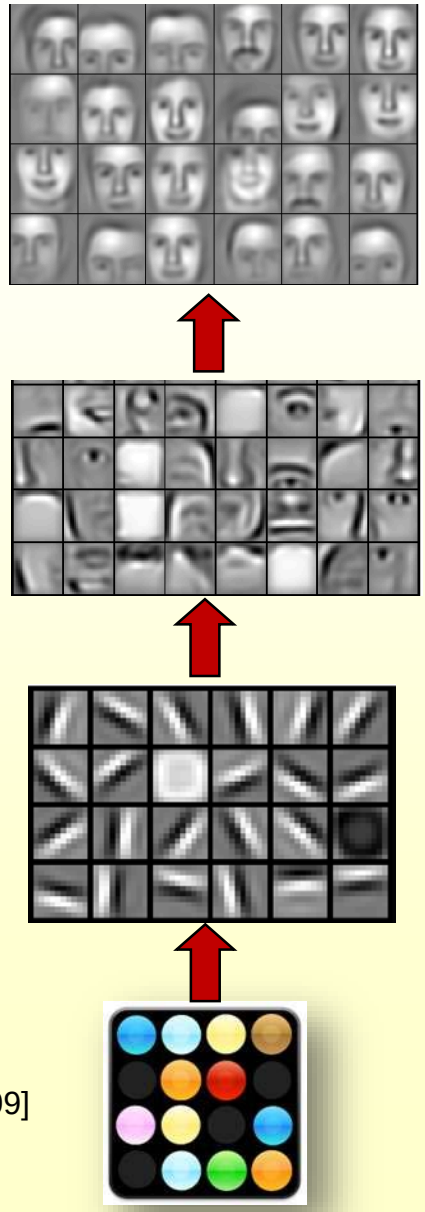

Data-driven feature learning at ascending abstraction layers

"Vertical" networks

[Lee et al., 2009]

Learning Pipelines

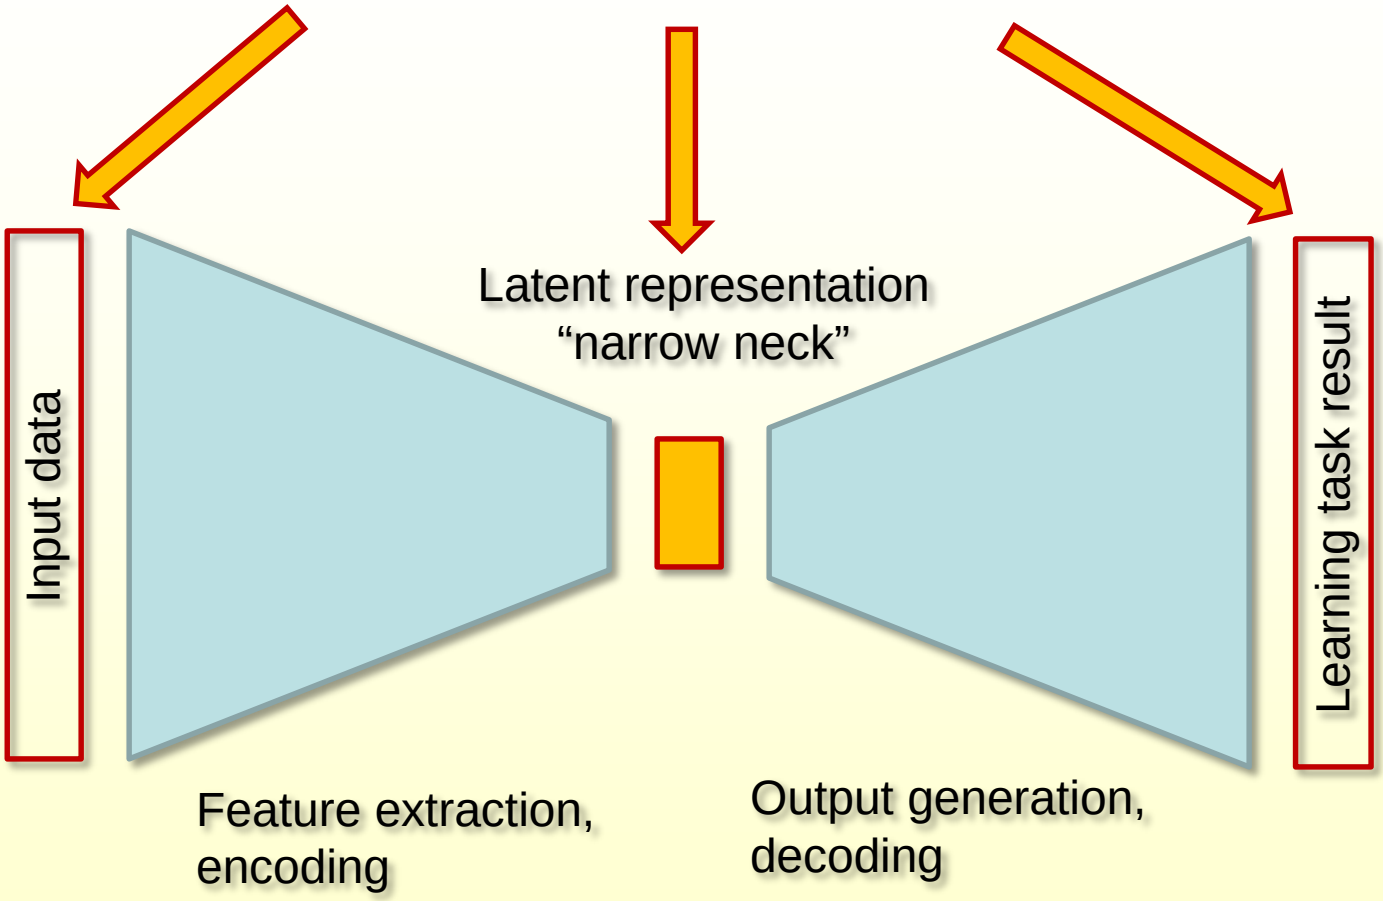

A latent code (= a point in a Euclidean feature space) acts as a low-d proxy for high-d input data, w.r.t. a learning task.

Input Diversity Limitations

Convolutional Networks

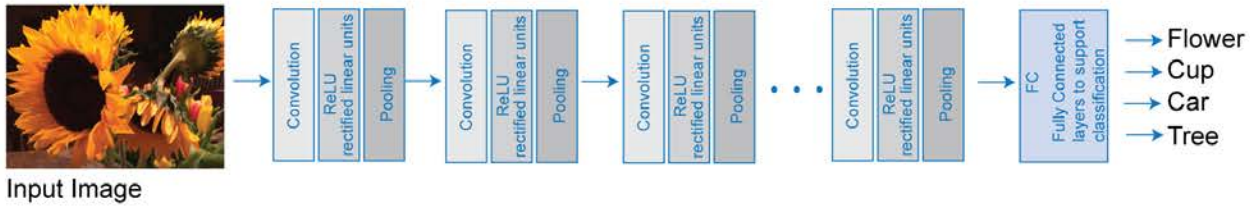

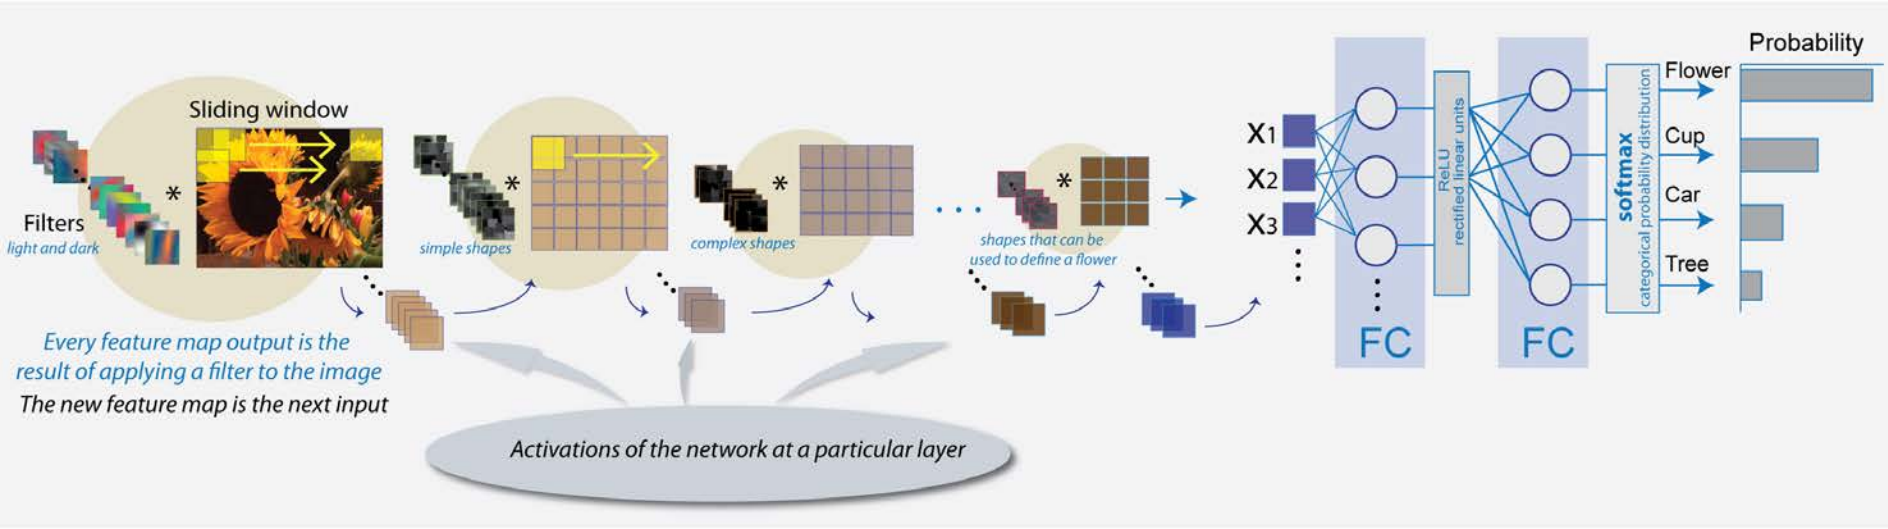

Convolutions require regular array type inputs for weight sharing and other optimizations

Deep Learning for 3D Geometry Analysis

Classification

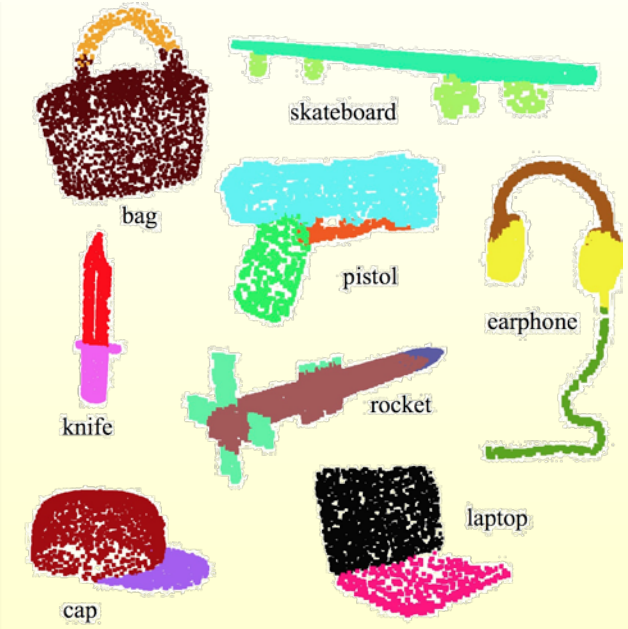

Segmentation/Parsing (object/scene)

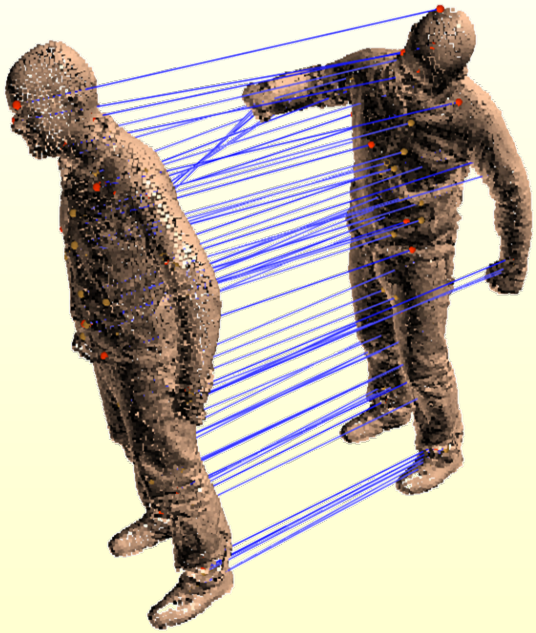

Correspondences

Multiple Representations for 3D

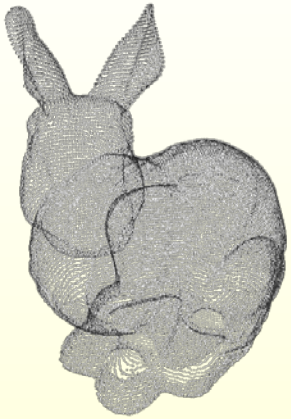

Point Cloud

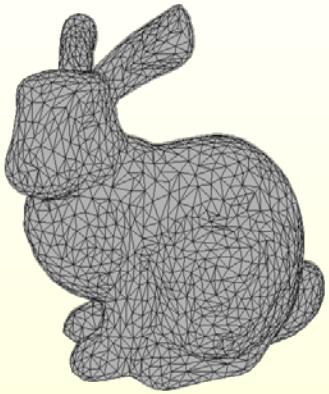

Mesh

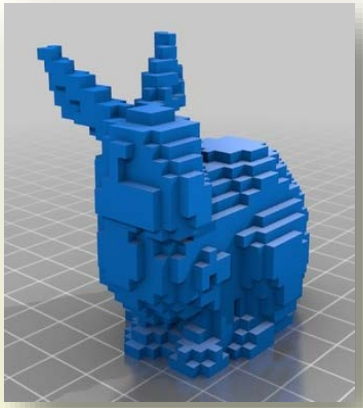

Volumetric

Projected View

RGB(D)

...

These are irregular representations

PointNet: Working Directly with Point Cloud Data

- Goal: design a NN architecture that can work directly with point clouds
- Must deal with unstructured, unordered data

