CS233, CME251: Geometric and Topological Data Analysis

Leonidas Guibas
Computer Science Department
Stanford University

Lecture 1
29 March 2021

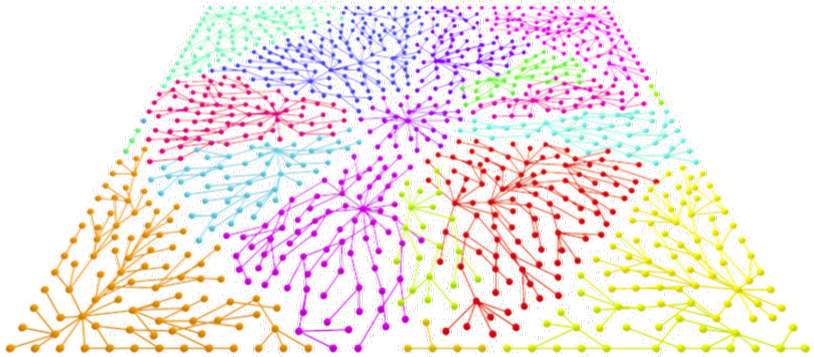

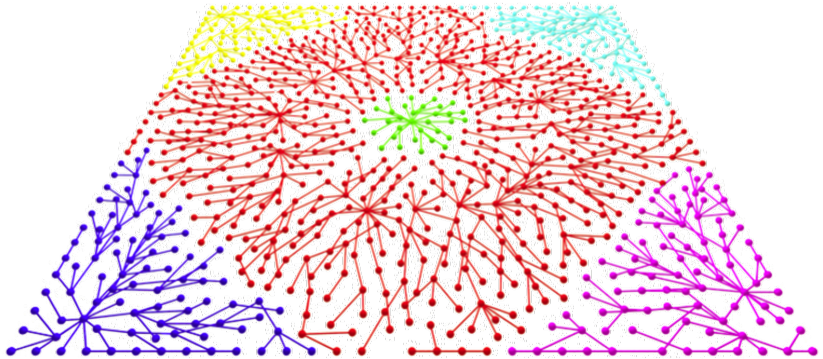

Introduction

Big Data Era

- Data from many kinds of sensors
- Data from simulations
- Data from the activities of individuals on the internet

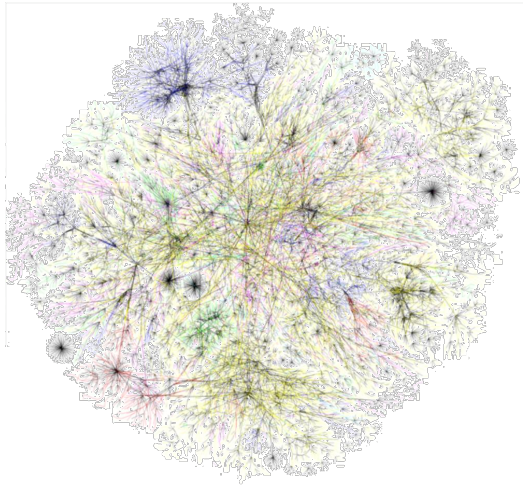

"Data Science"

Data Sets of Geometric Character

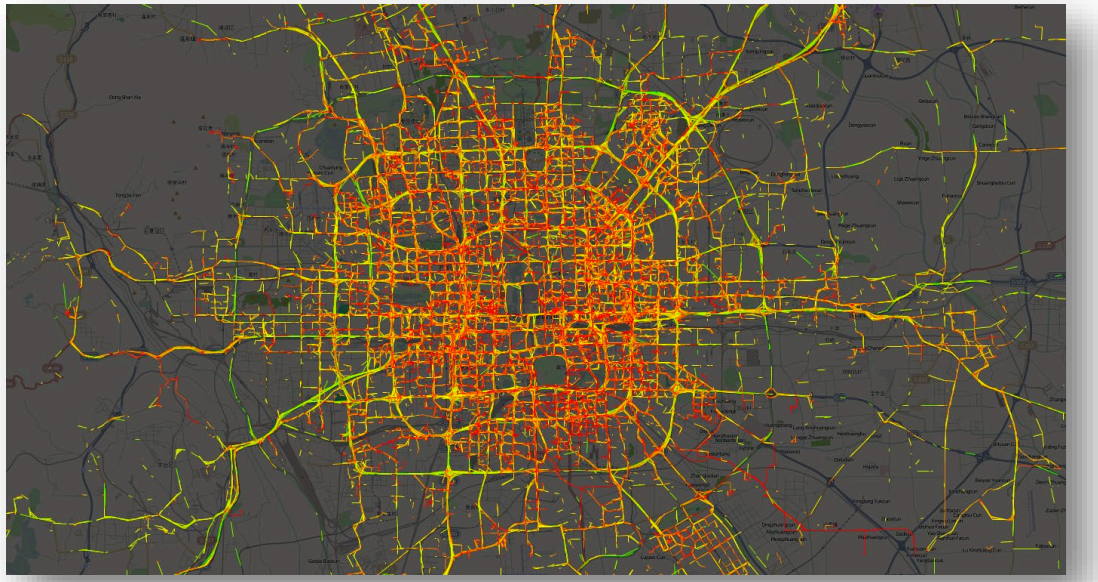

Vehicle GPS traces (Beijing)

3D city scans (Bremen)

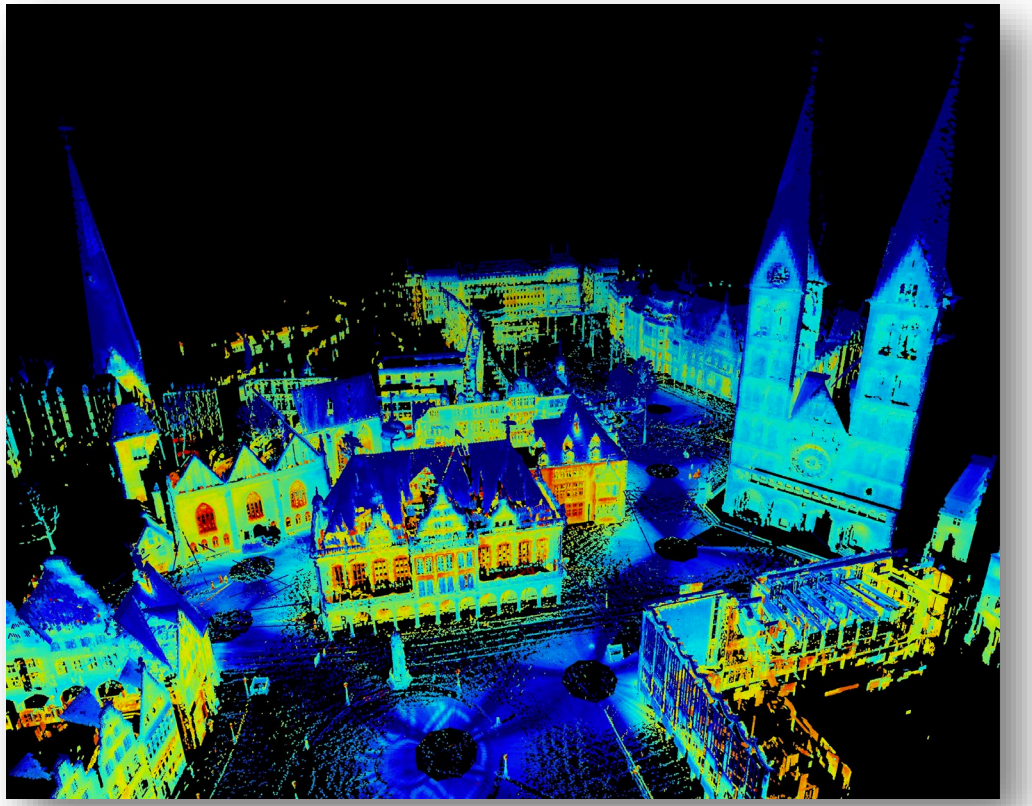

Geometry in Ordinary Images and Videos

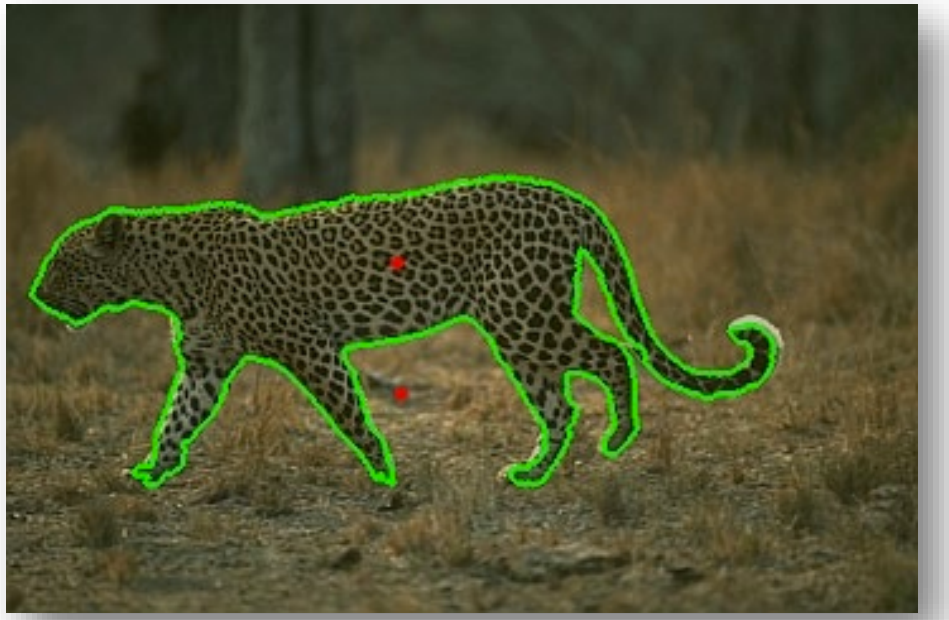

Object recognition

Segmentation involves geometry

Motion involves geometry

Optical flow

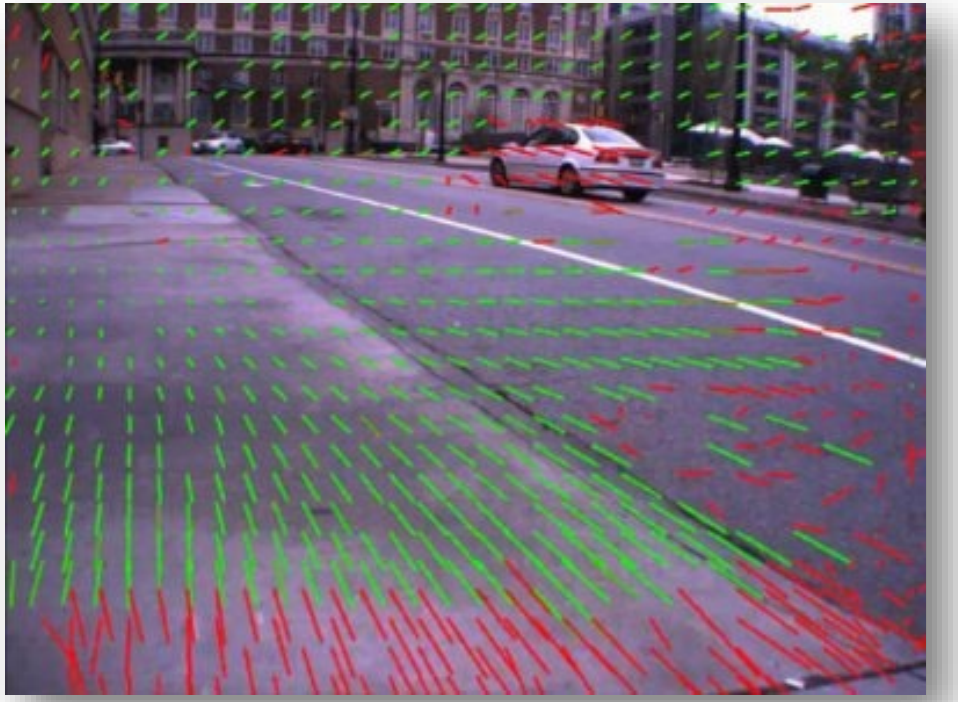

Non-Geometric Data

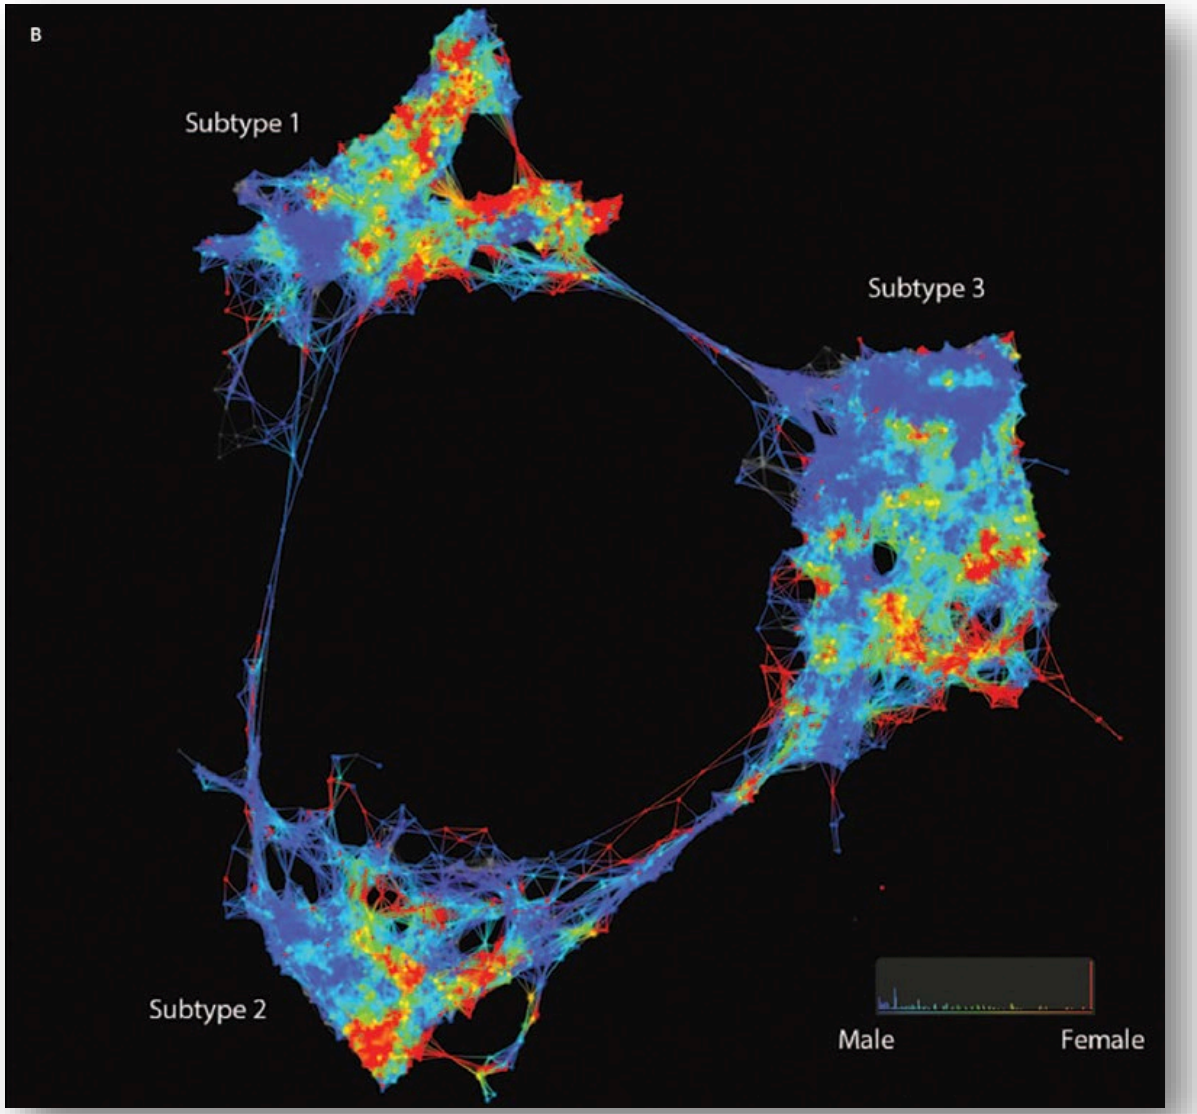

Diabetes II subtypes

Data Has Shape

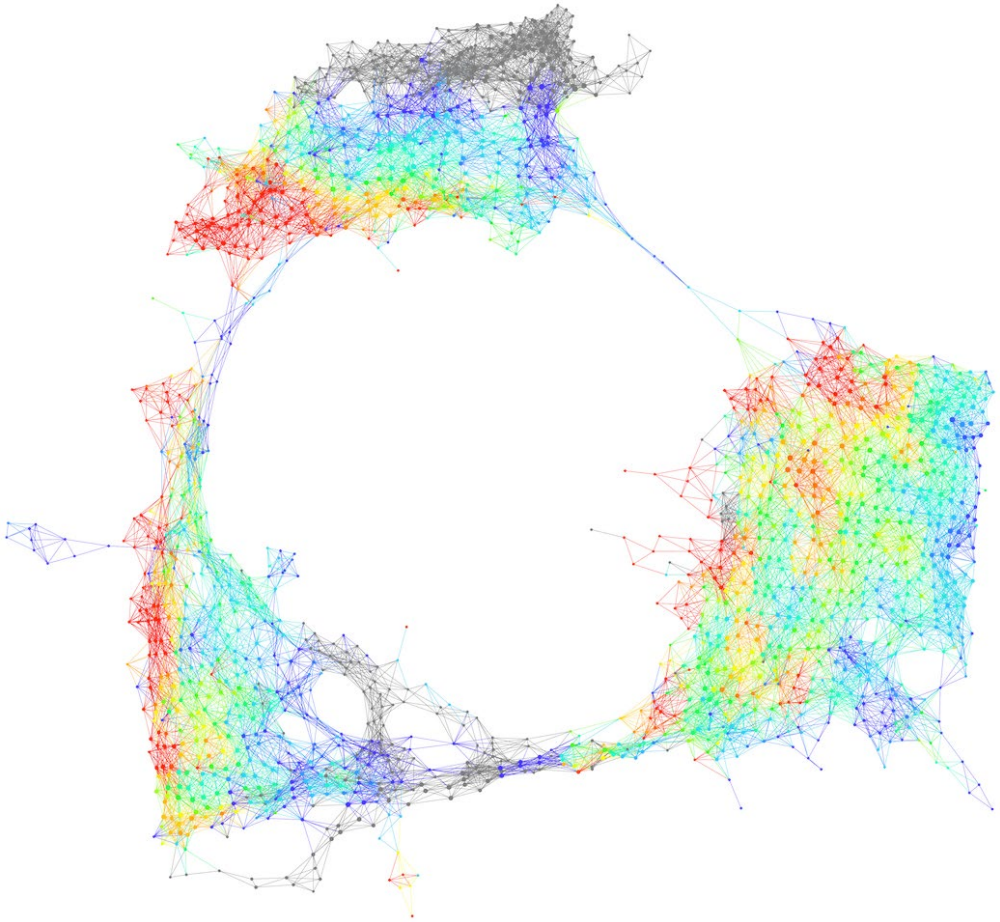

CS233 Key Course Goals

- Cover basic tools for geometric and topological data analysis, both supervised and unsupervised
- Discuss mathematical ways, based on geometry and topology, to encode and transfer knowledge about the data
- Introduce methods for joint data analysis and joint machine learning – benefiting from the “wisdom of the collection”

Integrate Data Across Multiple Modalities

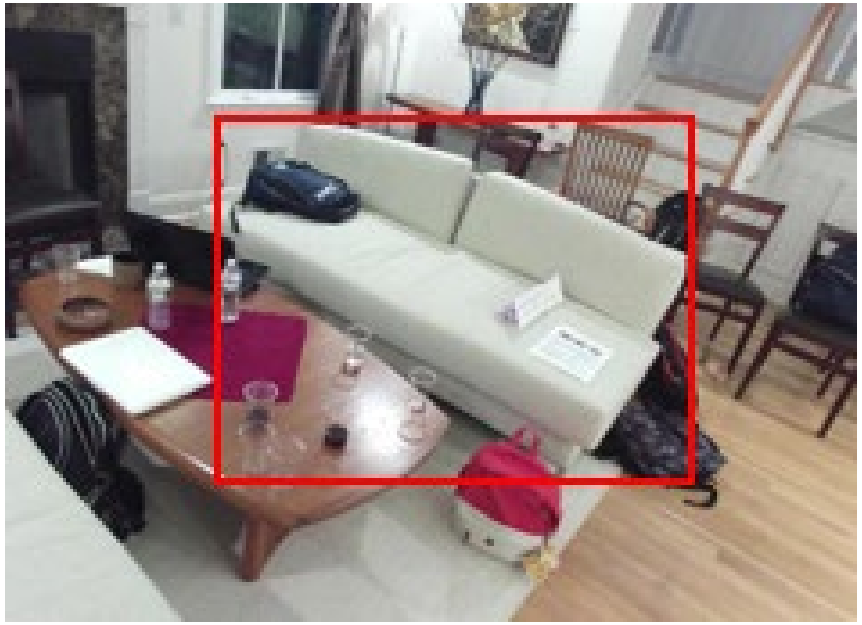

- + High resolution
- + Dense coverage
- - Subject to many imaging artifacts

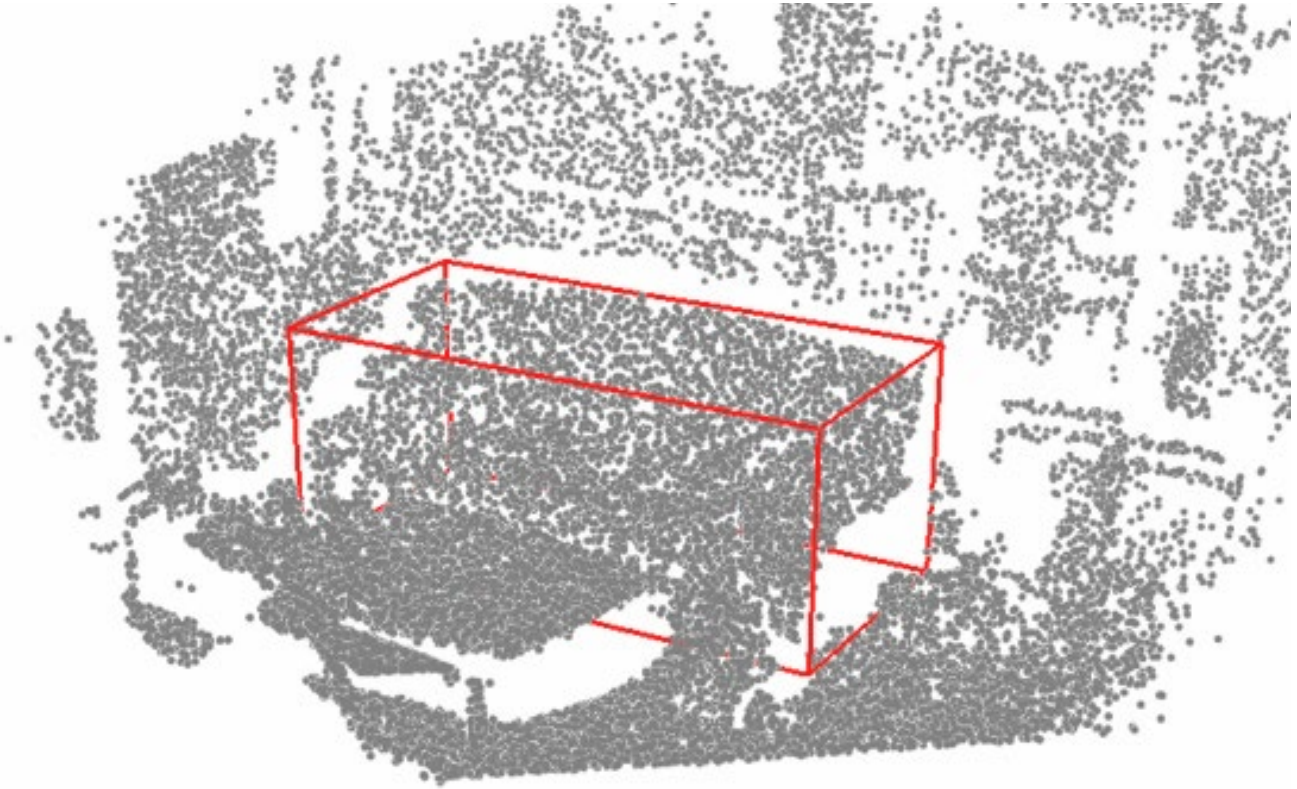

- + Absolute depth and scale
- - Sparse, low rez

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

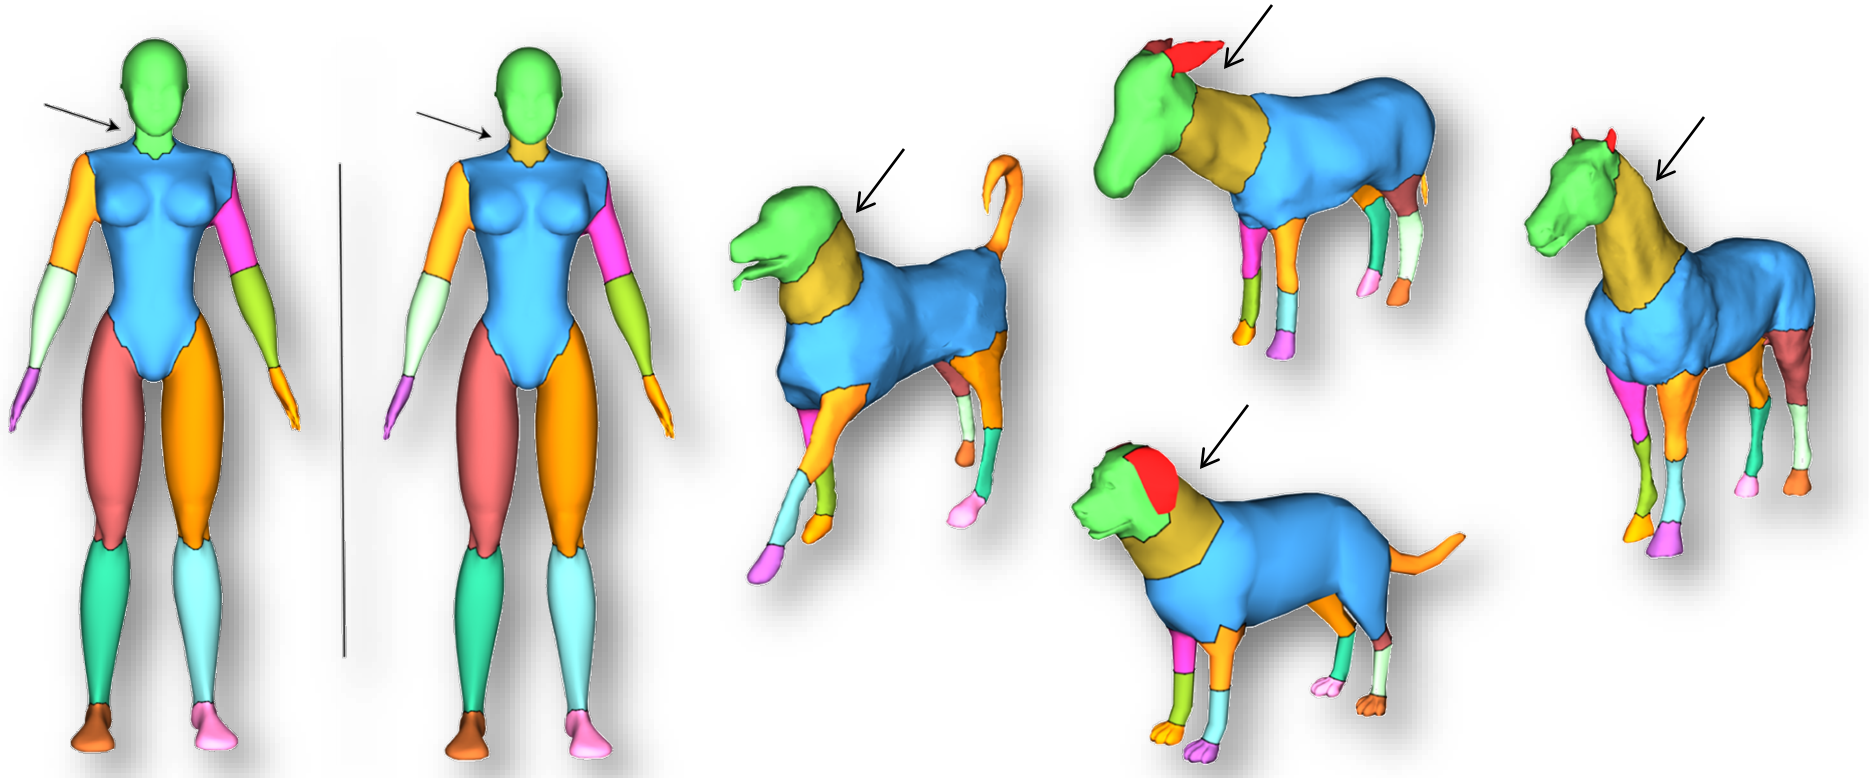

3D Segmentation

Look at Data from Many View Points

- Algebraic
- Graph theory based
- Geometric
- Topological

Getting to the semantics of data

Caveats

- Relatively new course in a fast evolving area – still some rough edges, especially in the current all virtual format

- Will cover a very wide variety of techniques: LA, ML, Stat, optimization, geometry processing, computer vision, algebraic topology …

geometric | data | analysis

- 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, irregular data types (e.g., point clouds, graphs) and unsupervised methods, in addition to the geometry / topology angle

The Team

The Principals

- Leonidas (Leo) Guibas (CS & EE)
  - Instructor
- Davis Rempe (CS)
  - Course Assistant
- Yueqi Duan (CS)
  - Course Assistant
- Kaichun Mo (CS)
  - Guest lecturer
- Samir Chowdhury (Neuropsychiatry)
  - Guest lecturer
- Carrie Petersen (CS)
  - Admin


