CS233 Lecture 2: Visual Datasets

Li Yi
Stanford University

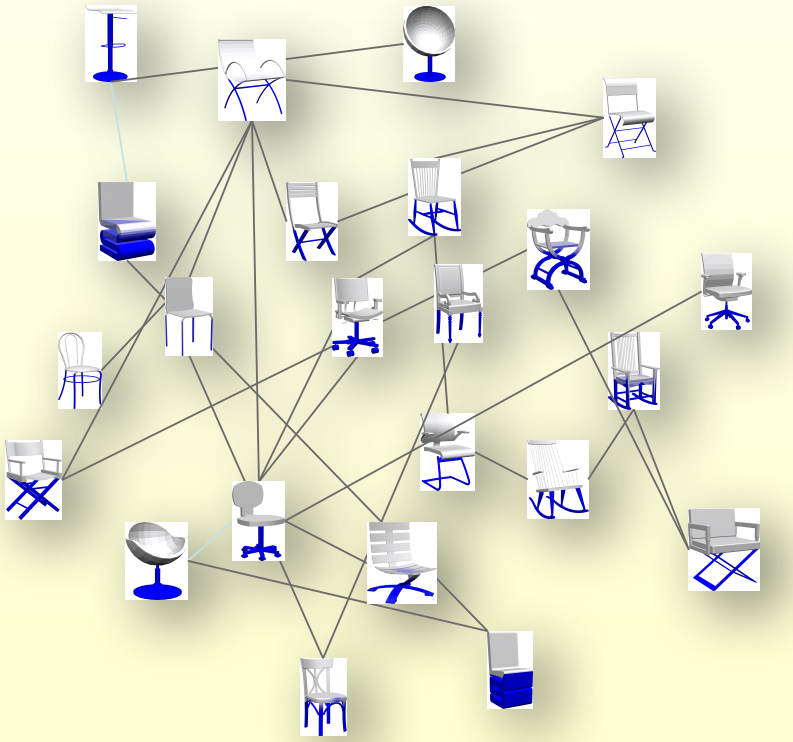

April 6, 2018

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontal networks
  - Annotation transportation in ShapeNet

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontal networks
  - Annotation transportation in ShapeNet

Goal of this Lecture

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

Goal of this Lecture

- Explain how ShapeNet are annotated

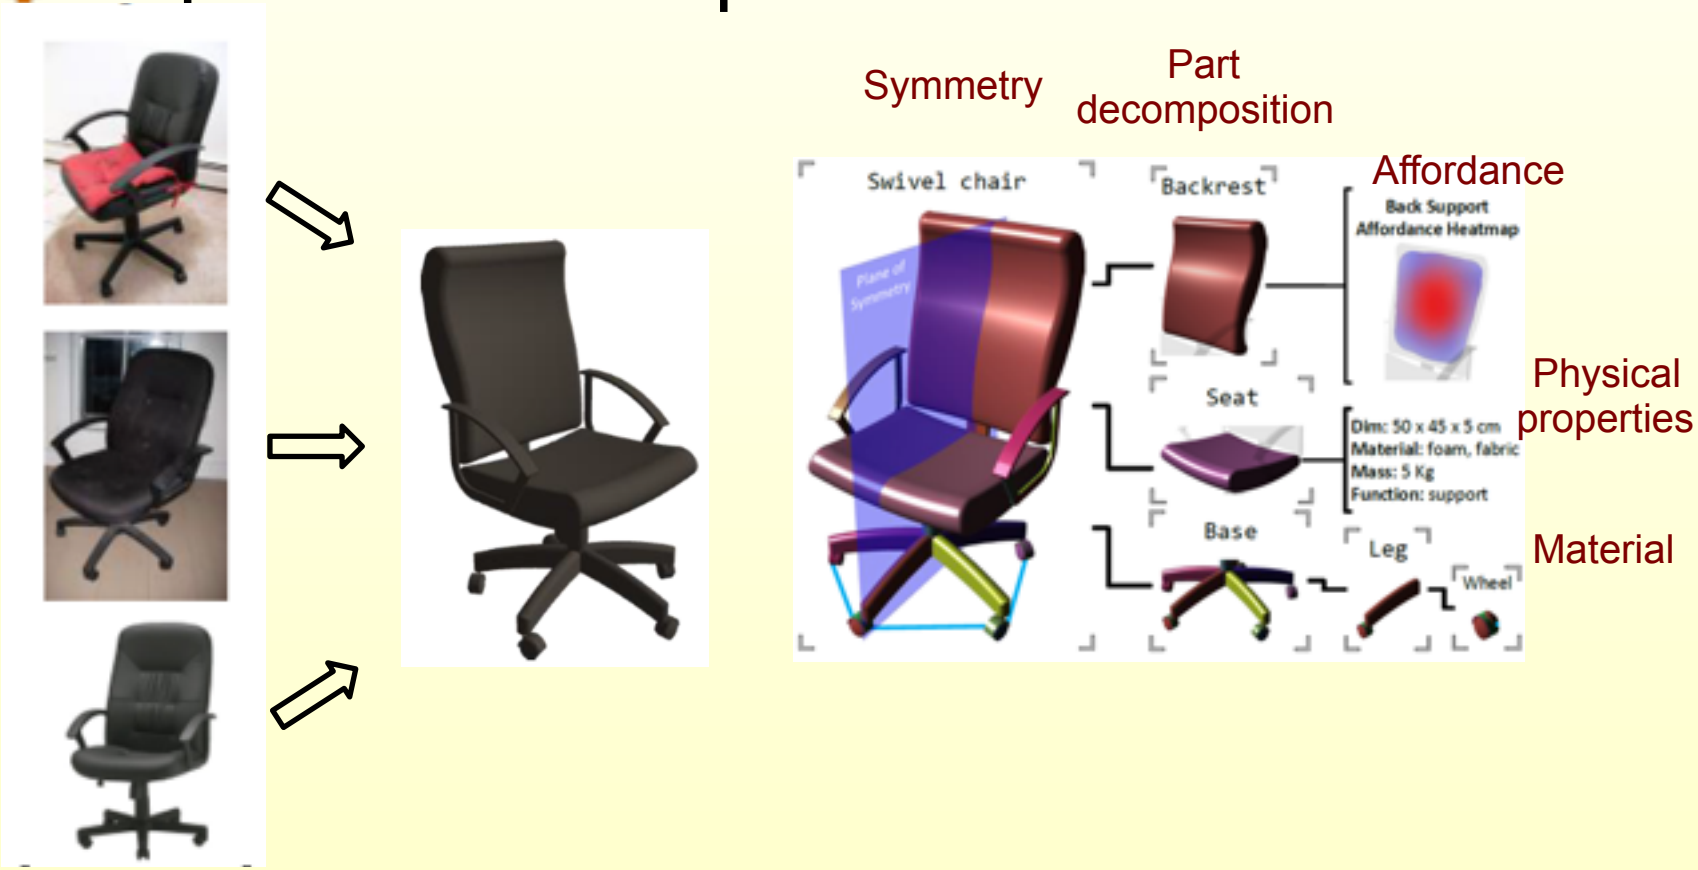

Goal of this Lecture

- Show examples of label transportation in a network

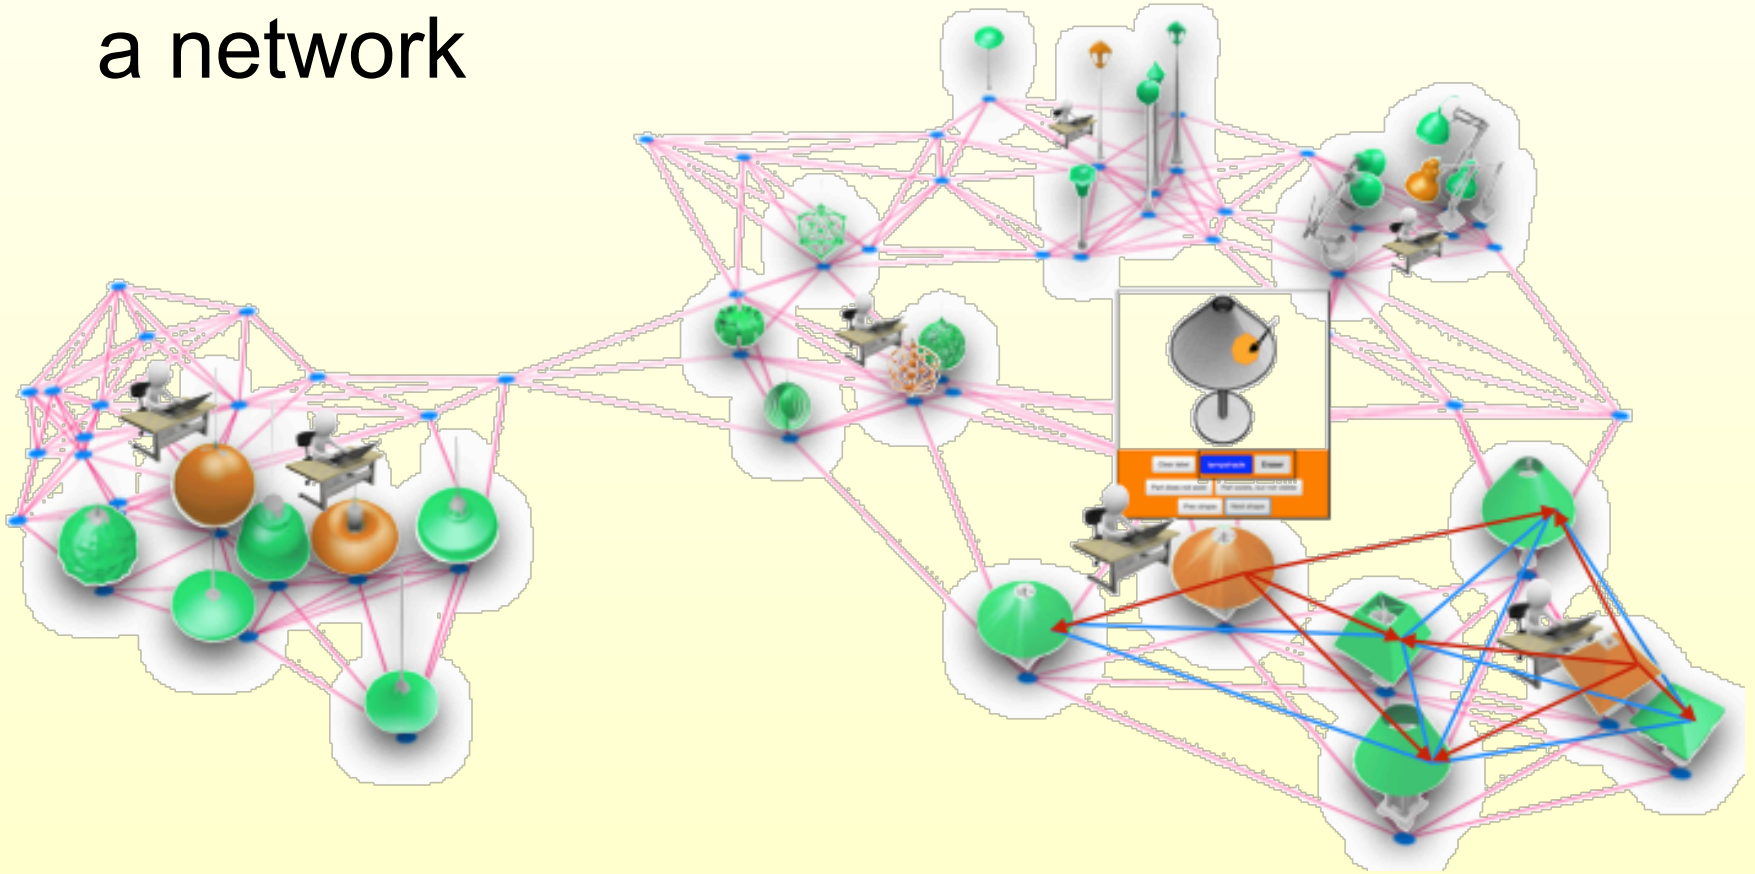

Goal of this Lecture

- Demo that networks can be an effective tool for organizing and annotating big data

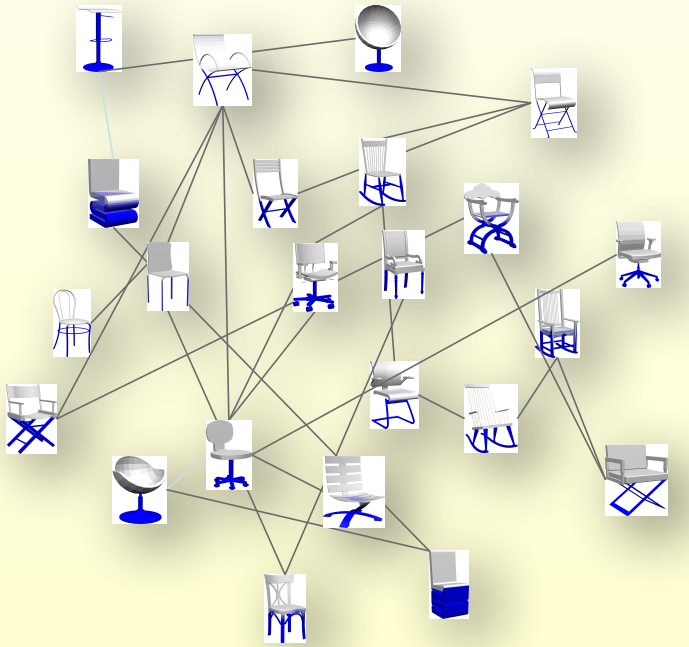

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontals network
  - Annotation transportation in ShapeNet

Semantic Networks

- Also known as **frame network**
- Encodes semantic relations between concepts.
- This is often used as a form of knowledge representation.
- It is a directed or undirected graph consisting of vertices, which represent concepts, and edges which represent concept relations