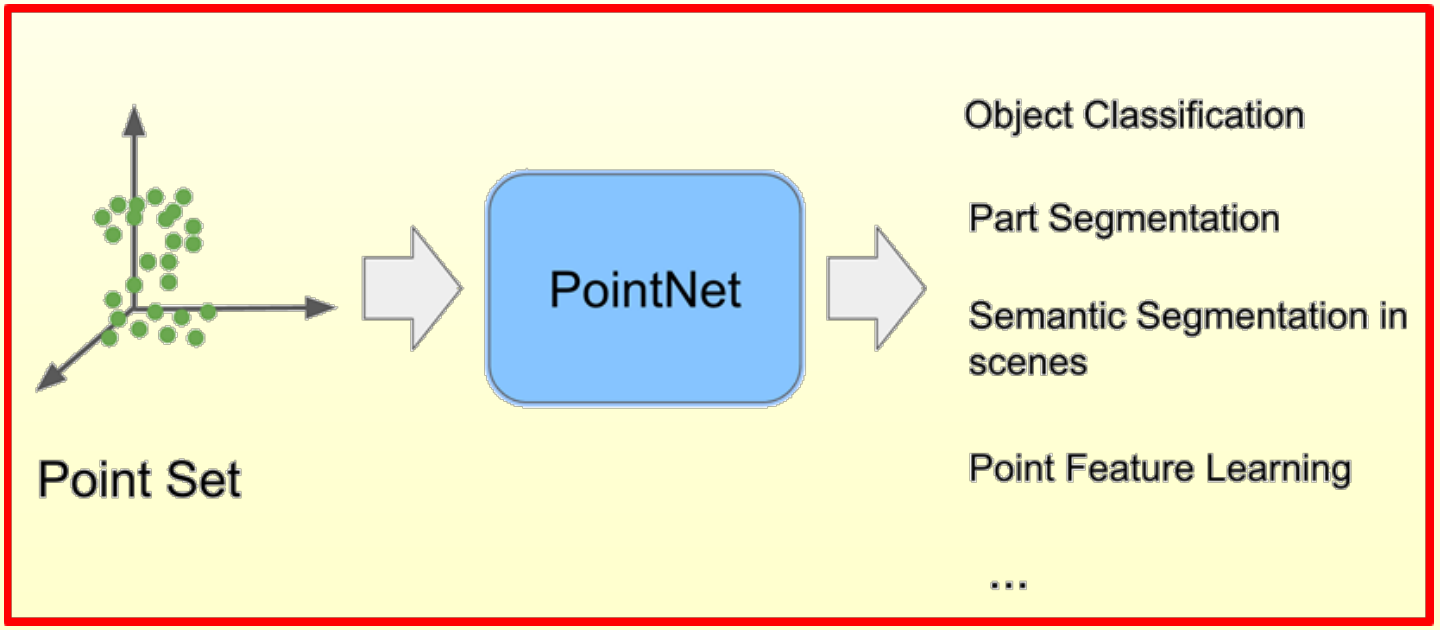

The Graph View of Data

Graphs and Simplicial Complexes

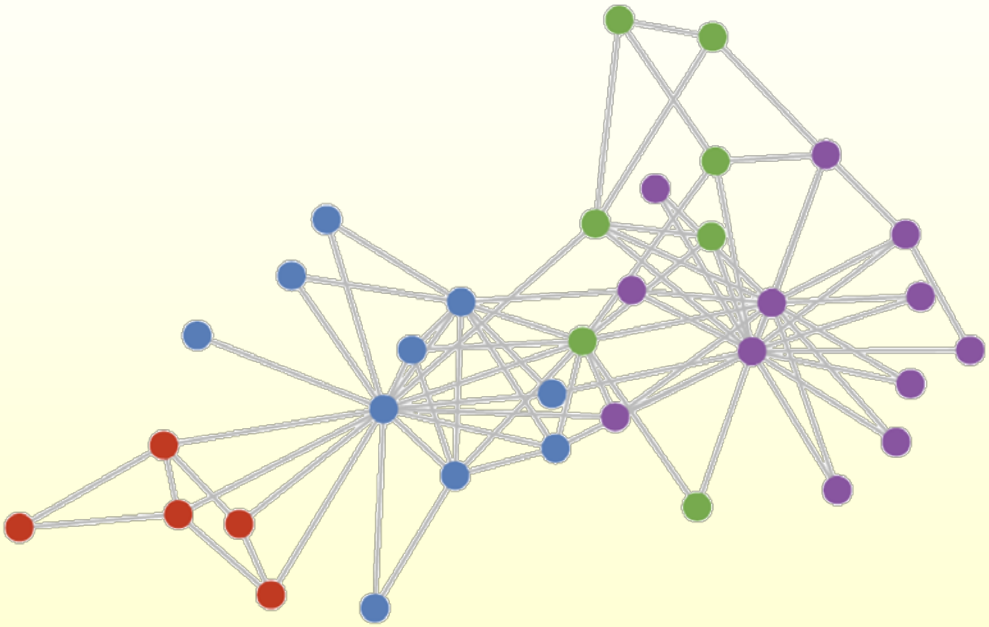

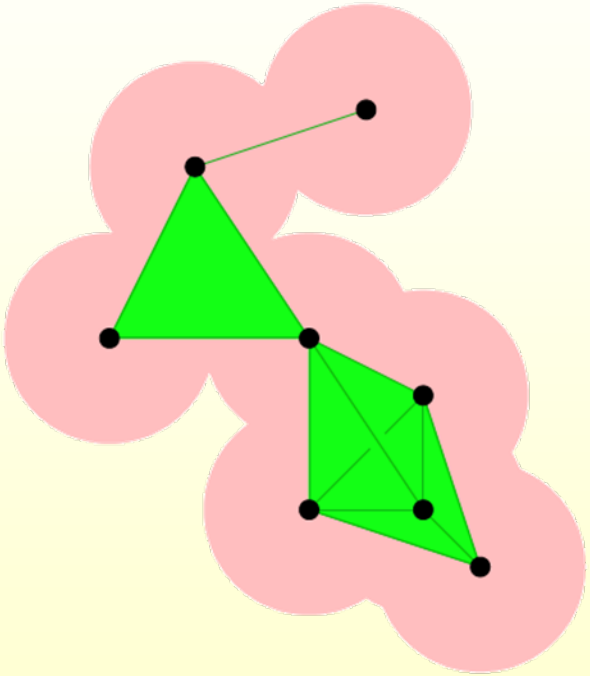

A Graph View of Data

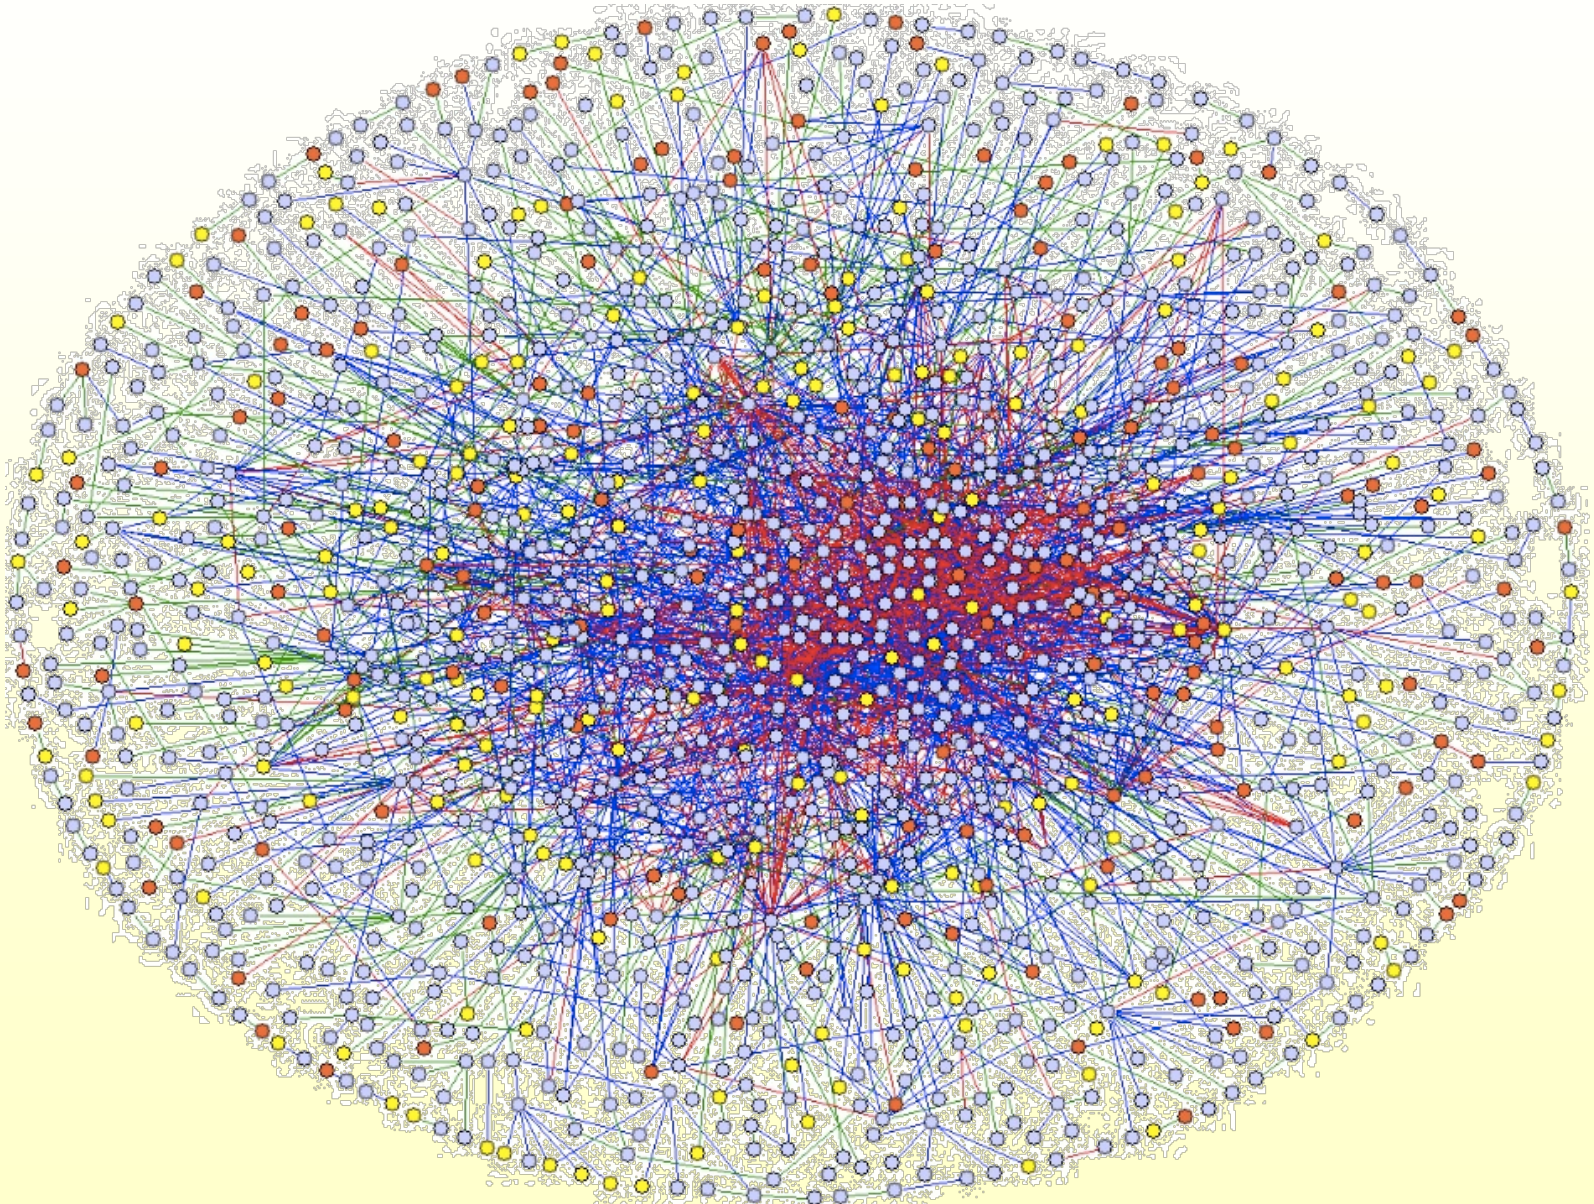

Spectral Graph Theory

- Linking the graph-theoretic and linear algebraic view of data.

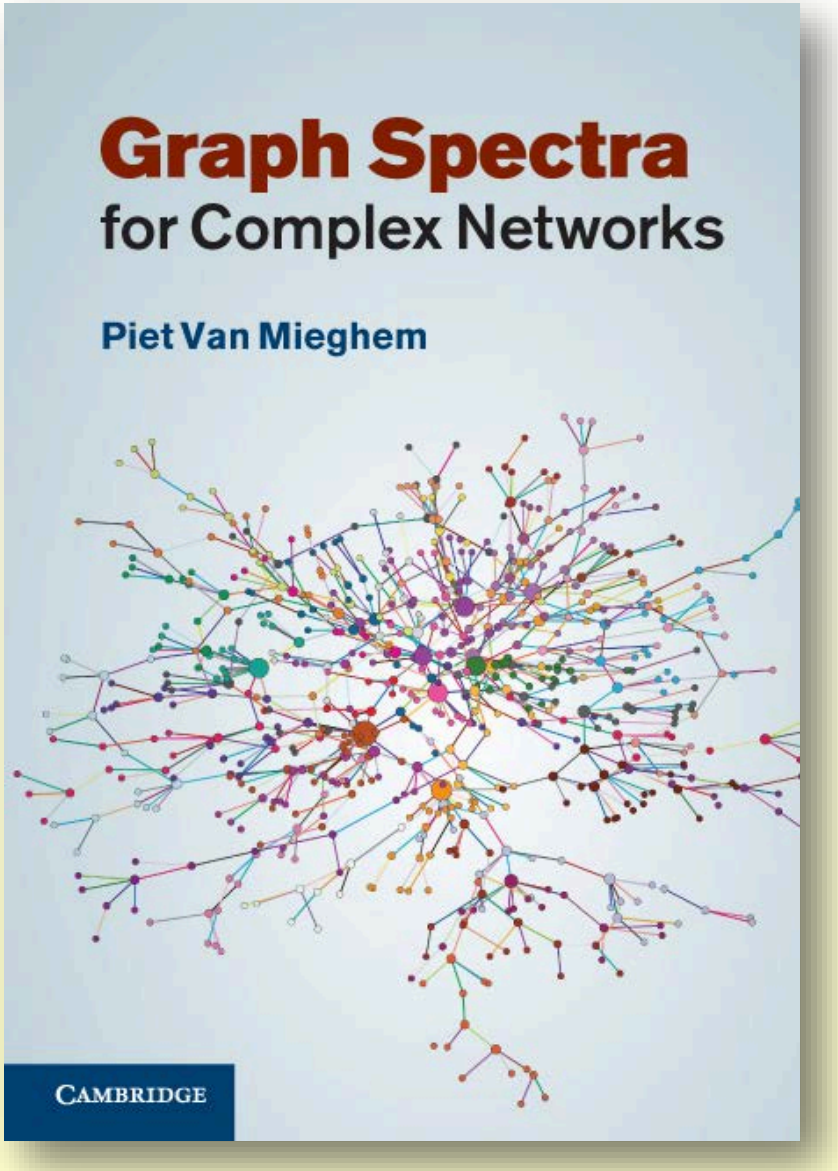

From Feature Vectors to Feature Sets

- What to do if the number of features per data set is variable?
- Distance or similarity metrics
- Multidimensional scaling

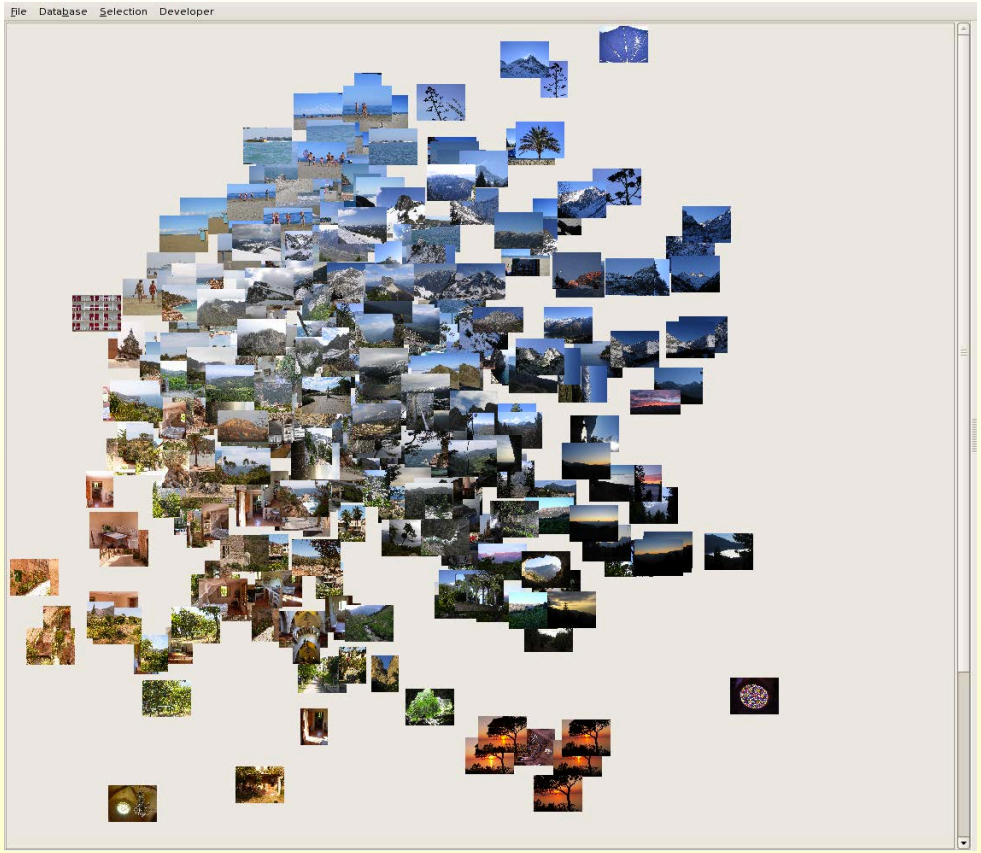

A “Procrustean” approach

Topological Data Analysis

Computational Topology

- Topology is the branch of mathematics that does not take distances too seriously.

G. Carlsson

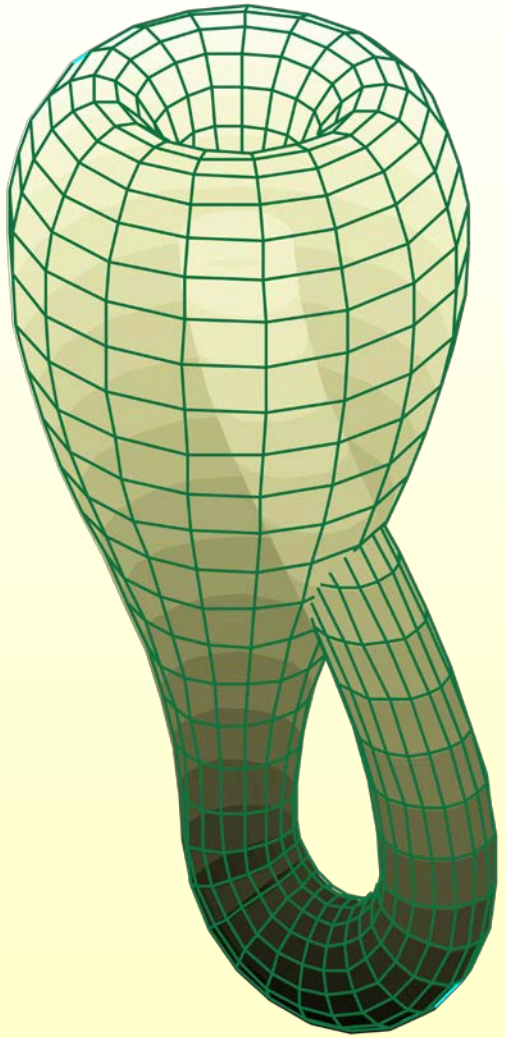

- Large distances are often suspect …

In TDA, Sampled Spaces

"The Shape of Data"

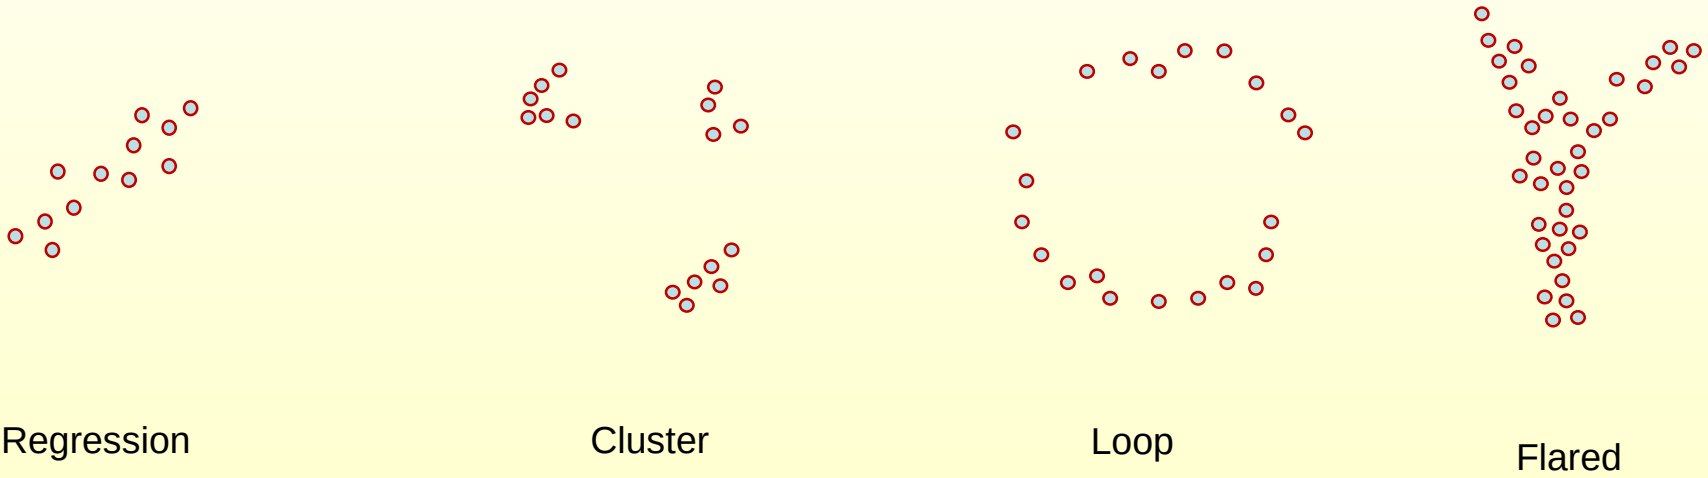

Regression

Cluster

Loop

Flared

A Multiscale View of Data

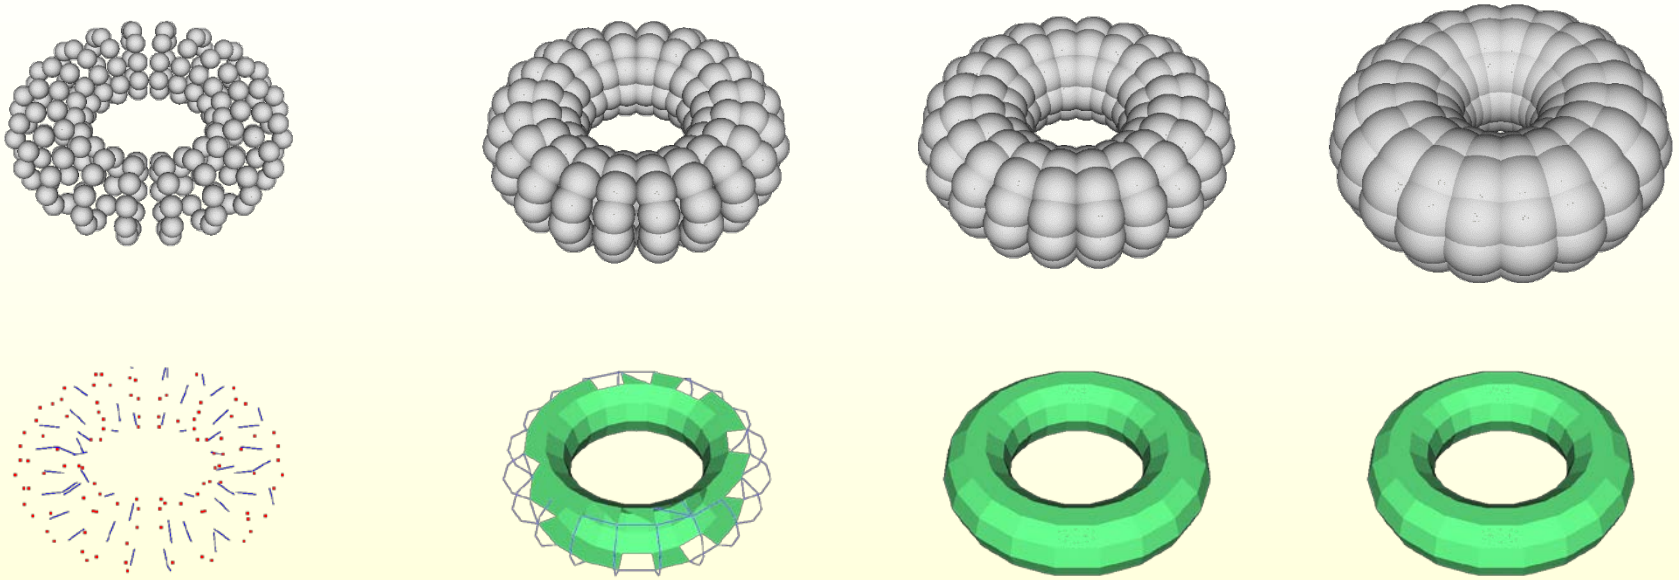

$\beta_0 = 150$
$\beta_1 = 0$

$\beta_0 = 1$
$\beta_1 = 37$

$\beta_0 = 1$
$\beta_1 = 2$

$\beta_0 = 1$
$\beta_1 = 1$

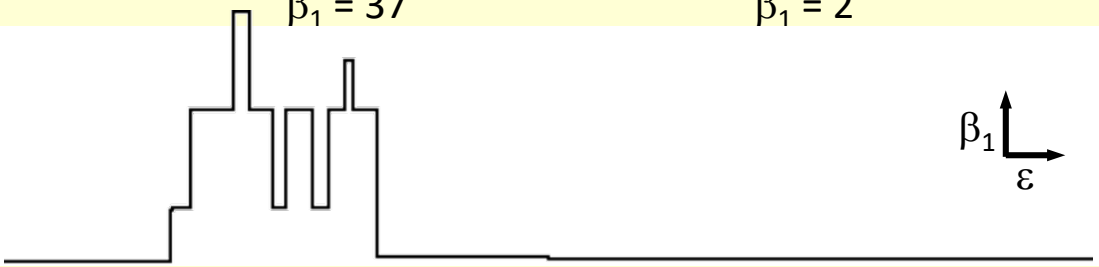

Persistent Homology

- Persistence of topological features
- Barcodes and persistence diagrams
- Data embeddings into non_Euclidean spaces