Yueqi


Davis


Leo

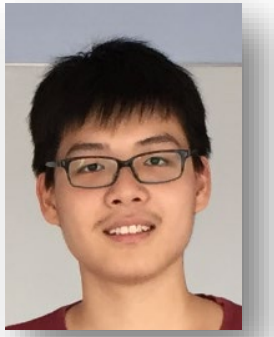
Kaichun


Samir


Carrie

Unsupervised Methods The Linear Space View of Data

Embedding Data into a Euclidean Space

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

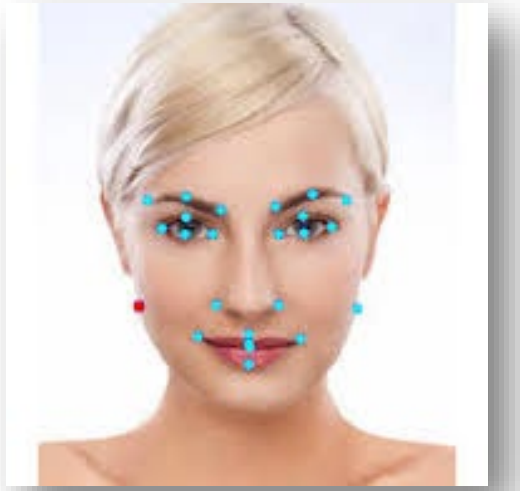

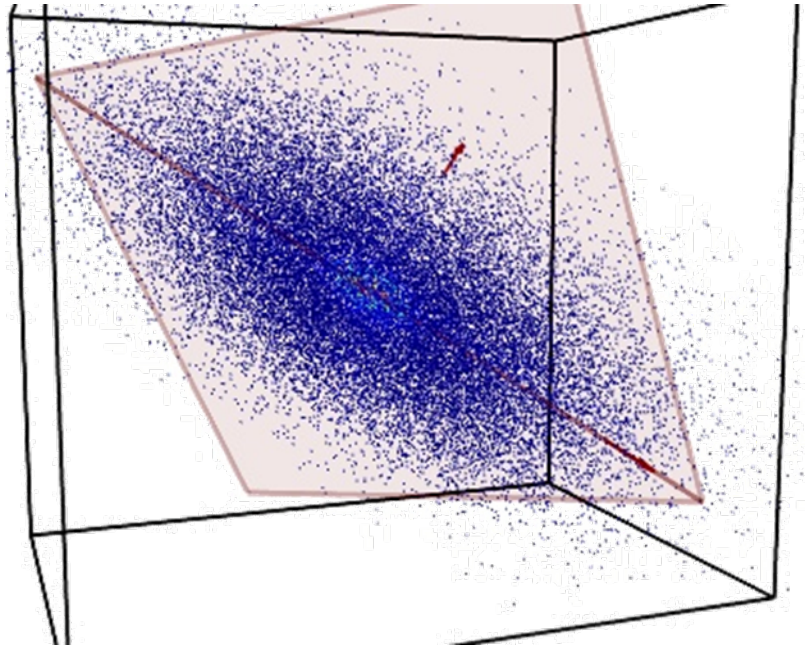

Direct Embeddings

- Input

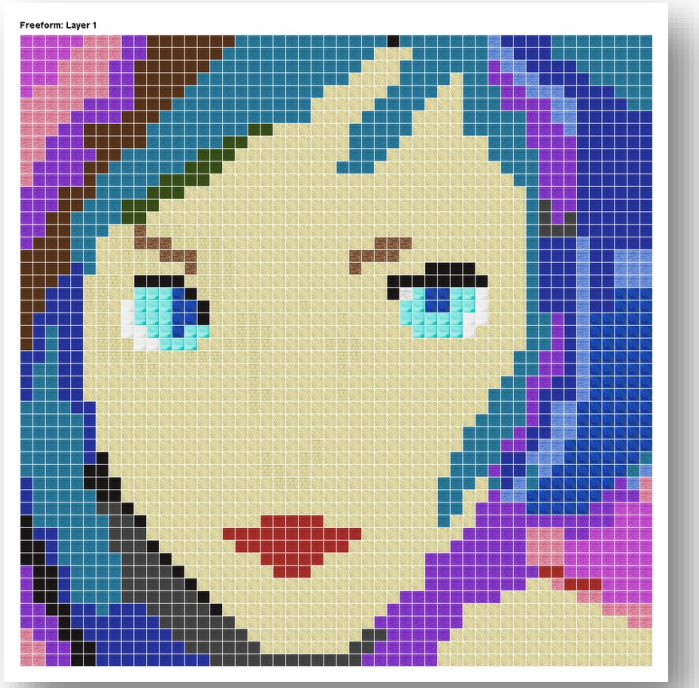

- Latent code

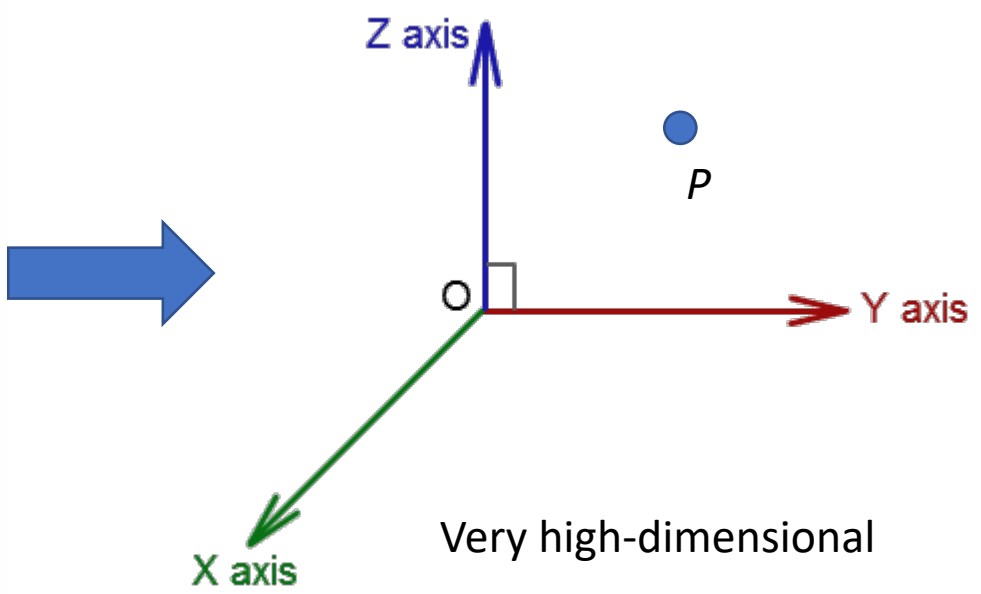

Very high-dimensional

Linear Space Methods

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

Dimensionality reduction

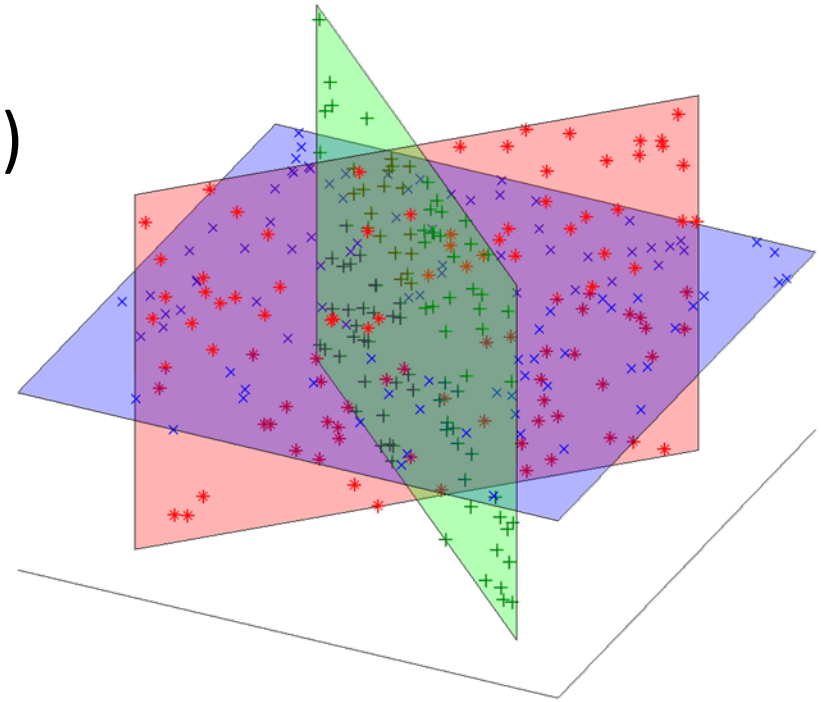

Low-d data (intrinsic dimension) living on a linear subspace, inside a high-d space (extrinsic dimension)

Data as Points on a Manifold

- Non-linear dimensionality reduction
- Low-d data inside high-d space may lie on a non-flat manifold

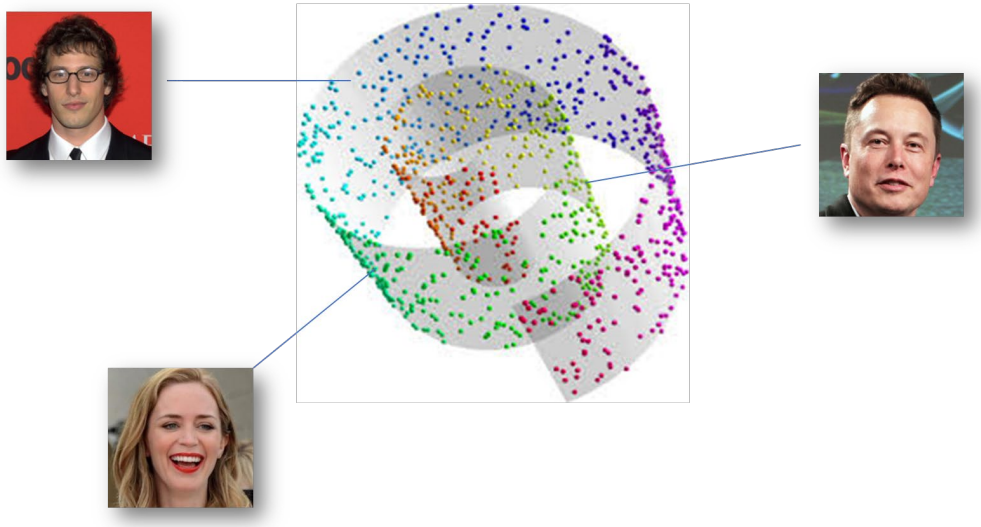

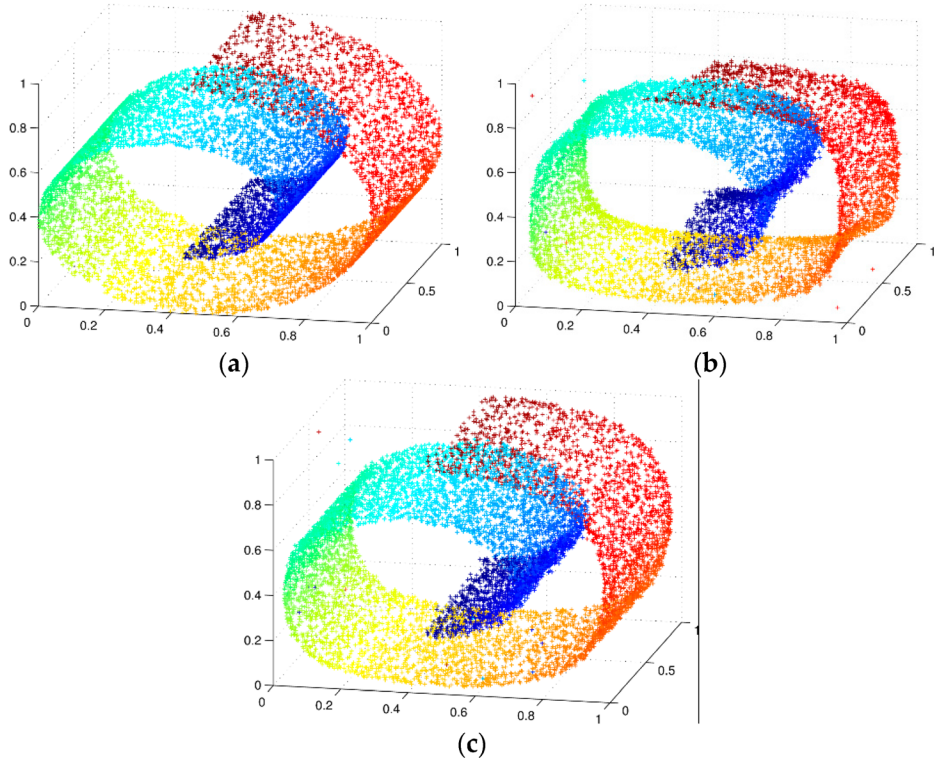

Isomap, locally linear embeddings, Laplacian eigenmaps, t-SNE

Supervised Methods
Deep Learning

Trainable Feed-Forward Learning Pipelines

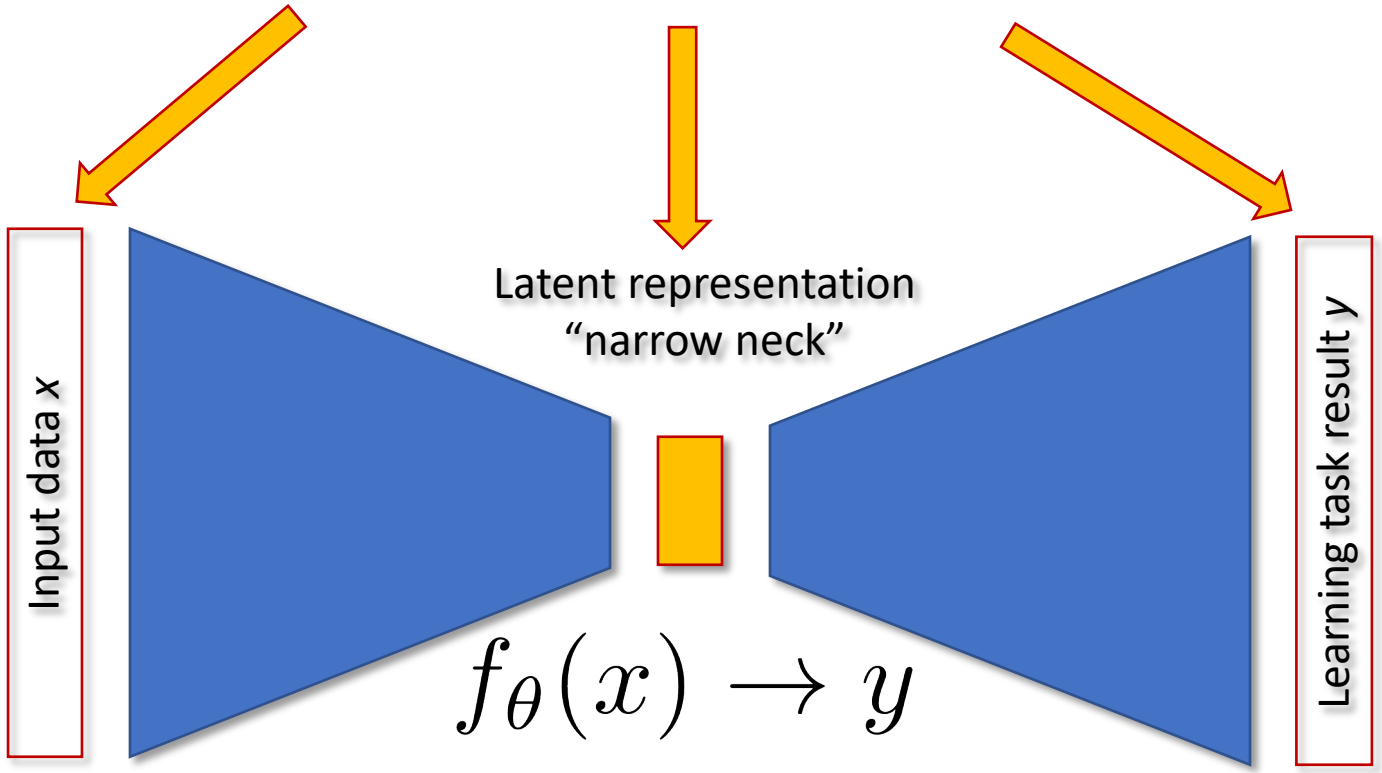

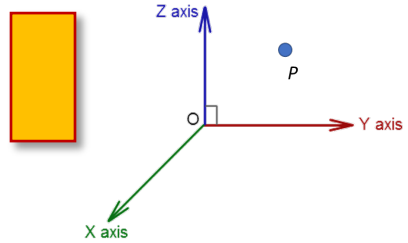

A latent code (= a point in a Euclidean feature space) acts as a low-d proxy for high-d input data, w.r.t. a learning task.

Image Deep Learning Networks

- Feed forward networks

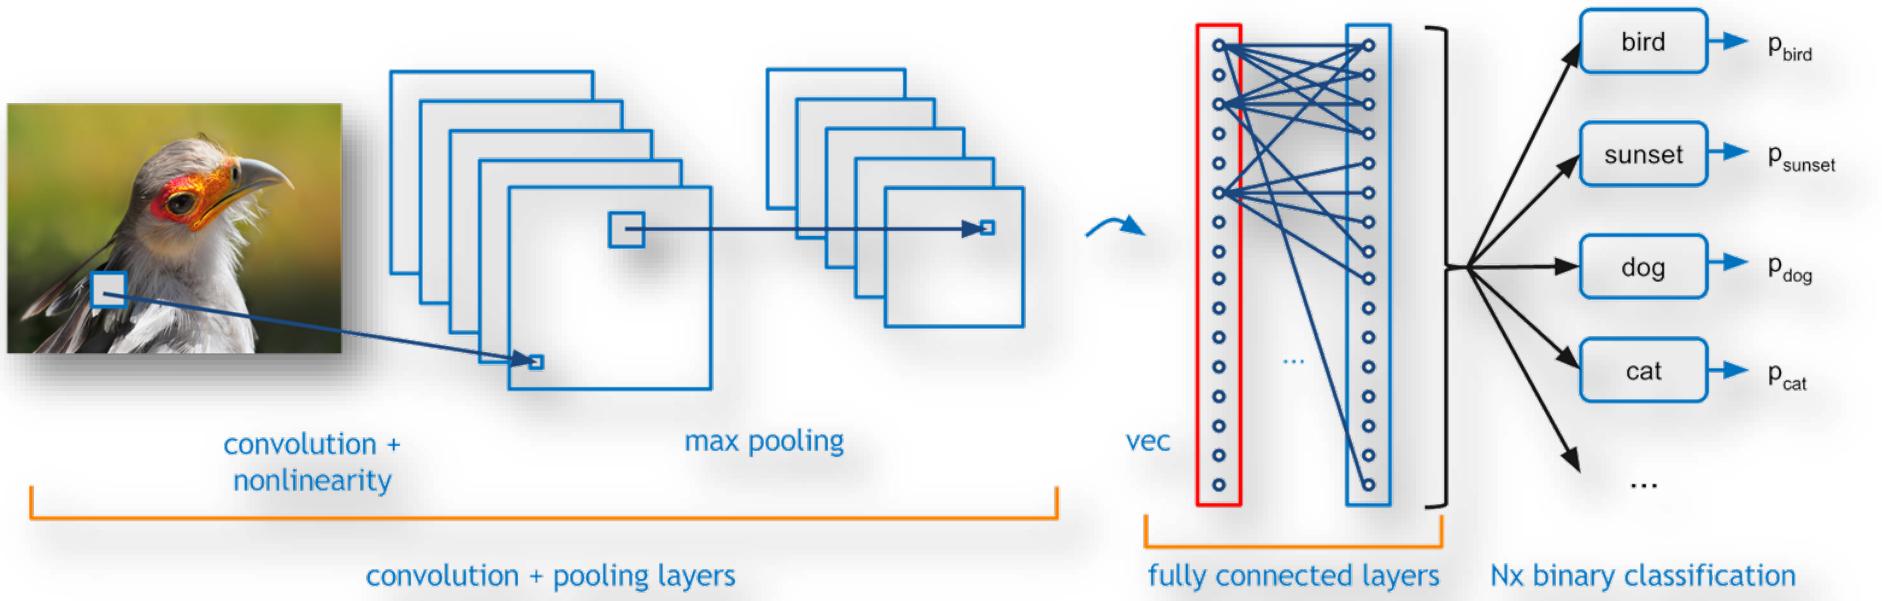

Classification

Convolutional Layers

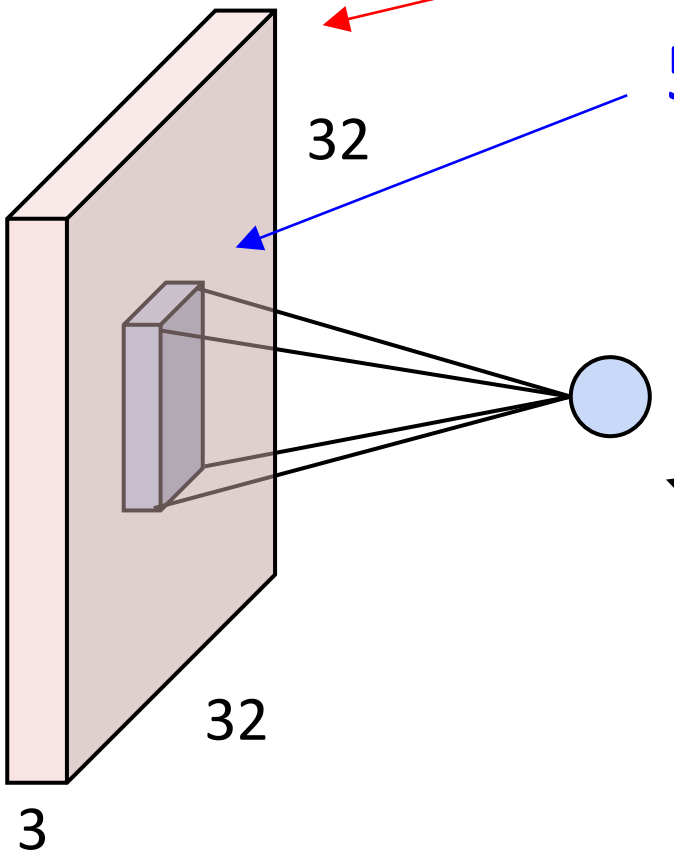

32x32x3 image

5x5x3 filter w

1 number:
the result of taking a dot product between the filter and a small 5x5x3 chunk of the image
(i.e. 5*5*3 = 75-dimensional dot product + bias)

$$w^T x + b$$

Convolutional Architectures

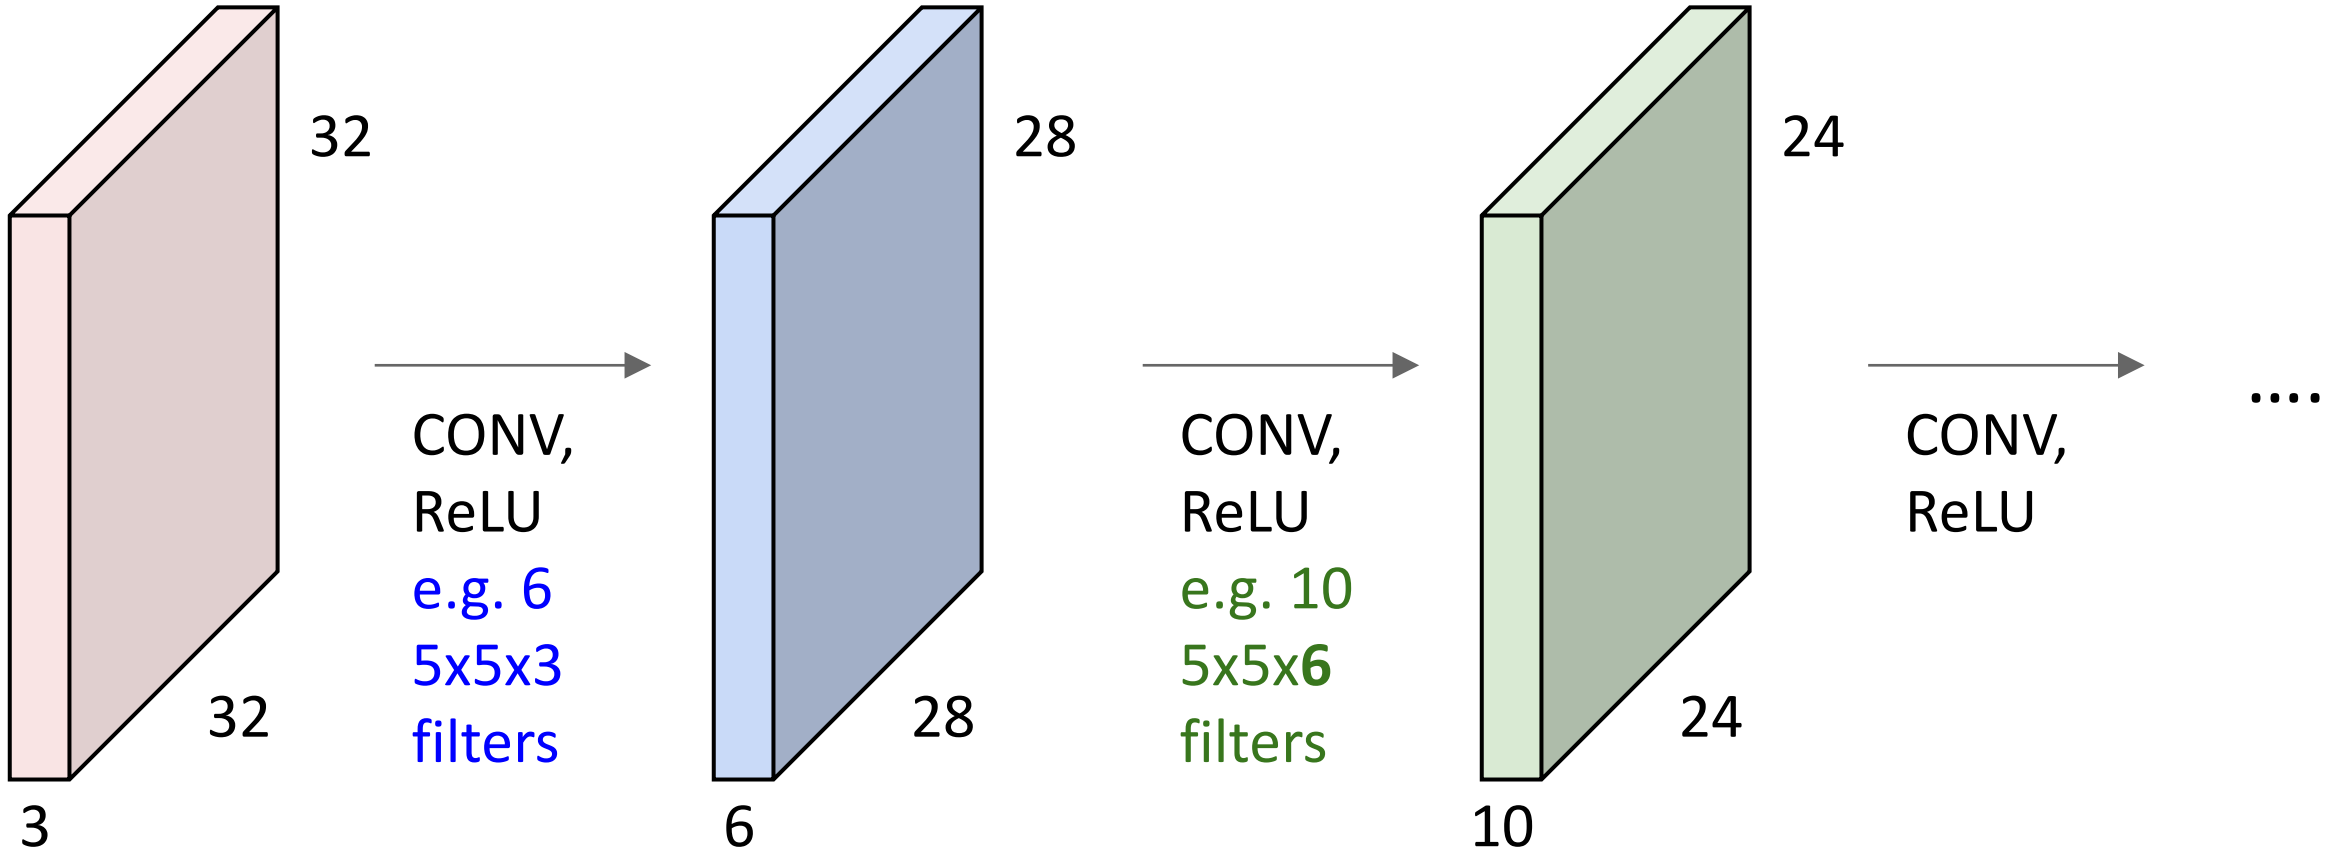

ConvNets are a sequence of convolutional layers, interspersed with activation functions

Pooling Layer

- makes the representations smaller and more manageable
- operates over each activation map independently:

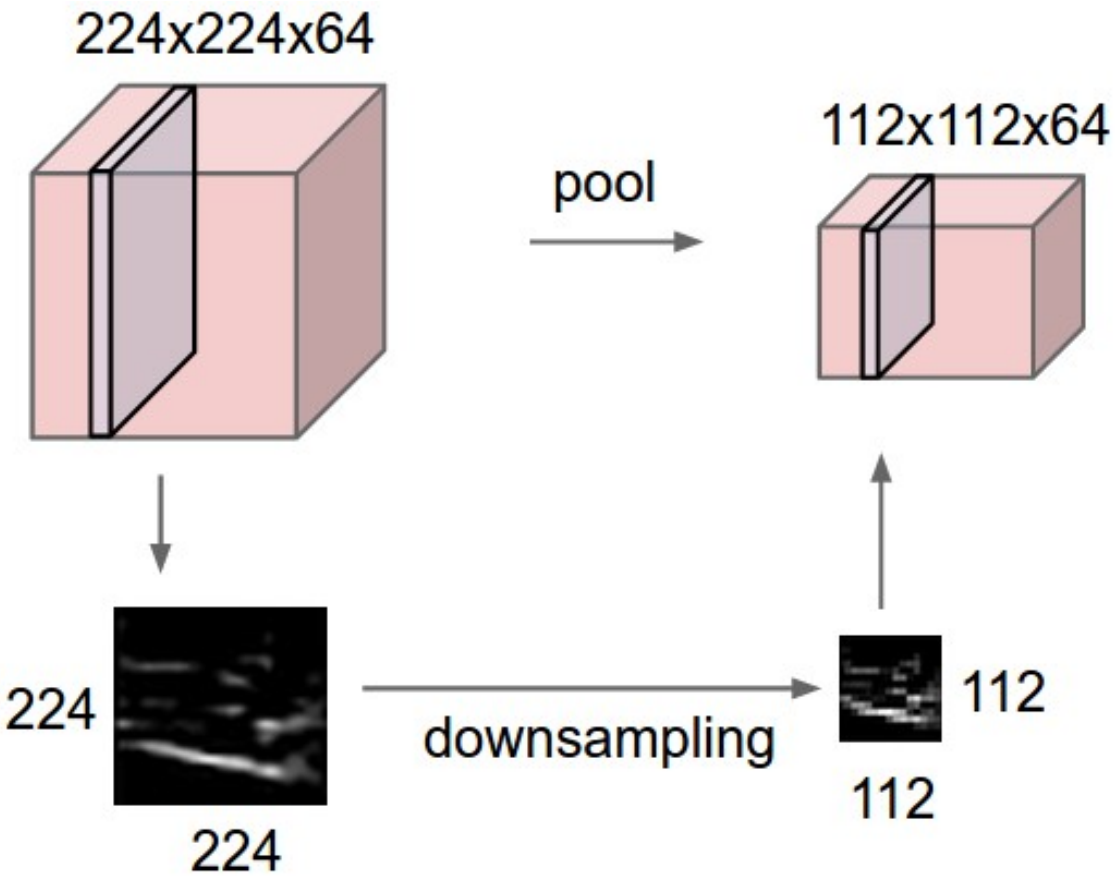

Data Abstraction Across Layers

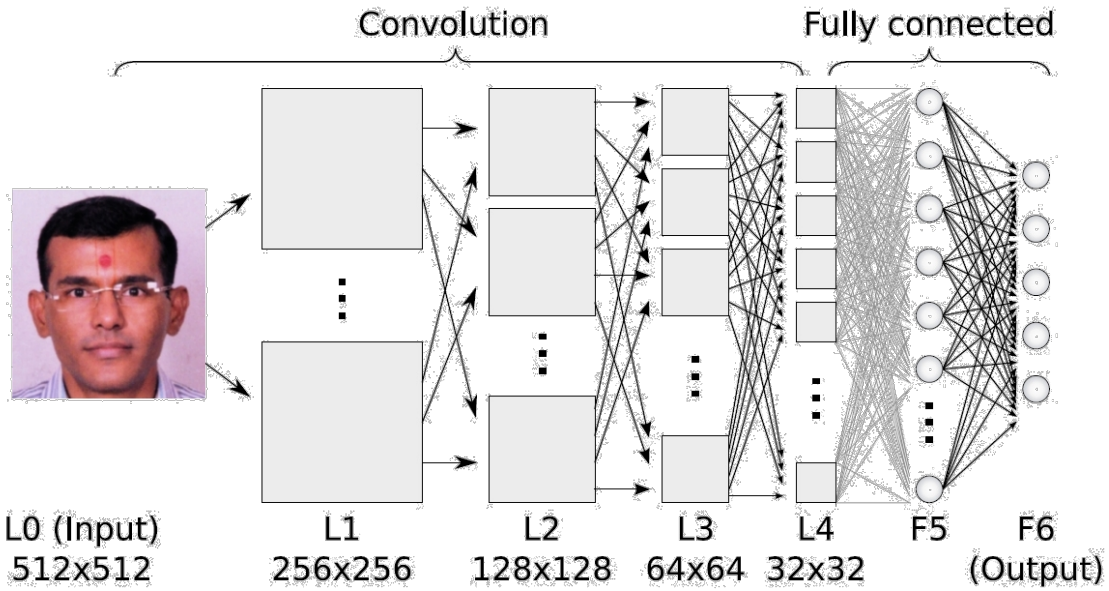

[Makwana, 2016]