Examples of Semantic Net

Semantic Net in Lisp

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

Graph representation

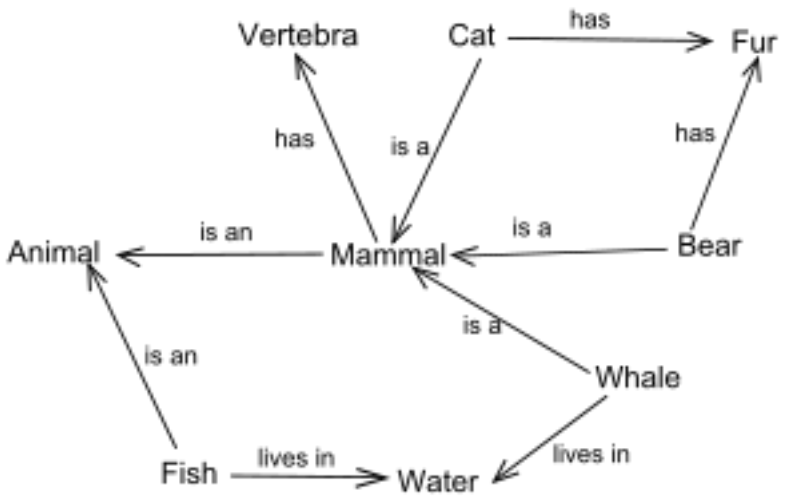

Examples of Semantic Net: WordNet

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

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

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

Examples of Semantic Net: 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

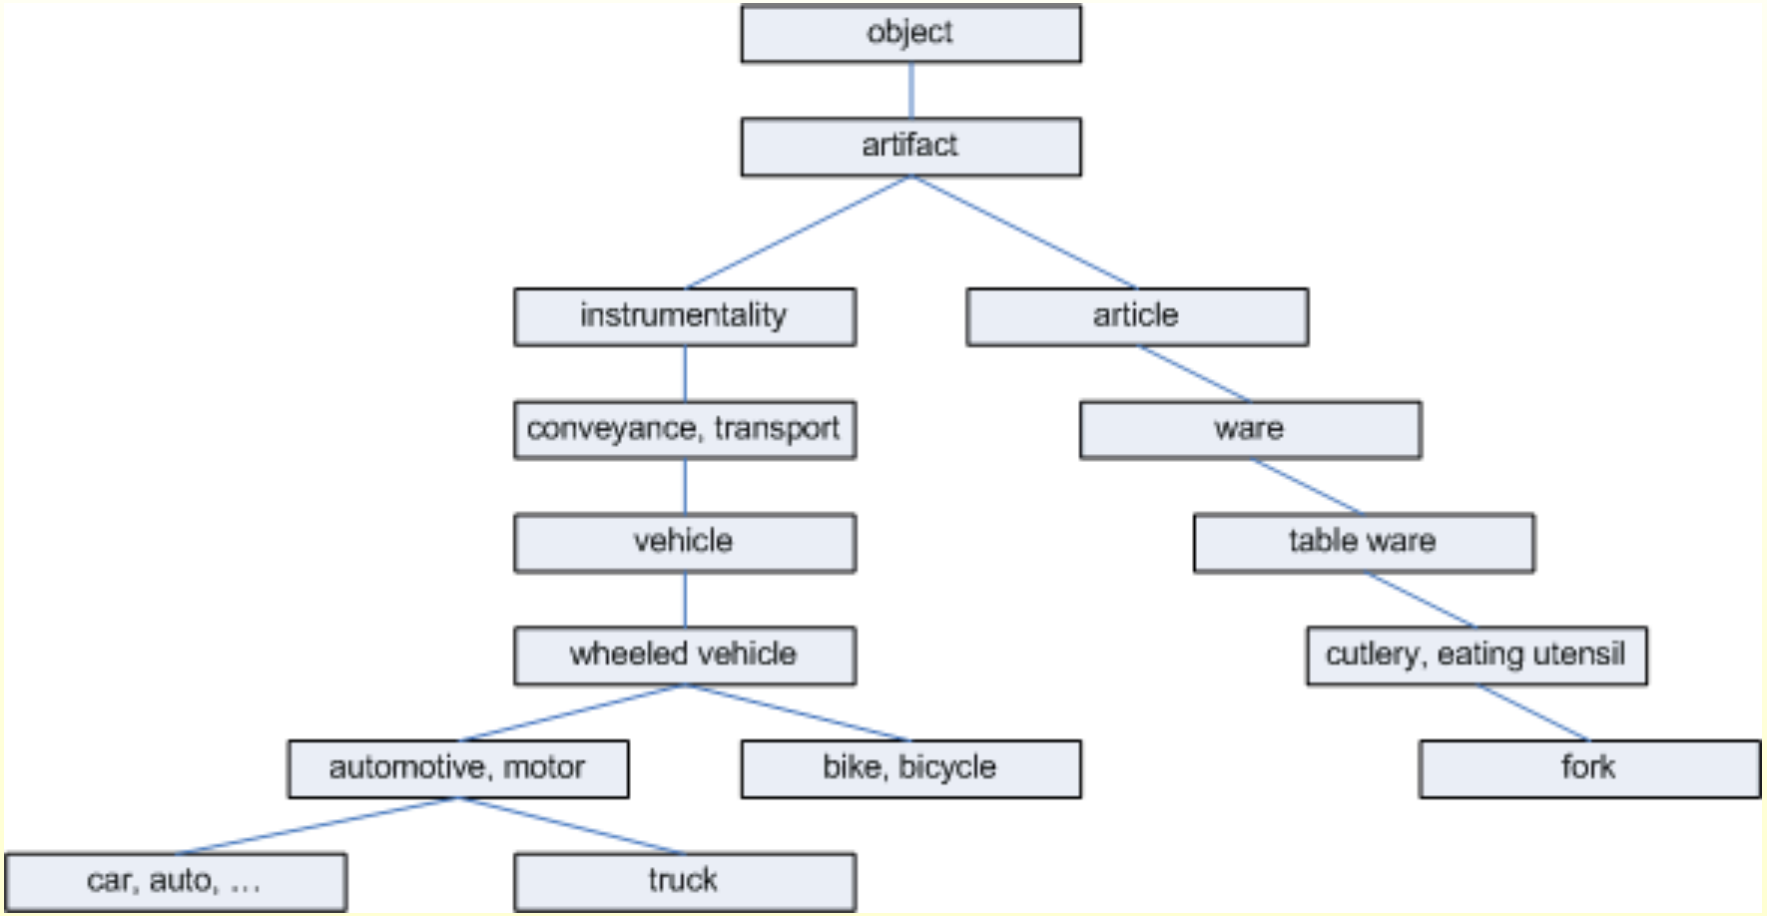

Partonomy: has-a Relationship

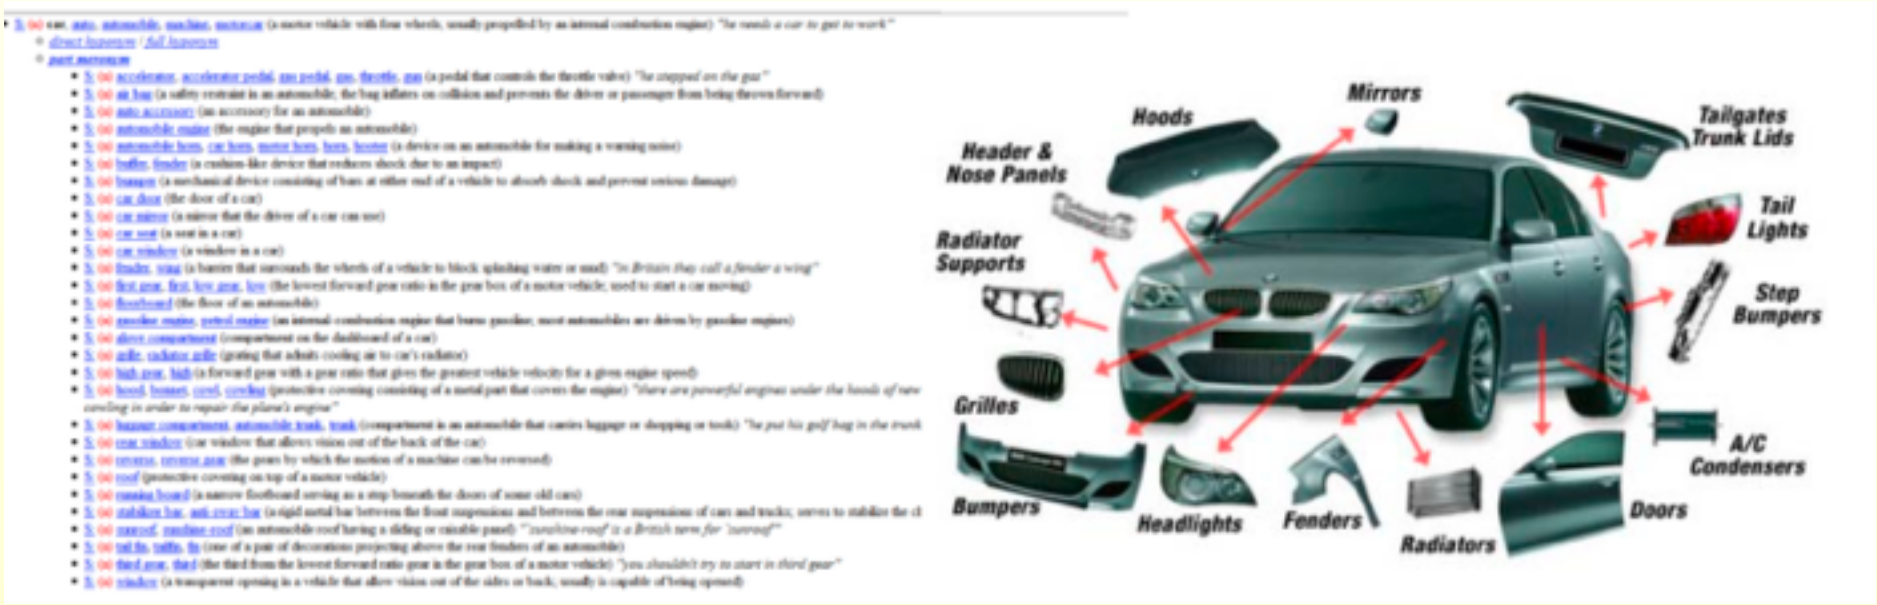

From Semantic Networks to Data Networks

- Instantiate *concepts* by *exemplars*
- Concepts from WordNet
  - Defined by properties (using language)

    • S: (n) chair (a seat for one person, with a support for the back)
- Exemplars from sensor data
  - images (ImageNet)
  - 3D shapes (ShapeNet)
  - videos

Why Go from a Semantic Network to a Data Network ?

- "A picture is worth a thousand words"
- Concepts and their relationships emerge directly from data

http://www.image-net.org

IMAGENET

14,197,122 images, 21841 synsets indexed

Explore Download Challenges Publications CoolStuff About

Not logged in. Login | Signup

ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. Currently we have an average of over five hundred images per node. We hope ImageNet will become a useful resource for researchers, educators, students and all of you who share our passion for pictures.

Click here to learn more about ImageNet, Click here to join the ImageNet mailing list.

What do these images have in common? *Find out!*

Check out the ImageNet Challenge 2015

© 2014 Stanford Vision Lab, Stanford University, Princeton University support@image-net.org Copyright infringement

Li Fei-Fei, Jia Deng, Hao Su, etc

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

IMAGENET is a knowledge ontology

- Taxonomy
- Partonomy

Slide Credit: Li Fei-Fei

Limitations of ImageNet
(From Knowledge Representation Perspective)

- Captures only shallow information in images

Object name, bounding box location

- Geometric and physical knowledge of objects is missing (e.g. ShapeNet)
- Relationships among objects are missing (e.g. VisualGenome)

Geometry and Physical Knowledge of Objects are Important

3D Models are Digital Representations Closest to the Physical Form of Objects

- Shapes as polygonal meshes

Knowledge as Properties Attached to 3D Shapes

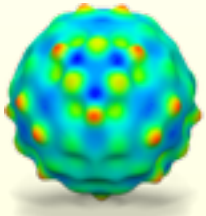

Curvature

Parts

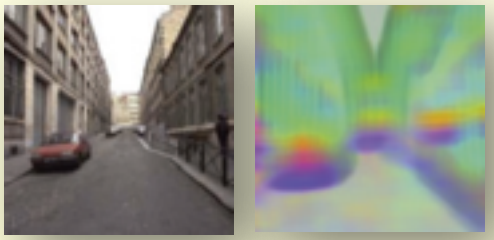

SIFT flow, C. Liu 2011

Knowledge as a tower over data

Unified Knowledge Representation in 3D

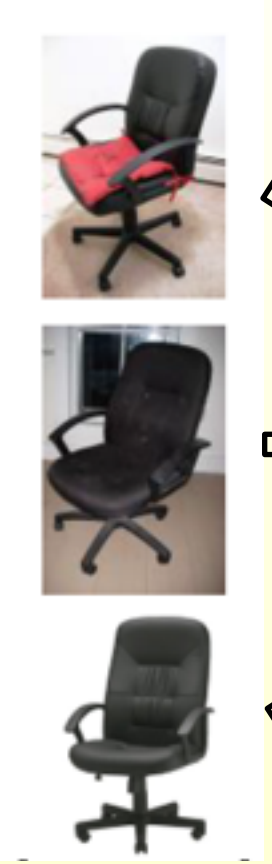

Symmetry

Unified Knowledge Representation in 3D

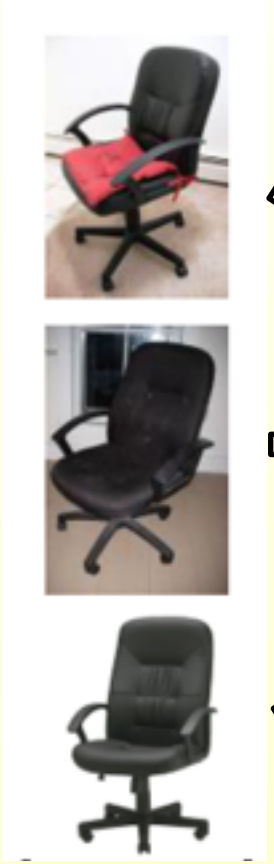

Symmetry

Part decomposition

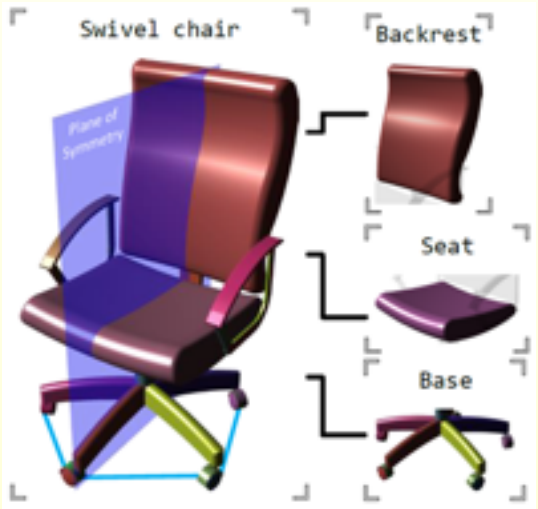

Unified Knowledge Representation in 3D

Symmetry

Part decomposition

Affordance

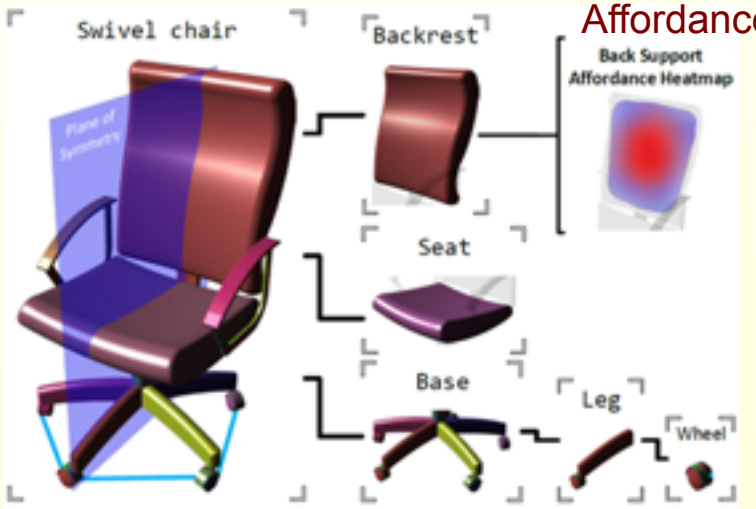

Unified Knowledge Representation in 3D

Symmetry

Part decomposition

Affordance

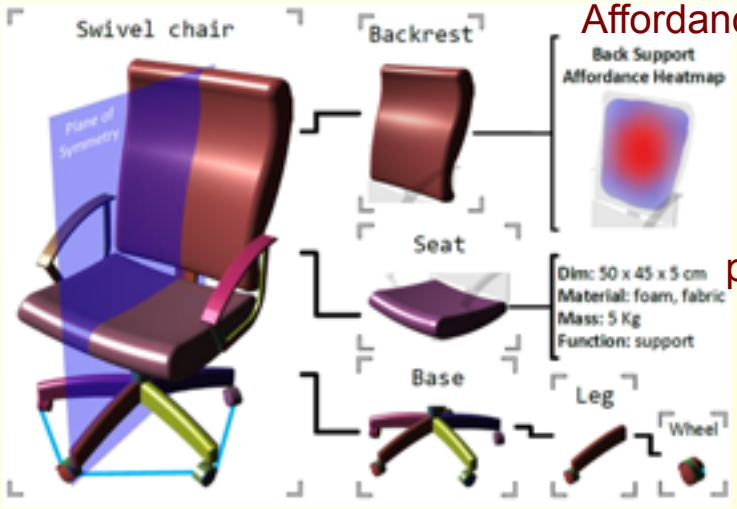

Physical properties

Material

ShapeNet: A Large-scale 3D Model Database

...