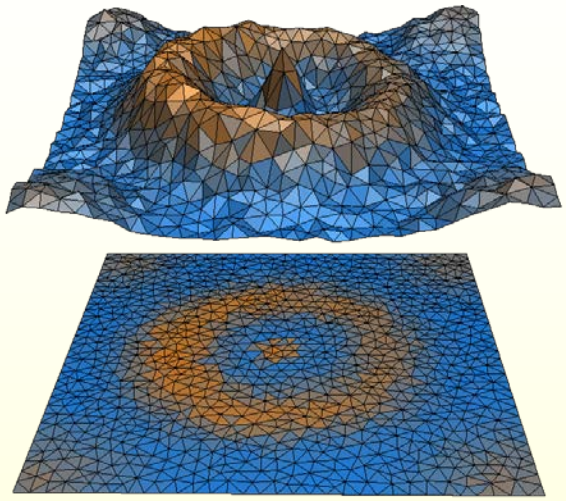

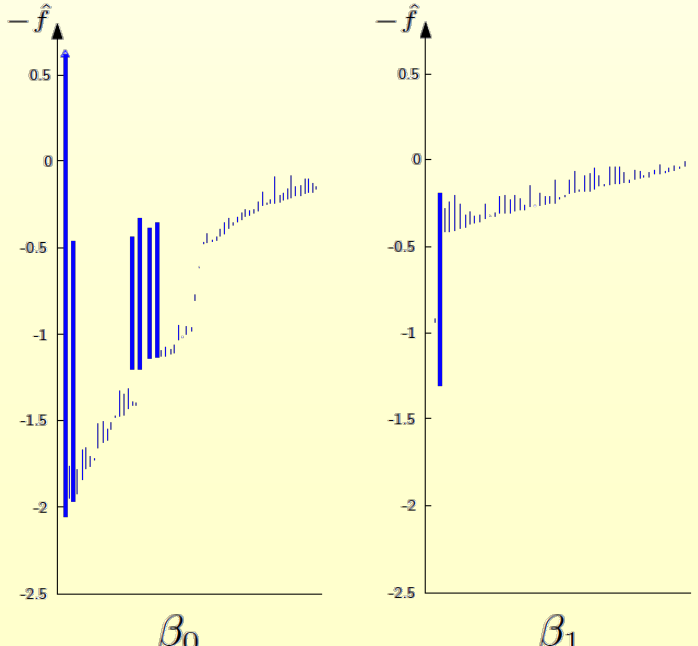

Functions as Filters

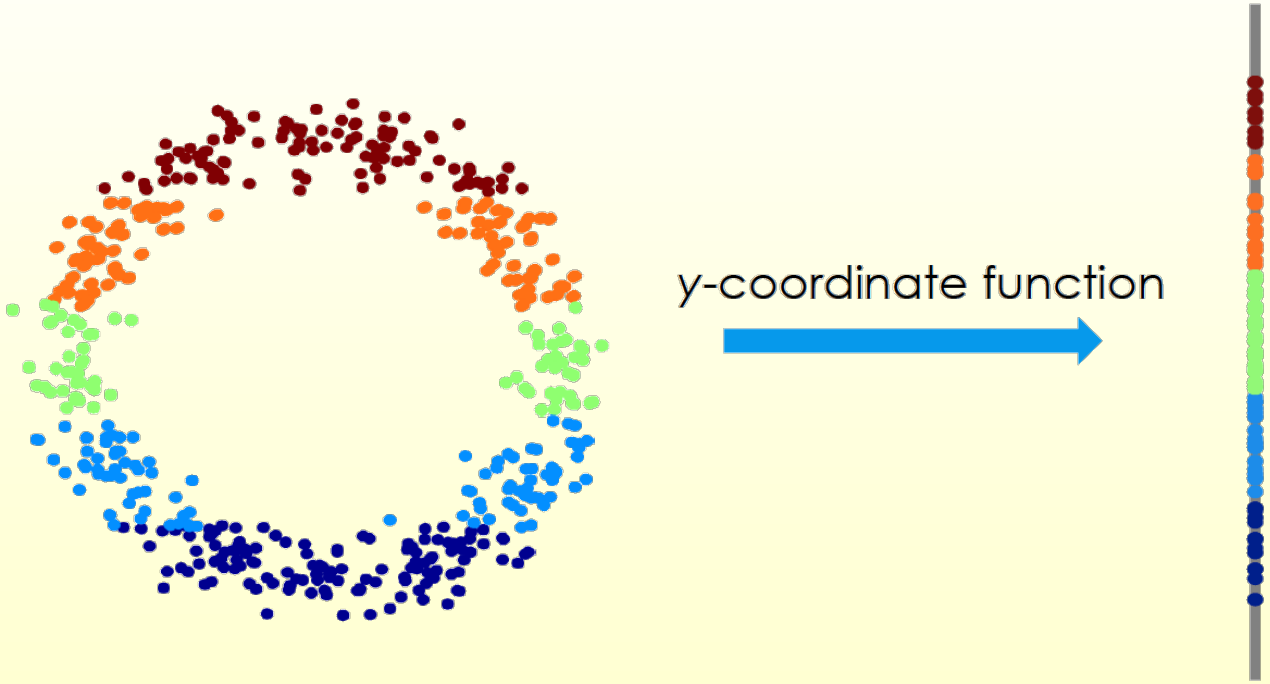

A Simplicial Complex

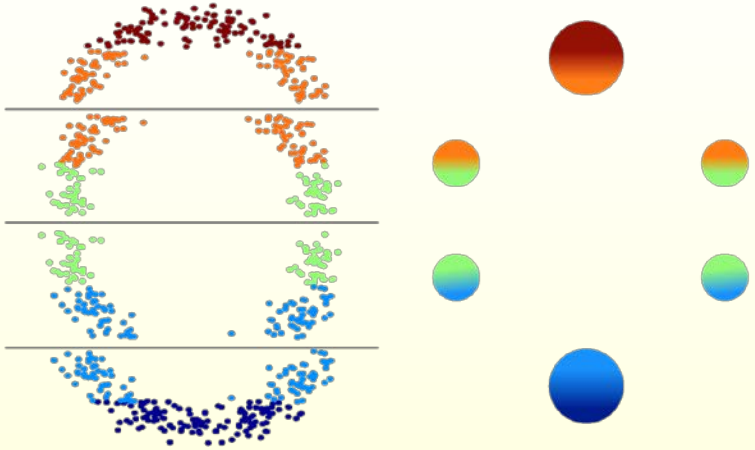

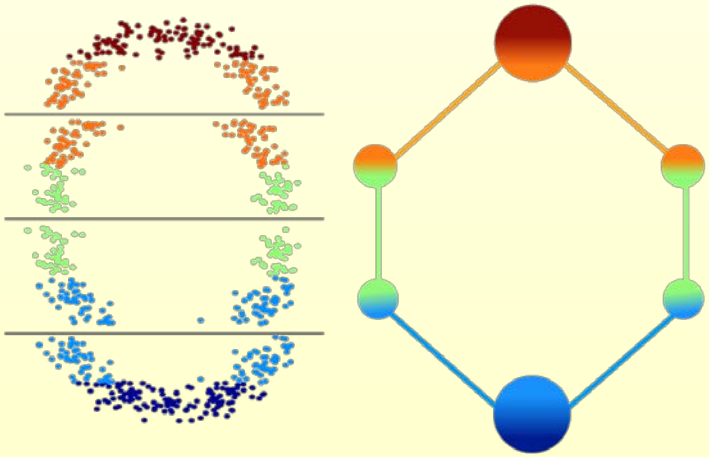

Functions as Lenses on Data

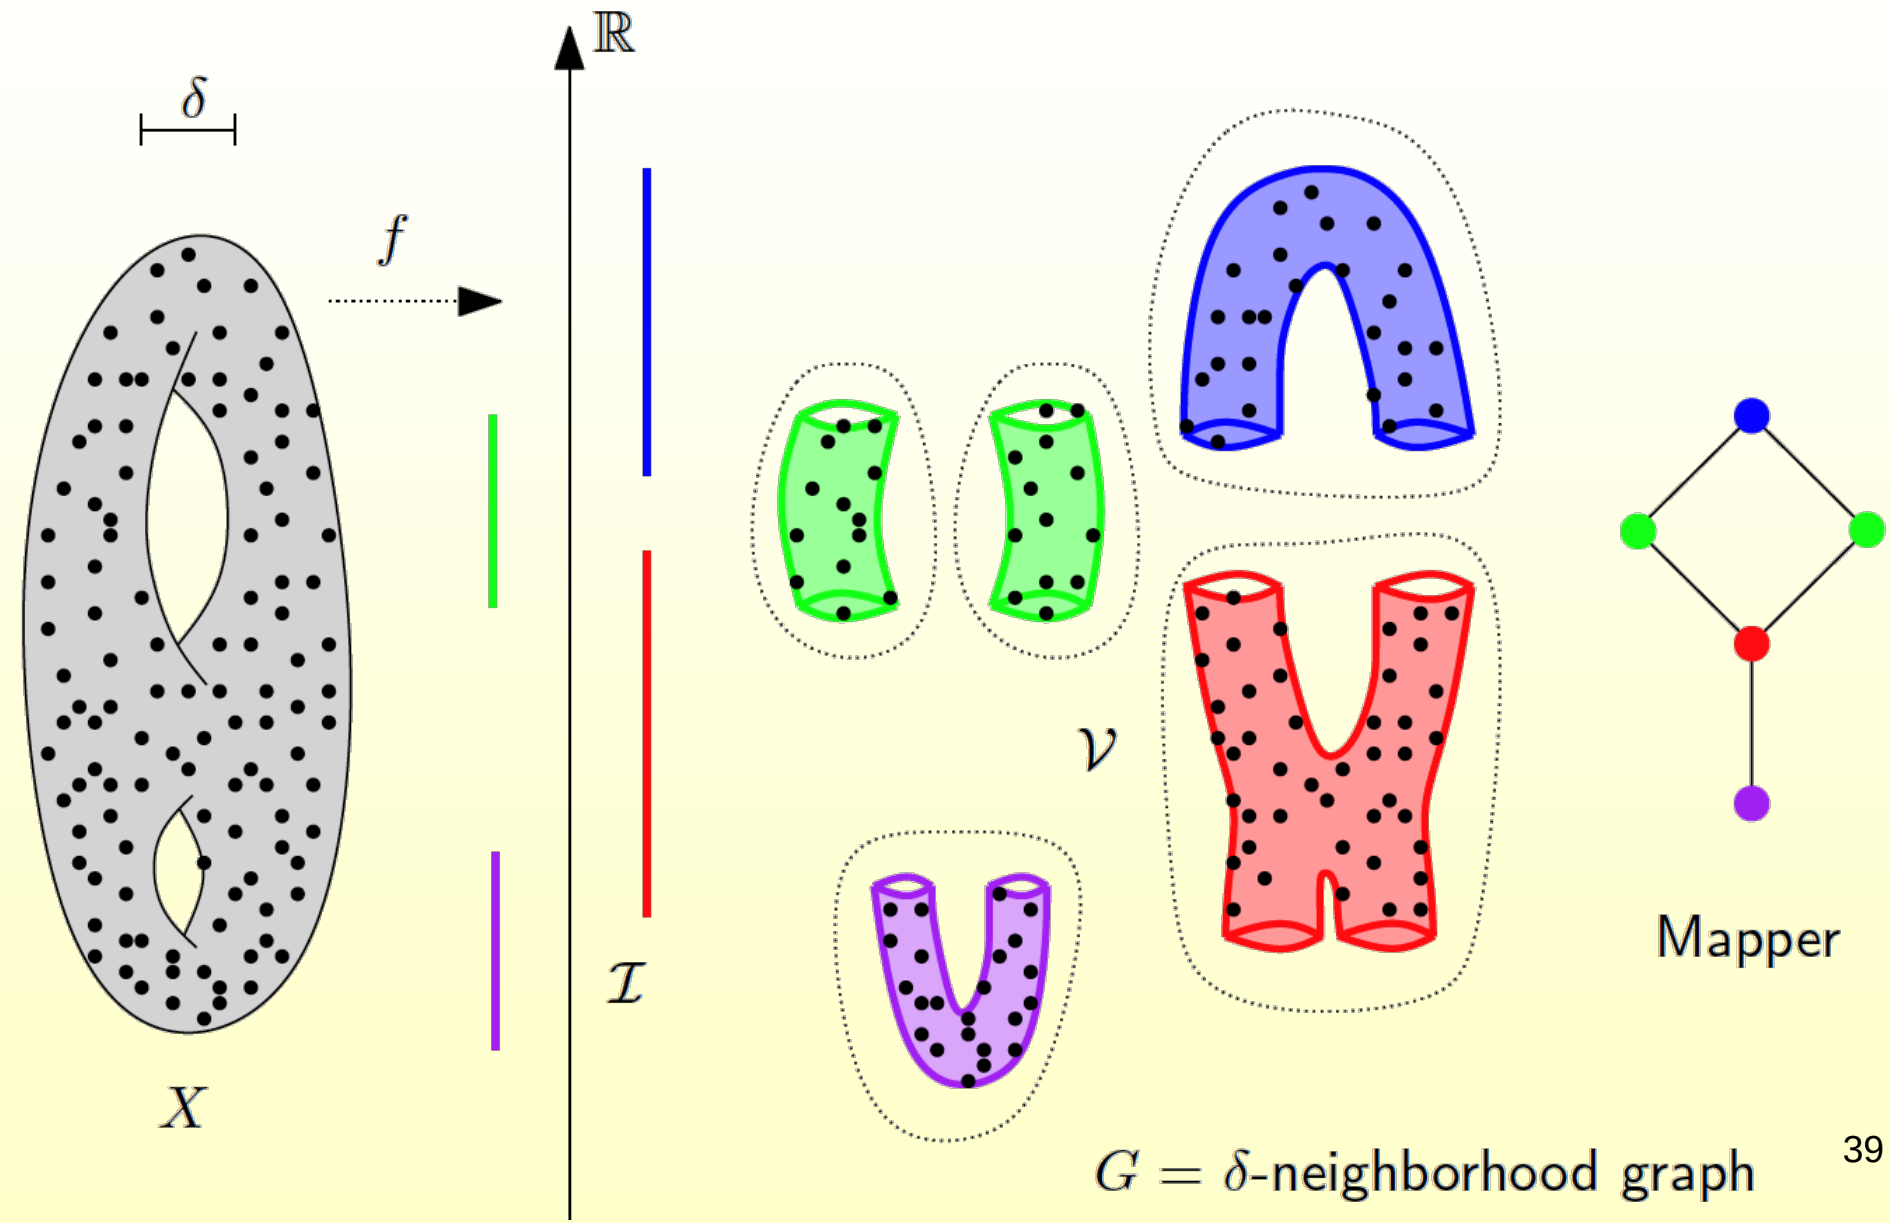

Some Limitations of the "Point" View of Data

(Too) Many Latent Spaces

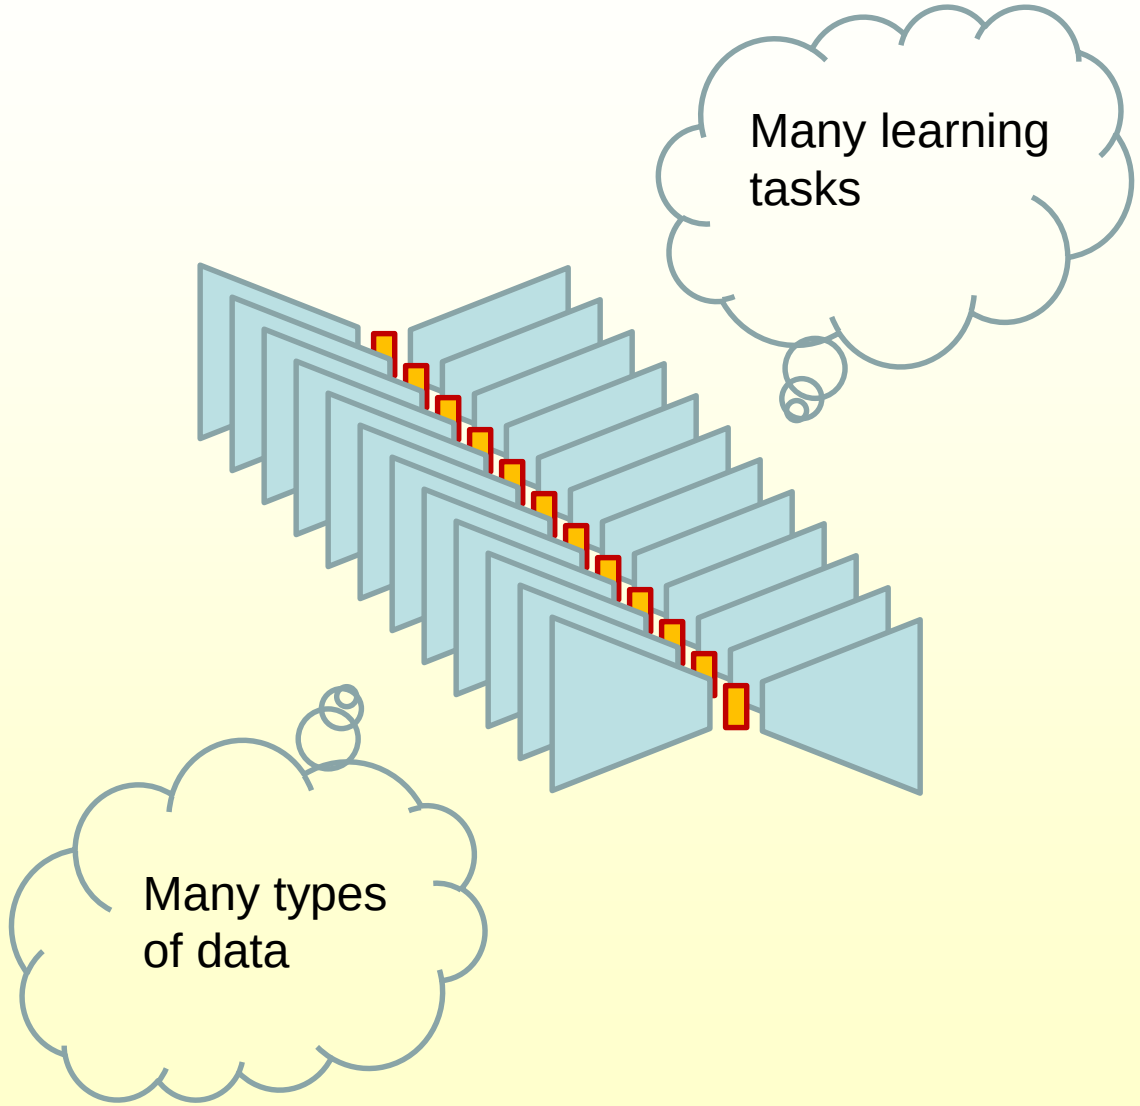

Codes

Latent vectors

Parametrizations

Representations

…

A Shared "Narrow Neck / Waist" in Other Fields

Power distribution

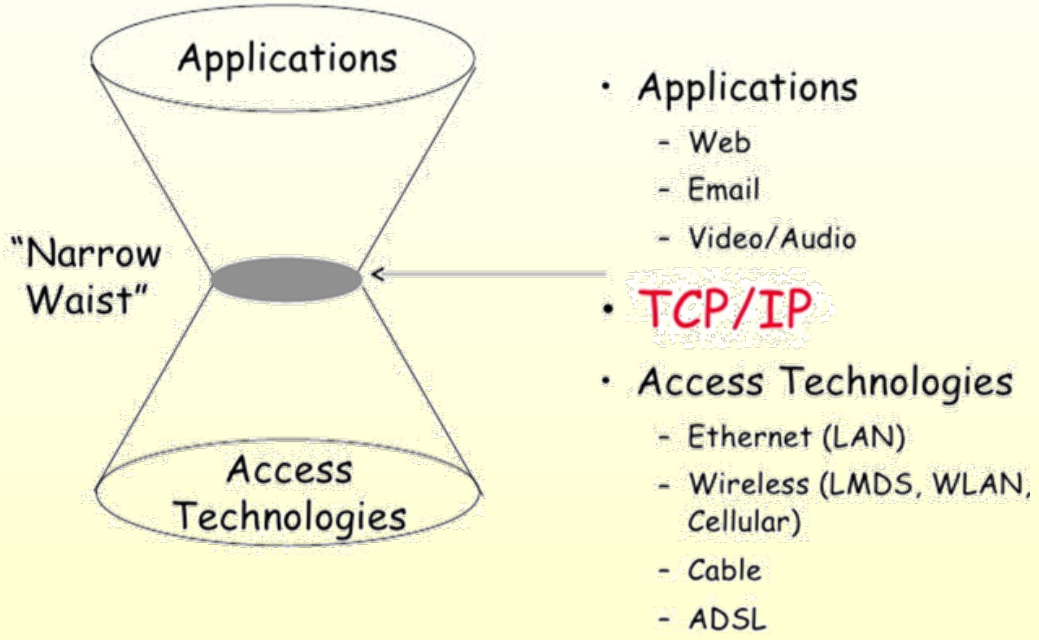

Network protocol stack

[Image: R. Katz]

The Structure of Latent Spaces

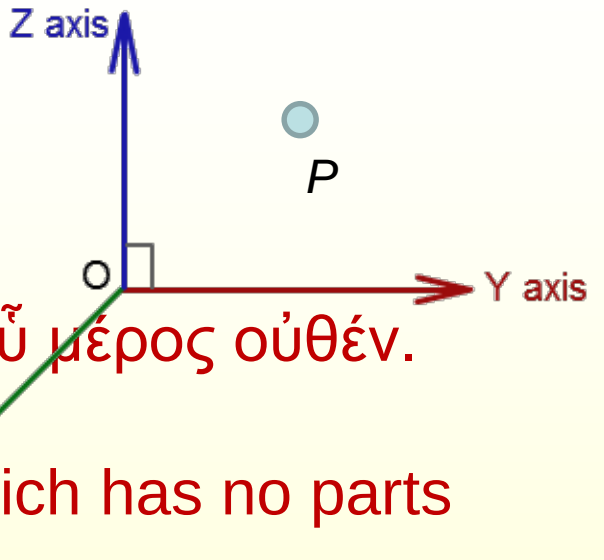

α΄. Σημεῖόν ἐστιν, οὗ μέρος οὐθέν.

a. A point is that which has no parts

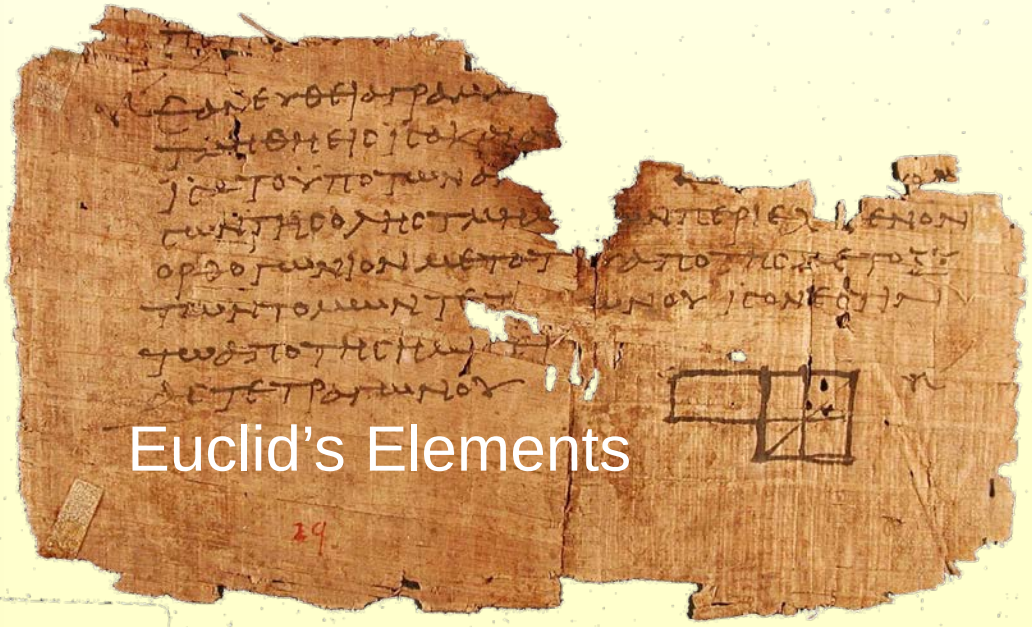

Euclid's Elements

Homomorphic Encyption

Chapter 2
Homomorphic Encryption

Abstract Homomorphic encryption is a form of encryption which allows specific types of computations to be carried out on ciphertexts and generate an encrypted result which, when decrypted, matches the result of operations performed on the plaintexts. This is a desirable feature in modern communication system architectures. RSA is the first public-key encryption scheme with a homomorphic property. However, for security, RSA has to pad a message with random bits before encryption to achieve semantic security. The padding results in RSA losing the homomorphic property. To avoid padding messages, many public-key encryption schemes with various homomorphic properties have been proposed in last three decades. In this chapter, we introduce basic homomorphic encryption techniques. It begins with a formal definition of homomorphic encryption, followed by some well-known homomorphic encryption schemes.