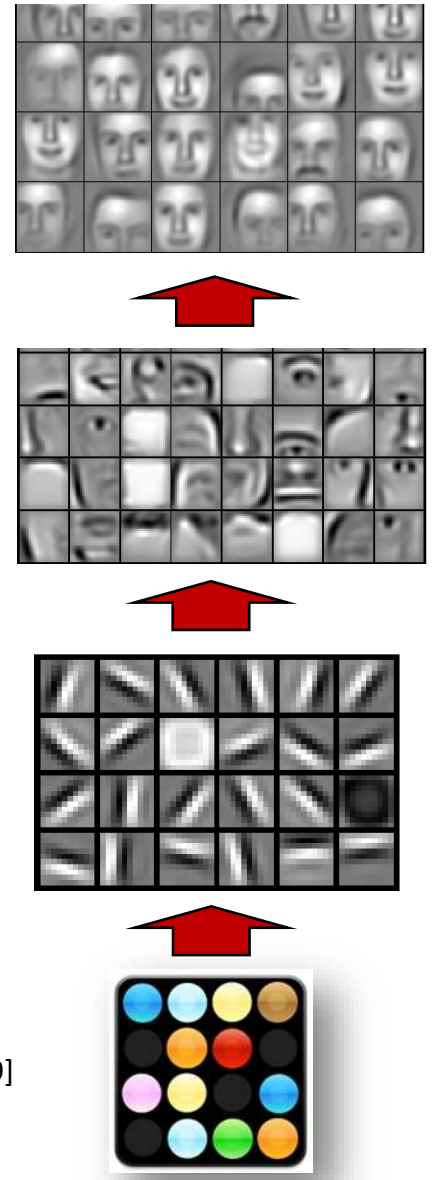

Data-driven feature learning at ascending abstraction layers

[Lee et al., 2009]

"Vertical" networks

ConvNets are Everywhere

Classification

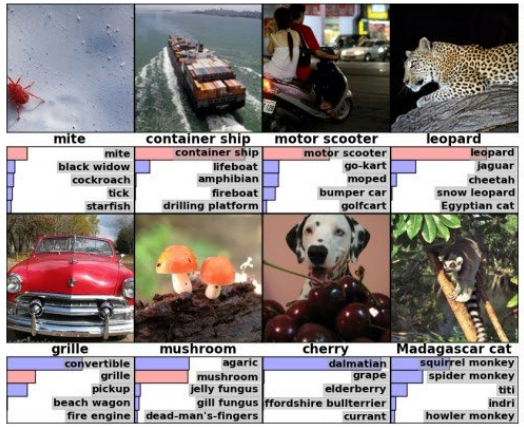

Retrieval

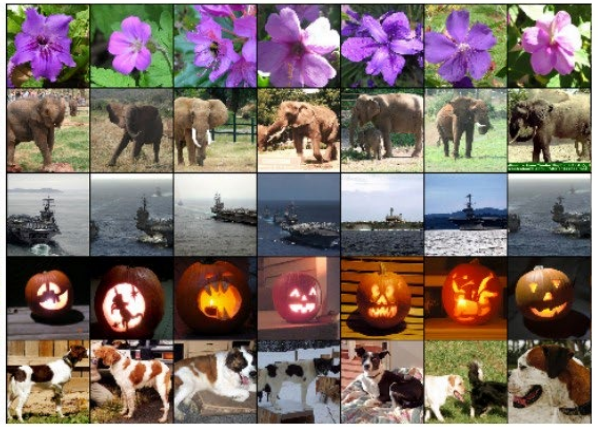

Detection

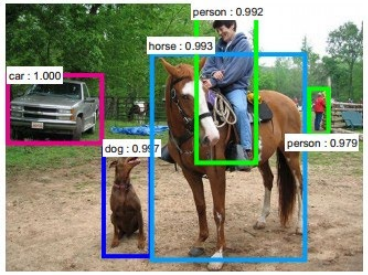

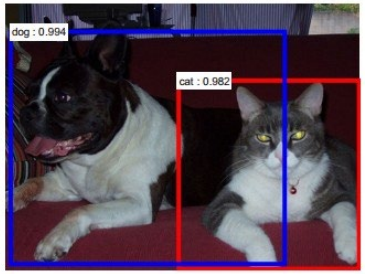

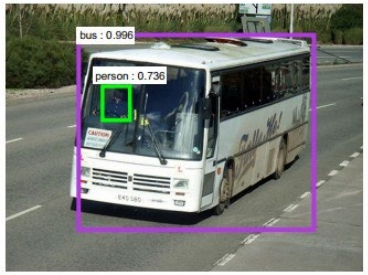

Text processing

I caught this movie on the Sci-Fi channel recently. It actually turned out to be pretty decent as far as B-list horror/suspense films go. Two guys (one naive and one loud mouthed a **) take a road trip to stop a wedding but have the worst possible luck when a maniac in a freaky, make-shift tank/truck hybrid decides to play cat-and-mouse with them. Things are further complicated when they pick up a ridiculously whorish hitchhiker. What makes this film unique is that the combination of comedy and terror actually work in this movie, unlike so many others. The two guys are likable enough and there are some good chase/suspense scenes. Nice pacing and comic timing make this movie more than passable for the horror/slasher buff. Definitely worth checking out.

I just saw this on a local independent station in the New York City area. The cast showed promise but when I saw the director, George Cosmotos, I became suspicious. And sure enough, it was every bit as bad, every bit as pointless and stupid as every George Cosmotos movie I ever saw. He's like a stupid man's Michael Bey -- with all the awfulness that accolade promises. There's no point to the conspiracy, no burning issues that urge the conspirators on. We are left to ourselves to connect the dots from one bit of graffiti on various walls in the film to the next. Thus, the current budget crisis, the war in Iraq, Islamic extremism, the fate of social security, 47 million Americans without health care, stagnating wages, and the death of the middle class are all subsumed by the sheer terror of graffiti. A truly, stunningly idiotic film.

Graphics is far from the best part of the game. This is the number one best TH game in the series. Next to Underground. It deserves strong love. It is an insane game. There are massive levels, massive unlockable characters... it's just a massive game. Waste your money on this game. This is the kind of money that is wasted properly. And even though graphics suck, thats doesn't make a game good. Actually, the graphics were good at the time. Today the graphics are crap. WHO CARES? As they say in Canada, This is the fun game, aye. (You get to go to Canada in THPS3) Well, I don't know if they say that, but they might. who knows. Well, Canadian people do. Wait a minute, I'm getting off topic. This game rocks. Buy it, play it, enjoy it, love it. It's PURE BRILLIANCE.

The first was good and original. I was a not bad horror/comedy movie. So I heard a second one was made and I had to watch it . What really makes this movie work is Judd Nelson's character and the sometimes clever script. A pretty good script for a person who wrote the Final Destination films and the direction was okay. Sometimes there's scenes where it looks like it was filmed using a home video camera with a grainy - look. Great made - for - TV movie. It was worth the rental and probably worth buying just to get that nice eerie feeling and watch Judd Nelson's Stanley doing what he does best. I suggest newcomers to watch the first one before watching the sequel, just so you'll have an idea what Stanley is like and get a little history background.

Segmentation

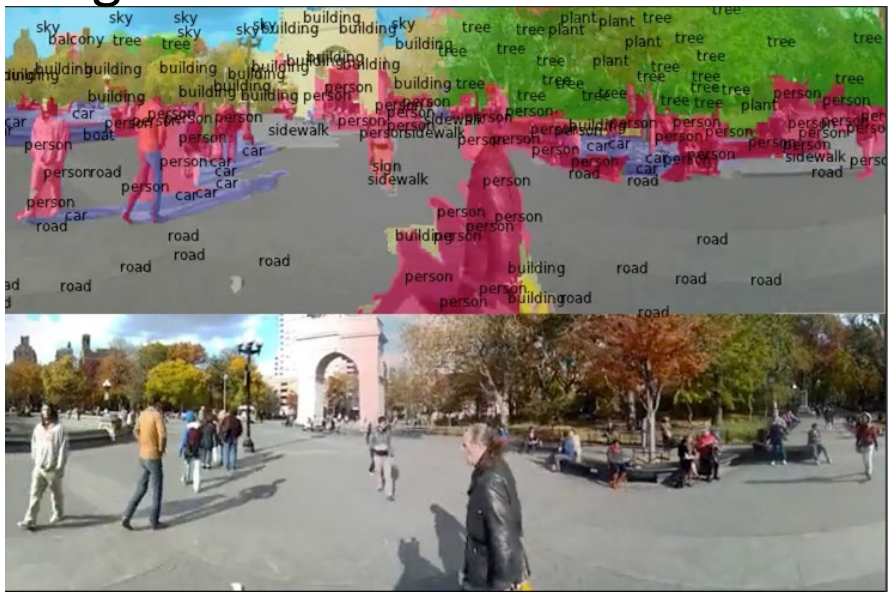

Control

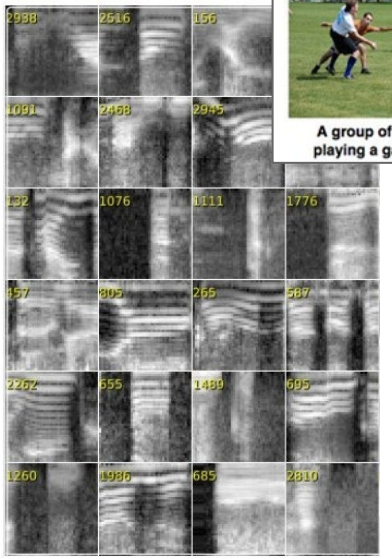

Image captioning

Speech processing

Success Made Possible By ….

[He et al., 2015]
Novel deep architectures

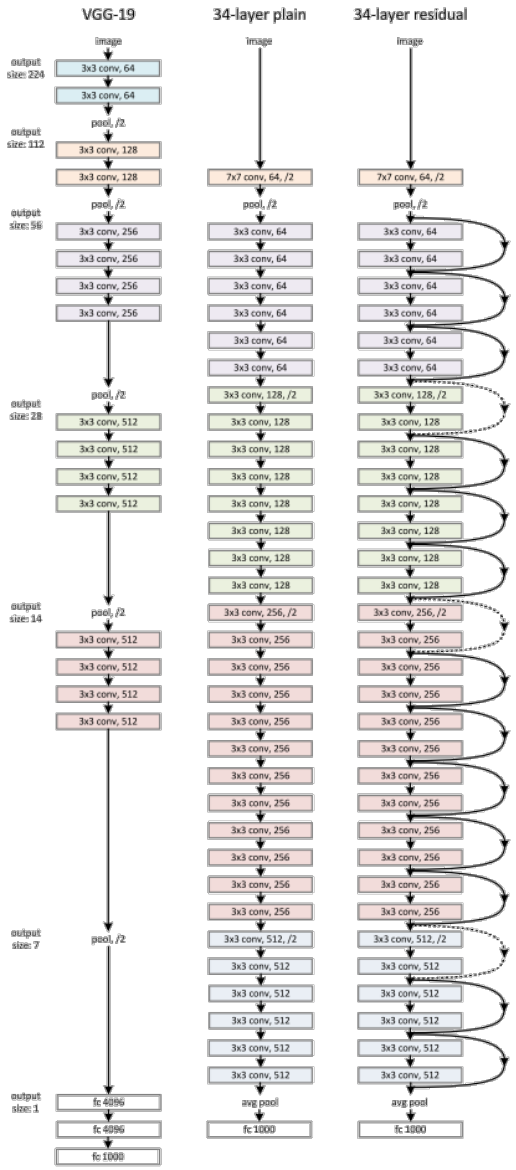

Plenty of annotated data

Lots of computing power

Learning on Irregular Data

Convolutional Networks

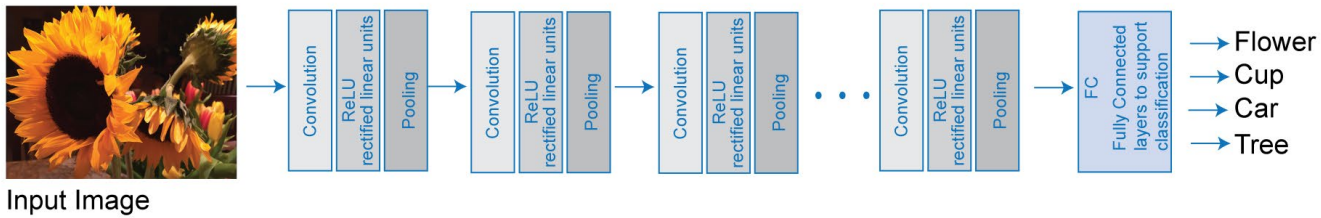

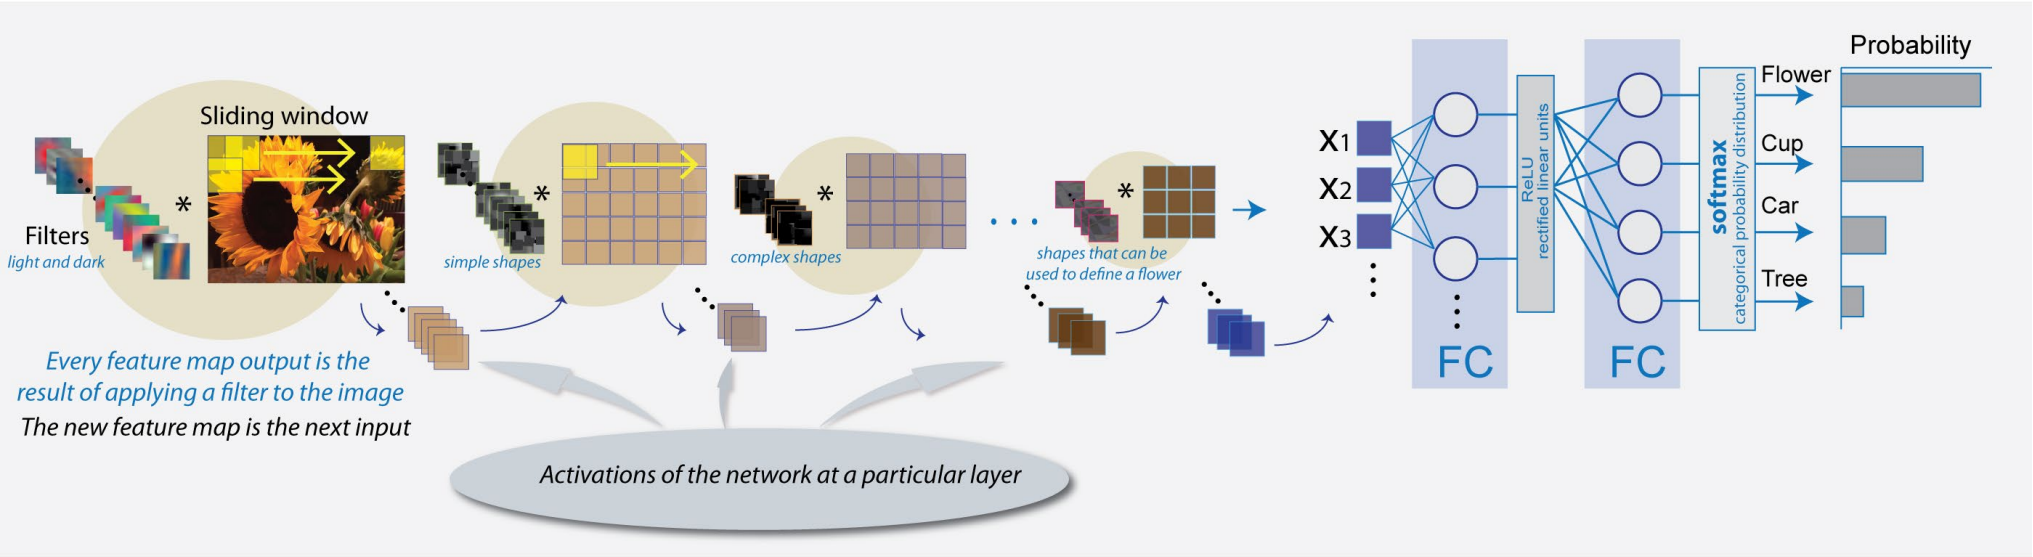

Convolutions require regular array type inputs for weight sharing and other optimizations

Multiple Representations for 3D Geometric Data

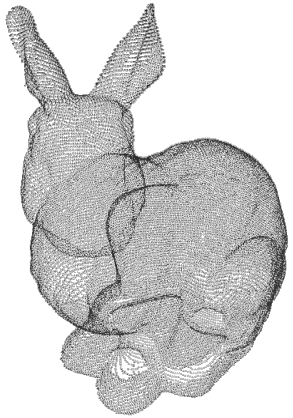

Point Cloud

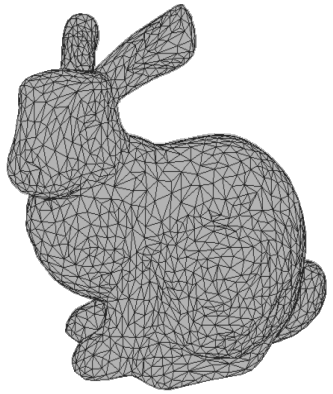

Mesh

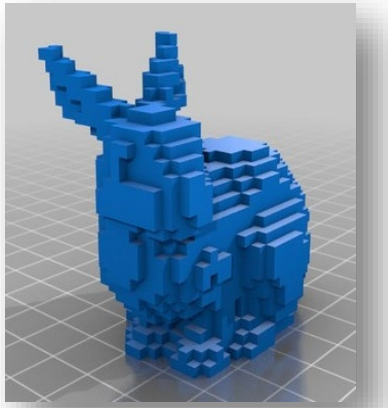

Volumetric

Projected View

RGB(D)

...

These are irregular representations

Deep Learning for 3D Geometry Analysis

Classification

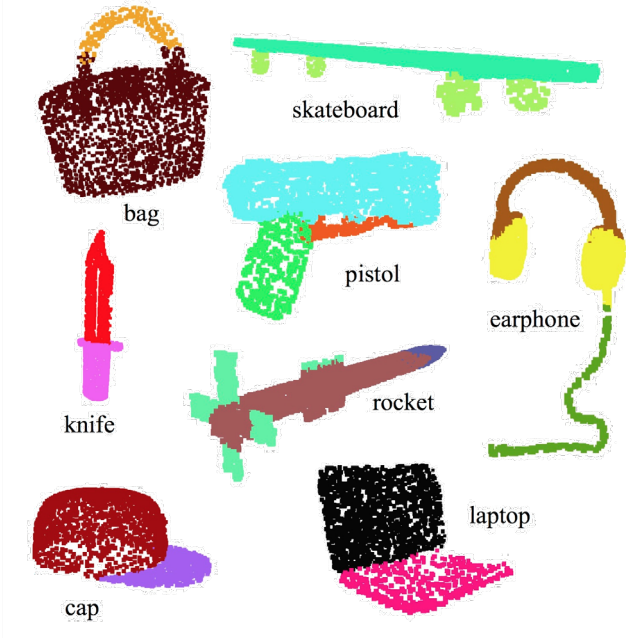

Segmentation/Parsing
(object/scene)

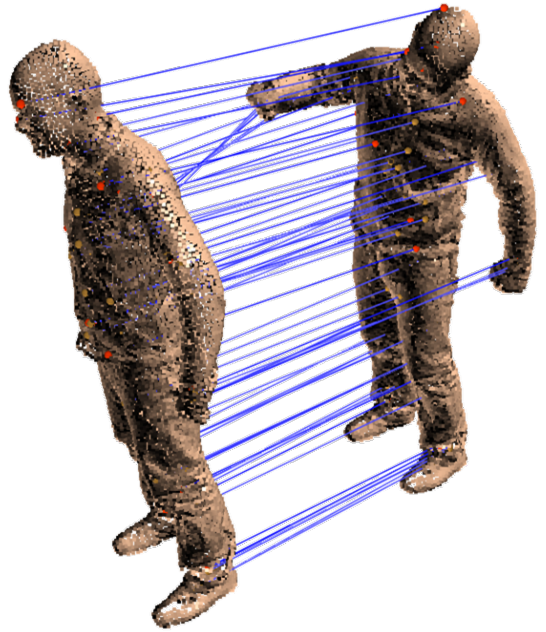

Correspondences

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

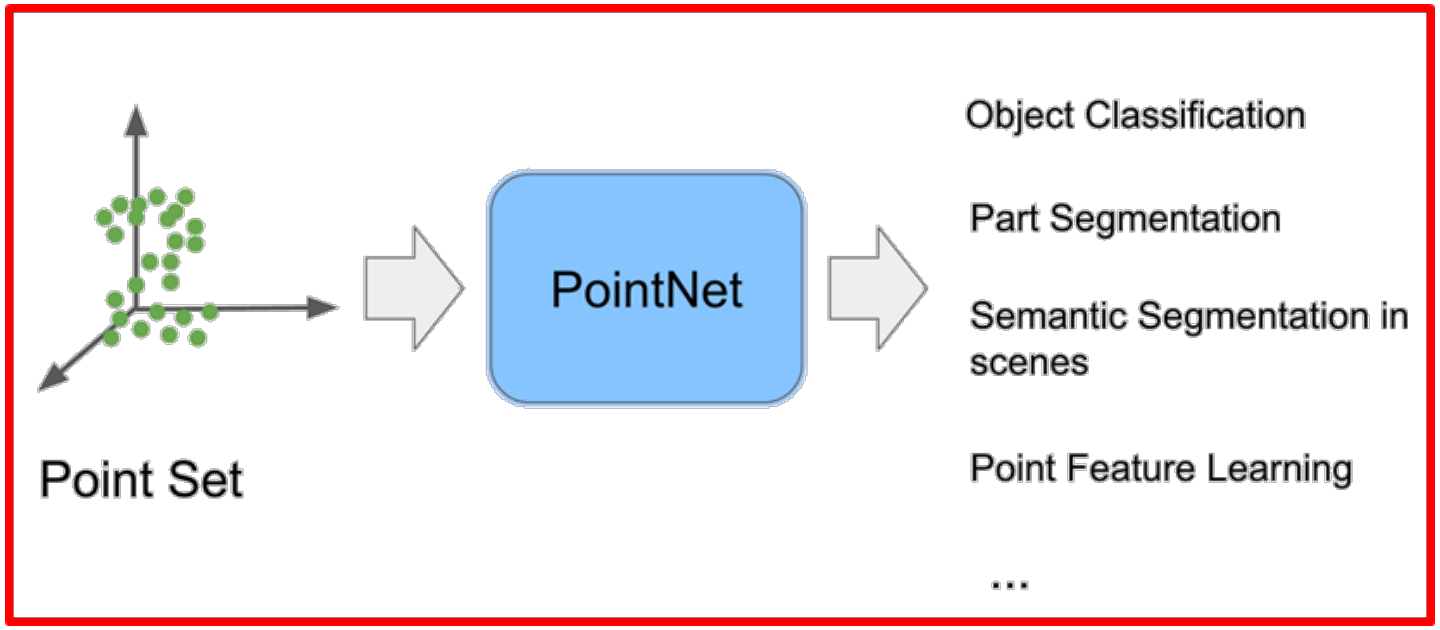

The Graph View of Data

Graphs and Simplicial Complexes

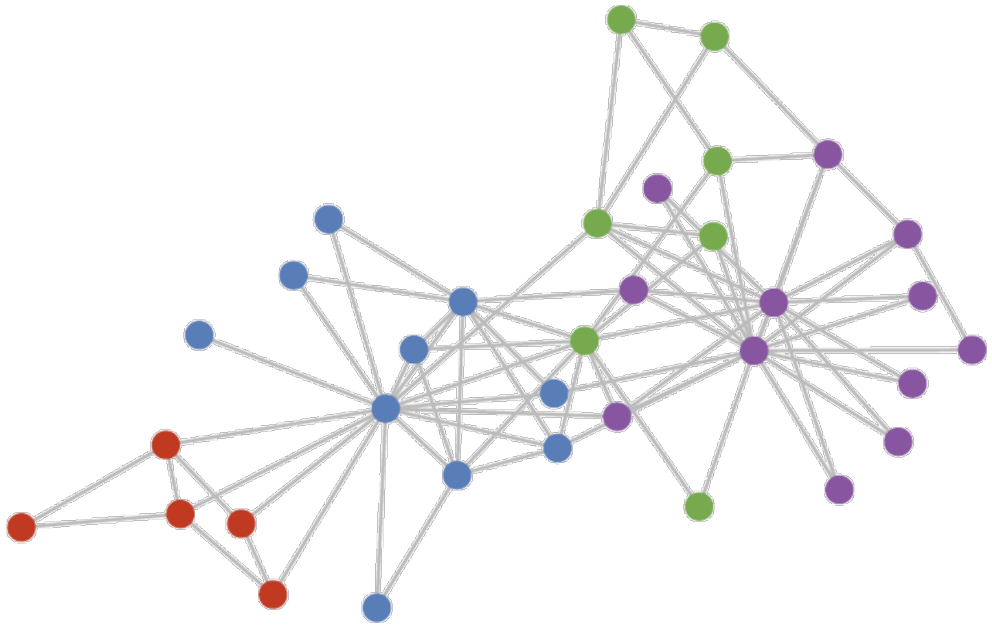

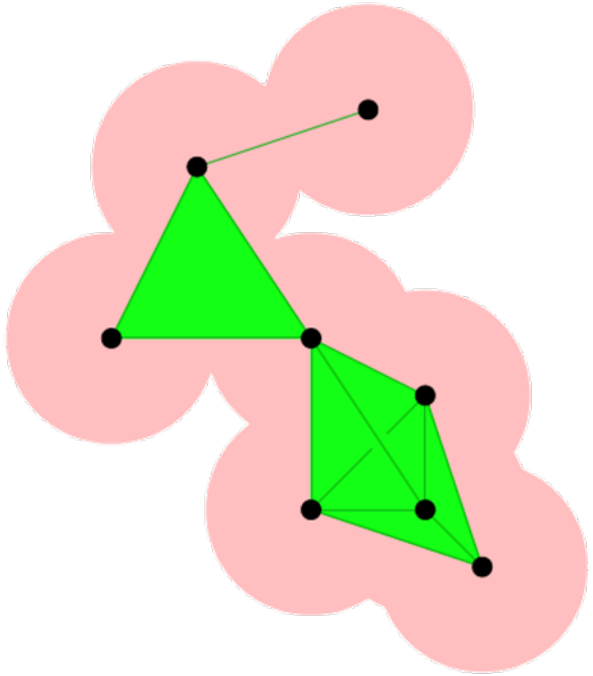

A Graph View of Data

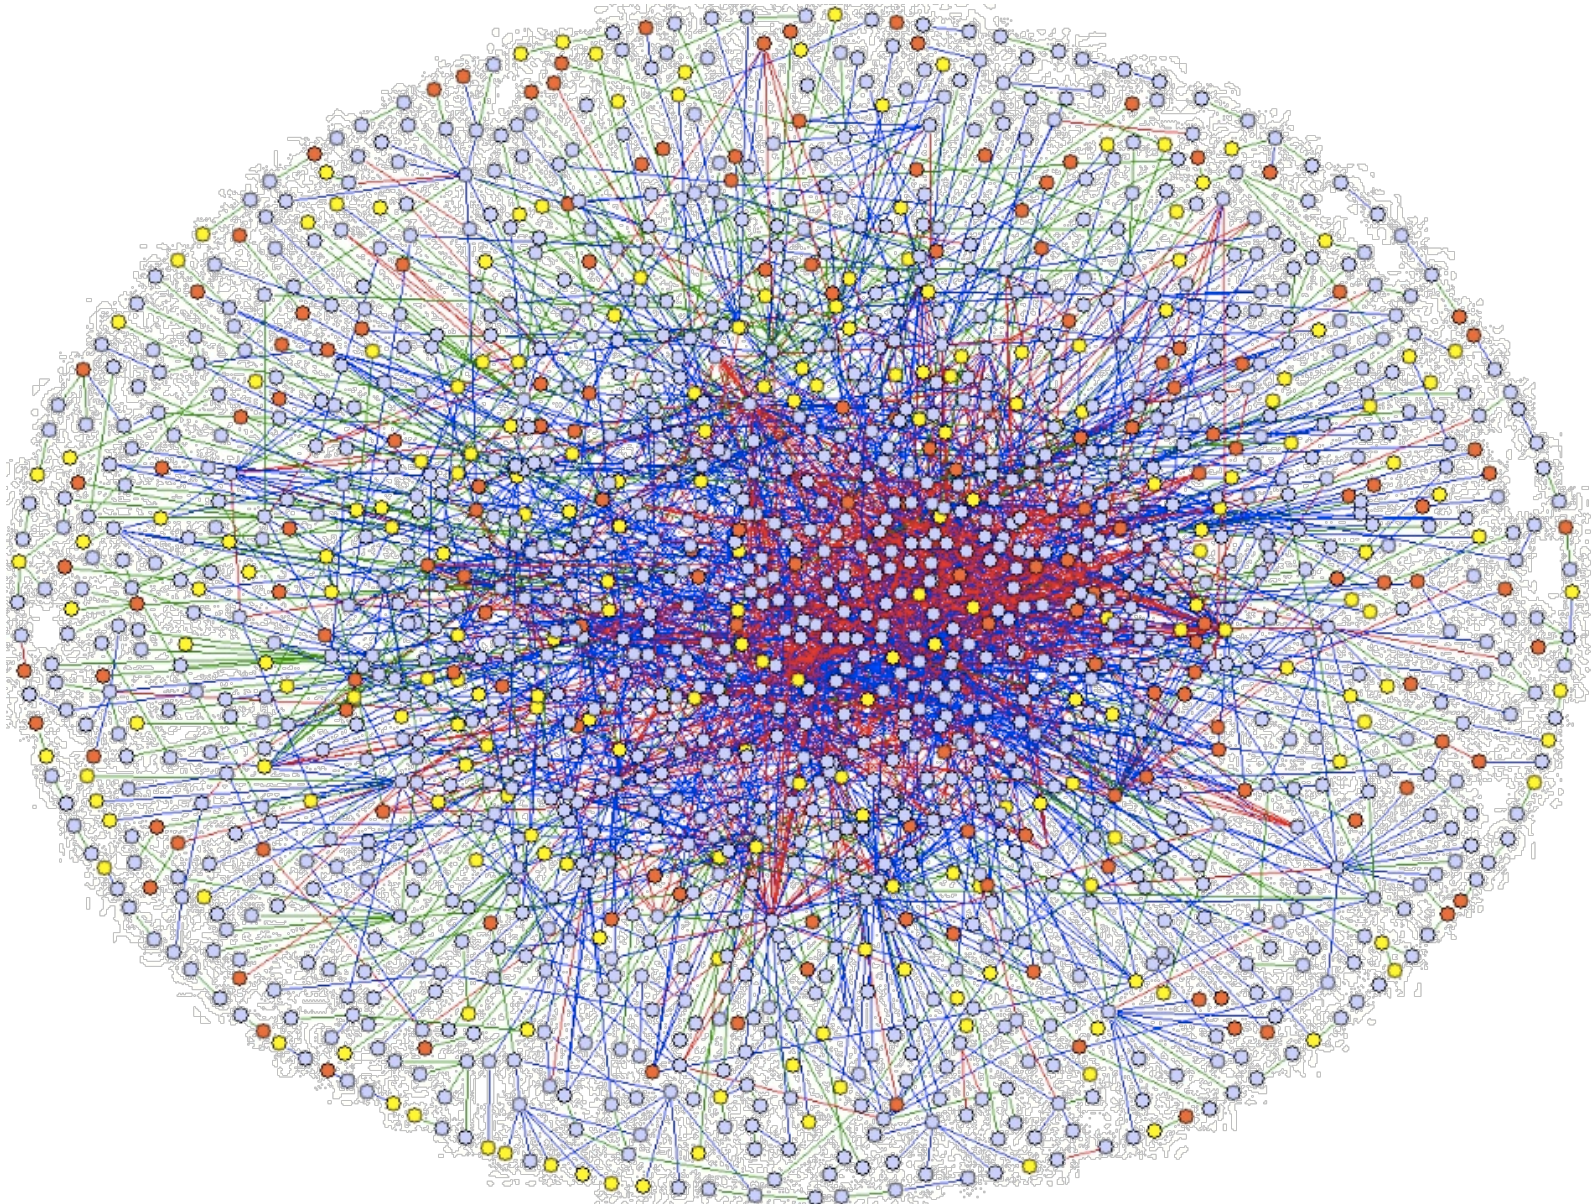

Spectral Methods in Graph Theory

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

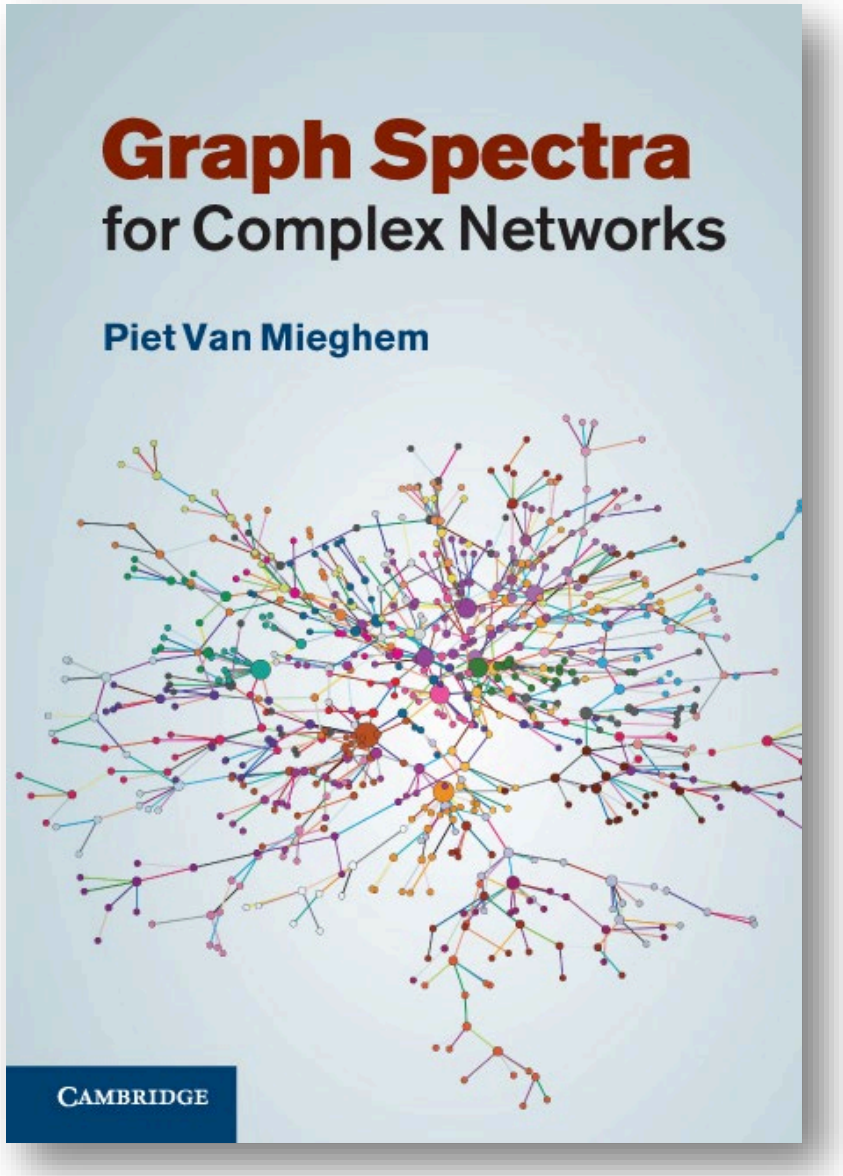

From Feature Vectors to Feature Sets

- What to do if the number of features per data set is variable?
  - “simple” vs “complex” objects