~3 million models in total

~2,000 classes

Rich annotations

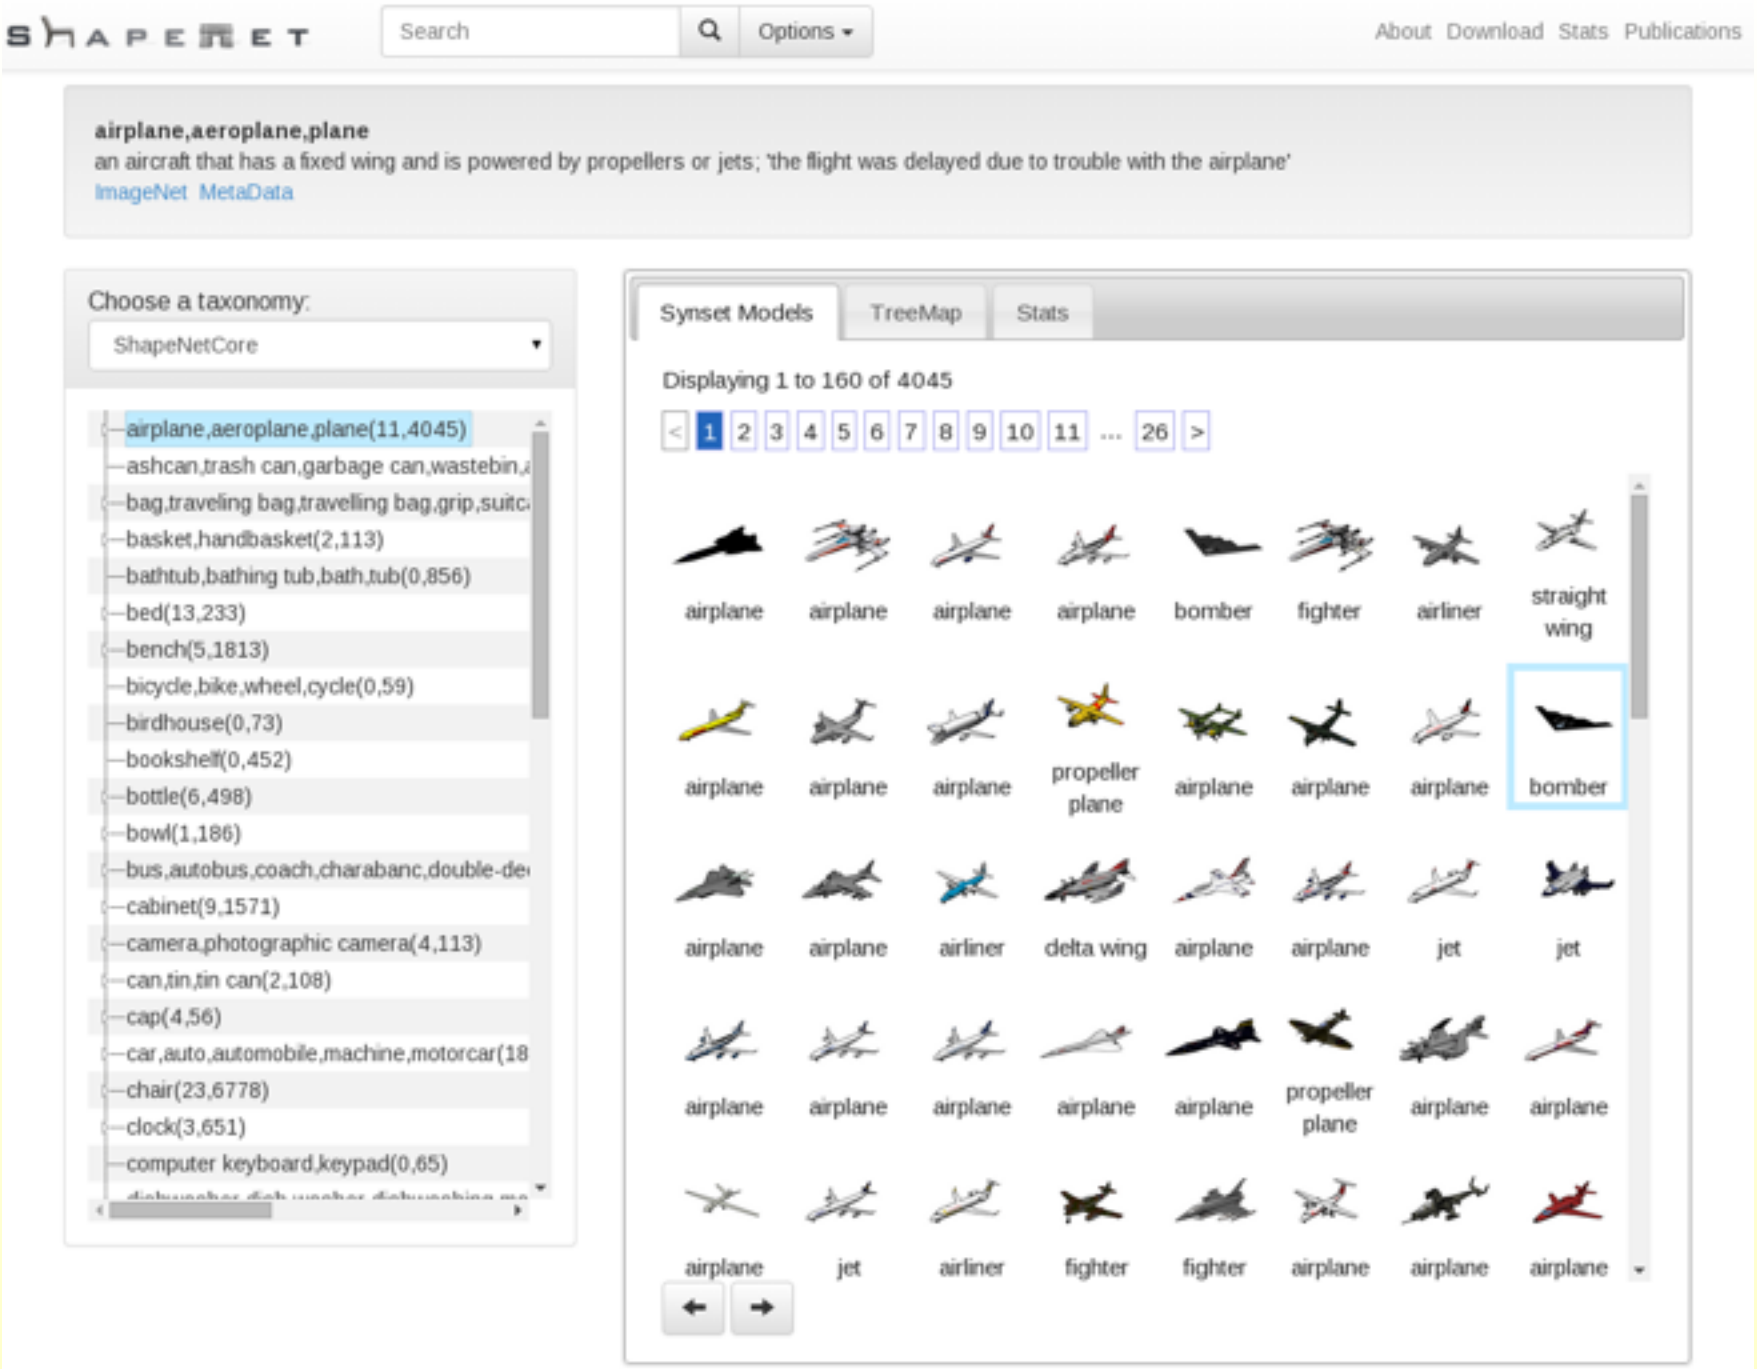
SHAPENET
Search
Options
About Download Stats Publications
airplane,aeroplane,plane
an aircraft that has a fixed wing and is powered by propellers or jets; 'the flight was delayed due to trouble with the airplane'
ImageNet MetaData
Choose a taxonomy:
ShapeNetCore
airplane,aeroplane,plane(11,4045)
ashcan,trash can,garbage can,wastebin,
bag,traveling bag,travelling bag,grip,suitc
basket,handbasket(2,113)
bathtub,bathing tub,bath,tub(0,856)
bed(13,233)
bench(5,1813)
bicycle,bike,wheel,cycle(0,59)
birdhouse(0,73)
bookshelf(0,452)
bottle(6,498)
bowl(1,186)
bus,autobus,coach,charabanc,double-de
cabinet(9,1571)
camera,photographic camera(4,113)
can,tin,tin can(2,108)
cap(4,56)
car,auto,automobile,machine,motorcar(18
chair(23,6778)
clock(3,651)
computer keyboard,keypad(0,65)
Synset Models
TreeMap
Stats
Displaying 1 to 160 of 4045
< 1 2 3 4 5 6 7 8 9 10 11 ... 26 >
airplane airplane airplane airplane bomber fighter airliner straight wing
airplane airplane airplane propeller plane airplane airplane airplane bomber
airplane airplane airliner delta wing airplane airplane jet jet
airplane airplane airplane airplane airplane propeller plane airplane airplane
airplane jet airliner fighter fighter airplane airplane airplane

Where is in ShapeNet currently?

- ShapeNetCore
  - 51,300 textured 3D models classified into 55 classes, mostly man-made objects
  - Mesh, point cloud, volumetric representations are provided
  - Consistent orientation within each class
  - Semantic part annotation for a subset
  - Physical dimensions and weights

Where is in ShapeNet currently?

- ShapeNetCore
- ShapeNetSem
  - 12,000 textured models classified into 270 categories, indoor objects
  - Mesh, volumetric representations are provided
  - Consistent orientation within each class
  - Physical dimensions and weights

How could these Annotations be Acquired?

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontal networks
  - Annotation transportation in ShapeNet

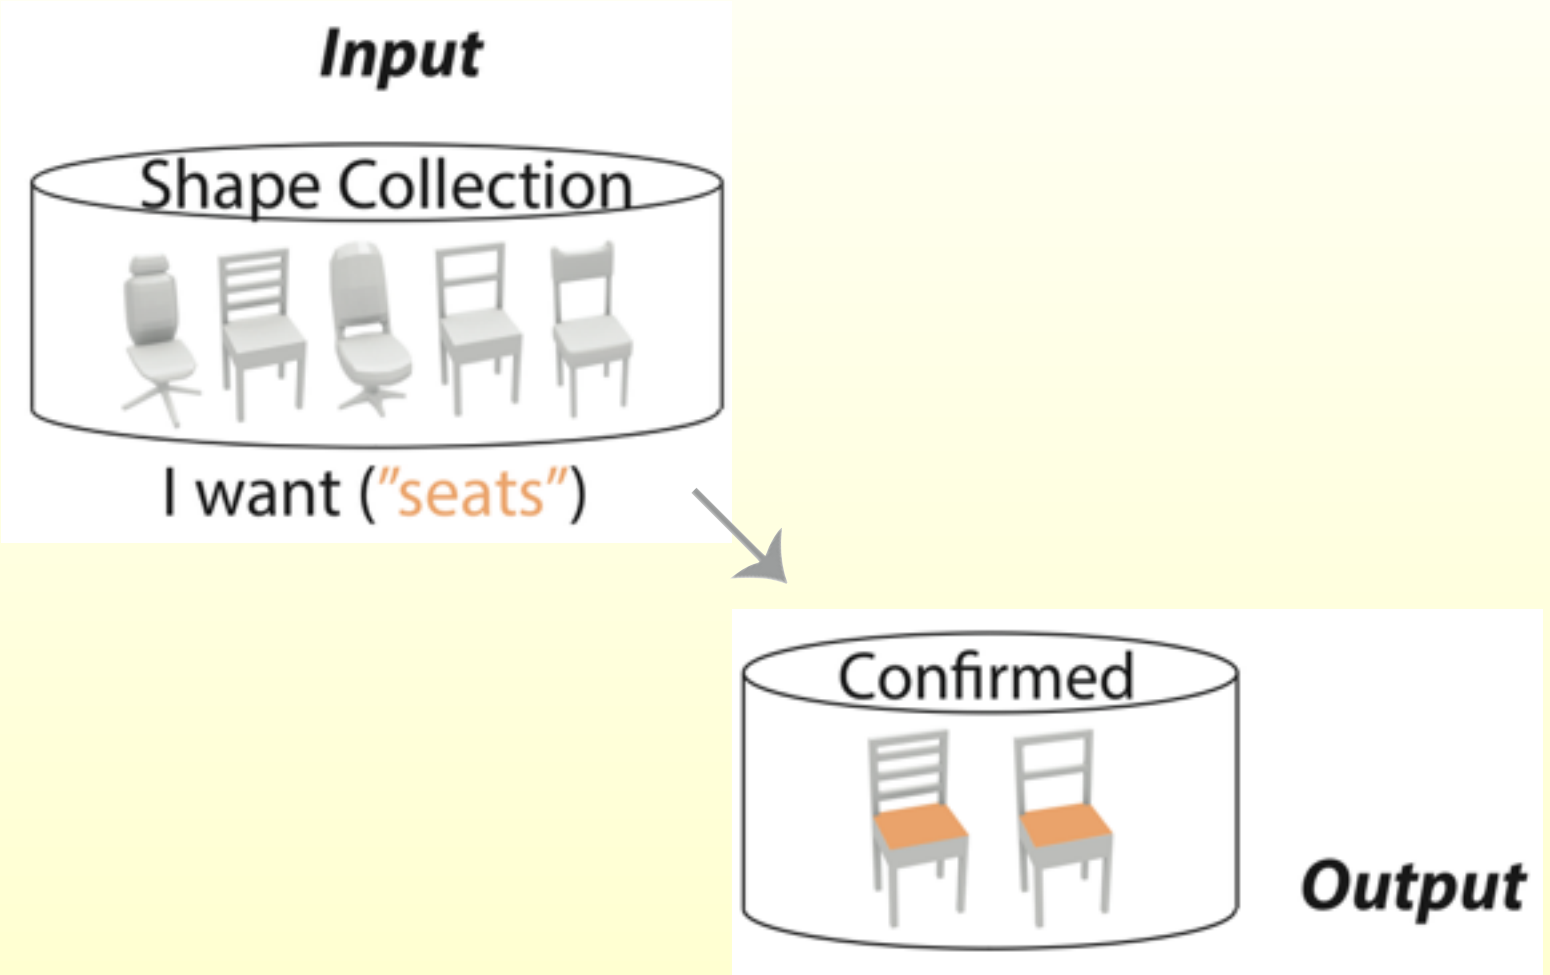
Input
Shape Collection
I want ("seats")
Confirmed
Output

Manually Annotate Everything?

