CS233, CME251: Geometric and Topological Data Analysis

Leonidas Guibas
Computer Science Department
Stanford University

Lecture 1
28 March 2022

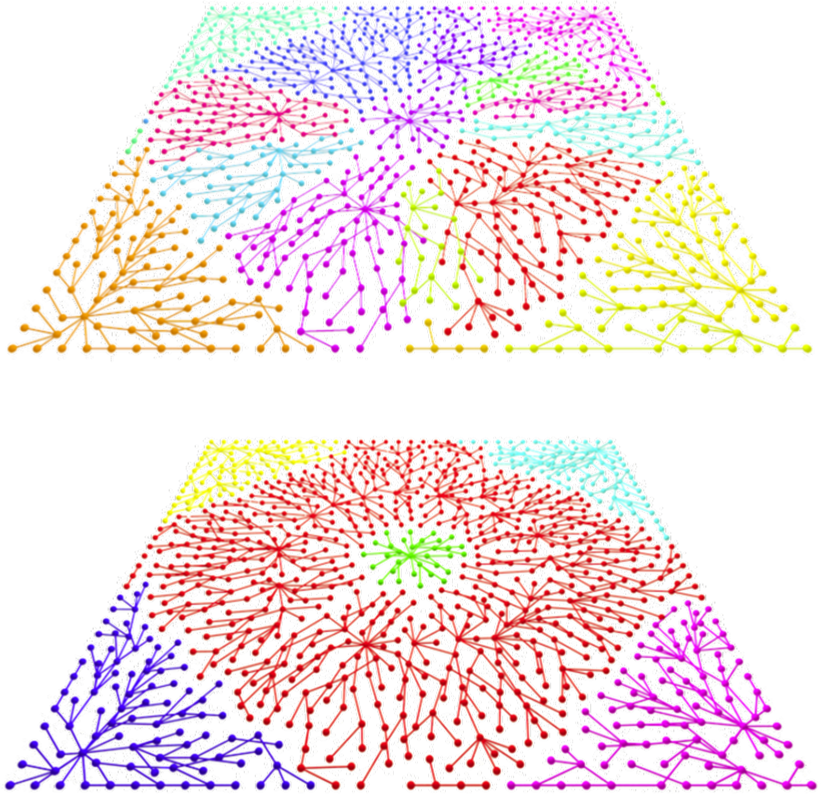

Leonidas Guibas Laboratory
Geometric Computing

Introduction

Big Data Era

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

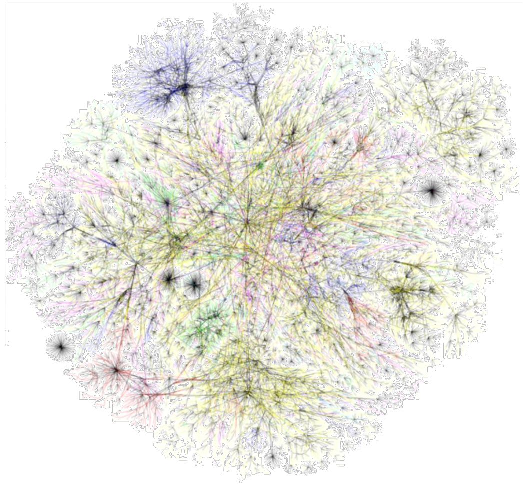

"Data Science"

Data Sets of Geometric Character

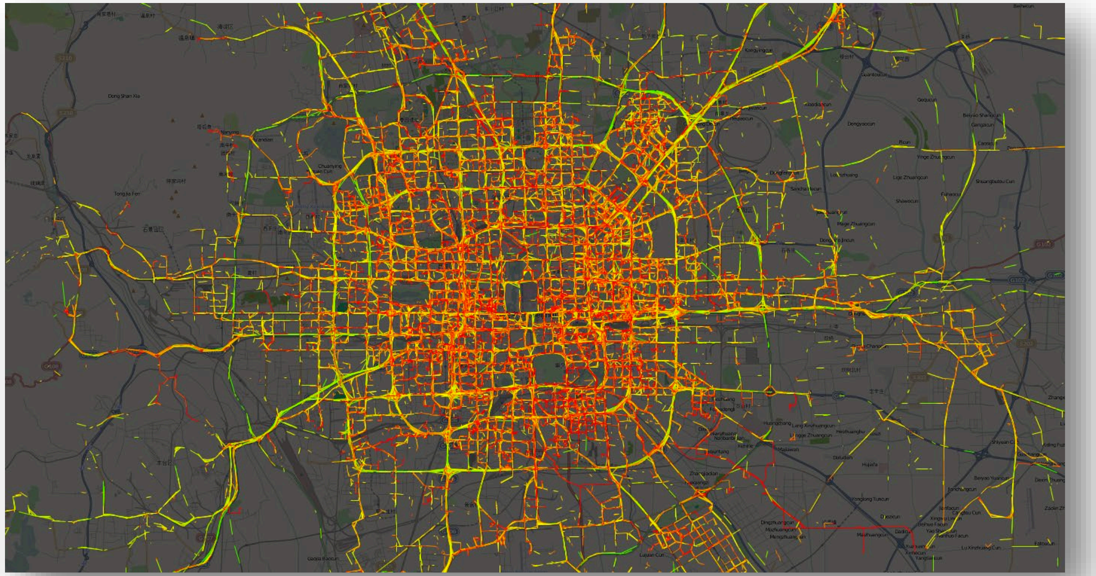

2D vehicle GPS traces (Beijing)

3D city scans (Bremen)

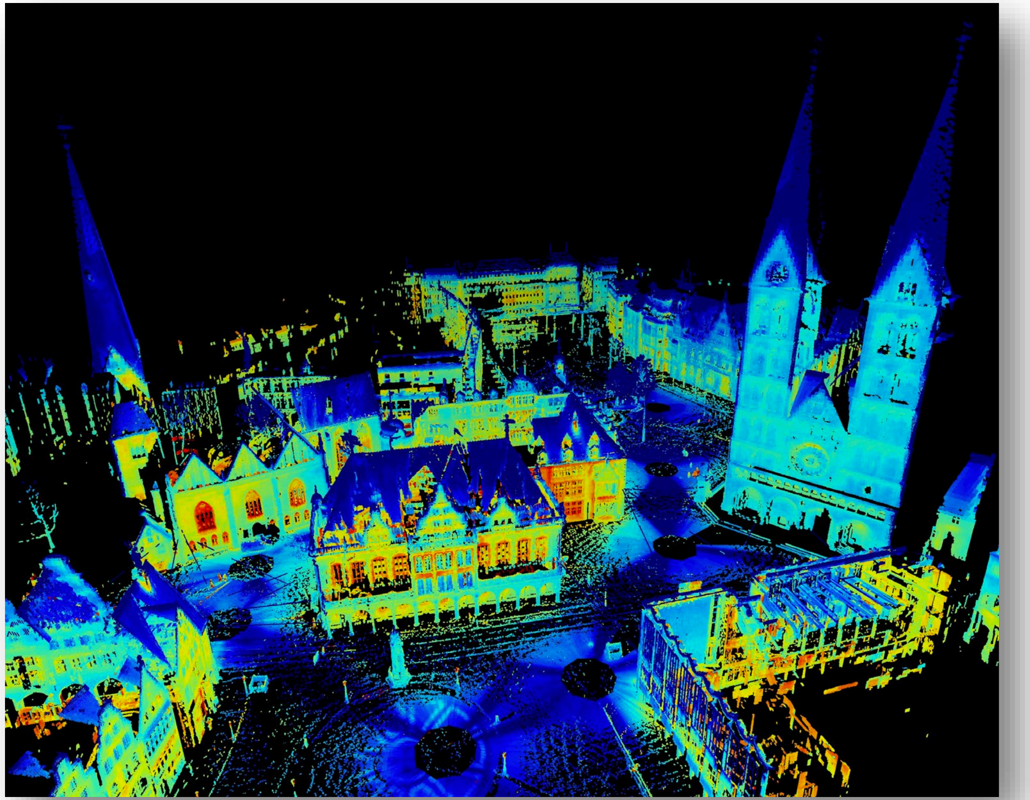

Semantics Requires Geometry

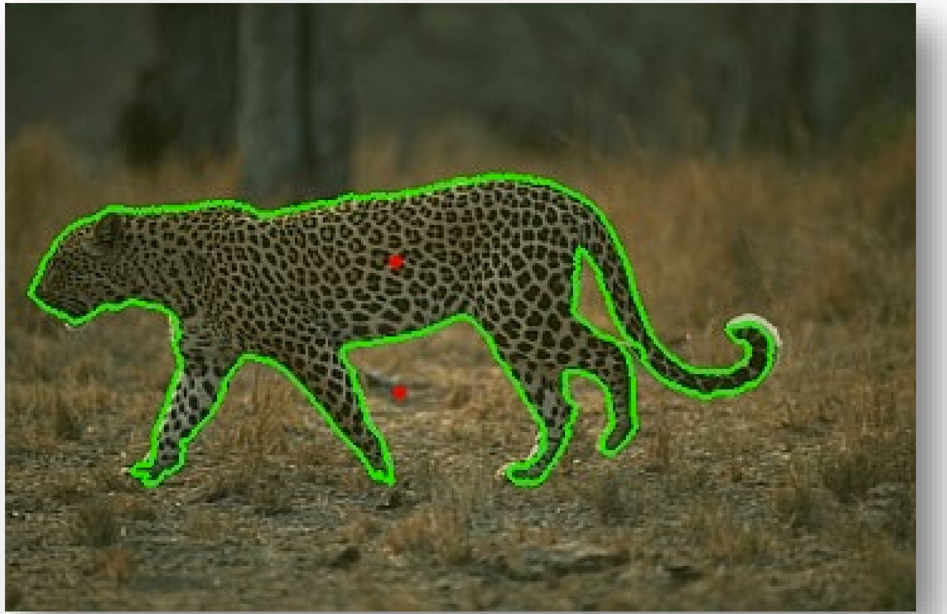

Object recognition and foreground extraction

Segmentation involves geometry

Motion involves geometry

Optical flow

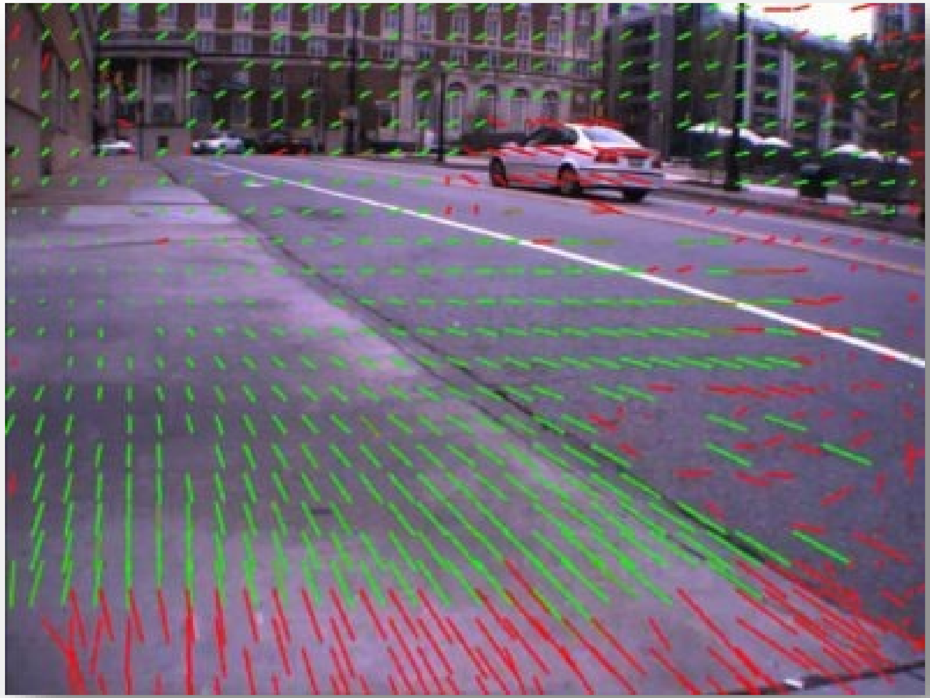

Non-Geometric Data

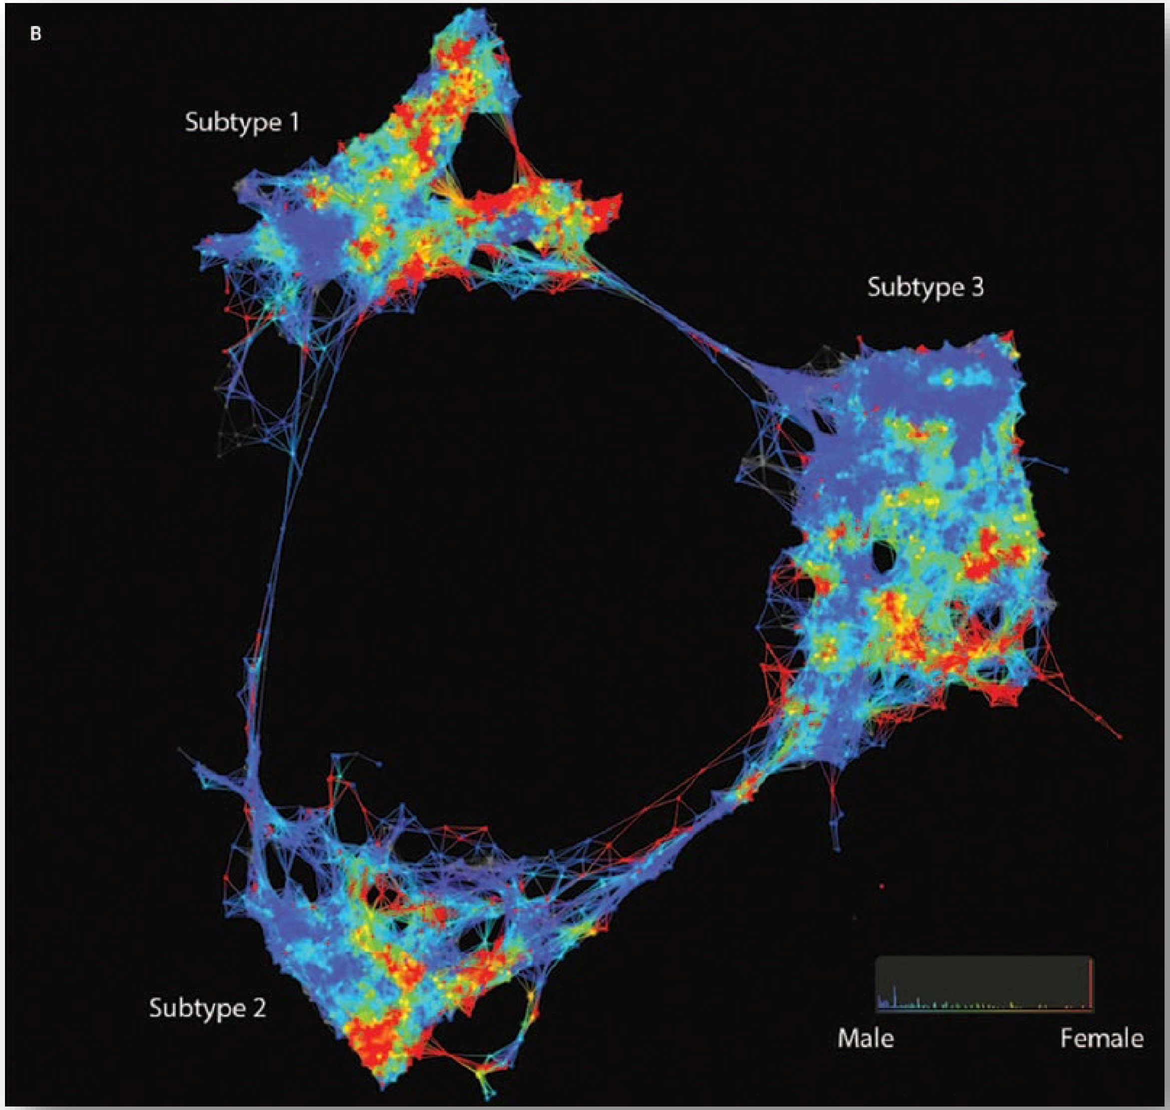

Diabetes II subtypes

Data Has Shape

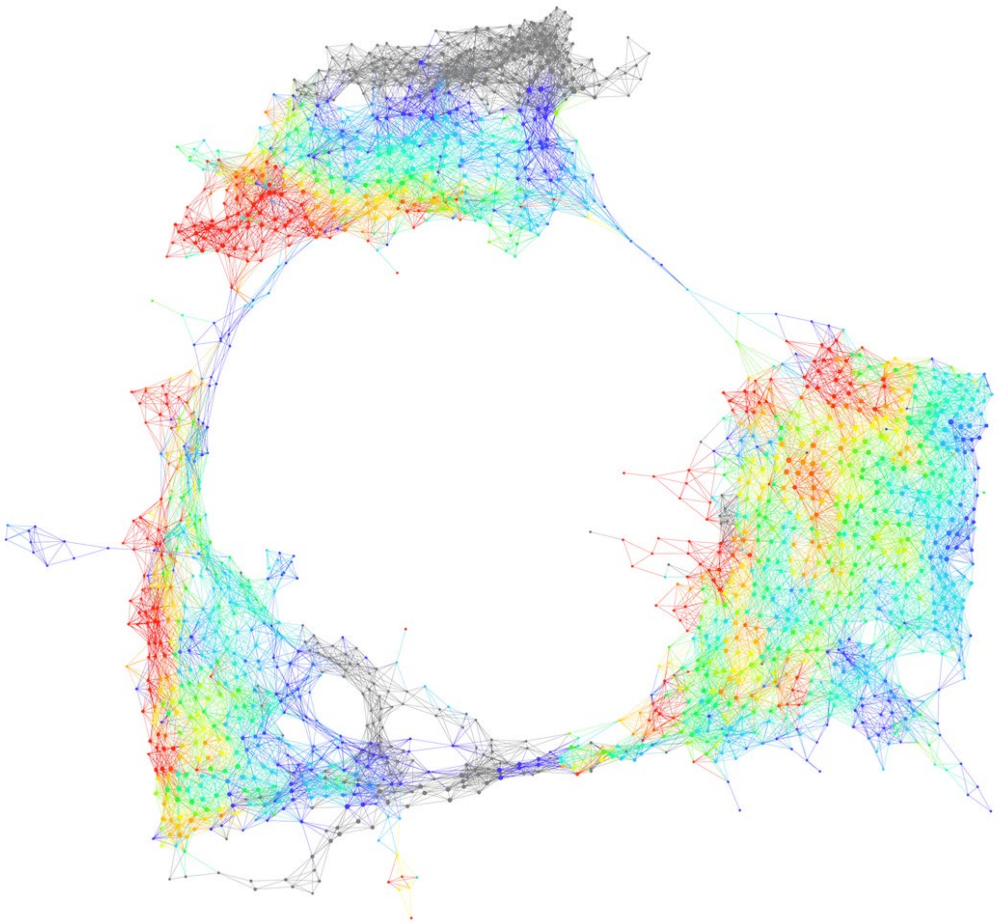

The Geometry of Data or the Data of Geometry

geometric	data	analysis

The Role of Supervision in Geometric ML

- The success of deep learning has to a large extent been made possible by huge annotated data sets and the availability of large amounts of computing power
- Yet, in many settings, obtaining quality training data at scale from human annotation is challenging – especially in geometric or 3D settings

ImageNet

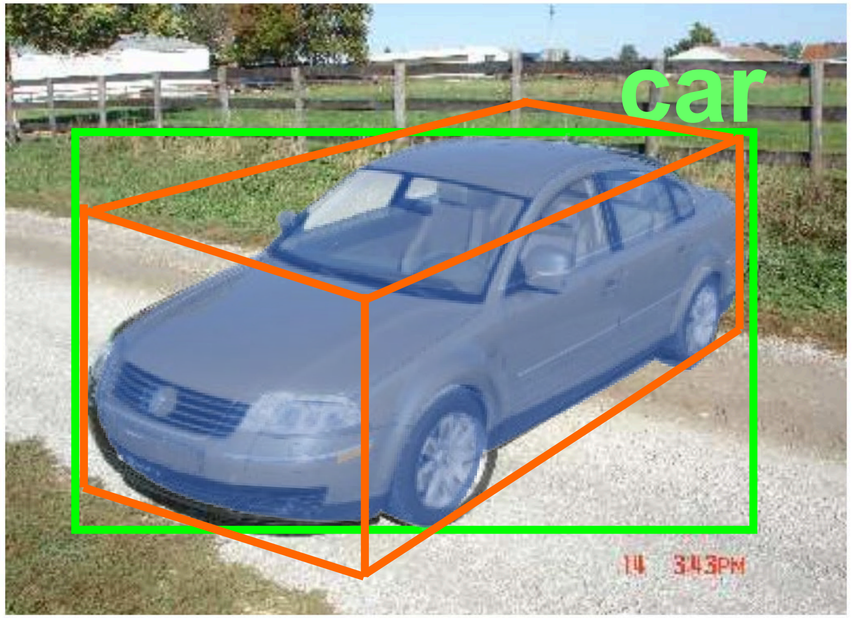

Towards Reduced or No Supervision

- Thus it is important to consider learning protocols and architectures that reduce, minimize, or eliminate supervision
  - Transfer learning, semi-supervised learning, few-shot learning, unsupervised learning

- Joint learning takes a more holistic approach towards the "social aspects" of learning
  - data, training or testing, have correlations
  - learning tasks have correlations
  - representations of data have correlations

ImageNet

Unsupervised Segmentation: “Wisdom of the Crowd”

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

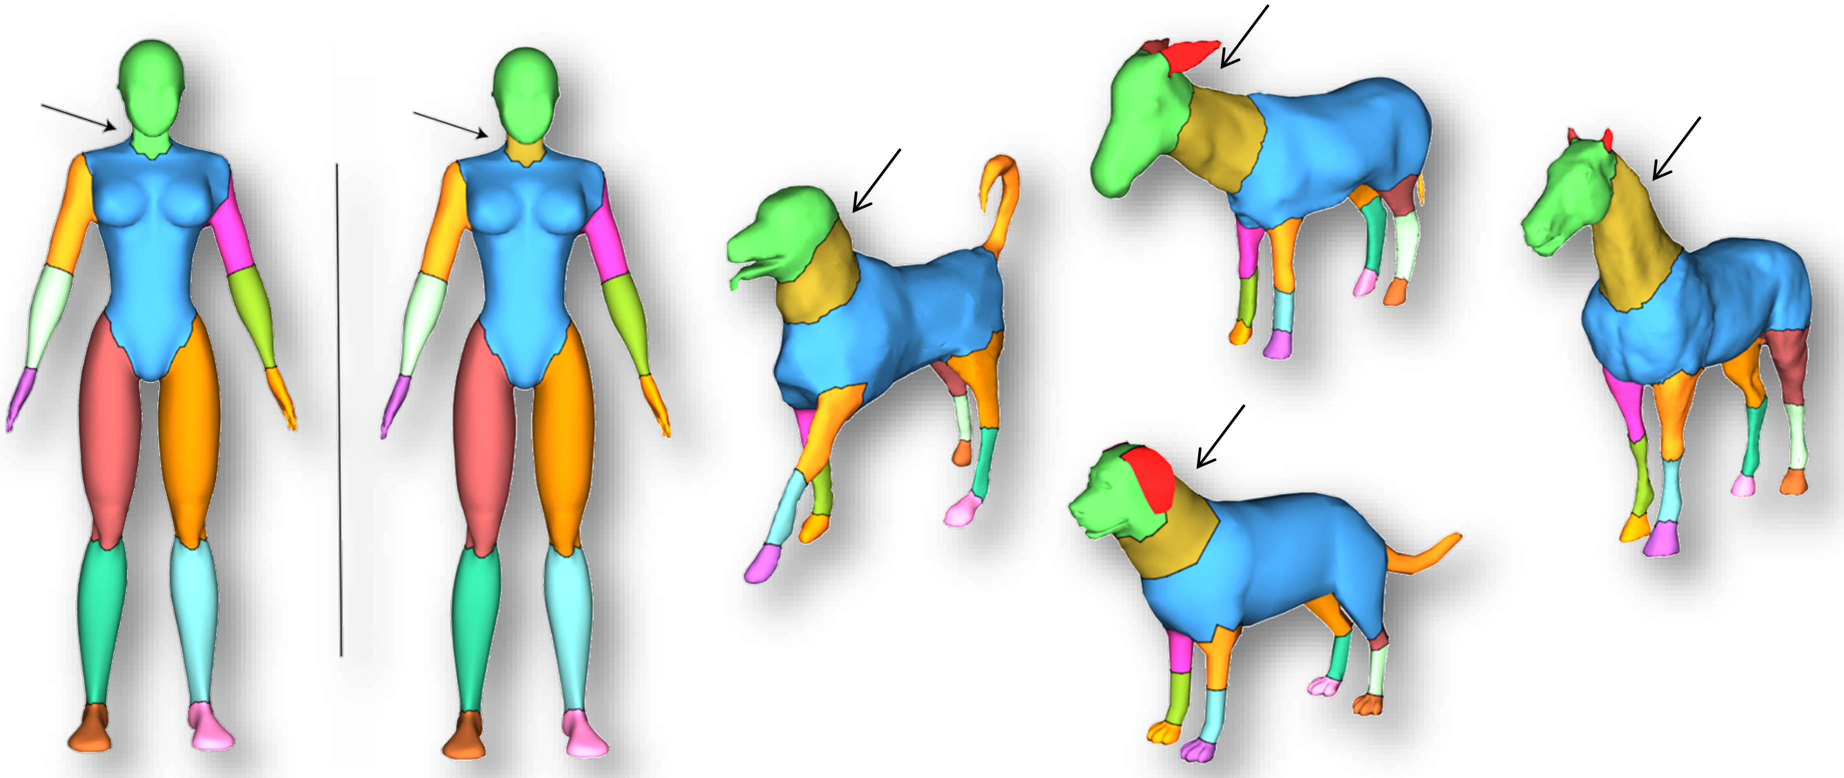

3D Segmentation

Data Integrtion Across Multiple Modalities

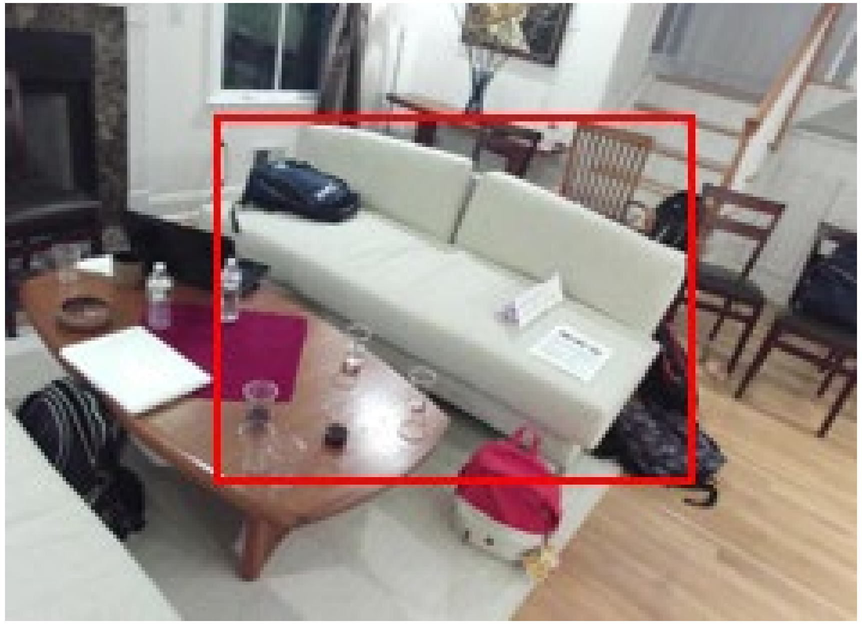

- + High resolution
- + Dense coverage
- - Subject to many imaging artifacts

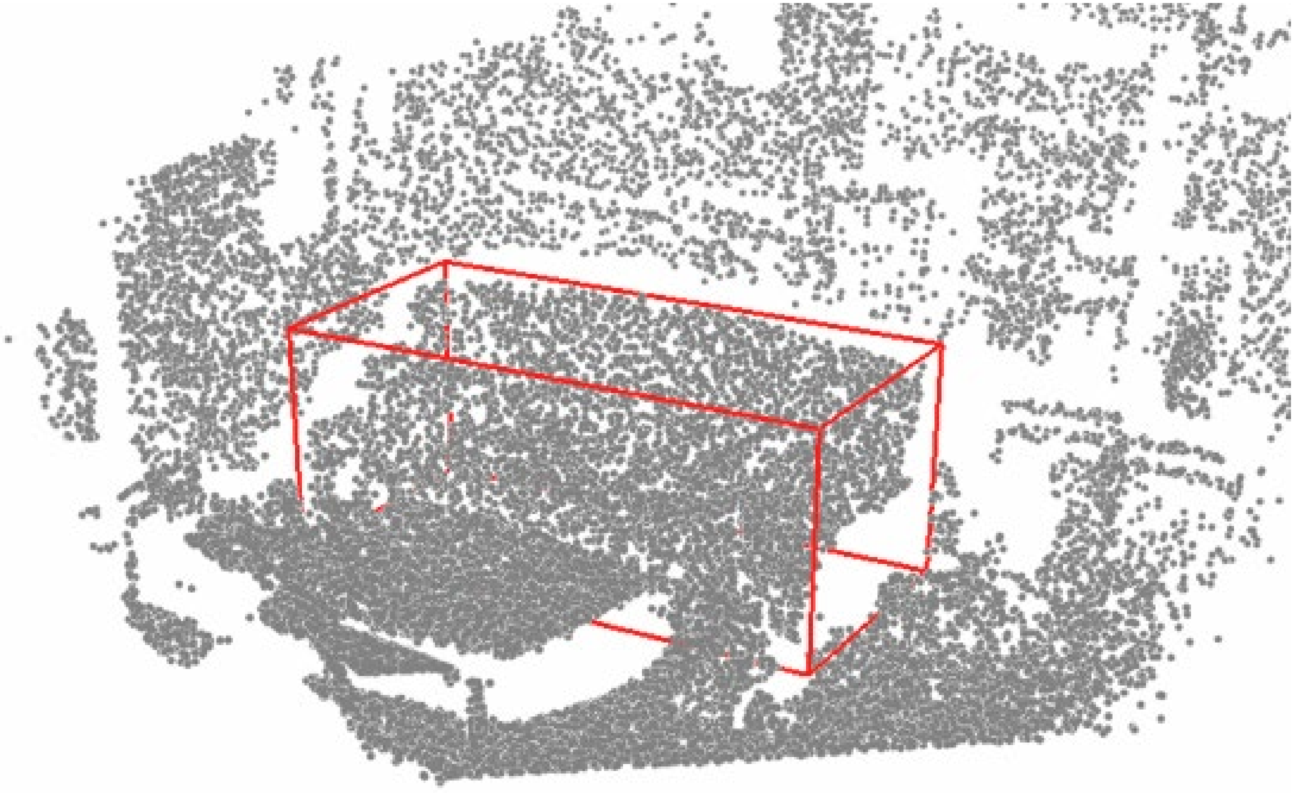

- + Absolute depth and scale
- - Sparse, low rez

Joint Learning

- 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

The Challenge: Information Transportation and Aggregation

- Aggregation needs
  - the ability to *transport information*
  - the ability to ignore *nuisance factors*
  - to ability to have a common language – *information representation consistency*

- Some powerful tools
  - voting mechanisms
  - factorization and canonicalization
  - path invariance and loop closure in transport

Tower of Babel

Look at Data from Many View Points

- Algebraic
- Graph theory based
- Geometric
- Topological

Getting to the semantics of data

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”
- Challenge: a diversity of tools … LA, ML, Stat, optimization, geometry processing, computer vision, algebraic topology …

Prerequisites / Overlaps

- Presumes some knowledge of linear algebra, optimization, algorithmic thinking, basic geometry …
- As compared to CS229, STATS216, greater emphasis on diverse, irregular data types (e.g., point clouds, graphs, meshes) and unsupervised methods, in addition to the geometry / topology angle
- Data sets will be primarily visual/geometric (images, 2D/3D point clouds, meshes, or CAD models)

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

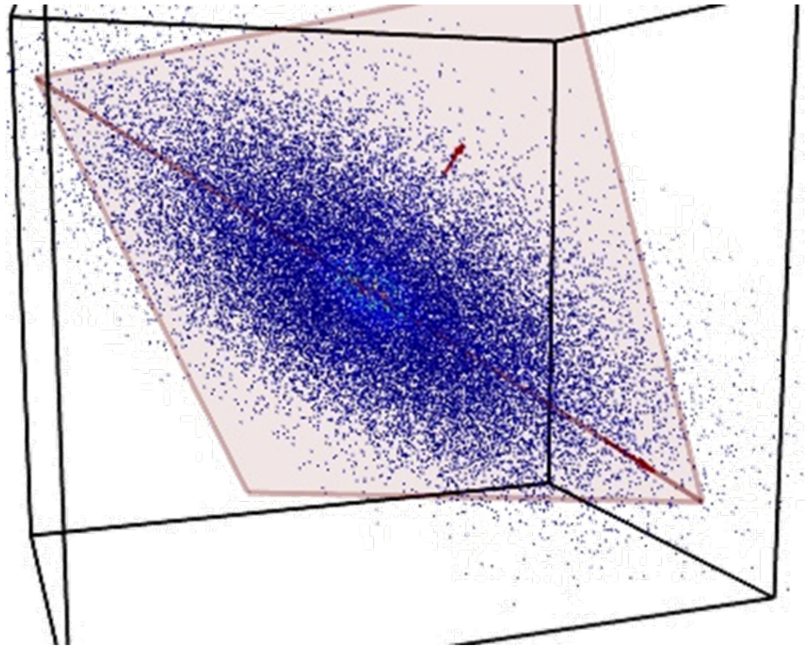

Direct and Indirect Embeddings

- Input

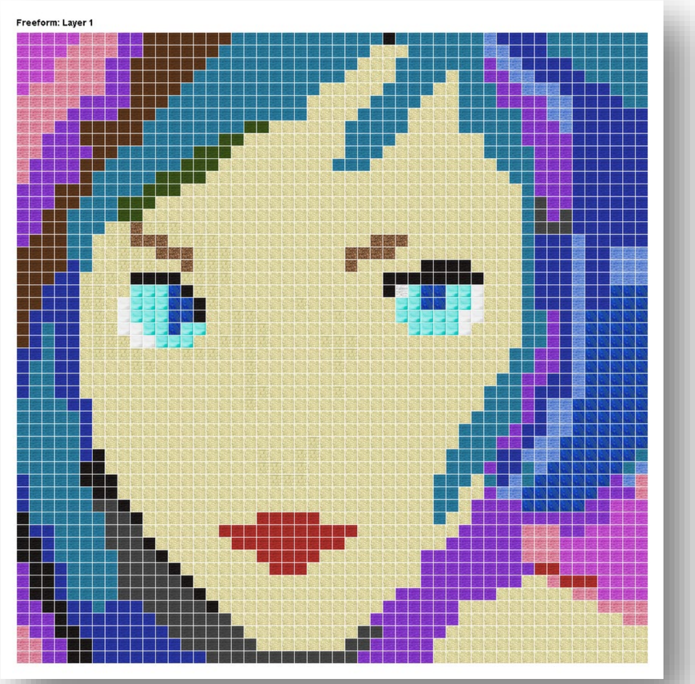

- Latent code

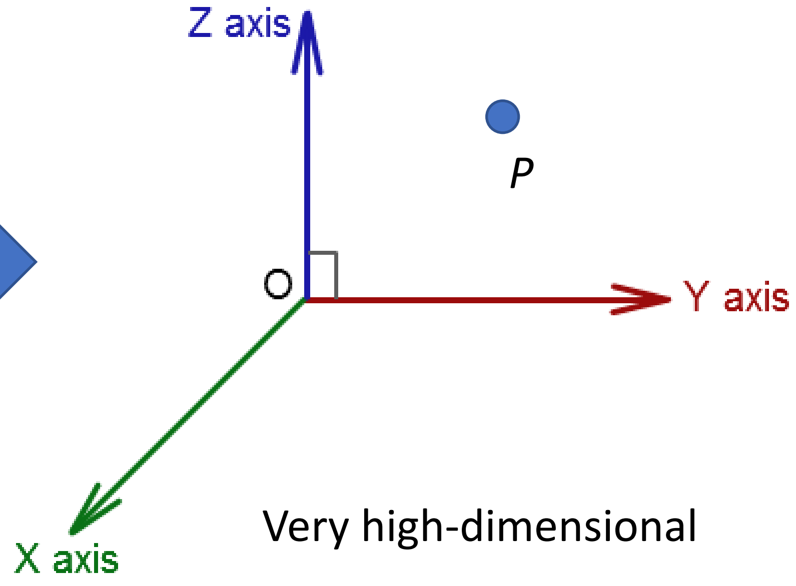

Very high-dimensional

Linear Space Methods

- Principal components analysis (PCA)

- Canonical correlation analysis (CCA)

Dimensionality reduction

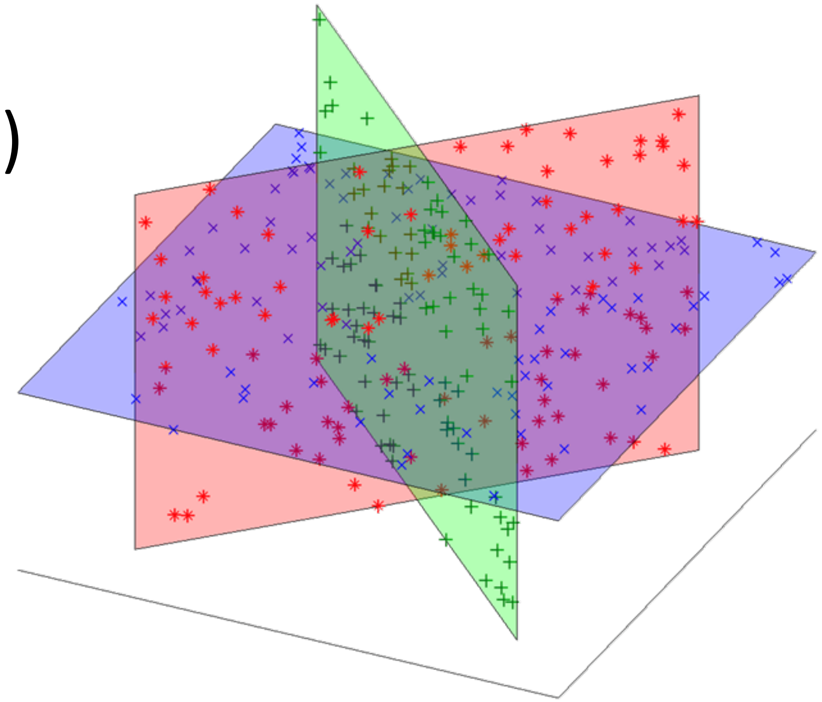

Low-d data (intrinsic dimension) living on a linear subspace, inside a high-d space (extrinsic dimension)