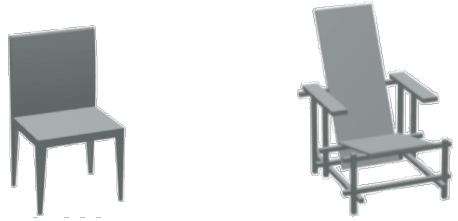

- Distance or similarity metrics

- Multidimensional scaling

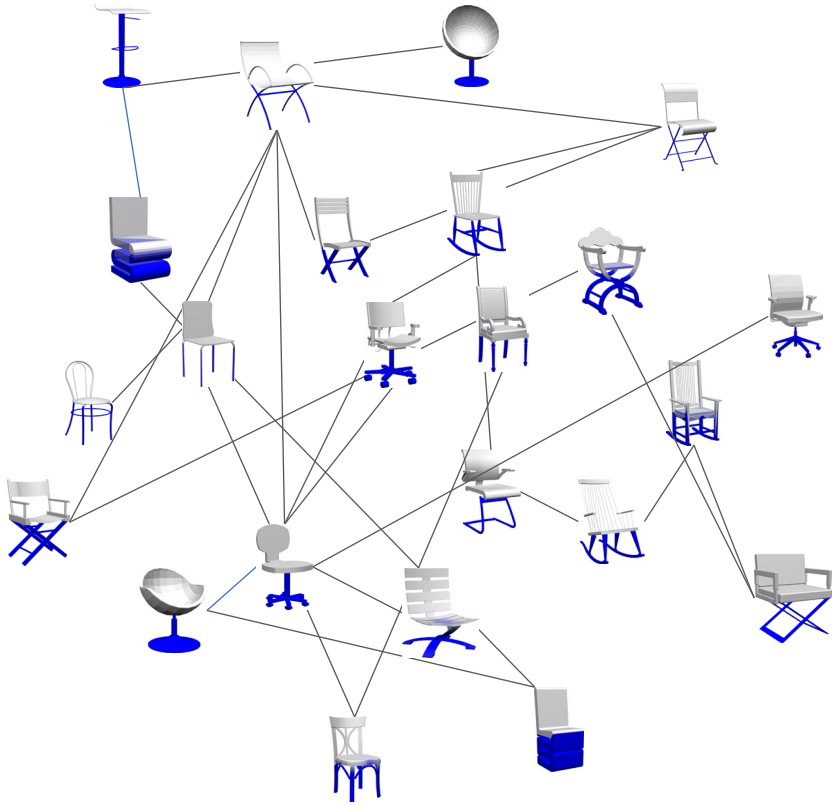

A “Procrustean” approach

Topological Data Analysis (TDA)

Computational Topology

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

- Large distances (aka “similarity metrics” are often suspect …

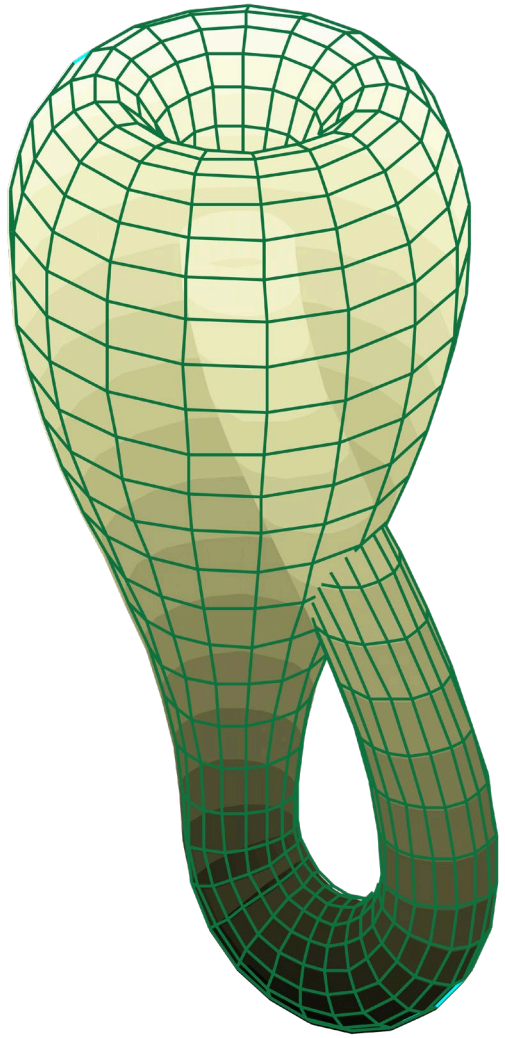

In TDA, Sampled Spaces: “The Shape of Data”

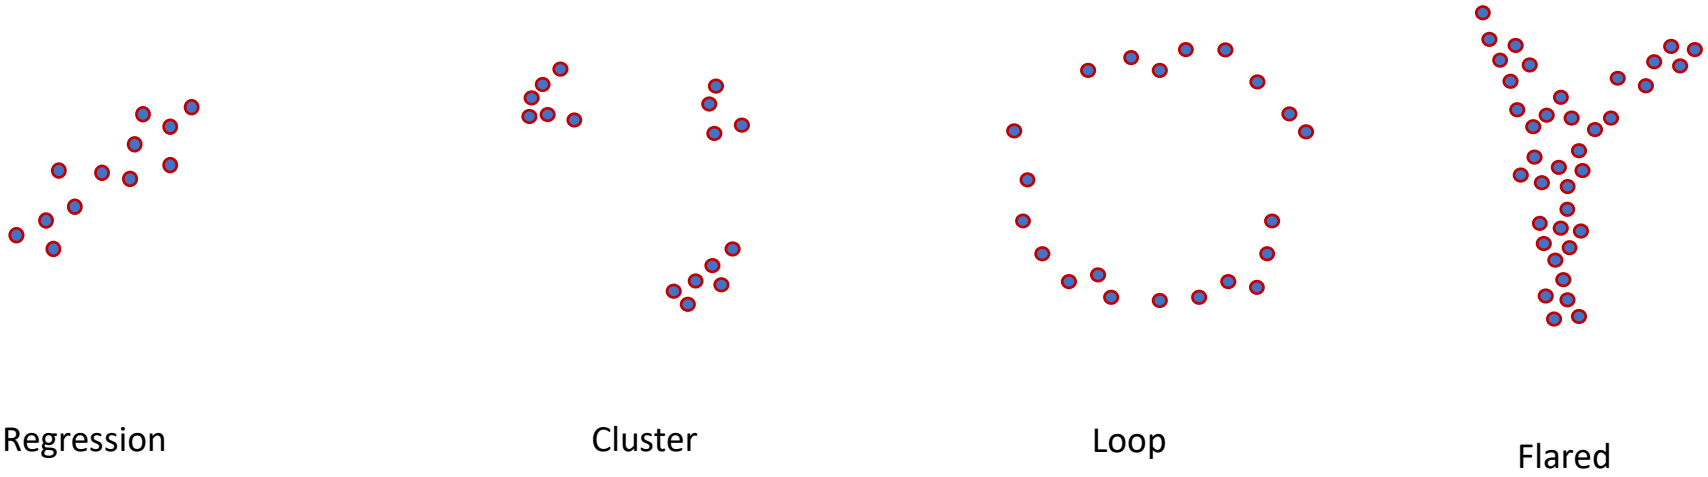

Regression

Cluster

Loop

Flared

Simplicial Complexes

- A simplicial complex K is a finite set of simplices such that
  1. $\sigma \in K, \tau \leq \sigma \Rightarrow \tau \in K$,
  2. $\sigma, \sigma' \in K \Rightarrow \sigma \cap \sigma' \leq \sigma, \sigma'$ or $\sigma \cap \sigma' = \emptyset$.
- The dimension of K is $\dim K = \max\{\dim \sigma \mid \sigma \in K\}$.
- The vertices of K are the zero-simplices in K.
- A simplex is principal if it has no proper coface in K.