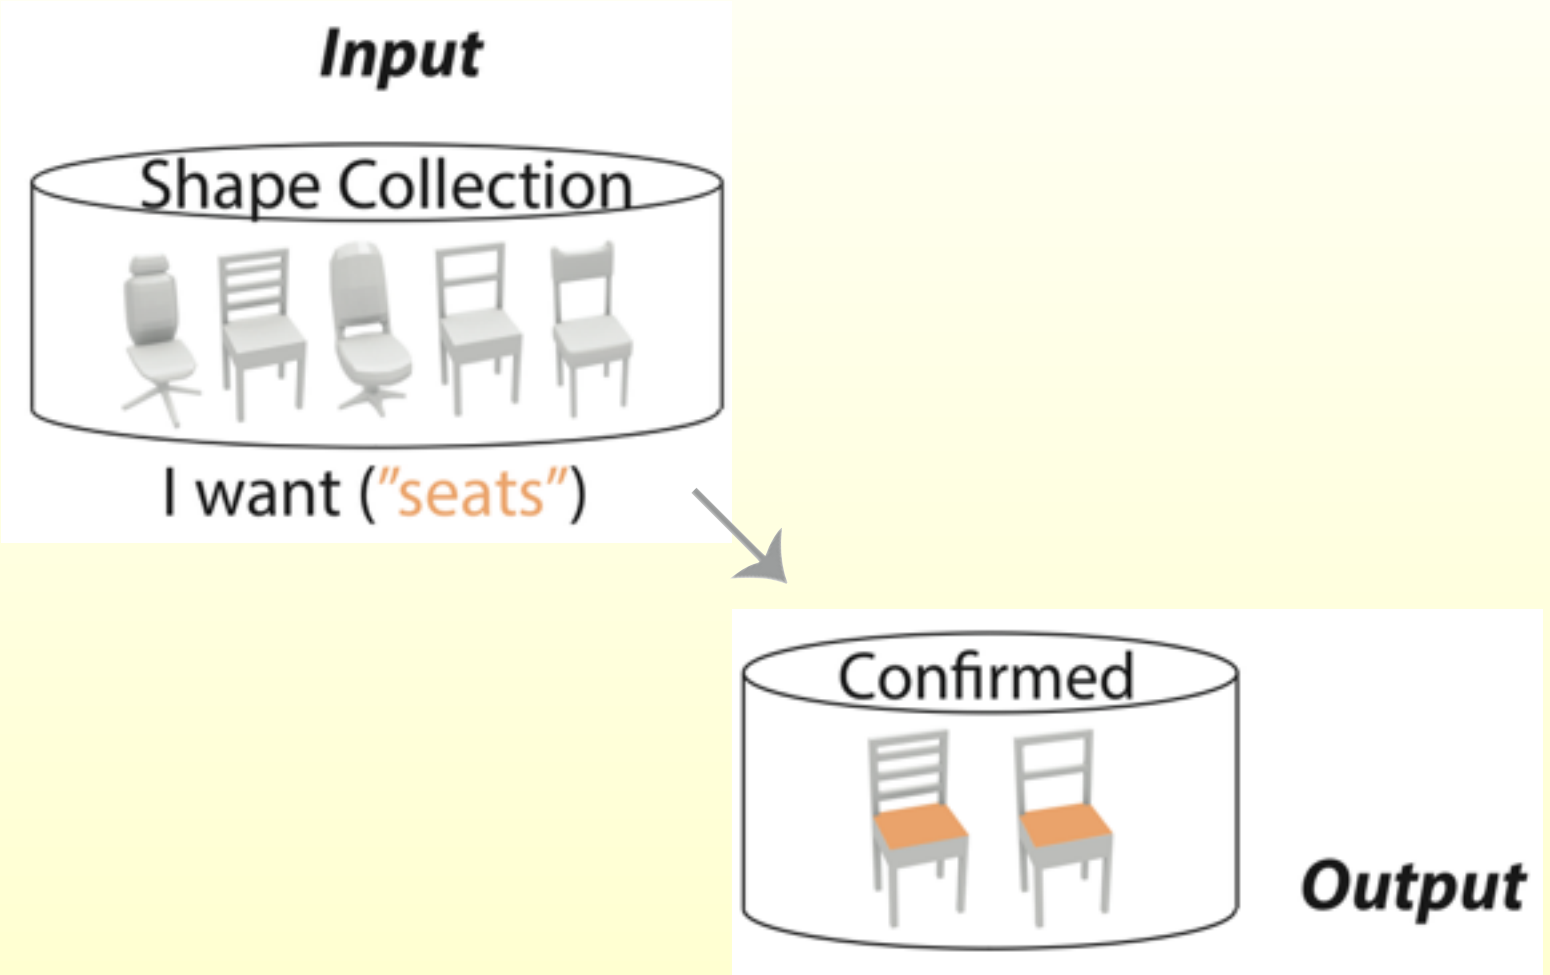

Manually Annotate Everything?

Time consuming

- *Millions of 3D models*
- *Constantly evolving*

Manually Annotate Everything?

Time consuming

Low efficiency

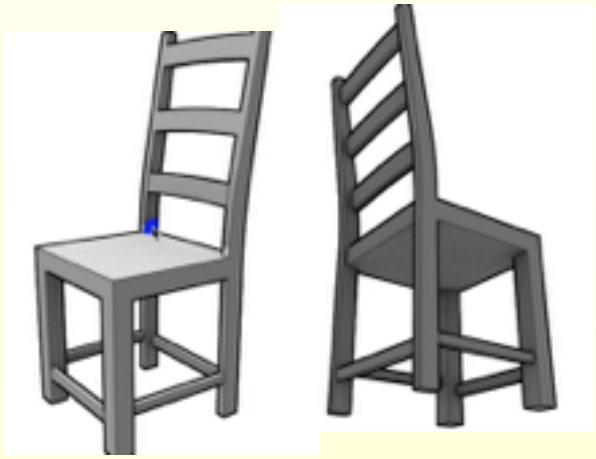

- *Millions of 3D models*
- *Constantly evolving*

- *3D annotation takes lots of time*

Manually Annotate Everything?

Time consuming

ShapeNet

3D Warehouse

yobi 3D

TurboSquid

- *Millions of 3D models*
- *Constantly evolving*

Low efficiency

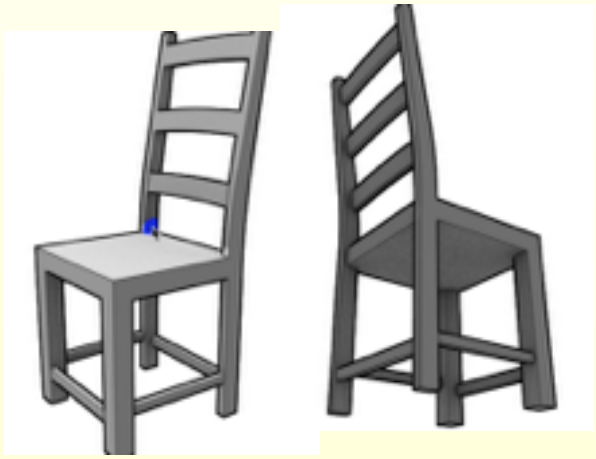

- *3D annotation takes lots of time*

Not satisfactory accuracy

Airplane tail

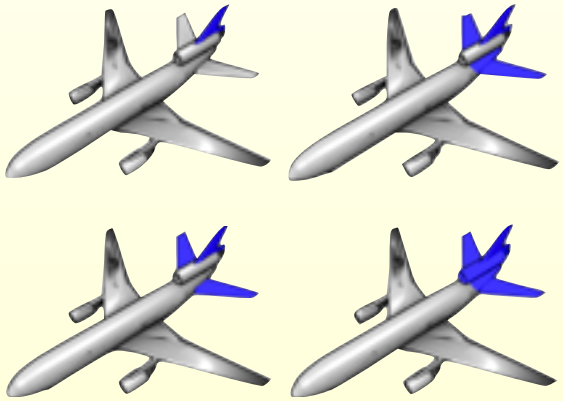

- *Semantic ambiguity*

M

Tim

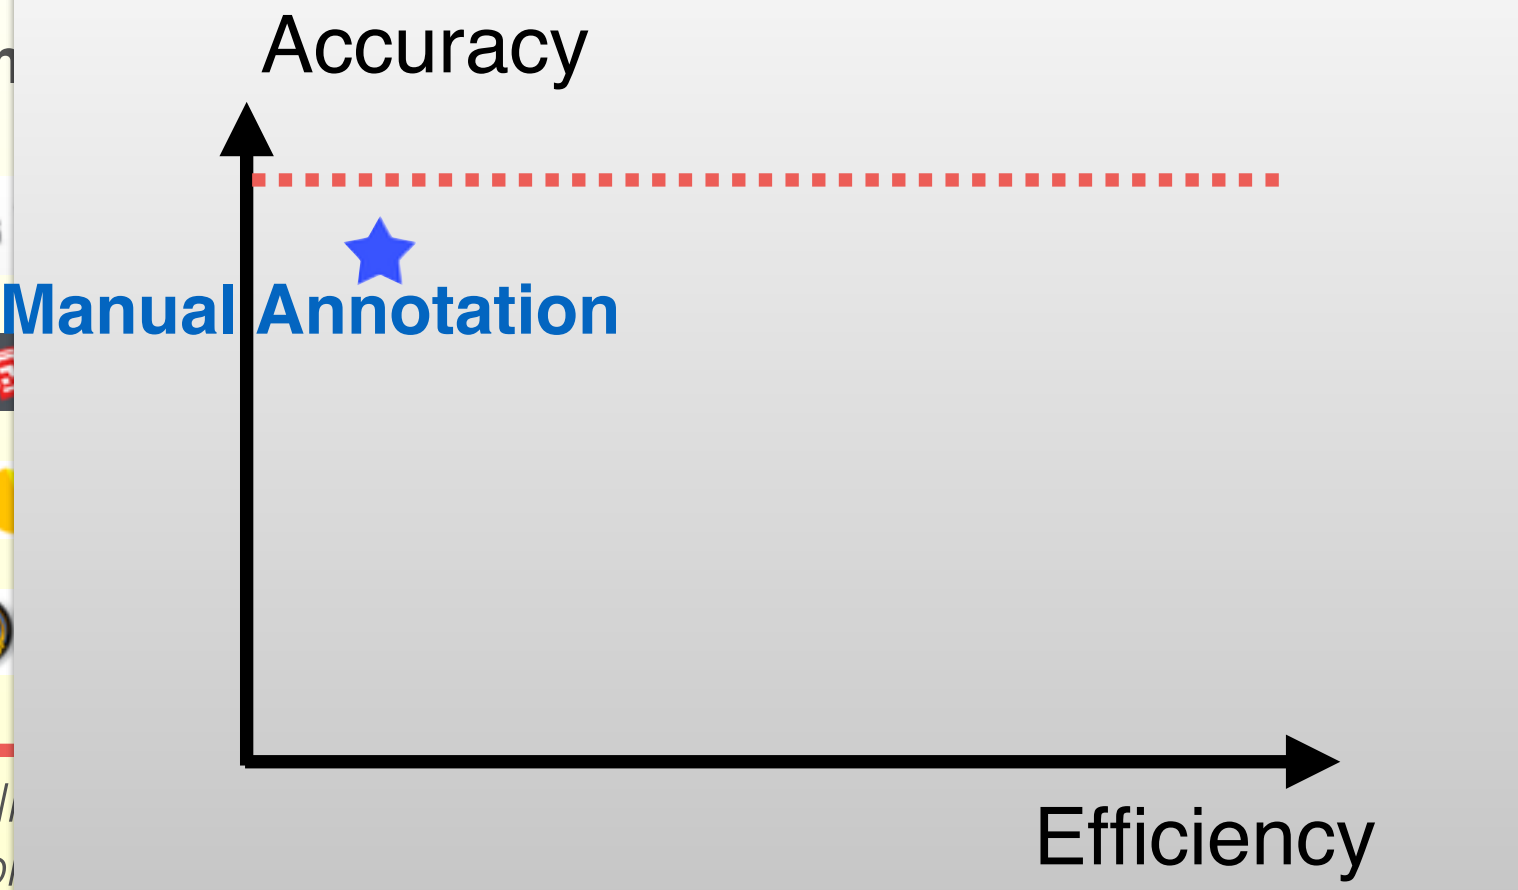

- *Mil*
- *Co*

ty

Algorithm Prediction?