2.1 Homomorphic Encryption Definition

In abstract algebra, a homomorphism is a structure-preserving map between two algebraic structures, such as groups.

A group is a set, G, together with an operation $\circ$ (called the group law of G) that combines any two elements a and b to form another element, denoted $a \circ b$. ...ify as a group, the set and operation, $(G, \circ)$, must satisfy four requirements ...up axioms:

...he result of the operation, $a \circ b$, is also in G.

... $(a \circ b) \circ c = a \circ (b \circ c)$.

... G, such that for every element a ...lement is unique, and thus $a \circ b =$

Craig Gentry

The Mathematics of Latent Codes

- What kinds of mathematical objects can serve well as latent space codes?
- How should such an object vary with
  - the input data?
  - the input modality?
  - the learning task?

Relating Data Sets

Horizontal Map Networks

Horizontal Networks

Similarity as a communications channel

34 *The Mathematical Theory of Communication*

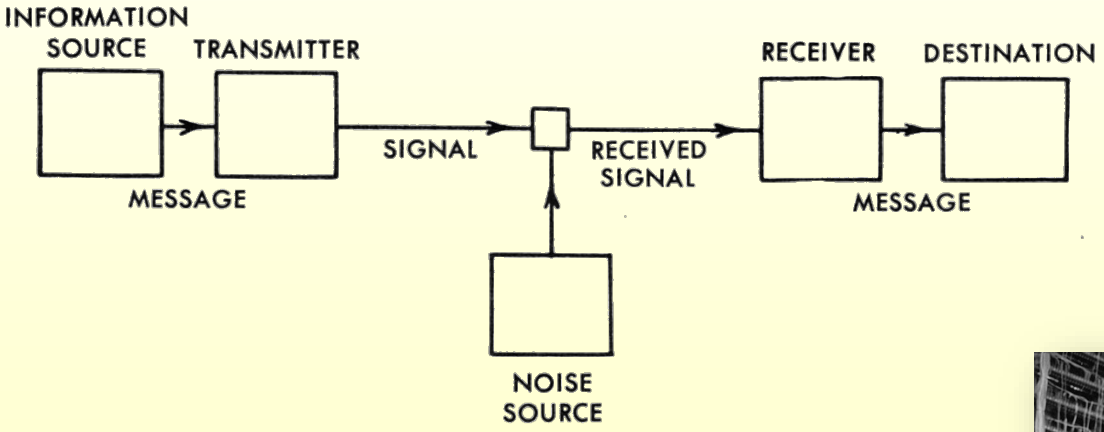

Fig. 1. — Schematic diagram of a general communication system.

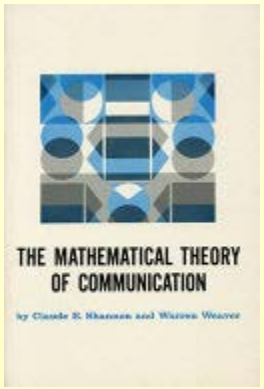

Claude Shannon

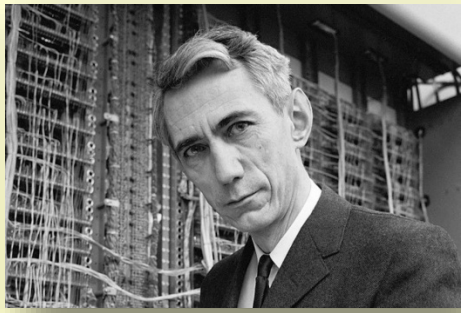

Networks of Visual Data

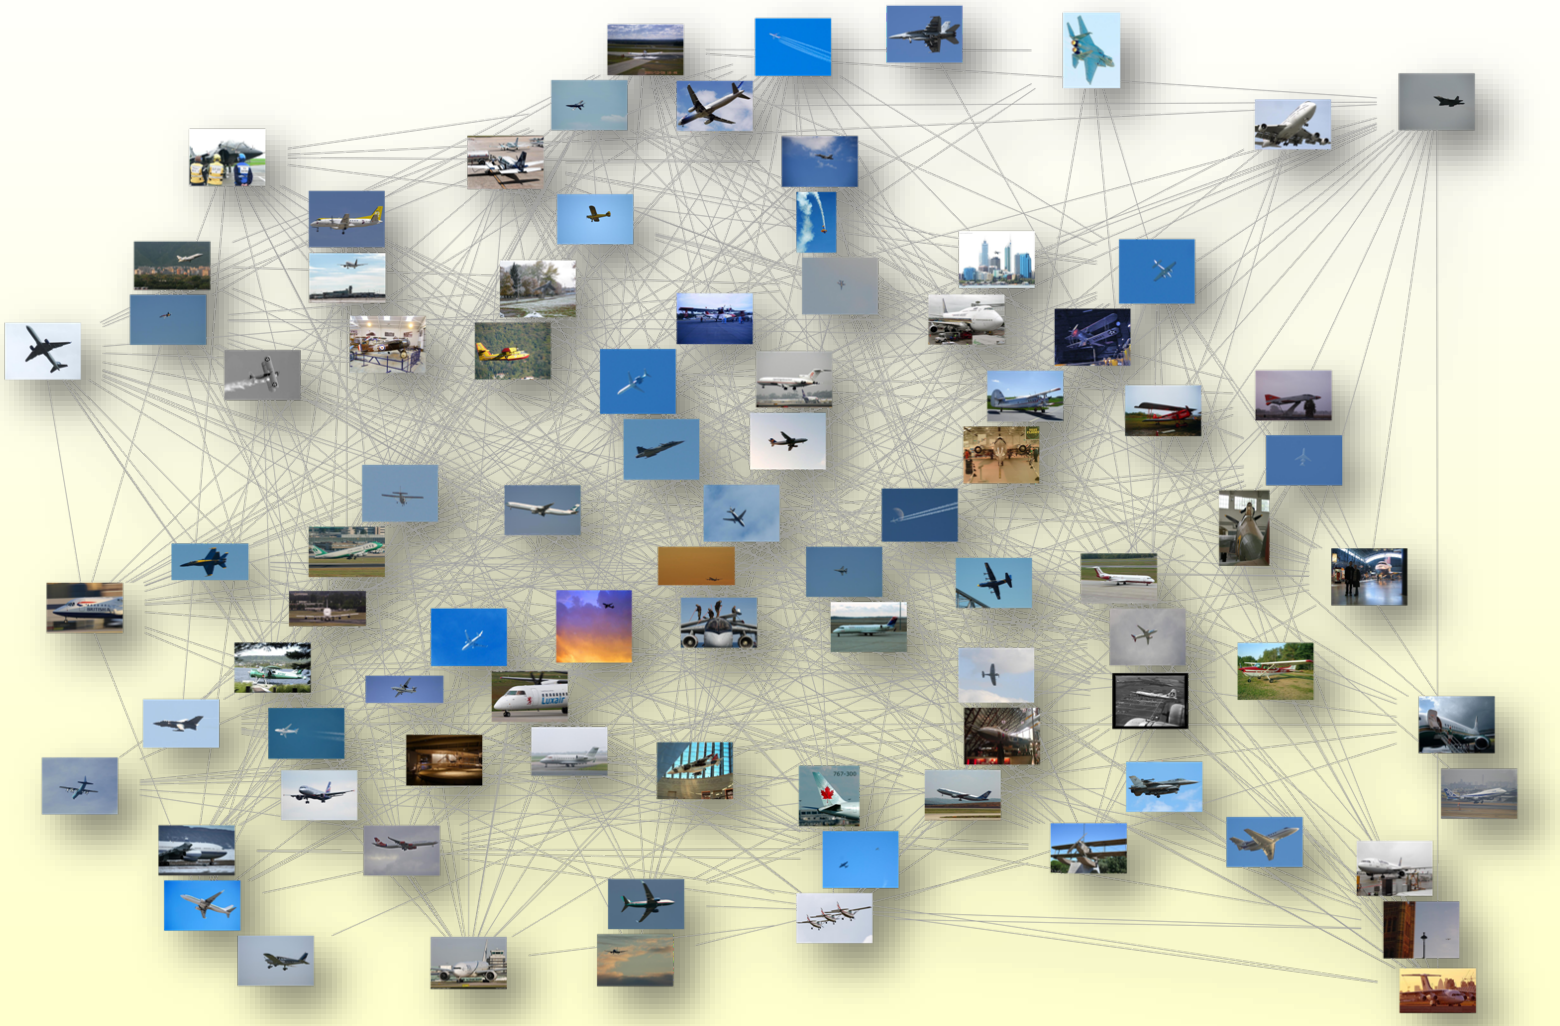

Alignment and Registration

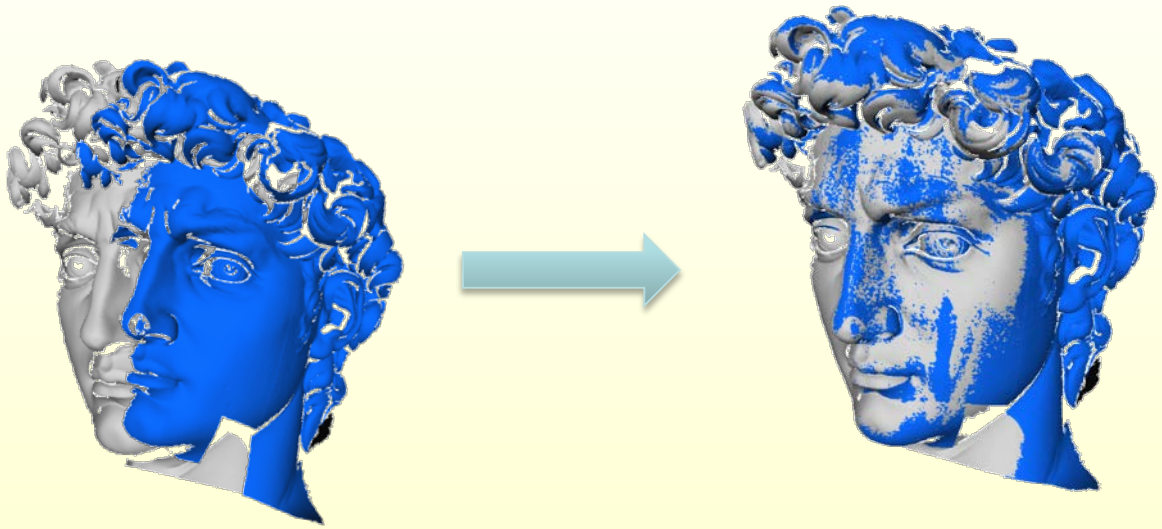

Rigid Registration

Low-dimensional transformation group

Maps and Correspondences

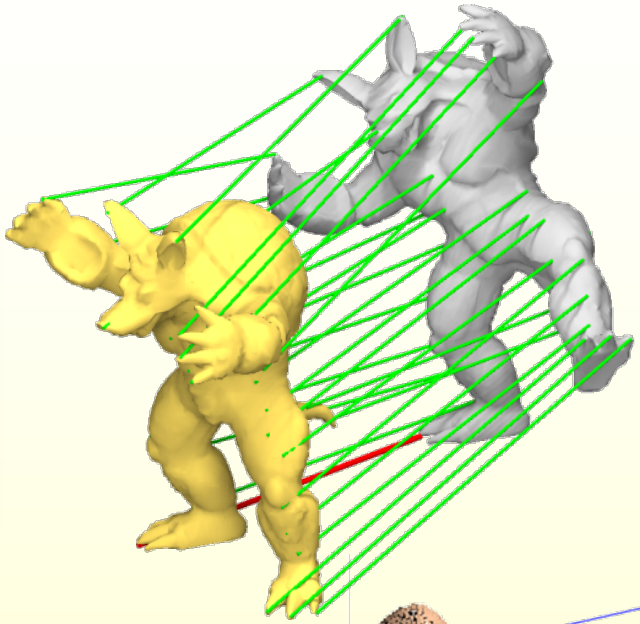

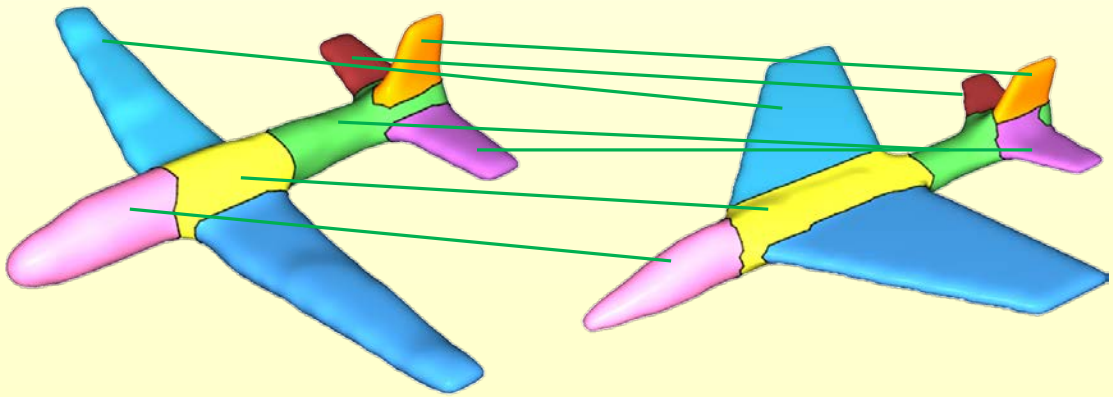

Transport Operators

Societies, or Social Networks of Data Sets

Our understanding of data can greatly benefit from extracting these relations and building relational networks.

We can exploit the relational network to

- transport information around the network
- assess the validity of operations or interpretations of data (by checking consistency against related data)
- assess the quality of the relations themselves (by checking consistency against other relations through cycle closure, etc.)
- extract shared structure among the data

Thus the network becomes the great regularizer in joint data analysis.

Functional Over Data and Functional Maps

Knowledge as Functions over Data

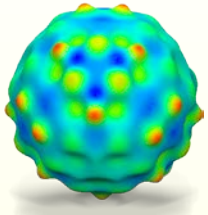

Curvature

Parts

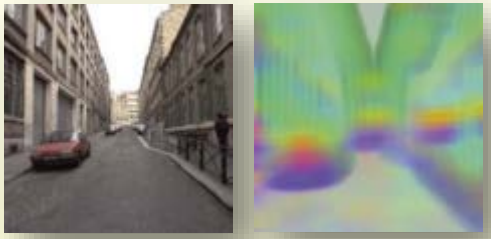

SIFT flow, C. Liu 2011

Knowledge towers over visual data: function spaces

Hierarchical Bases for a Function Space

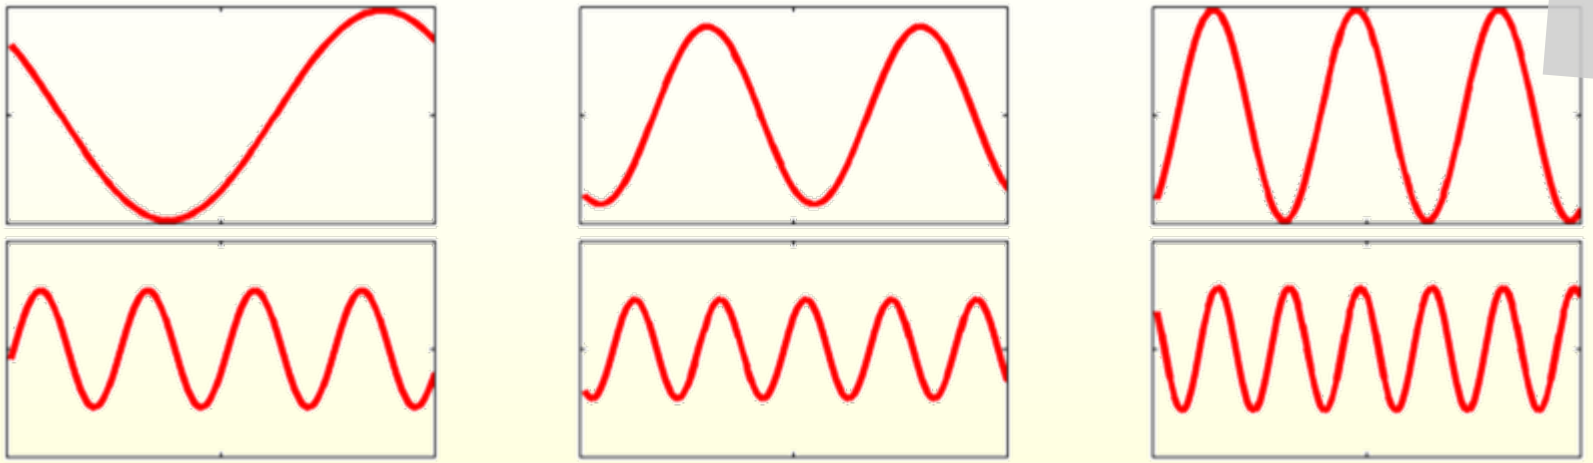

Fourier

Laplace-Beltrami

global support

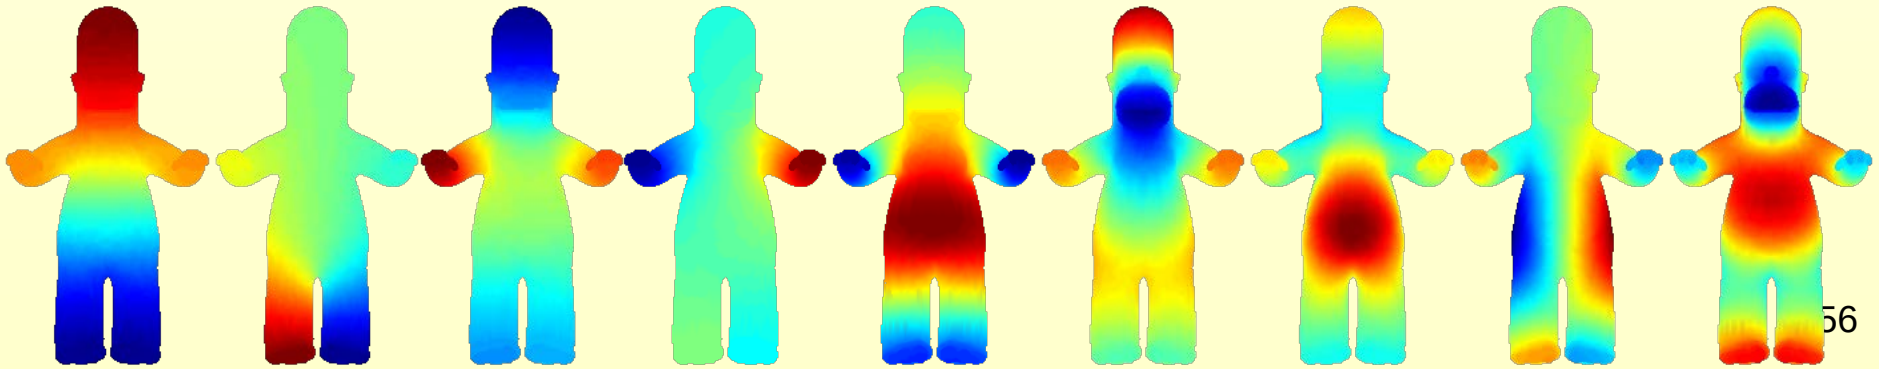

Functional Maps

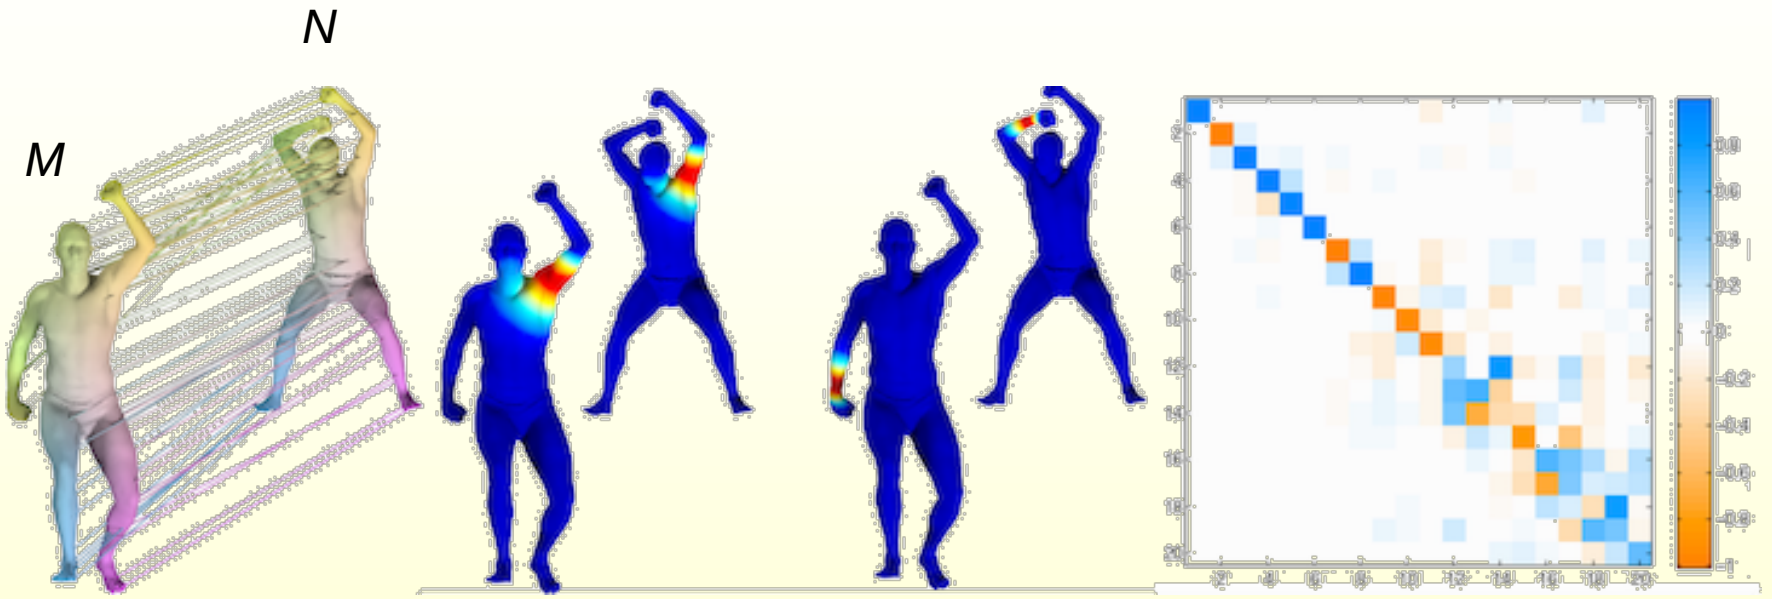

$$T : N \to M$$

$$C_T : \begin{array}{ccc} L^2(M) & \to & L^2(N) \\ f & \mapsto & f \circ T \end{array}$$