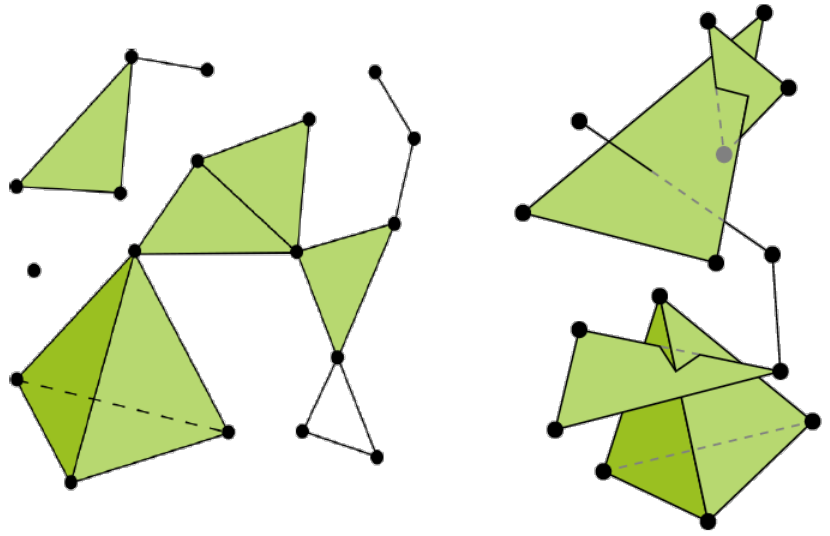

(left) an example (right) a non example

A Multiscale View of Data

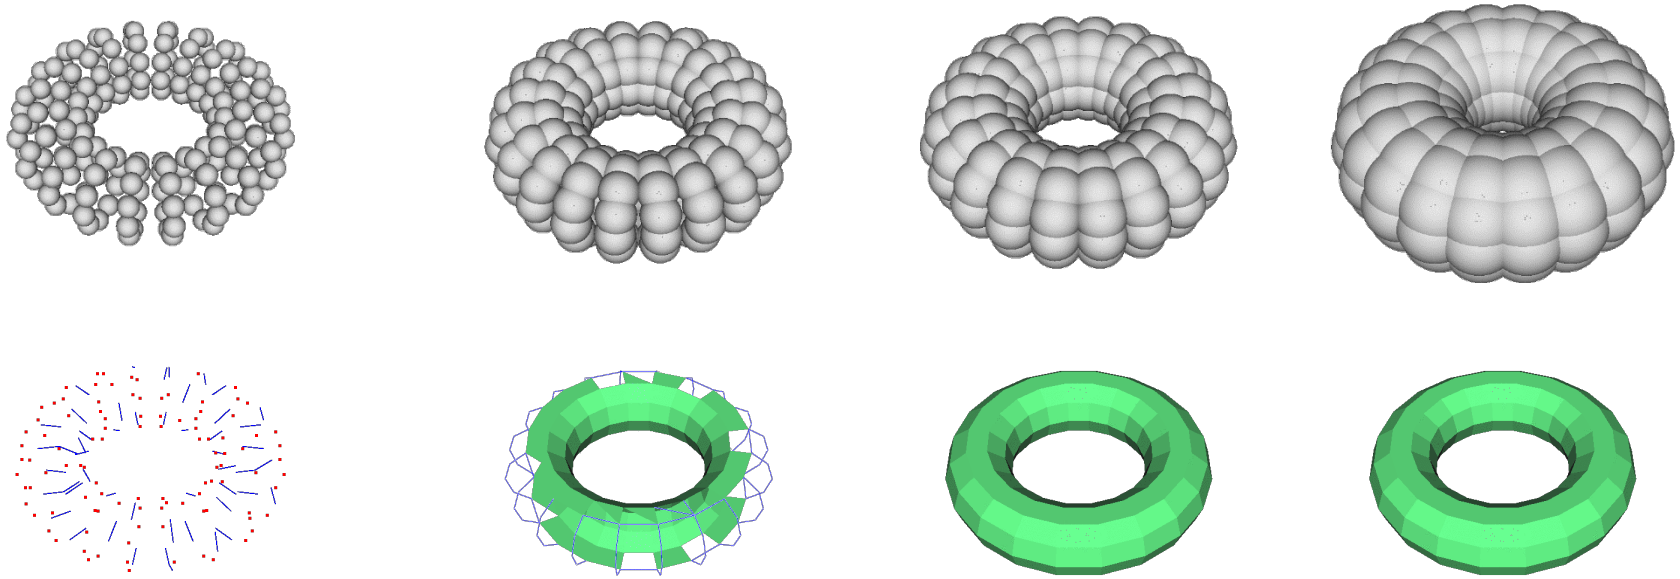

$\beta_0 = 150$
$\beta_1 = 0$

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

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

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

Persistent Homology

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

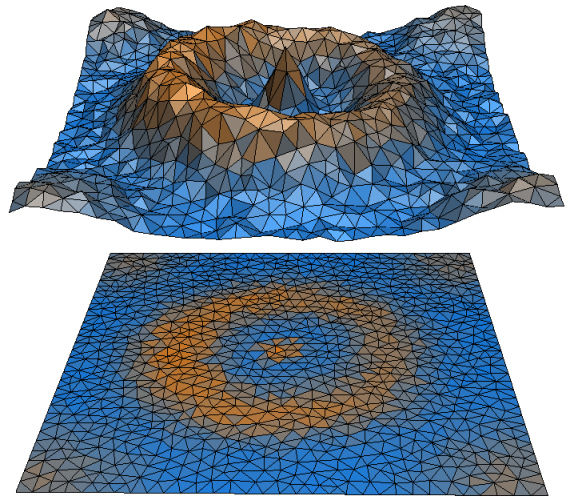

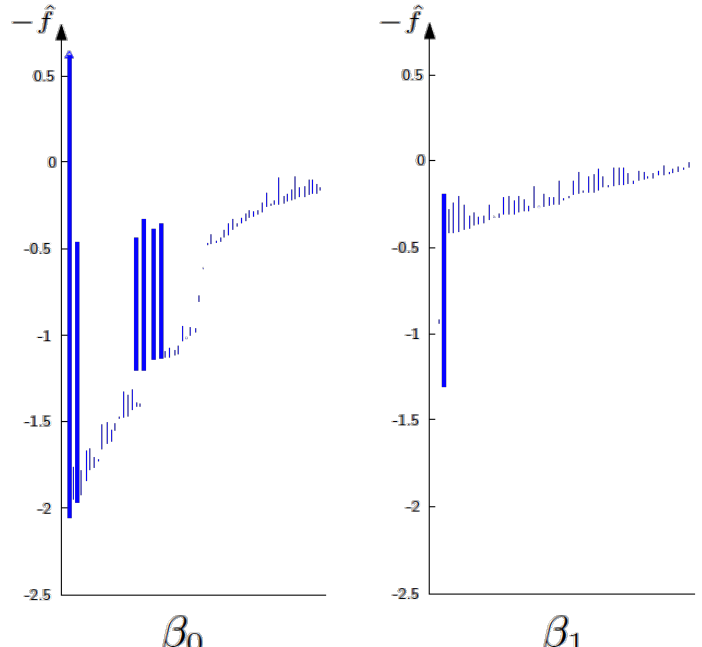

Functions on the Data as Filters

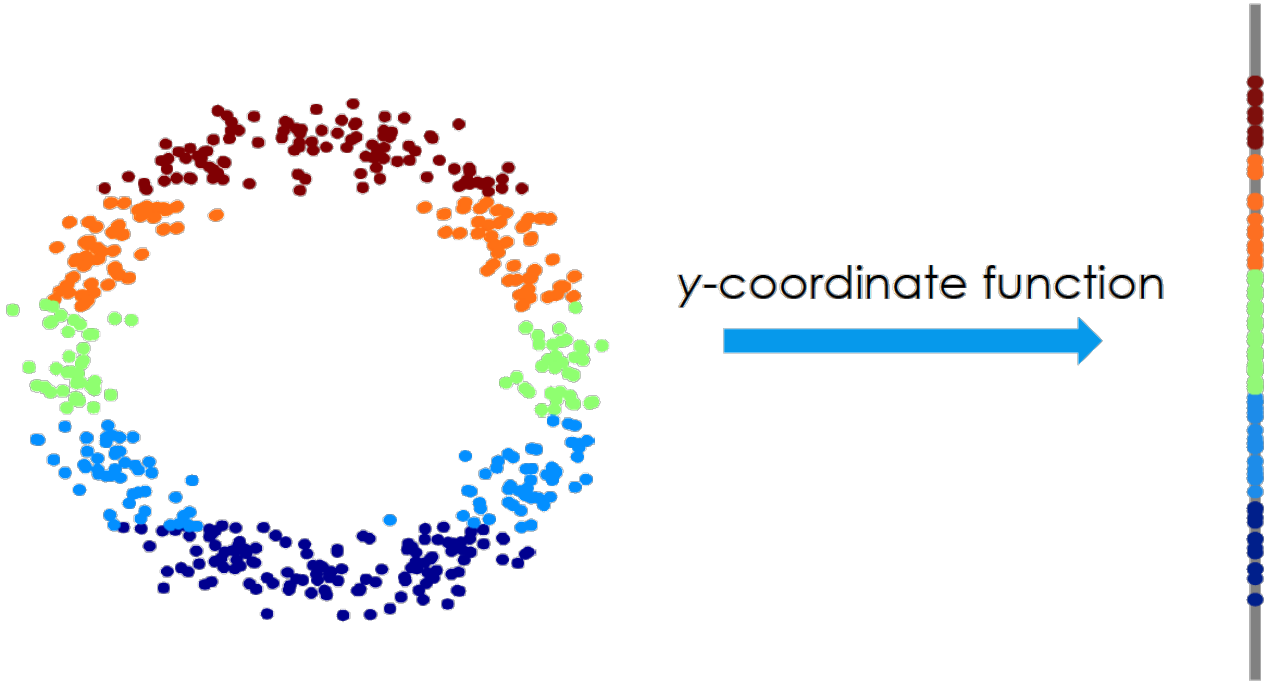

A Simplicial Complex

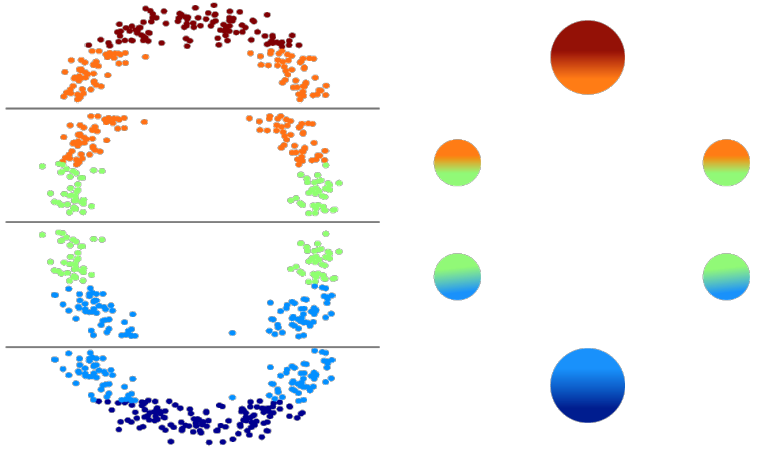

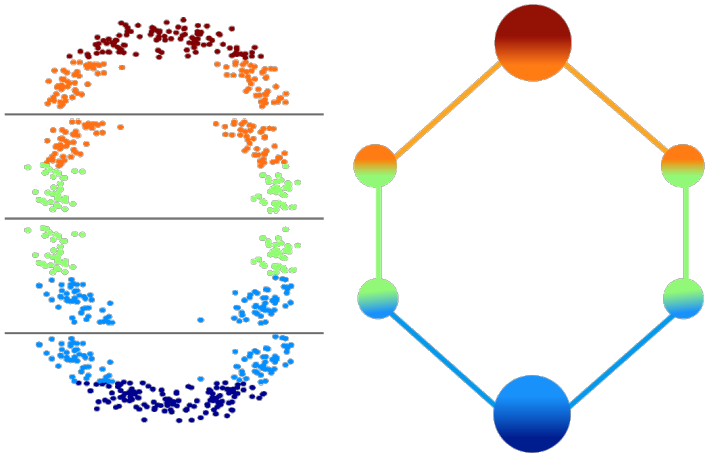

Functions as Lenses on Data

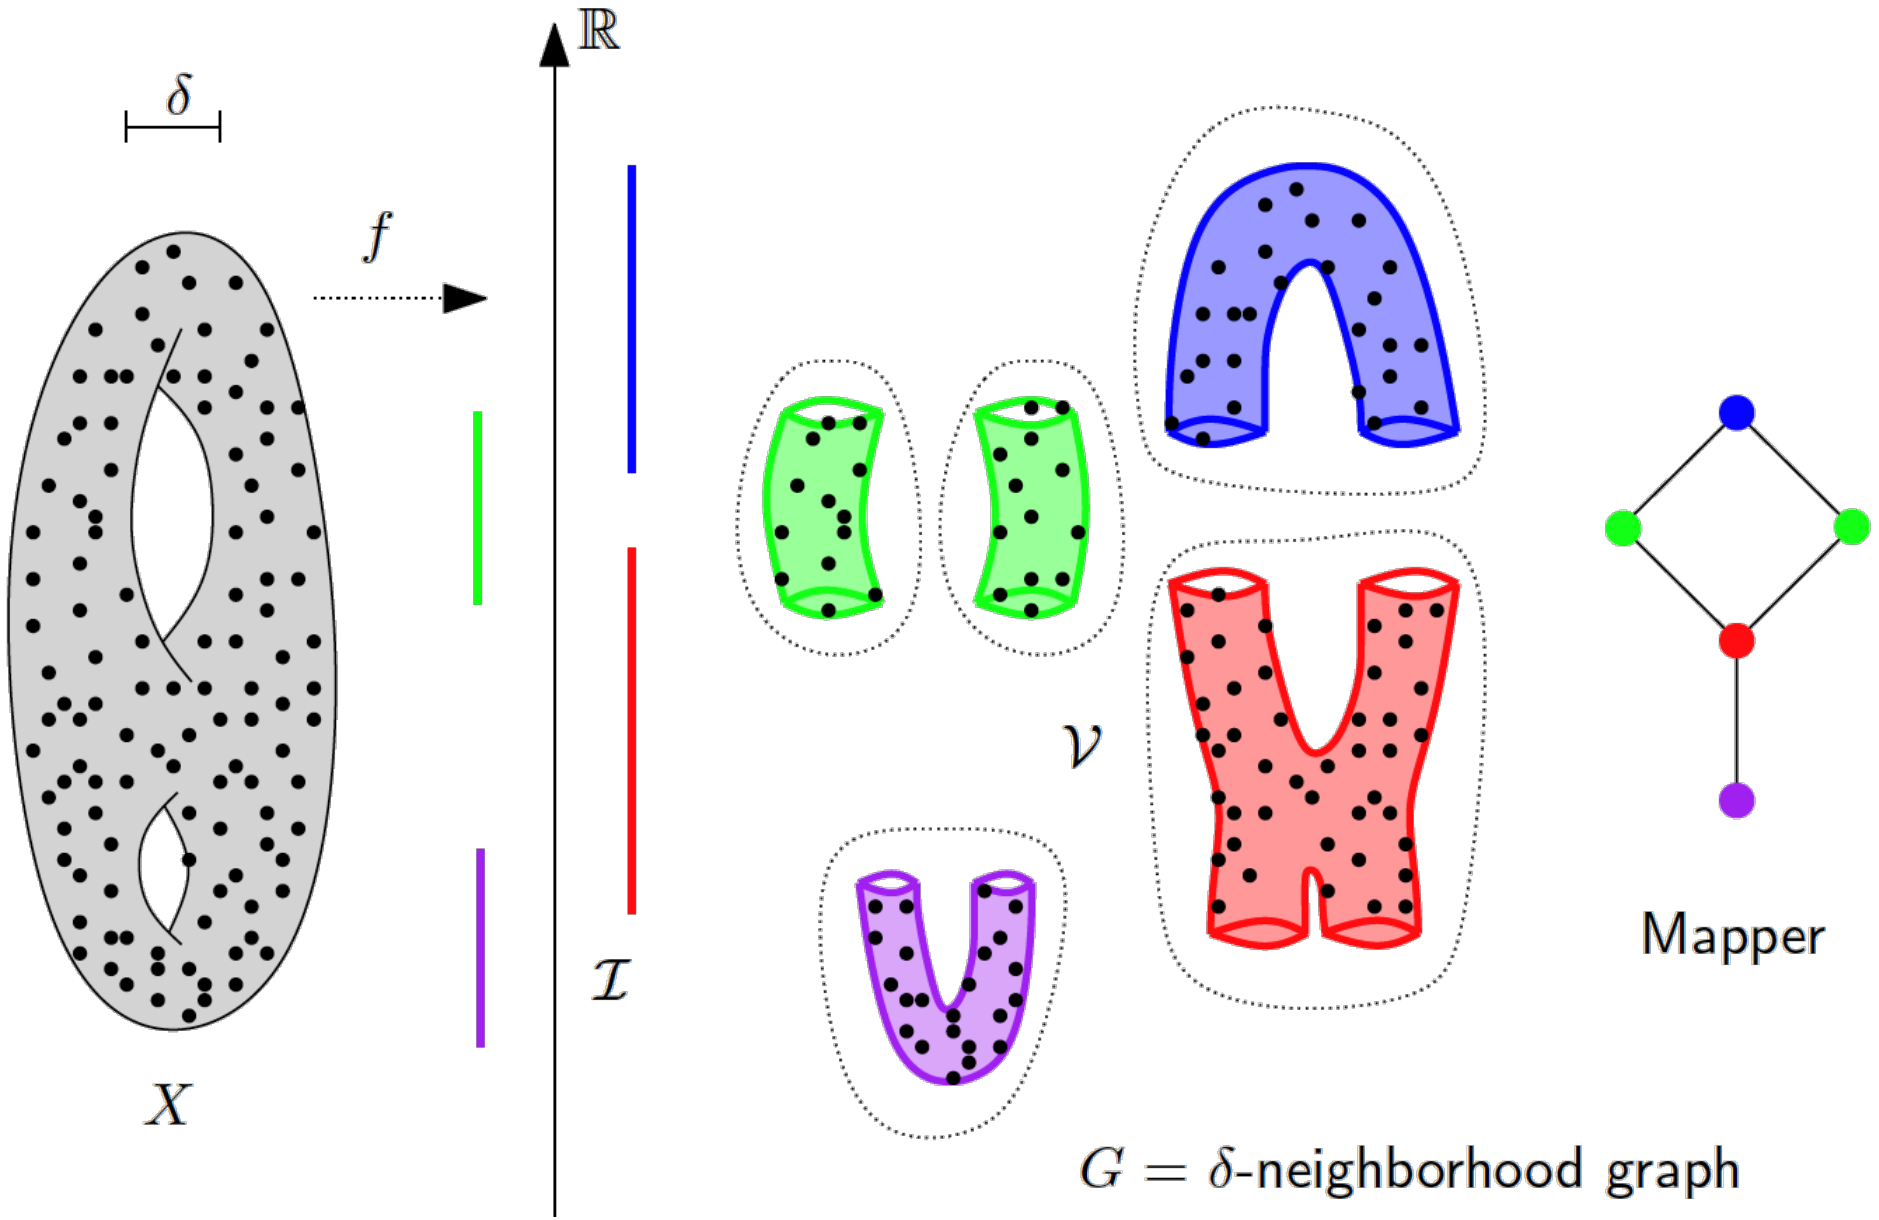

Data Representations and Latent Spaces

(Too) Many Latent Spaces

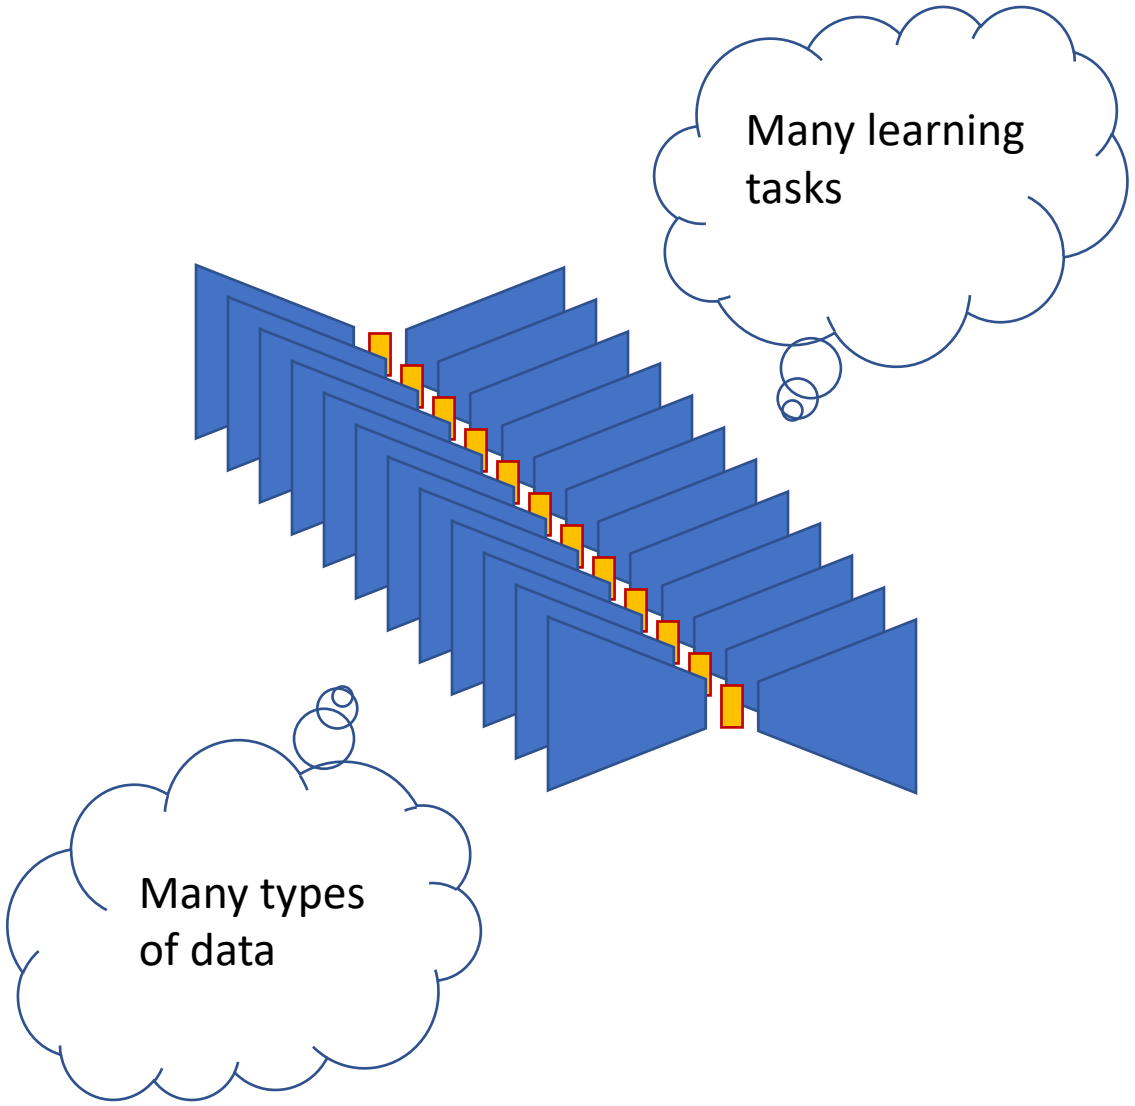

Codes

Latent vectors

Parametrizations

Representations

…

Correlated Data Sets

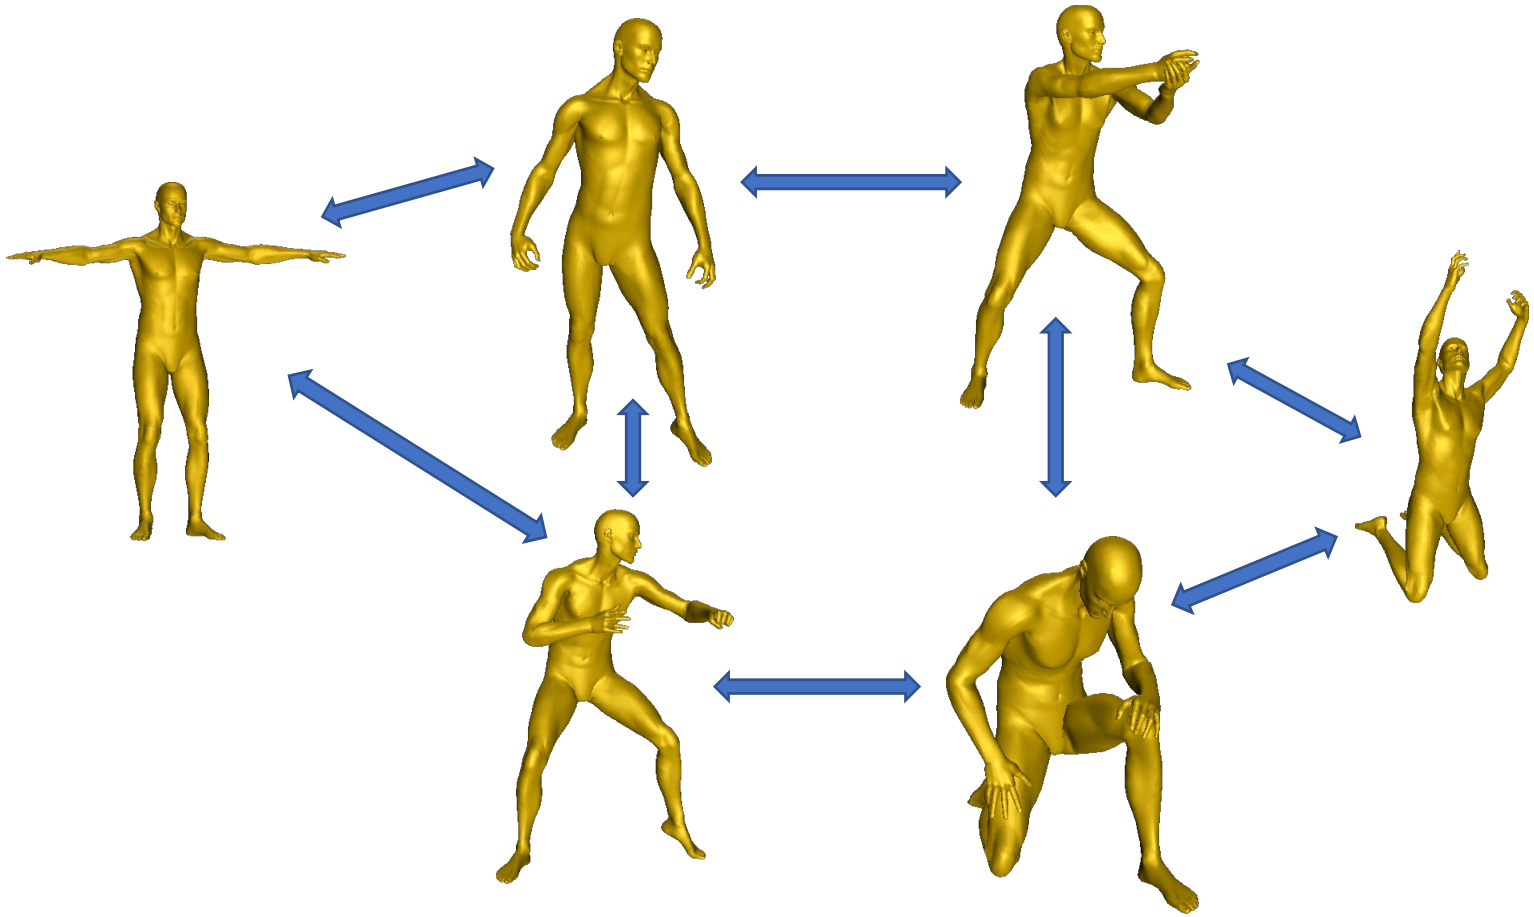

Correlated Tasks

The Structure of Latent Spaces

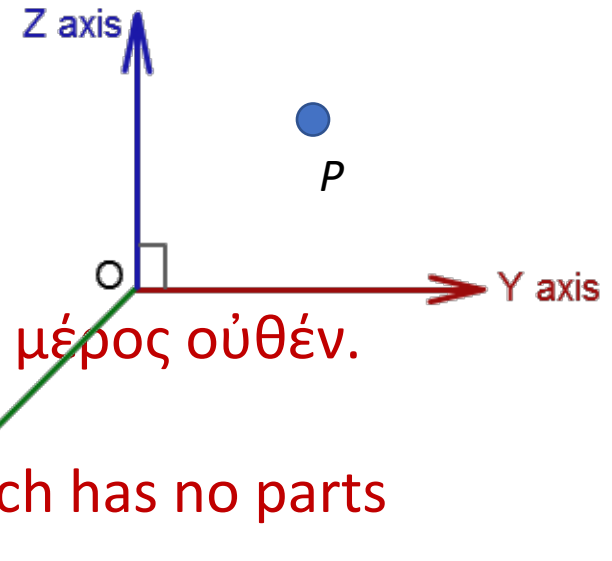

α'. Σημεῖόν ἐστιν, οὗ μέρος οὐθέν.

a. A point is that which has no parts

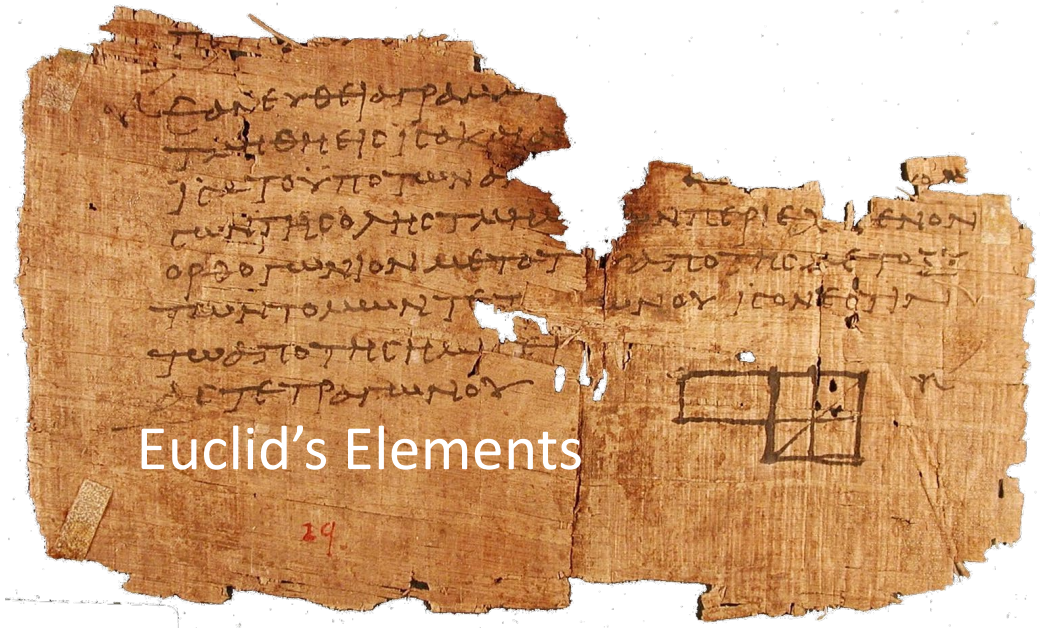

Homomorphic Encyption

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 $a \circ b$. To qualify as a group, the set and operation, $(G, \circ)$, must satisfy four requirements known as the group axioms:

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?

A Shared “Narrow Neck / Waist” in Other Fields

Power distribution

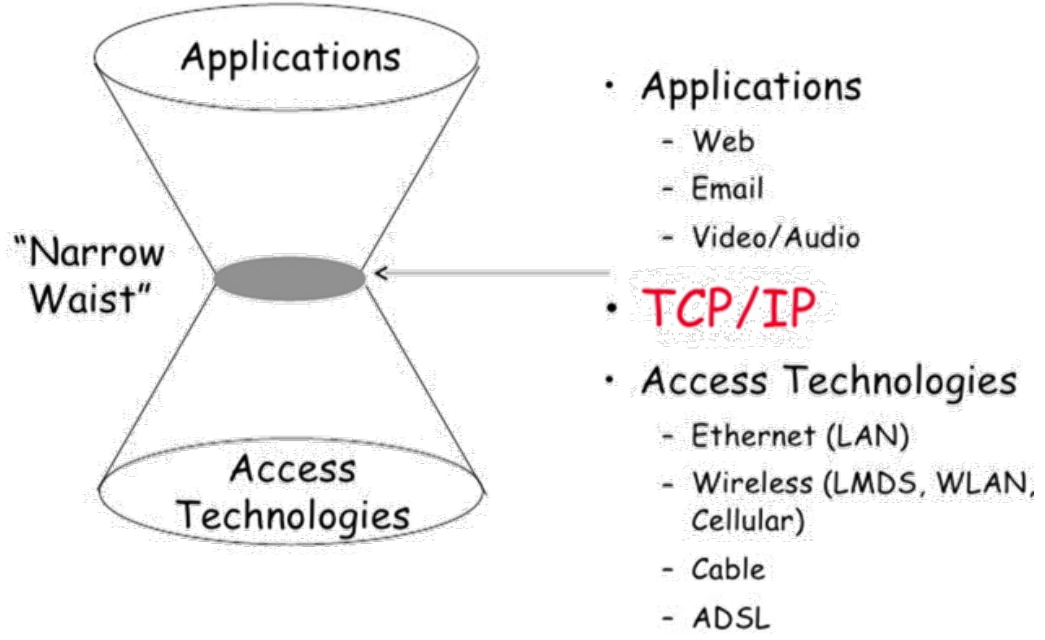

Network protocol stack [Image: R. Katz]

Relating Data Sets: Horizontal Map Networks

Horizontal Networks & Joint Learning

Similarity as a communications channel

34 *The Mathematical Theory of Communication*

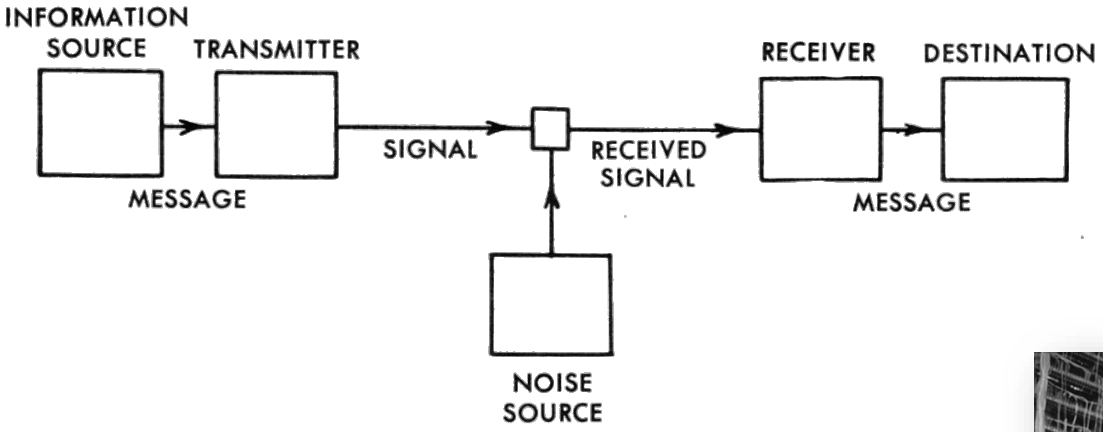

Fig. 1. — Schematic diagram of a general communication system.

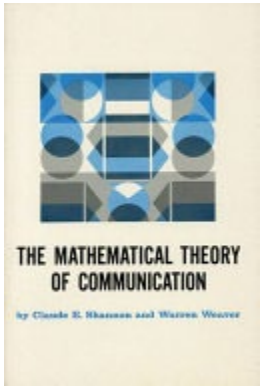

Claude Shannon

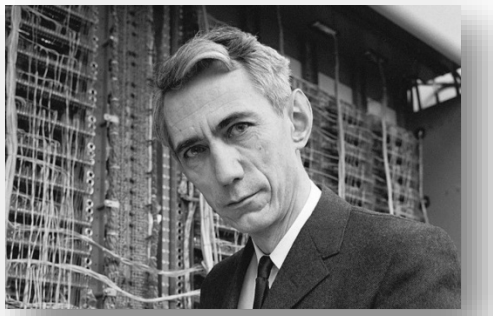

Networks of Visual Data

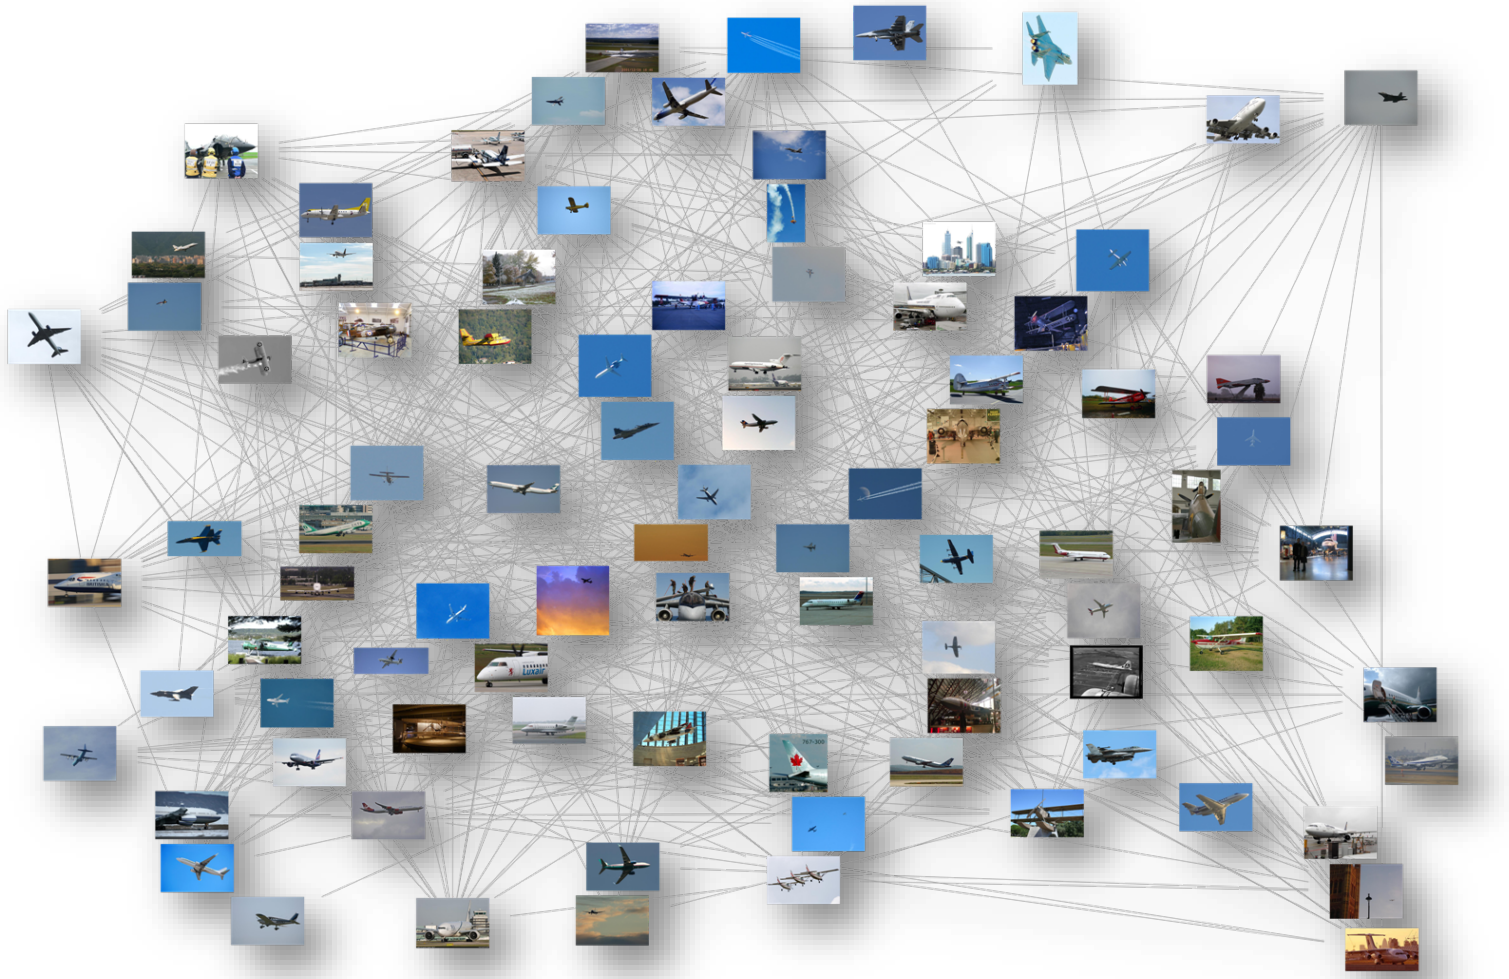

Alignment and Registration

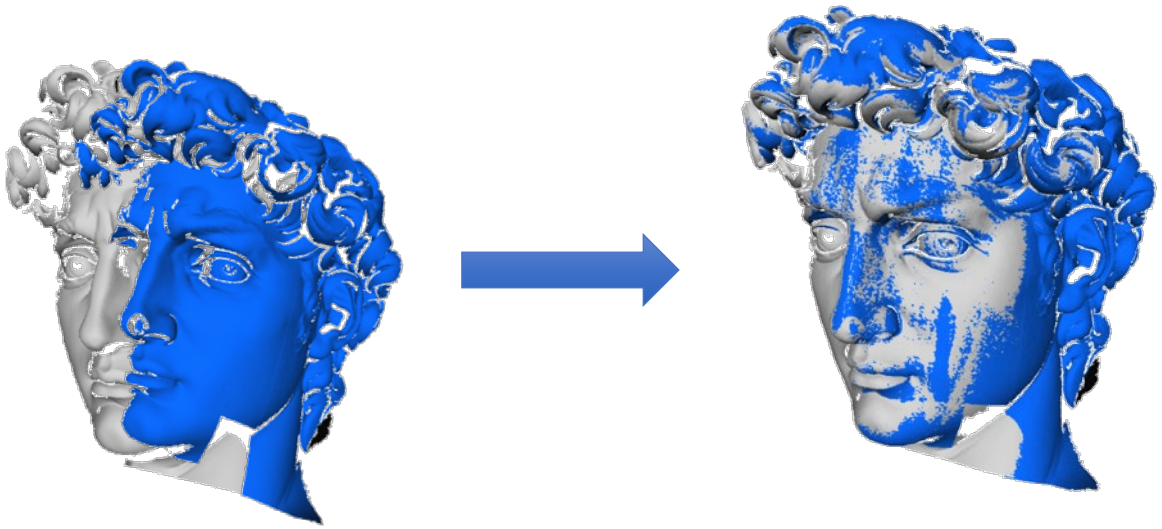

Rigid Registration

Low-dimensional transformation group

Maps and Correspondences

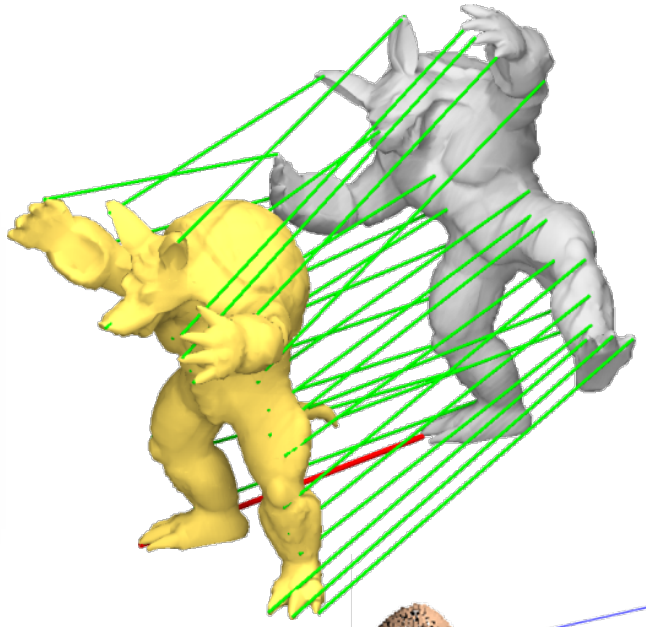

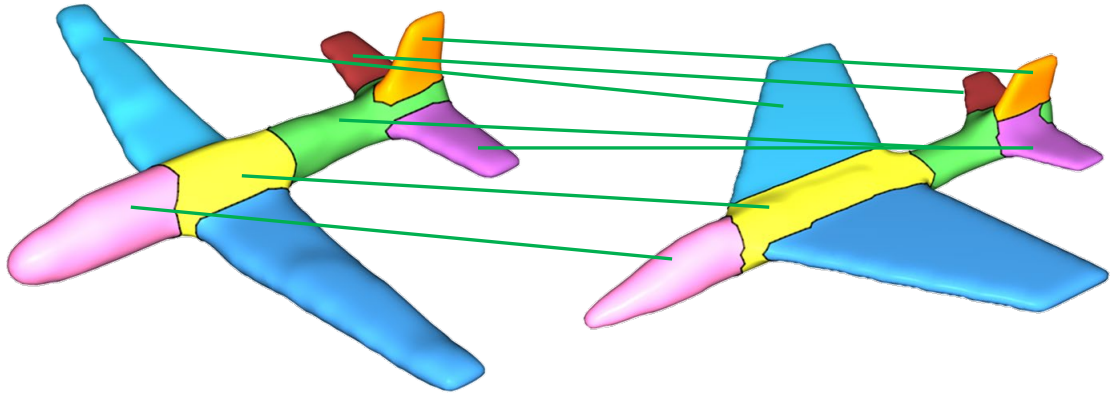

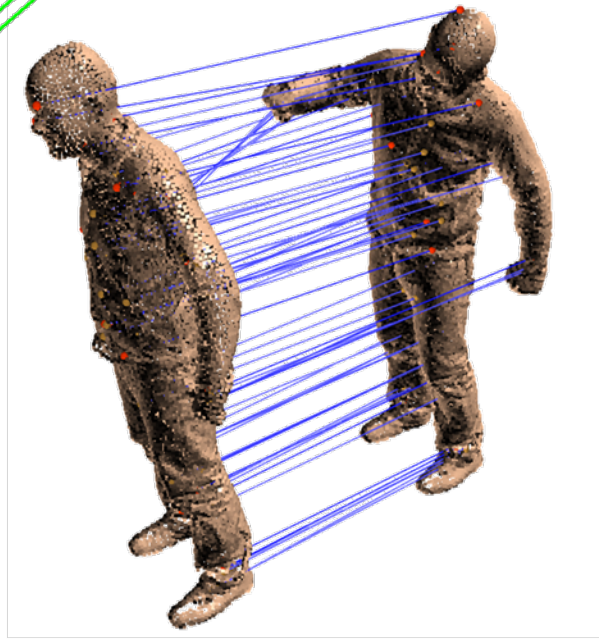

Transport Operators

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.