Data as Points on a Manifold

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

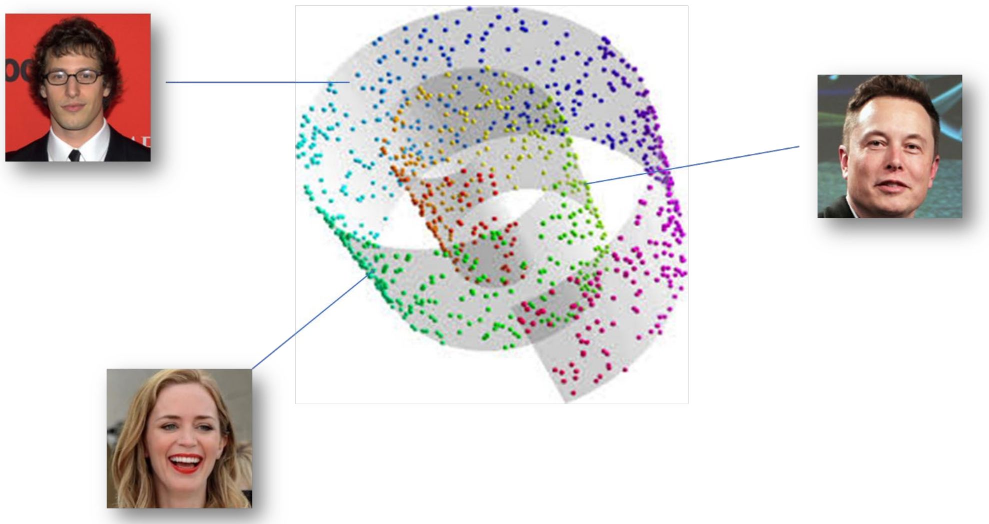

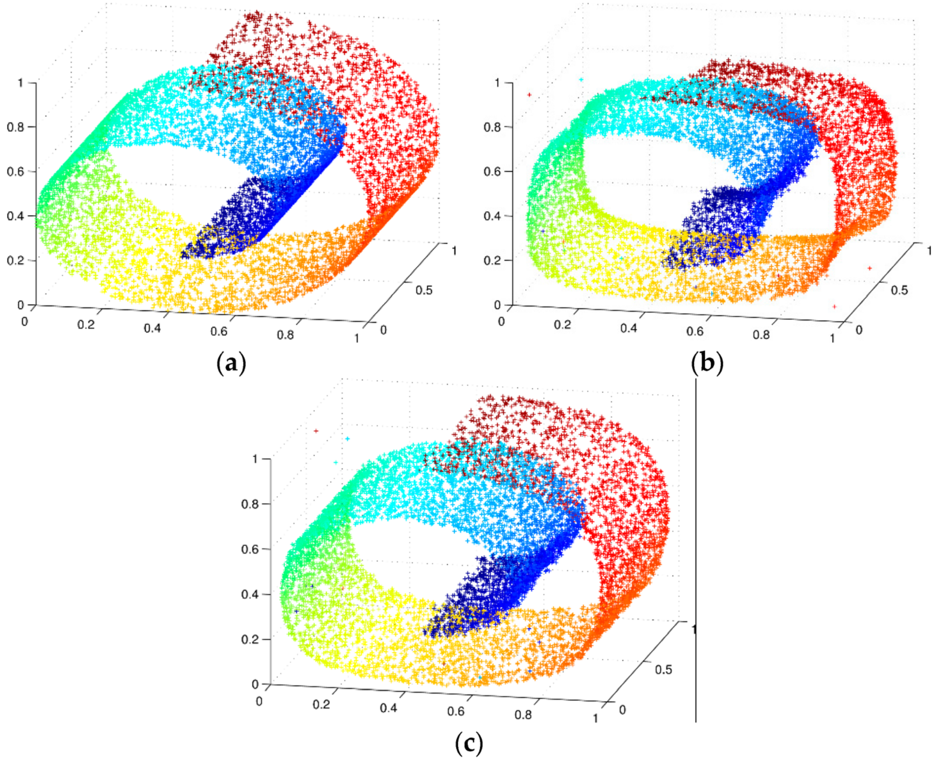

Isomap, locally linear embeddings, Laplacian eigenmaps, t-SNE

Supervised Methods, Deep Learning

Machine Learning

- Traditional Programming

Machine Learning

- Traditional Programming

- Machine Learning

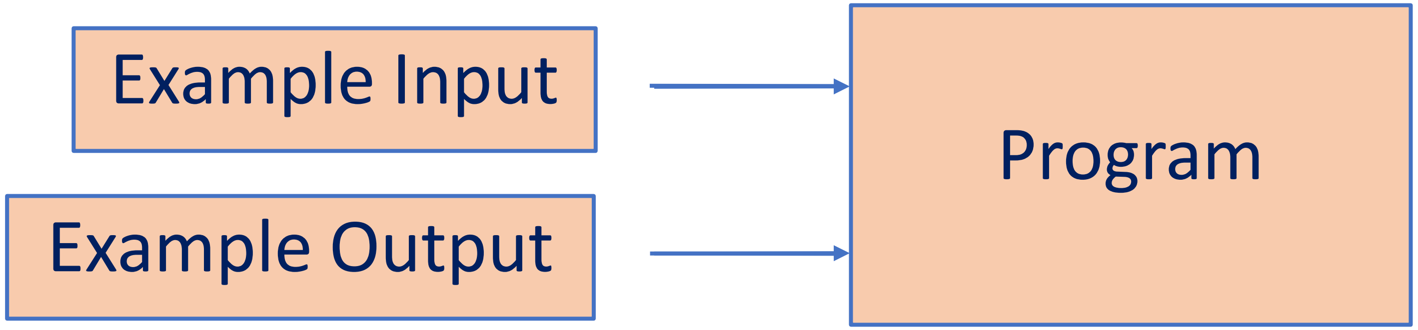

Machine Learning

f()=cat

Parametrized by many learnable parameters

Model Fitting

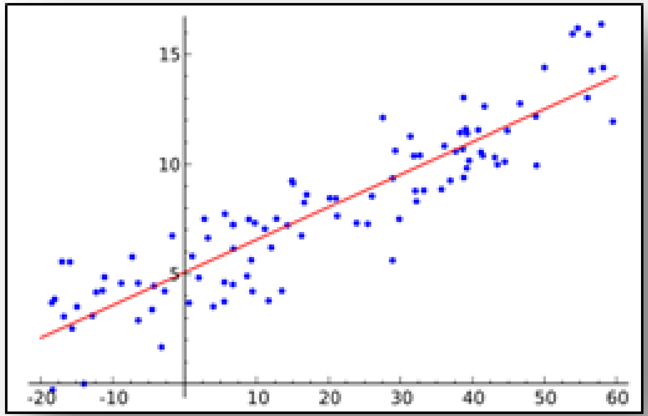

Regression

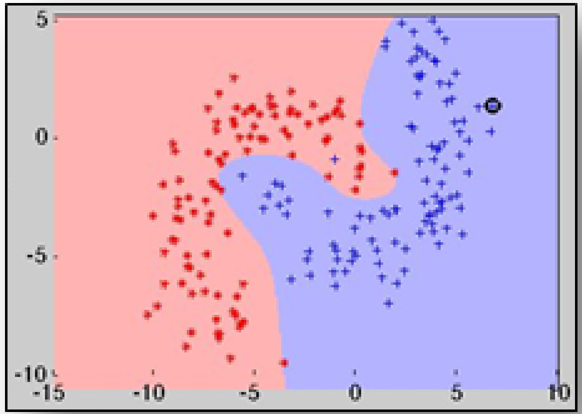

Classification

Deep Learning

- Deep learning allows computational models that are composed of **multiple processing layers** **to learn representations of data** with **multiple levels of abstraction**.

Deep Learning by Y. LeCun et al. Nature 2015

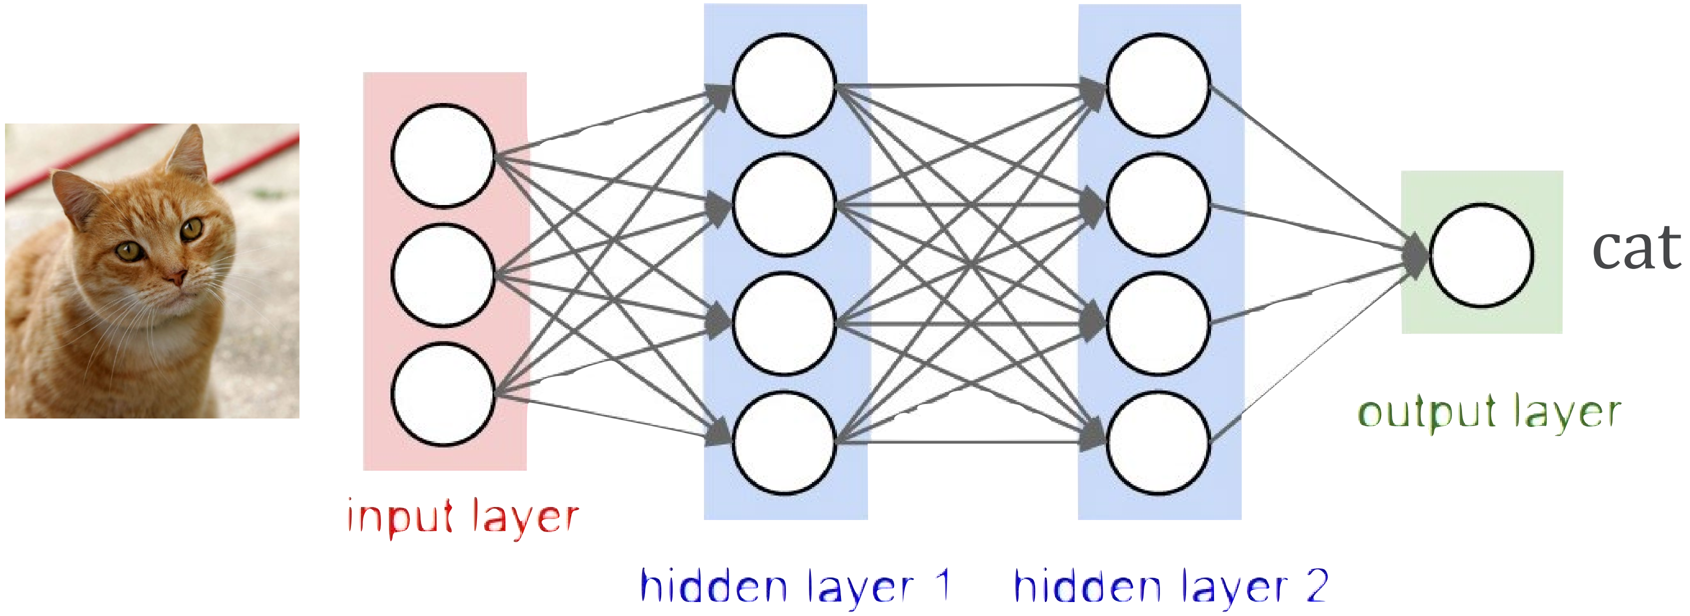

Inside Neural Networks

h1=f(w11 * x1 + w21 * x2 + w31 * x3 + b1)

f: non-linear activation function

$$\mathrm{ReLU}(x) = \begin{cases} 0 \text{ if } x < 0 \\ x \text{ if } x \geq 0 \end{cases}$$

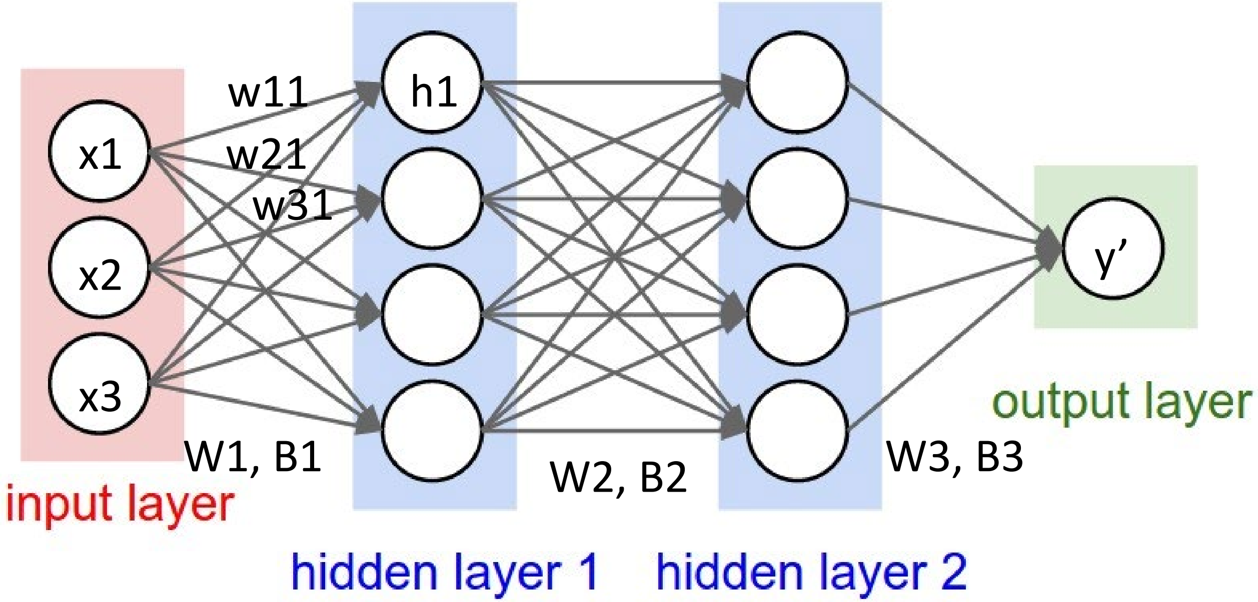

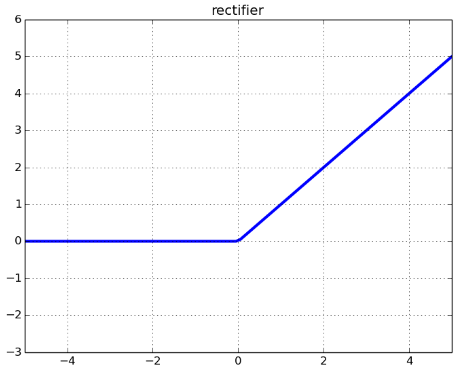

Model: Multi-Layer Perceptron (MLP)

$$y' = W_3 f(W_2 f(W_1 x + b_1) + b_2) + b_3)$$

Neural Networks as Feature Extractors

- Neural networks extract powerful features from data

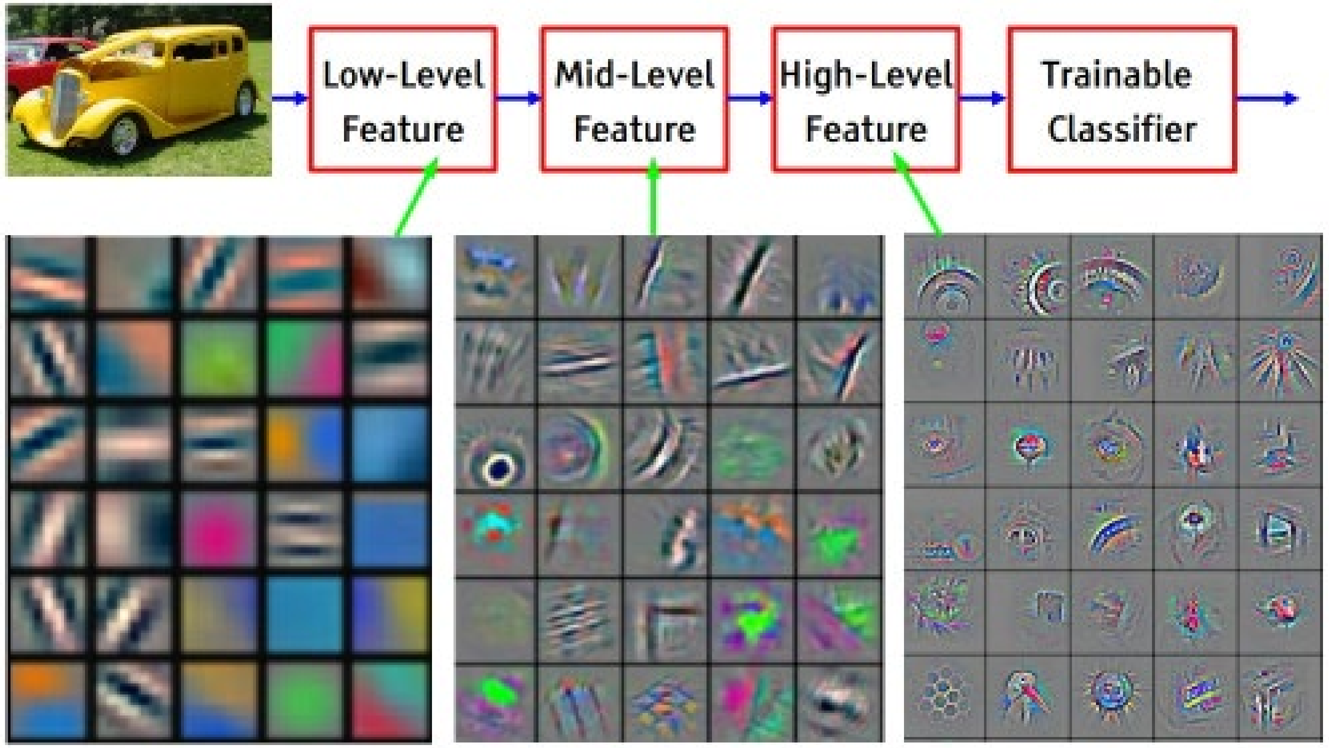

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Image Credits: Yan LeCun

Trainable Feed-Forward Learning Pipelines

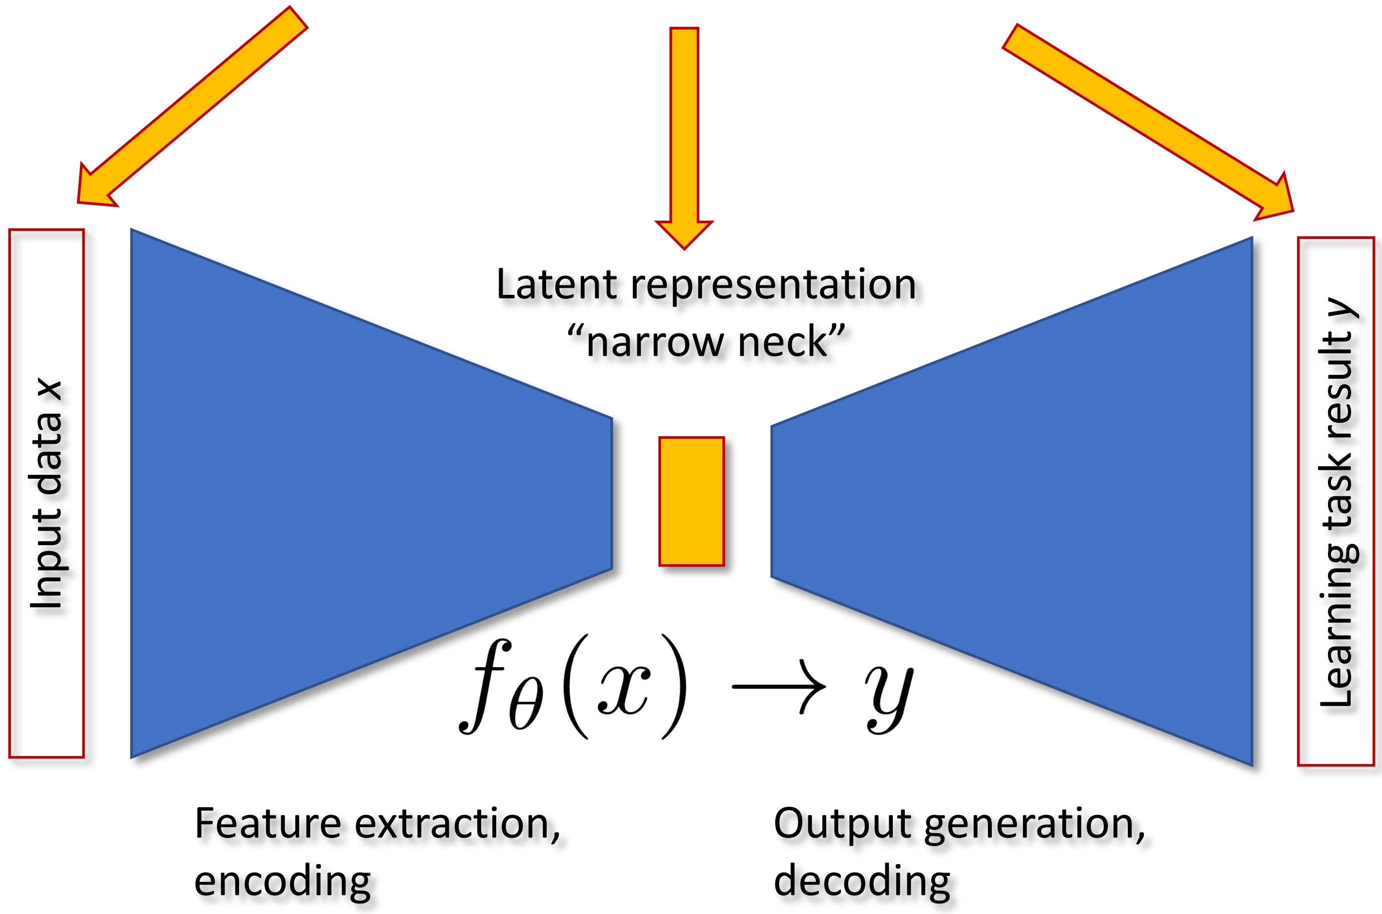

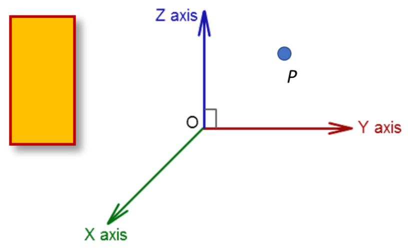

A latent code (= a point in a Euclidean feature space) acts as a low-d proxy for high-d input data, w.r.t. a learning task.

Image Deep Learning Networks

- Feed forward networks

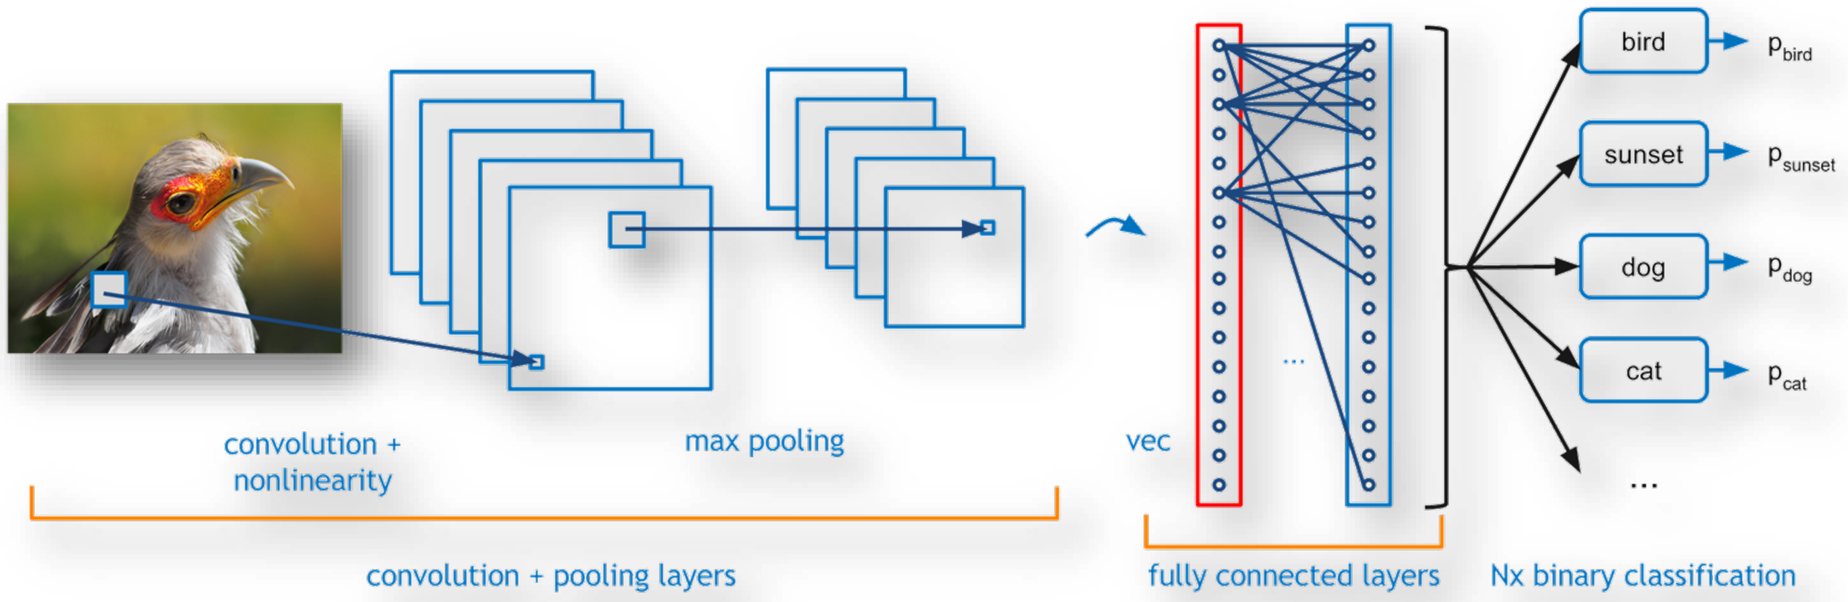

Classification

Convolutional Layers

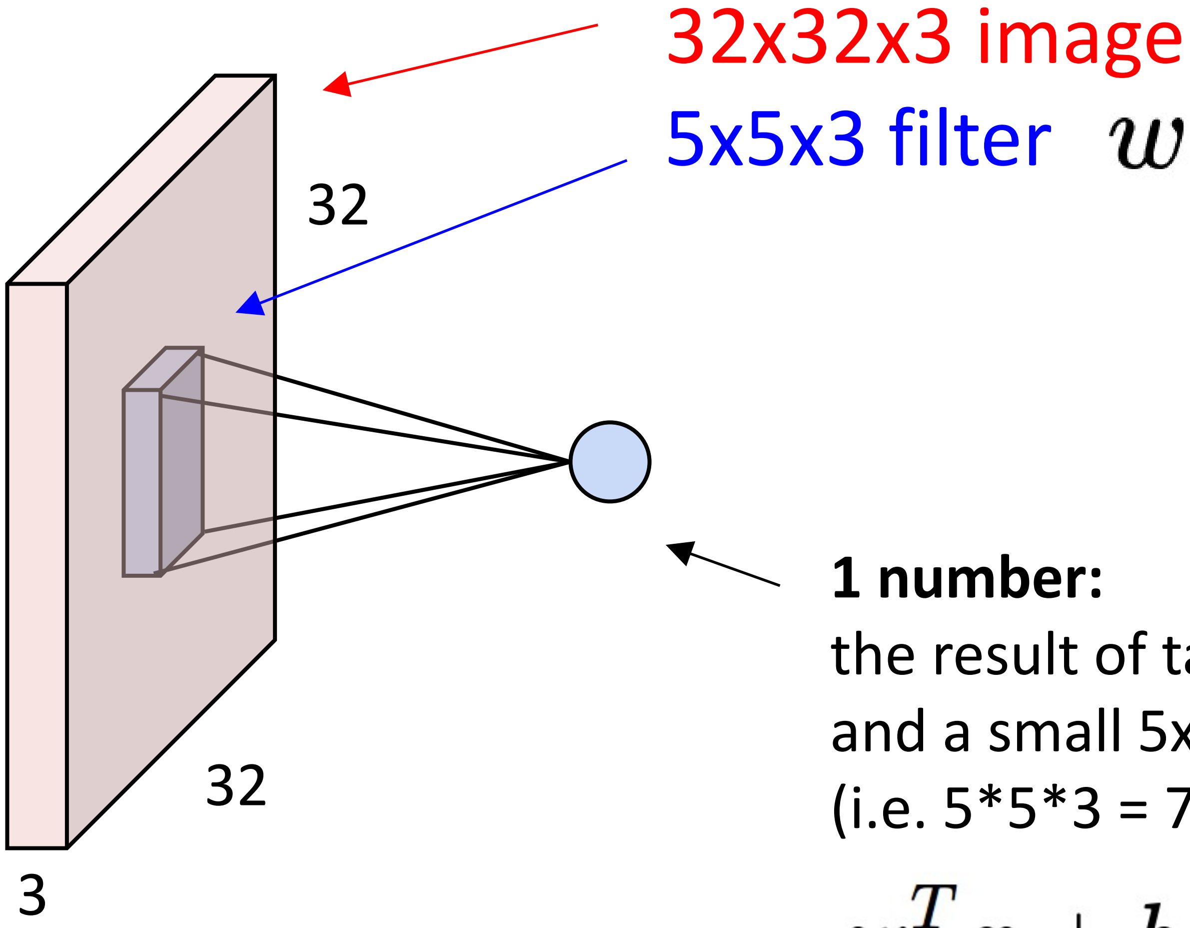

1 number:
the result of taking a dot product between the filter and a small 5x5x3 chunk of the image
(i.e. 5*5*3 = 75-dimensional dot product + bias)

$$w^T x + b$$

Convolutional Architectures

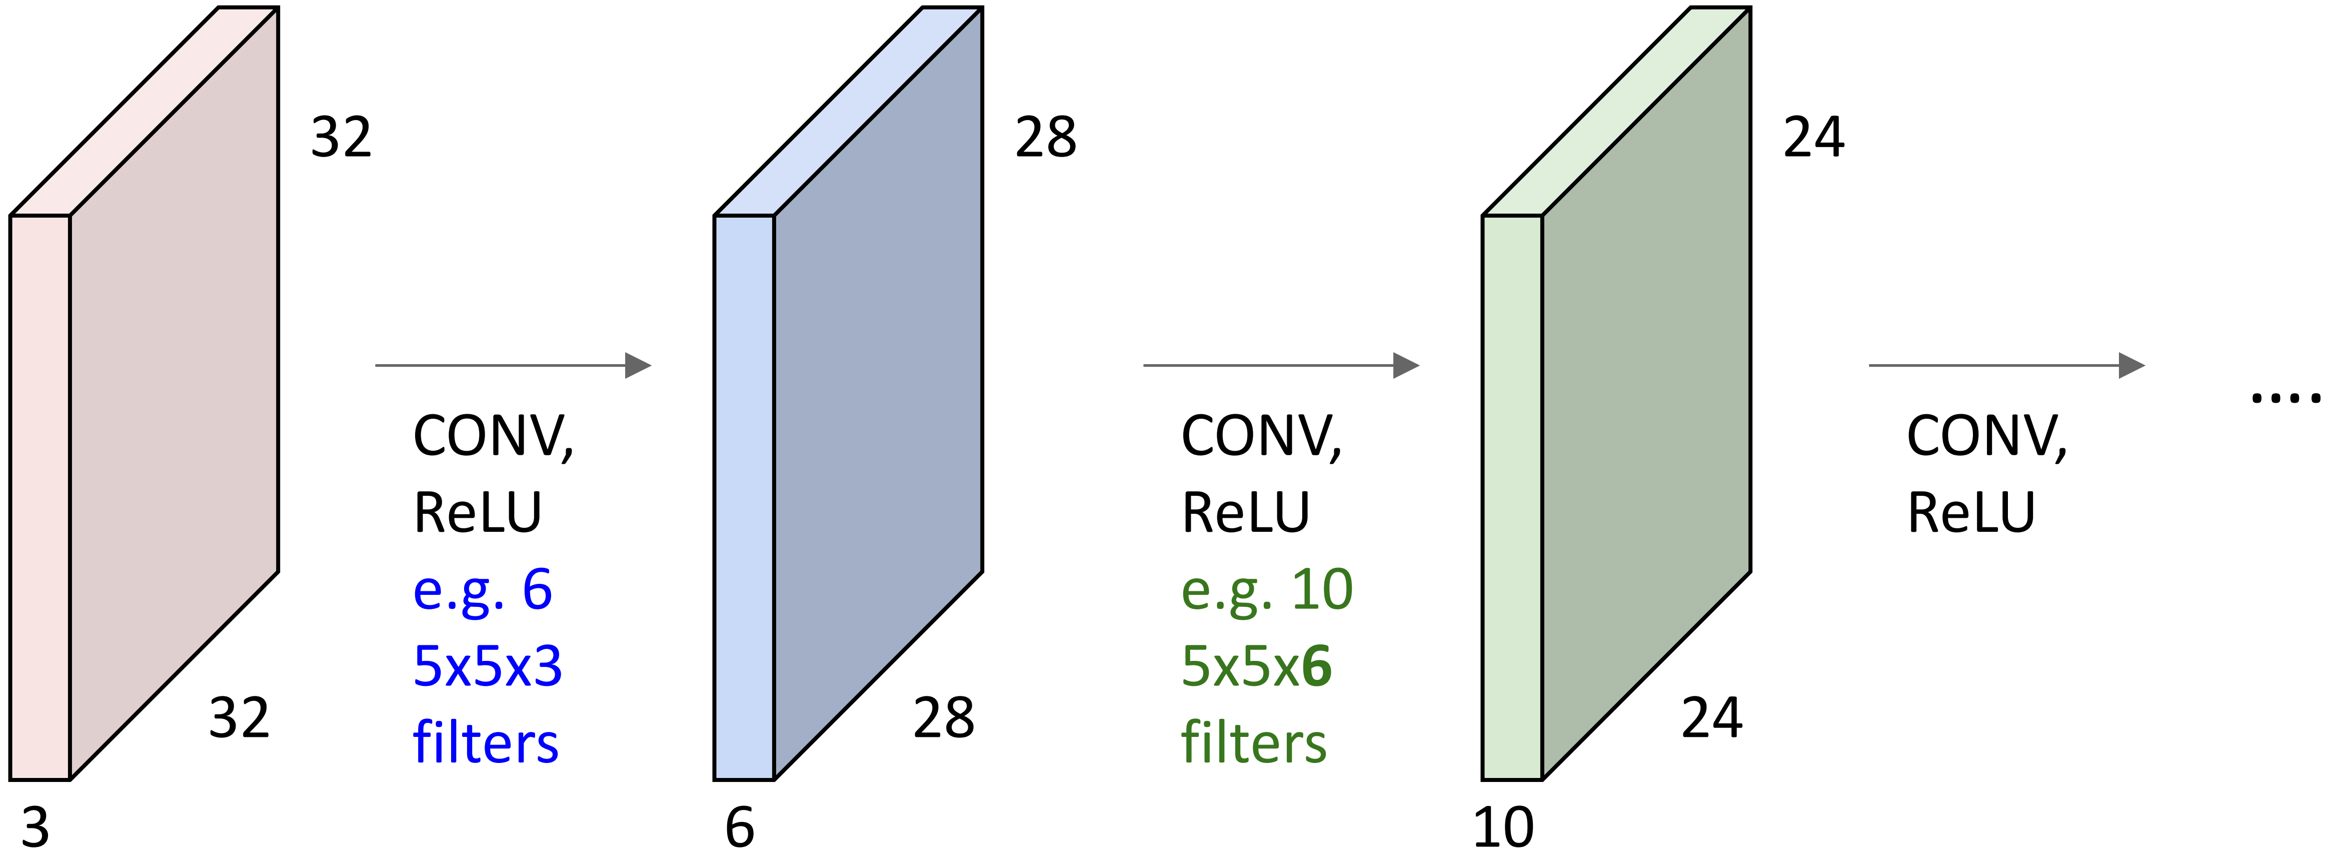

ConvNets are a sequence of convolutional layers, interspersed with activation functions

Pooling Layer

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

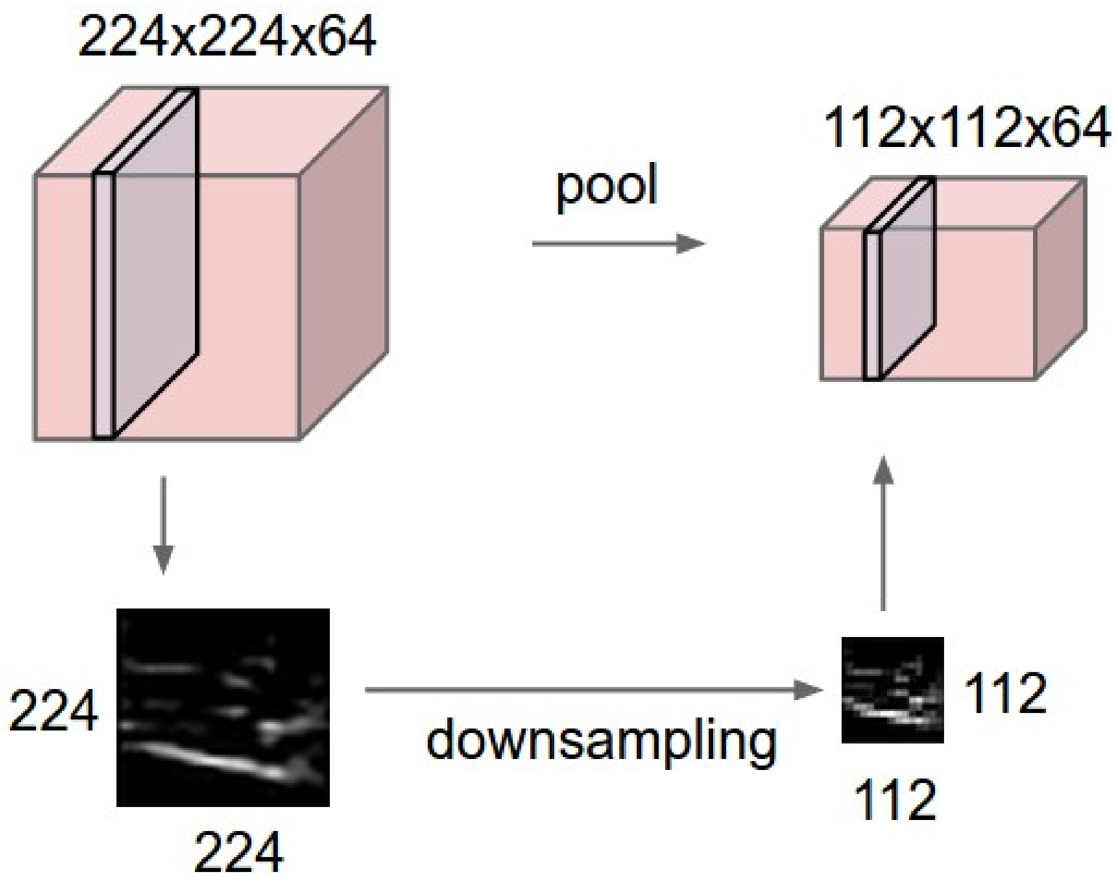

Data Abstraction Across Layers

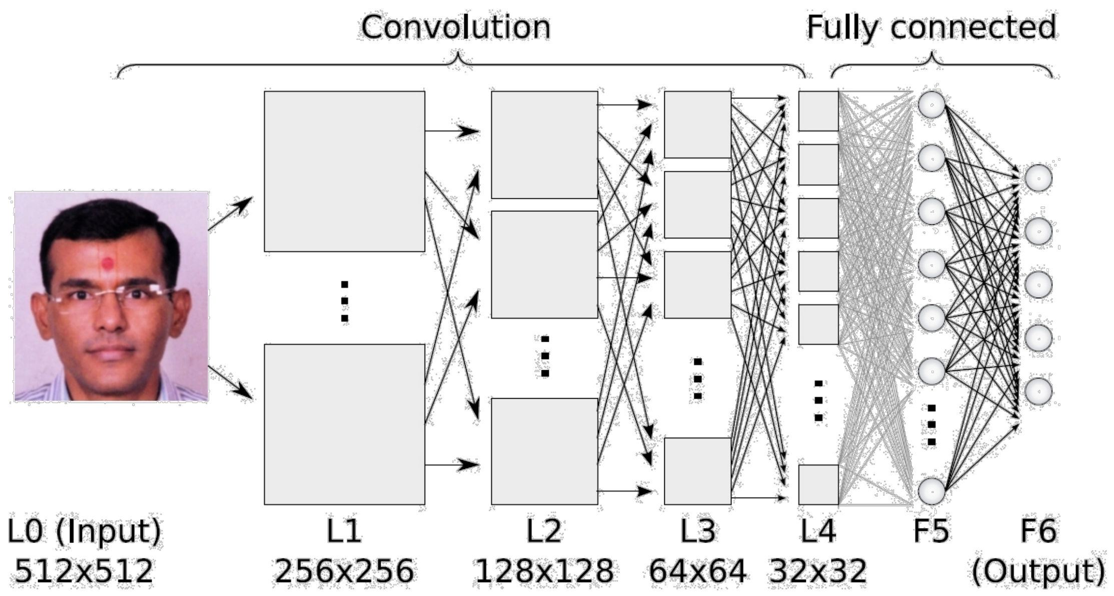

[Makwana, 2016]

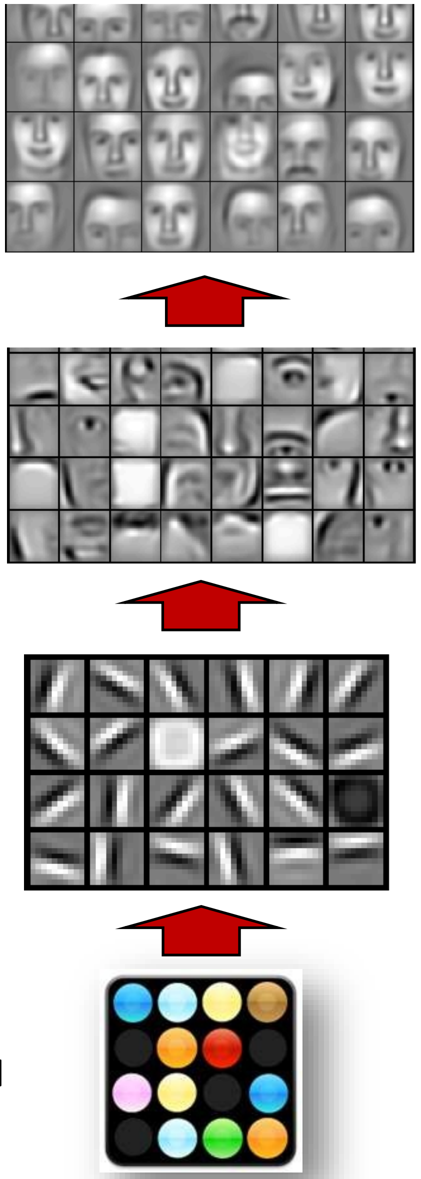

[Lee et al., 2009]

Data-driven feature learning at ascending abstraction layers

“Vertical” networks

Success Made Possible By ….