A contravariant functor

Starting from a Regular Map φ

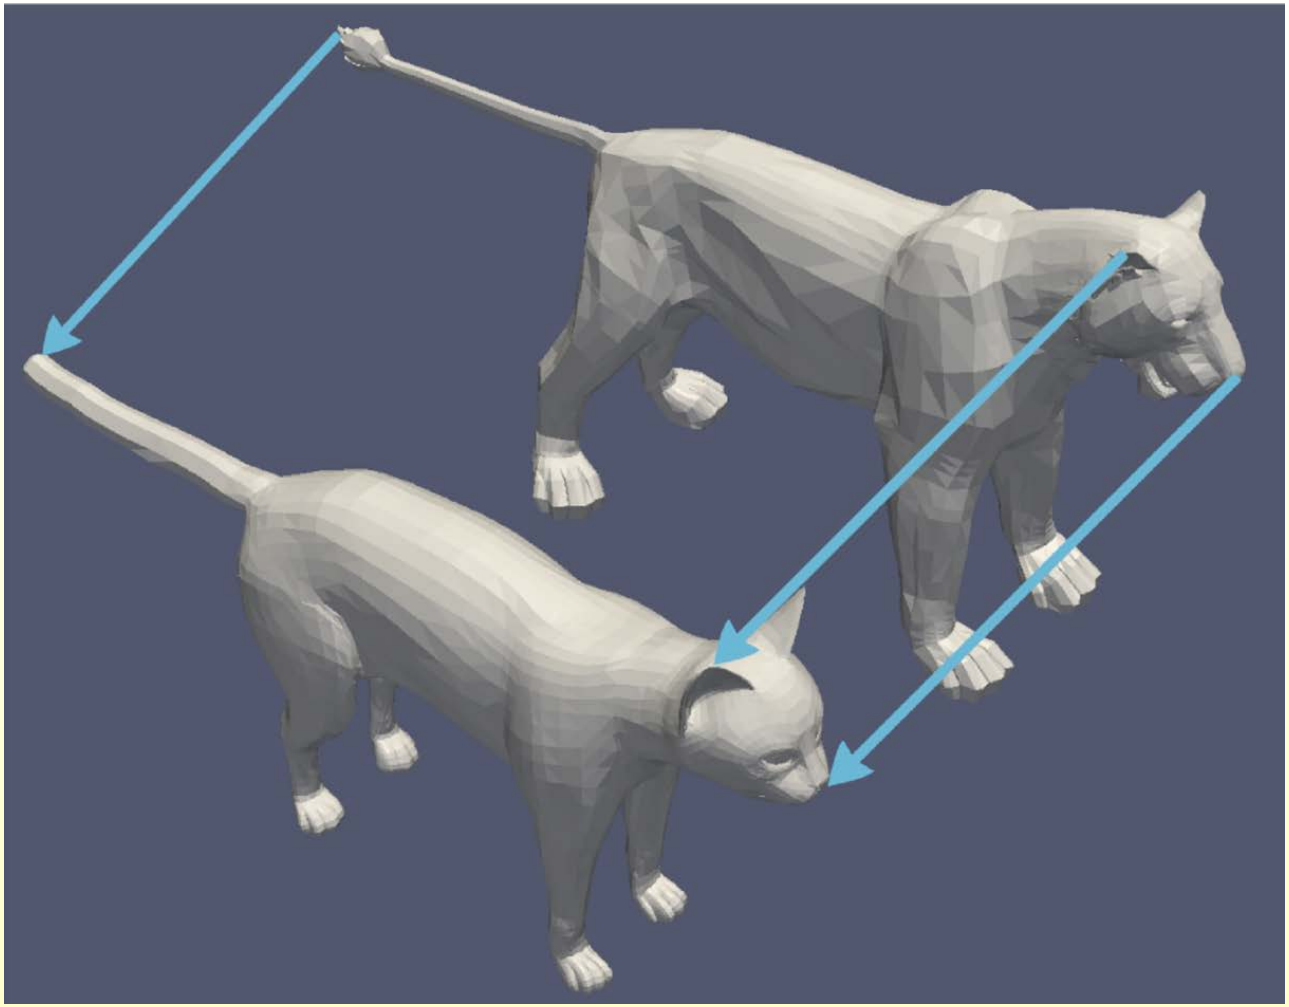

φ: lion → cat

A Contravariant Functor

from cat to lion

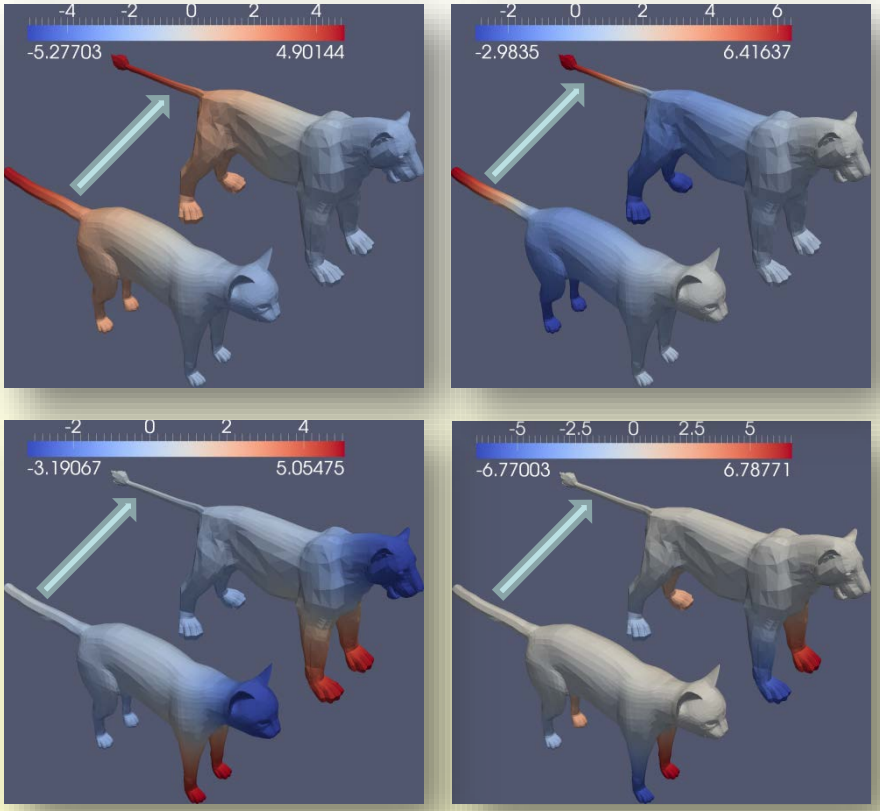

Functions on cat are transferred to lion using T_φ

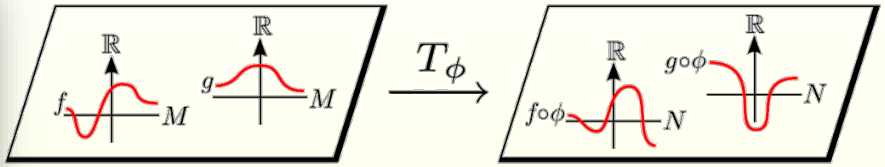

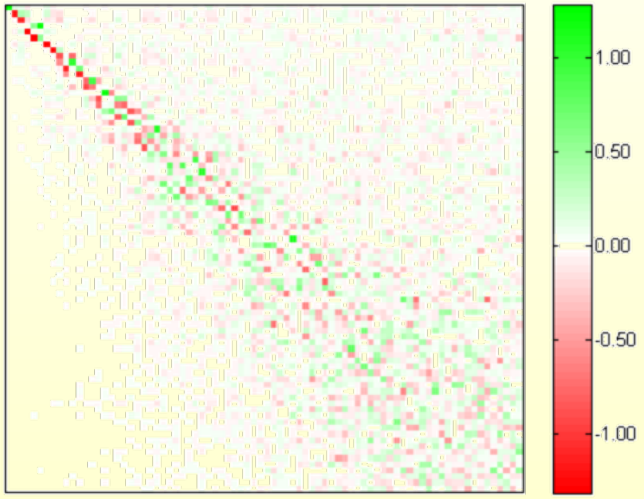

T_φ is a linear operator (matrix)

$$T_{\acute{A}} : L^2(\text{cat}) \; ! \quad L^2(\text{lion})$$

Compact Encoding of Maps as Matrices

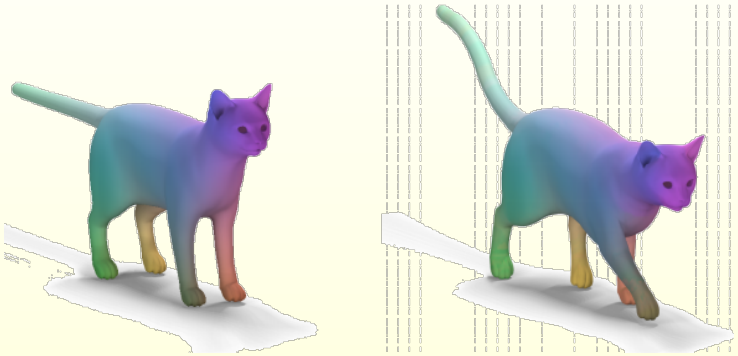

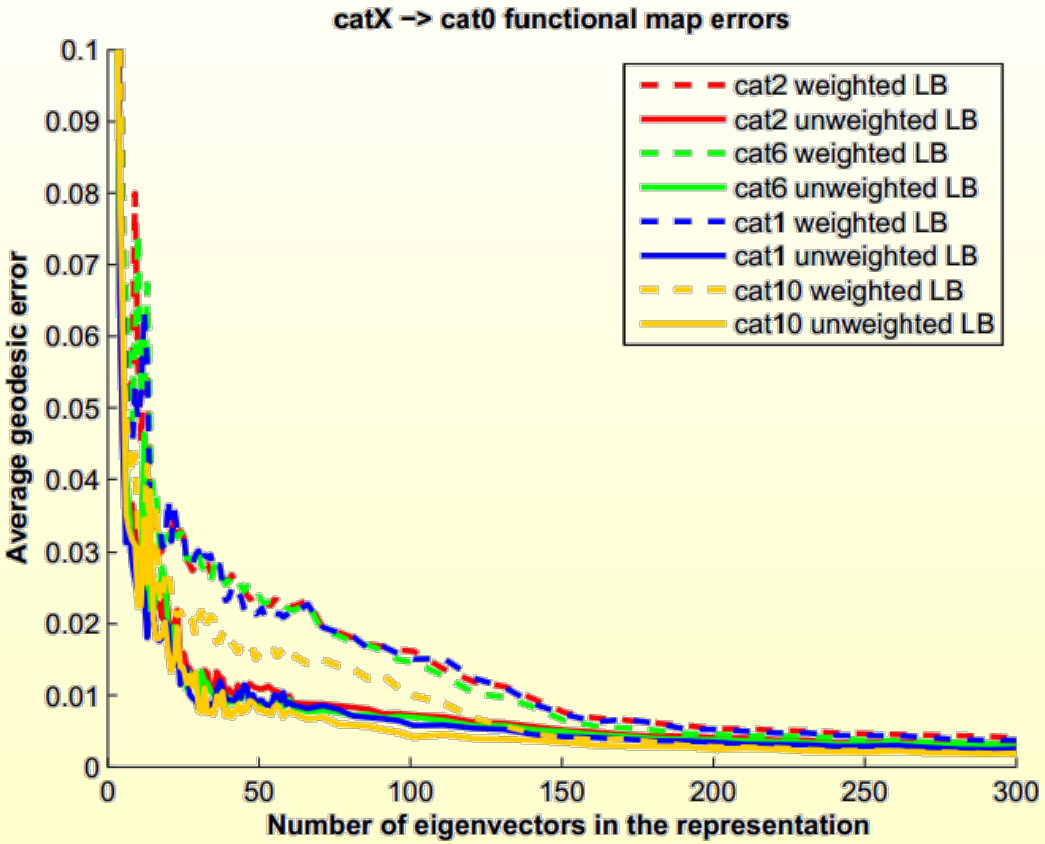

Changes to Measure and Metric under a Map

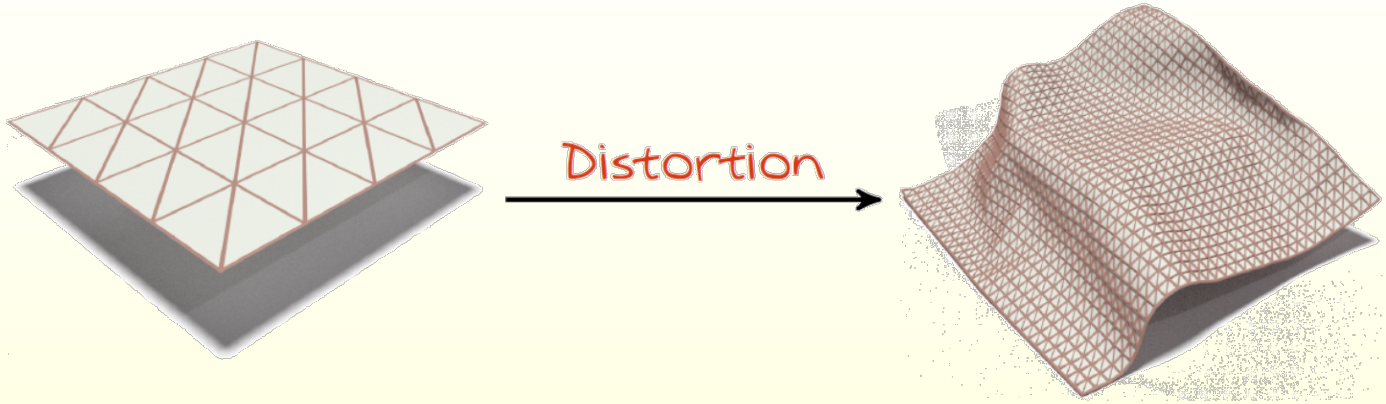

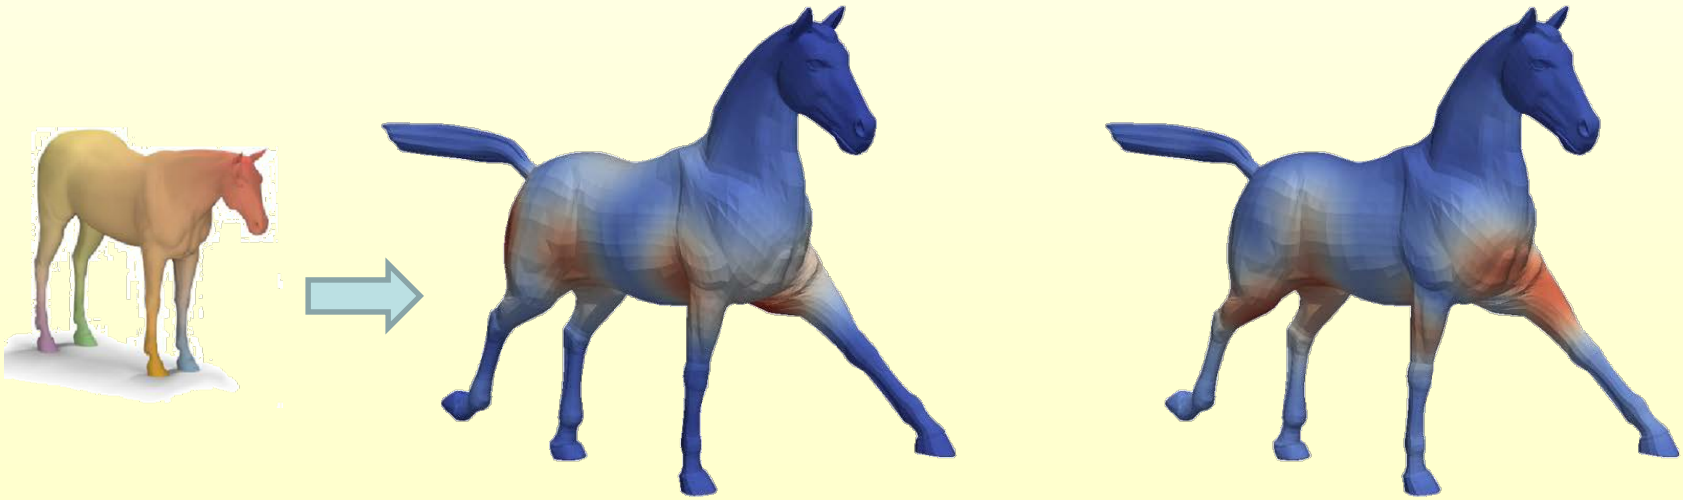

Area distortion

Conformal distortion

Shape Differences are a Change Recipe

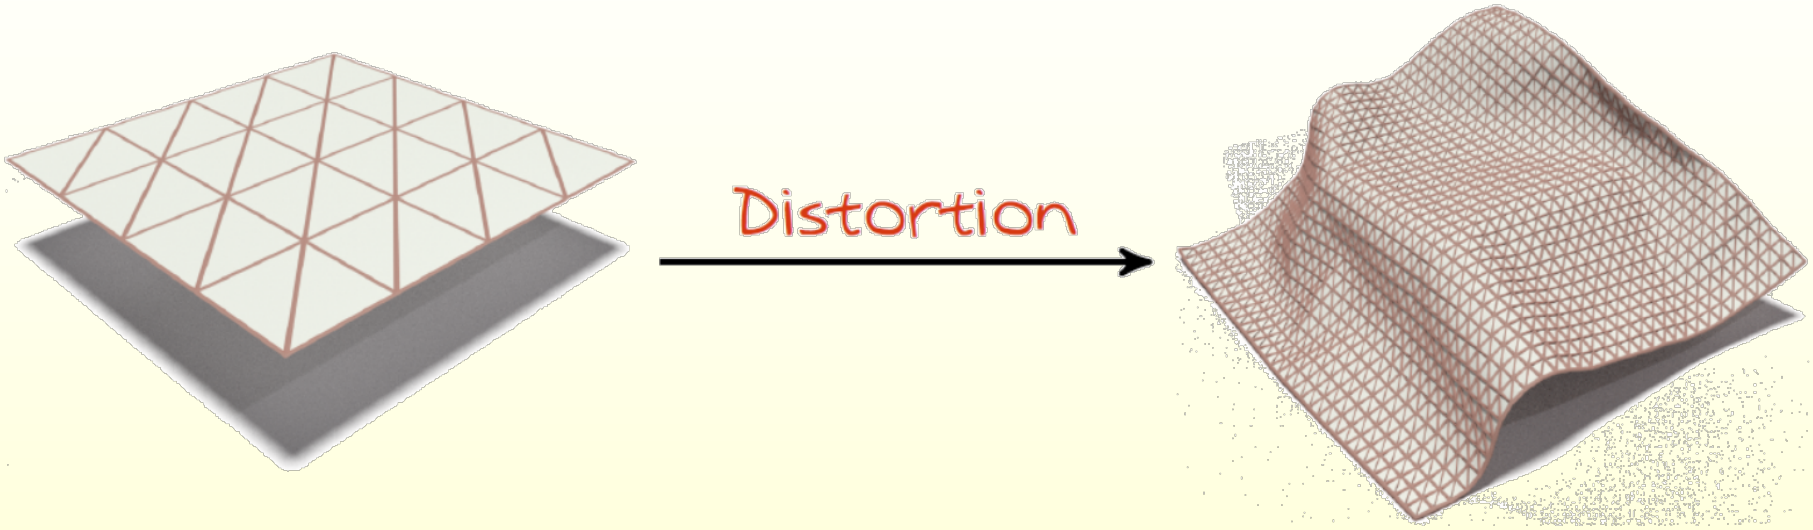

A recipe encoded as a matrix:

Area distortion
Conformal distortion

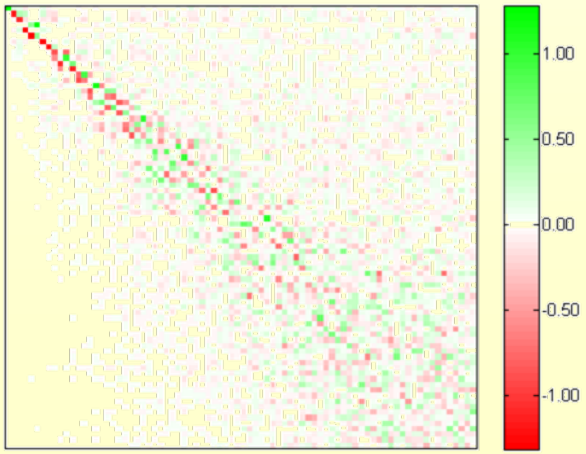

A novel type of latent space representation for 3D data

Under some conditions, lossless

The Space of Shapes

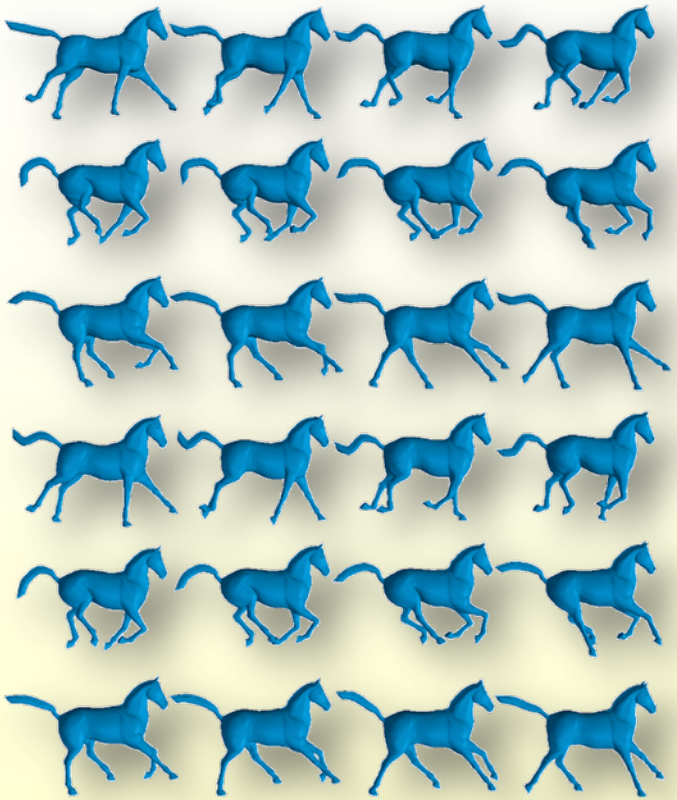

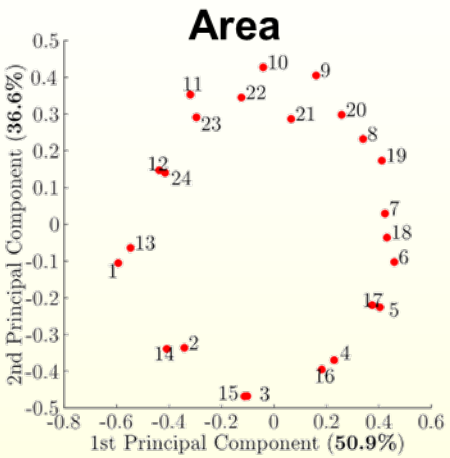

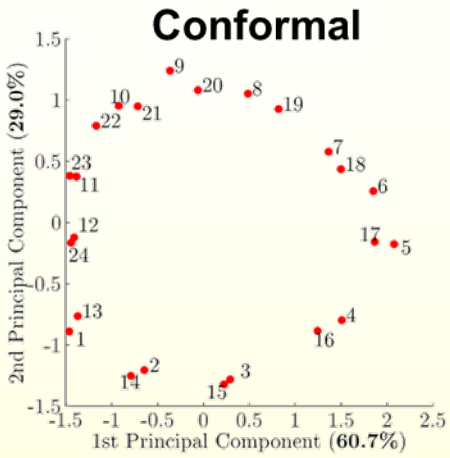

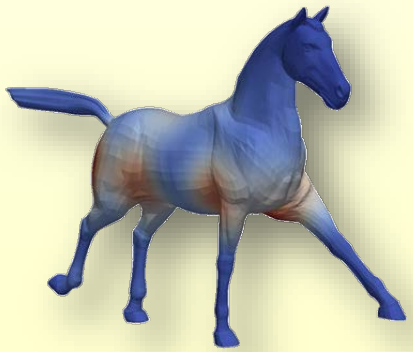

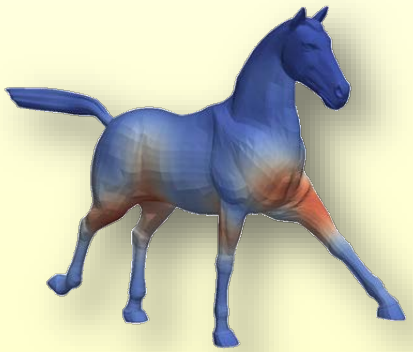

Map Networks

Networks Between Data

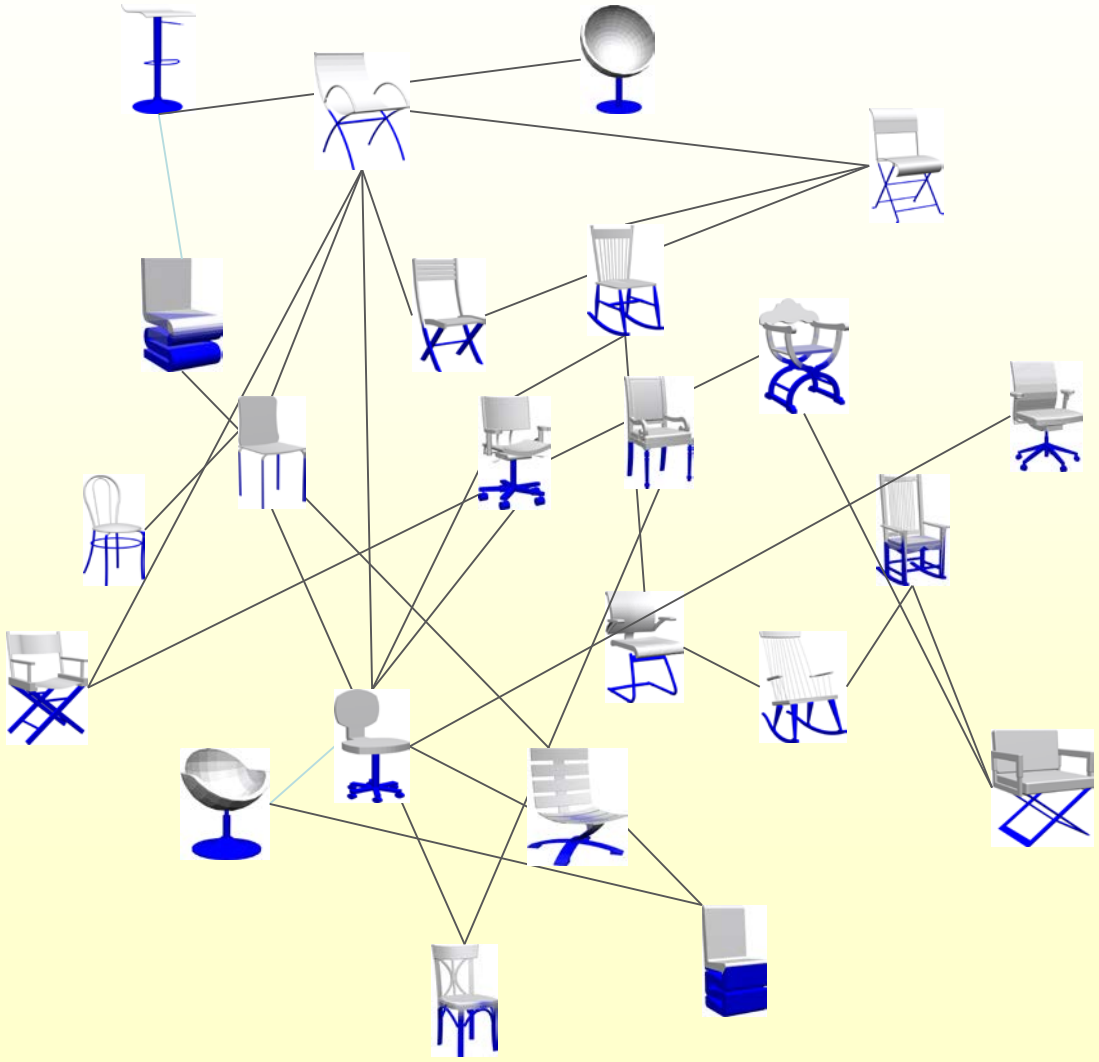

Joint Analysis: Co-Segmentation

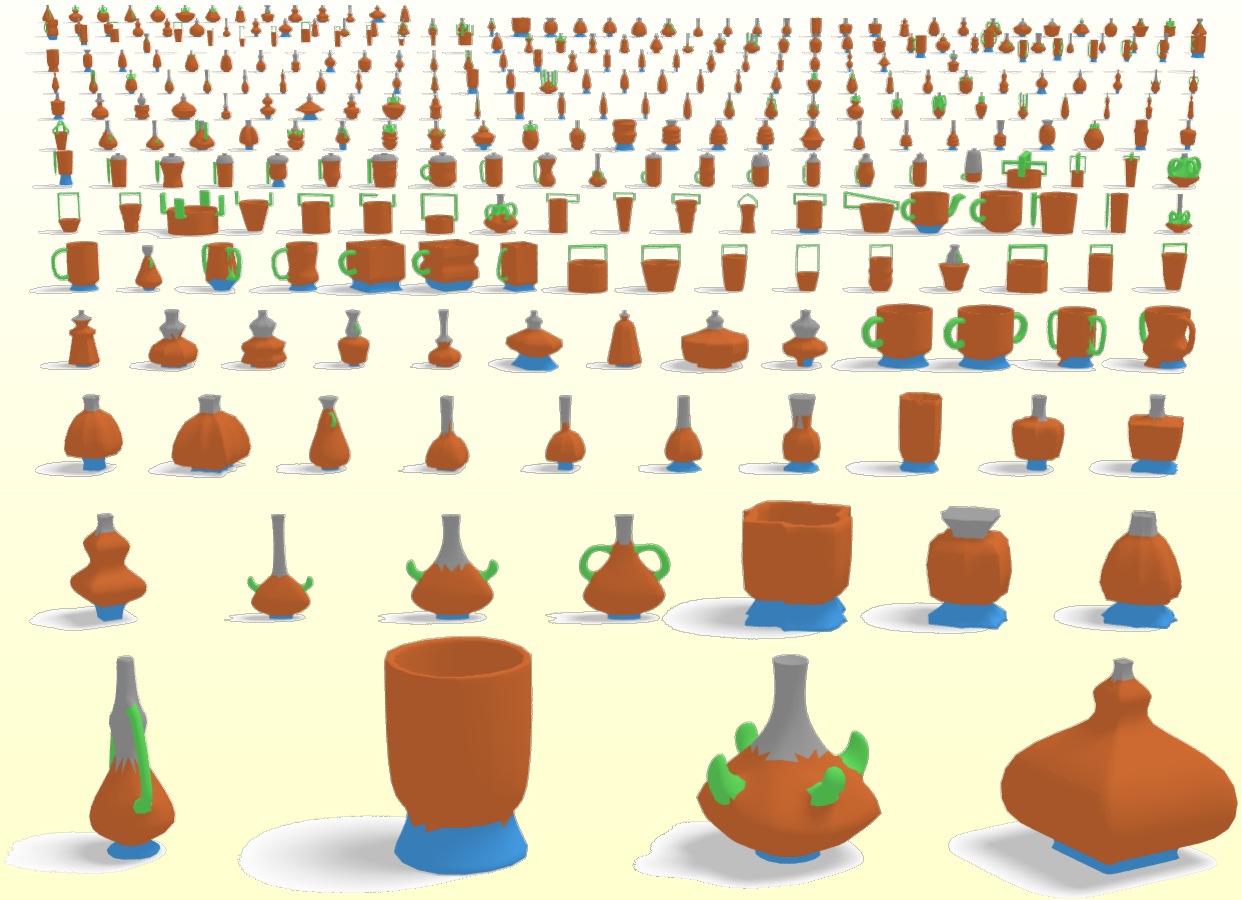

Unsupervised Structure Extraction

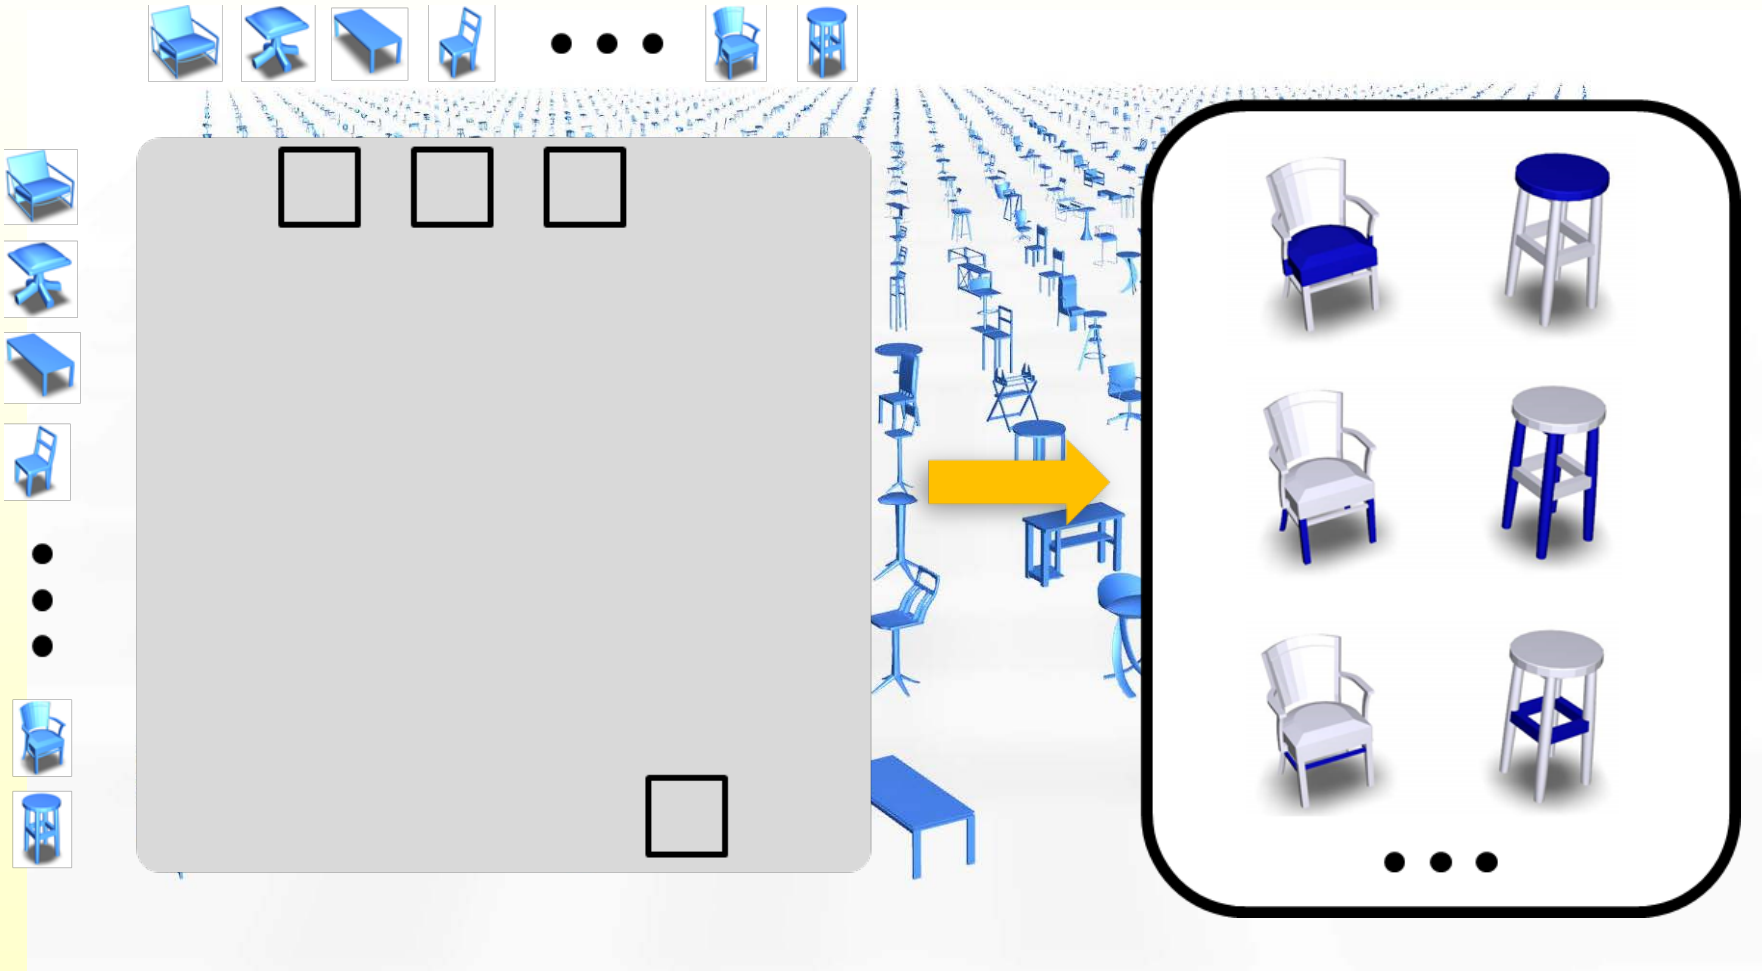

[Q. Huang, F. Wang, L. Guibas, ’14]