Functionals Over Data and Functional Maps

Knowledge as Functions over Data

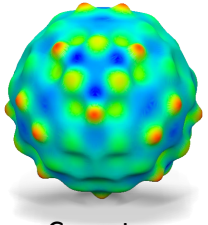

Curvature

Parts

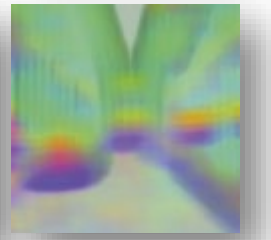

SIFT flow, C. Liu 2011

Knowledge towers over visual data: function spaces

Hierarchical Bases for a Function Space

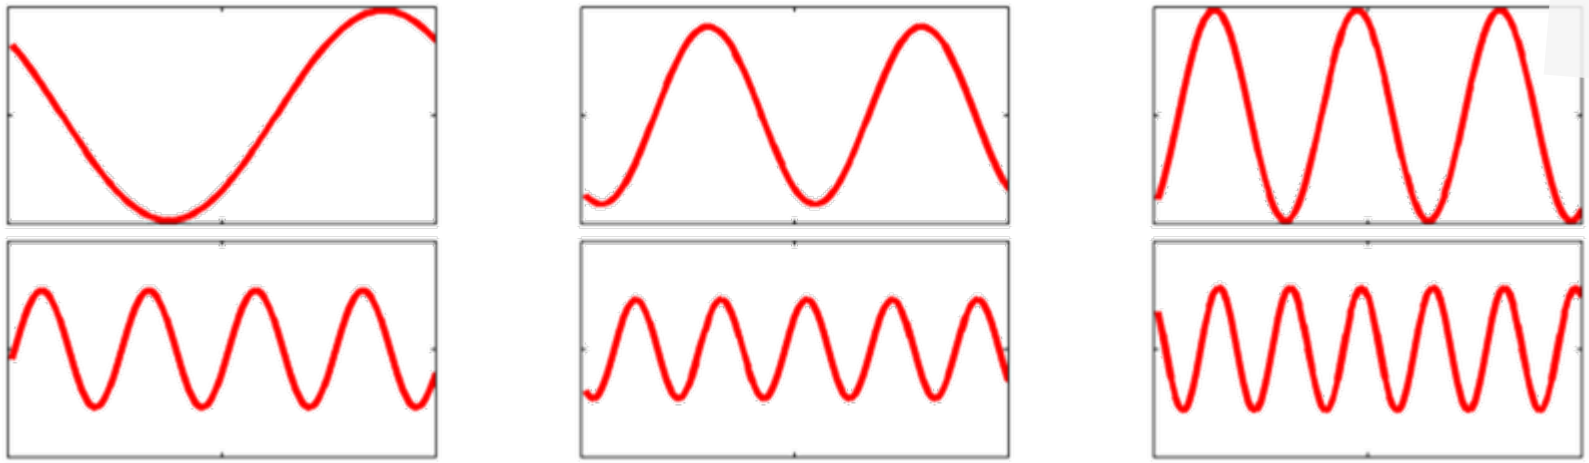

Fourier

Laplace-Beltrami

global support

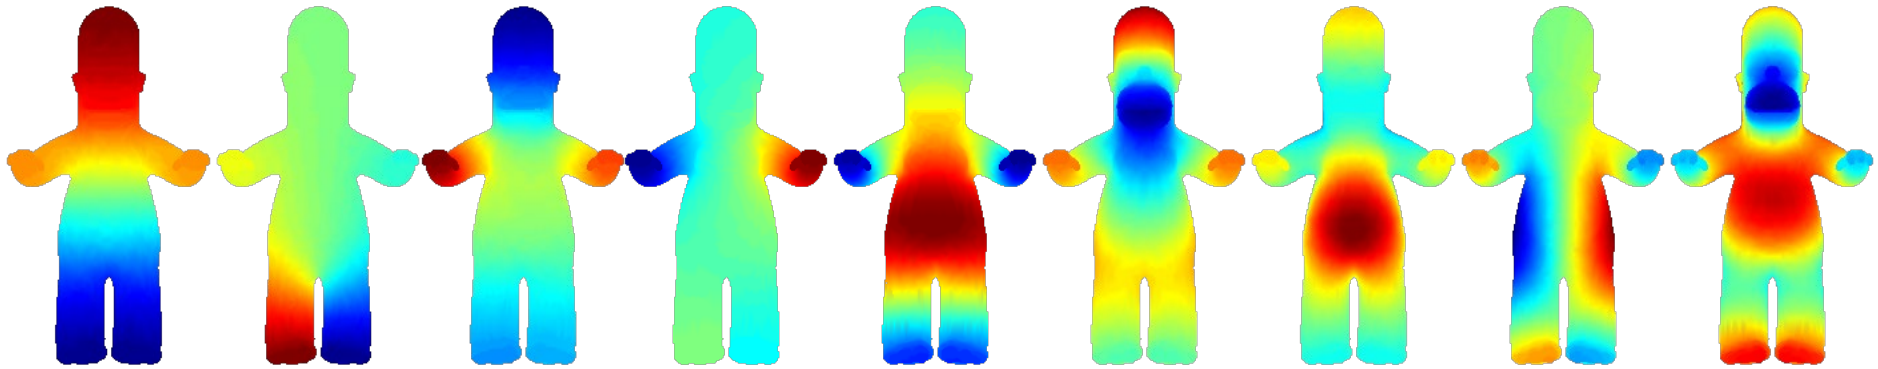

Functional Maps

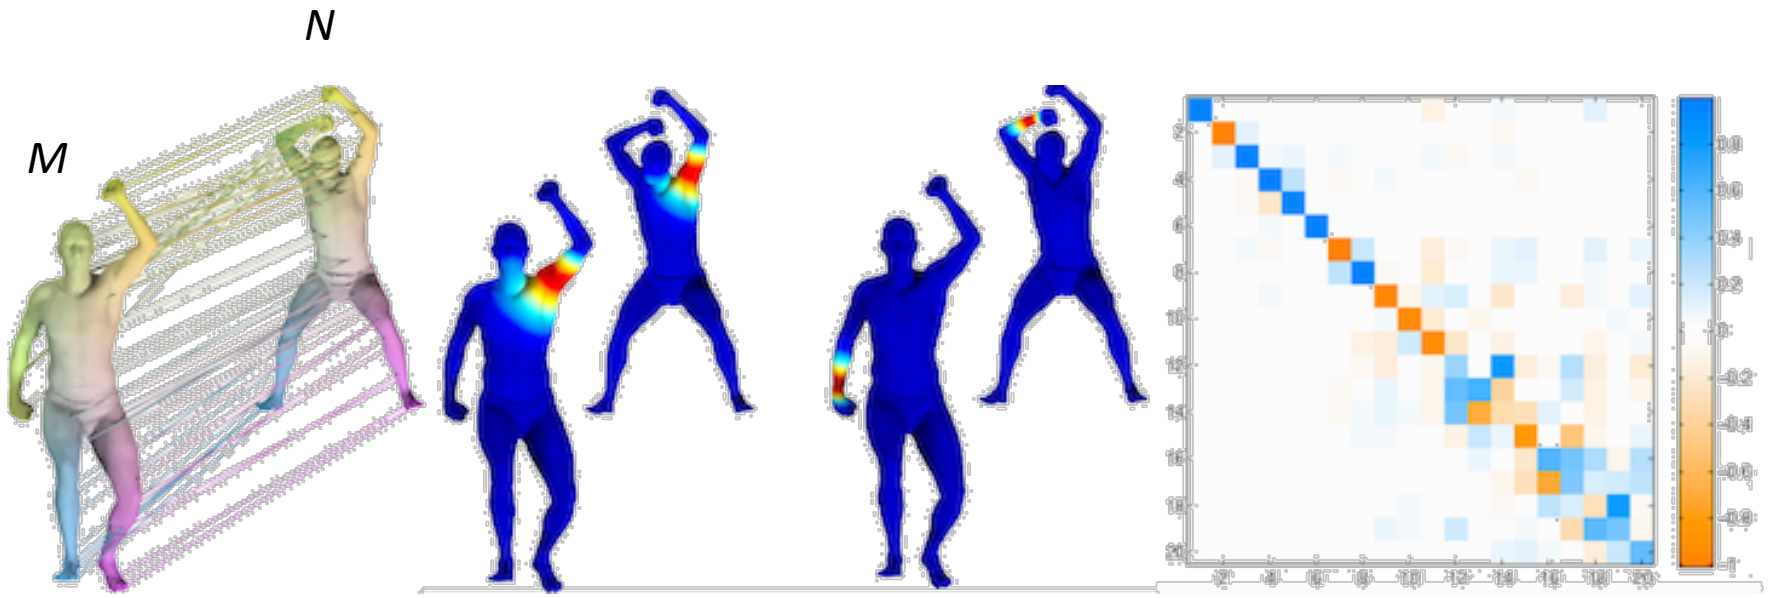

$$T : N \to M$$

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

A contravariant functor

Starting from a Regular Map ϕ

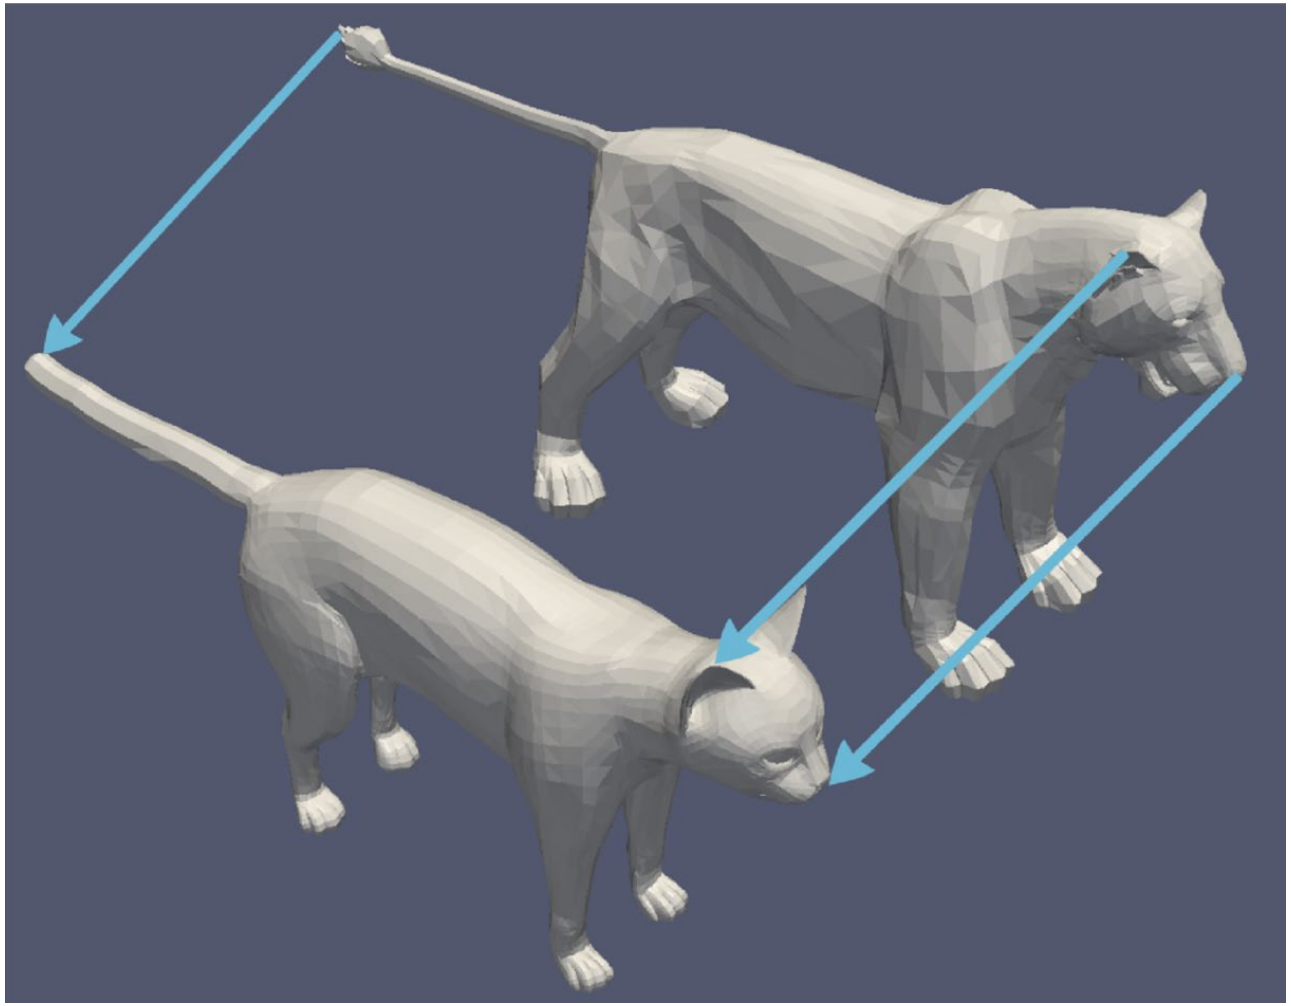

ϕ: lion → cat

A Contravariant Functor

from cat to lion

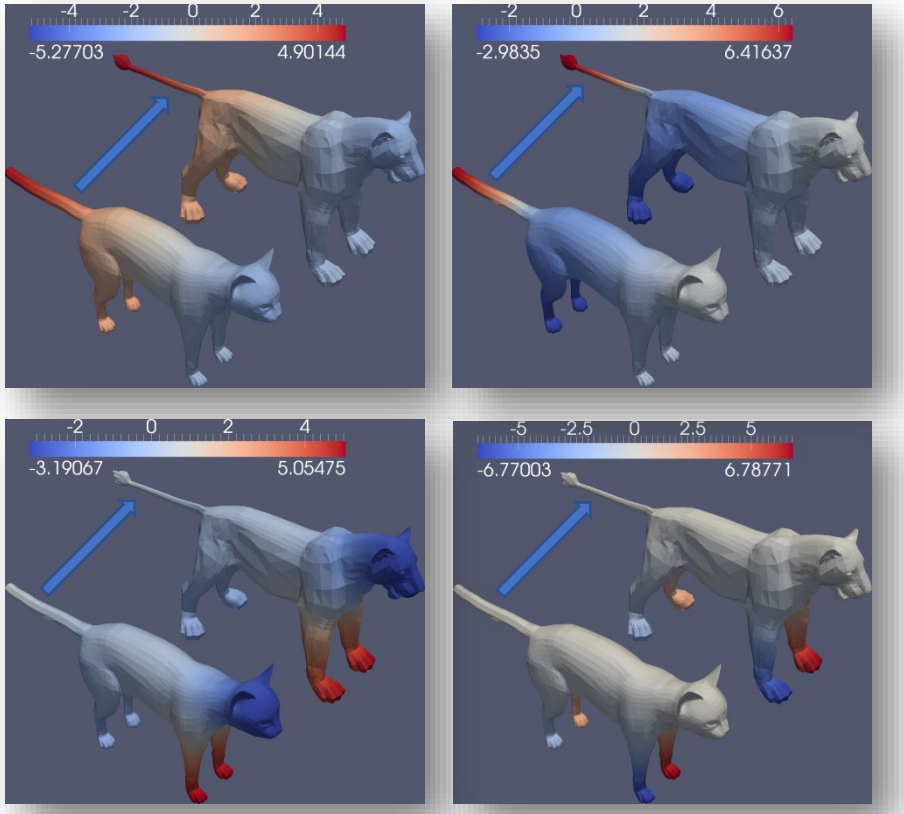

Functions on cat are transferred to lion using T_ϕ

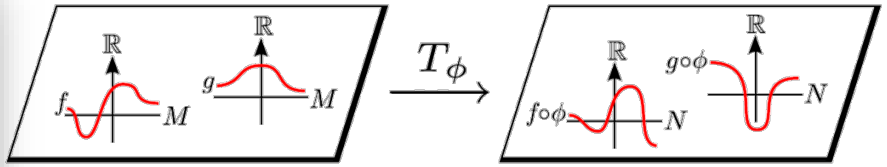

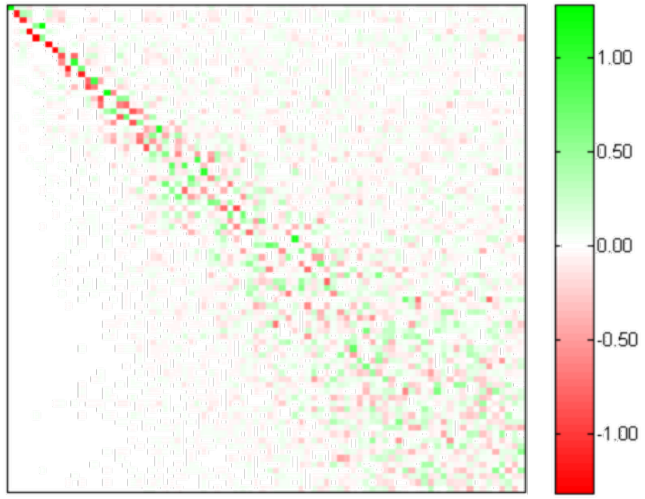

T_ϕ is a linear operator (matrix)

$$T_\phi : L^2(cat) \to L^2(lion)$$

Compact Encoding of Maps as Matrices

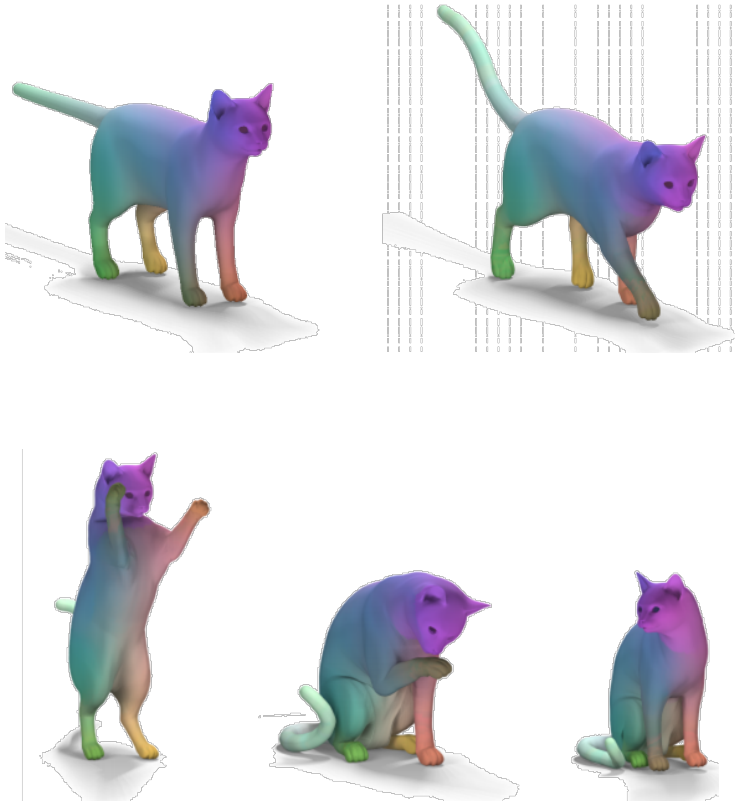

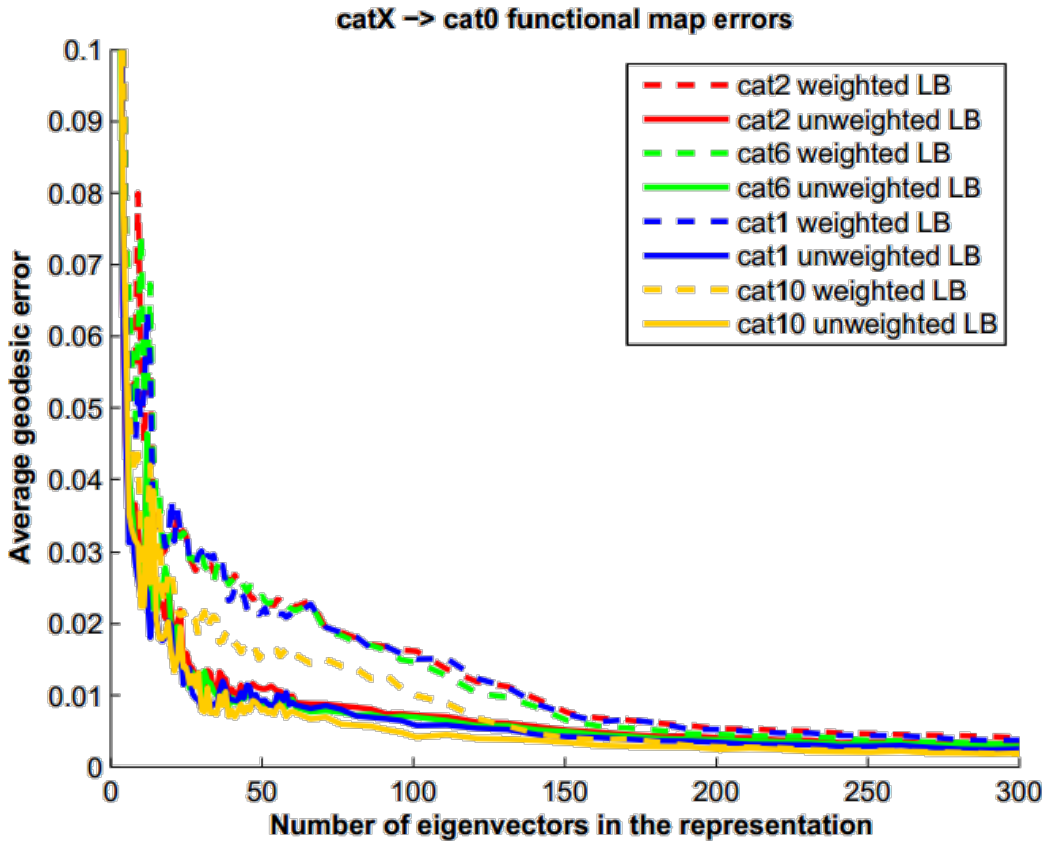

Changes to Measure and Metric under a Map

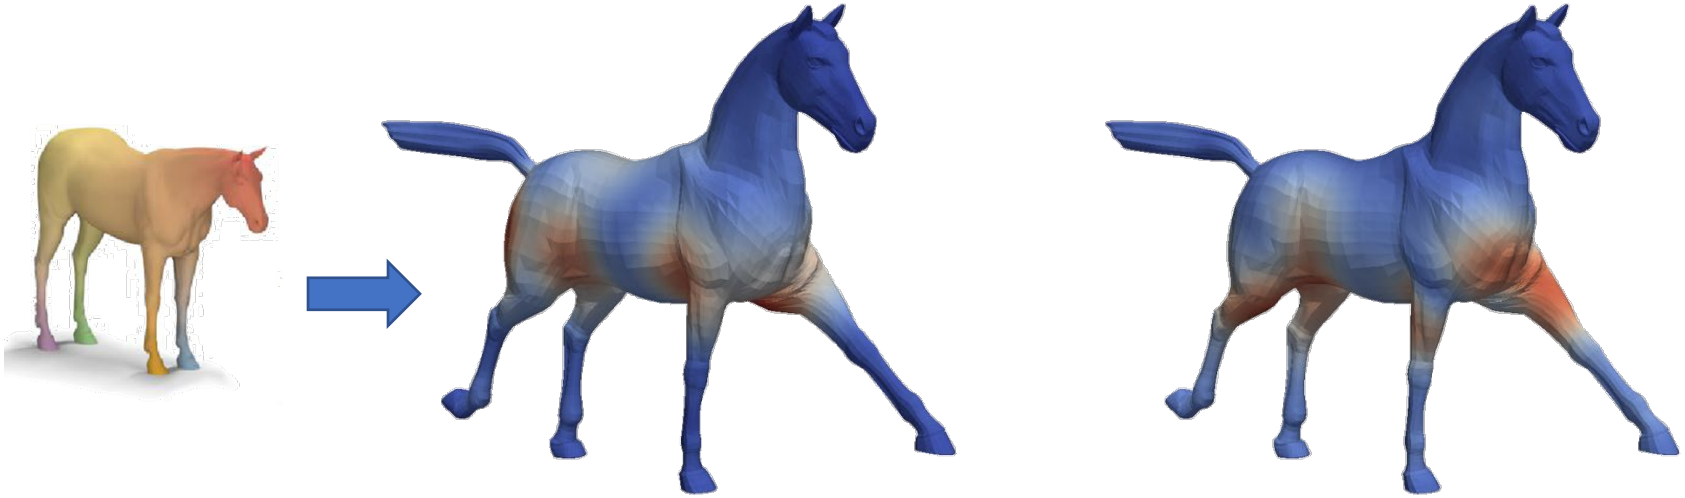

Area distortion

Conformal distortion

Shape Differences are a Change Recipe

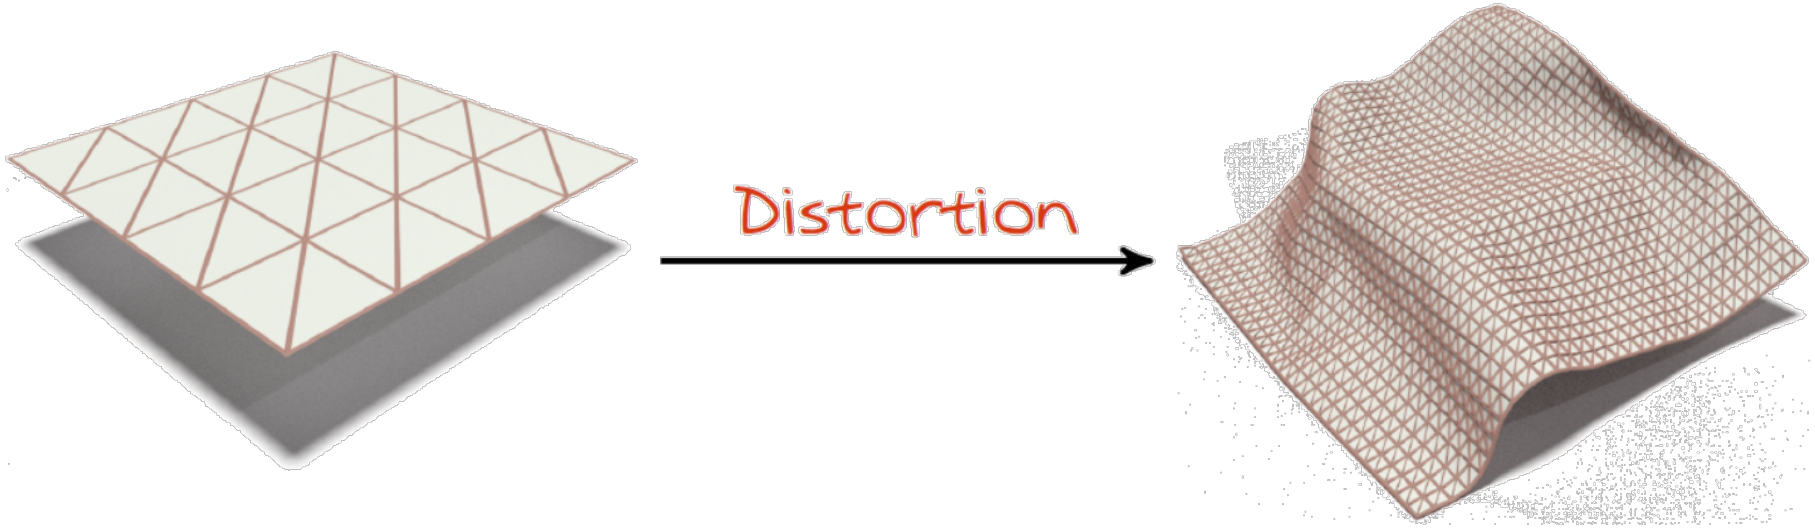

A recipe encoded as a matrix:

Area distortion
Conformal distortion

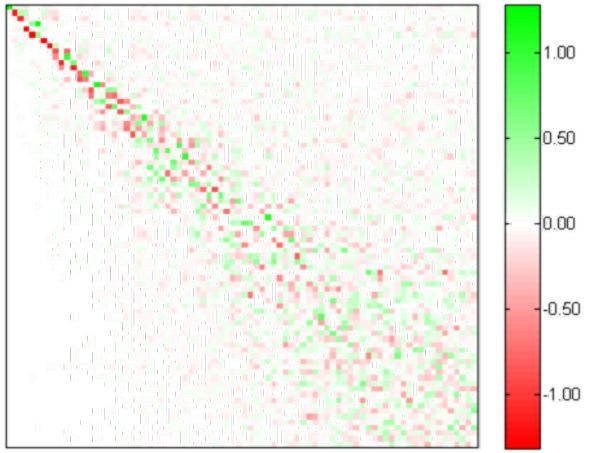

A novel type of latent space representation for 3D data

Under some conditions, lossless

The Space of Shapes

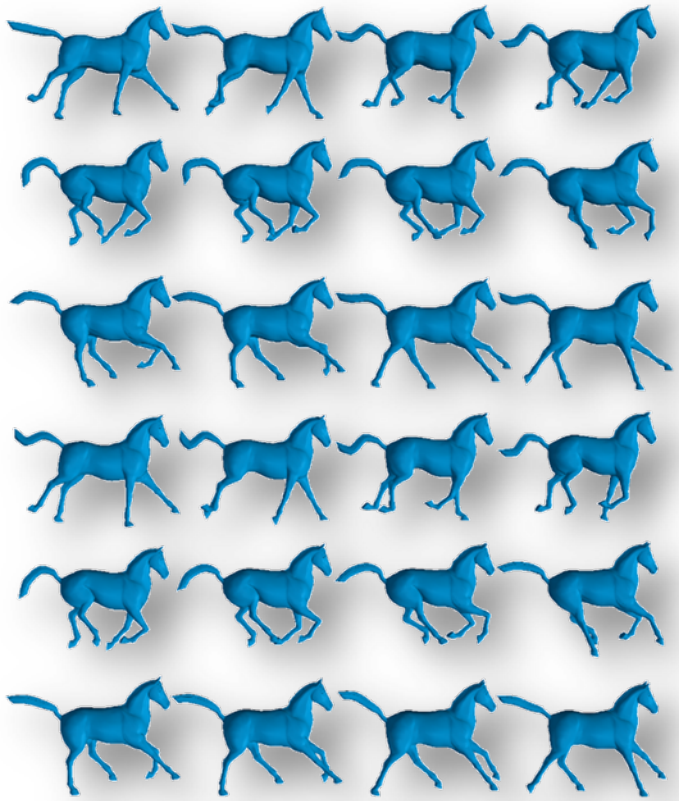

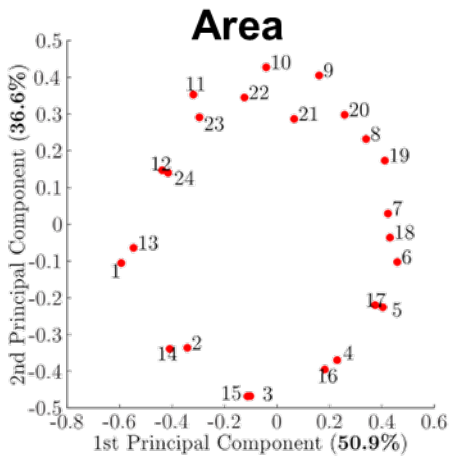

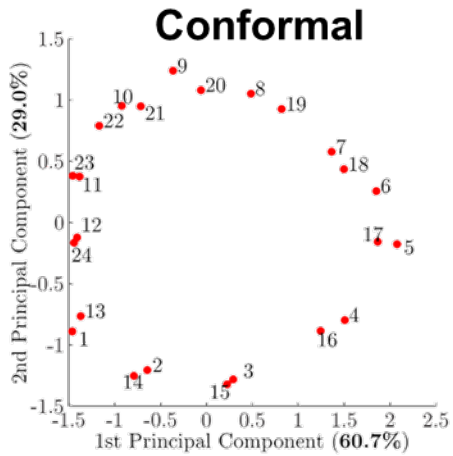

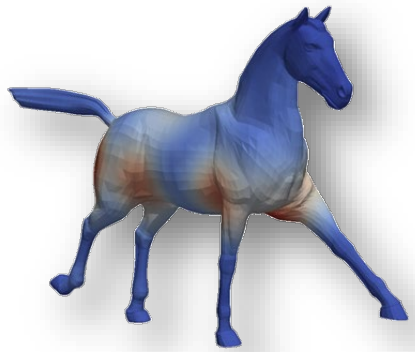

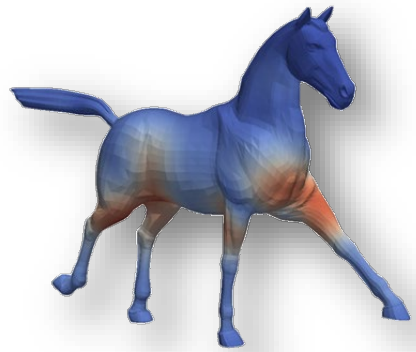

Horizontal Map Networks

Networks Between Data

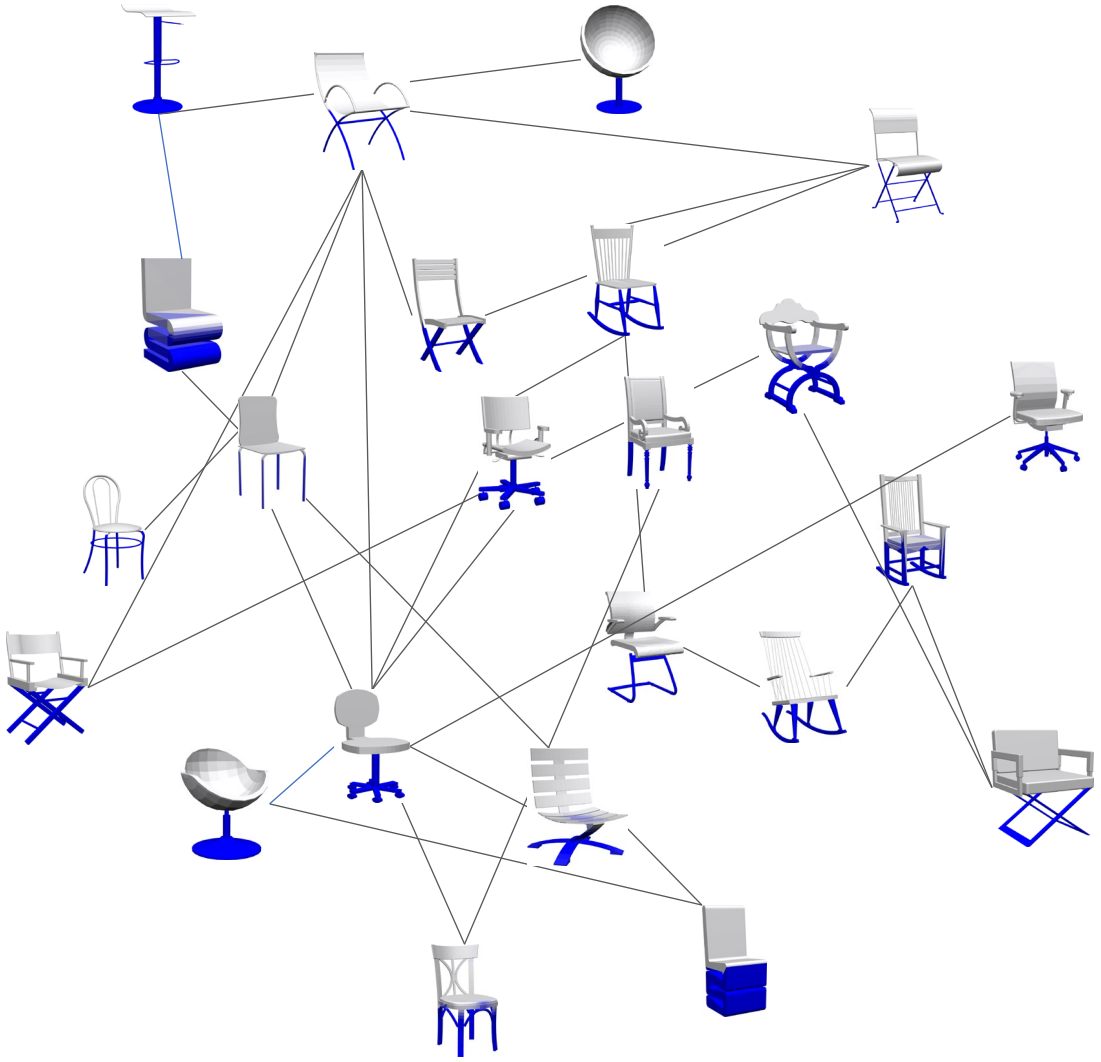

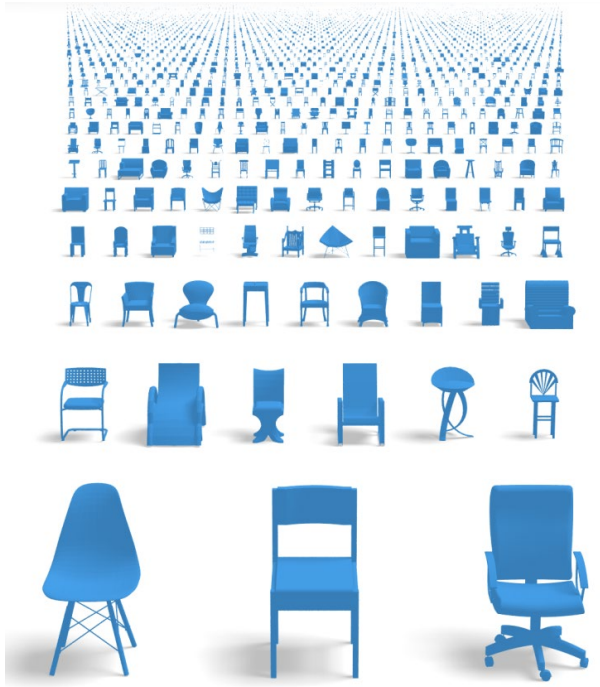

Joint Analysis: Co-Segmentation

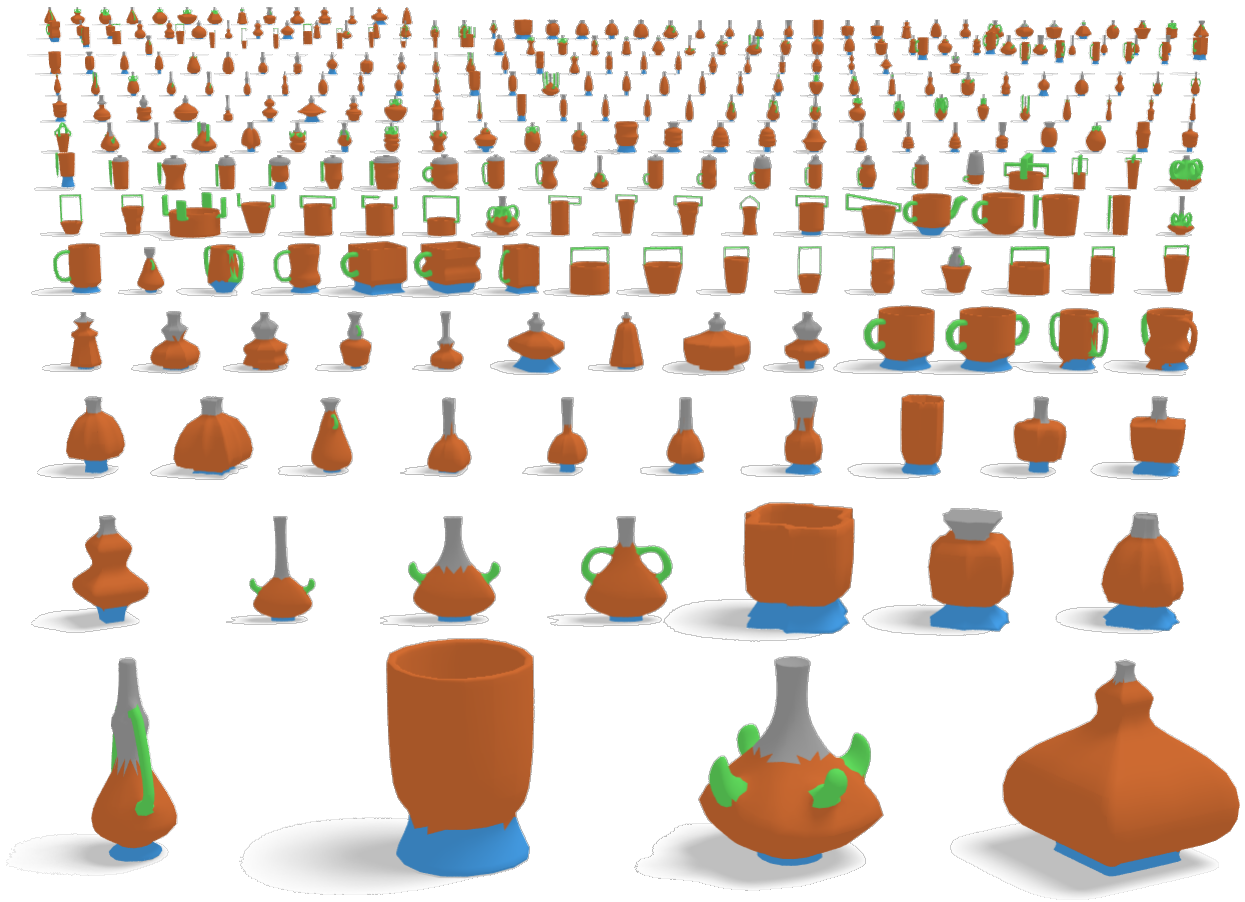

Unsupervised Structure Extraction

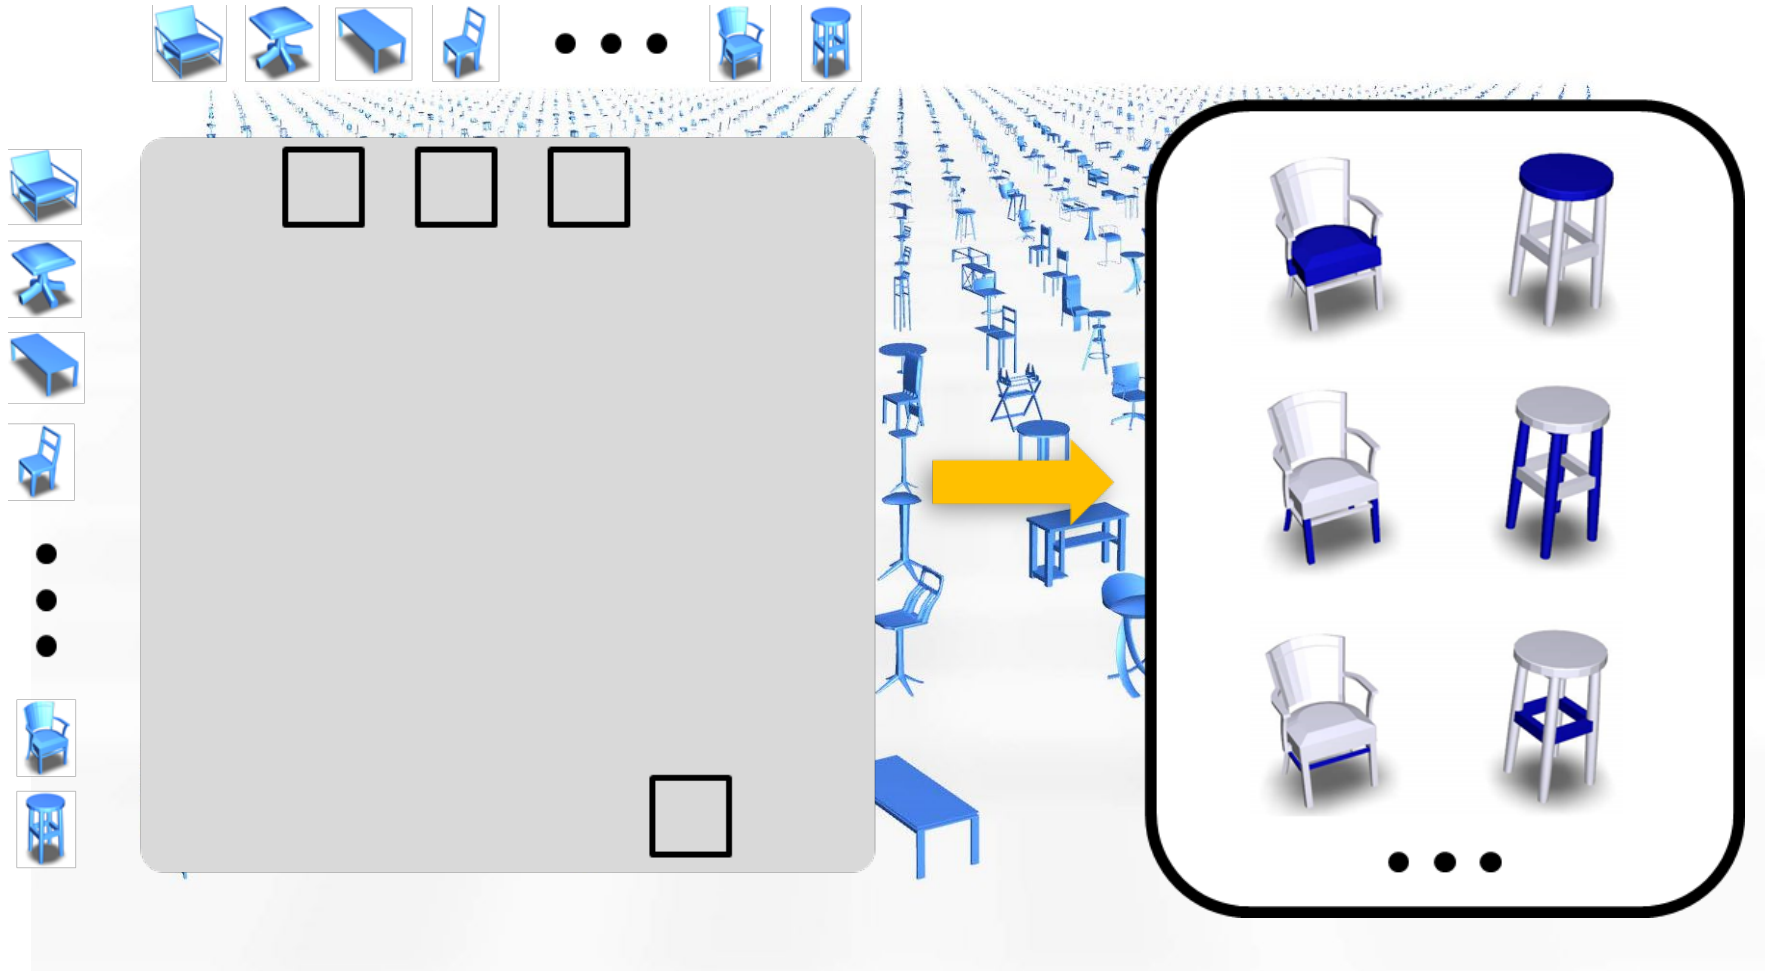

[Q. Huang, F. Wang, L. Guibas, ’14]

MSRC: 5 images per class are shown

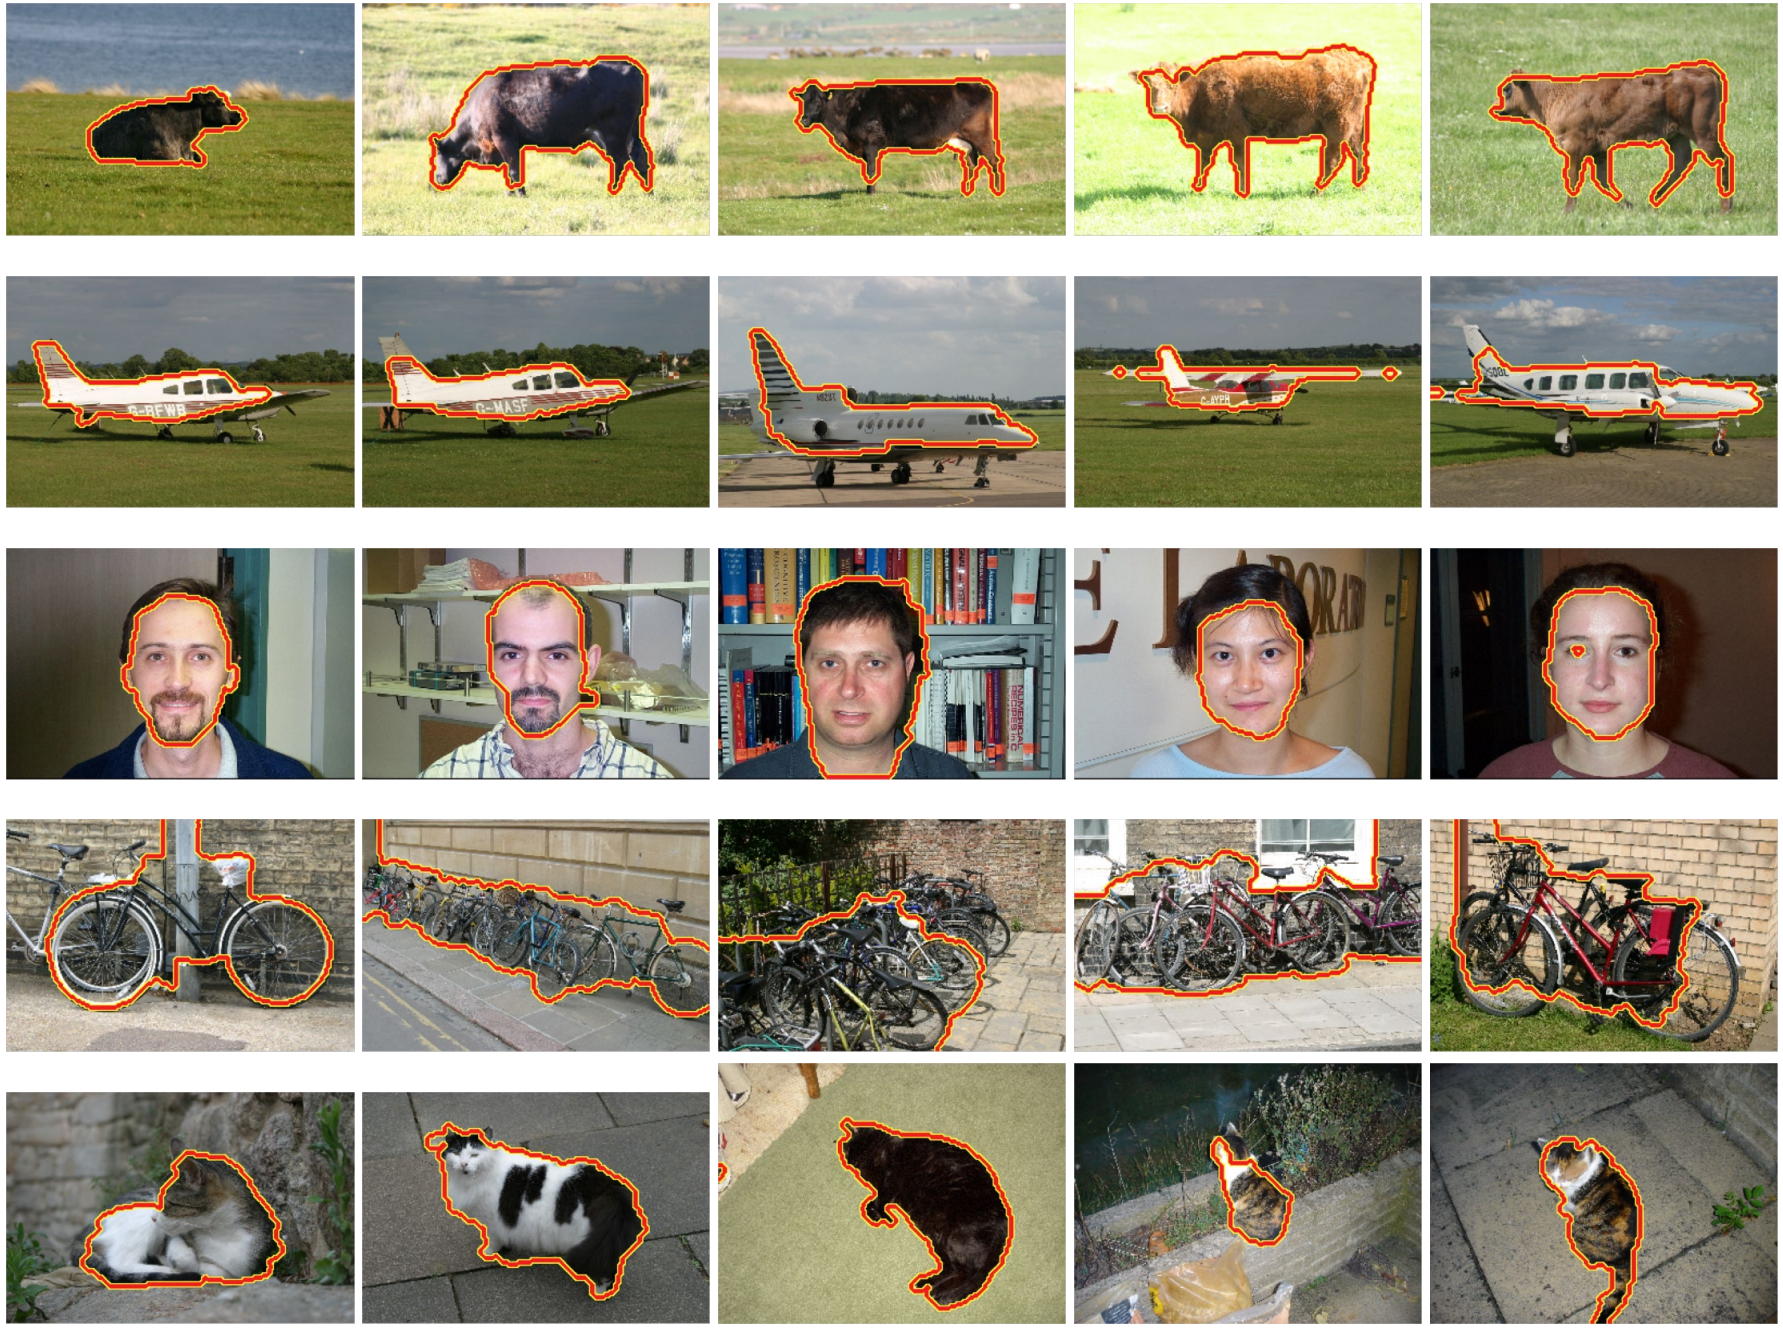

The Network is the Abstraction

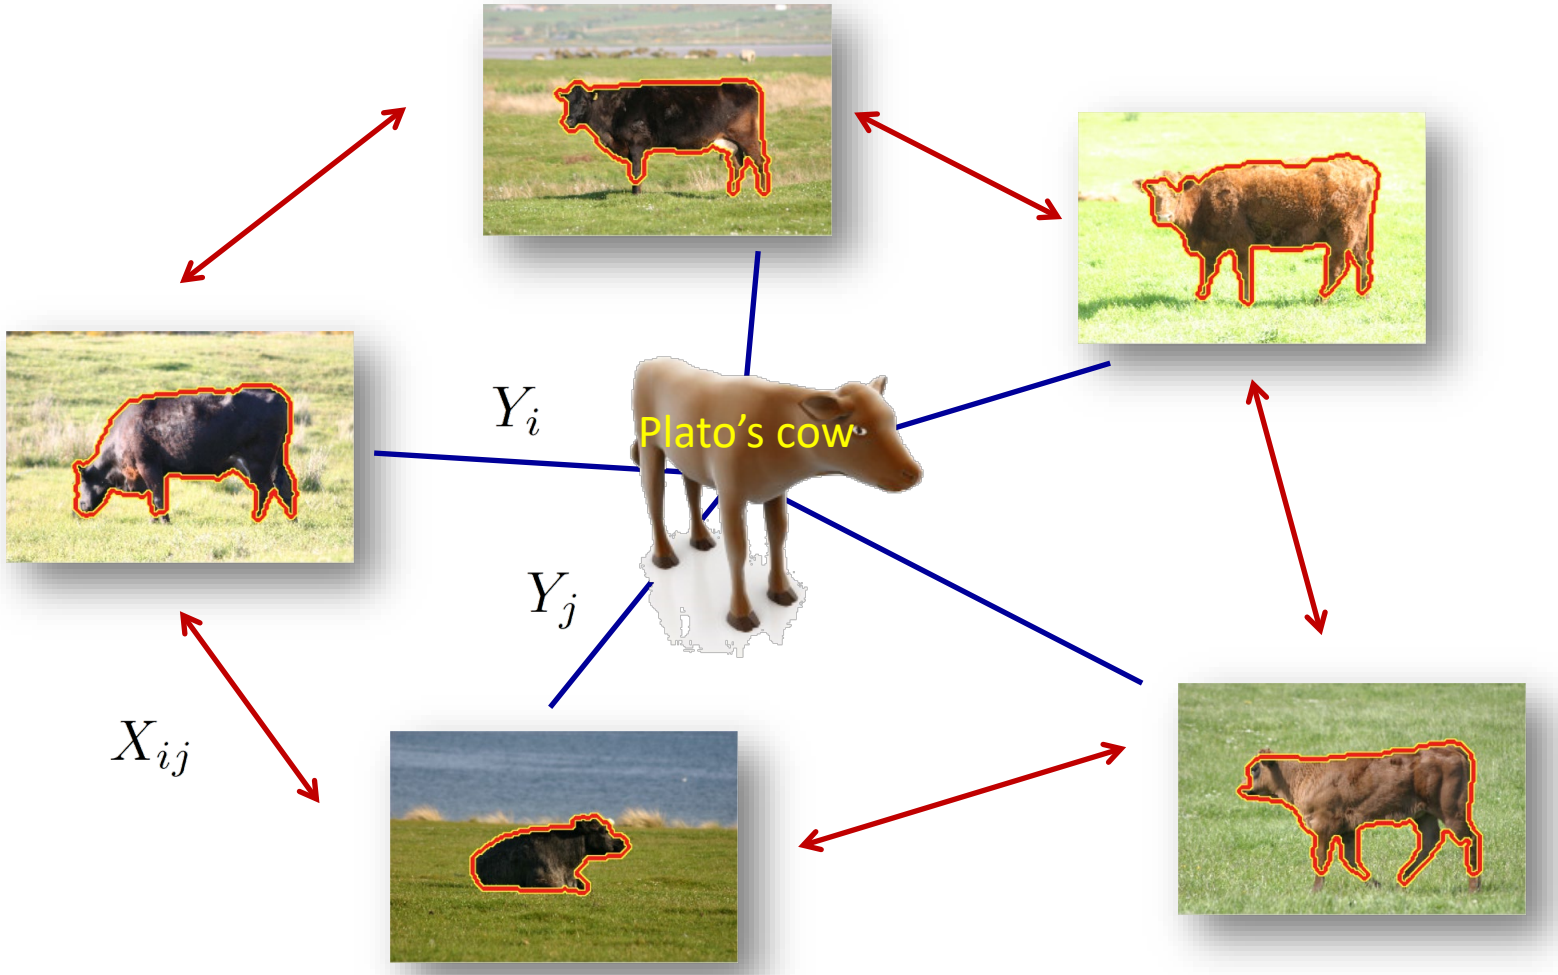

Mosaicking or SLAM at the Level of Functions

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

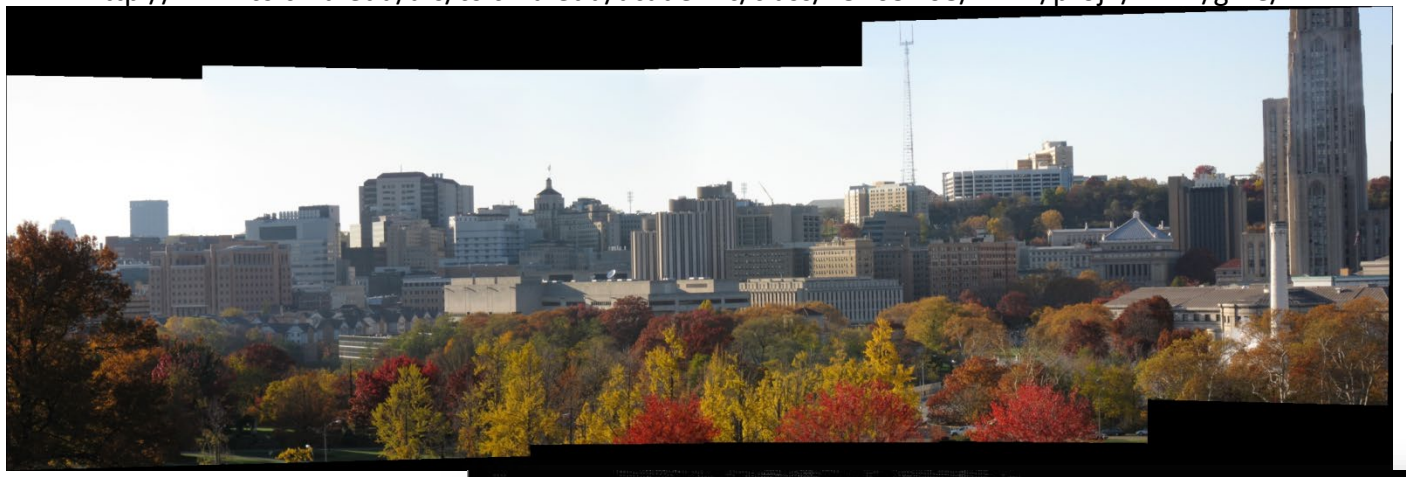

robotics.ait.kyushu-u.ac.jp

Joint Learning and Information Aggregation

- Need to aggregate information:
  - *over different data sets*
  - *over different modalities (geometry, appearance, language)*
  - *over space and time*
  - *over different representations*
  - *over different predictions*
  - *over different tasks*

in settings where the above refer to same entities in the world – and are thus correlated

Visual Data Repositories for Storing Semantic Knowledge

Semantic Networks

- Also known as **frame networks**
- Encode semantic relations between concepts
- Often used as a form of knowledge representation
- A directed or undirected graph consisting of vertices, which represent concepts, and edges which represent concept relations

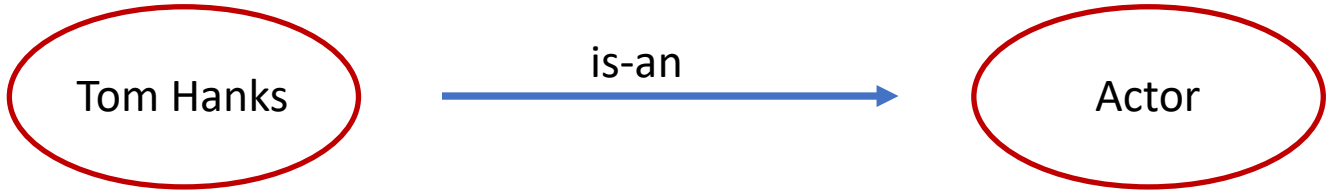

Example of a Semantic Net

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

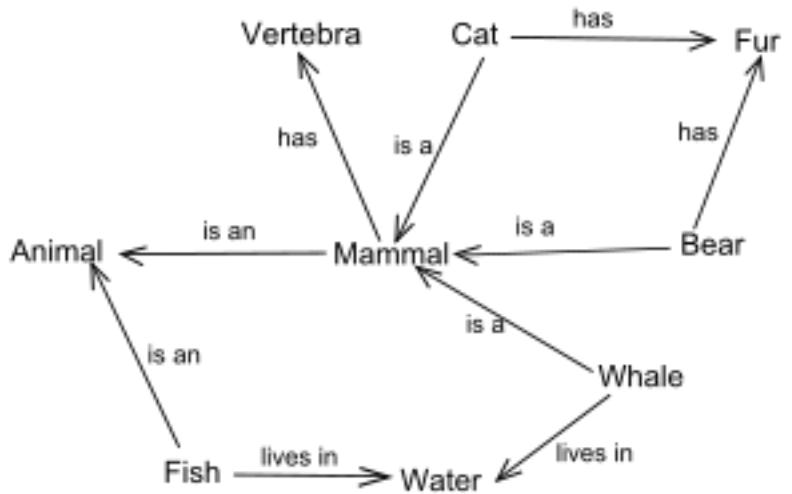

Google Knowledge Graph

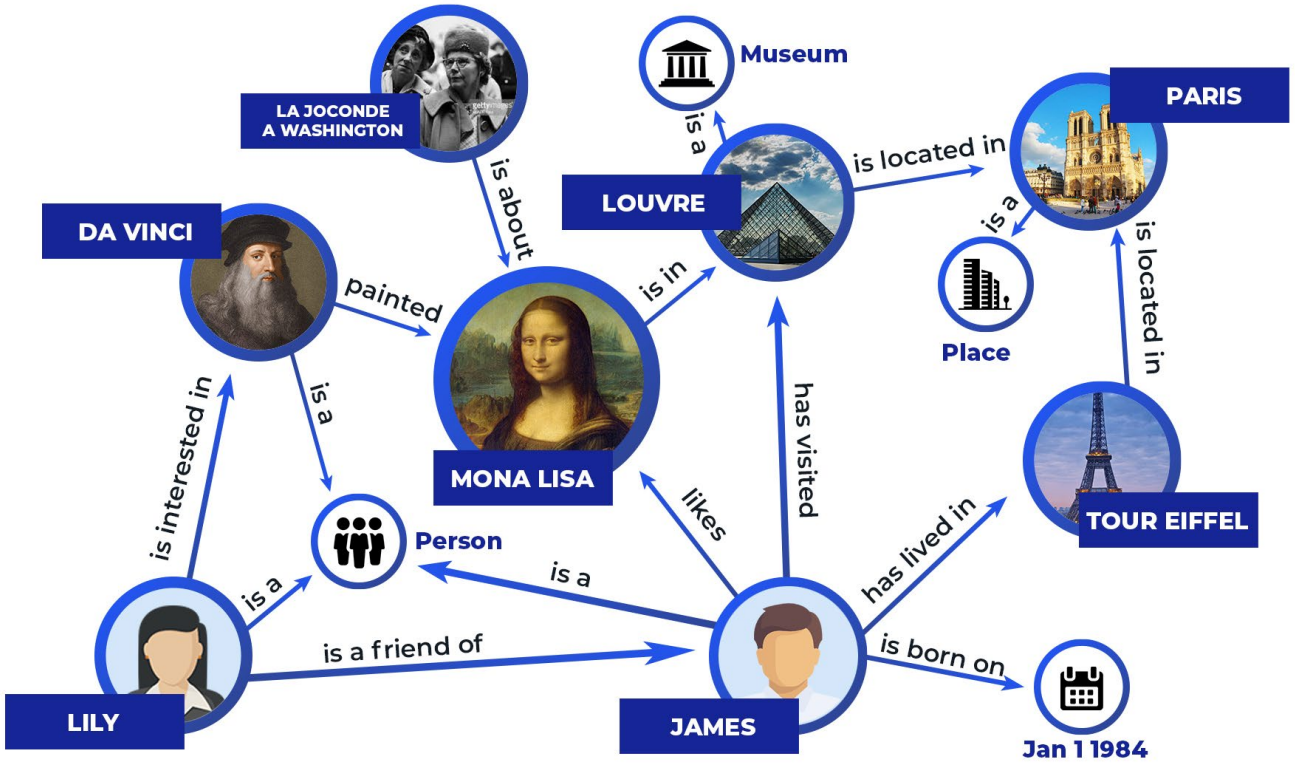

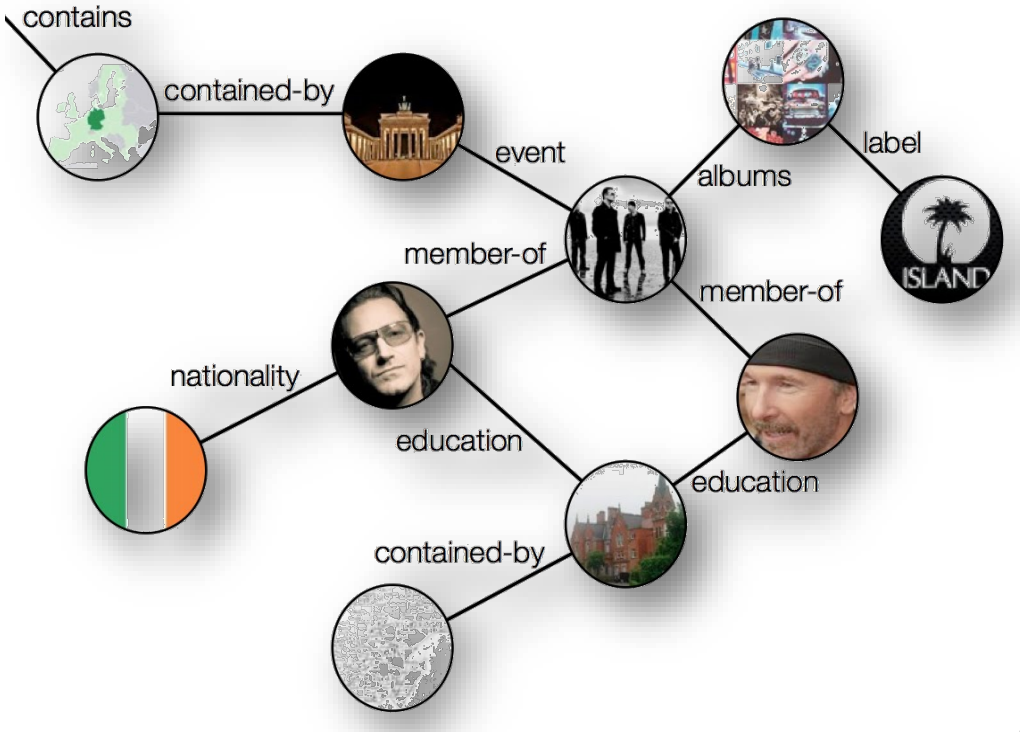

WordNet

What is WordNet?

Original paper by **[George Miller, et al 1990]** cited over 5,000 times

Organizes over 150,000 words into 117,000 categories called *synsets*.

Establishes ontological and lexical relationships in NLP and related tasks.

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.