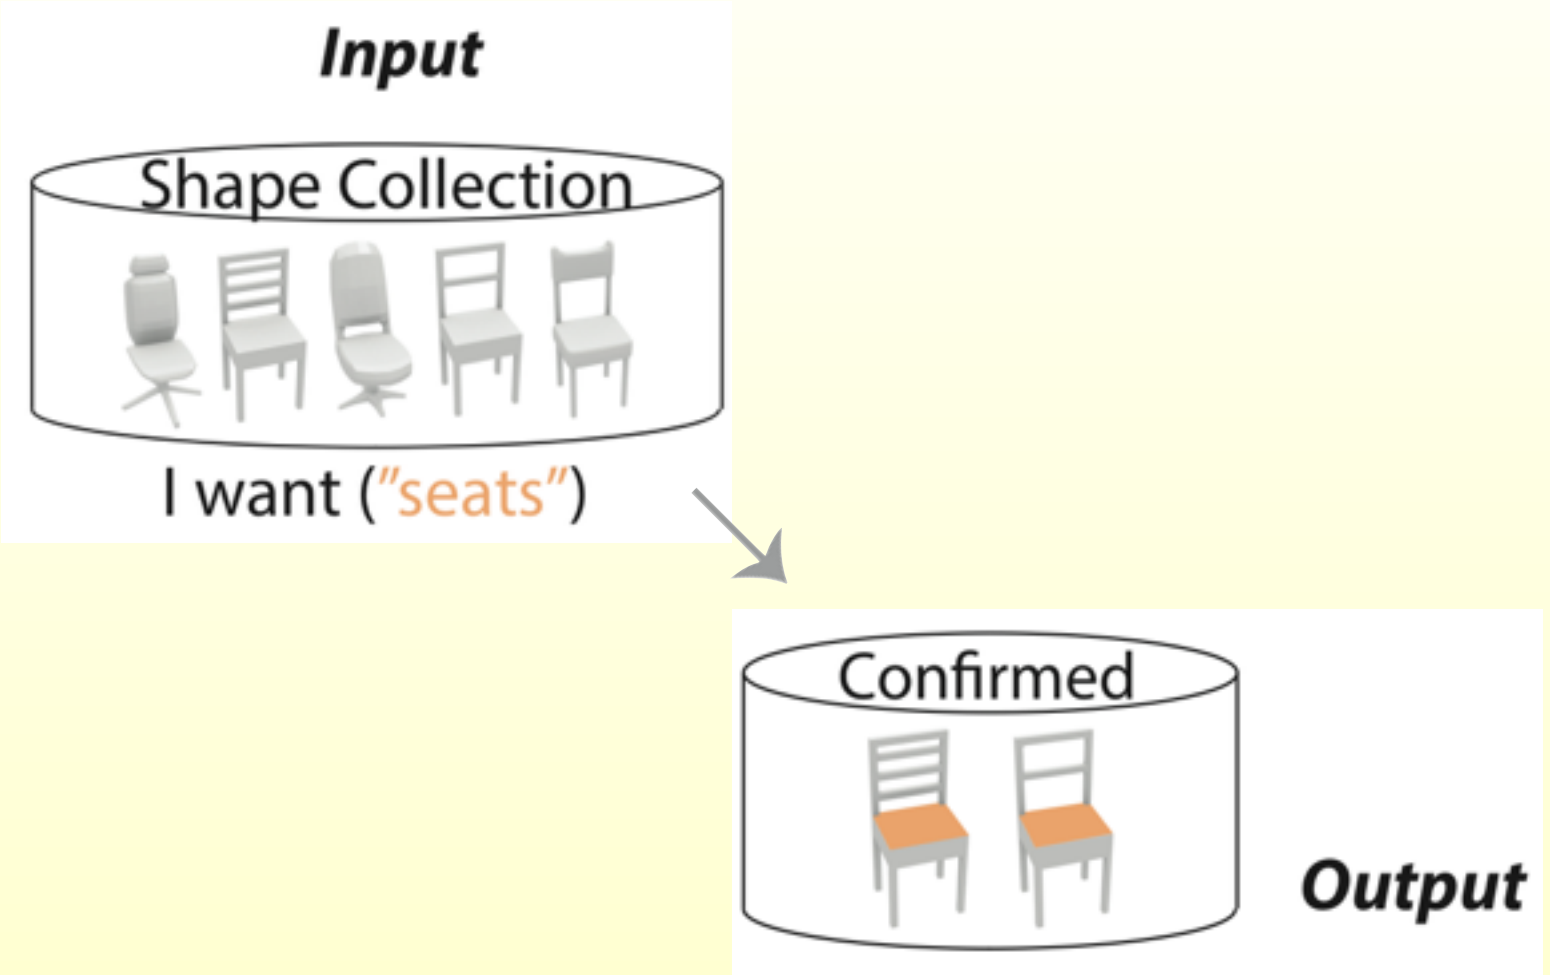

Algorithm Prediction?

Not reliable

- ***Large variation***
- ***Different levels of semantic knowledge are needed***

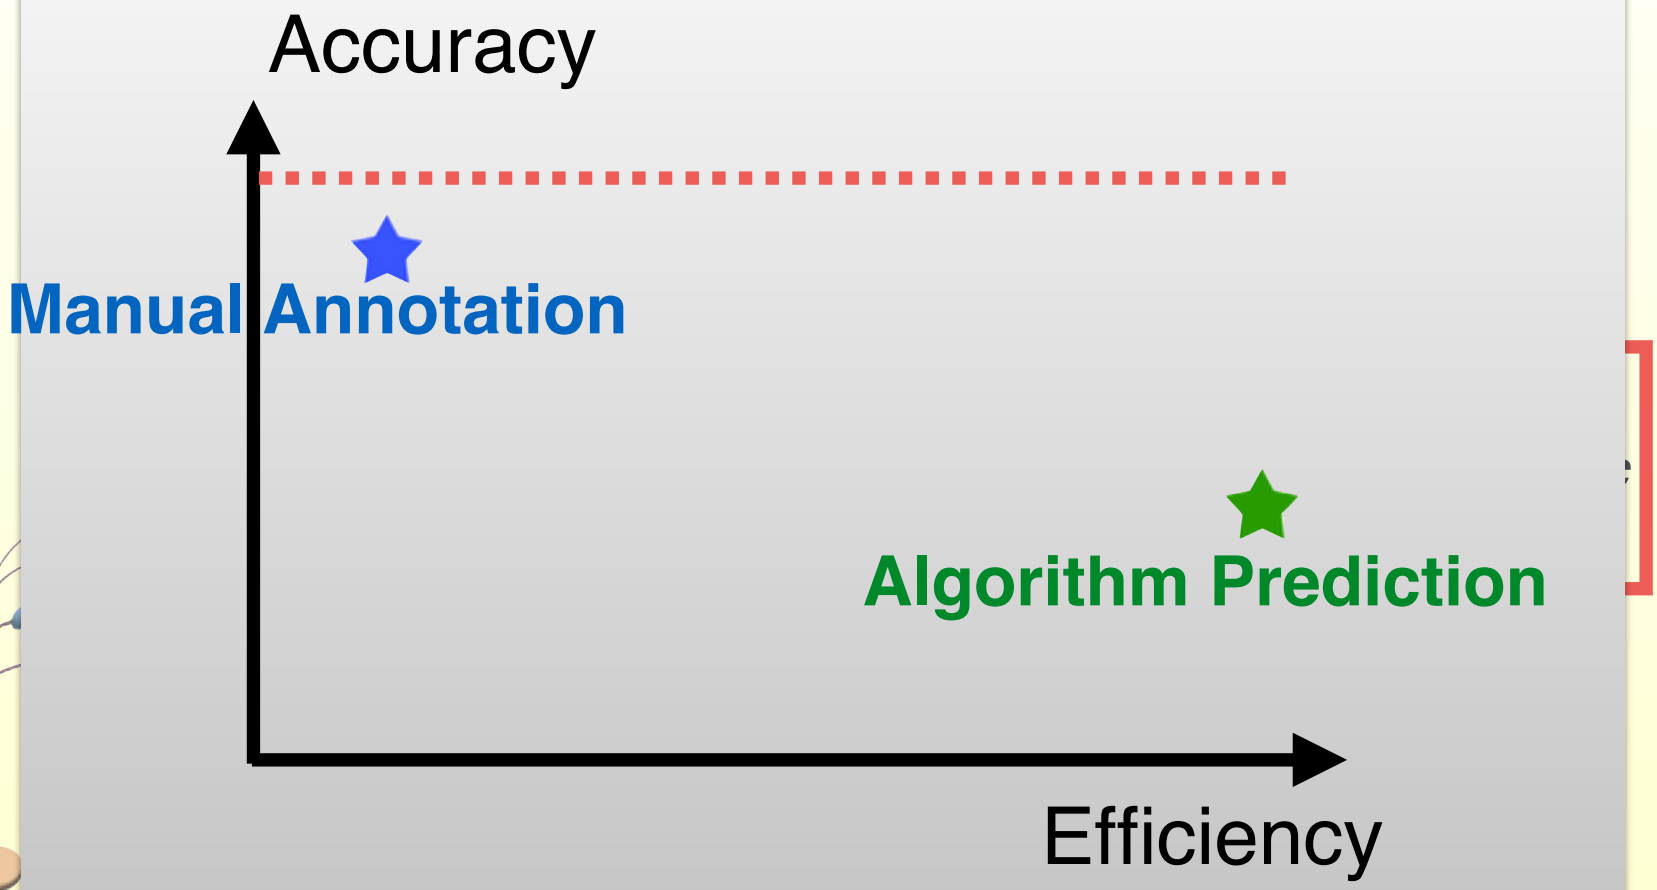
Accuracy
Manual Annotation
Algorithm Prediction
Efficiency

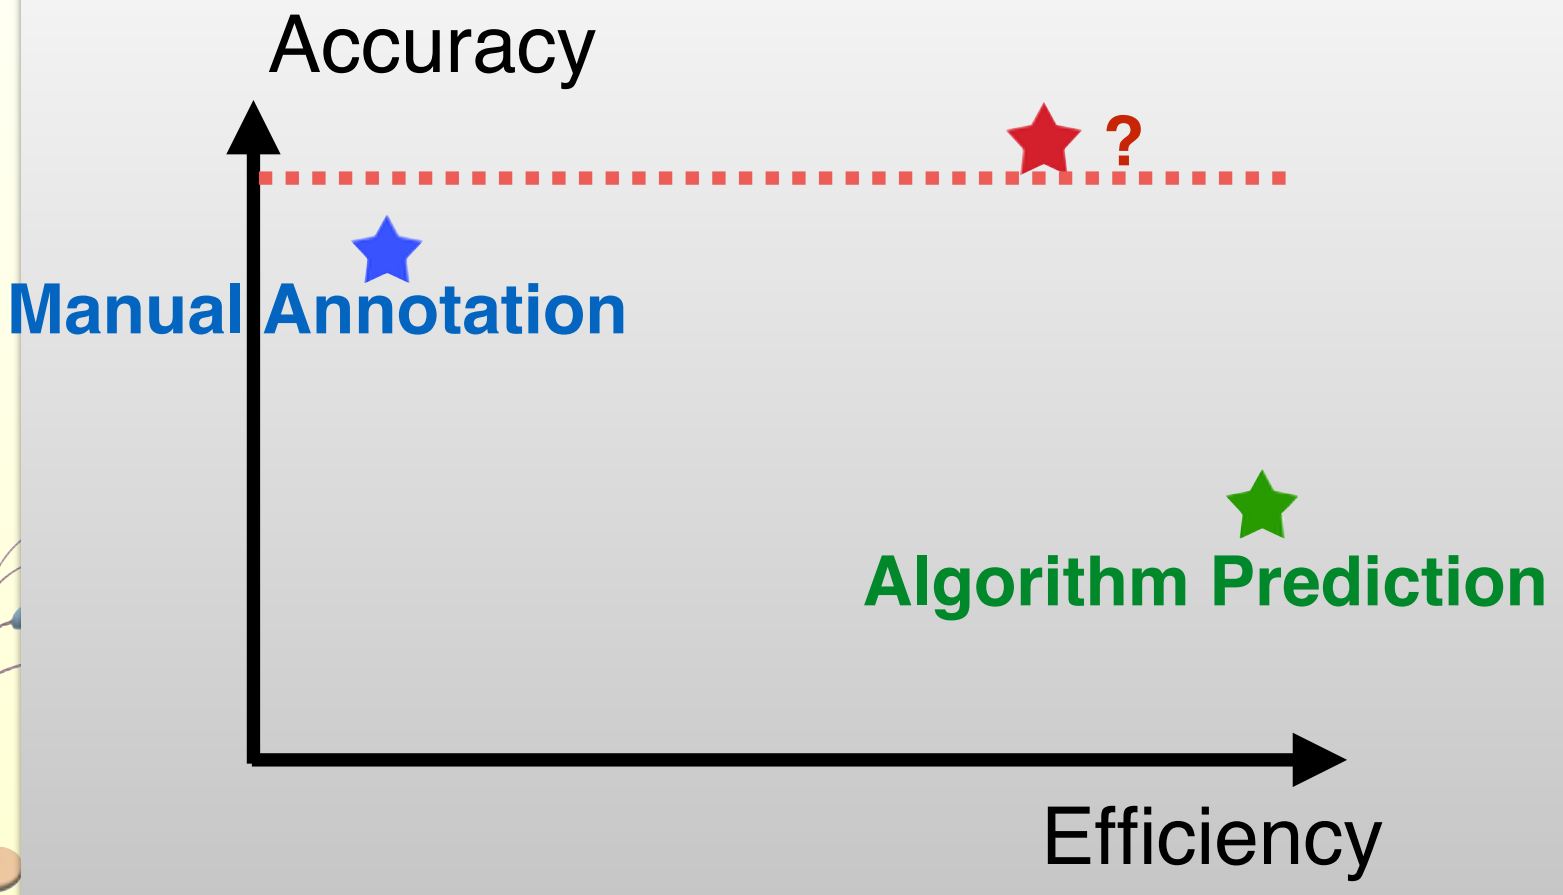
Accuracy
?
Manual Annotation
Algorithm Prediction
Efficiency