MSRC: 5 images per class are shown

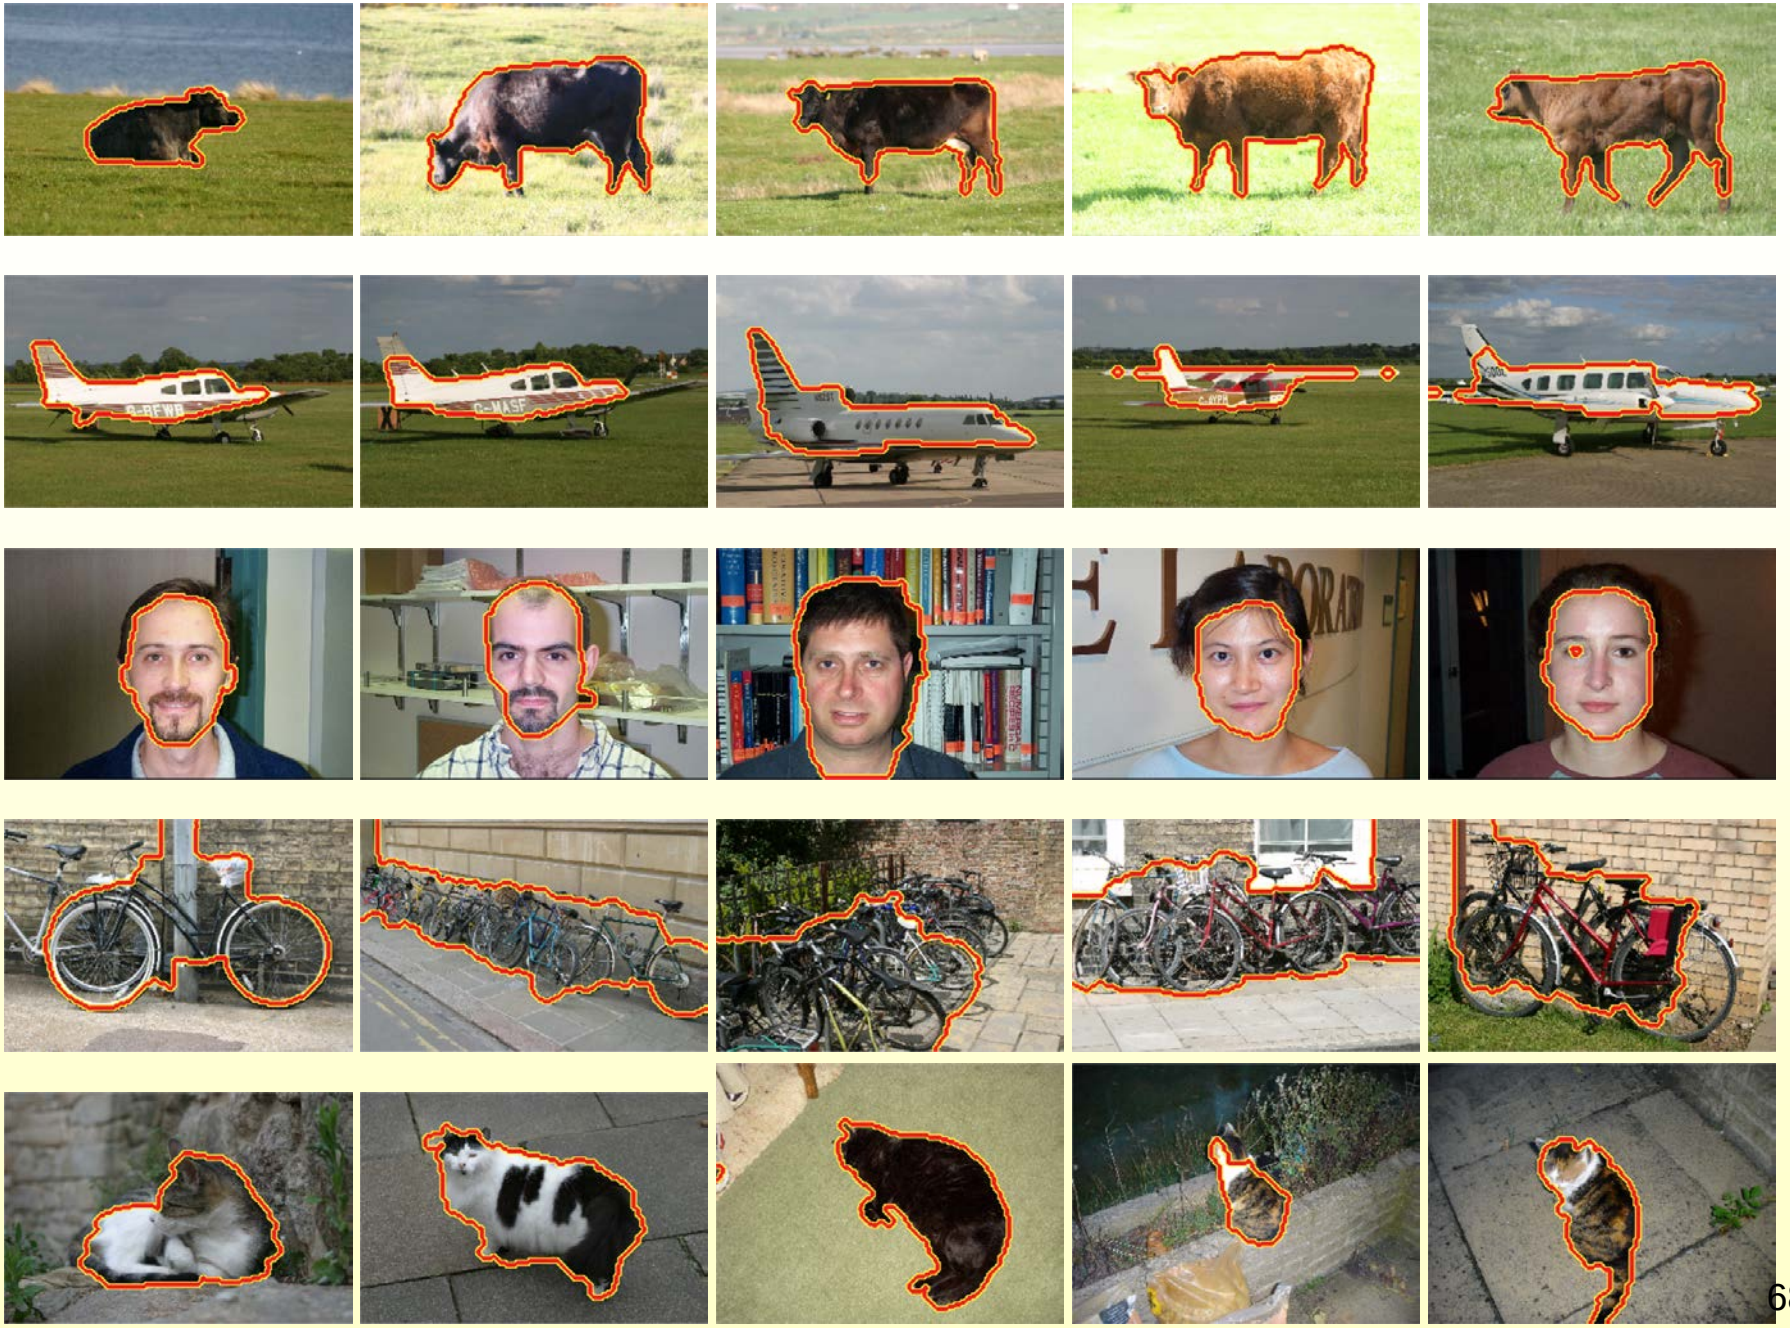

The Network is the Abstraction

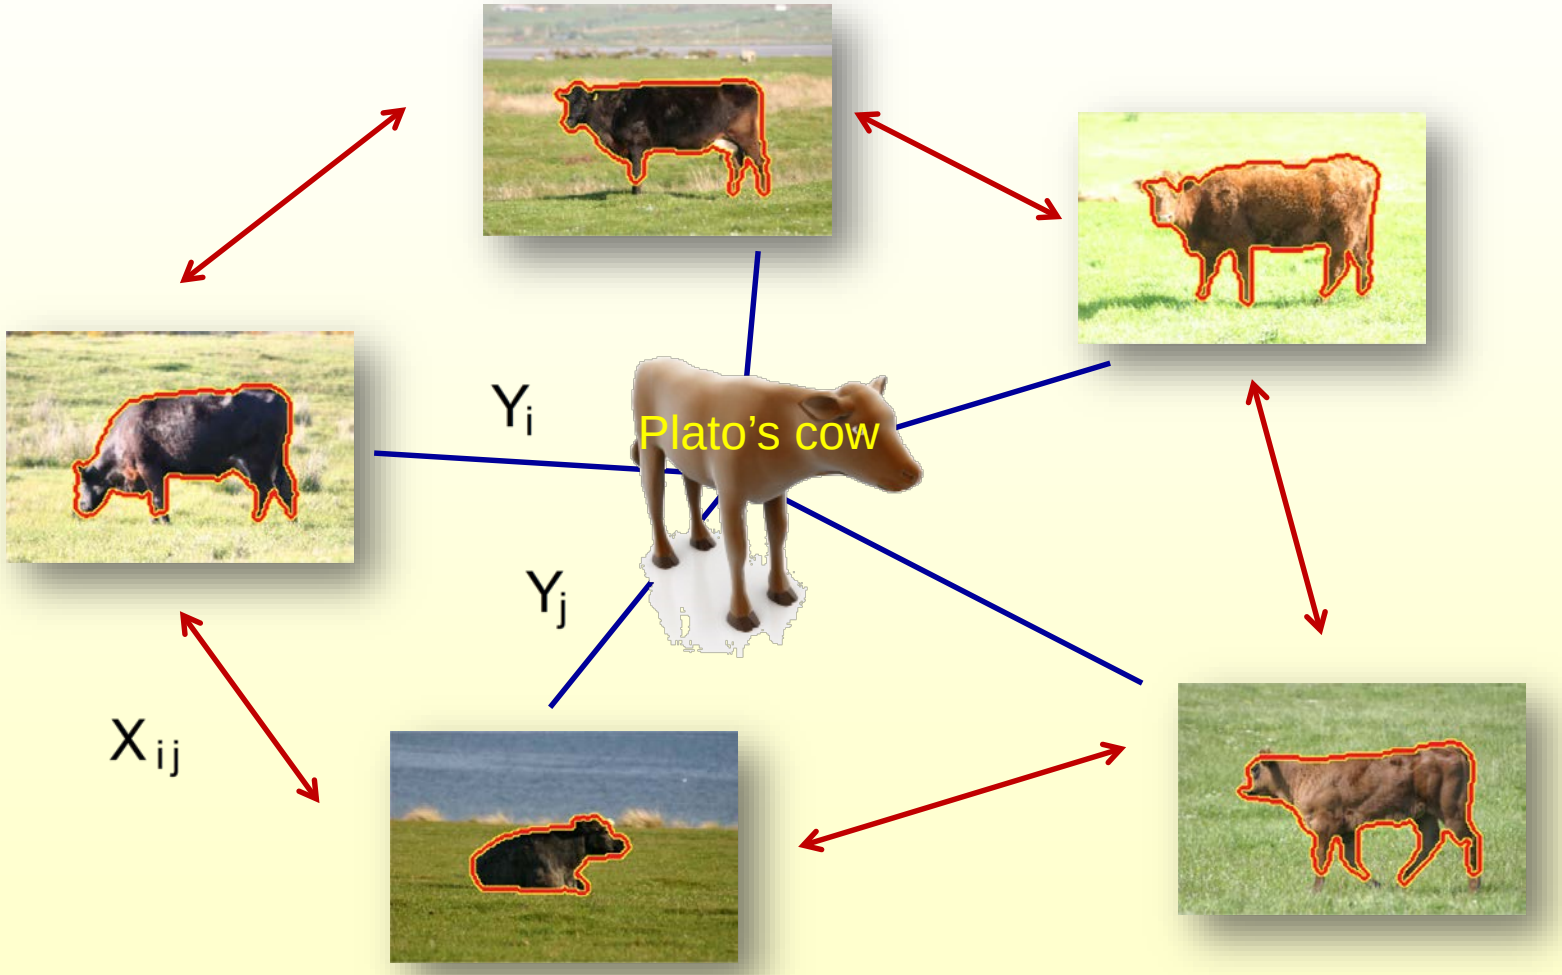

Mosaicking or SLAM at the Level of Functions

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

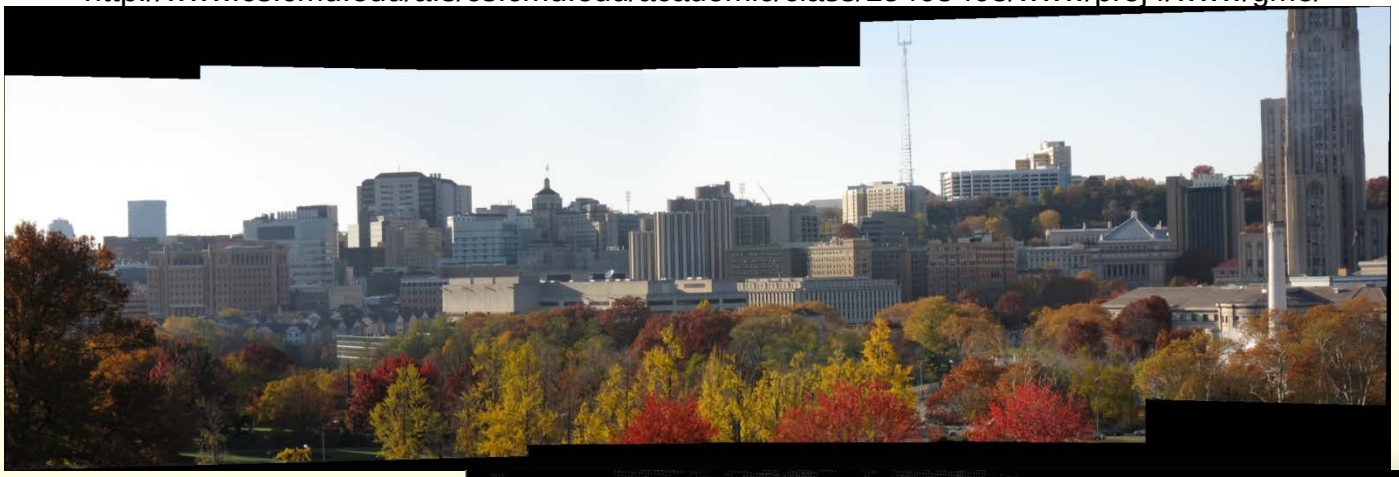

Visual Data Repositories for Storing Semantic Knowledge

Semantic Networks

- Also known as **frame networks**
- They represent semantic relations between concepts
- Often used as a form of knowledge representation
- Stored as a directed or undirected graph consisting of vertices (concepts) and edges (concept relations)

Examples of Semantic Net: WordNet

- a lexical database of English
- words -> synonym sets (synsets)

```
dog, domestic dog, Canis familiaris
    => canine, canid
       => carnivore
          => placental, placental mammal, eutherian, eutherian mammal
             => mammal
                => vertebrate, craniate
                   => chordate
                      => animal, animate being, beast, brute, creature, fauna
                         => ...
```

G. A. Miller, R. Beckwith, C. D. Fellbaum, D. Gross, K. Miller. 1990. WordNet: An online lexical database. Int. J. Lexicograph.

Examples of Semantic Net

Semantic Net in Lisp

```
(defun *database* ()
'((canary  (is-a bird)
           (color yellow)
           (size small))
  (penguin (is-a bird)
           (movement swim))
  (bird    (is-a vertebrate)
           (has-part wings)
           (reproduction egg-laying))))
```

Graph representation

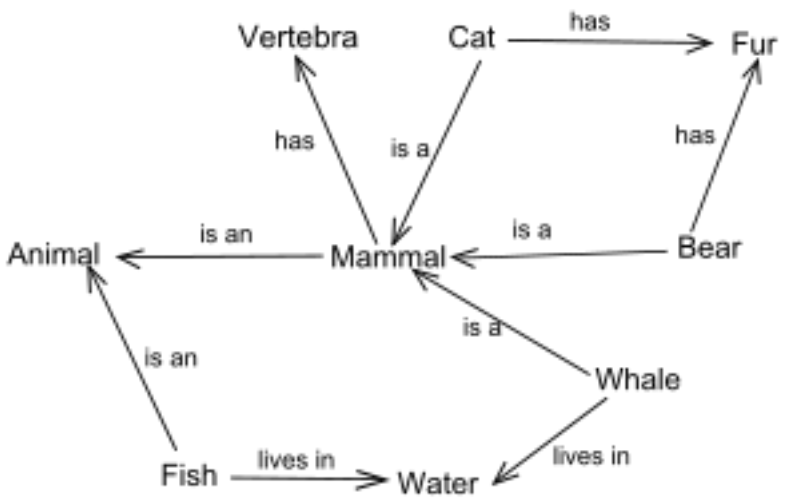

Taxonomy: is-a relationship

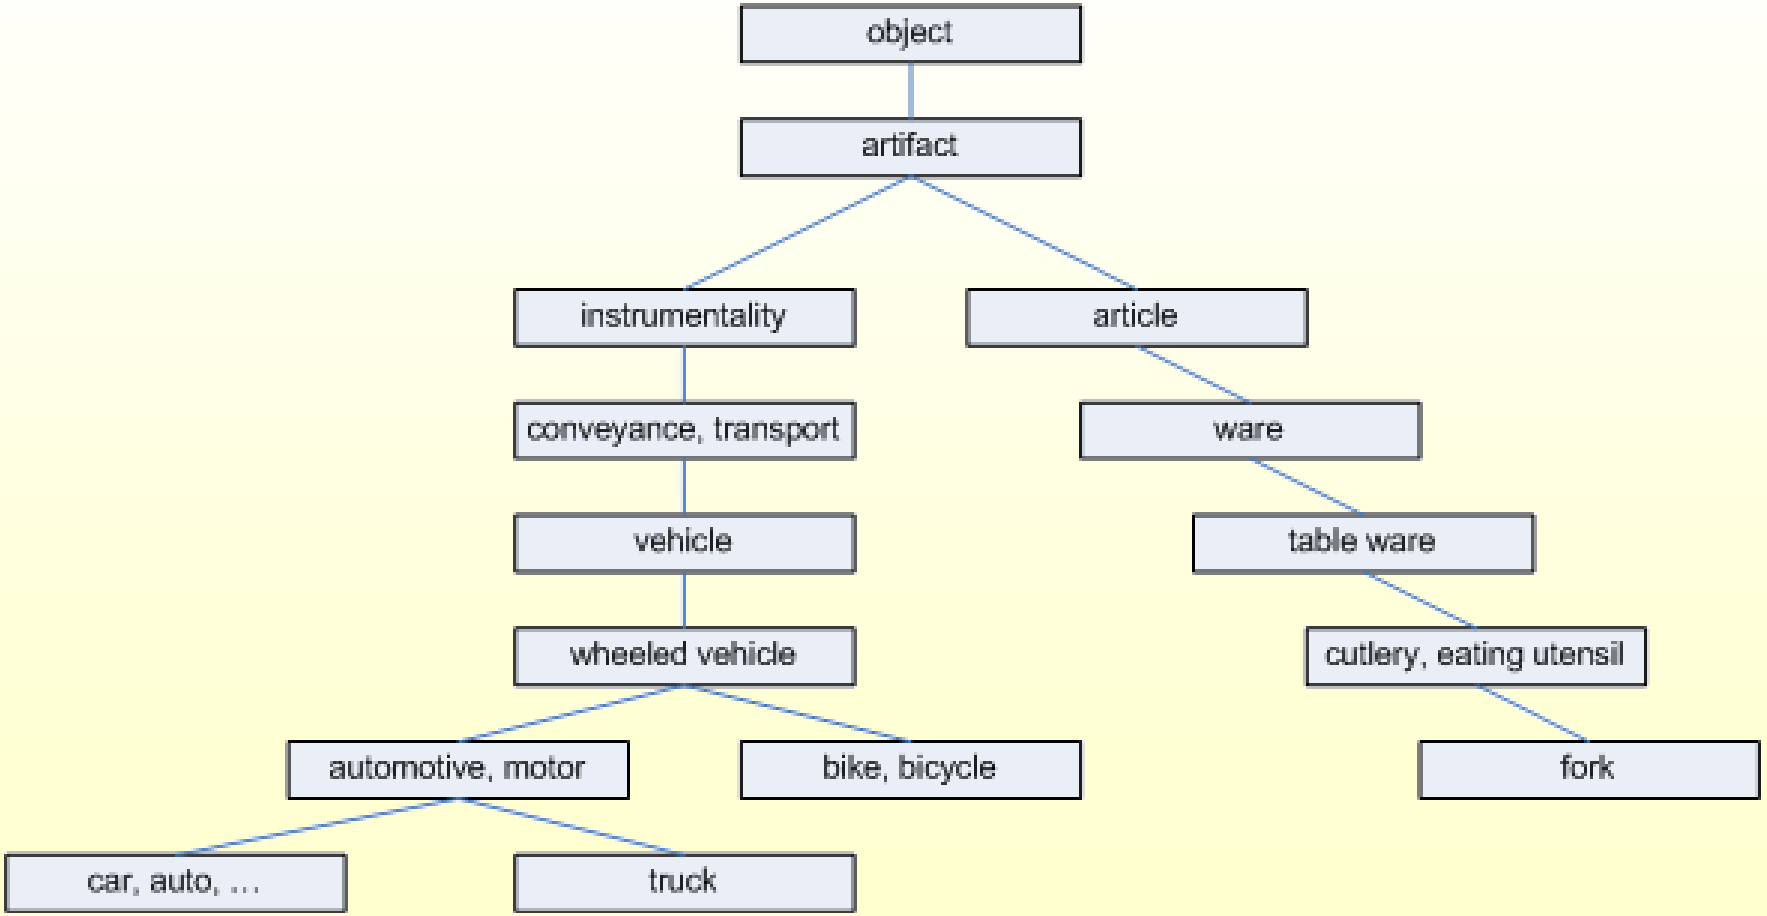

IMAGENET is a knowledge ontology

- Taxonomy (with WordNet backbone)

- S: (n) Eskimo dog, **husky** (breed of heavy-coated Arctic sled dog)
  - *direct hypernym* / ***inherited hypernym*** / *sister term*
    - S: (n) working dog (any of several breeds of usually large powerful dogs bred to work as draft animals and guard and guide dogs)
      - S: (n) dog, domestic dog, Canis familiaris (a member of the genus Canis (probably descended from the common wolf) that has been domesticated by man since prehistoric times; occurs in many breeds) *"the dog barked all night"*
        - S: (n) canine, canid (any of various fissiped mammals with nonretractile claws and typically long muzzles)
          - S: (n) carnivore (a terrestrial or aquatic flesh-eating mammal) *"terrestrial carnivores have four or five clawed digits on each limb"*
            - S: (n) placental, placental mammal, eutherian, eutherian mammal (mammals having a placenta; all mammals except monotremes and marsupials)
              - S: (n) mammal, mammalian (any warm-blooded vertebrate having the skin more or less covered with hair; young are born alive except for the small subclass of monotremes and nourished with milk)
                - S: (n) vertebrate, craniate (animals having a bony or cartilaginous skeleton with a segmented spinal column and a large brain enclosed in a skull or cranium)
                  - S: (n) chordate (any animal of the phylum Chordata having a notochord or spinal column)
                    - S: (n) animal, animate being, beast, brute, creature, fauna (a living organism characterized by voluntary movement)
                      - S: (n) organism, being (a living thing that has (or can develop) the ability to act or function independently)
                        - S: (n) living thing, animate thing (a living (or once living) entity)
                          - S: (n) whole, unit (an assemblage of parts that is regarded as a single entity) *"how big is that part compared to the whole?"; "the team is a unit"*
                            - S: (n) object, physical object (a tangible and visible entity; an entity that can cast a shadow) *"it was full of rackets, balls and other objects"*
                              - S: (n) physical entity (an entity that has physical existence)
                                - S: (n) entity (that which is perceived or known or inferred to have its own distinct existence (living or nonliving))