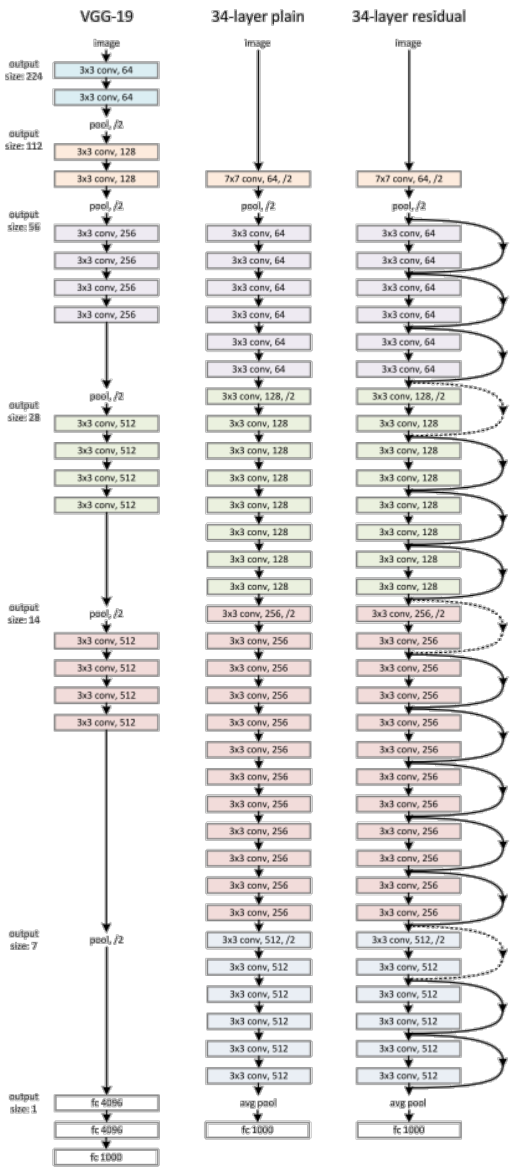

Novel deep architectures
[He et al., 2015]

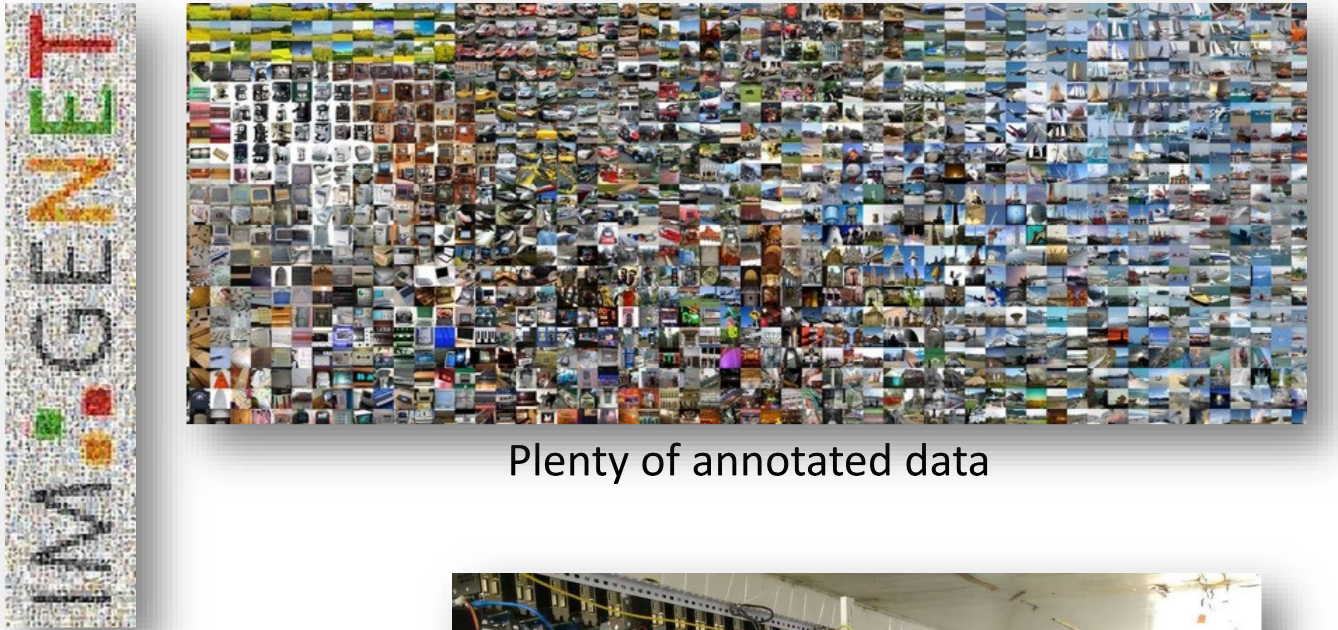

Plenty of annotated data

Lots of computing power

Transformers

Attention Is All You Need

Ashish Vaswani[*]
Google Brain
avaswani@google.com

Noam Shazeer[*]
Google Brain
noam@google.com

Niki Parmar[*]
Google Research
nikip@google.com

Jakob Uszkoreit[*]
Google Research
usz@google.com

Llion Jones[*]
Google Research
llion@google.com

Aidan N. Gomez[*] [†]
University of Toronto
aidan@cs.toronto.edu

Łukasz Kaiser[*]
Google Brain
lukaszkaiser@google.com

Illia Polosukhin[*] [‡]
illia.polosukhin@gmail.com

Transformers

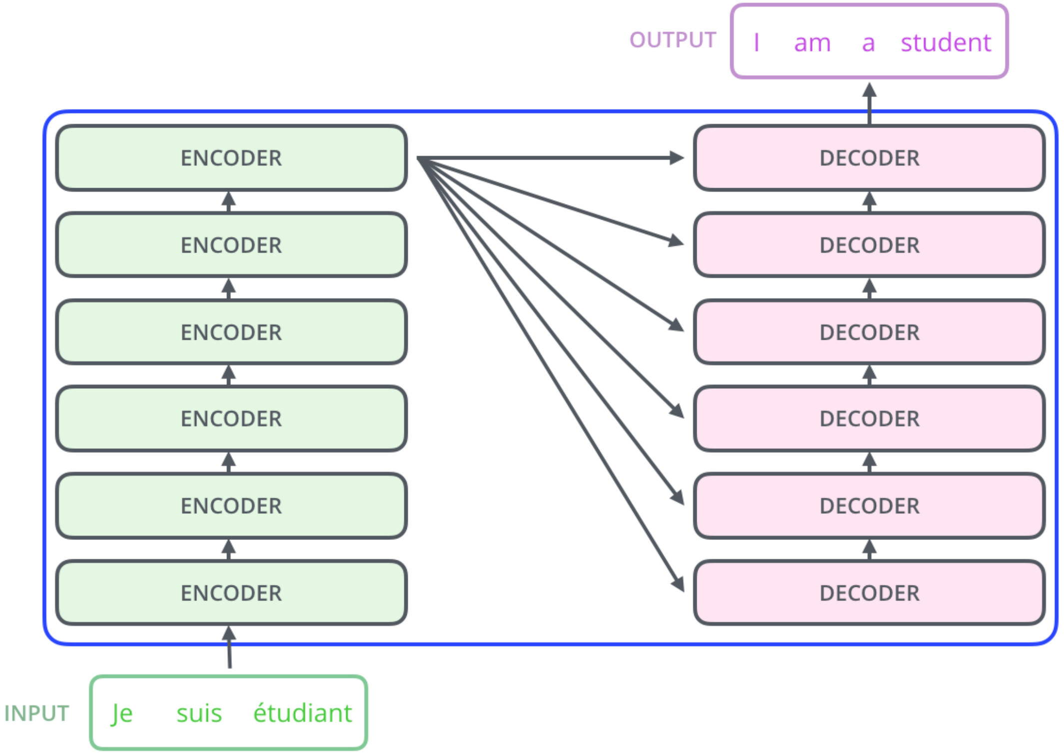

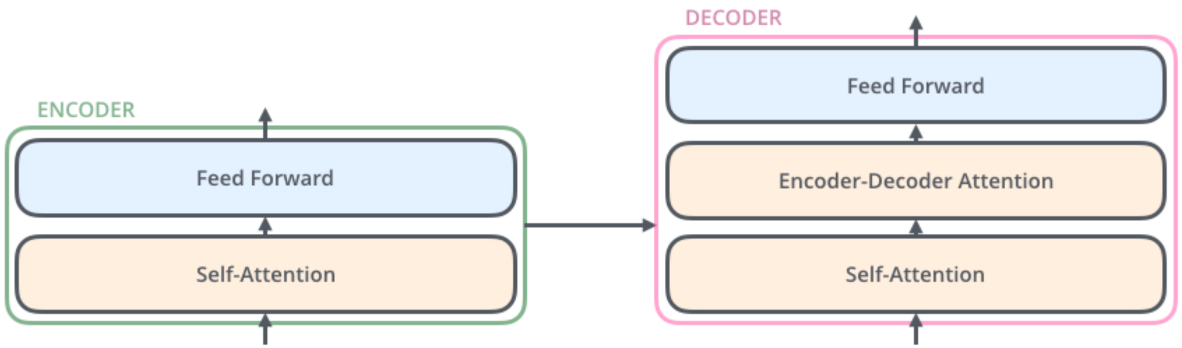

Credits: Jay Alammar's blog (http://jalammar.github.io/illustrated-transformer/)

Transformers for NLP: Attention Mechanism

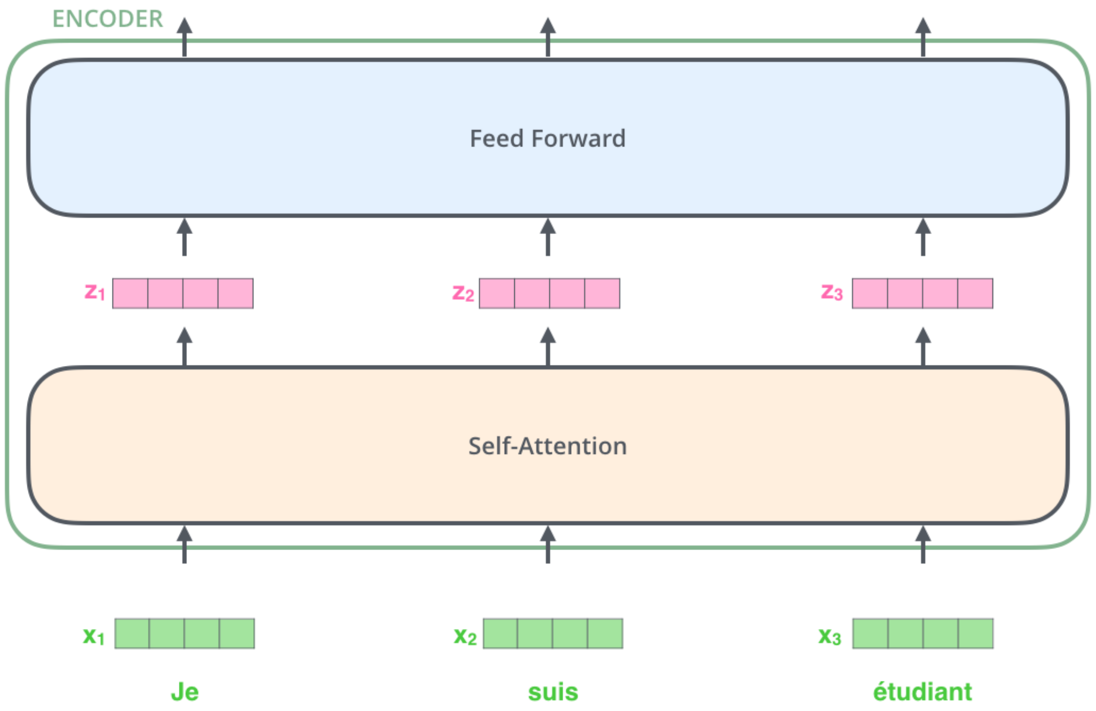

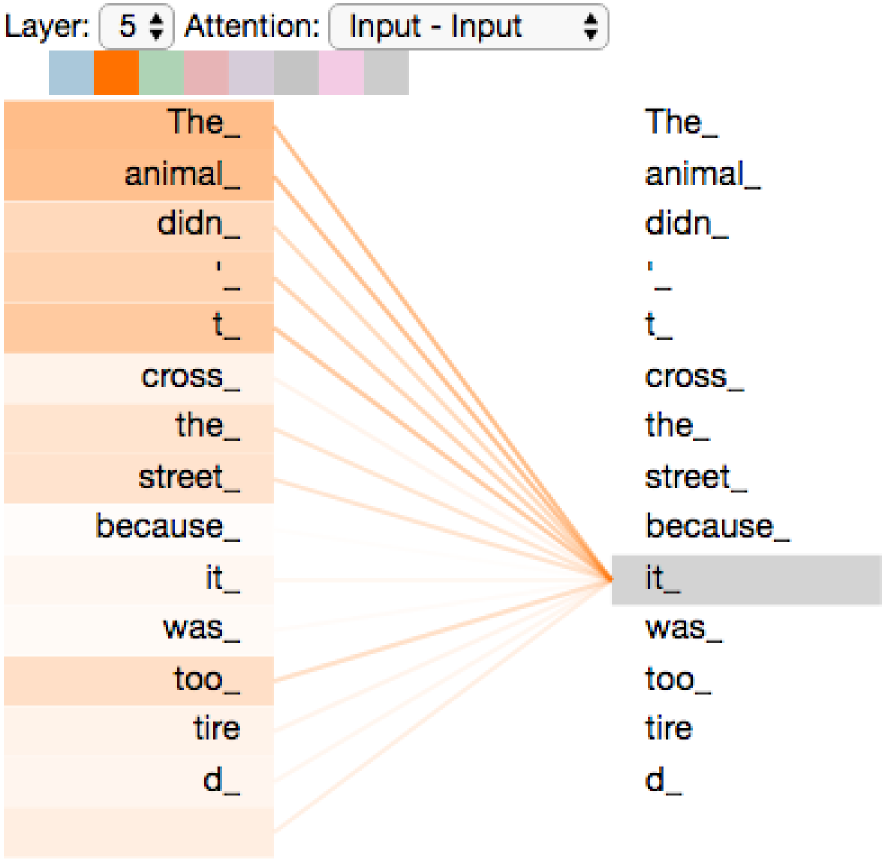

Credits: Jay Alammar's blog
(http://jalammar.github.io/illustrated-transformer/)

A Closer Look at Self-Attention

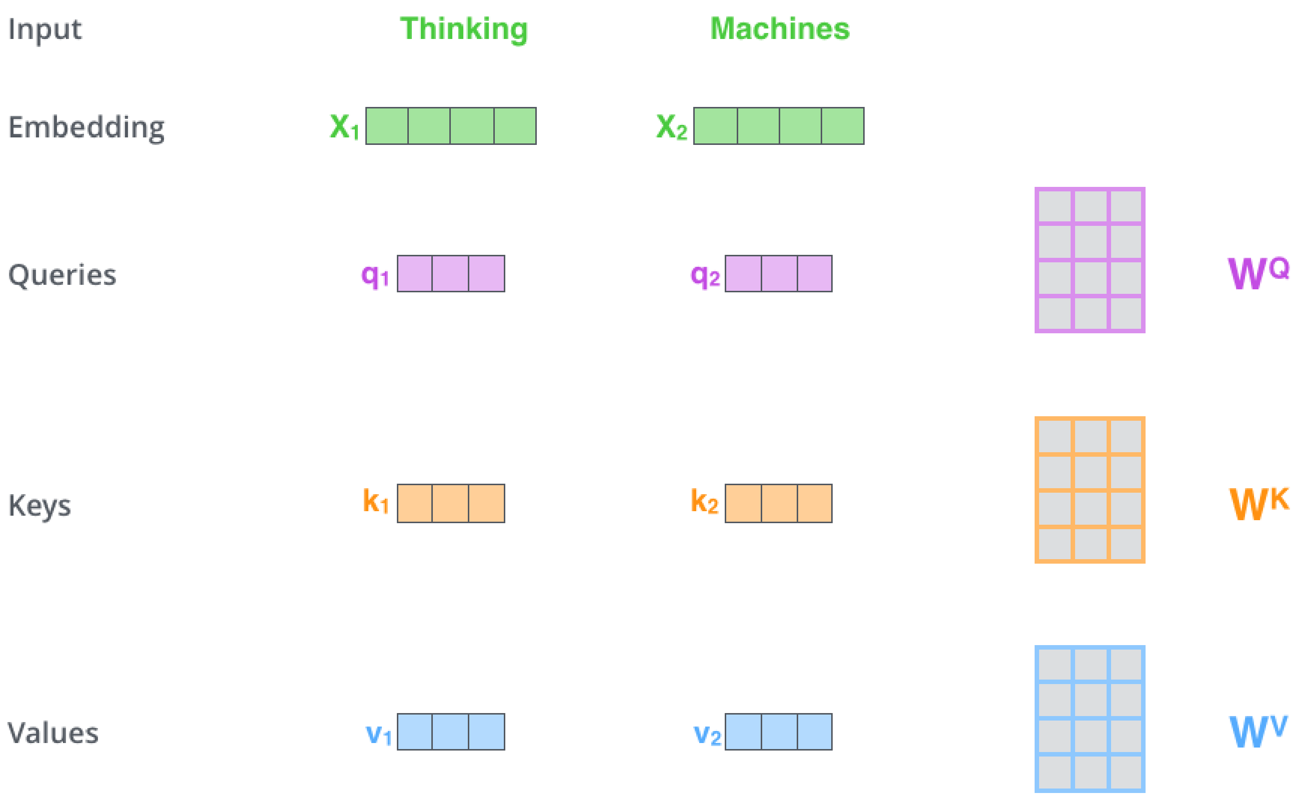

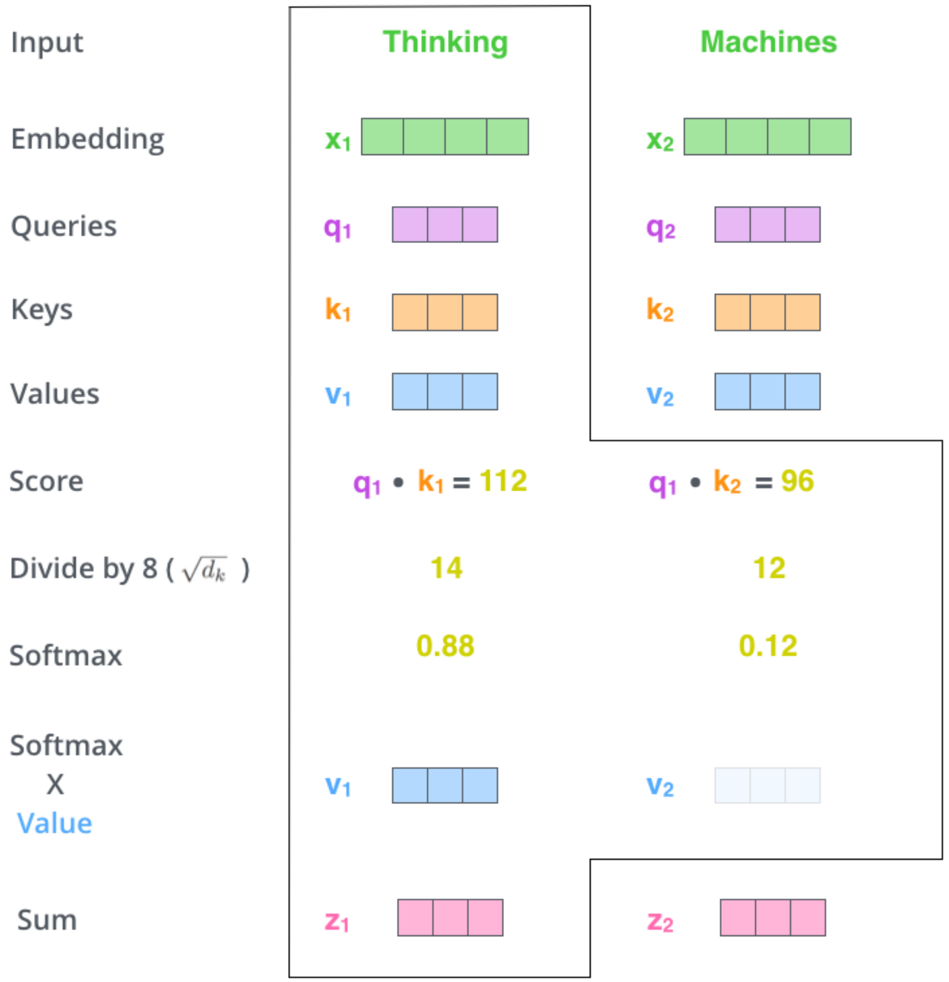

$Q = W^Q \times X$; $K = W^K \times X$; $V = X^V \times X$

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Credits: Jay Alammar's blog
(http://jalammar.github.io/illustrated-transformer/)

Visual Transformer (ViT)

Encoder: $\mathbf{c} = T_W(\mathbf{z})$
where $\mathbf{z}$ is spatial CNN features
$T_W(.)$ is the transformer encoder

<u>The general attention layer is a new type of layer (like Convolutional layer) that can be used to design new neural network architectures.</u>

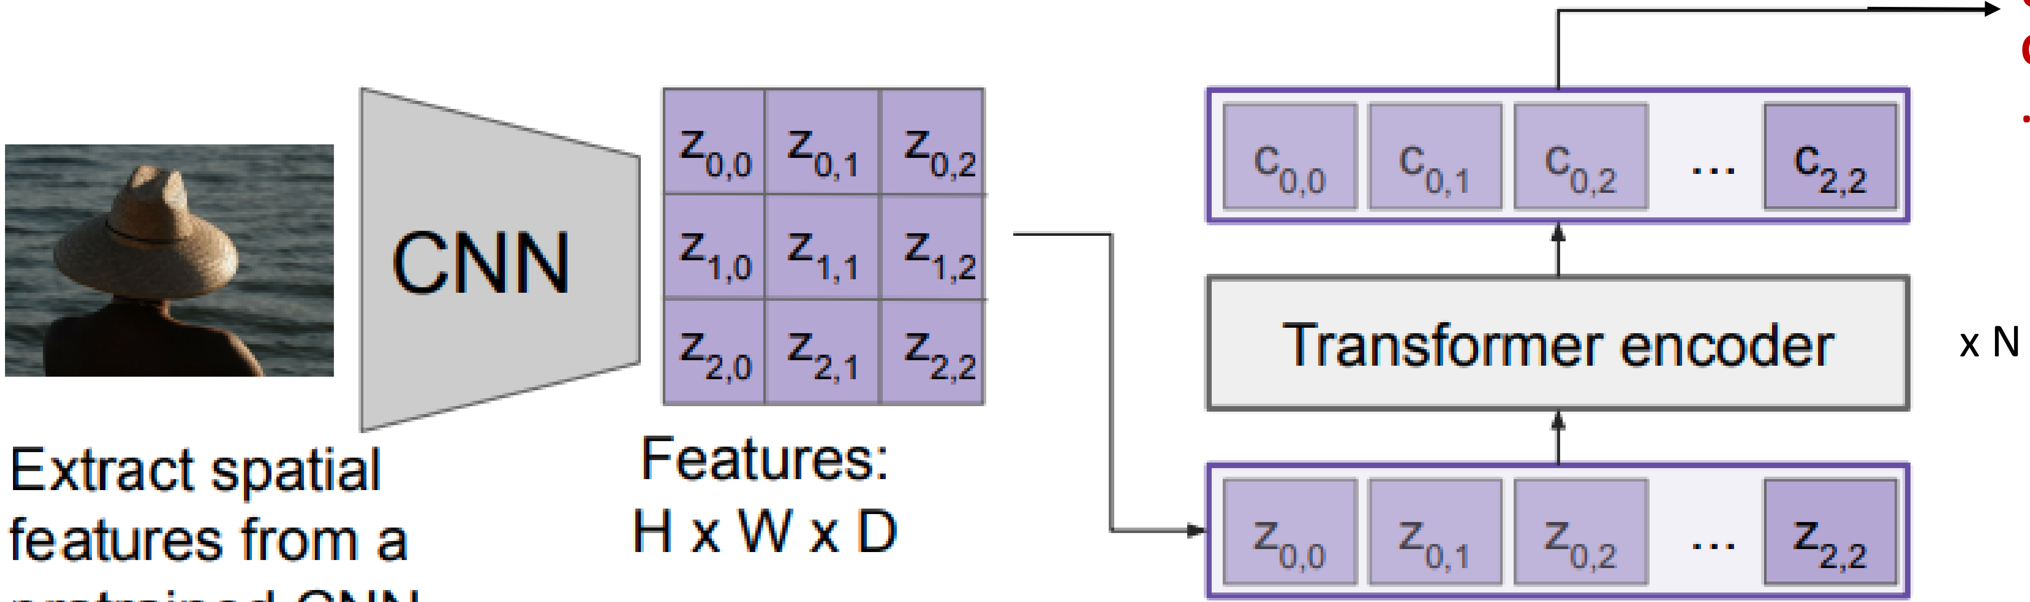

Classification
Object Detection
Caption Generation
……

Extract spatial features from a pretrained CNN

Features: H x W x D

Flattened into tokens, but with positional encodings

Image Credits: CS 231n

Deep Nets are Everywhere

Classification

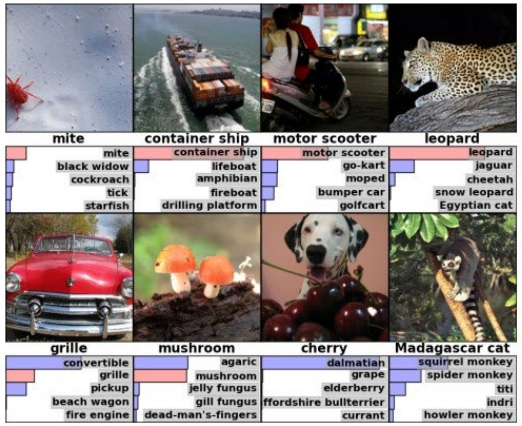

Retrieval

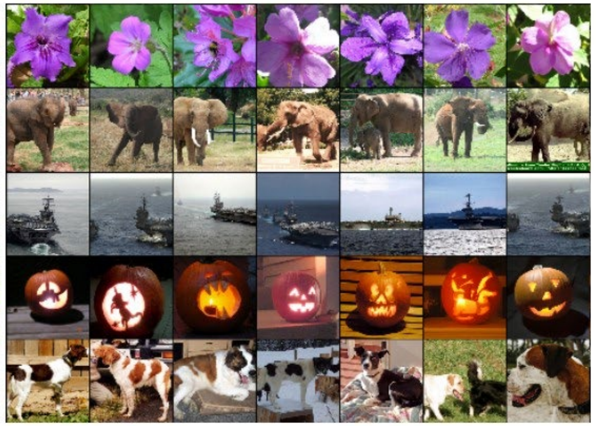

Detection

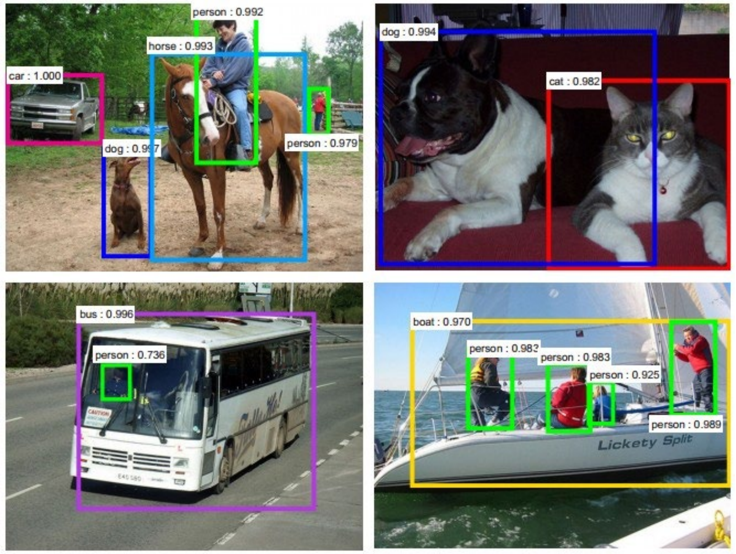

Text processing

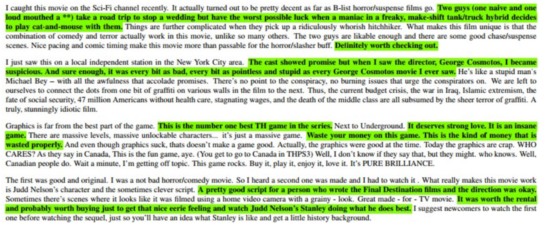

Image captioning

Segmentation

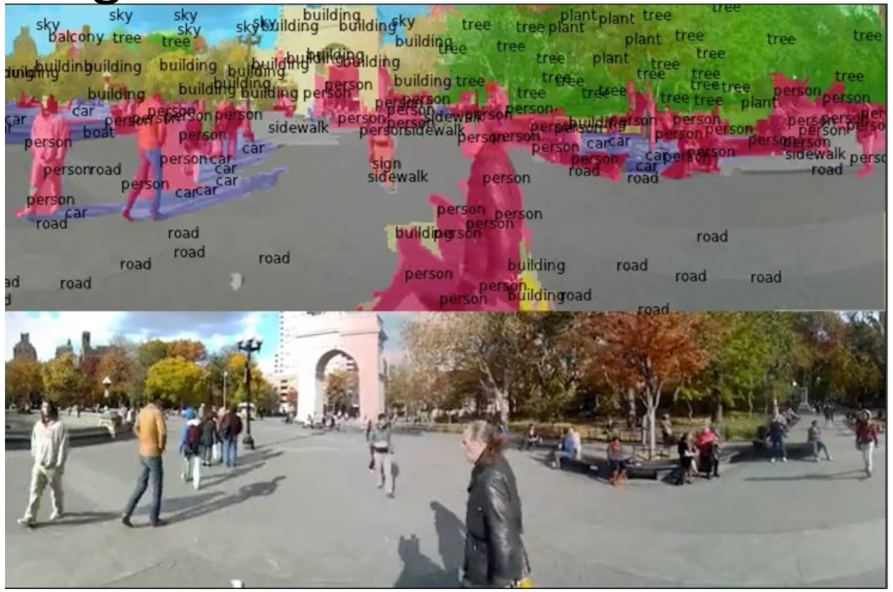

Control

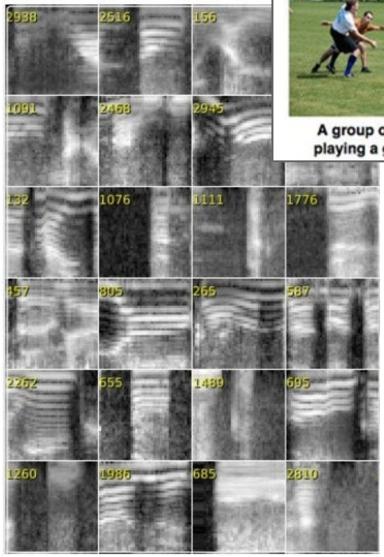

Speech processing

Learning on Irregular Data

Multiple Representations for 3D Geometric Data

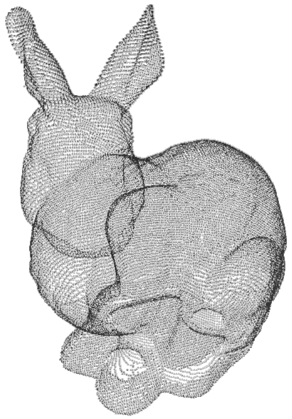

Point Cloud

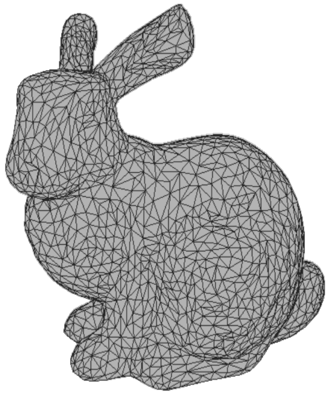

Mesh

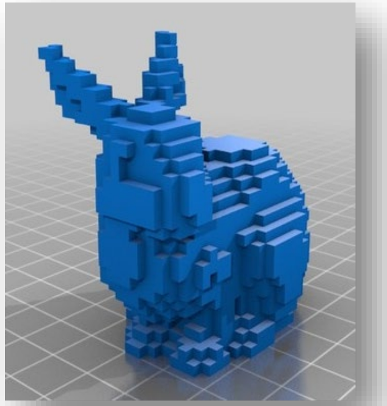

Volumetric

Projected View

RGB(D)

...