WordNet

- Important relations between synsets (nouns):

Relation	Definition	Example
Hypernym	From concepts to superordinates	water^1 →liquid
Hyponym	From concepts to subtypes	water^1 →seawater
Has-Part	From groups to their members	water^1 → oxygen
Part-of	From members to their groups	water^1 →ice
Antonym	Opposites	leader→follower

Taxonomy: is-a Relationship

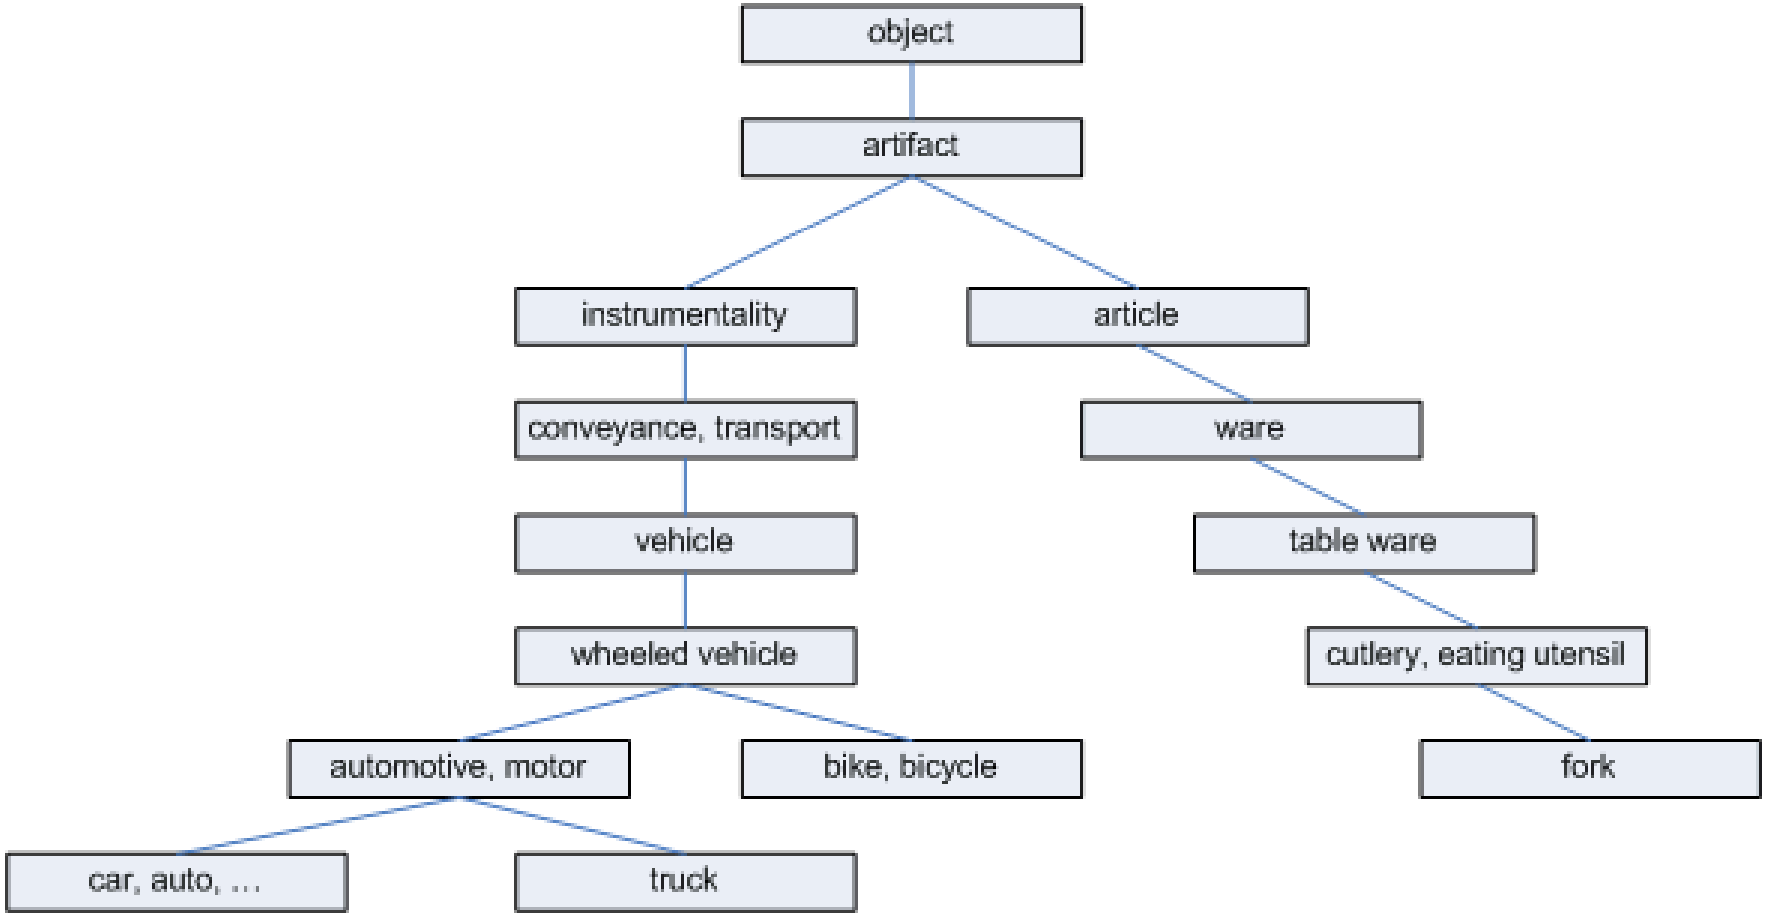

Partonomy: has-a Relationship

- S: (n) **car**, auto, automobile, machine, motorcar (a motor vehicle with four wheels; usually propelled by an internal combustion engine) *"he needs a car to get to work"*
  - direct hyponym / full hyponym
  - **part meronym**
    - S: (n) accelerator, accelerator pedal, gas pedal, gas, throttle, gun (a pedal that controls the throttle valve) *"he stepped on the gas"*
    - S: (n) air bag (a safety restraint in an automobile; the bag inflates on collision and prevents the driver or passenger from being thrown forward)
    - S: (n) auto accessory (an accessory for an automobile)
    - S: (n) automobile engine (the engine that propels an automobile)
    - S: (n) automobile horn, car horn, motor horn, horn, hooter (a device on an automobile for making a warning noise)
    - S: (n) buffer, fender (a cushion-like device that reduces shock due to an impact)
    - S: (n) bumper (a mechanical device consisting of bars at either end of a vehicle to absorb shock and prevent serious damage)
    - S: (n) car door (the door of a car)
    - S: (n) car mirror (a mirror that the driver of a car can use)
    - S: (n) car seat (a seat in a car)
    - S: (n) car window (a window in a car)
    - S: (n) fender, wing (a barrier that surrounds the wheels of a vehicle to block splashing water or mud) *"in Britain they call a fender a wing"*
    - S: (n) first gear, first, low gear, low (the lowest forward gear ratio in the gear box of a motor vehicle; used to start a car moving)
    - S: (n) floorboard (the floor of an automobile)
    - S: (n) gasoline engine, petrol engine (an internal-combustion engine that burns gasoline; most automobiles are driven by gasoline engines)
    - S: (n) glove compartment (compartment on the dashboard of a car)
    - S: (n) grille, radiator grille (grating that admits cooling air to car's radiator)
    - S: (n) high gear, high (a forward gear with a gear ratio that gives the greatest vehicle velocity for a given engine speed)
    - S: (n) hood, bonnet, cowl, cowling (protective covering consisting of a metal part that covers the engine) *"there are powerful engines under the hoods of new cowlings in order to repair the plane's engine"*
    - S: (n) luggage compartment, automobile trunk, trunk (compartment in an automobile that carries luggage or shopping or tools) *"he put his golf bag in the trunk"*
    - S: (n) rear window (car window that allows vision out of the back of the car)
    - S: (n) reverse, reverse gear (the gears by which the motion of a machine can be reversed)
    - S: (n) roof (protective covering on top of a motor vehicle)
    - S: (n) running board (a narrow footboard serving as a step beneath the doors of some old cars)
    - S: (n) stabilizer bar, anti-sway bar (a rigid metal bar between the front suspensions and between the rear suspensions of cars and trucks; serves to stabilize the ch
    - S: (n) sunroof, sunshine-roof (an automobile roof having a sliding or raisable panel) *"'sunshine-roof' is a British term for 'sunroof'"*
    - S: (n) tail fin, tailfin, fin (one of a pair of decorations projecting above the rear fenders of an automobile)
    - S: (n) third gear, third (the third from the lowest forward ratio gear in the gear box of a motor vehicle) *"you shouldn't try to start in third gear"*
    - S: (n) window (a transparent opening in a vehicle that allow vision out of the sides or back; usually is capable of being opened)

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)

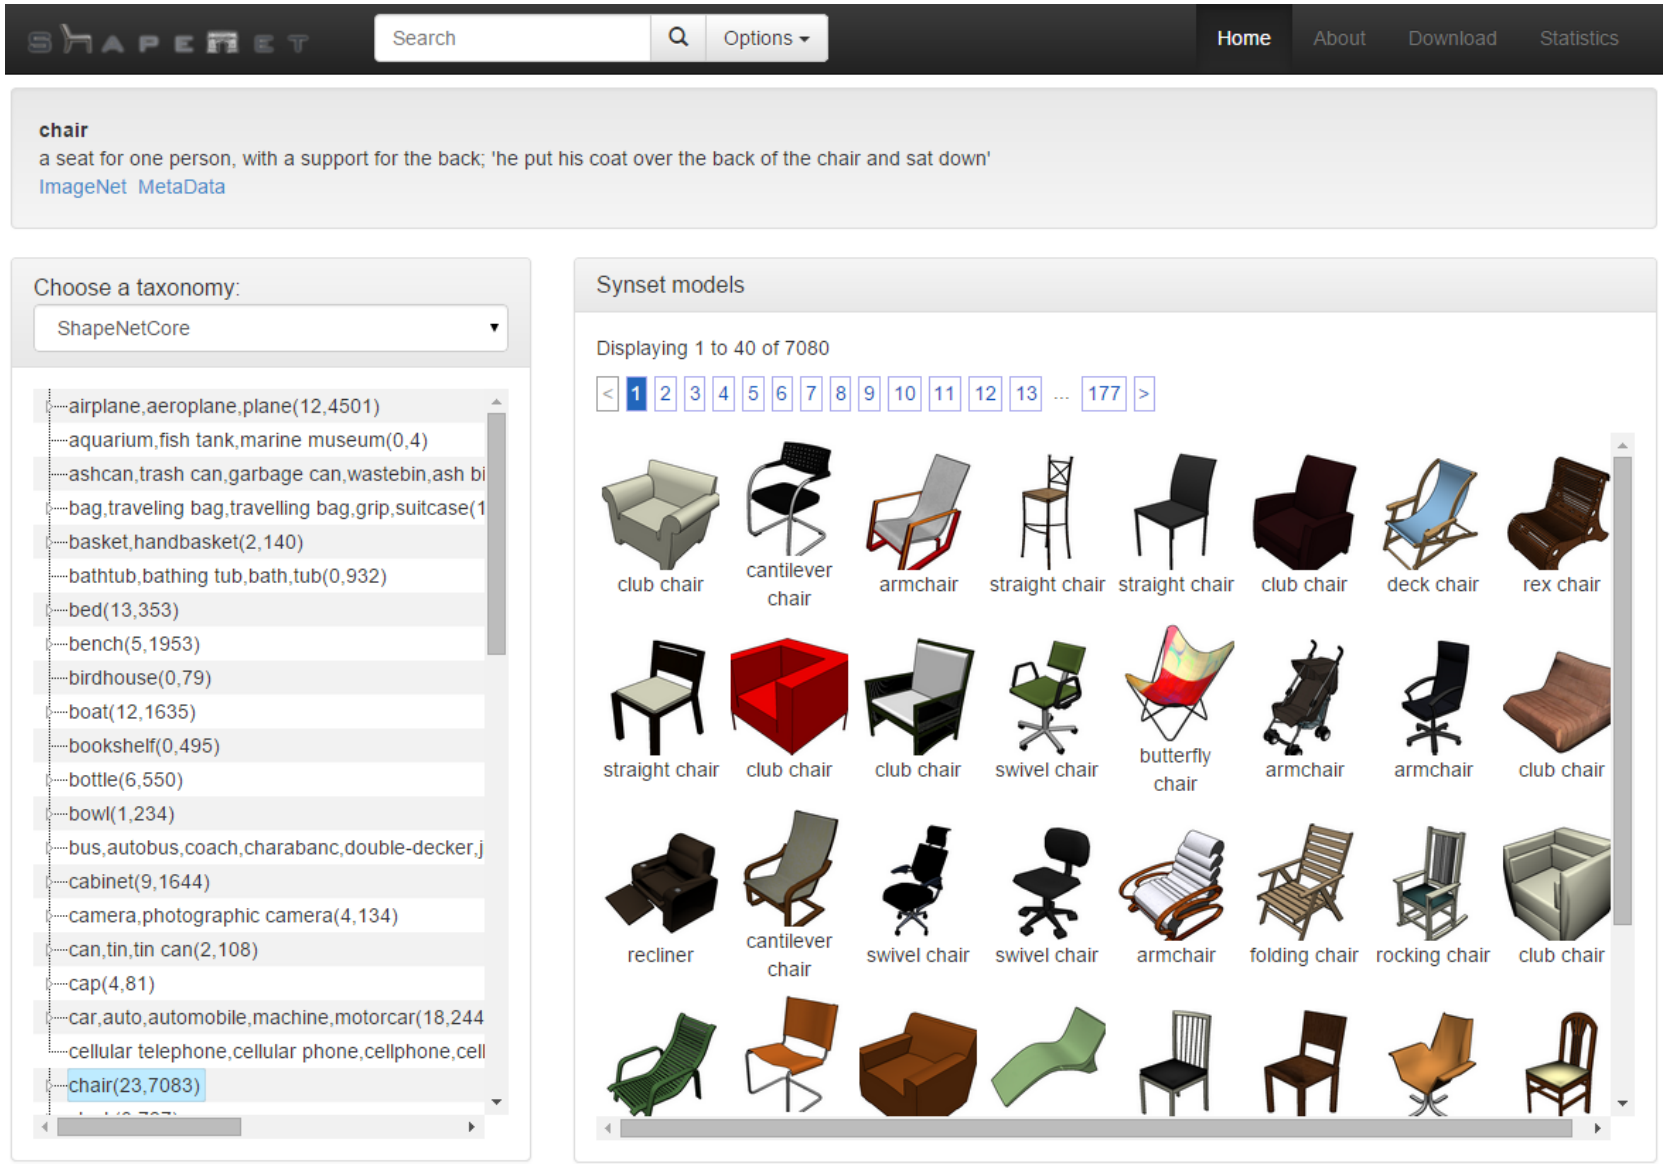

Object Knowledge

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

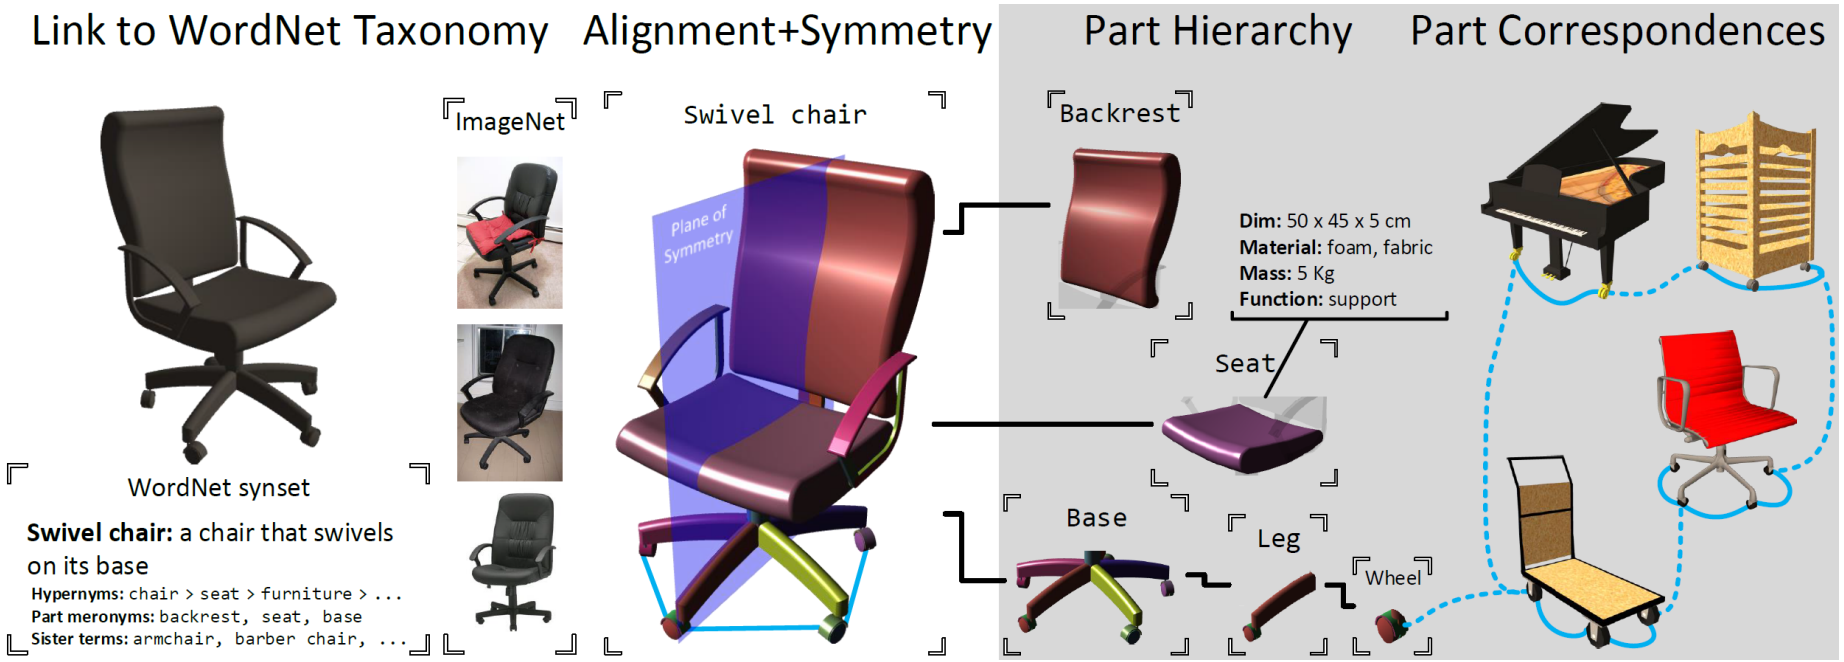

Object Interaction Knowledge

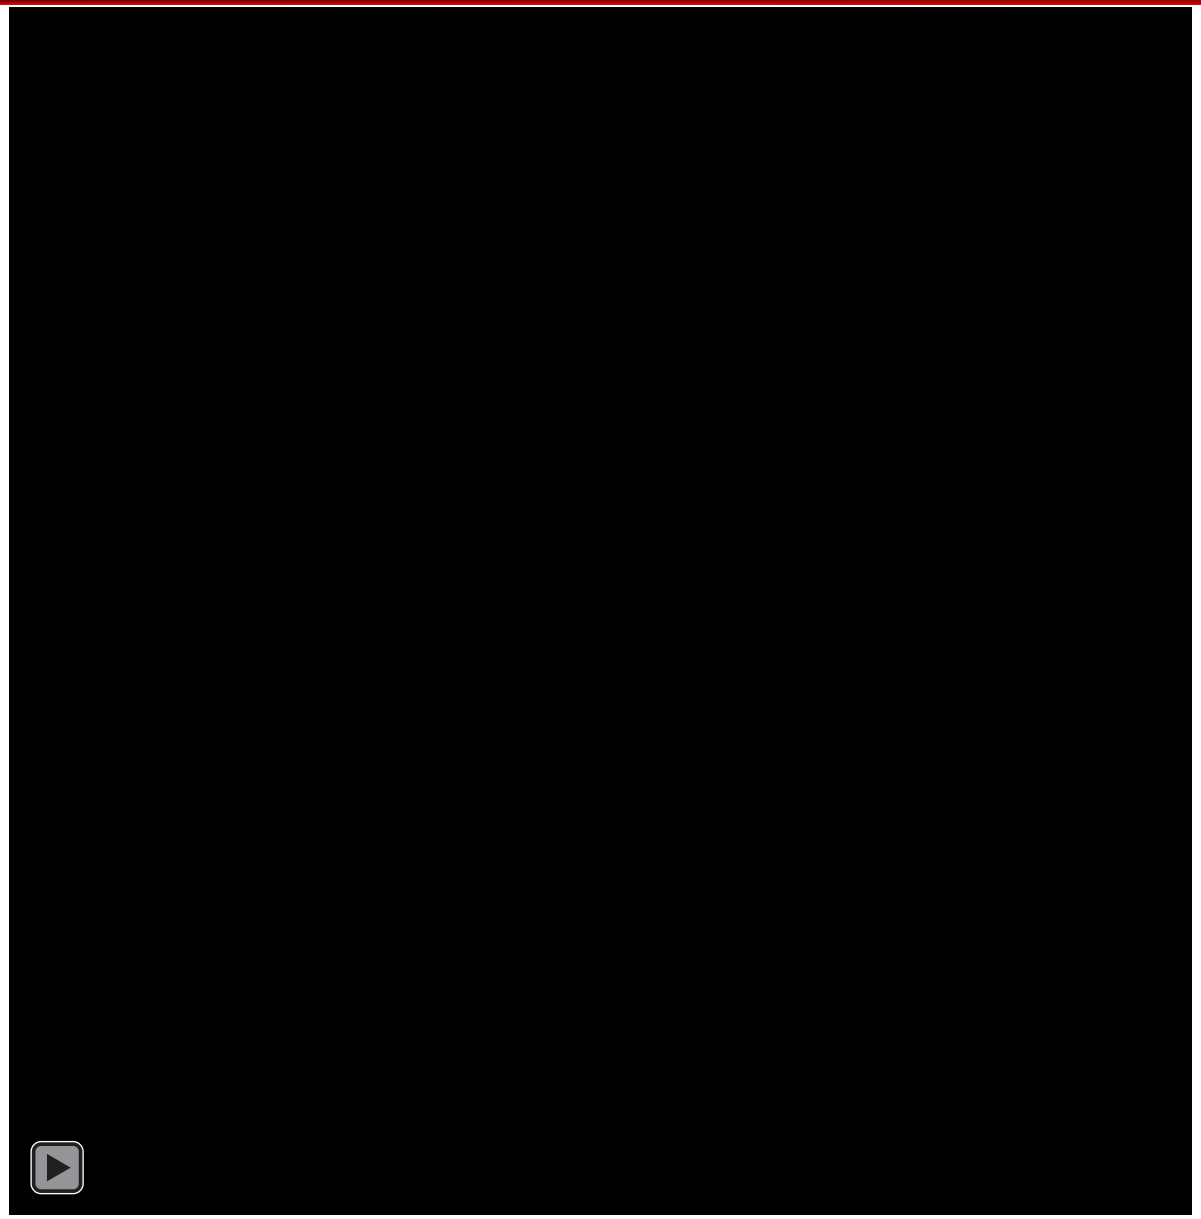

Vector Field to Histograms

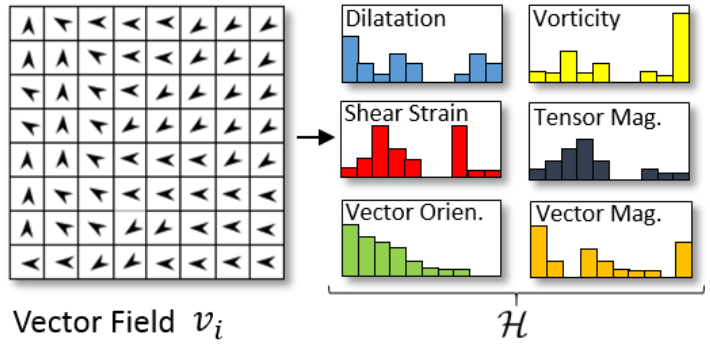

Supervised Methods
Visual DataSets

ImageNet and ShapeNet

- Explain how big visual datasets including ImageNet and ShapeNet are organized

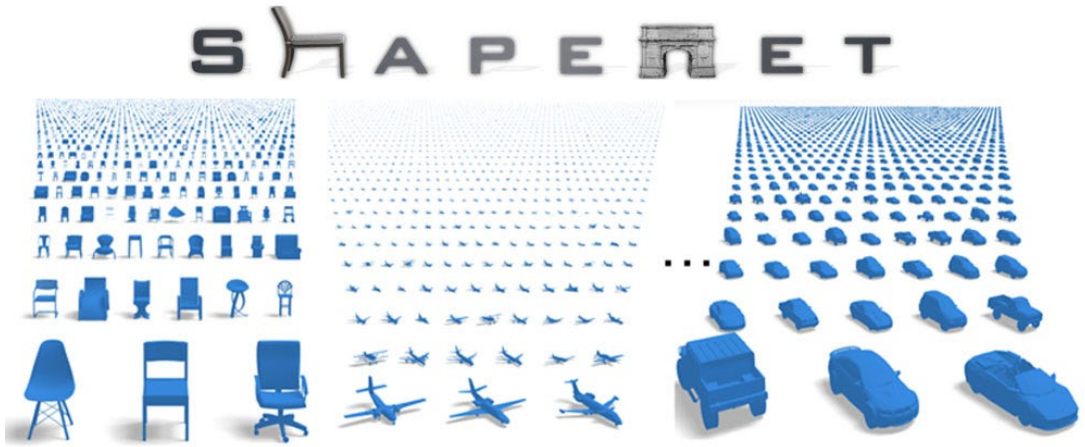

Lecture on Wednesday, March 31

- Live lecture by Kaichun Mo will be from 10:30 – 11:50 am (of course it will be recorded)

Course Mechanics

Class Mechanics

- Two online weekly lectures
- Online office hours
- Class web site http://cs233.stanford.edu

  http://graphics.stanford.edu/courses/cs233-21-spring
- Use Canvas for lecture videos
- Use Piazza (discussion group), Gradescope (for homework submissions)

Course Work

- Four assignments
  - modest programming in MATLAB/Python for three, one in JavaPlex
  - working in groups OK, up to three students
  - Use Google Cloud for Education
- An open book midterm
- Class participation
- No final

CS233 / CME 251 Class Schedule

Monday	Wednesday
March 29	**March 31**
Introduction; Geometric and topological perspective on data analysis; Data representations; Learning on point clouds and graphs; Joint data analysis. Lecture Slides: Reading:	Visual data sets: ImageNet and ShapeNet; Techniques for annotation and annotation transport. Lecture Slides: Reading:
April 5	**April 7**
Linear algebraic techniques: principal components analysis (PCA), Kernel PCA. Lecture Slides: Reading:	Linear algebraic techniques: canonical correlation analysis (CCA). Multidimensional scaling (MDS). Lecture Slides: Reading: Homework 1 out.
April 12	**April 14**
Graph methods; spectral approaches, graph Laplacians, Laplacian embeddings, spectral clustering. Lecture Slides: Reading:	Non-linear dimensionality reduction: locally linear embeddings, Laplacian eignemaps, Isomap, t-SNE. Lecture Slides: Reading:
April 19	**April 21**
Computational topology: topology review, complexes, homology groups. Lecture Slides: Reading:	Persistent homology, barcodes and persistence diagrams. Reading: Lecture Slides: Homework 1 due. Homework 2 out.

Monday	Wednesday
April 26	**April 28**
Topological inference; the Mapper algorithm. Applications. Lecture Slides: Reading:	Representations of 3D Geometry: Voxel-Grids, Point Clouds, Meshes and Other Boundary Models, Solid Models. Lecture Slides: Reading:
May 3	**May 5**
Geometry processing; Laplace-Beltrami and other operators on meshes. Lecture Slides: Reading:	Rigid and non-rigid shape alignment. Global and local shape descriptors; intrinsic descriptors, heat and wave kernel signatures. Lecture Slides: Reading: Homework 2 due. Homework 3 out.
May 10	**May 12**
Class Midterm	Geometric deep learning; Volumetric and multi-view CNNs for 3D geometry Lecture Slides: Reading:
May 17	**May 19**
Deep nets for pointclouds and applications to classification and segmentation. Lecture Slides: Reading:	Functional spaces and functional maps, variations; map visualization. Lecture Slides: Reading: Homework 3 due. Homework 4 out.

Monday	Wednesday
May 24	**May 26**
Networks of shapes and images; cycle consistency; map processing and latent spaces. Lecture Slides: Reading:	Deep nets for graphs and meshes. Lecture Slides: Reading:
May 31	**June 2**
Memorial day holiday -- no class	Encoding shape differences and shape variability. Class summary. Lecture Slides: Reading: Homework 4 due Friday, June 4.

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 the data
- Present methods for joint data analysis – benefiting from the “wisdom of the collection”

Data Has Shape

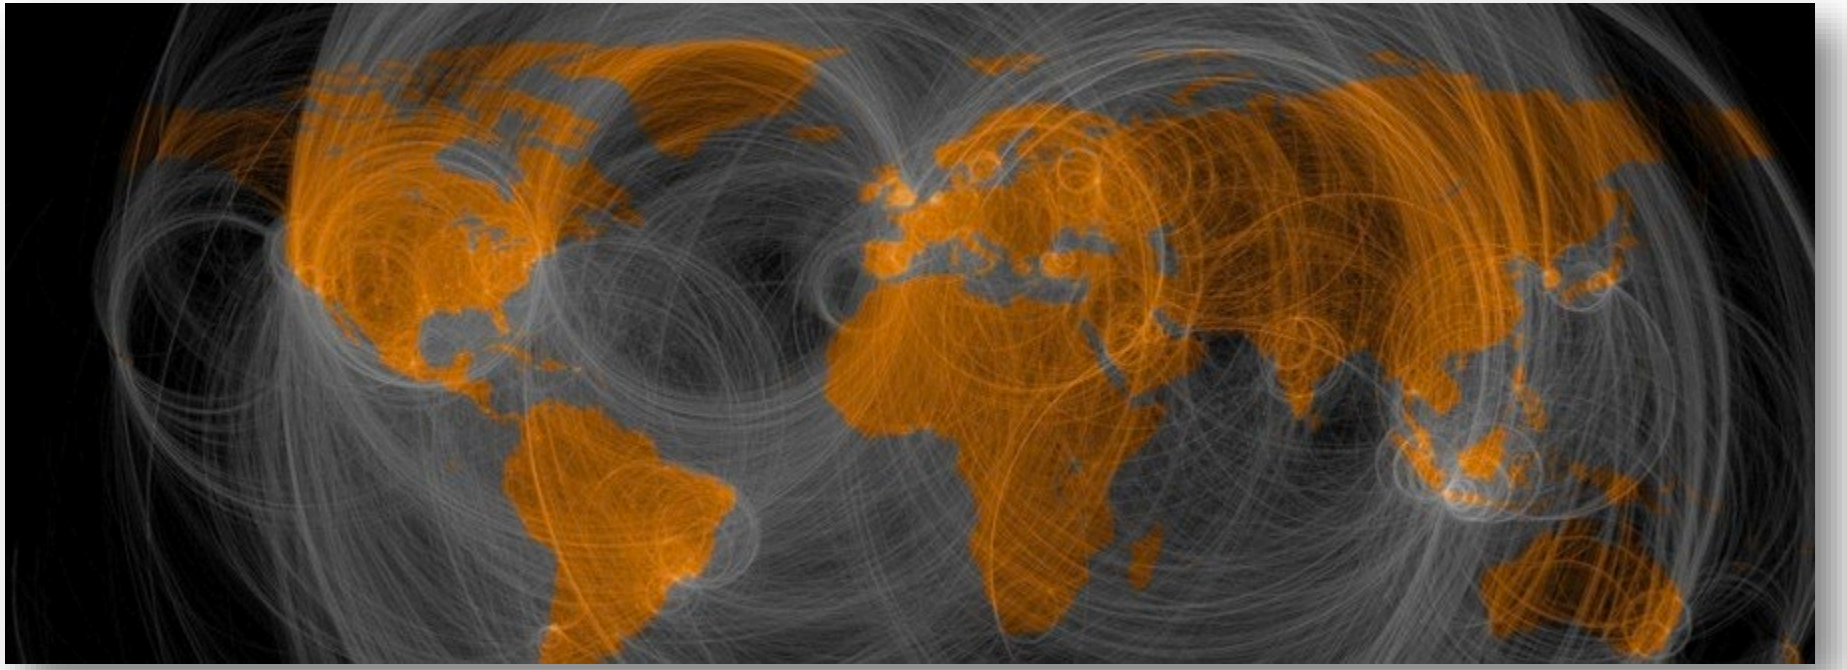

That’s All