Slide Credit: Li Fei-Fei

ShapeNet (>3M Models)

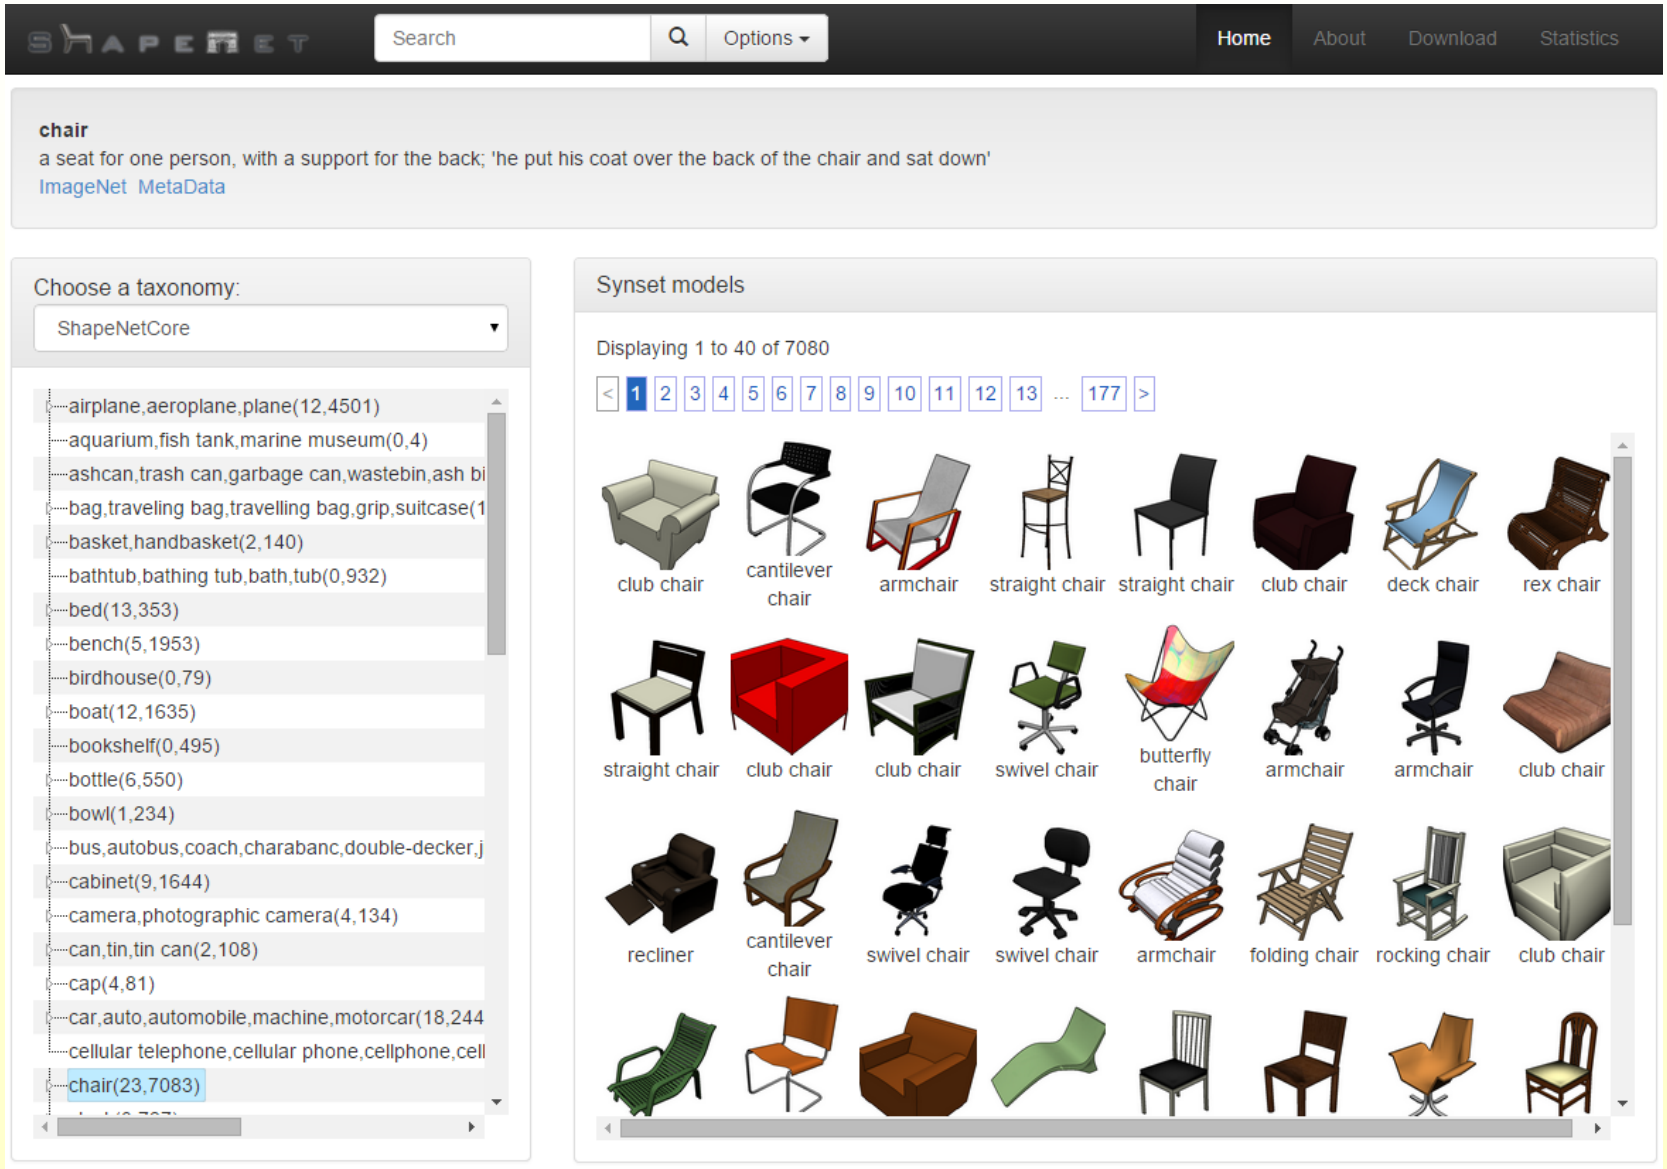

Object Knowledge

Parts, symmetries, keywords, physical properties, materials, affordances, …

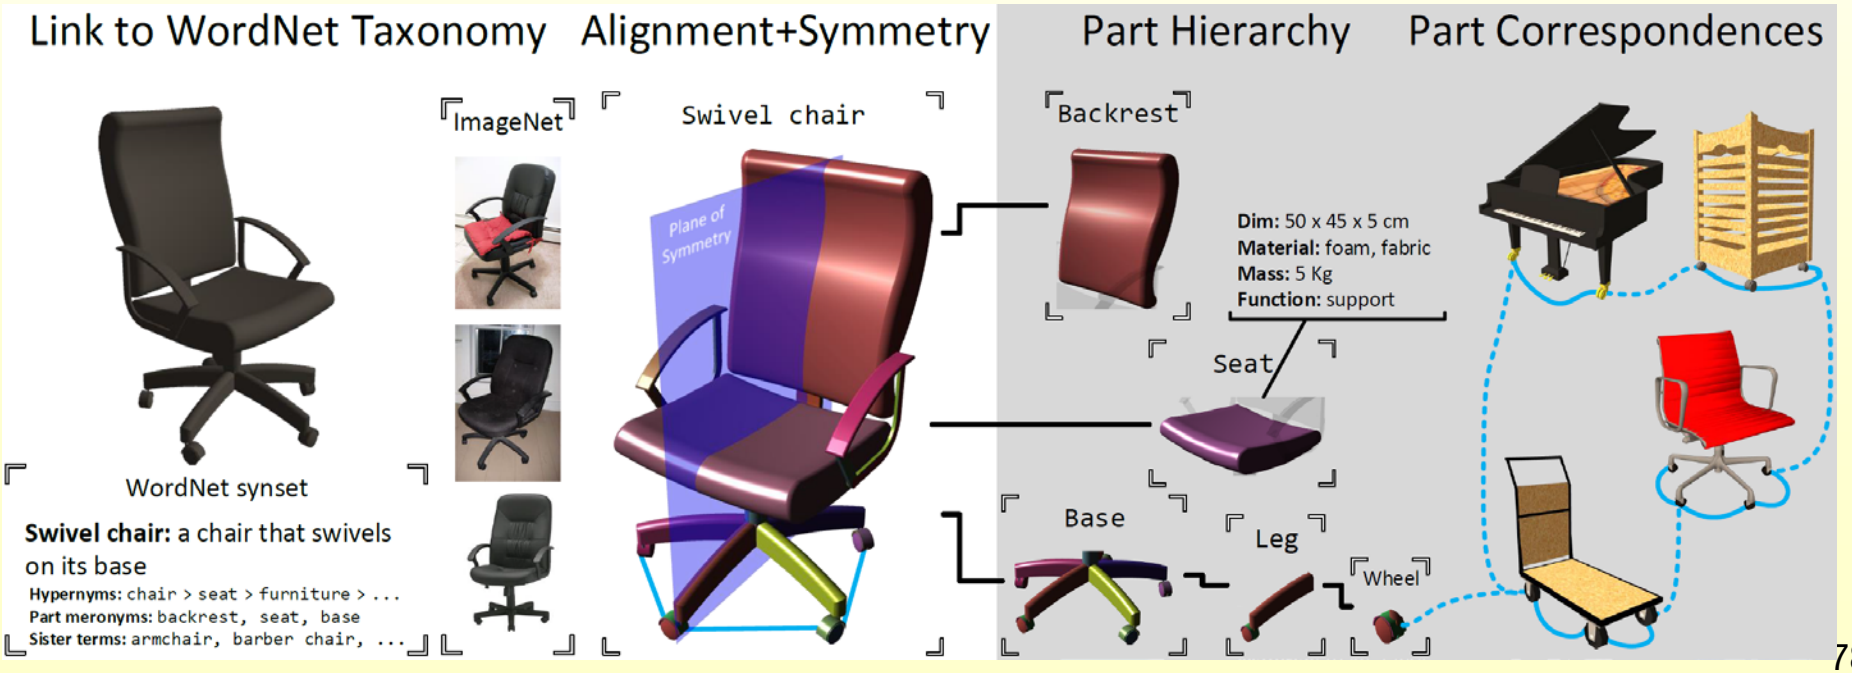

How to Generate Semantic Annotations

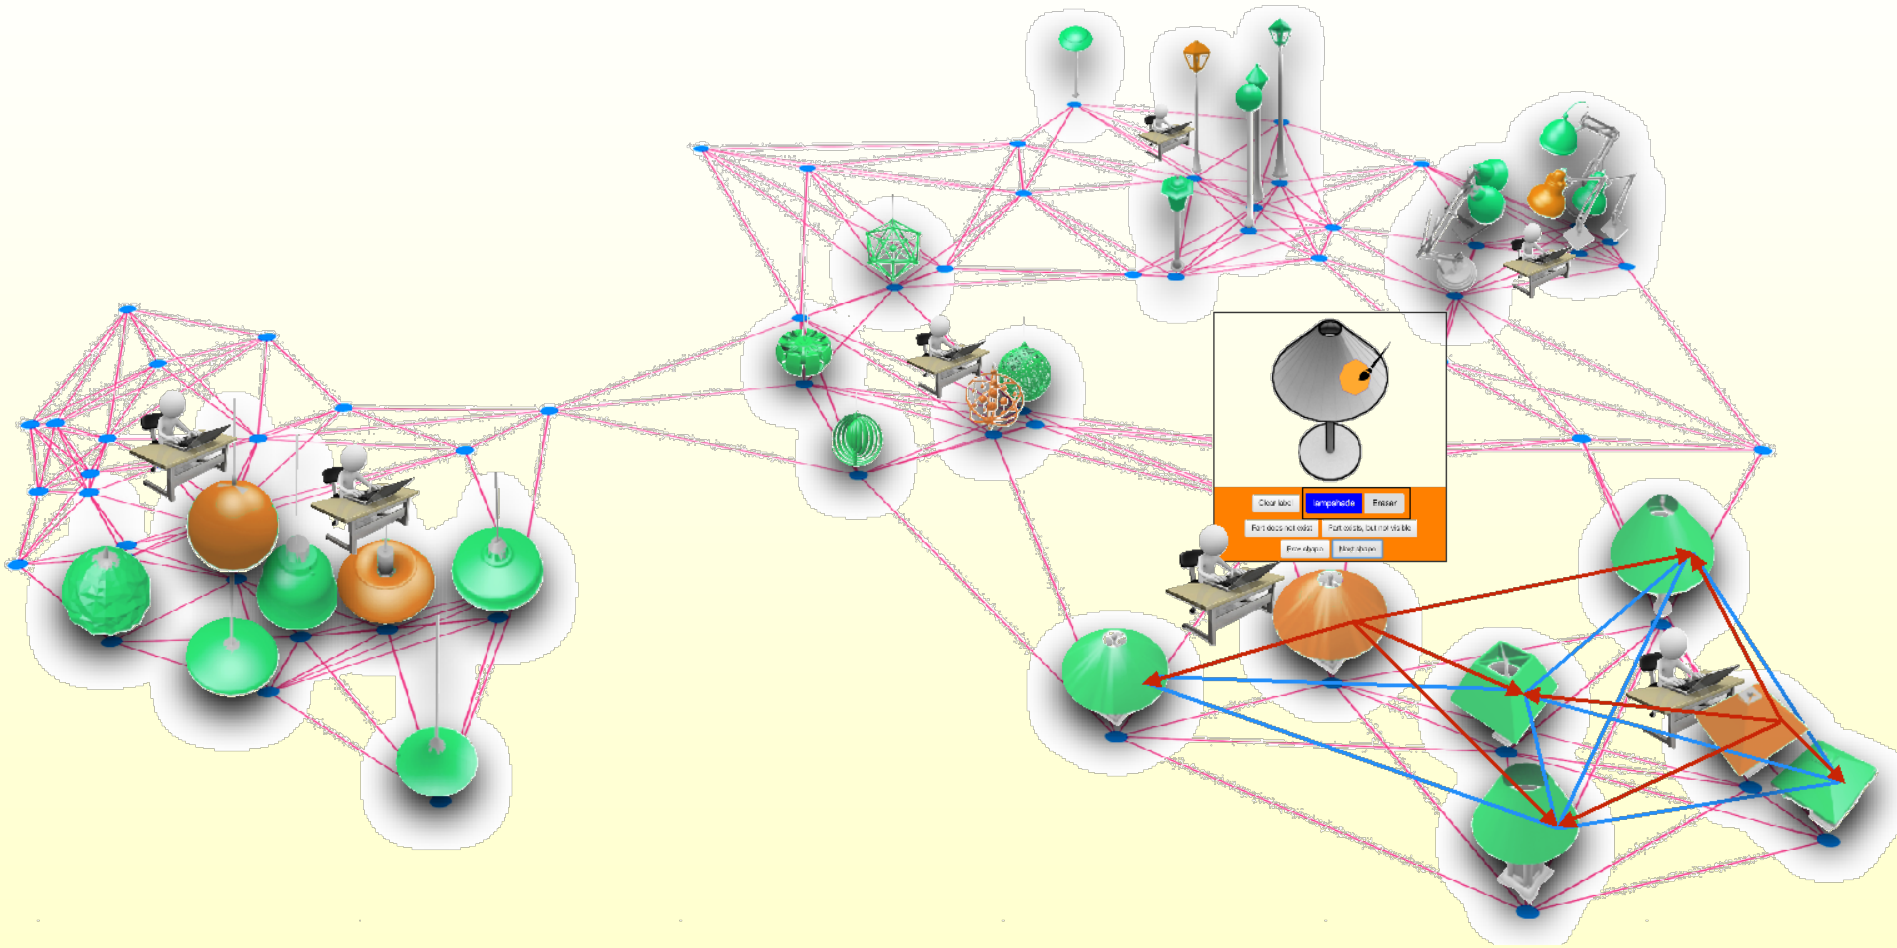

Networks for Performing Knowledge Transfer and Inference

- Bring knowledge to sensor data – and make it available to humans and to computers (or robots)
- Making it possible to "See the Unseen"

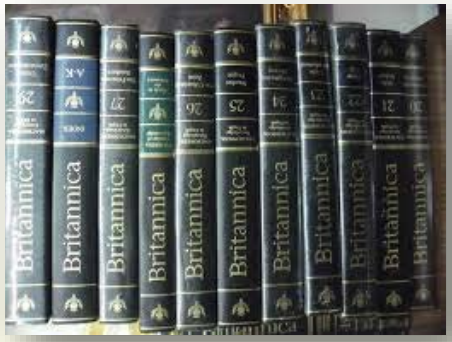

a digital encyclopedia

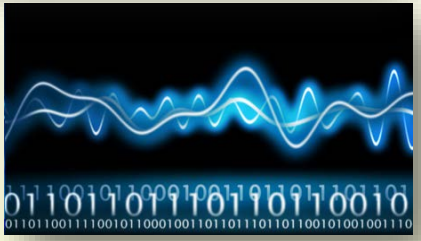

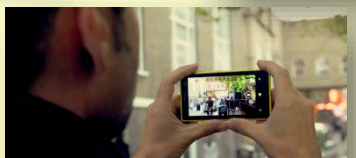

Course Outline

Course Outline, I

April 2	April 4
Introduction; Geometric and topological perspective on data analysis; Data representations; Learning on point clouds and graphs; Joint data analysis. Reading: Lecture 1 Slides	Visual data sets: ImageNet and ShapeNet; annotation and annotation transport. Reading: Lecture 2 Slides
April 9	**April 11**
Linear algebraic techniques: principal components analysis (PCA), Kernel PCA. Reading: Lecture 3 Slides	Linear algebraic techniques: canonical correlation analysis (CCA). Reading: Homework 1 out. Lecture 4 Slides
April 16	**April 18**
Graph methods; spectral approaches, graph Laplacians, Laplacian embeddings, spectral clustering. Reading: Lecture 5 Slides	Multidimensional scaling. Non-linear dimensionality reduction: locally linear embeddings, Laplacian eignemaps, Isomap. Reading: Lecture 6 Slides
April 23	**April 25**
Computational topology: topology review, complexes, homology groups. Reading: I Lecture 7 Slides	Persistent homology, barcodes and persistence diagrams. Reading: Homework 1 due. Homework 2 out. Lecture 8 Slides
April 30	**May 2**
Topological inference; the Mapper algorithm. Applications. Reading: Lecture 9 Slides	Representations of 3D Geometry: Voxel-Grids, Point Clouds, Meshes and Other Boundary Models, Solid Models Reading: Lecture 10 Slides

Course Outline, II

Topological inference; the Mapper algorithm. Applications. Reading: Lecture 9 Slides	Representations of 3D Geometry: Voxel-Grids, Point Clouds, Meshes and Other Boundary Models, Solid Models Reading: Lecture 10 Slides
May 7	**May 9**
Geometry processing; Laplace-Beltrami and other operators on meshes; Reading: Lecture 11 Slides	Global and local shape descriptors; intrinsic descriptors, heat and wave kernel signatures. Reading: Homework 2 due. Homework 3 out. Lecture 12 Slides
May 14	**May 16**
Rigid and non-rigid shape alignment; isometric matching, conformal maps, Möbius voting, blended intrinsic maps Reading: Lecture 13 Slides	Volumetric and multi-view CNNs for 3D geometry Reading: Lecture 14 Slides
May 21	**May 23**
Deep nets for pointclouds Reading: Lecture 15 Slides	Deep nets for graphs and meshes Reading: Homework 3 due. Homework 4 out. Lecture 16 Slides
May 28	**May 30**
Memorial day holiday -- no class	Functional spaces and functional maps, variations; map visualization Reading: Lecture 17 Slides
June 4	**June 6**
Shape differences and shape variability. Reading: s Lecture 18 Slides	Networks of shapes and images; cycle consistency; map processing and latent spaces. Reading: Lecture 19 Slides Homework 4 due.

Class Mechanics

The Team

- Instructor:

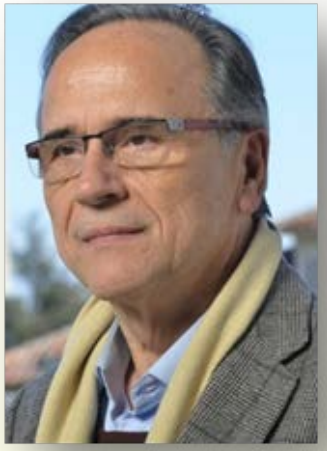

Leonidas Guibas

- CAs:

Panos Achlioptas

Vignesh Ganapathi-Subramanian

- Contributors:

Anastasia Dubrovina

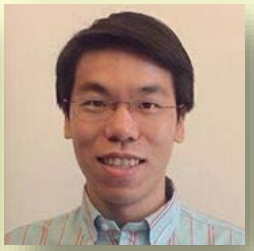

Charles Qi

Minhyuk Sung

Eric Yi

Class Mechanics

- Two weekly lectures
- Office hours
- Class web site http://cs233.stanford.edu

  http://graphics.stanford.edu/courses/cs233-18-spring
- Use Piazza, Gradescope

Course Work

- Four assignments (modest programming in MATLAB for three, one in JavaPlex)
- A short (90 min) final
- Class participation and material contributions

Key Course Goals

- Cover basic tools for geometric and topological data analysis, both supervised and unsupervised
- Present mathematical ways to encode and transfer “knowledge” about data
- Present methods for joint data analysis – benefiting from the “wisdom of the collection”


SHAPE OF WATER

Data Has Shape