These are irregular representations

Regular vs Irregular

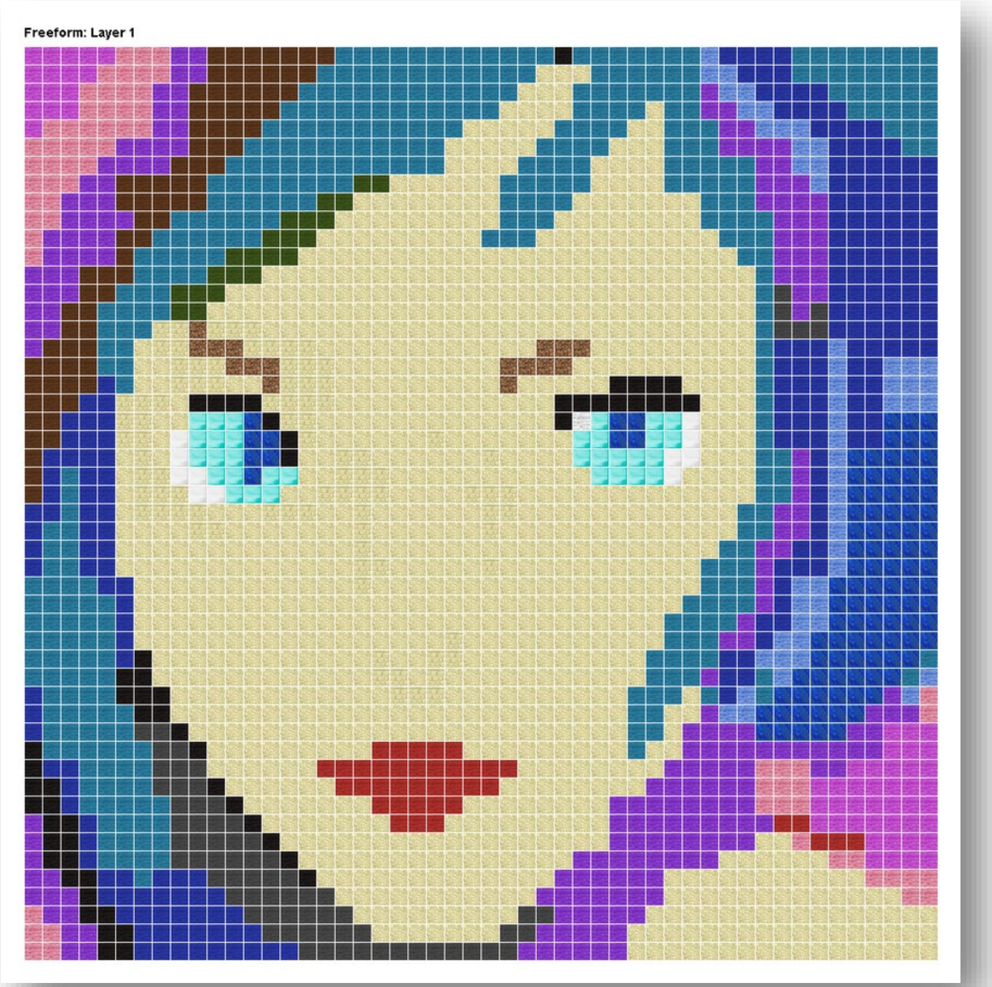

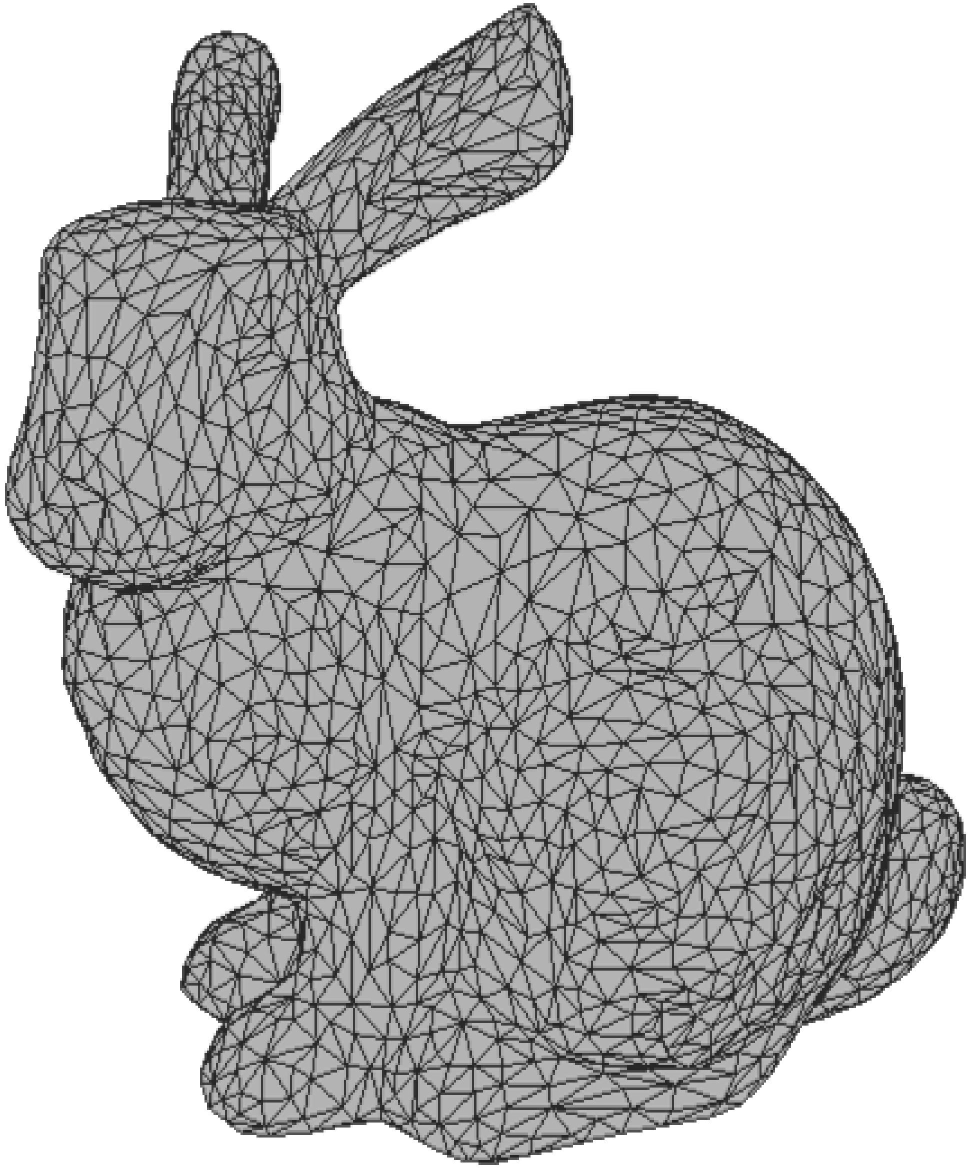

Deep Learning for 3D Geometry Analysis

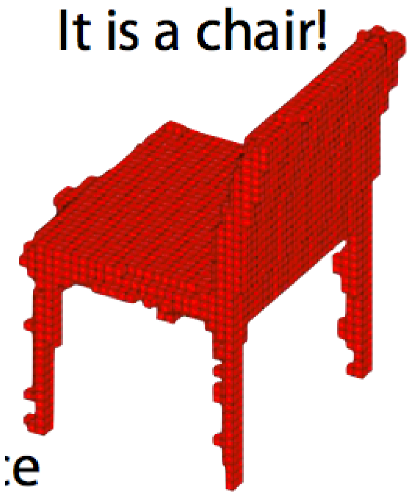

Classification

Segmentation/Parsing (object/scene)

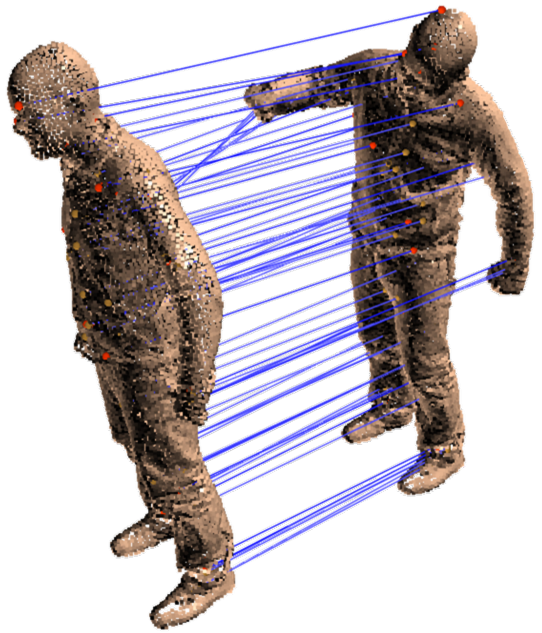

Correspondences

Convolutional Networks

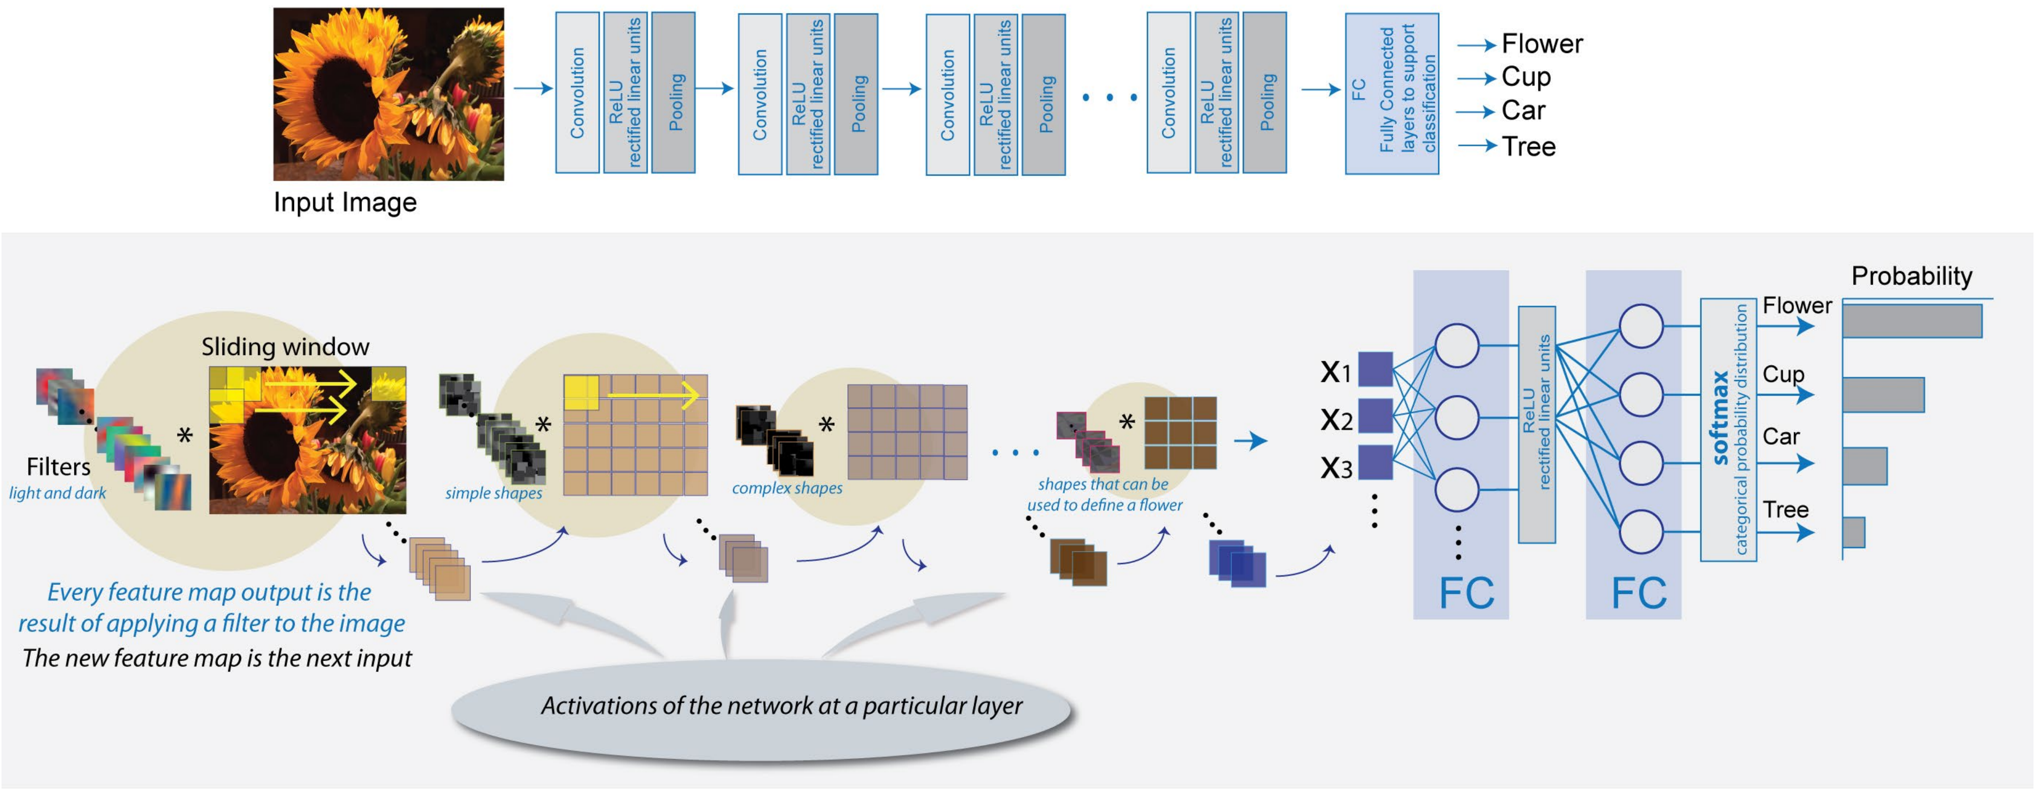

Convolutions require regular array type inputs for weight sharing and other optimizations

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

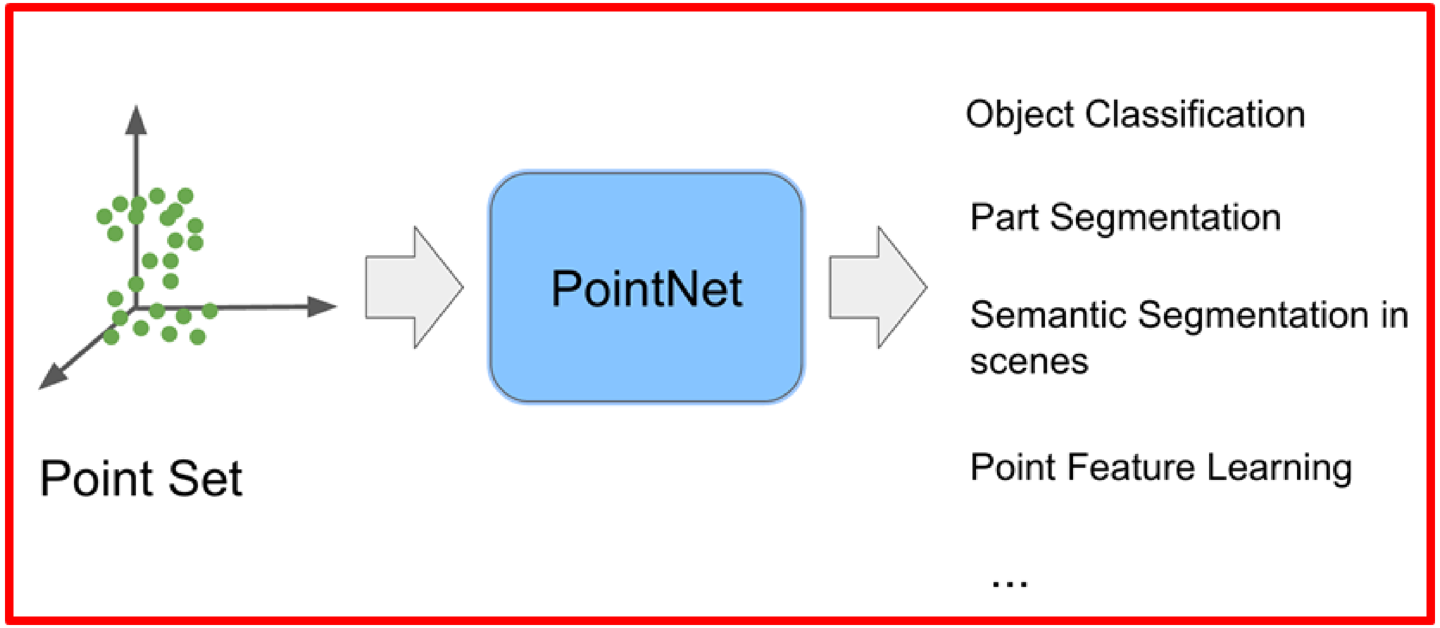

The Graph View of Data

A Graph View of Data

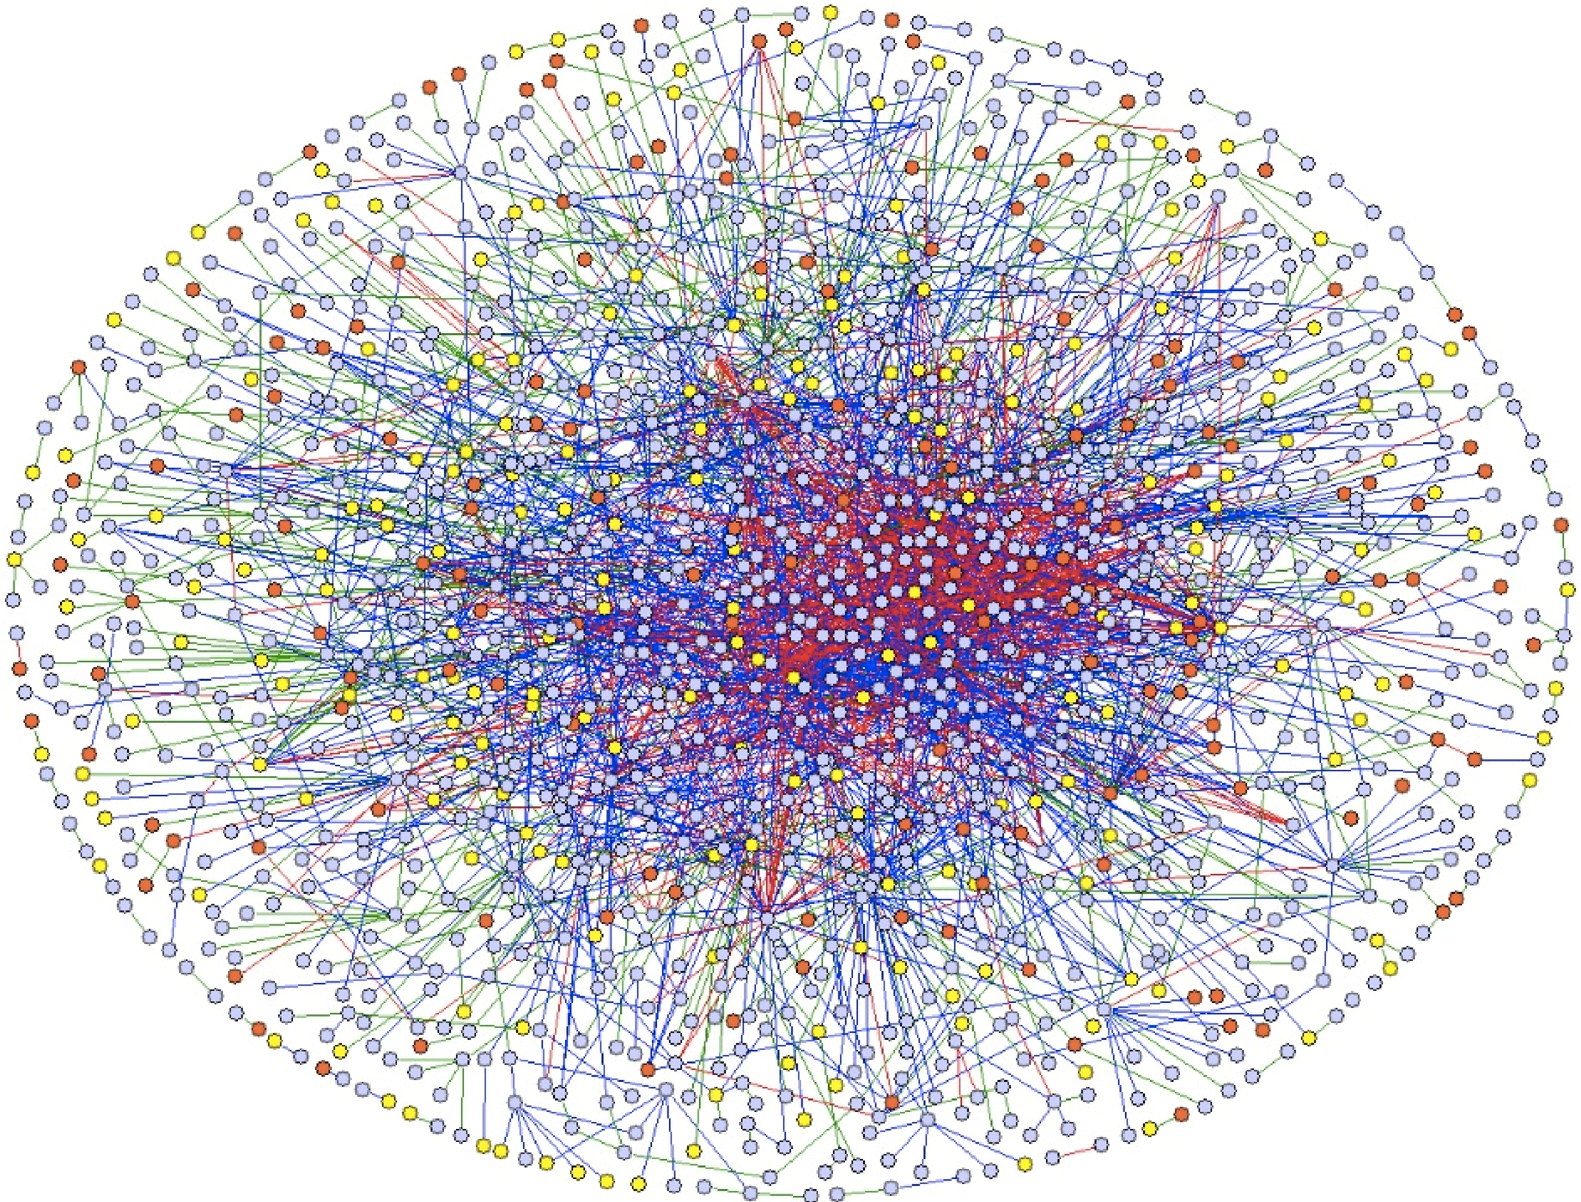

Graphs and Simplicial Complexes

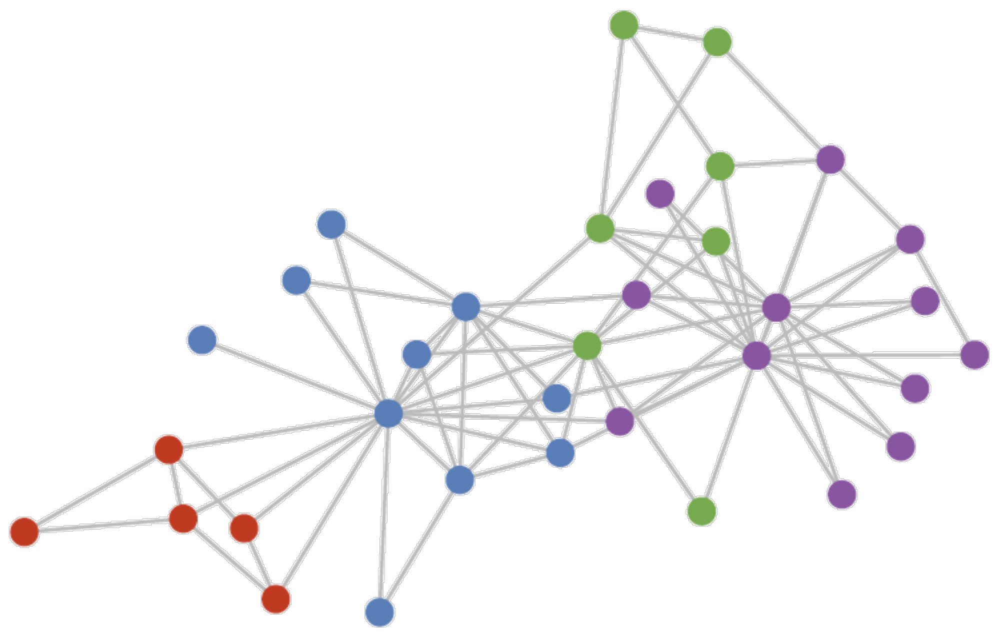

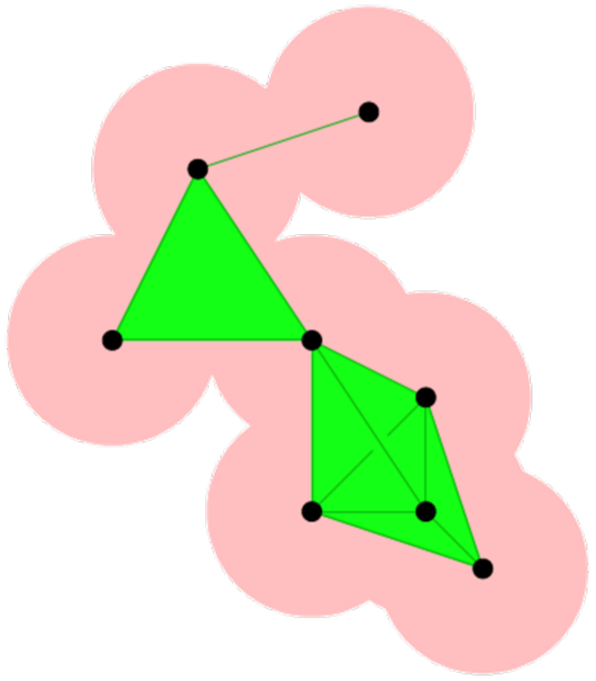

Spectral Methods in Graph Theory

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

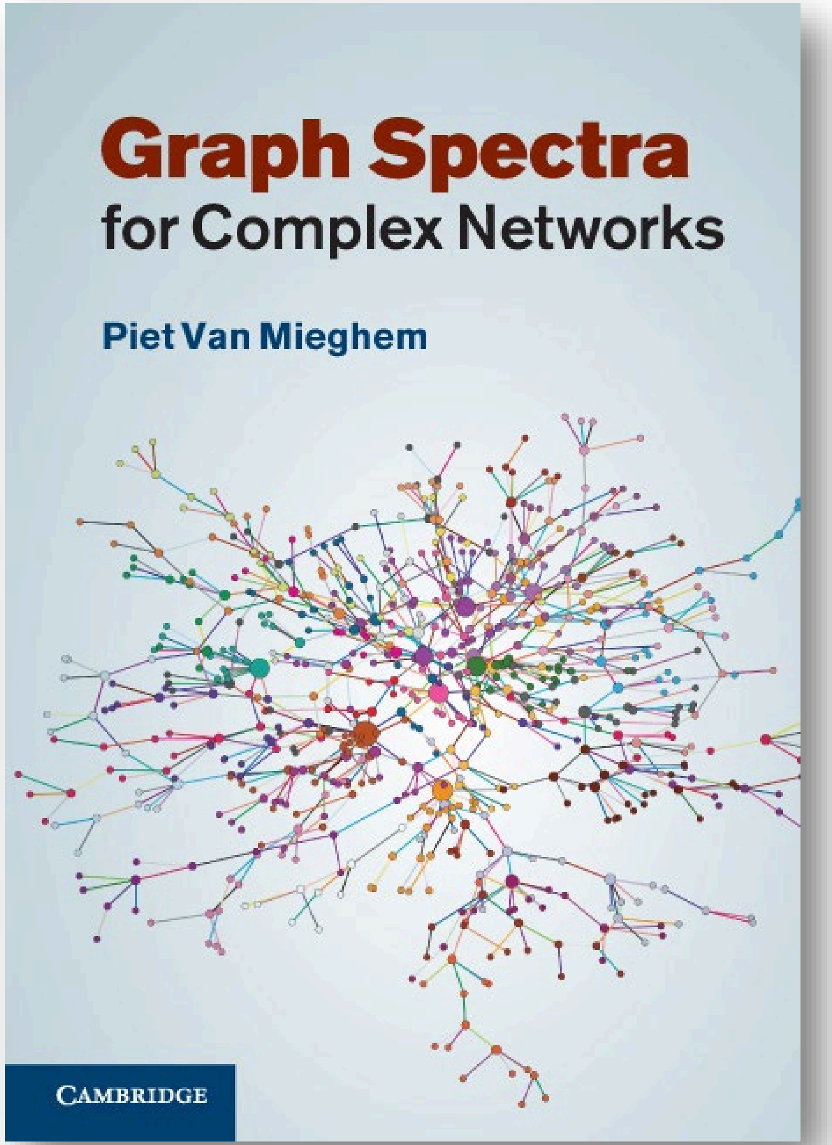

Multidimensional Scaling (MDS)

- A “distance preserving” embedding of the data into a Euclidean space
  - Sometimes distances are observed directly (e.g., similarity ratings)
  - Sometimes they can be calculated from a data table (e.g., Euclidean distances, correlations)

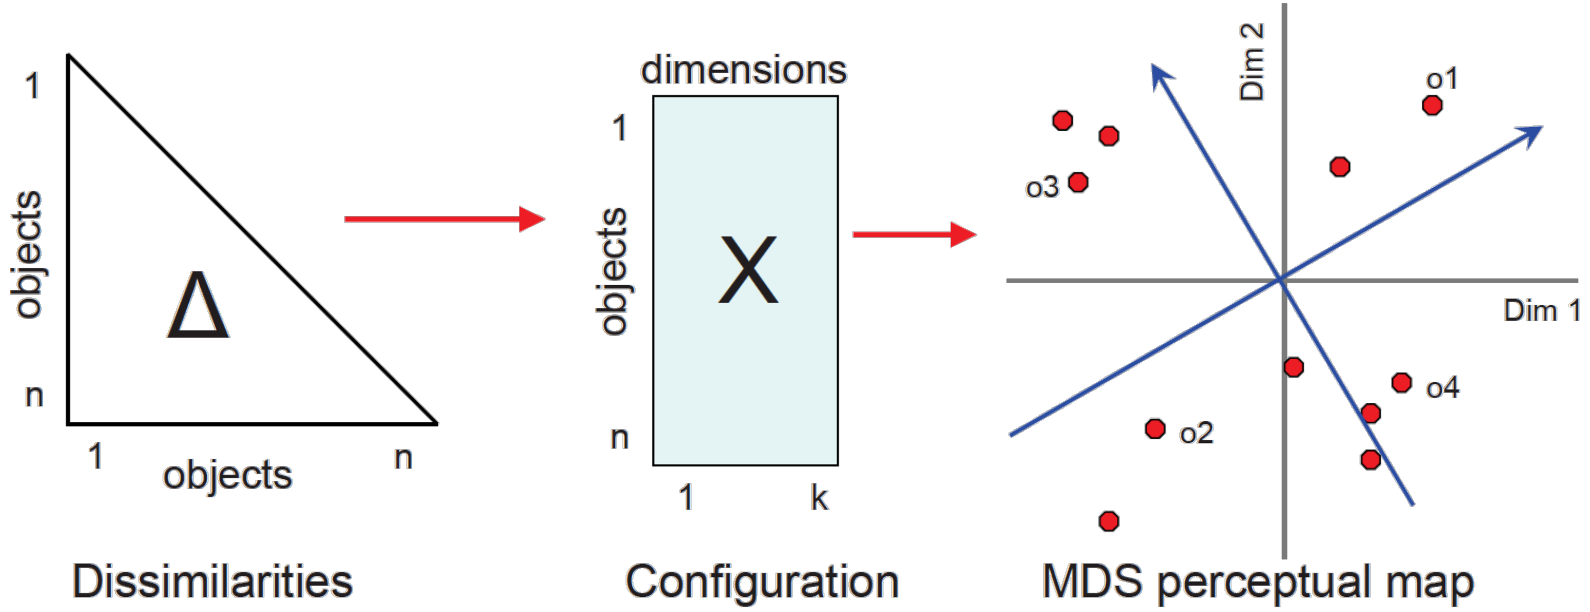

Topological Data Analysis (TDA)

From Data to Algebraic Objects

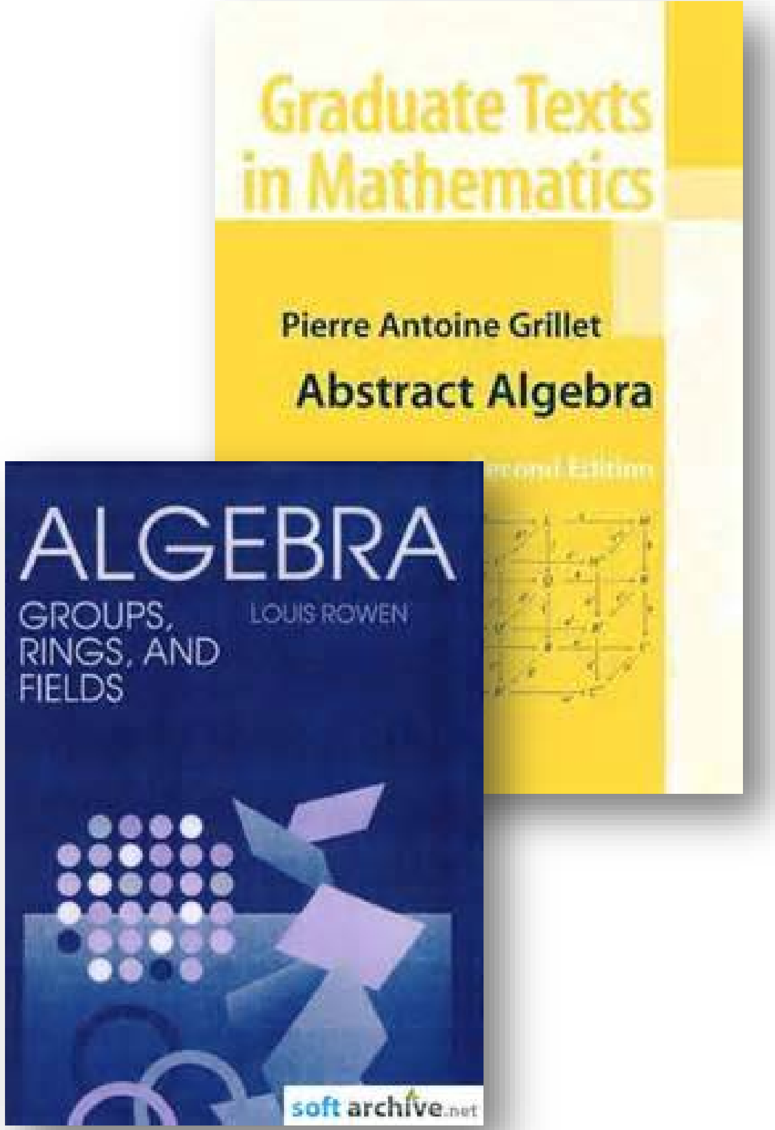

Algebraic entities as data descriptors

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 …

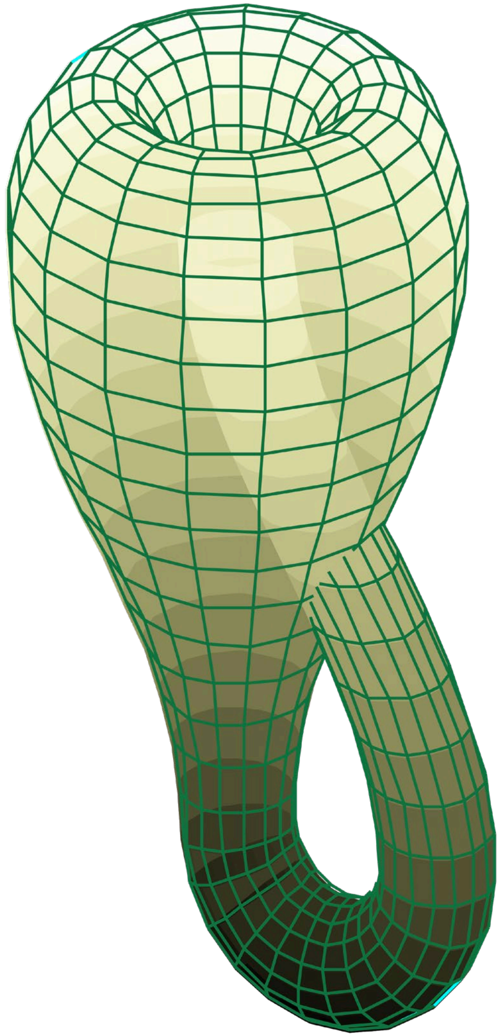

In TDA, Sampled Spaces: “The Shape of Data”

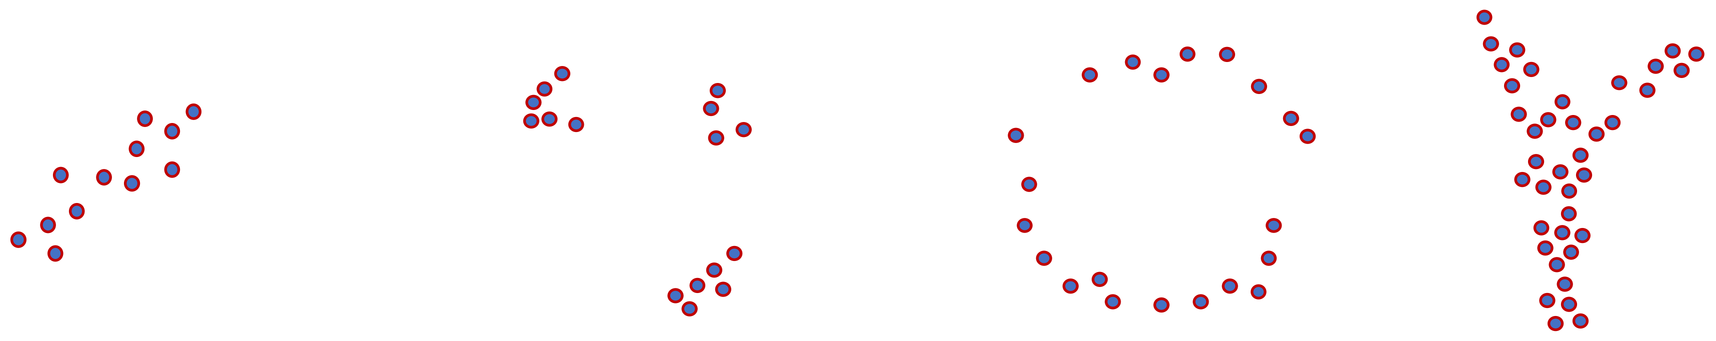

Regression Cluster Loop Flared

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.