A Hybrid Approach

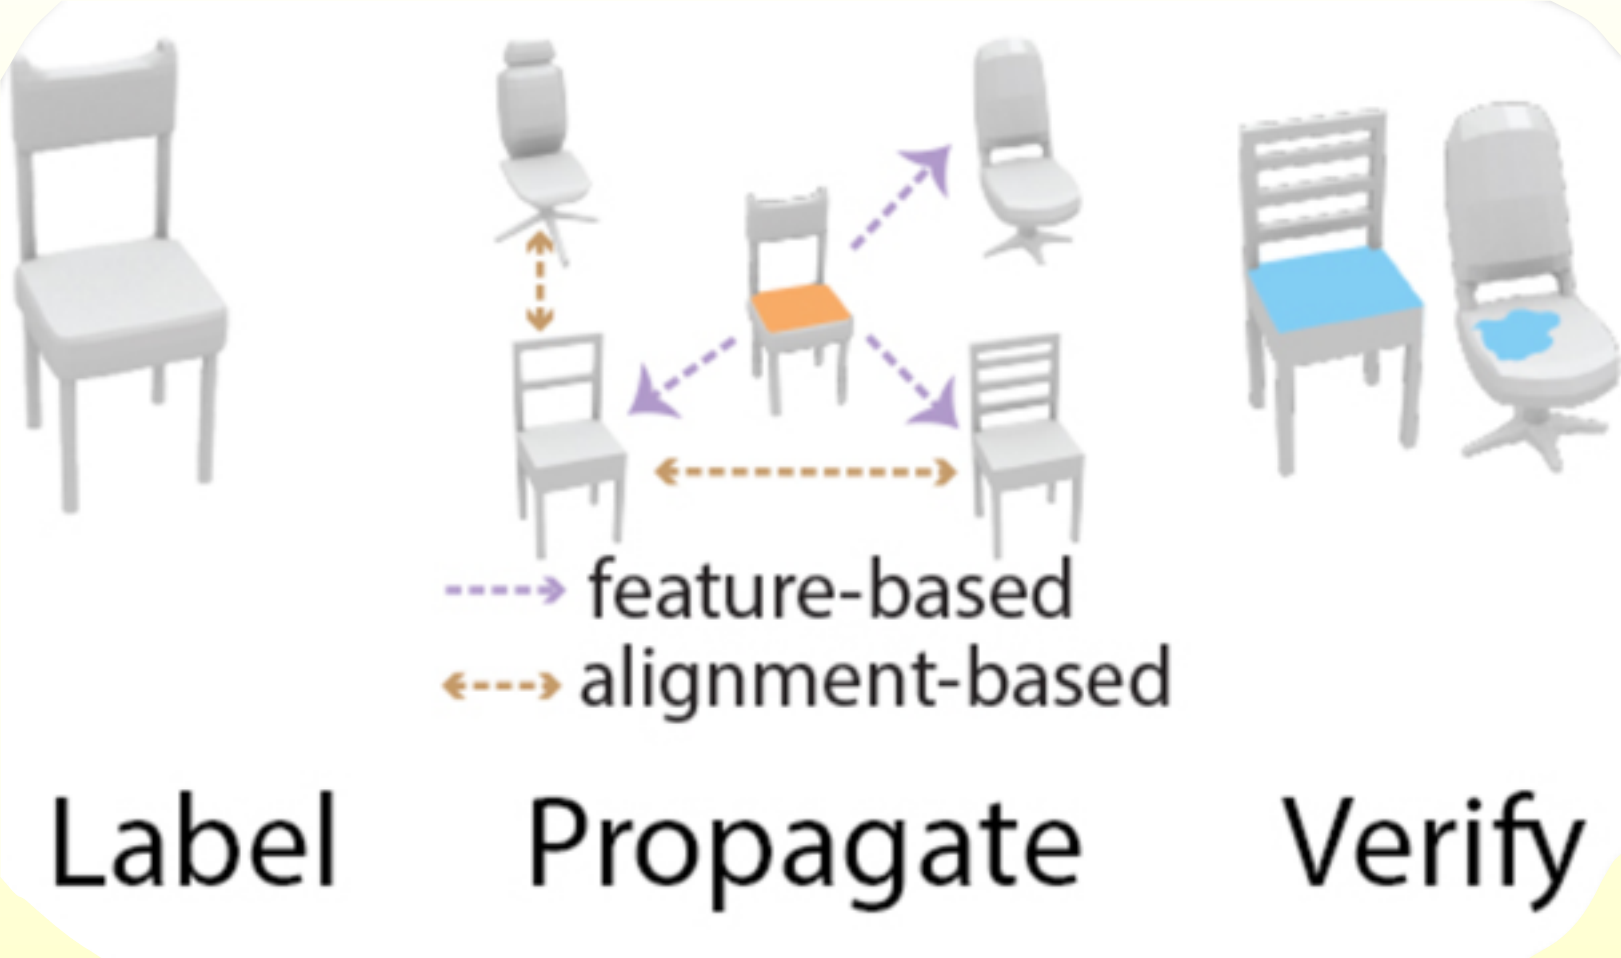

A Hybrid Approach

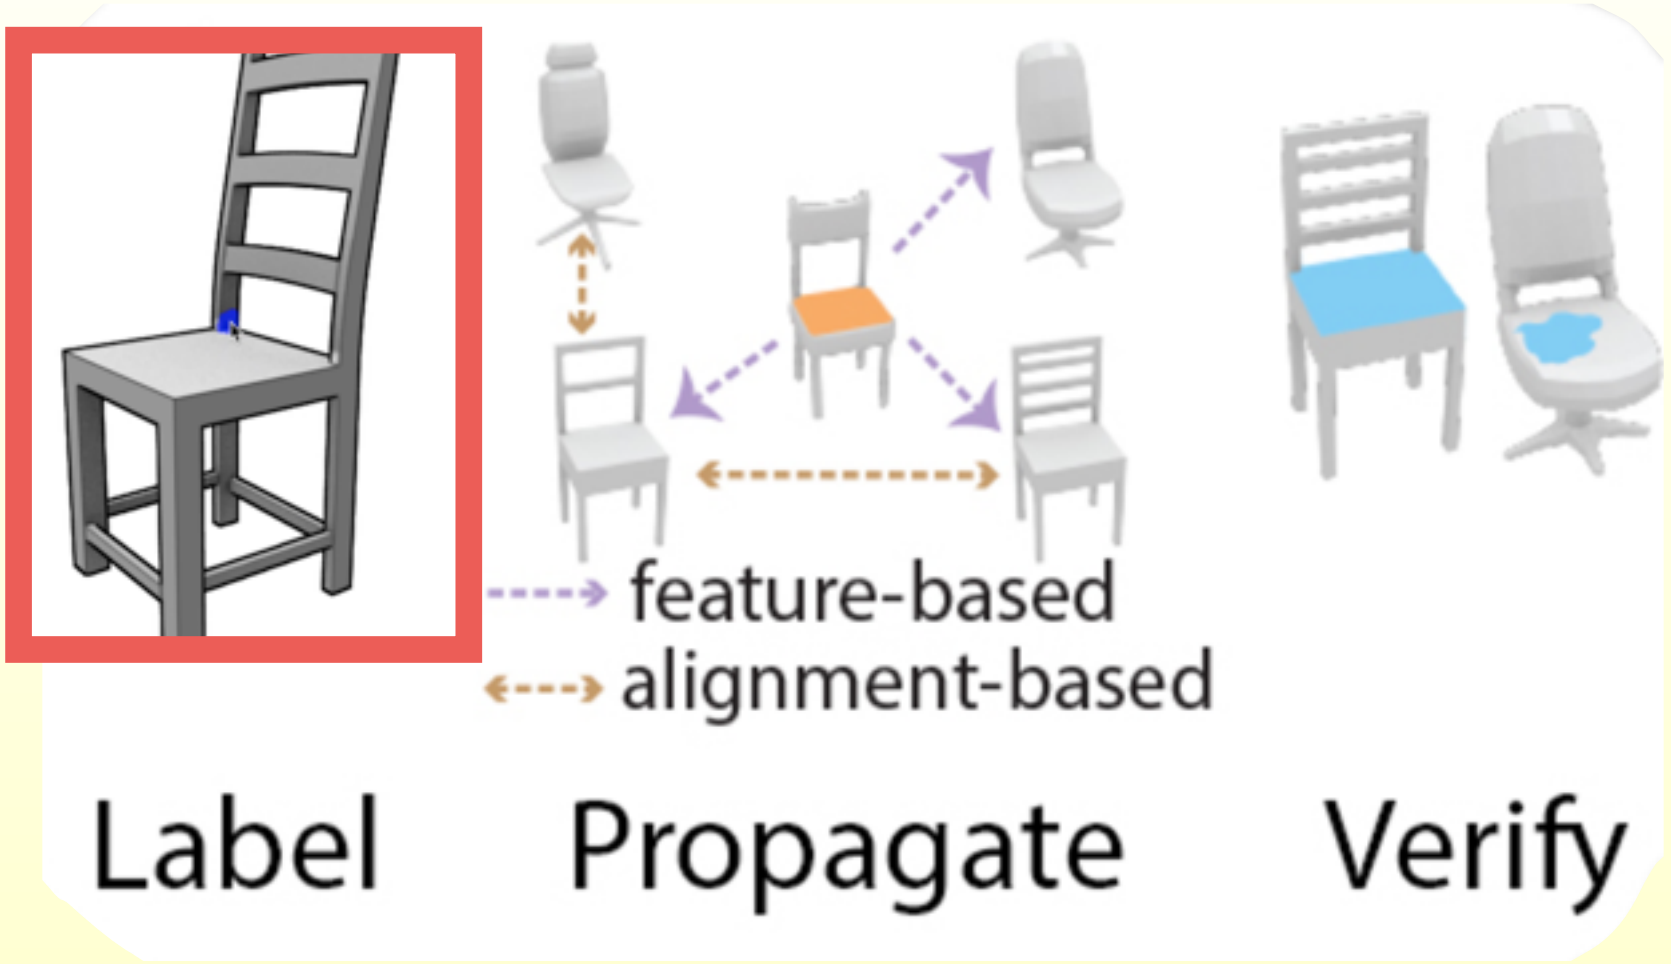

A Hybrid Approach

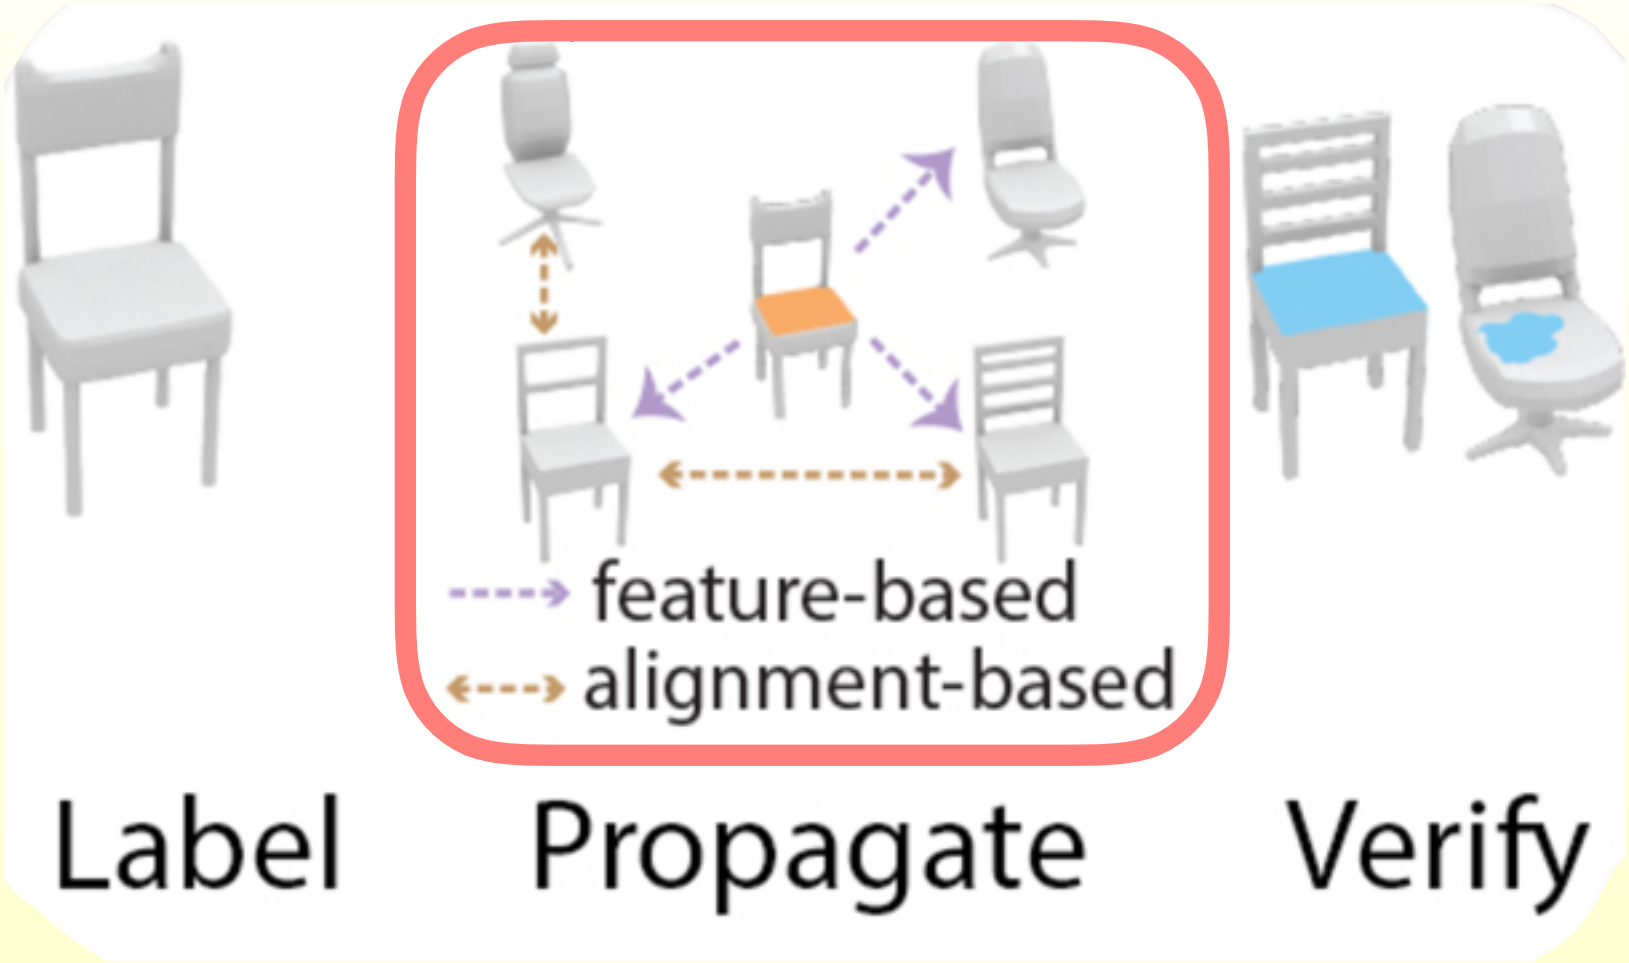

A Hybrid Approach

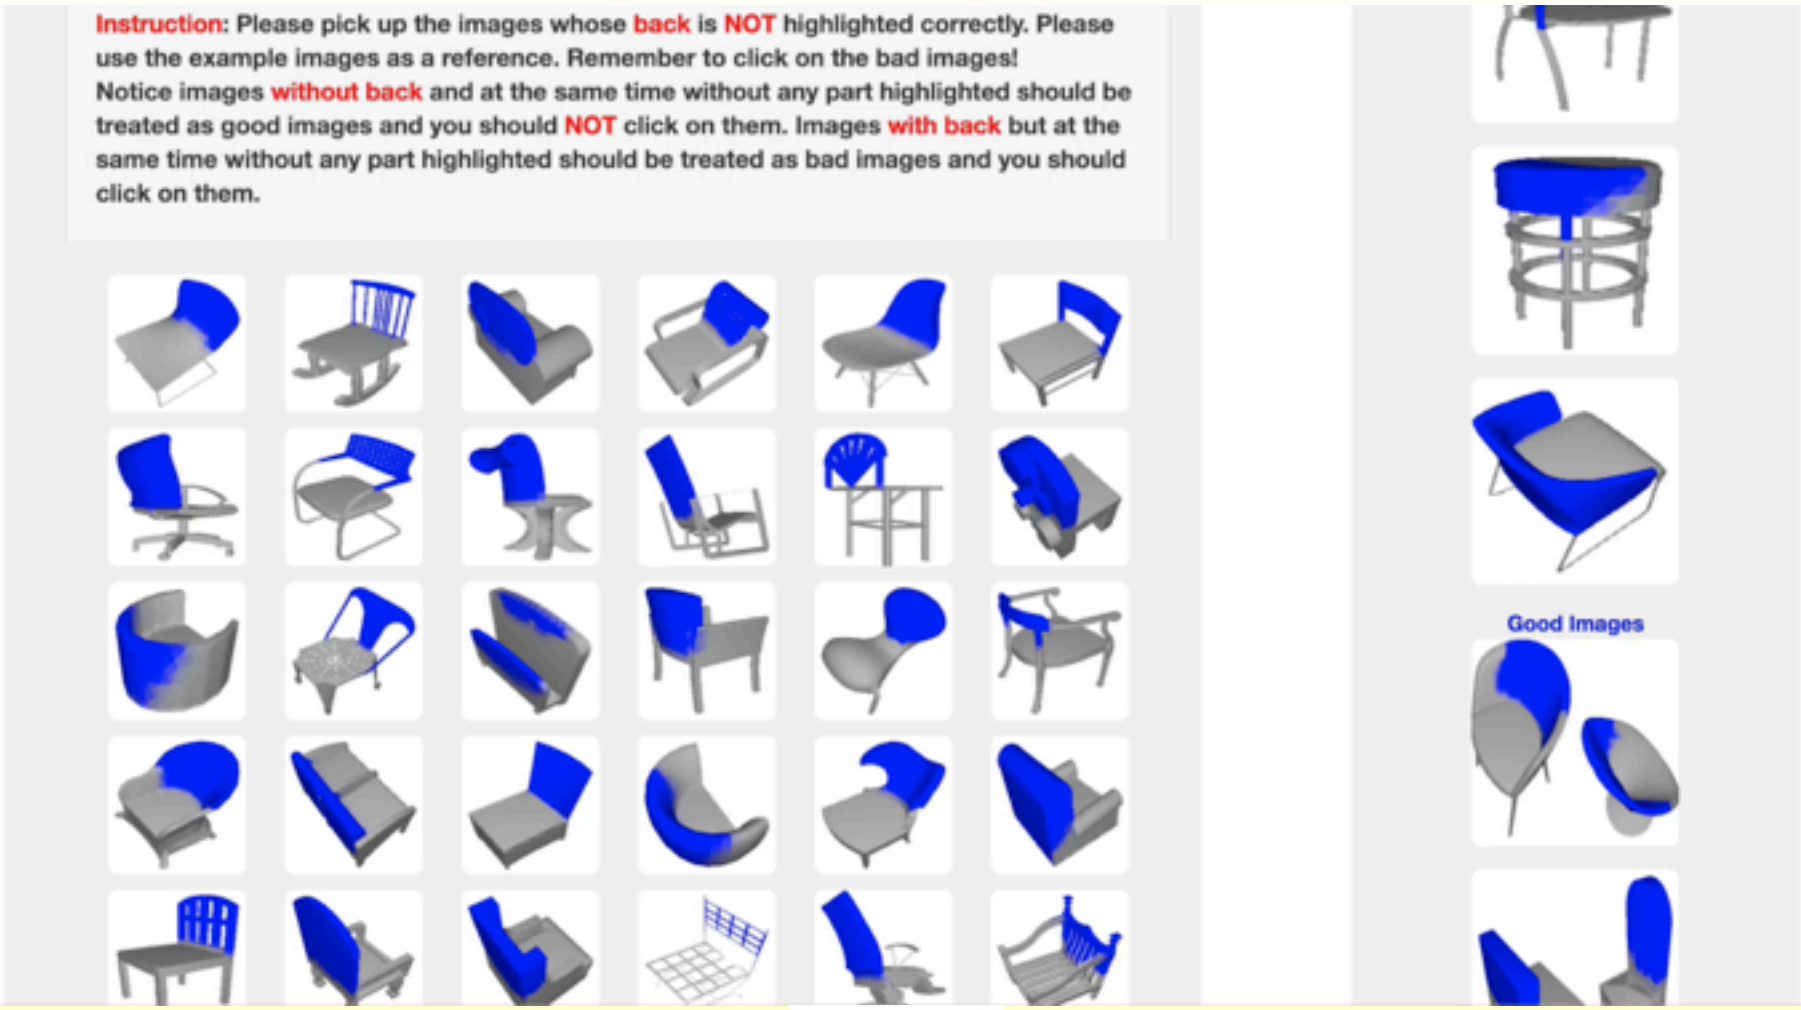

A Hybrid Approach

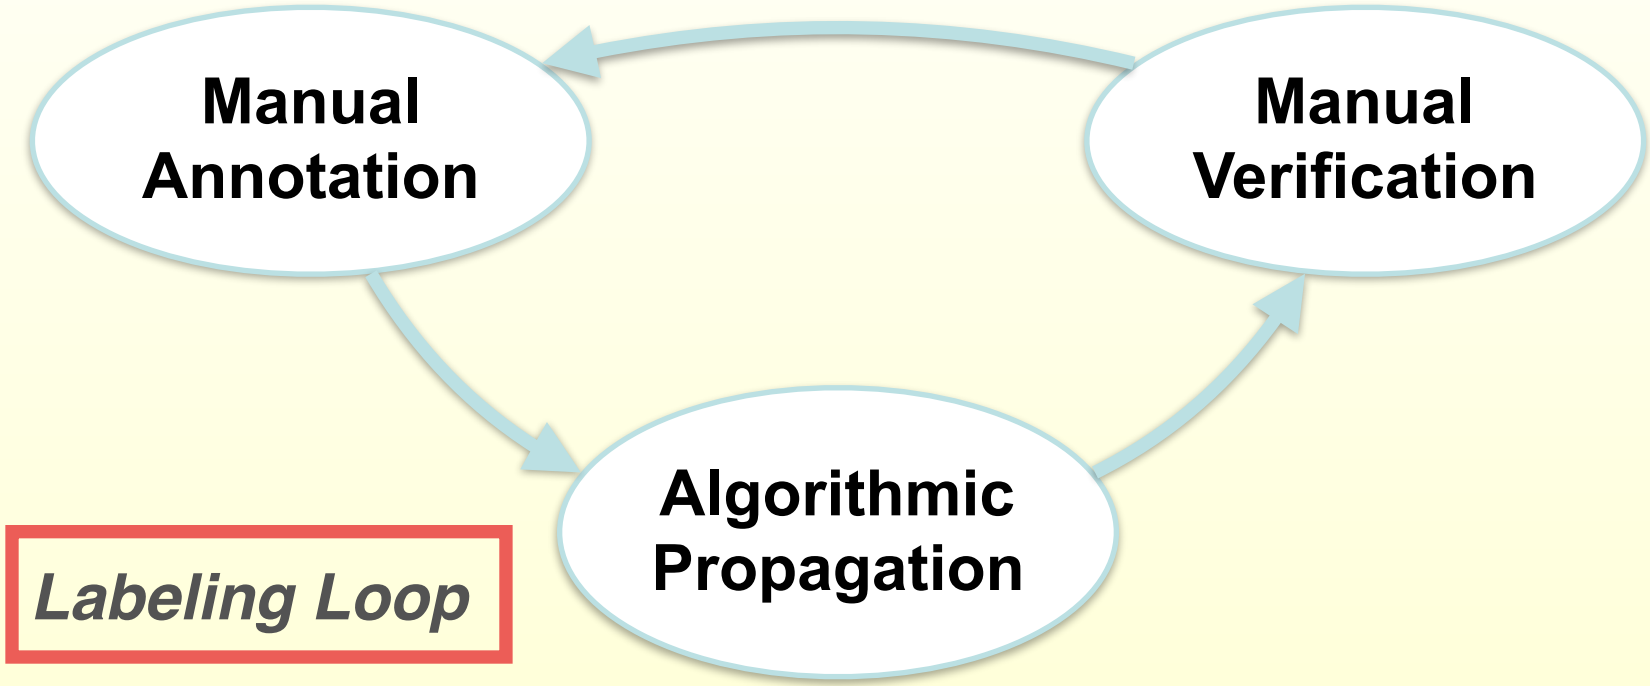

A Hybrid Approach

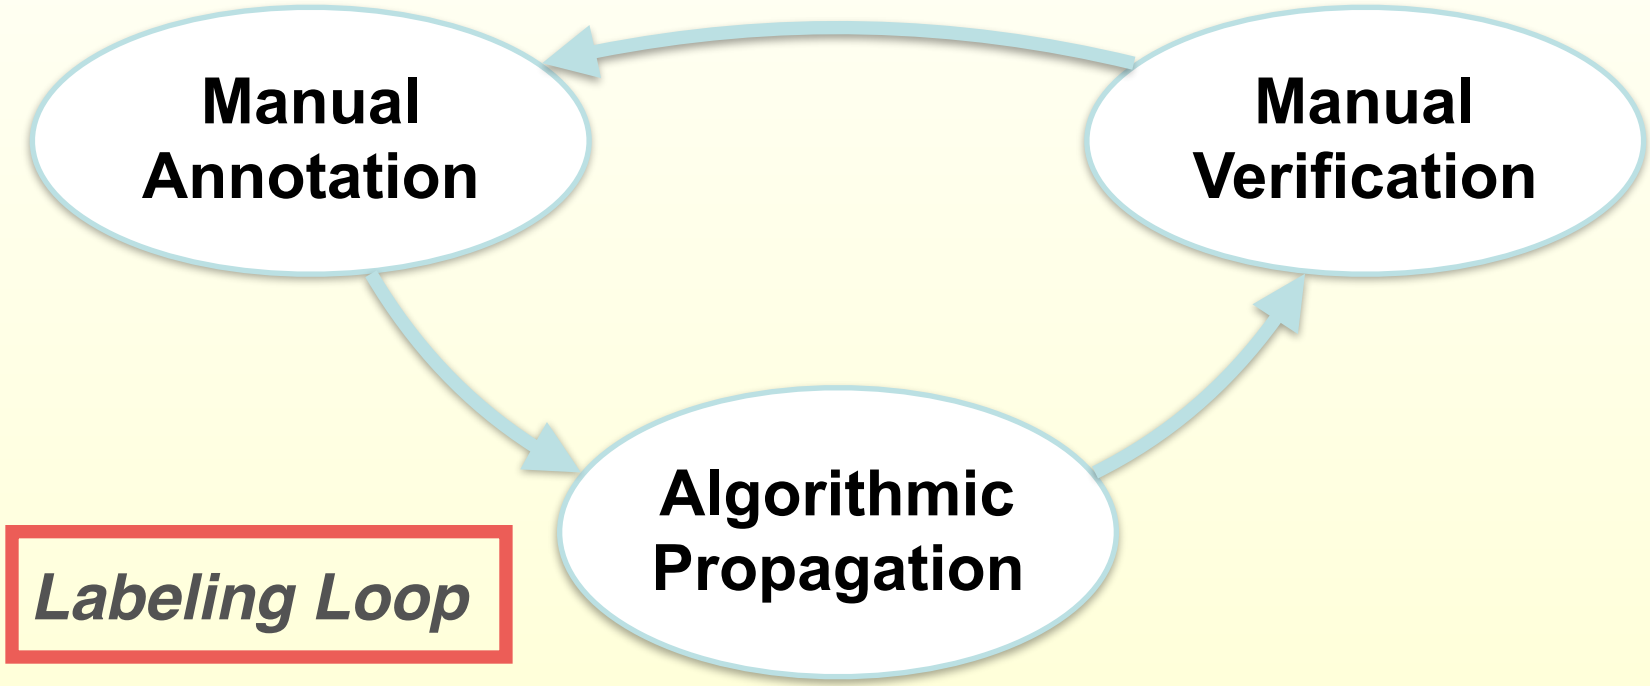

- Categorization, rigid alignment, part and key points annotations, weight and size annotations

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontal networks