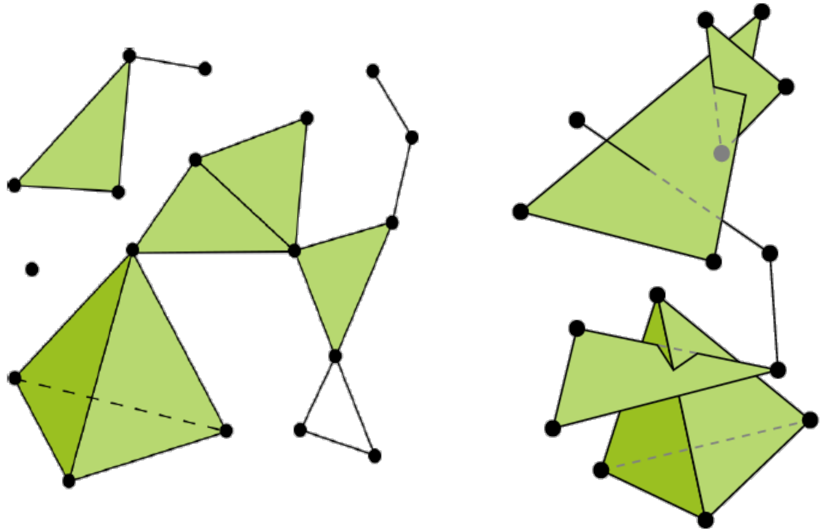

(left) an example (right) a non example

Complexes via Proximity

Must choose which simplices to introduce

Čech

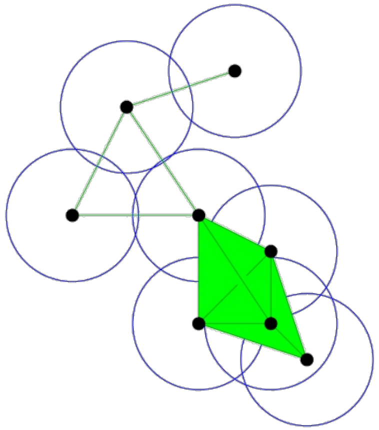

Alpha

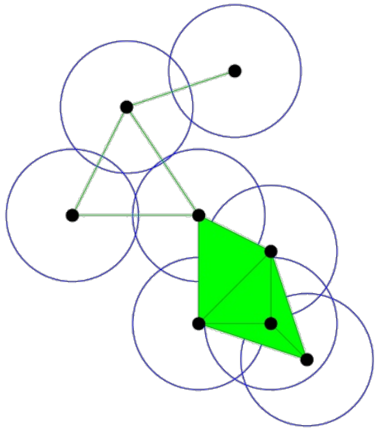

Rips

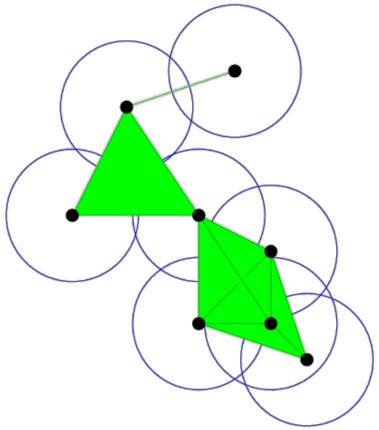

Combinatorial complexes provide discrete representations of the underlying space

Algebraic Topology and Topology Inference

- Unlike geometry, topology studies the global structure of spaces
- Getting to this structure via algebra

algebraic topology

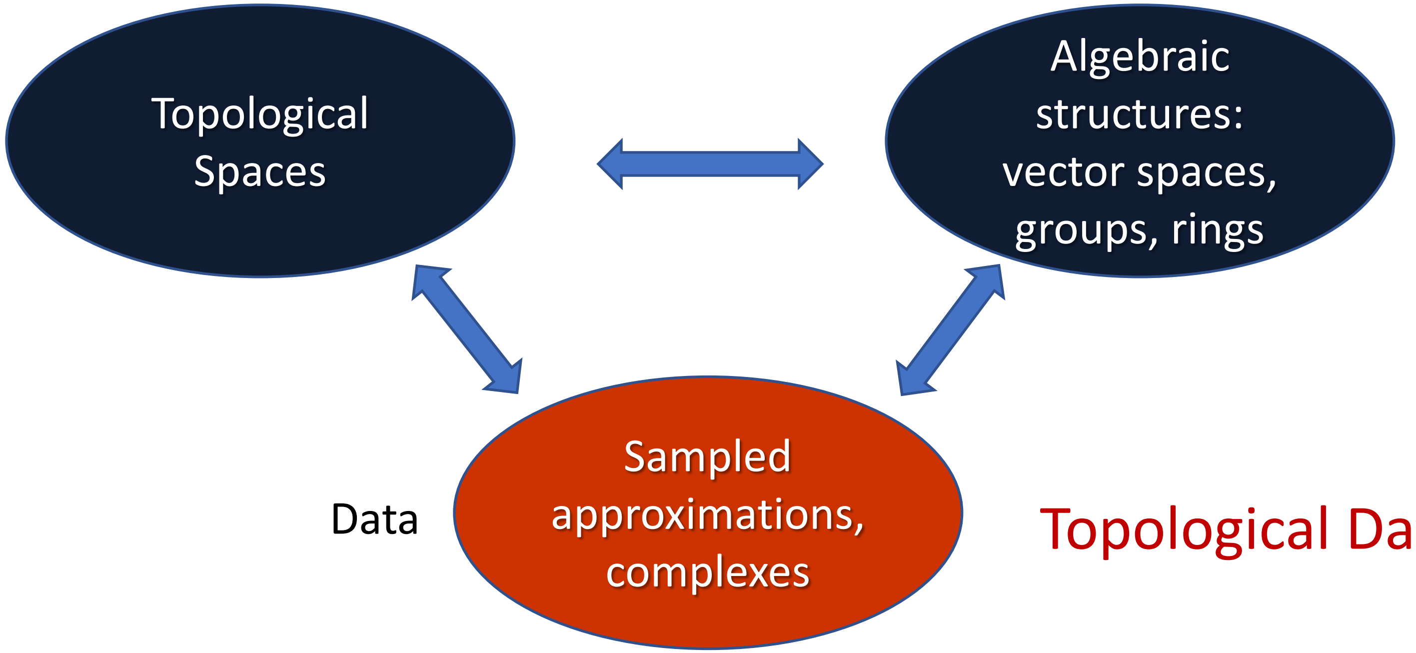

Topological Data Analysis (TDA)

Algebraic Topology: Simplicial Homology

- The kth homology group is

$$\mathsf{H}_k = \mathsf{Z}_k / \mathsf{B}_k = \ker \partial_k / \operatorname{im} \partial_{k+1}.$$

- If $z_1 = z_2 + \mathsf{B}_k$, $z_1, z_2 \in \mathsf{Z}_k$, we say z_1 and z_2 are homologous
- $z_1 \sim z_2$.

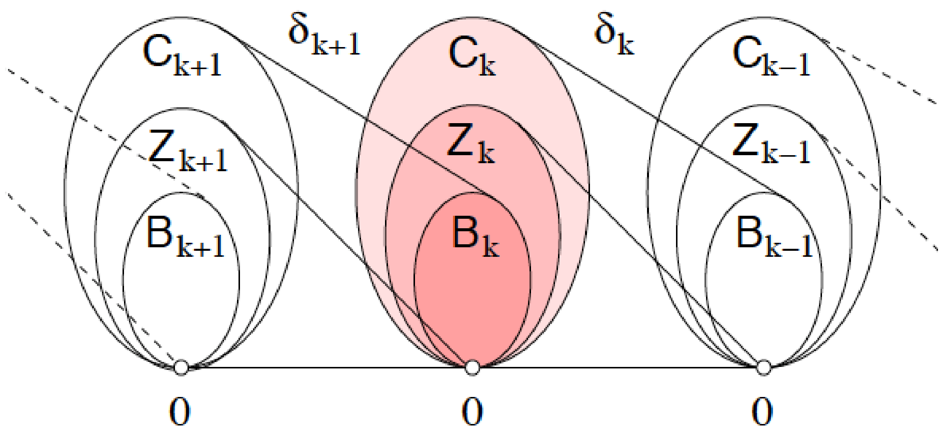

A Multiscale View of Data

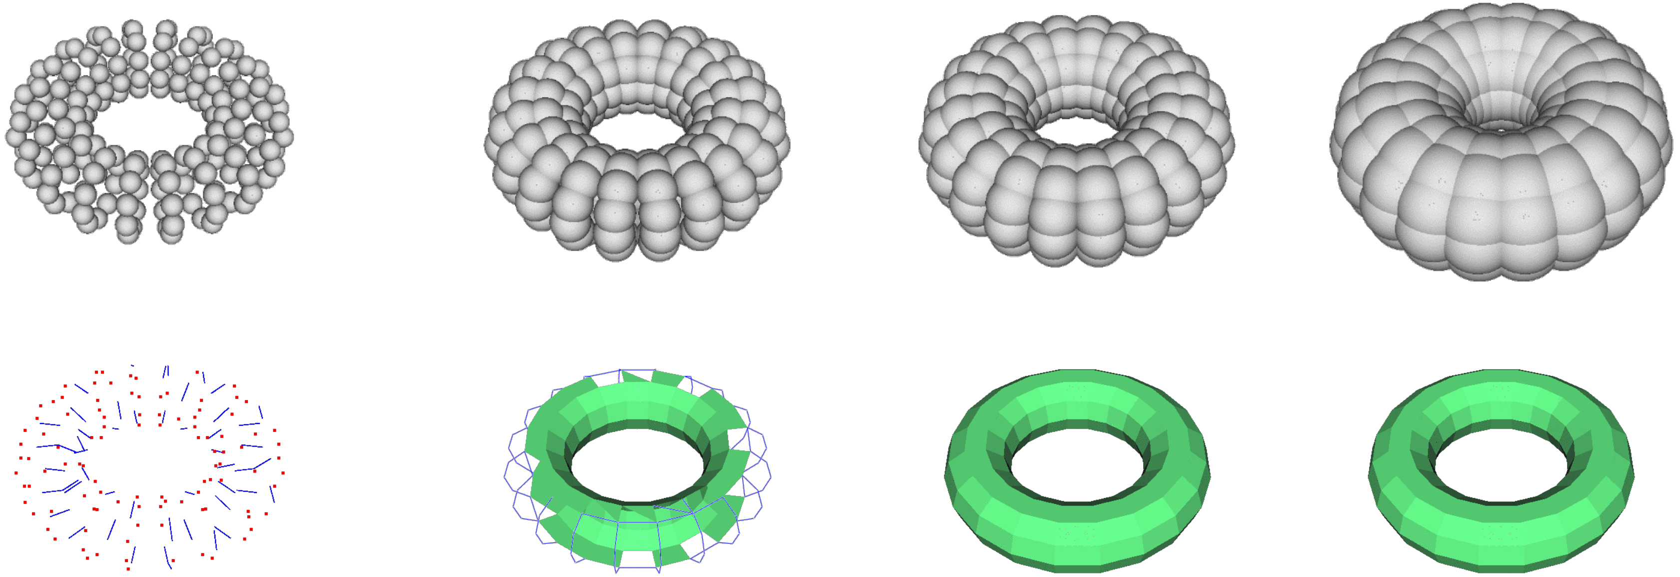

$\beta_0 = 150$
$\beta_1 = 0$

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

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

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

Persistent Homology

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

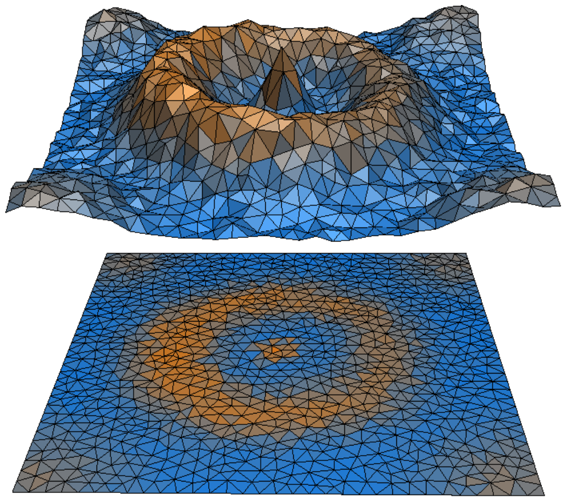

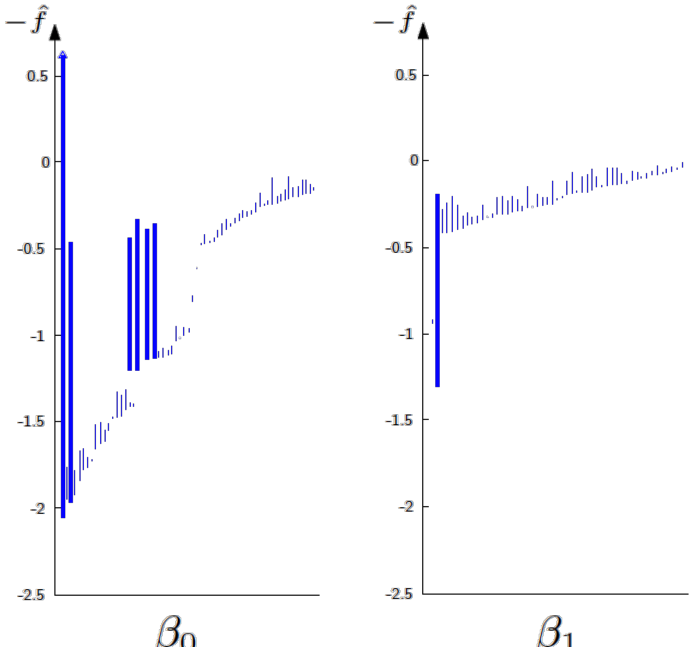

Functions on the Data as Filters

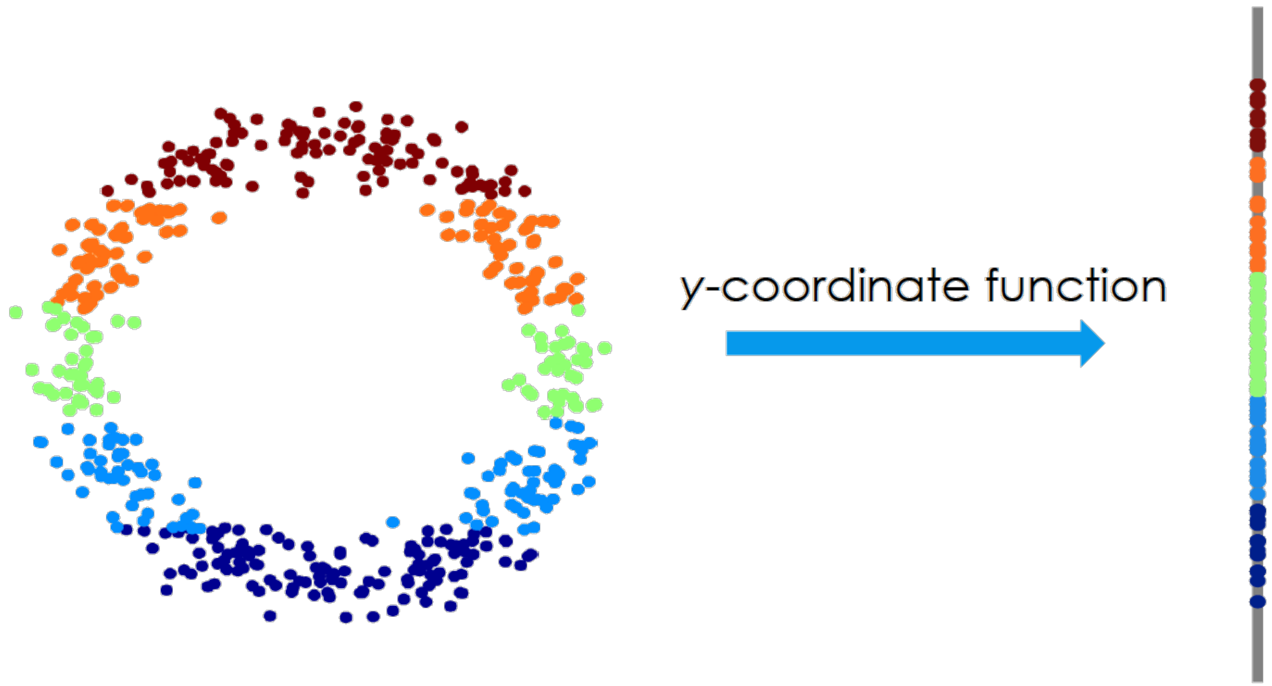

Complexes via Function Filters

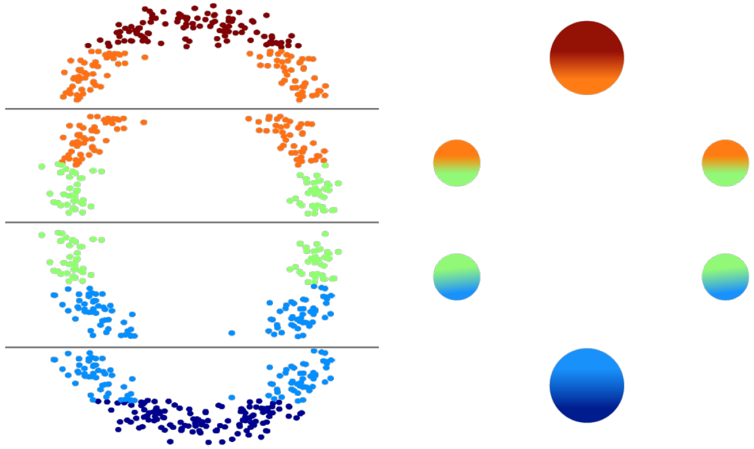

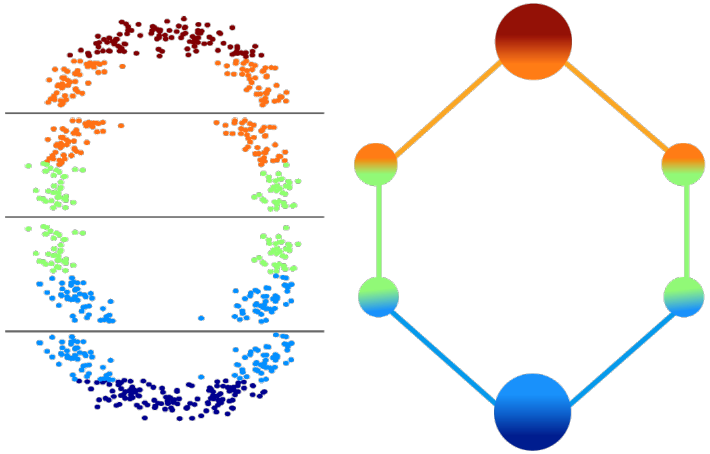

Functions as Lenses on Data

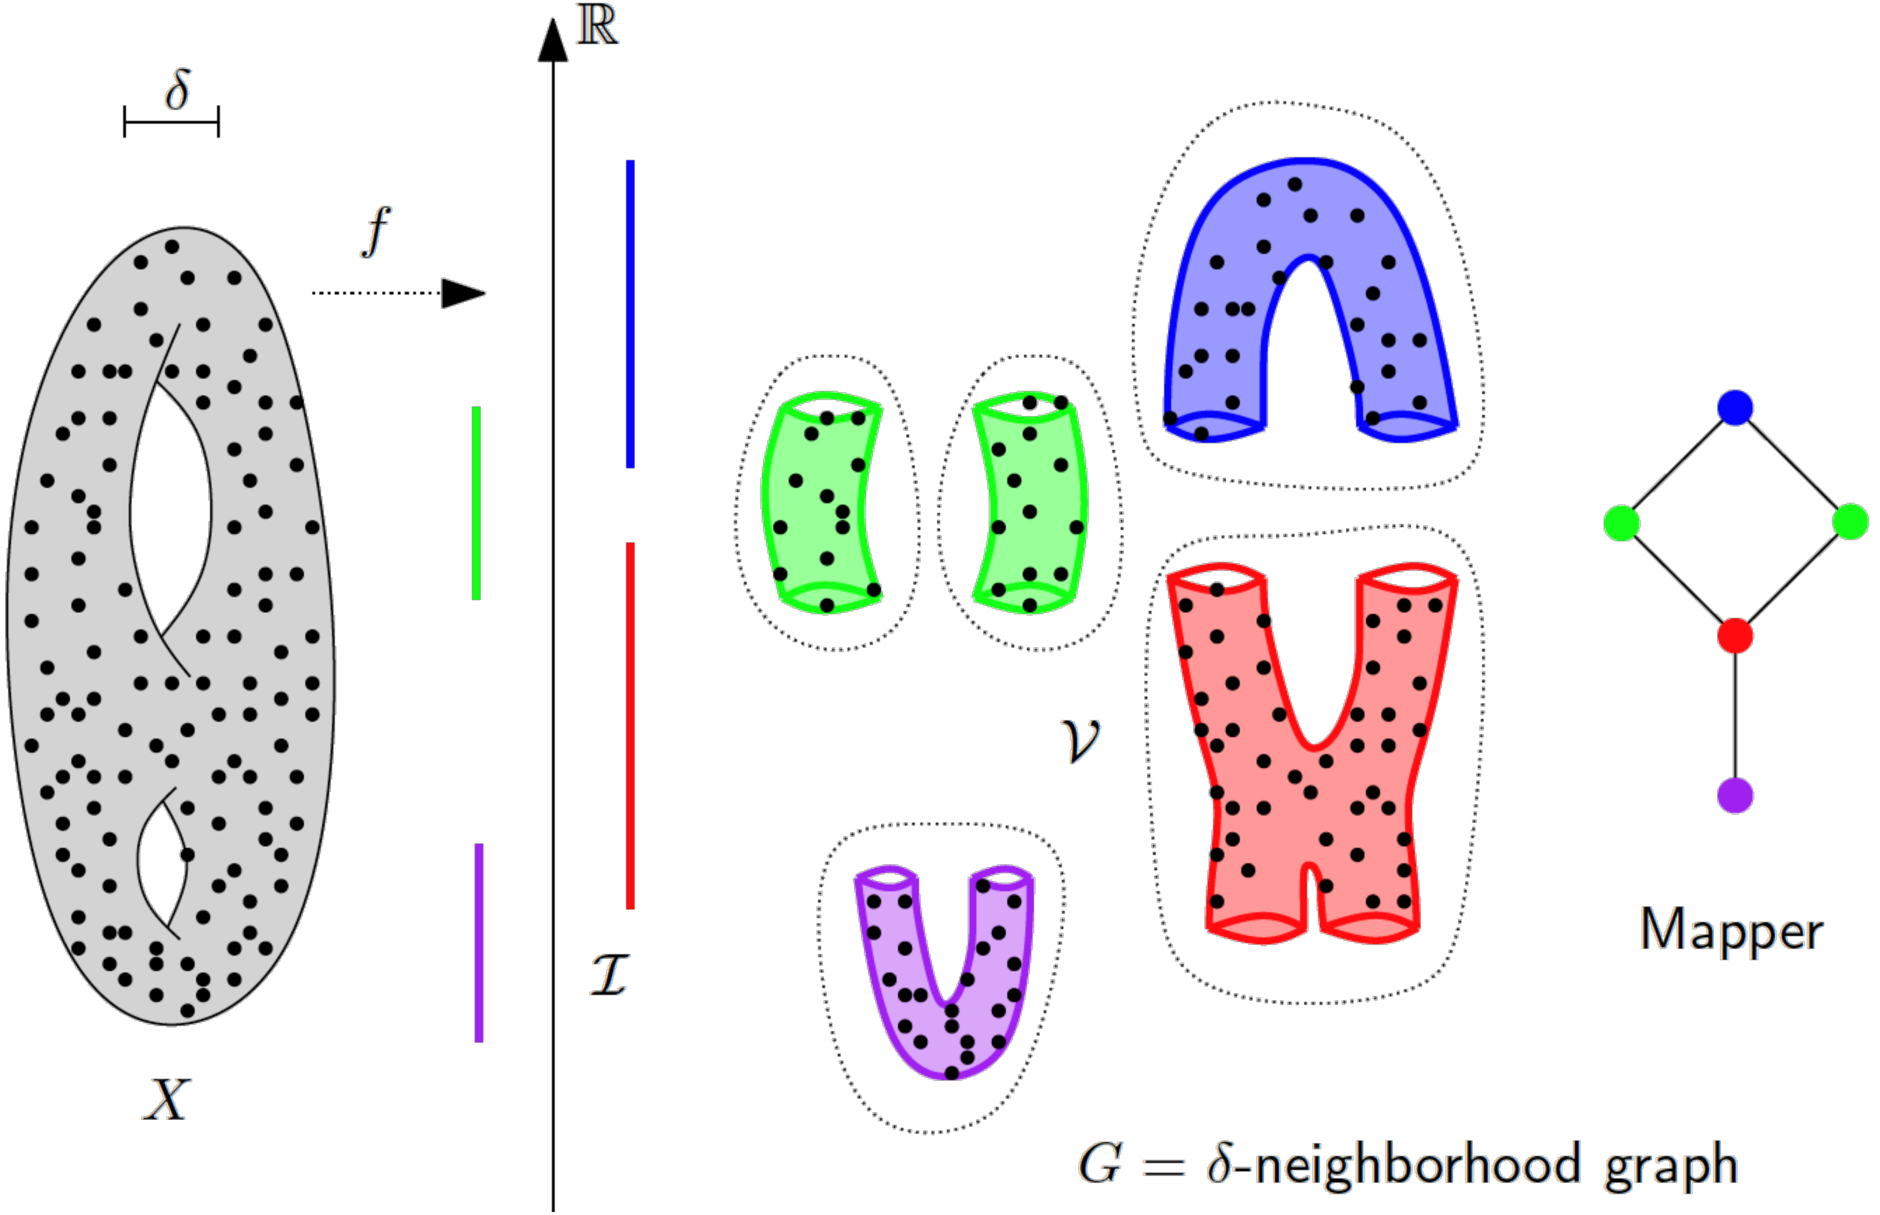

Data Representations and Latent Spaces

(Too) Many Latent Spaces

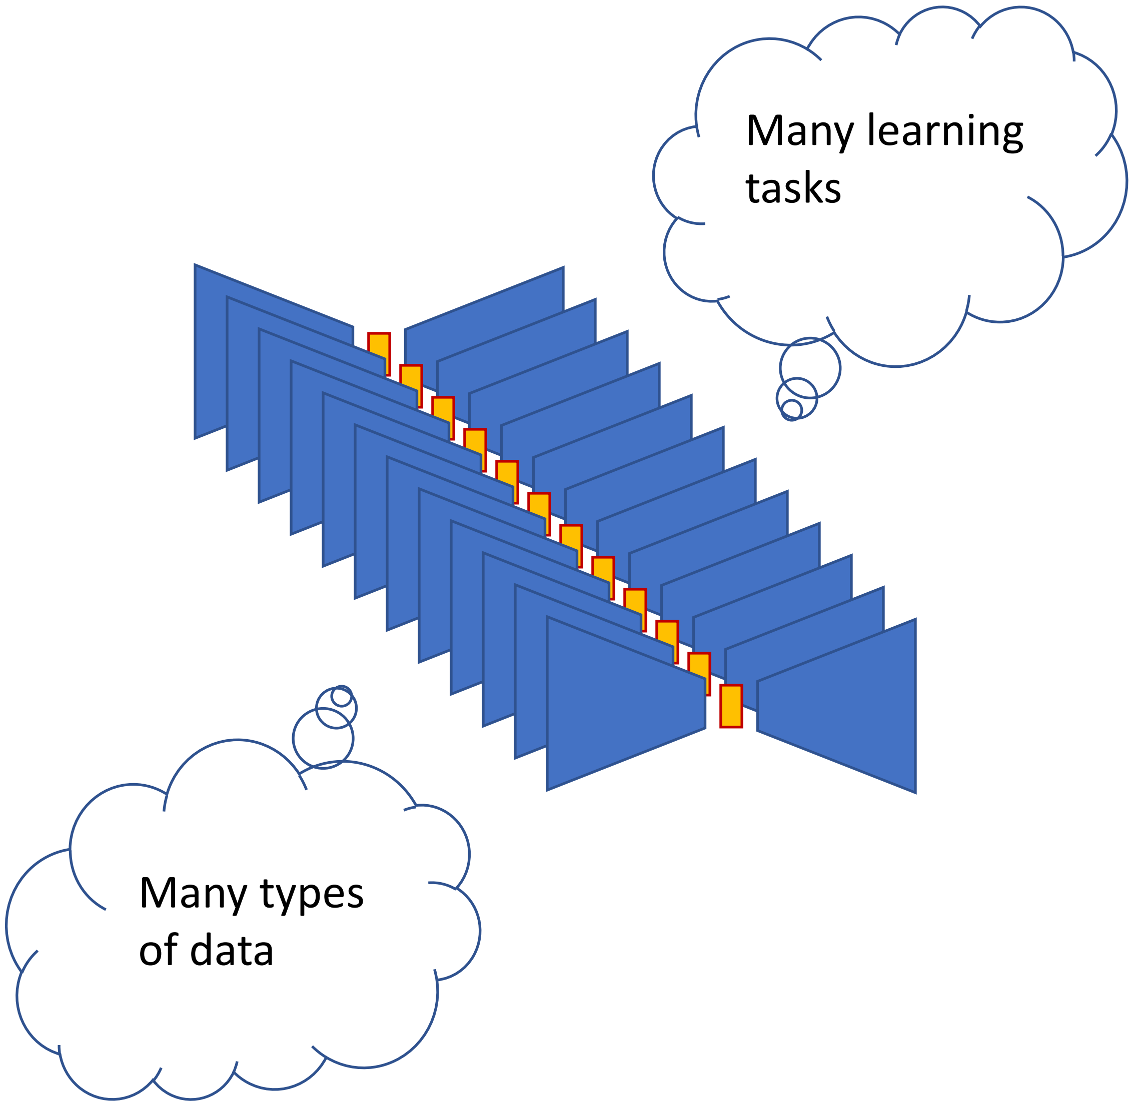

Codes

Latent vectors

Parametrizations

Representations

…

Correlated Data Sets

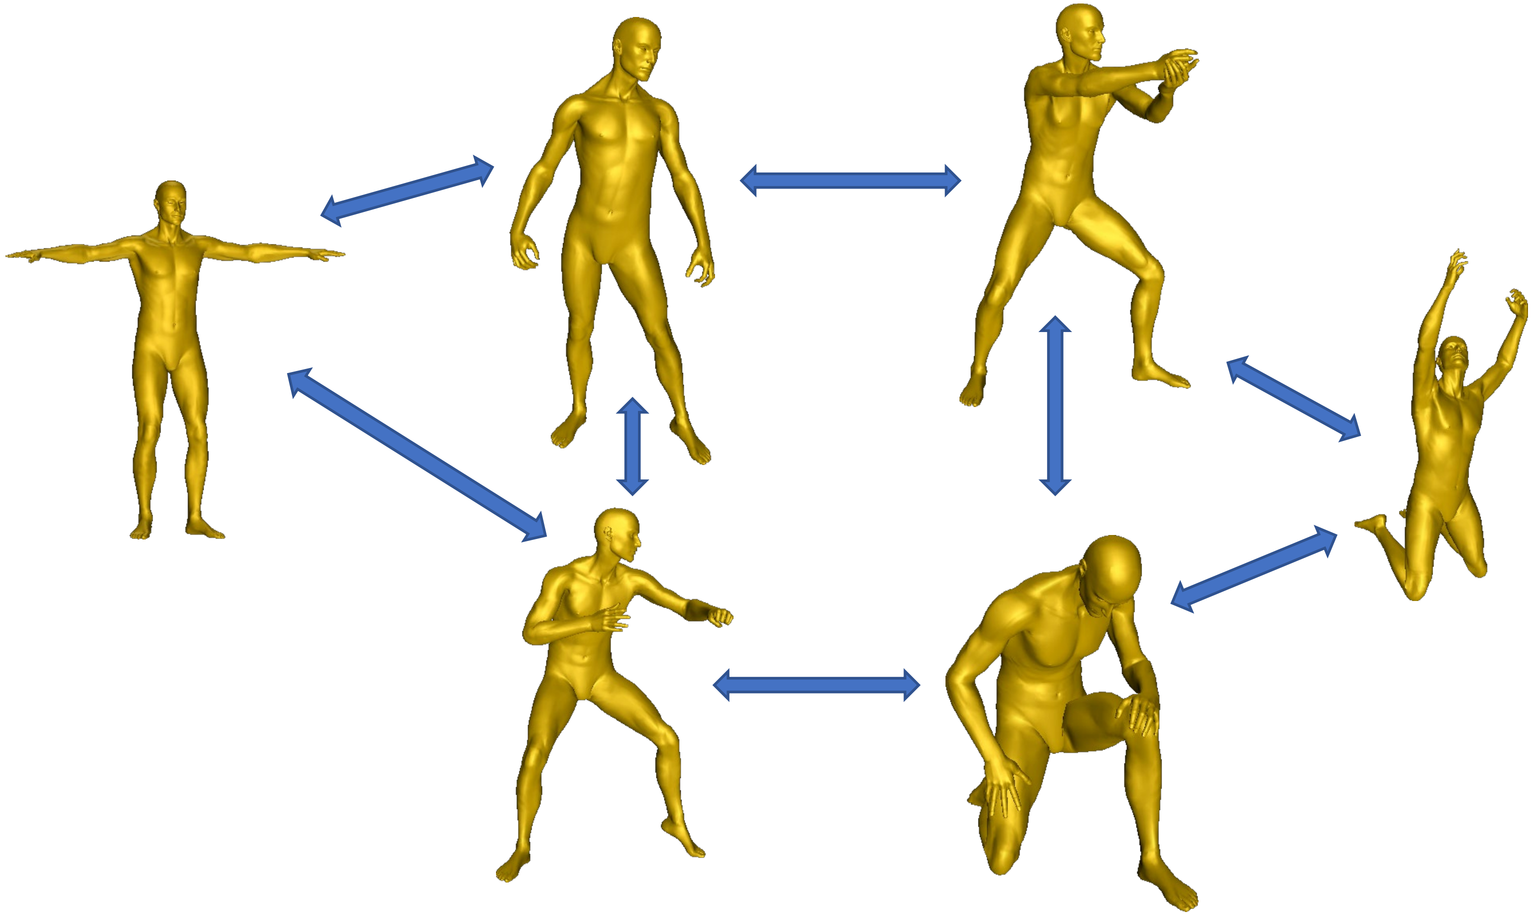

Correlated Tasks

RGB Image

Predicted Normals

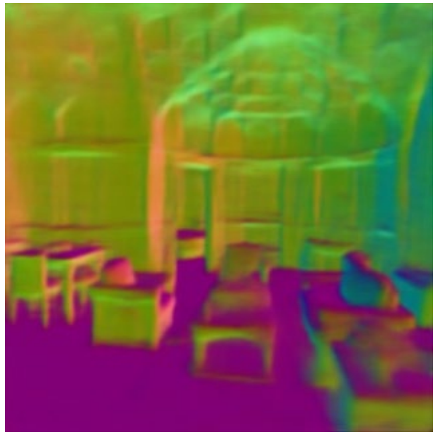

Predicted Depth

Predicted (re)Shading

Predicted Curvature

The Structure of Latent Spaces

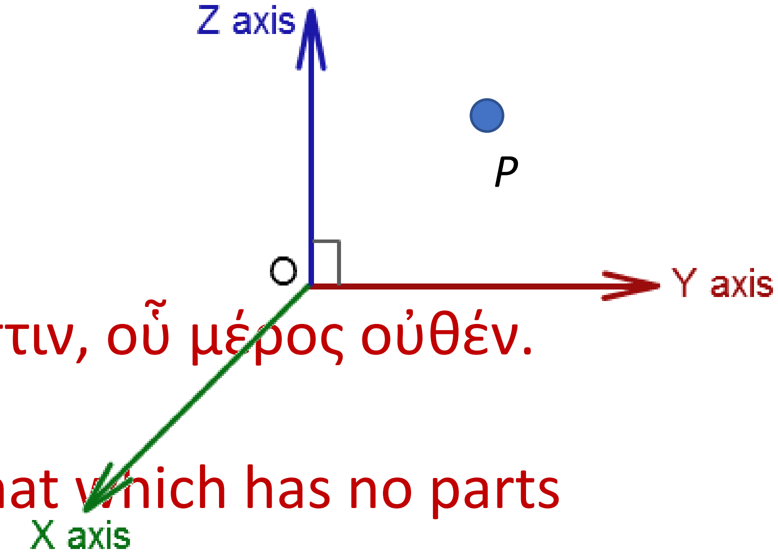

α'. Σημεῖόν ἐστιν, οὗ μέρος οὐθέν.

a. A point is that which has no parts

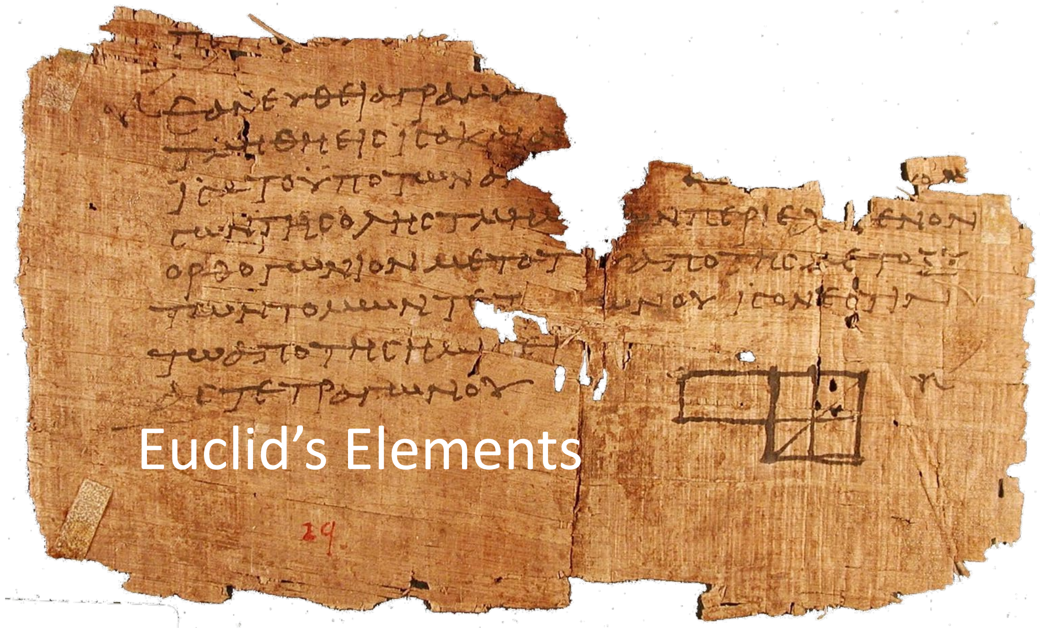
Euclid’s Elements

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?

Homomorphic Encyption

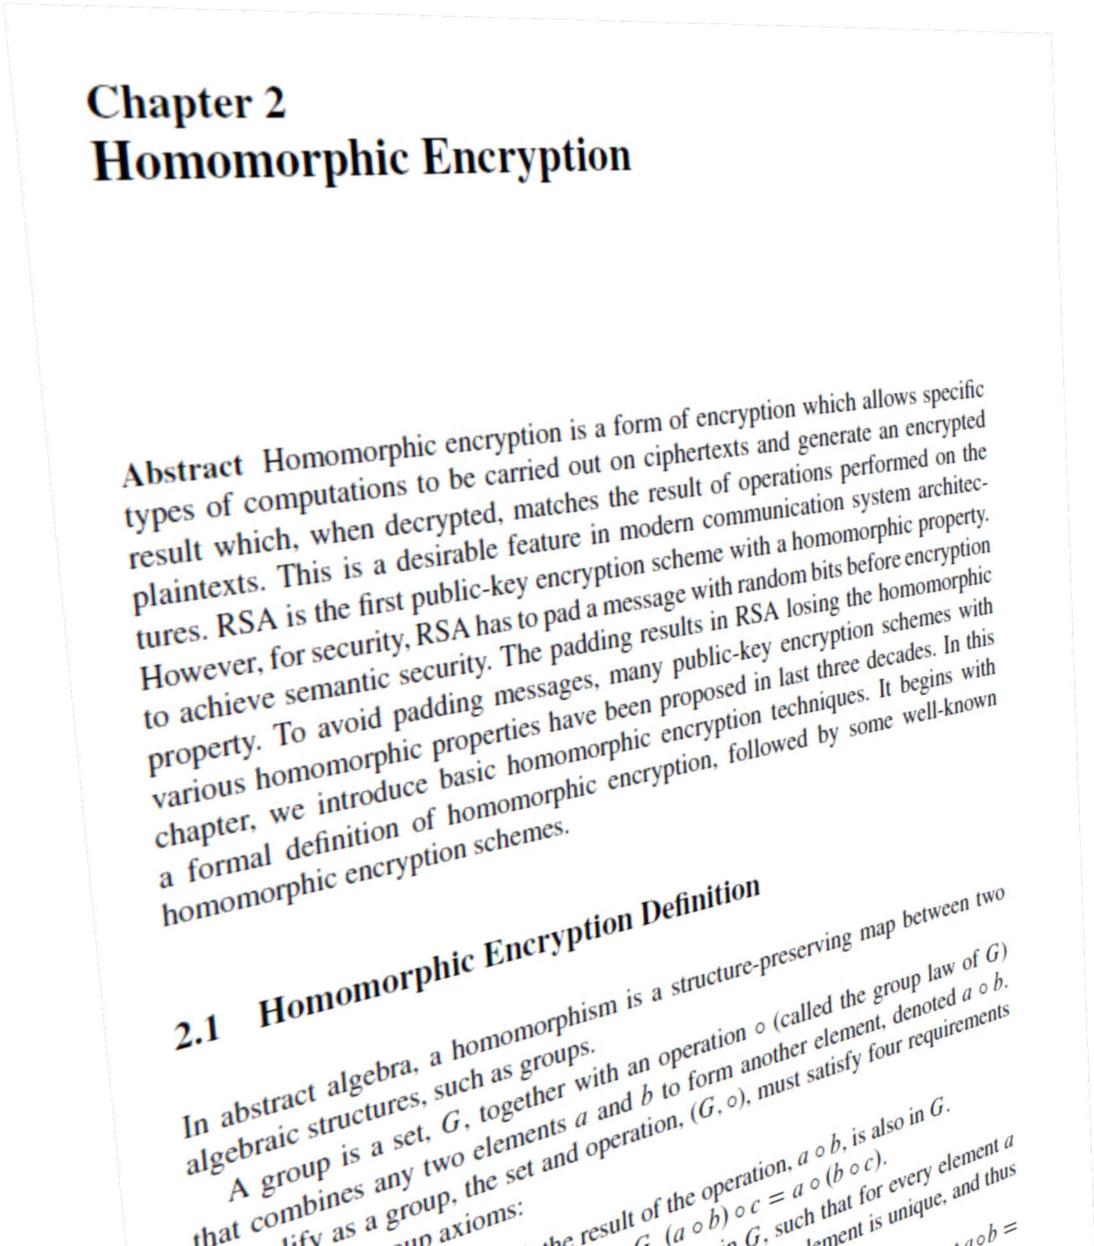
Chapter 2

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 group axioms:


Craig Gentry

Relating Data Sets: Horizontal Map Networks

Horizontal Networks & Joint Learning

Similarity as a communications channel

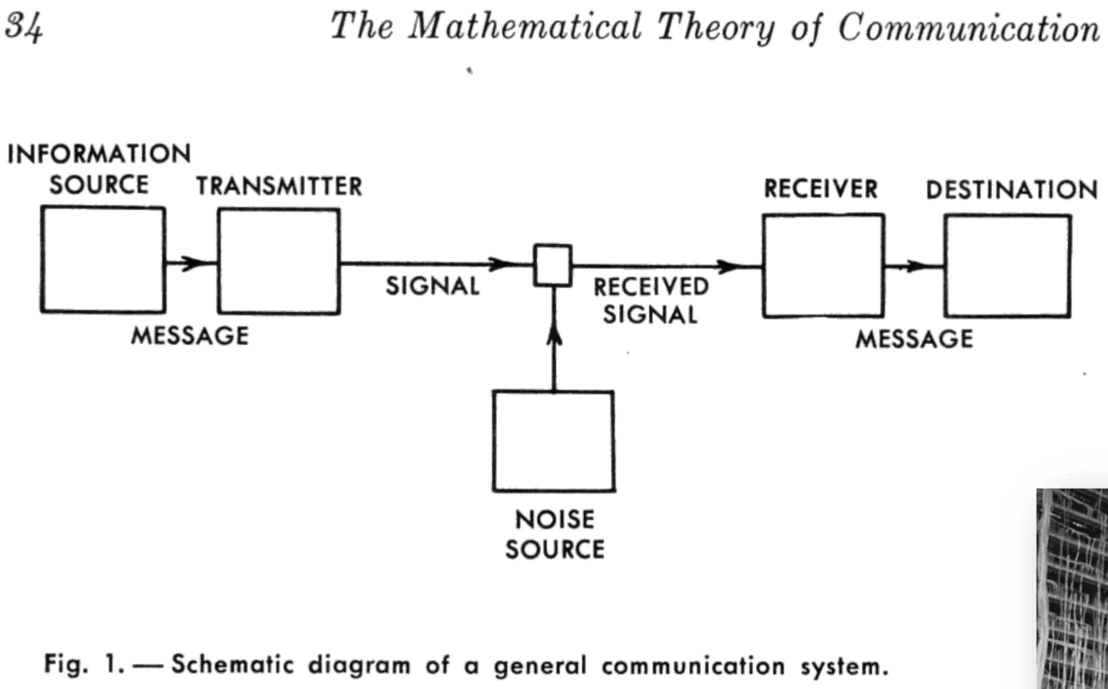
34 *The Mathematical Theory of Communication*

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

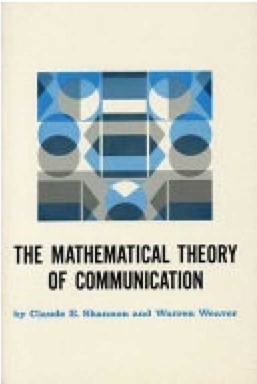

Claude Shannon

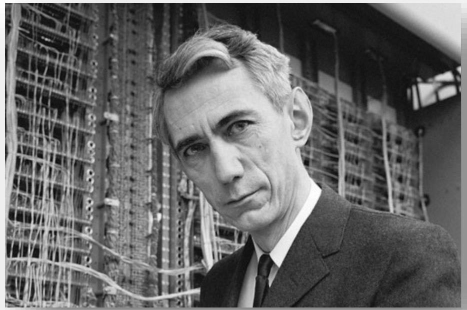

Networks of Visual Data

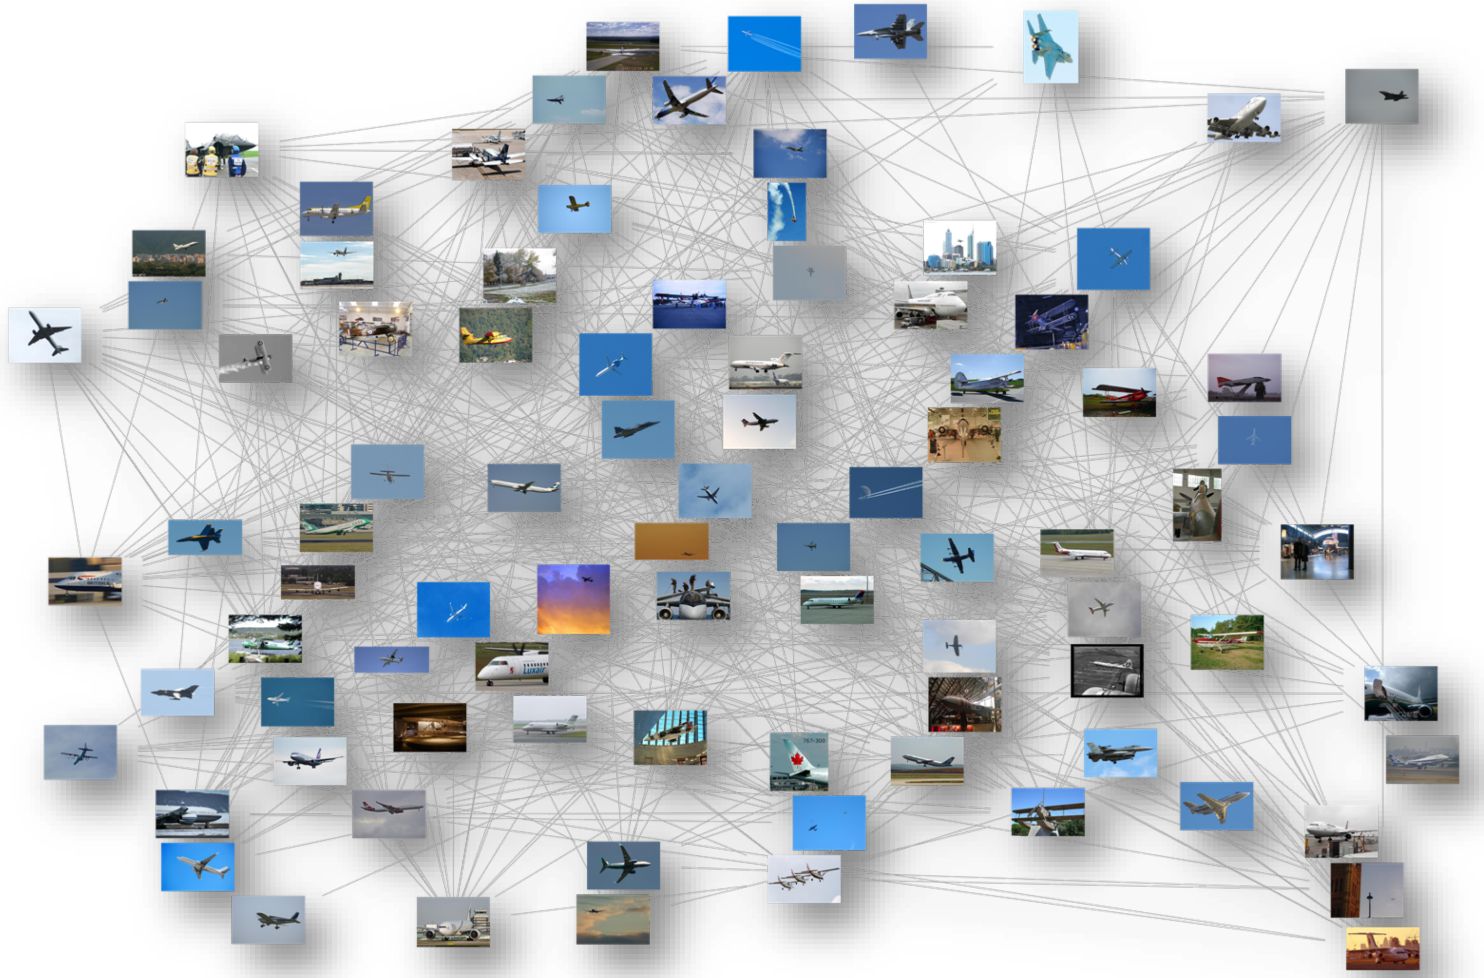

Alignment and Registration

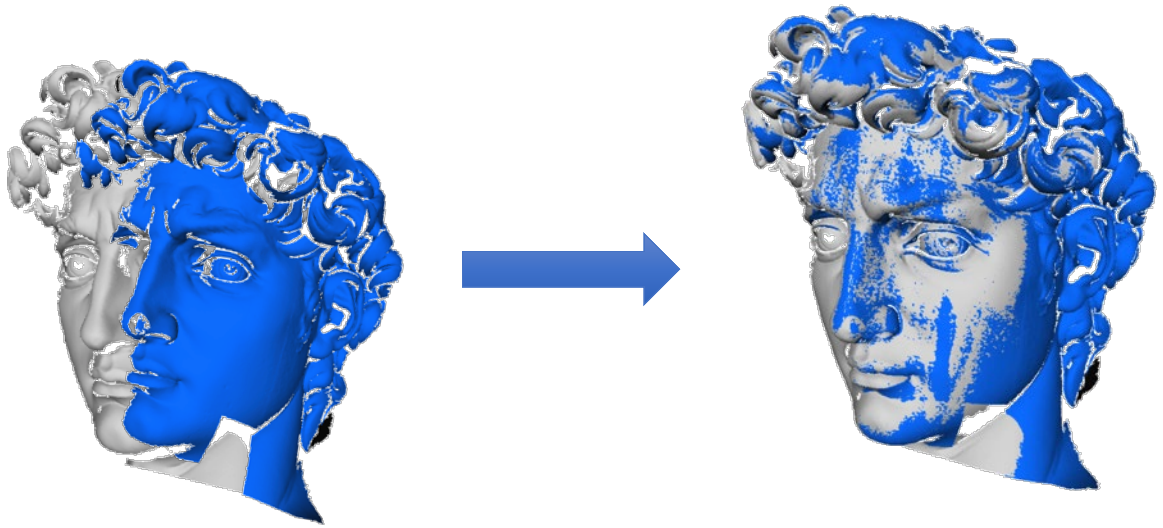

Rigid Registration

Low-dimensional transformation group

Maps and Correspondences

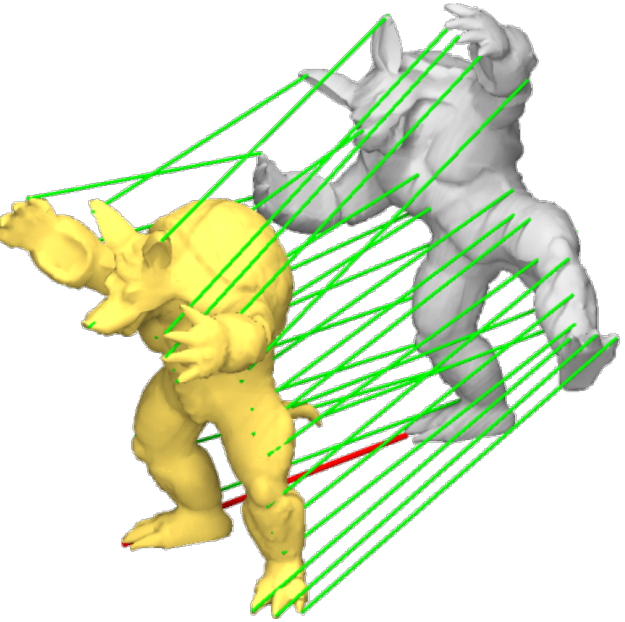

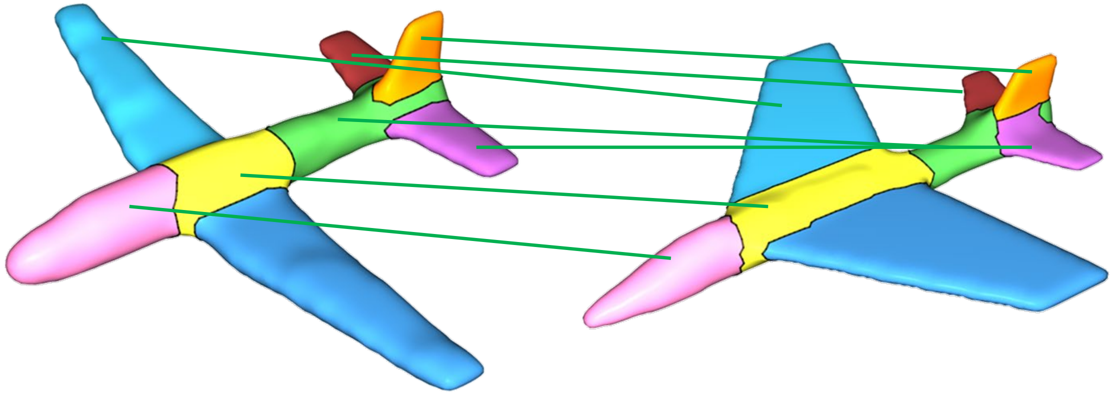

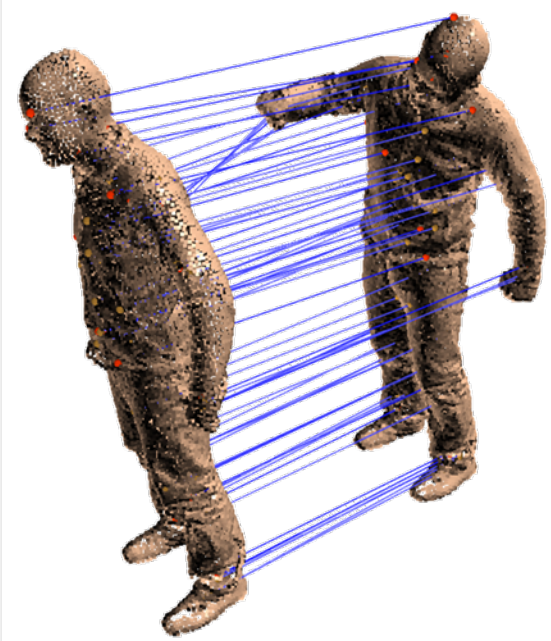

Transport Operators

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 an important regularizer in joint data analysis.

Functionals Over Data and Functional Maps

Knowledge as Functions over Data

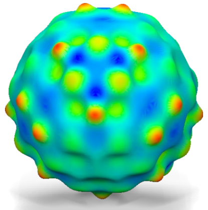

Curvature

Parts

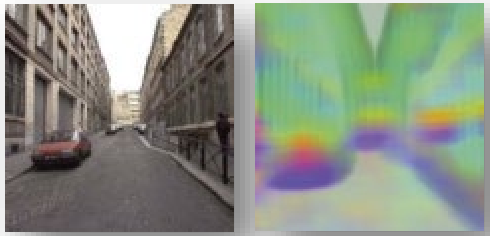

SIFT flow, C. Liu 2011

Knowledge towers over visual data: function spaces

Hierarchical Bases for a Function Space

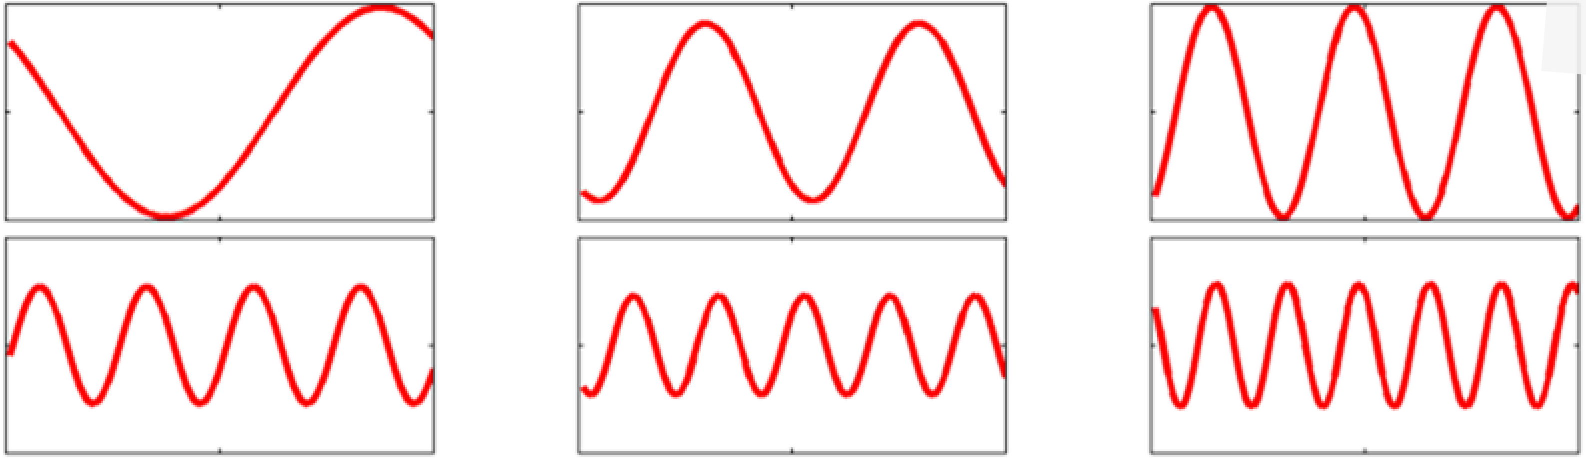

Fourier

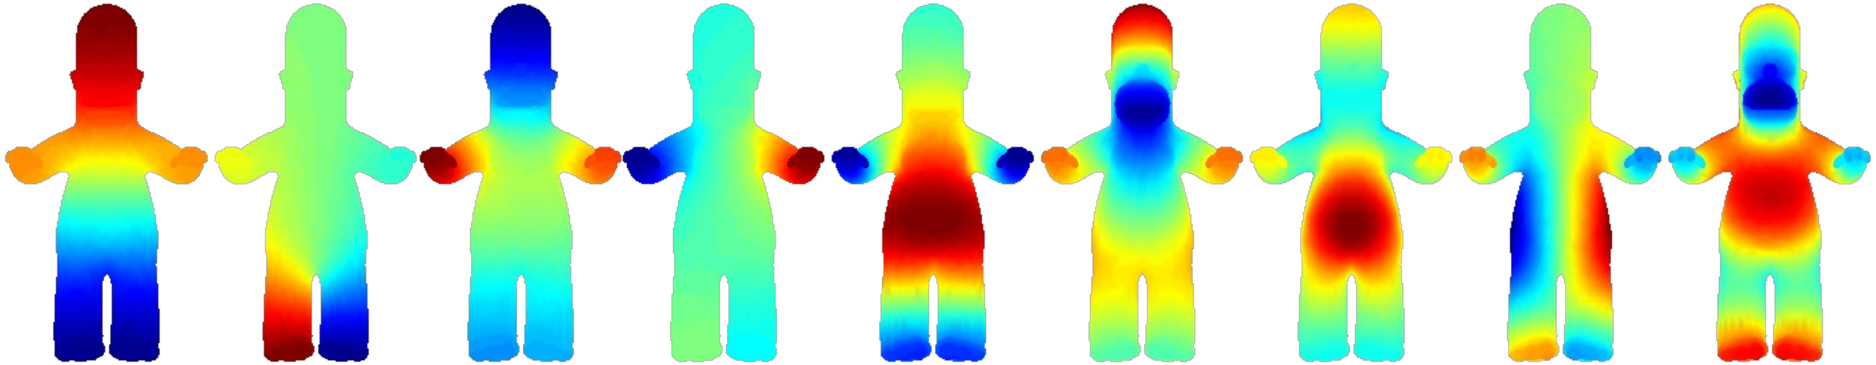

Laplace-Beltrami

global support

Functional Maps

$$T : N \to M$$

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

A contravariant functor

Starting from a Regular Map ϕ

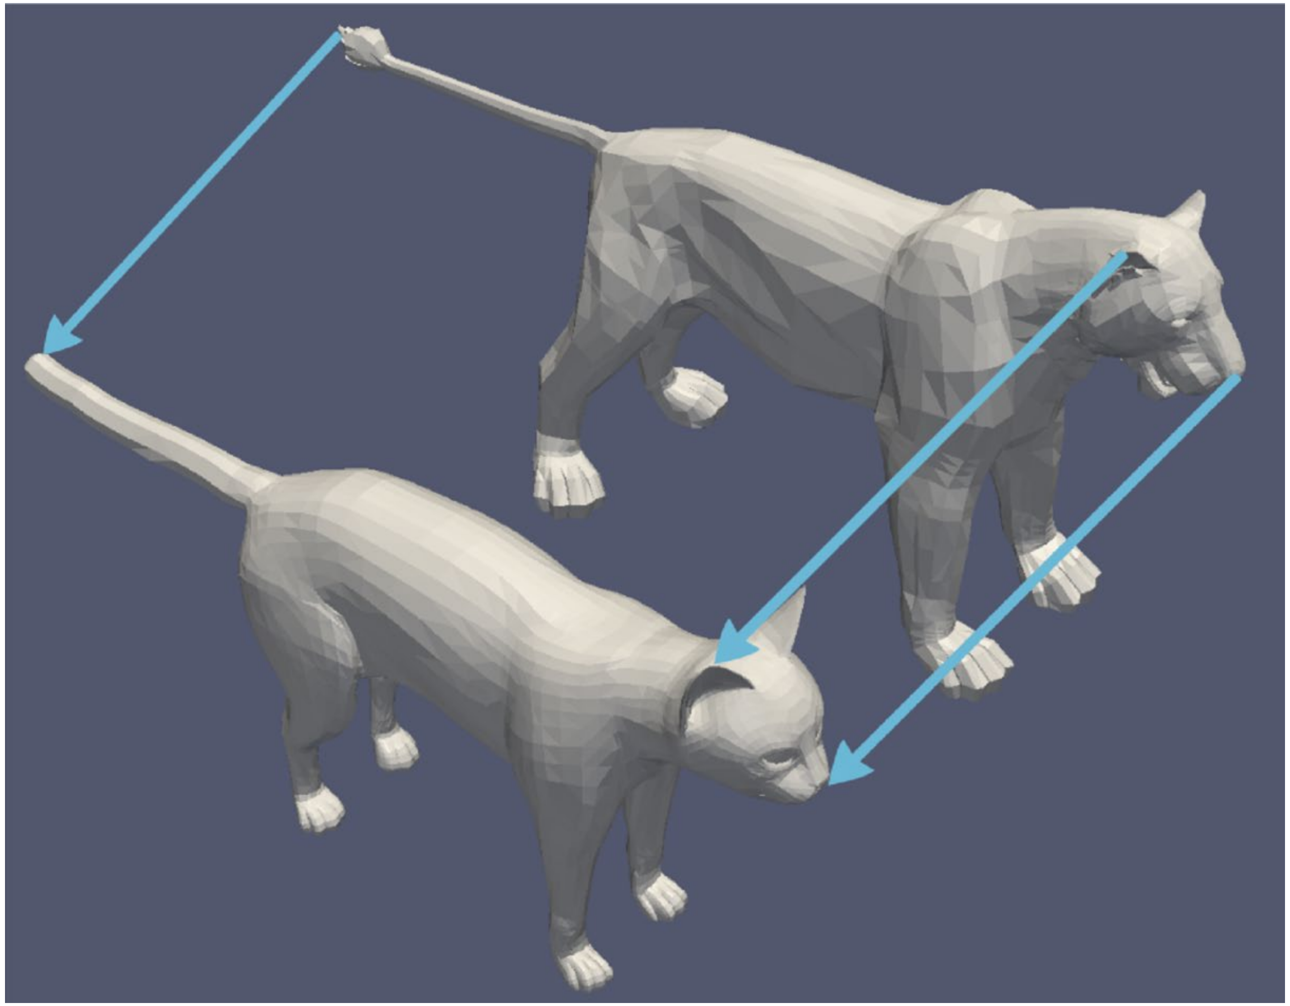

ϕ: lion → cat

A Contravariant Functor

from cat to lion

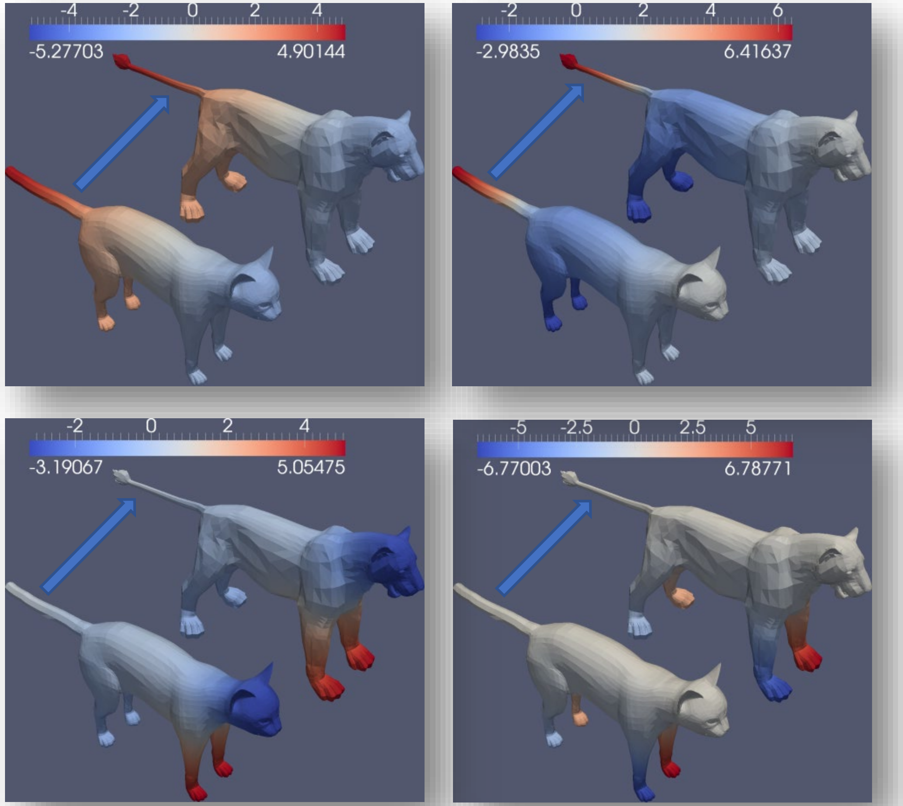

Functions on cat are transferred to lion using T_ϕ

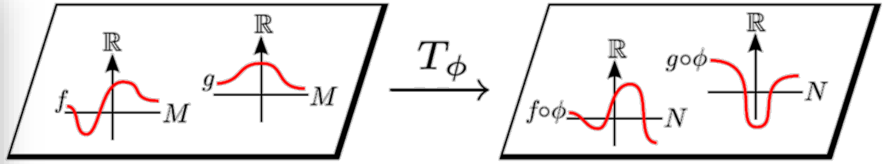

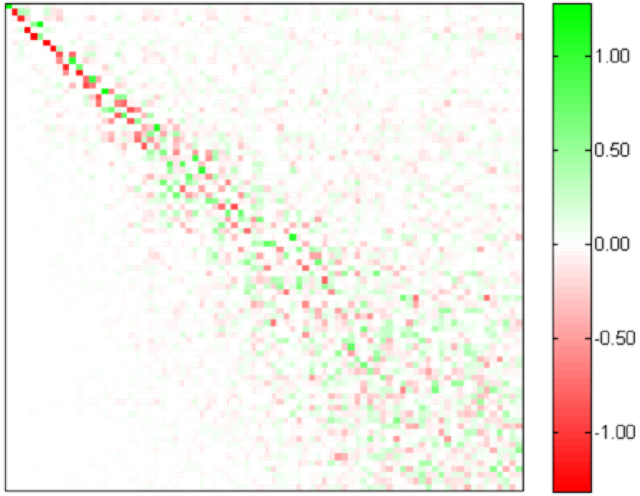

T_ϕ is a linear operator (matrix)

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

Compact Encoding of Maps as Matrices

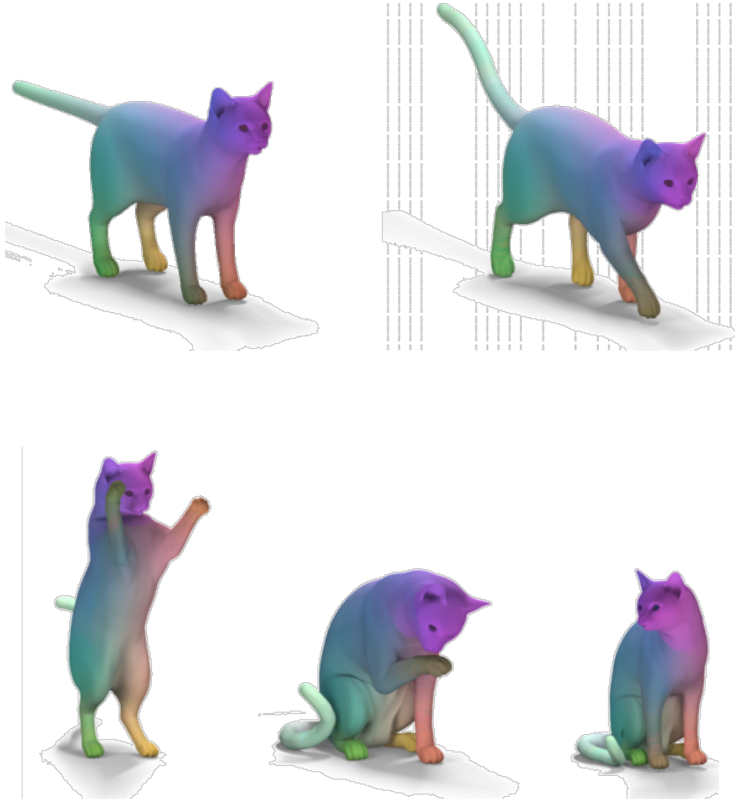

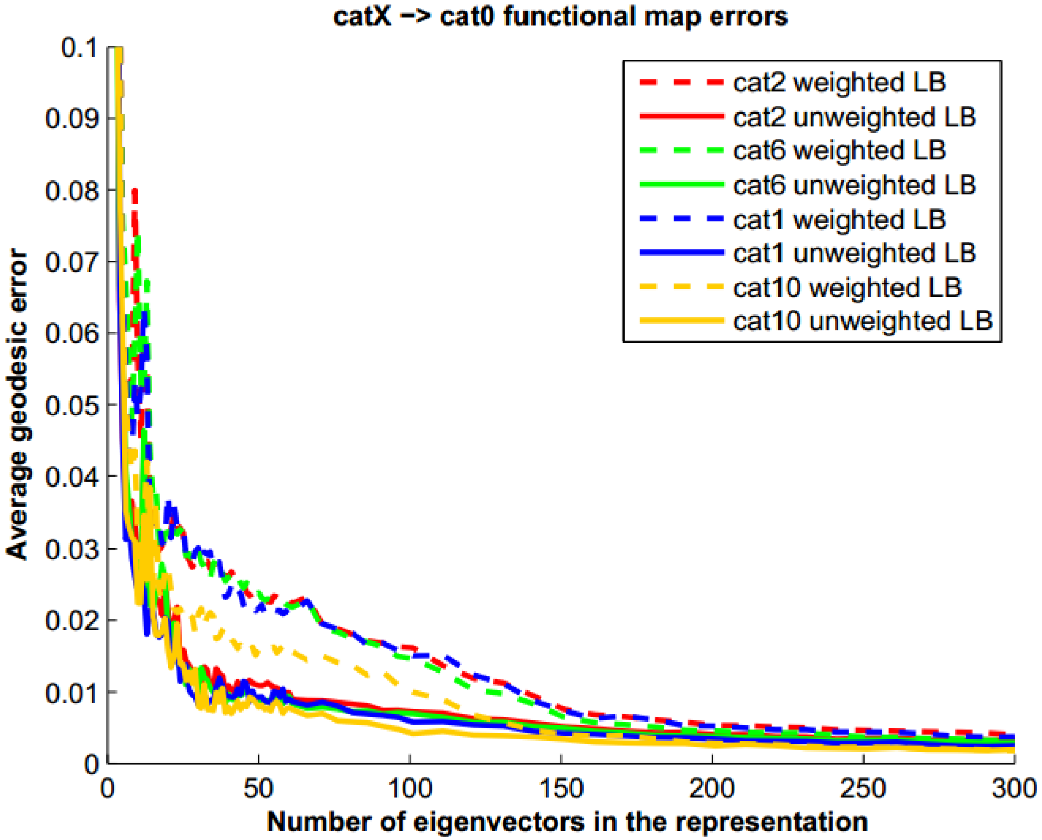

Changes to Measure and Metric under a Map

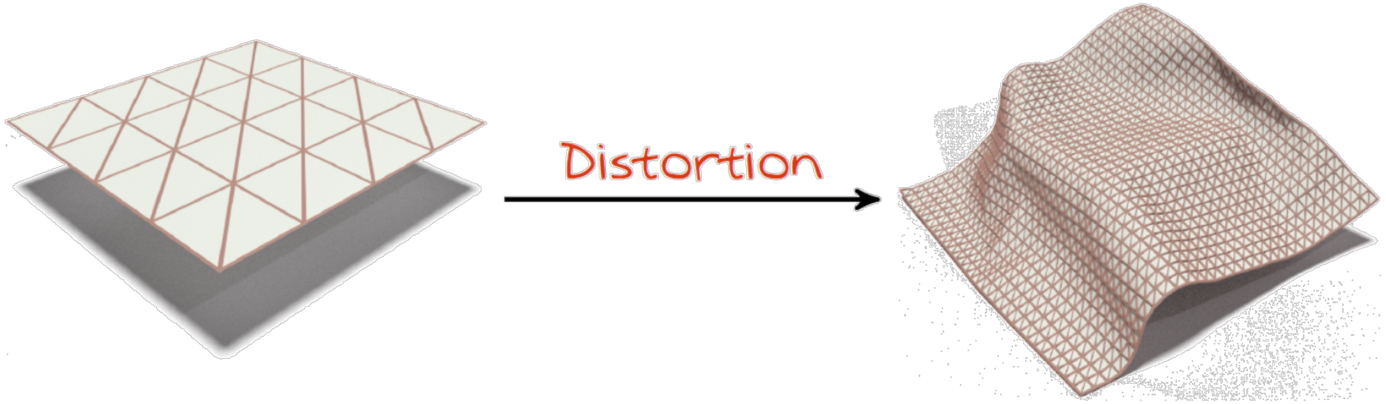

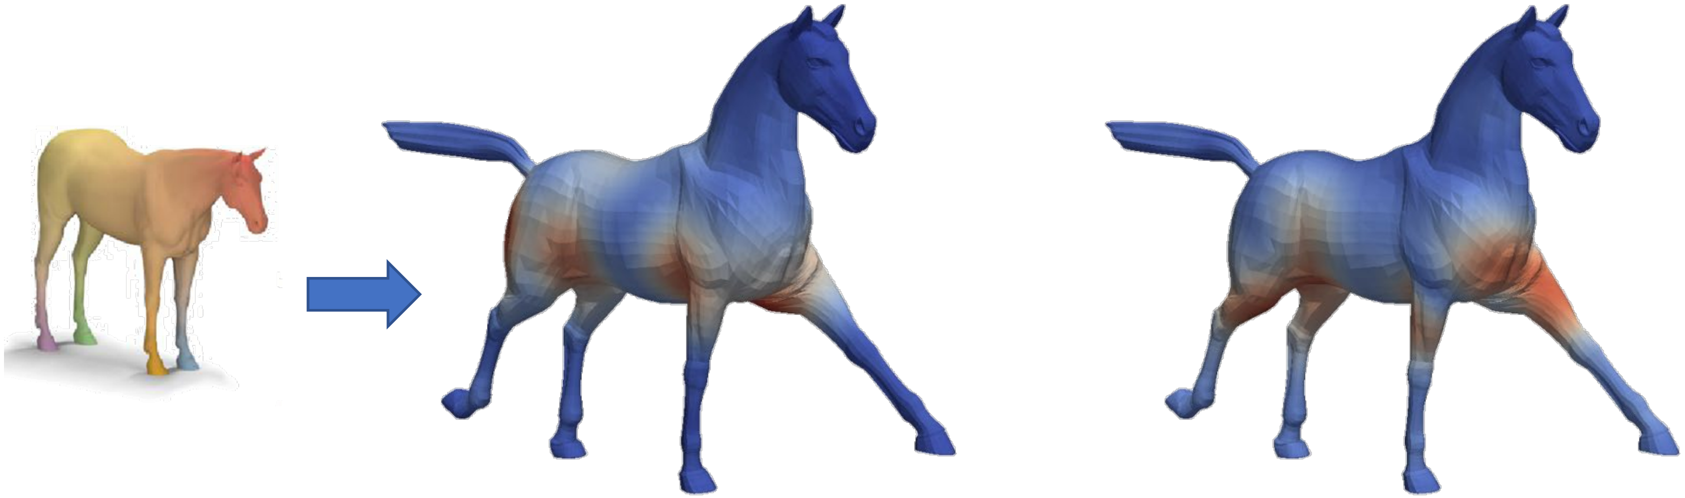

Area distortion

Conformal distortion

Shape Differences are a Change Recipe

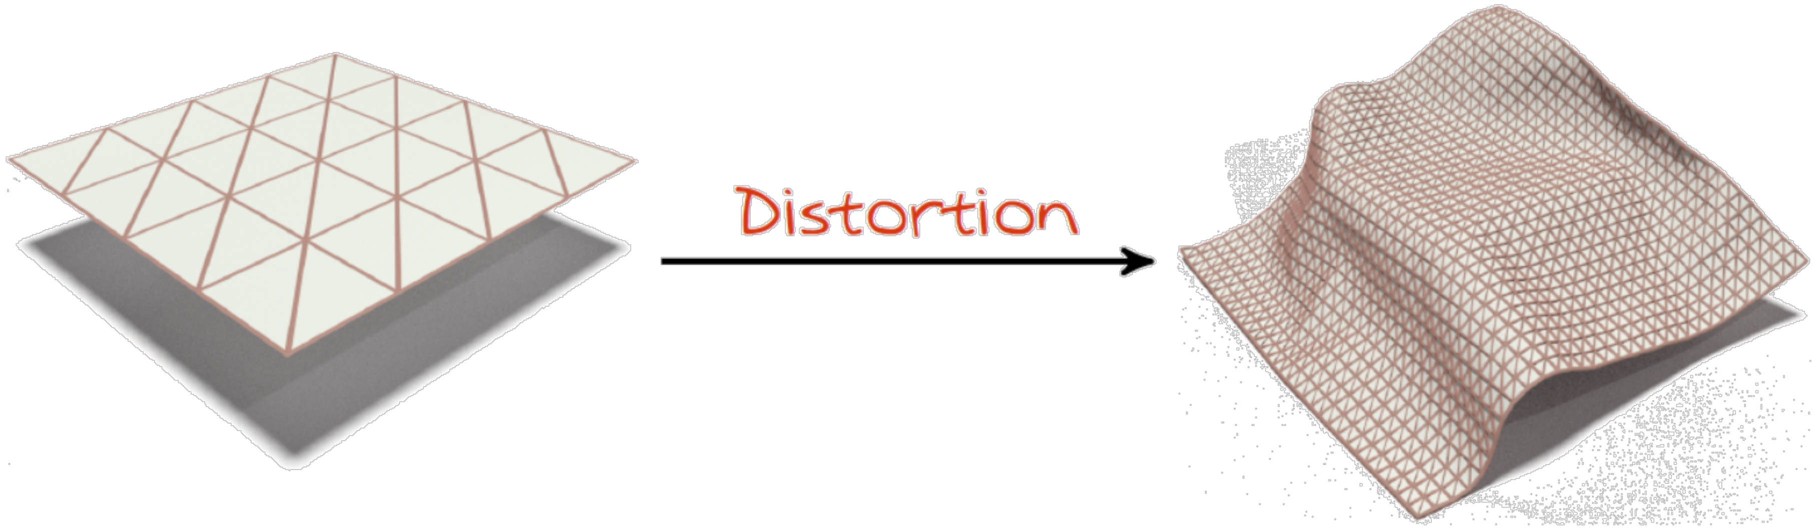

A recipe encoded as a matrix:

Area distortion
Conformal distortion

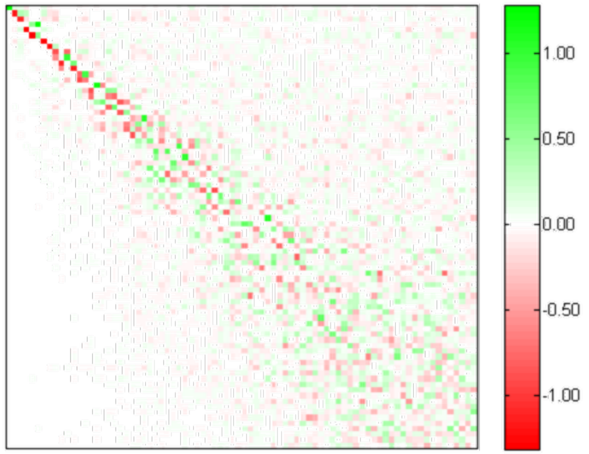

A novel type of latent space representation for 3D data

Under some conditions, lossless

The Space of Shapes

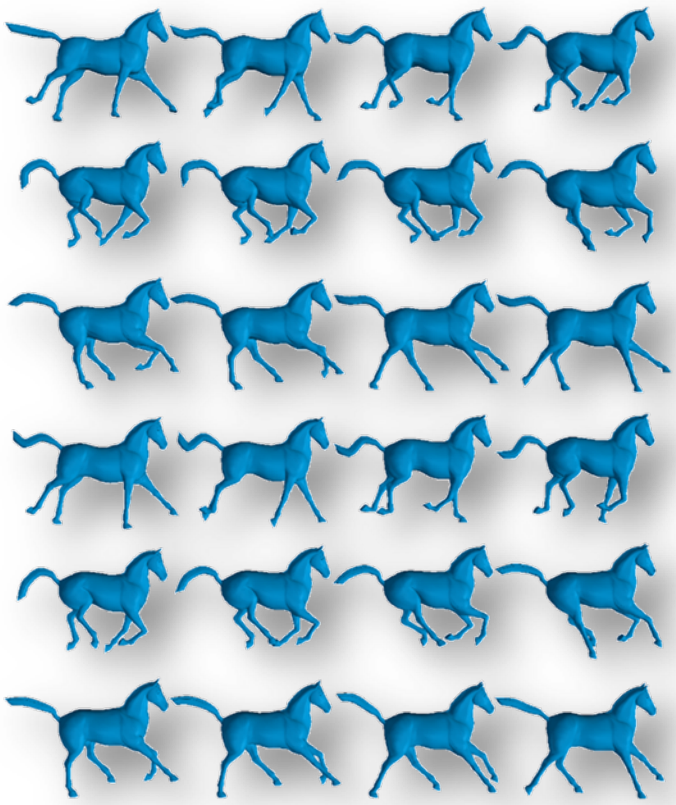

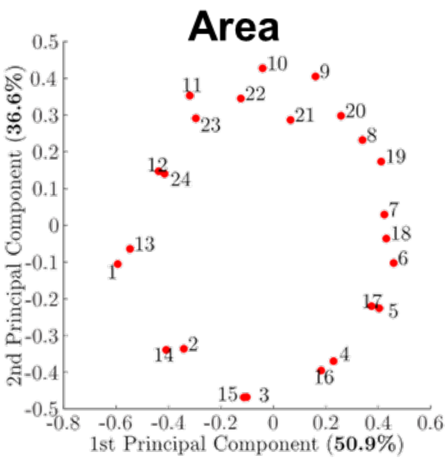

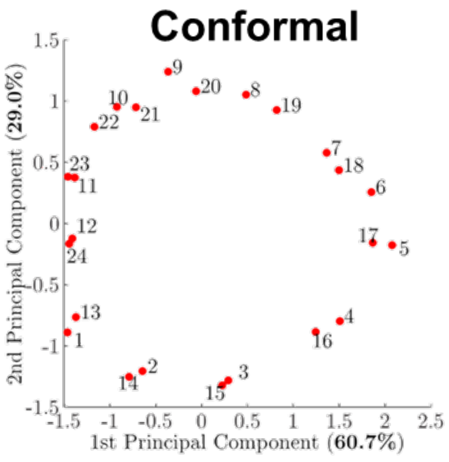

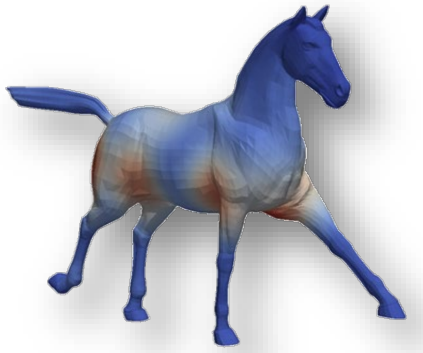

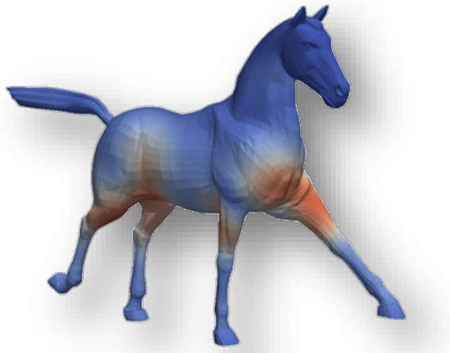

Horizontal Map Networks

Networks Between Data

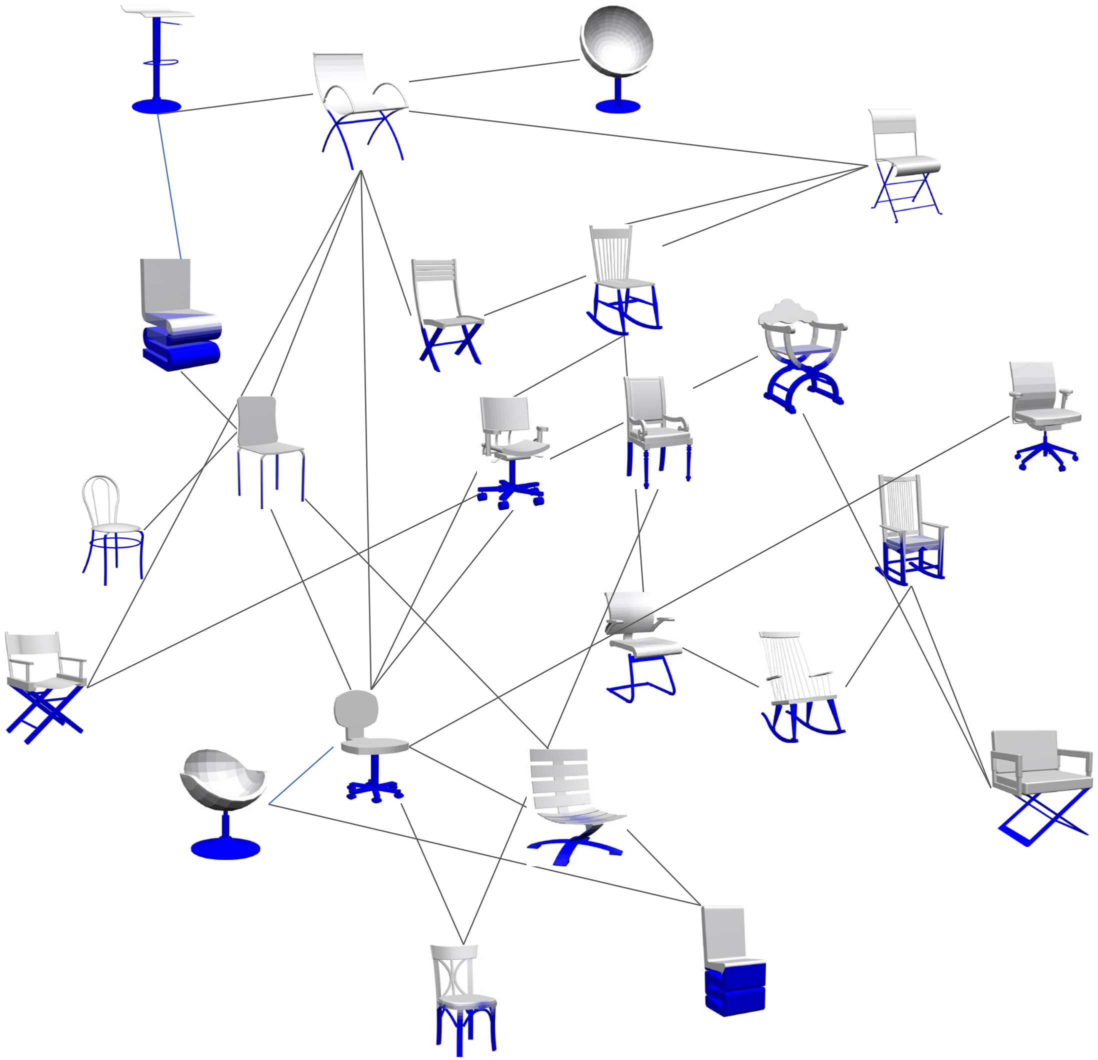

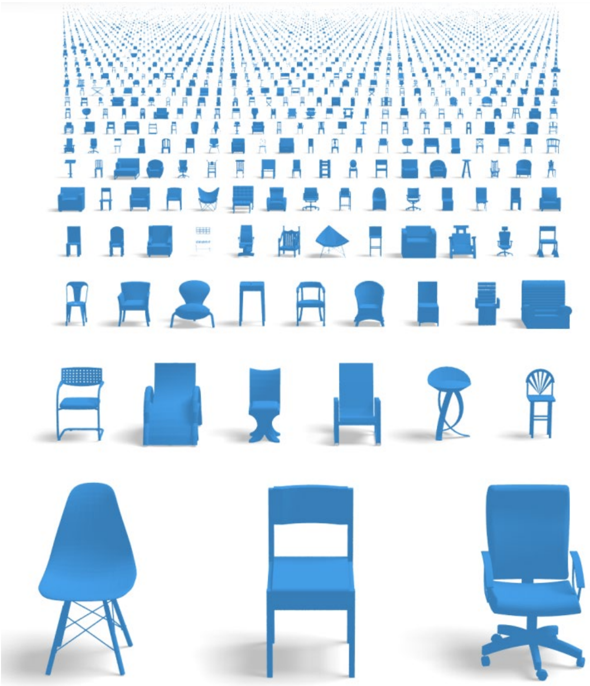

Joint Analysis: Co-Segmentation

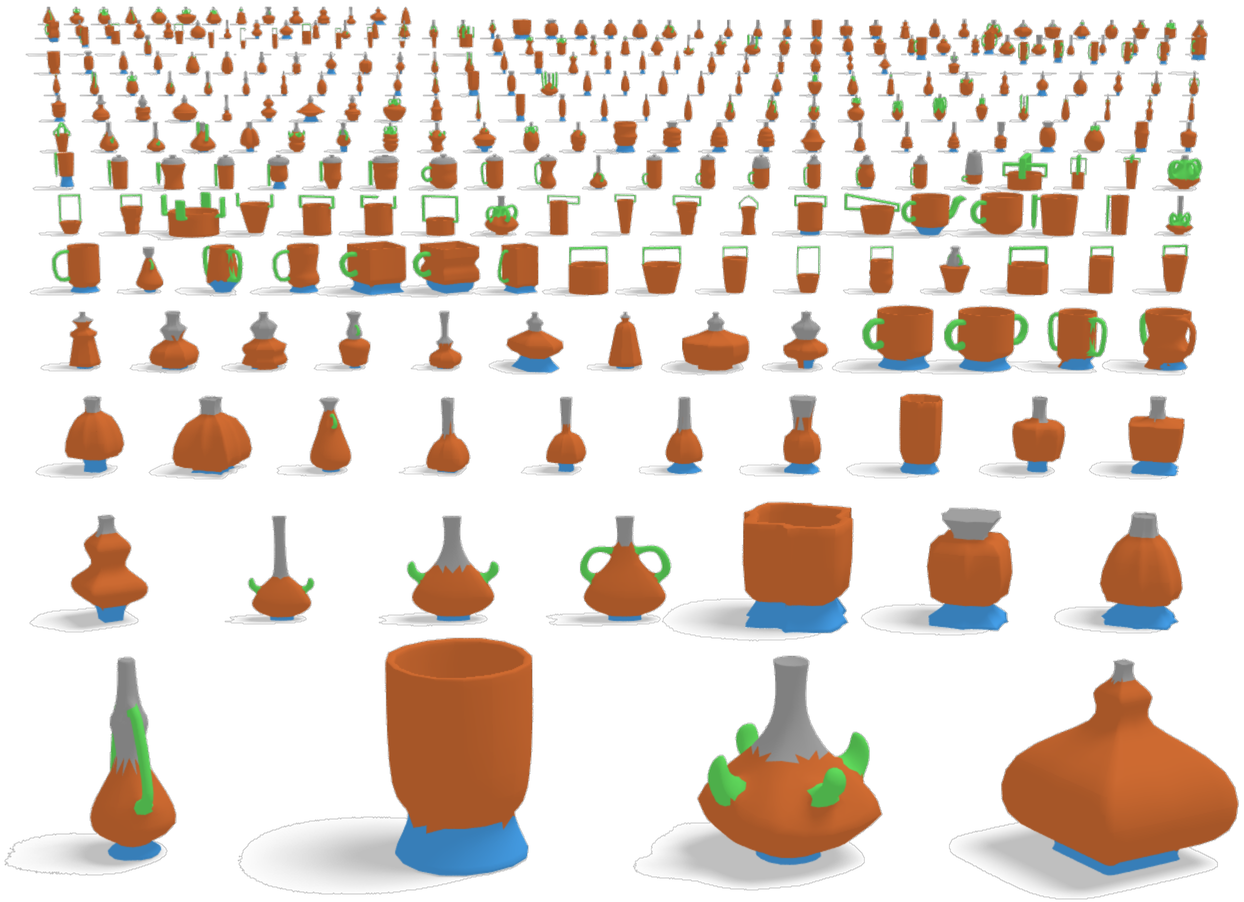

Unsupervised Structure Extraction

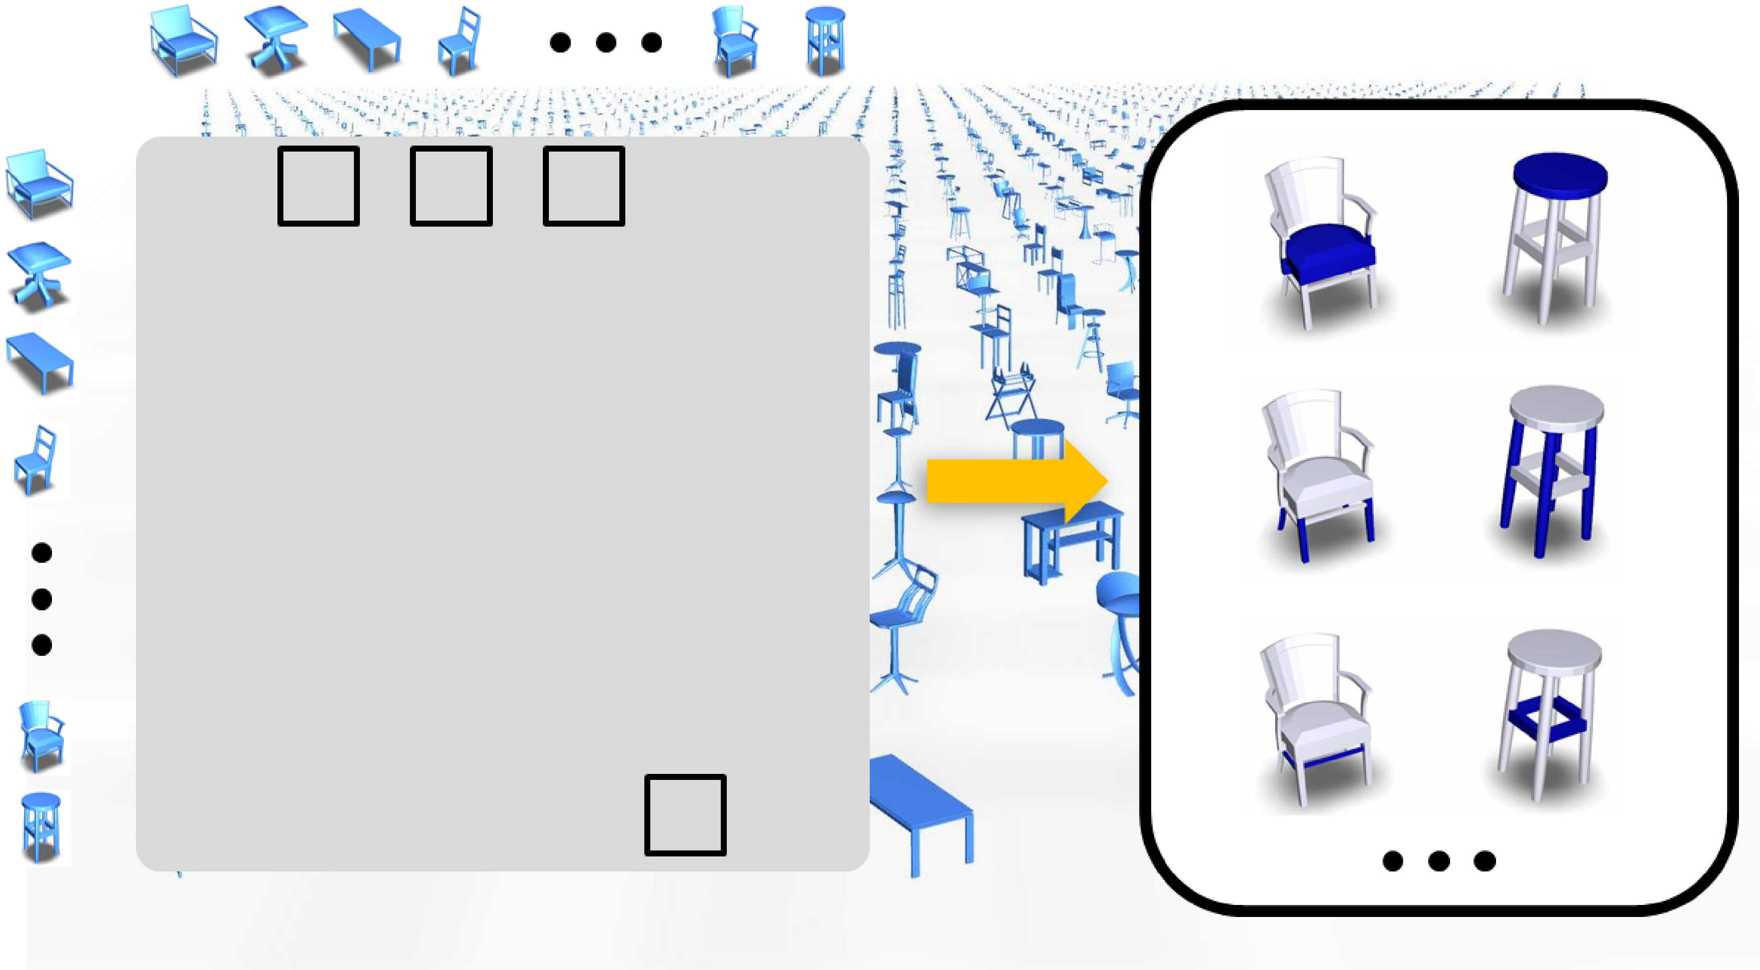

[Q. Huang, F. Wang, L. Guibas, '14]