  - Annotation transportation in ShapeNet

“Vertical” networks

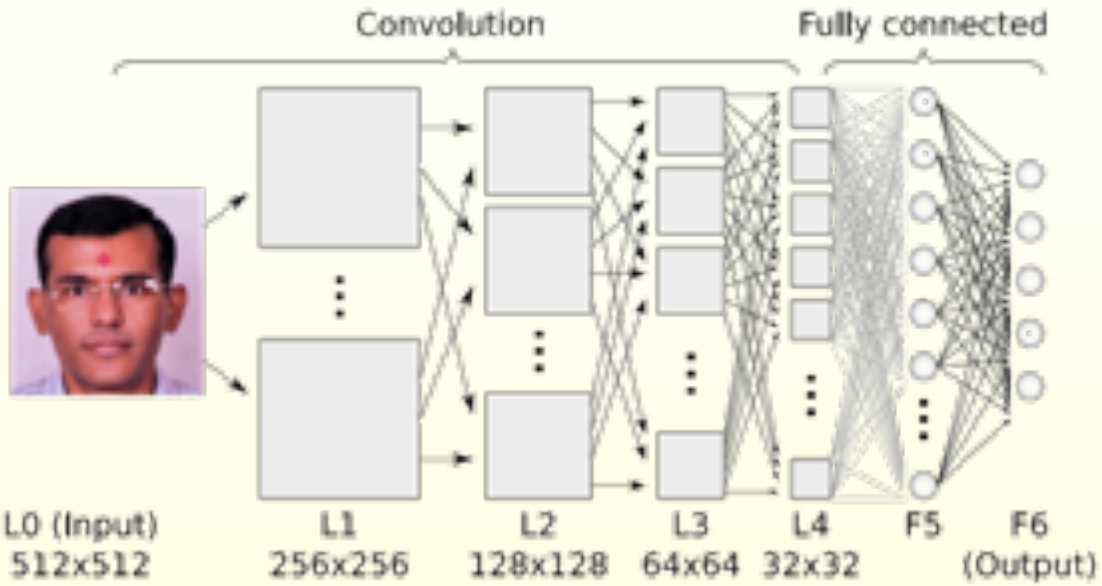

[Makwana, 2016]

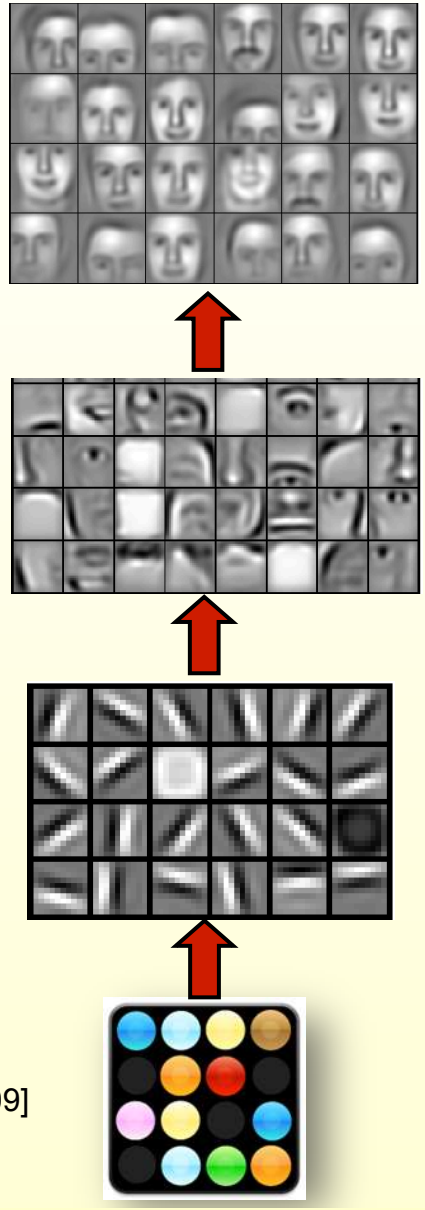

[Lee et al., 2009]

Data-driven feature learning at ascending abstraction layers

Success Made Possible By ….

[He et al., 2015]

Novel deep architectures

Plenty of annotated data

Lots of computing power

Success Made Possible By ….

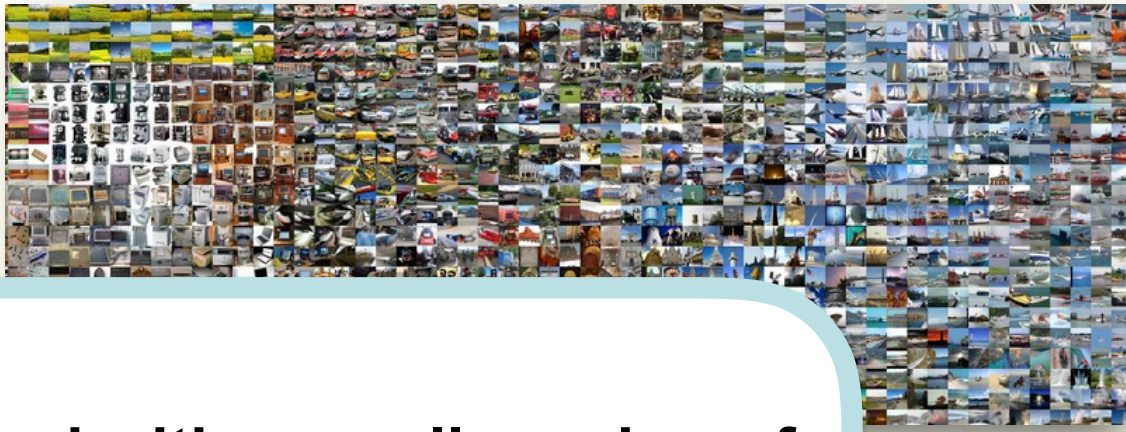

tated data

Cannot be trained with a small number of annotations to start the labeling loop.