MSRC: 5 images per class are shown

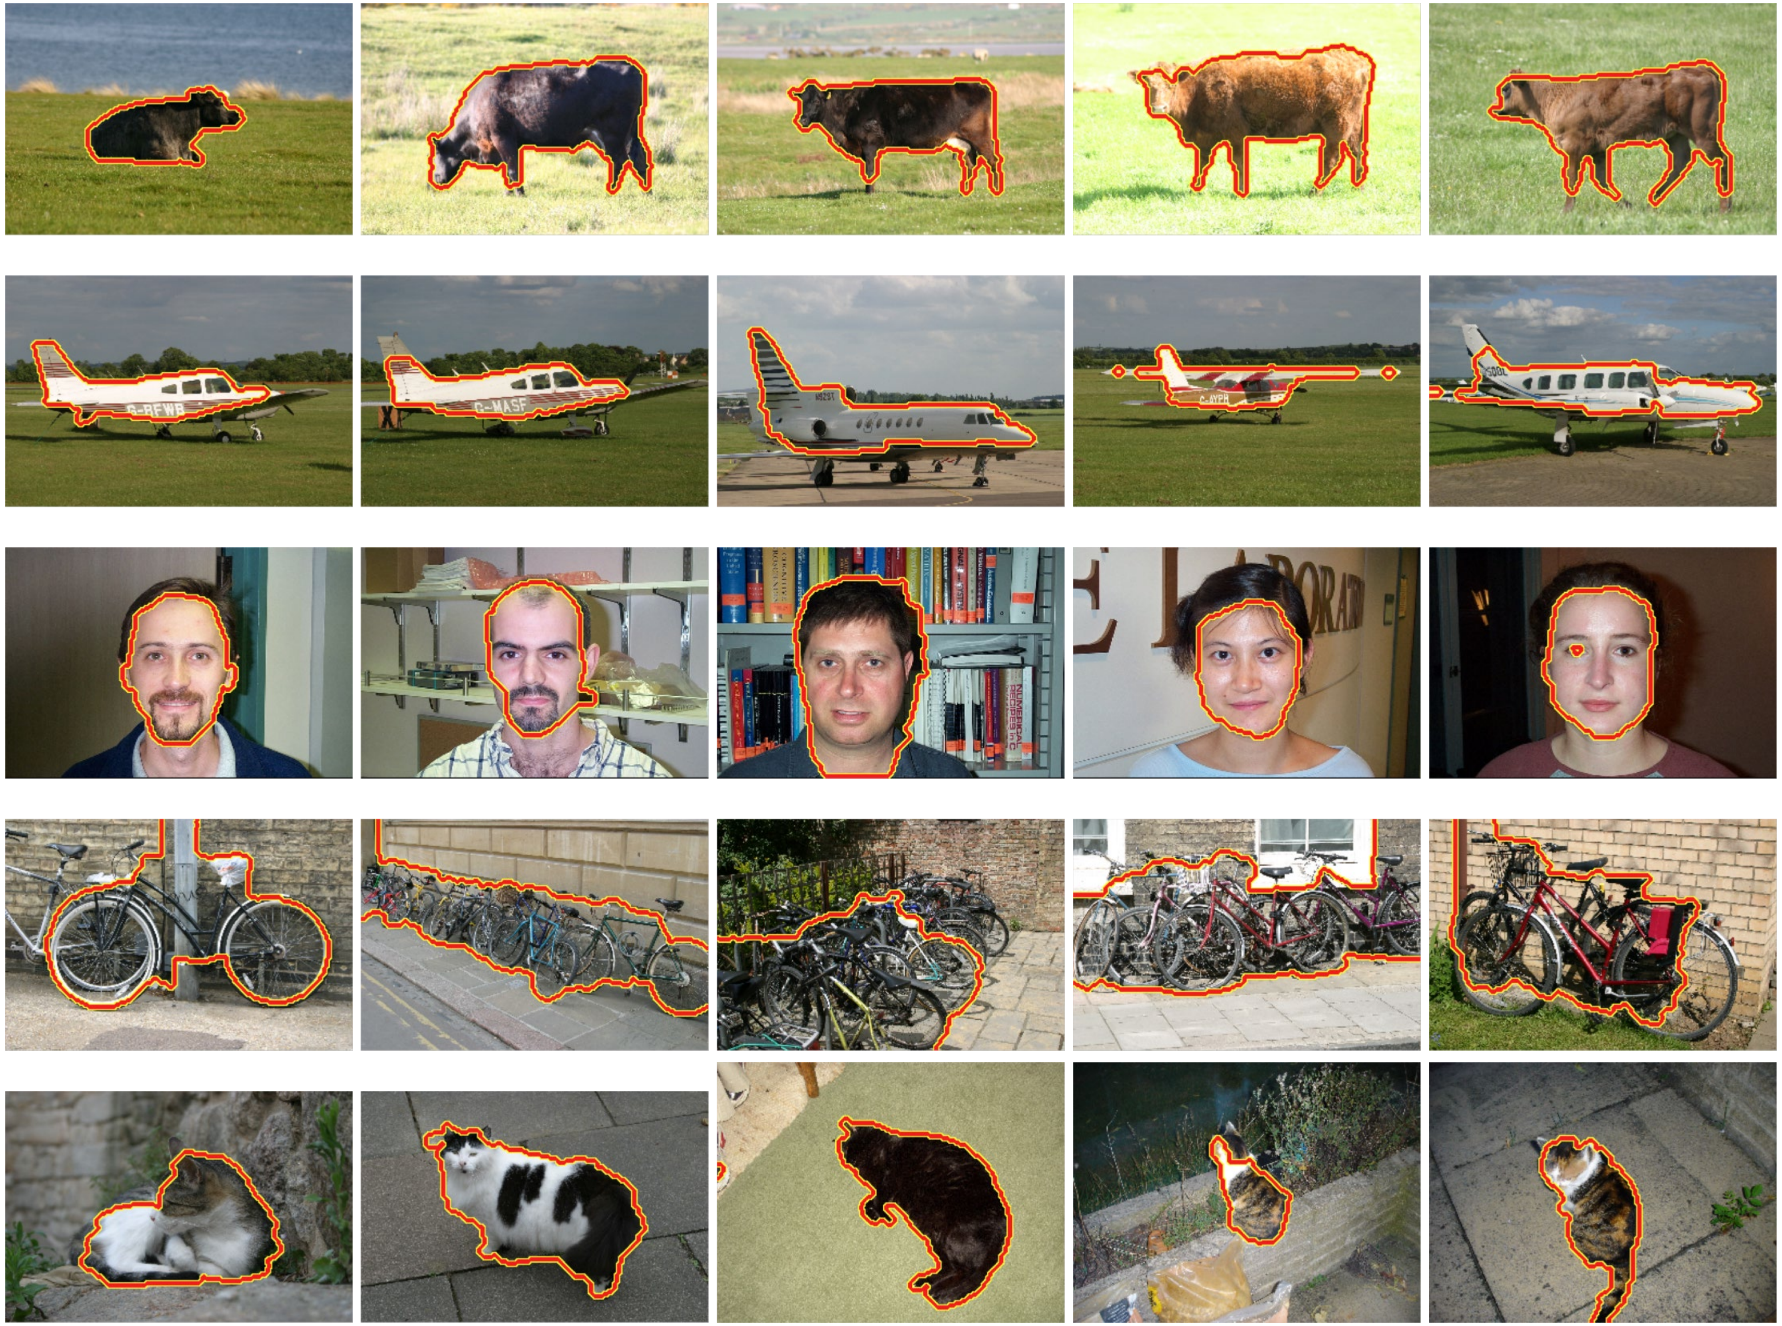

The Network is the Abstraction

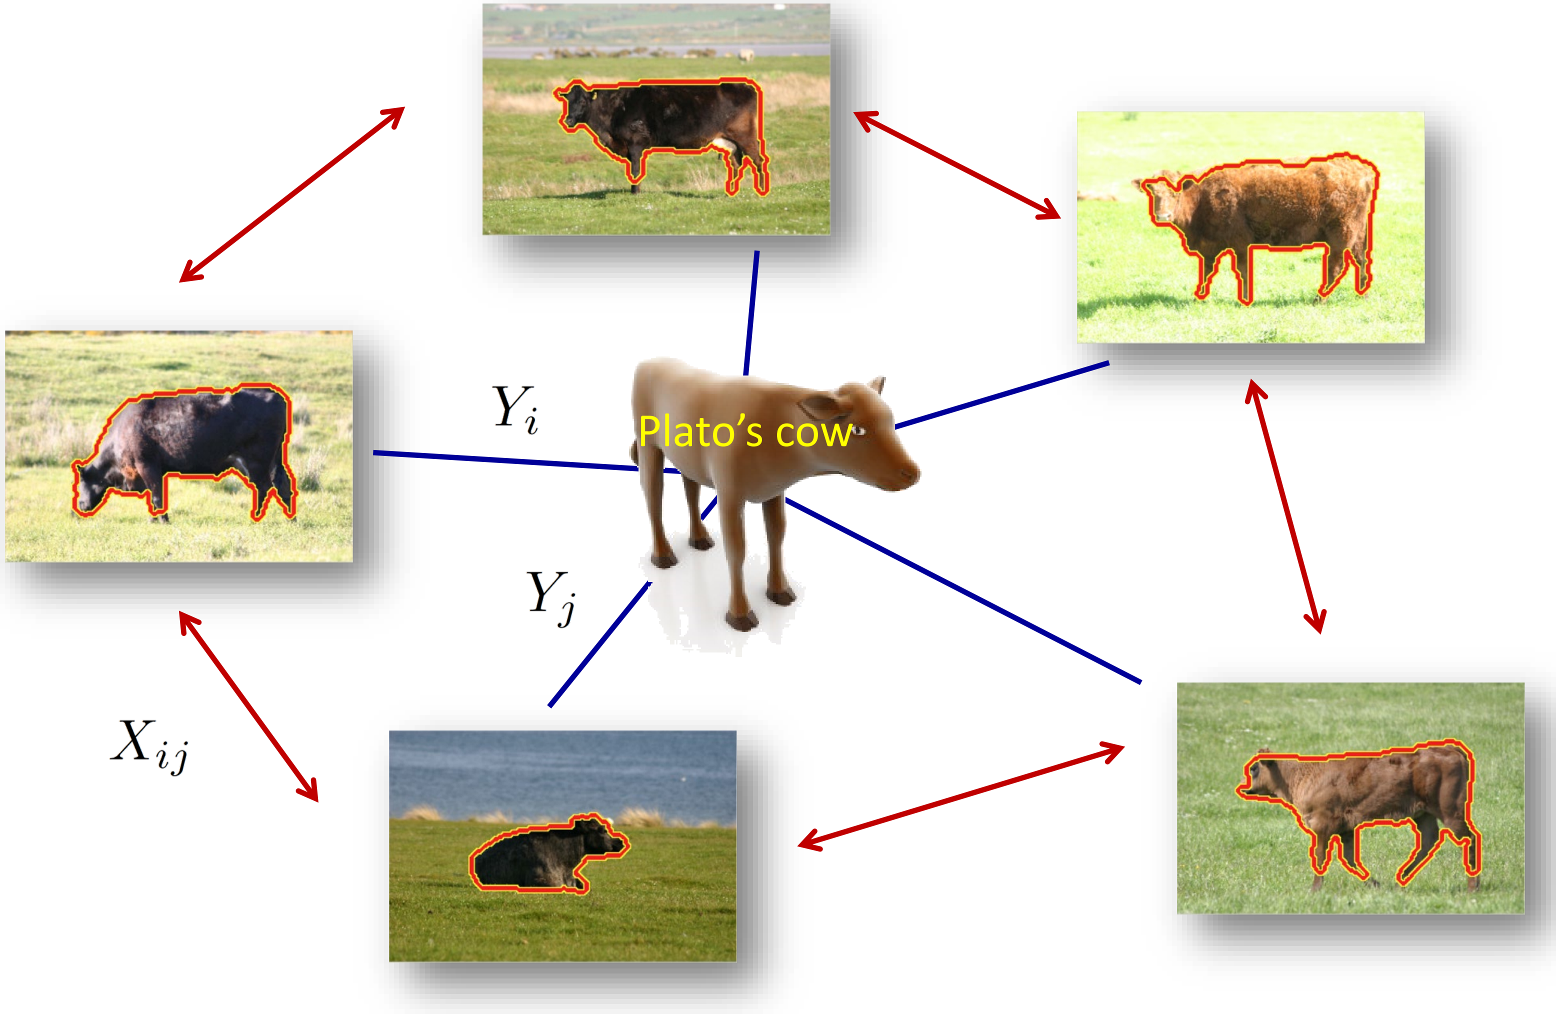

Mosaicking or SLAM at the Level of Functions

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

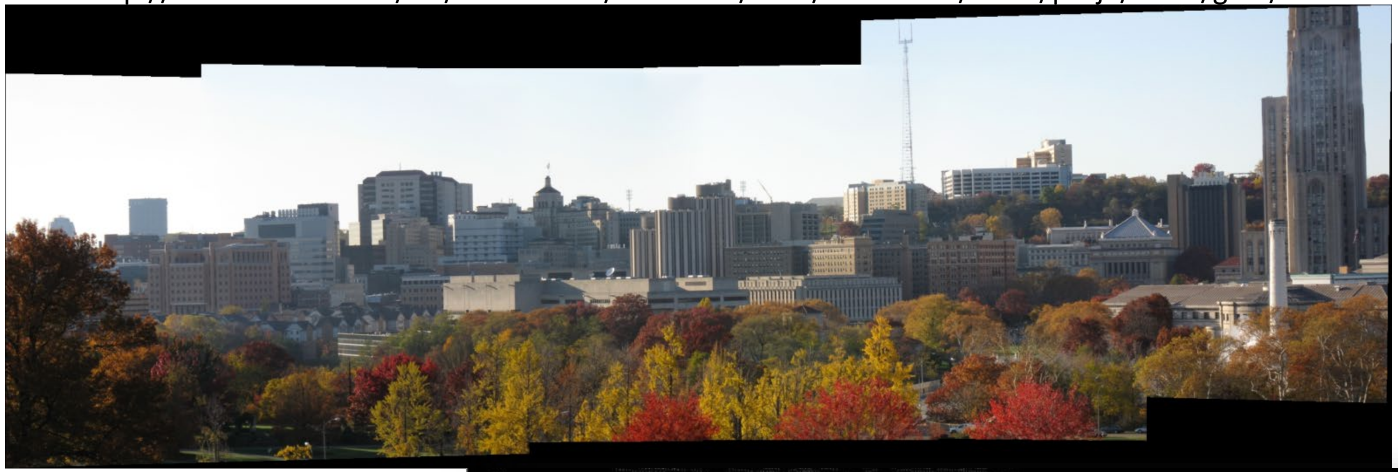

robotics.ait.kyushu-u.ac.jp

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

ImageNet and ShapeNet

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

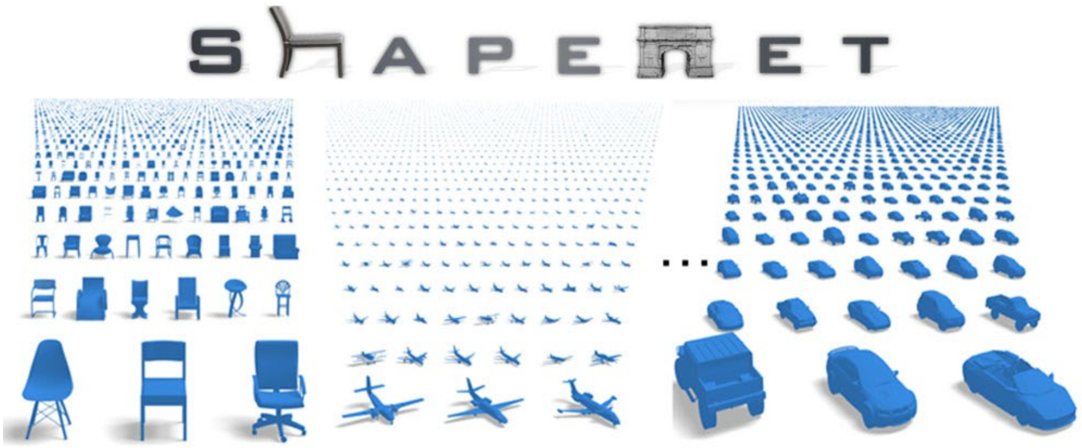

Semantic Networks

- 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

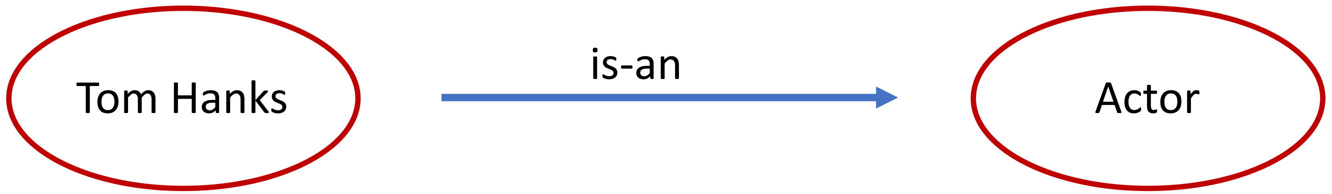

Example of a Semantic Net

Semantic Net in Lisp

```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

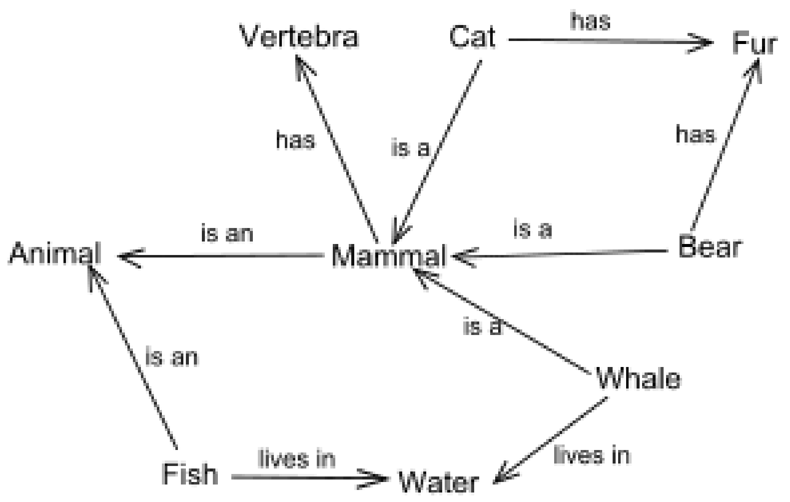

Google Knowledge Graph

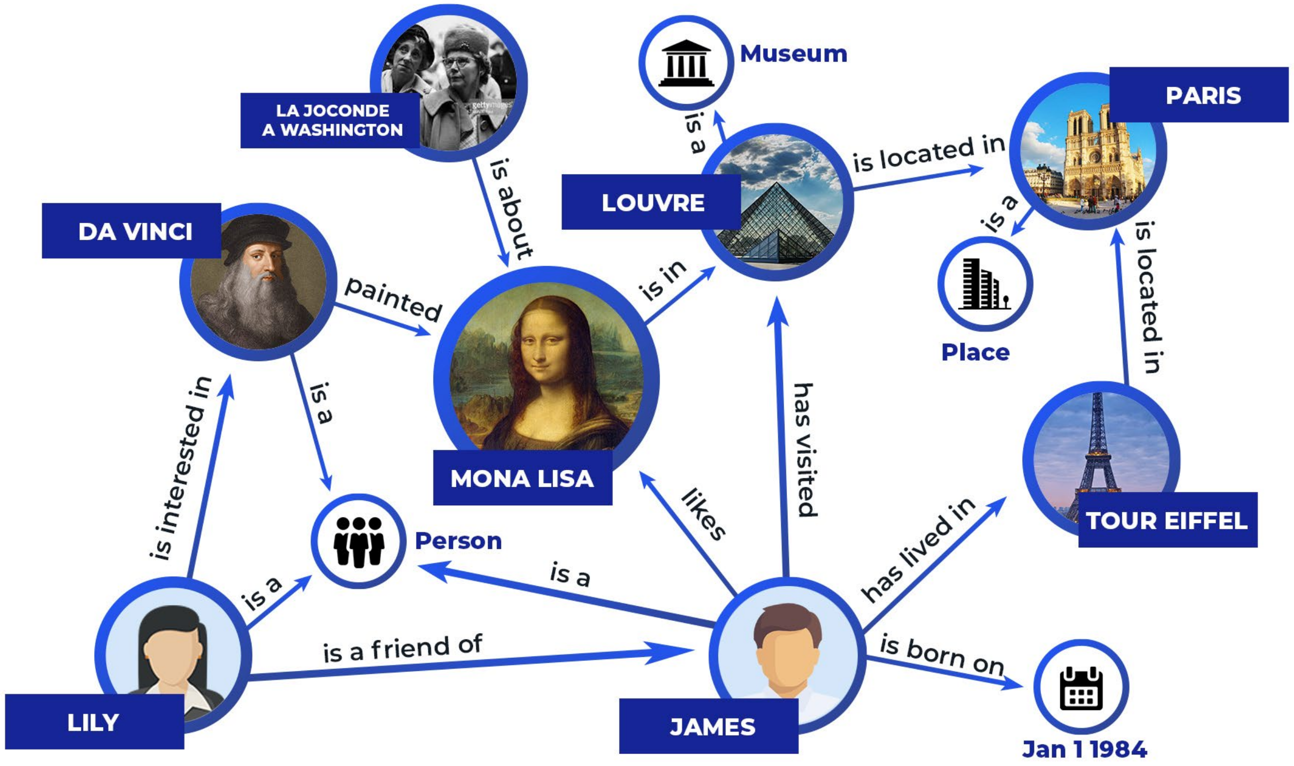

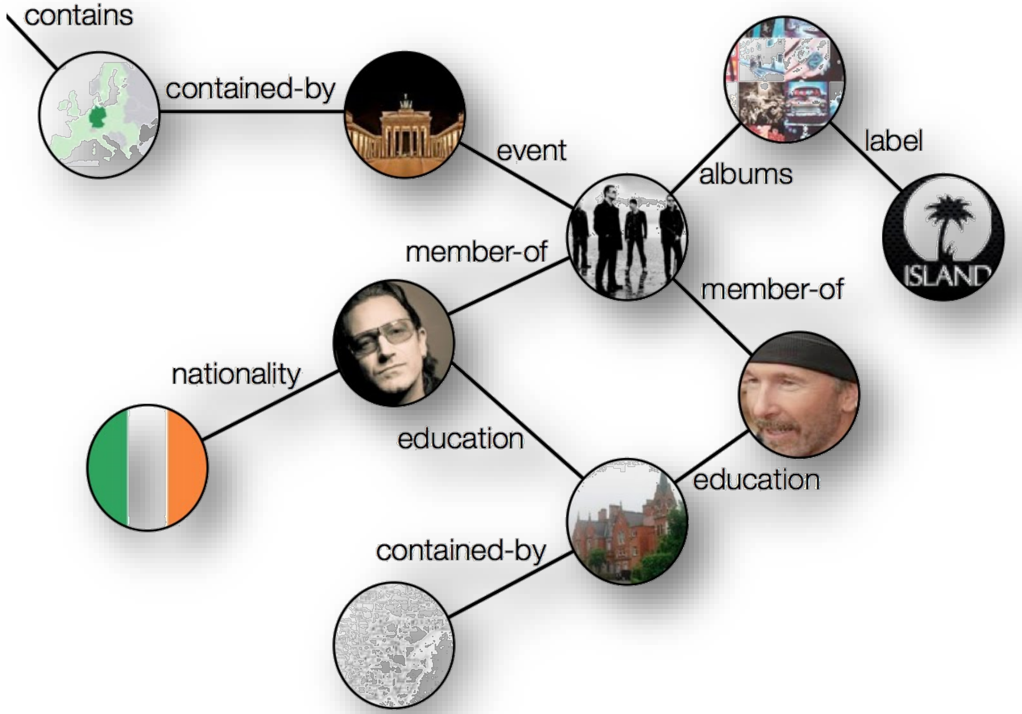

WordNet

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	$\text{water}^1 \rightarrow$ liquid
Hyponym	From concepts to subtypes	$\text{water}^1 \rightarrow$ seawater
Has-Part	From groups to their members	$\text{water}^1 \rightarrow$ oxygen
Part-of	From members to their groups	$\text{water}^1 \rightarrow$ ice
Antonym	Opposites	leader $\rightarrow$ follower

Taxonomy: is-a Relationship

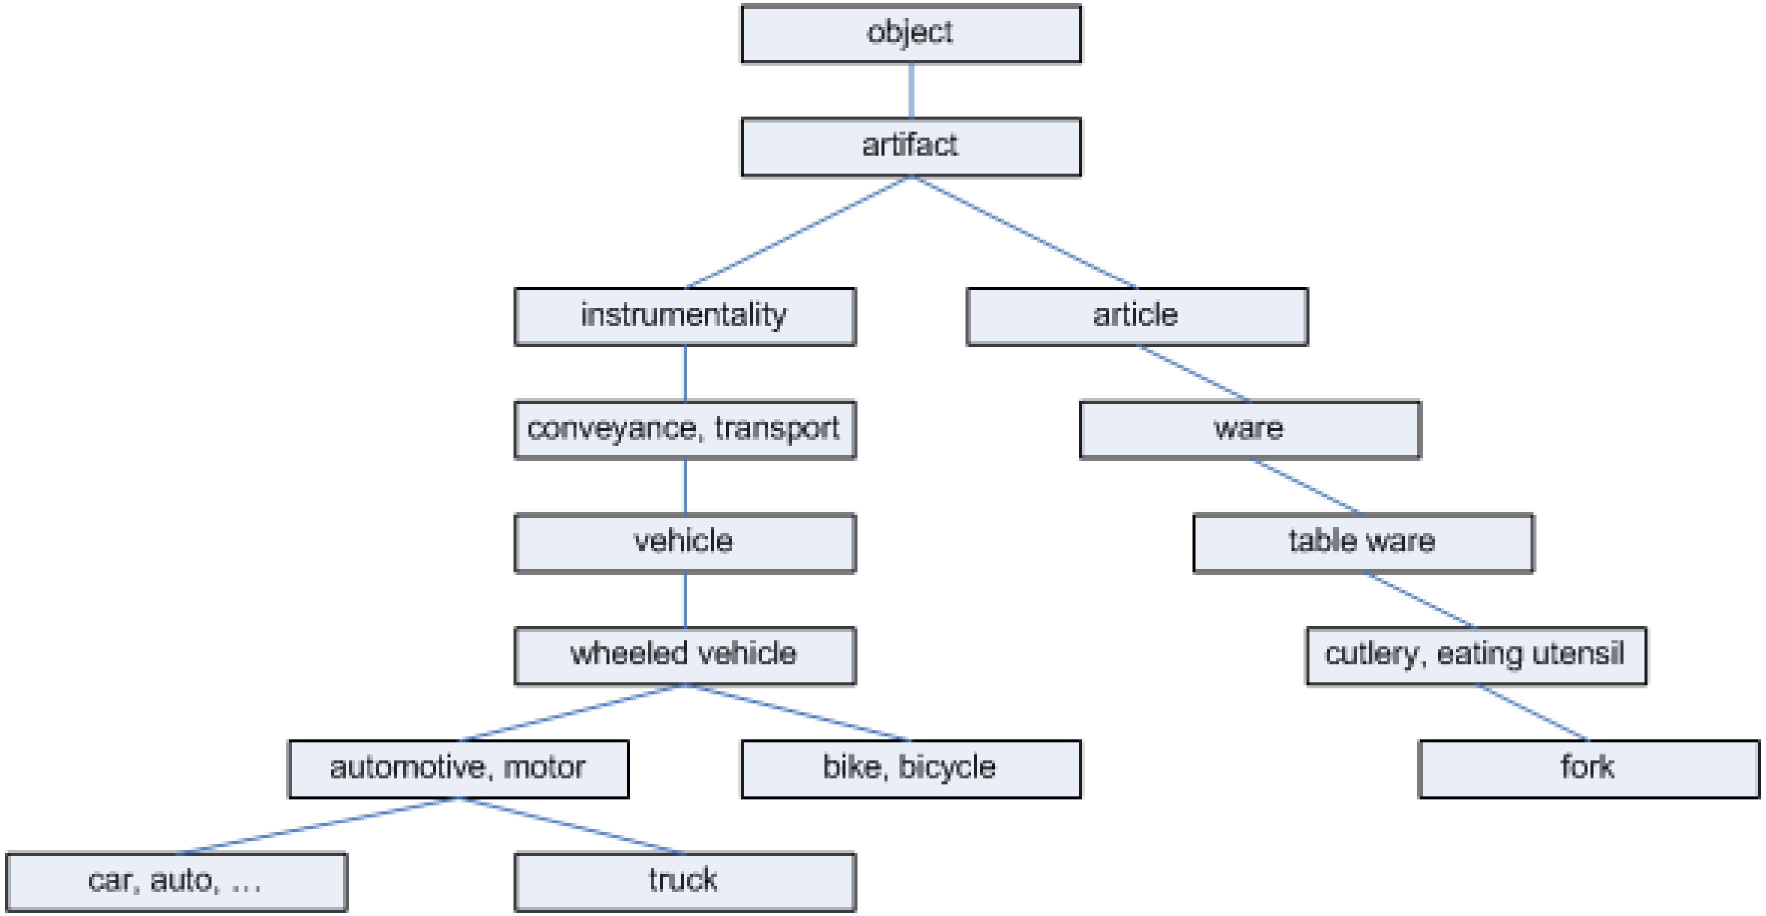

Partonomy: has-a Relationship

- 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 cowling 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 c
    - 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)

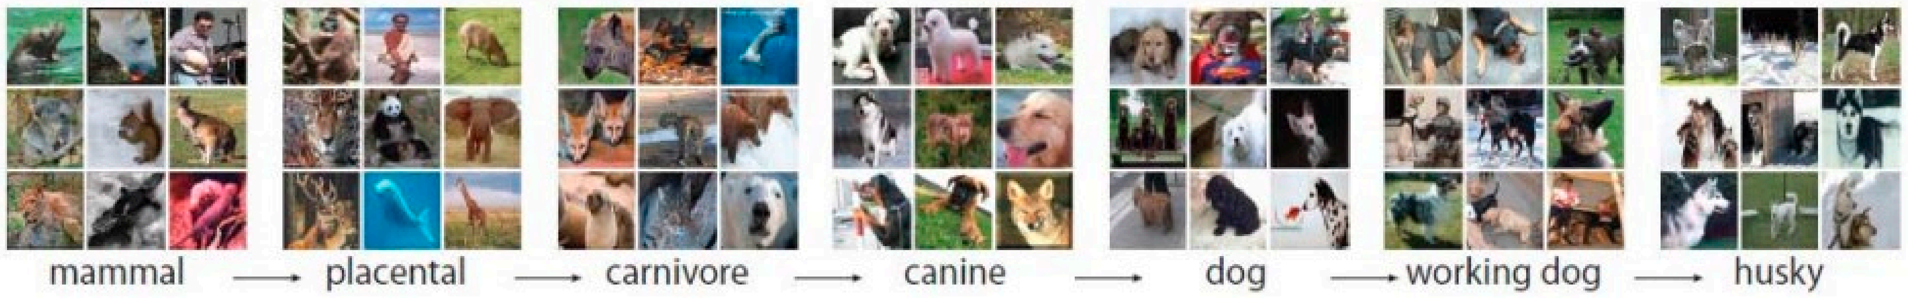

- S: (n) Eskimo dog, **husky** (breed of heavy-coated Arctic sled dog)
  - *direct hypernym* / ***inherited hypernym*** / *sister term*
    - S: (n) working dog (any of several breeds of usually large powerful dogs bred to work as draft animals and guard and guide dogs)
      - S: (n) dog, domestic dog, 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)

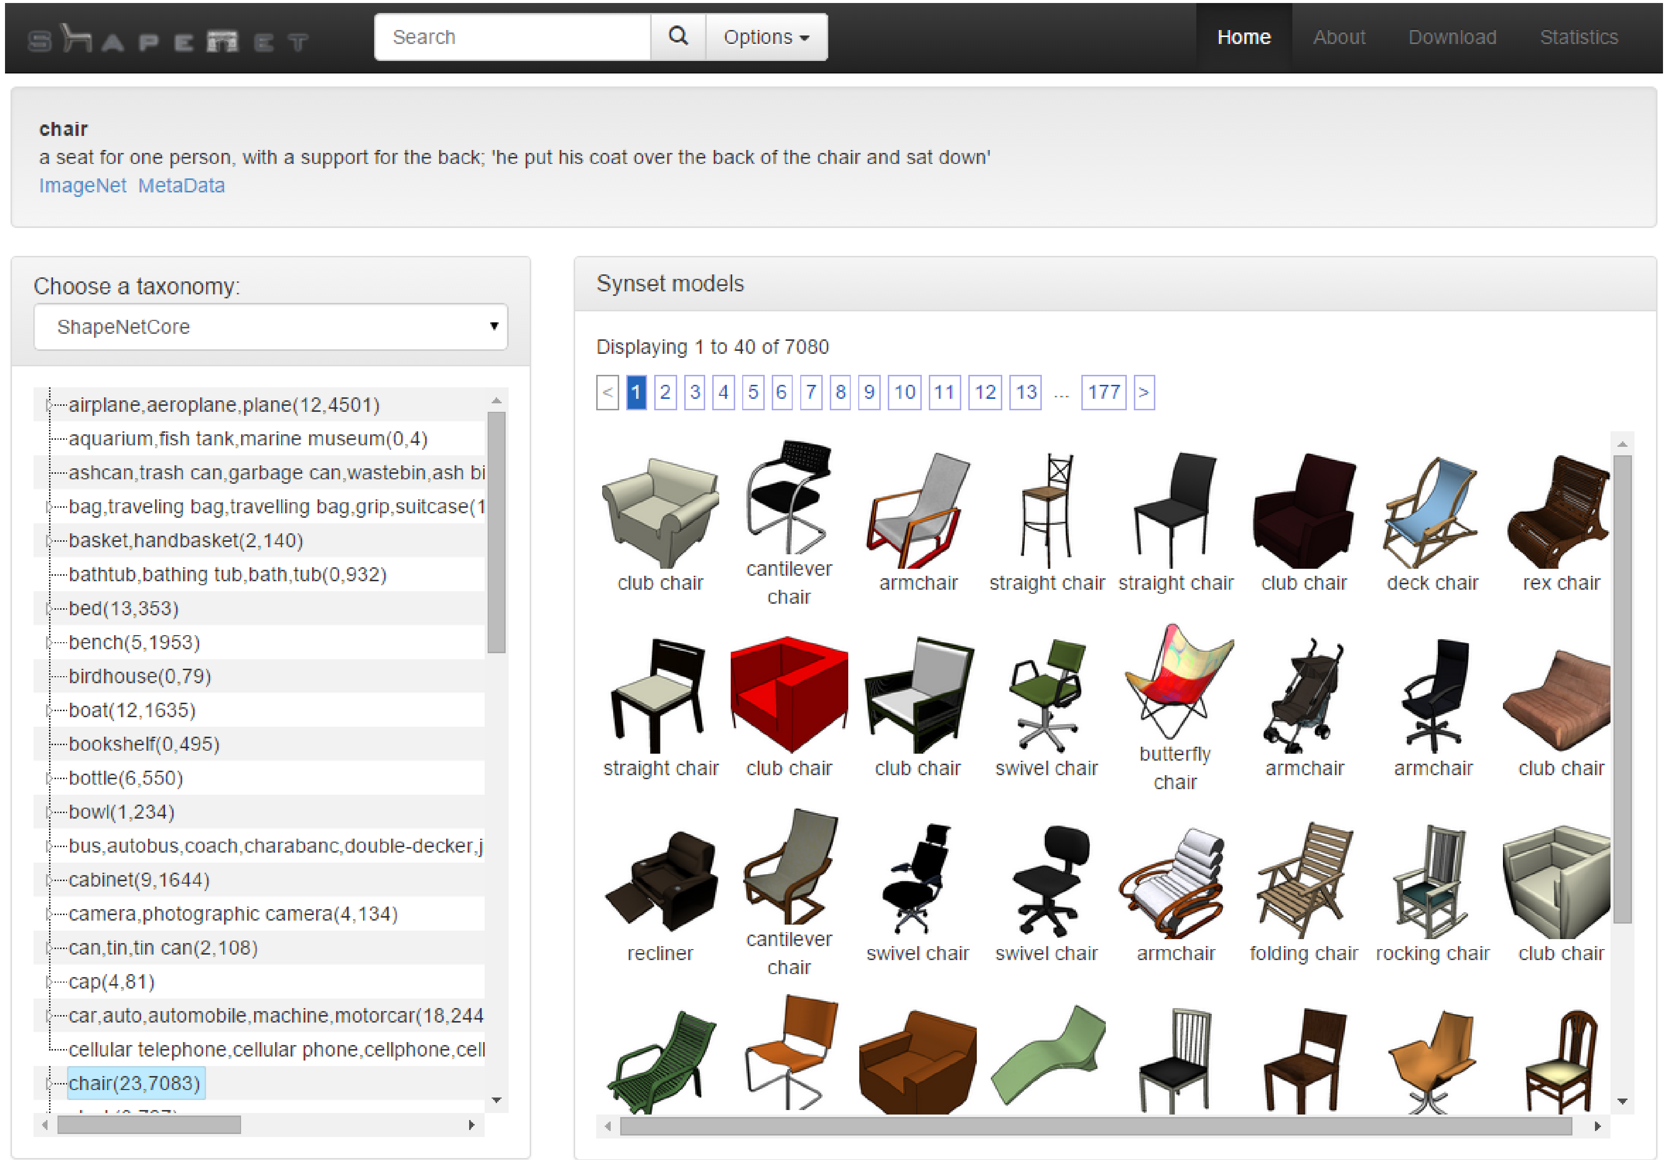

Object Knowledge

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

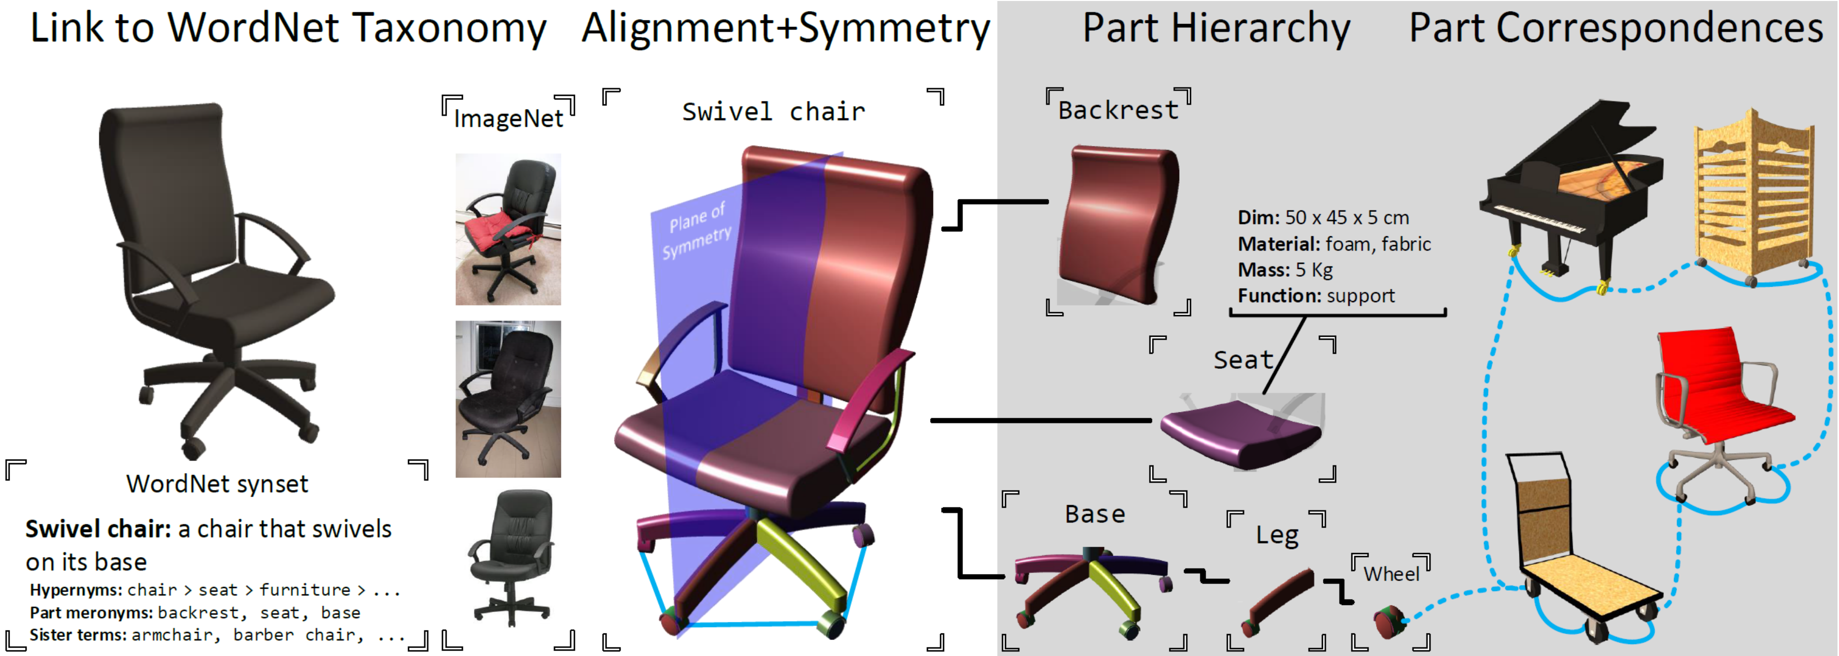

Object Interaction Knowledge

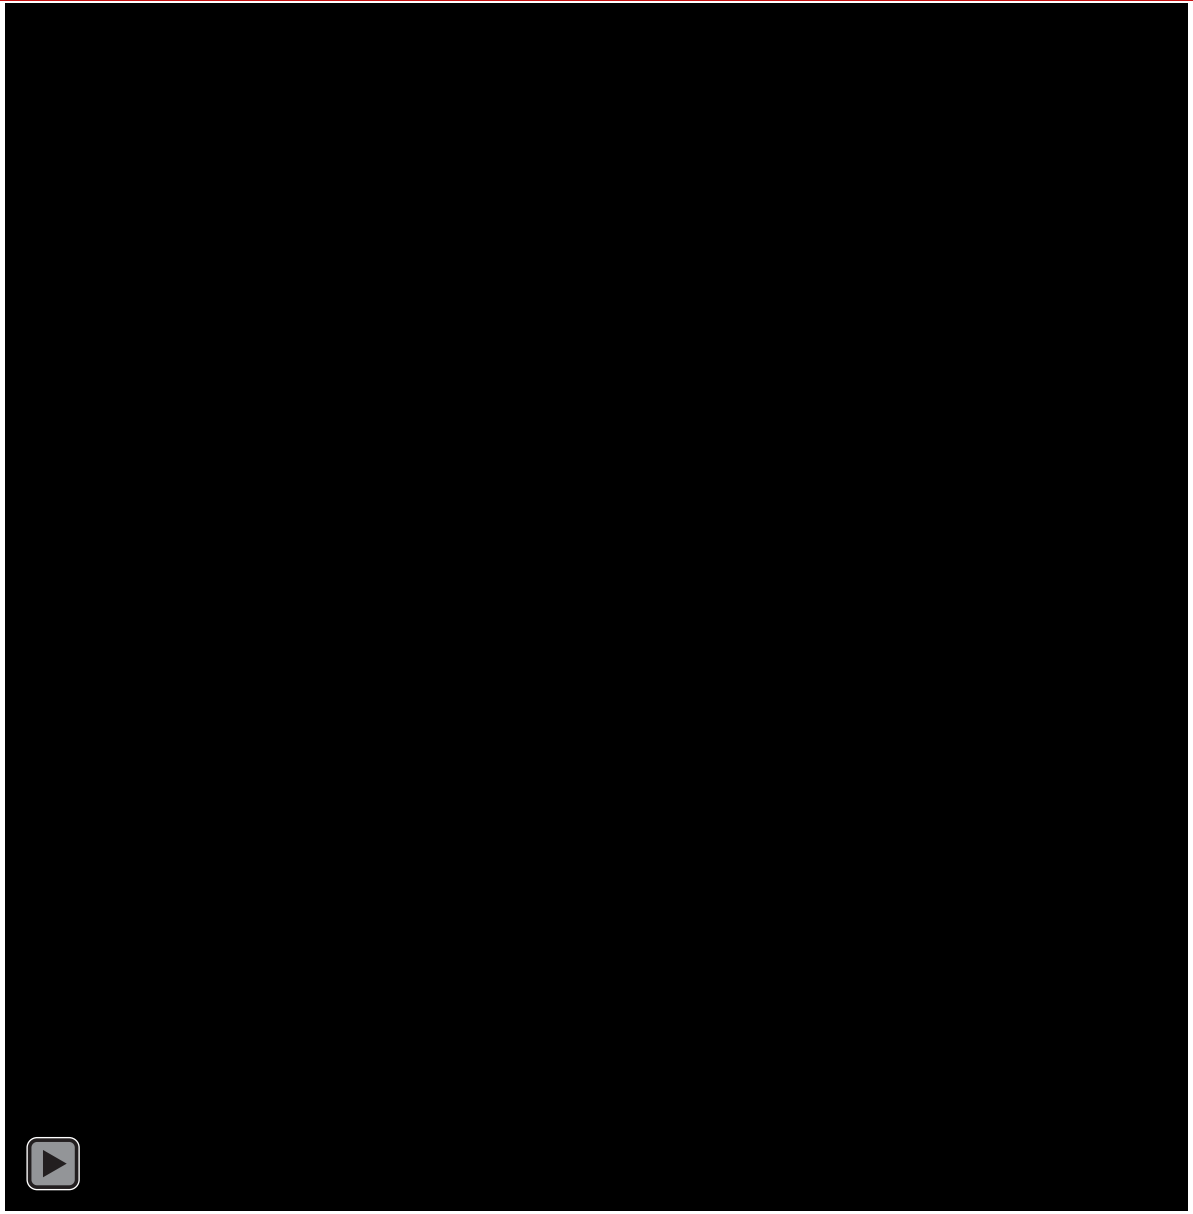

Vector Field to Histograms

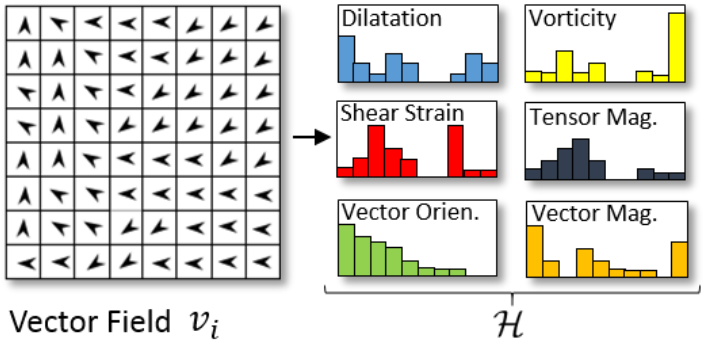

The Team

The Principals

- Leonidas (Leo) Guibas (CS & EE)
  - Instructor
- Davis Rempe (CS)
  - Course Assistant
- Despoina Paschalidou (CS)
  - Course Assistant
- Carrie Petersen (CS)
  - Admin


Leo


Davis


Despoina

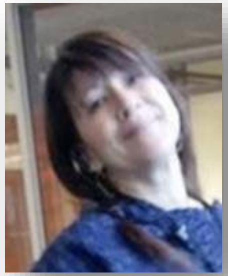
Carrie

Course Mechanics

CS233 / CME 251 Class Schedule

Monday	Wednesday
March 28	**March 30**
Introduction; Geometric and topological perspective on data analysis; Data representations; Learning on point clouds and graphs; Joint data analysis. Lecture Slides: Reading:	Linear algebraic techniques: principal components analysis (PCA), Kernel PCA. Lecture Slides: Reading:
April 4	**April 6**
Visual data sets: ImageNet and ShapeNet; Techniques for annotation and annotation transport. Lecture Slides: Reading:	Linear algebraic techniques: canonical correlation analysis (CCA). Multidimensional scaling (MDS). Lecture Slides: Reading: Homework 1 out.
April 11	**April 13**
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 18	**April 20**
Computational topology: topology review, complexes, homology groups. Lecture Slides: Reading:	Persistent homology, barcodes and persistence diagrams. Lecture Slides: Reading: Homework 1 due. Homework 2 out.

Monday	Wednesday
April 25	**April 27**
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 2	**May 4**
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 9	**May 11**
Class Midterm	Geometric deep learning; Volumetric and mesh CNNs for 3D geometry Lecture Slides: Reading:
May 16	**May 18**
Deep nets for pointclouds and applications to classification and segmentation. Lecture Slides: Pointnets Reading:	Functional spaces and functional maps, variations; map visualization. Lecture Slides: Reading: Homework 3 due. Homework 4 out.

CS233 / CME 251 Class Schedule

Monday	Wednesday
May 23	**May 25**
Networks of shapes and images; cycle consistency; map processing and latent spaces. Lecture Slides: Reading:	Encoding shape differences and shape variability. Lecture Slides: Reading:
May 30	**June 1**
Memorial day holiday -- no class	Class summary. Lecture Slides: Reading: Homework 4 due.

Class Mechanics

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

  http://graphics.stanford.edu/courses/cs233-22-spring
- Use Piazza (discussion group), Gradescope (for homework submissions)

Course Work: Assignments

- Assignment 1: Principal Components Analysis (PCA) and Canonical Correlation analysis (CCA)
- Assignment 2: Computational Topology - Simplicial Complexes, Homology, Persistence, Mapper Algorithm
- Assignment 3: Geometry Correspondences, Shape Matching, Multi-way Alignments
- Assignment 4: Deep Learning with 3D Point-Clouds

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 coupons for deep net training
- A midterm
- Class participation
- No final
- Special consideration for those taking the course credit / no credit

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”
- Challenge: a diversity of tools … LA, ML, Stat, optimization, geometry processing, computer vision, algebraic topology …

Data Has Shape

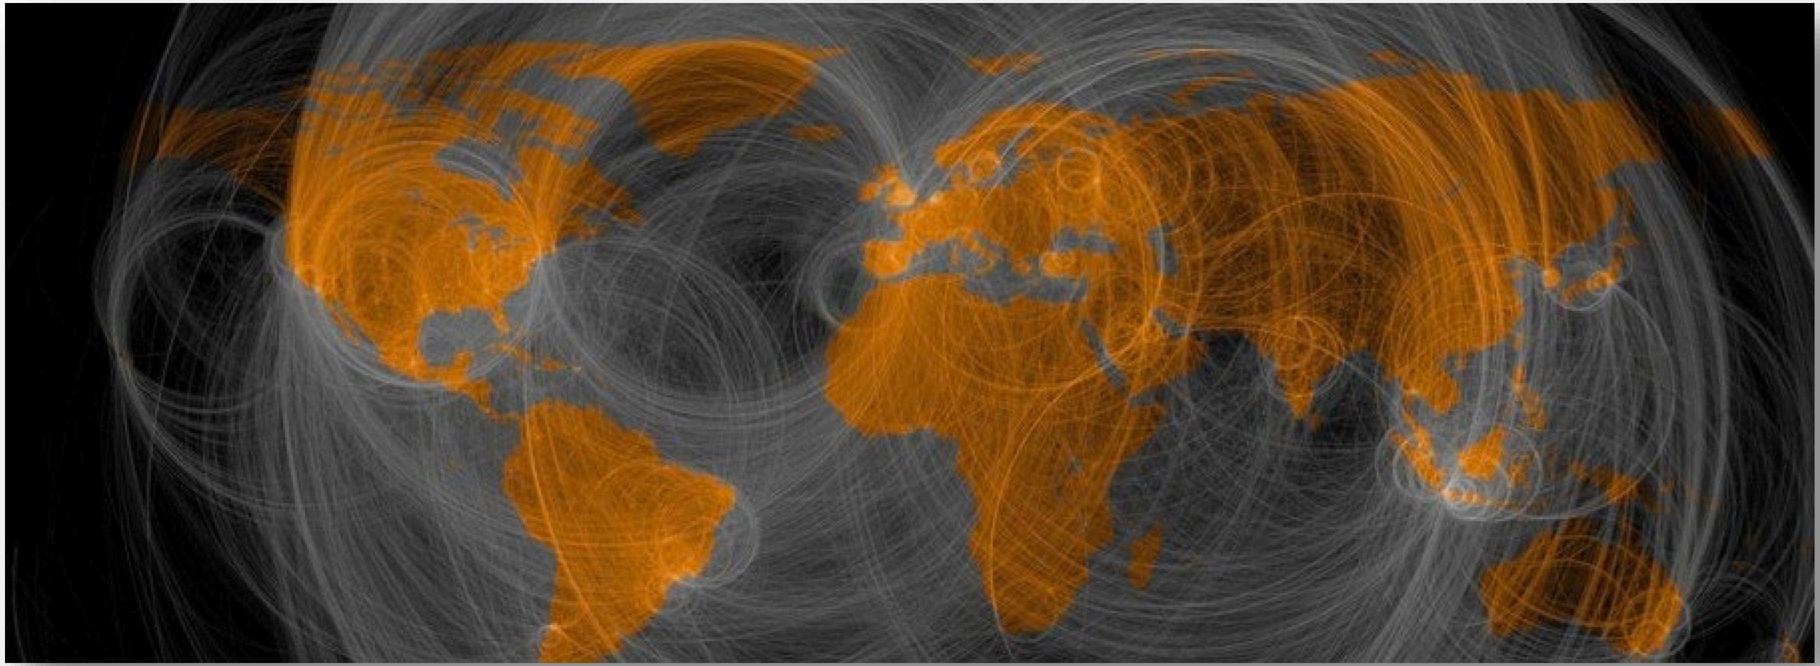

That’s All