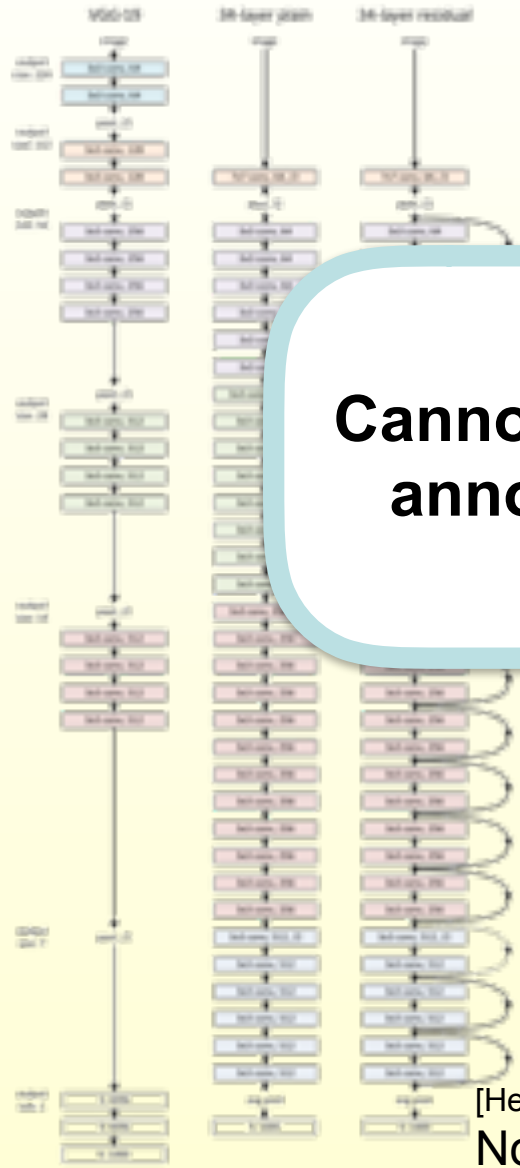

[He et al., 2015]

Novel deep architectures

Lots of computing pov

What's missing?

- The structural relationship among different objects are not well exploited
- Correspondences allow information propagation

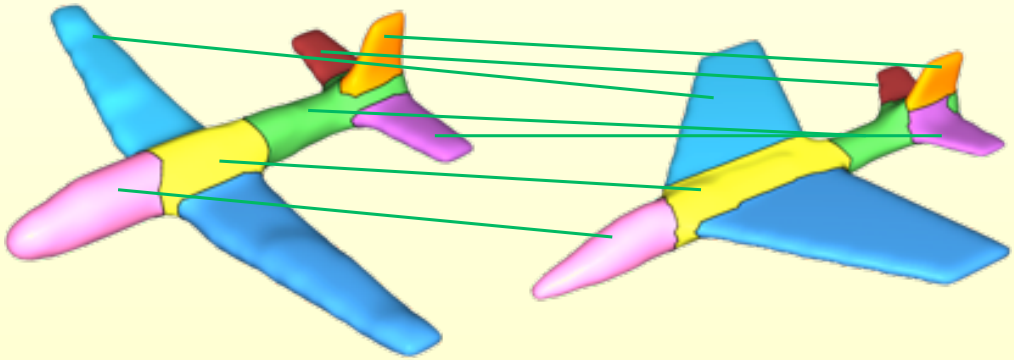

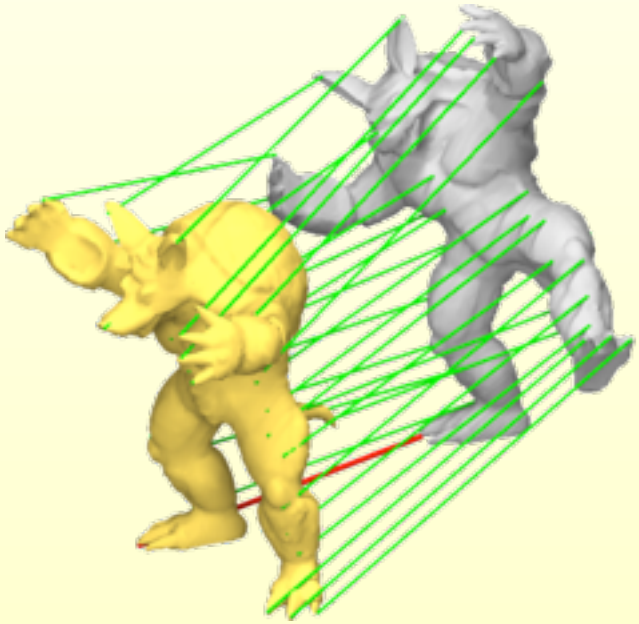

Why are Relationships Important?

- Consistency resolves ambiguities
- Structures emerge from consistency

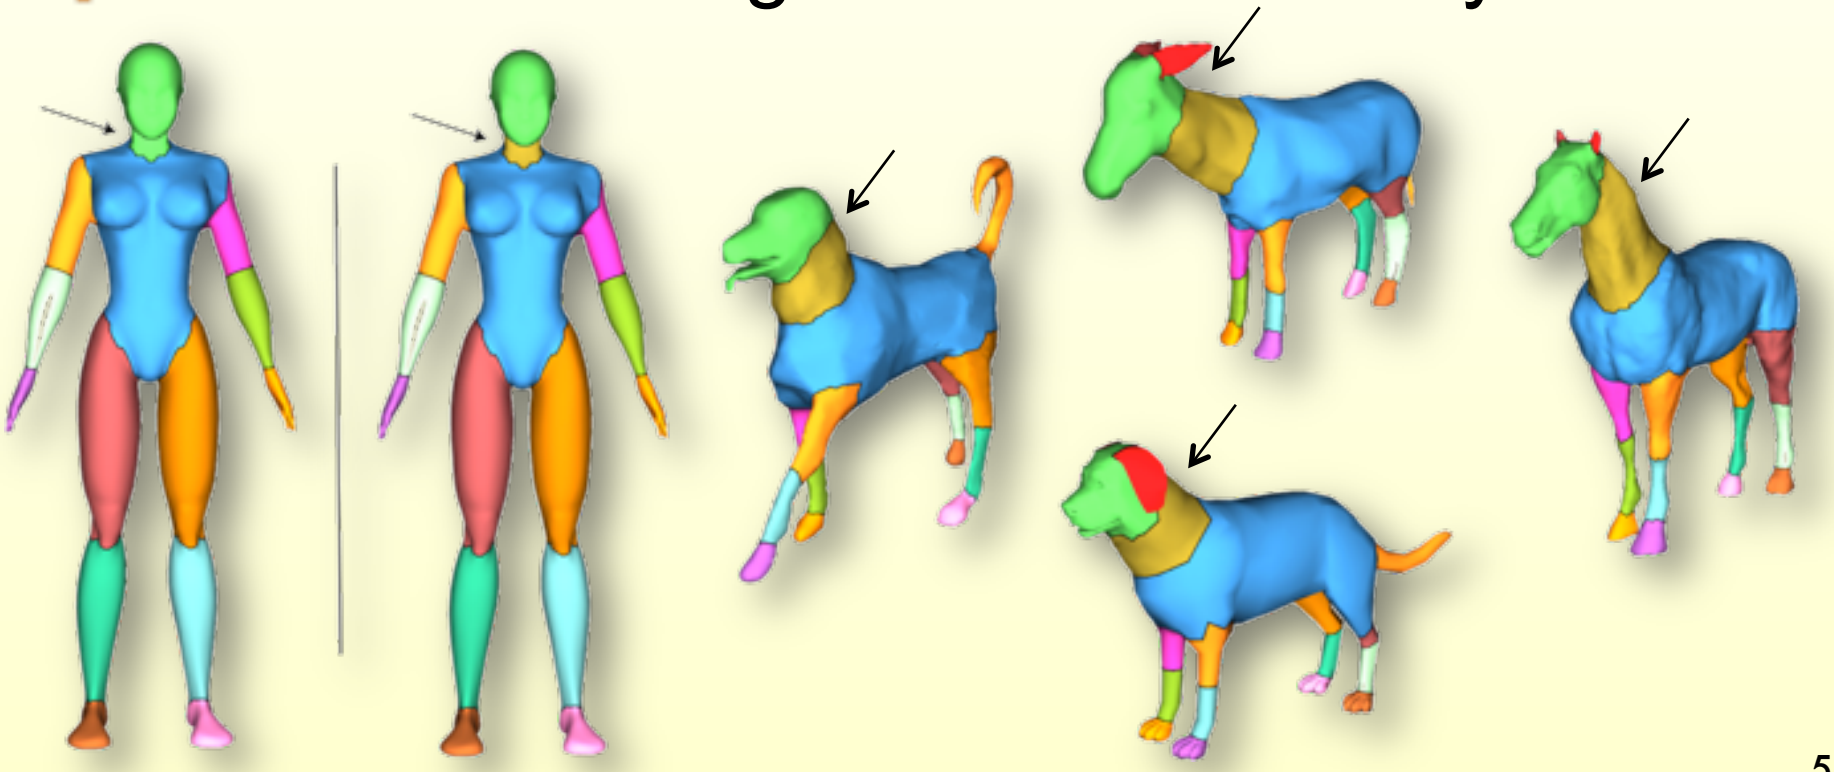

Horizontal Networks

- Explicitly model the structural relationships among data

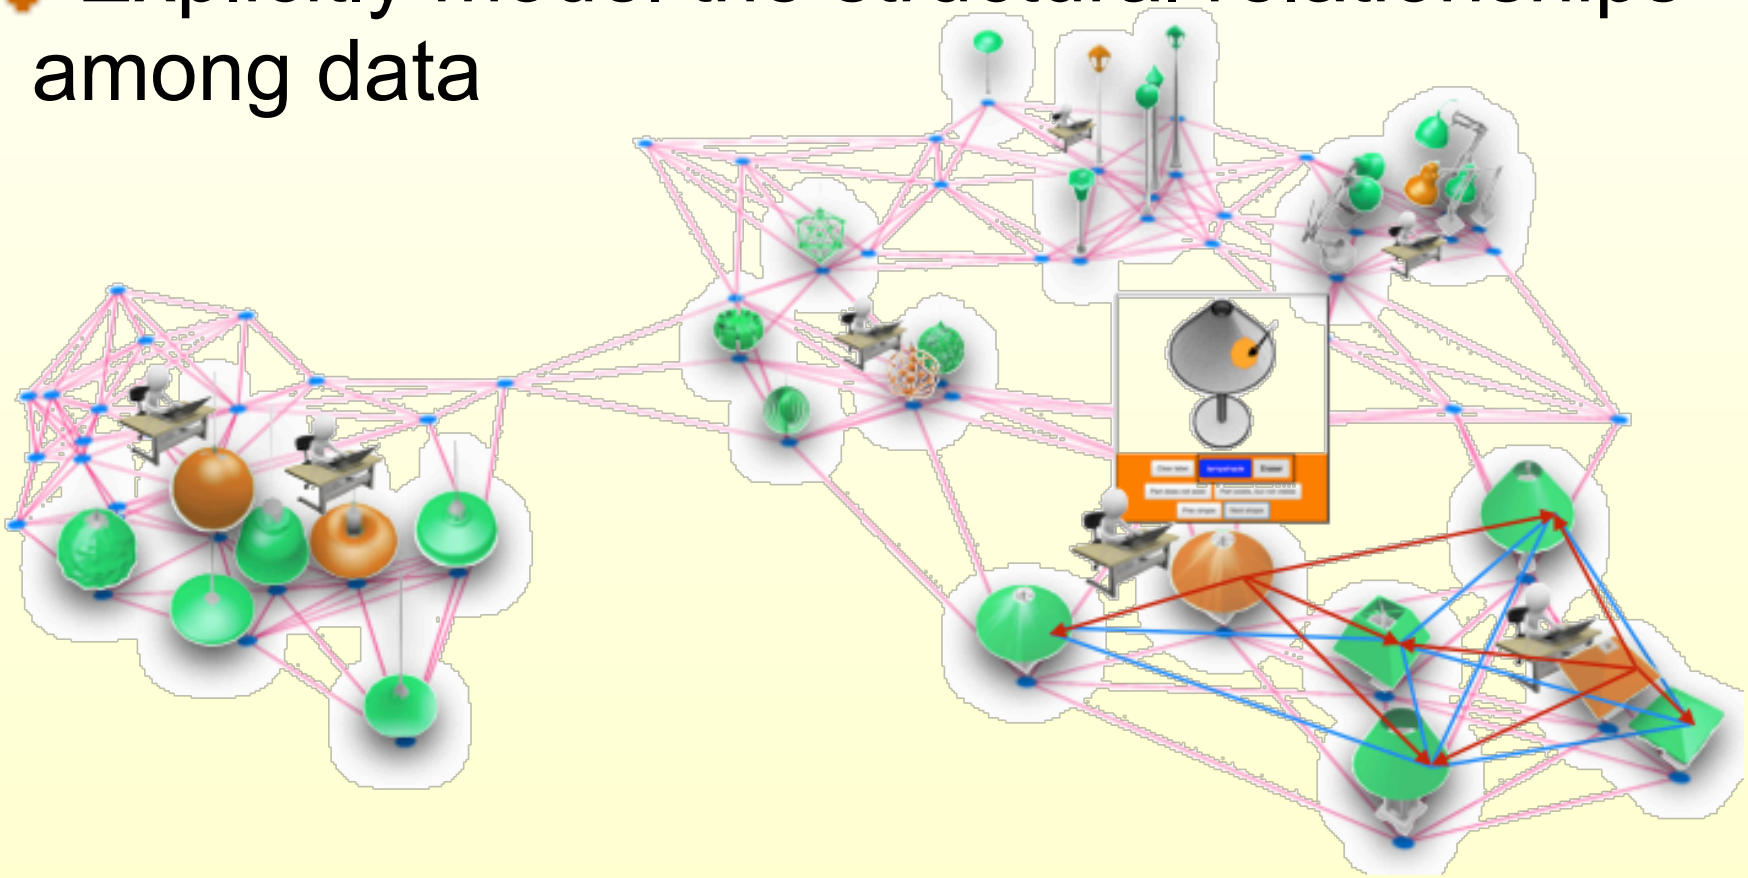

Applications of Horizontal Networks

- A scalable active framework for part annotation in ShapeNet
- Learning hierarchical shape segmentation and labeling in a weakly supervised manner from online repositories

Applications of Horizontal Networks

- A scalable active framework for part annotation in ShapeNet
- Learning hierarchical shape segmentation and labeling in a weakly supervised manner from online repositories

A Hybrid Part Annotation Approach

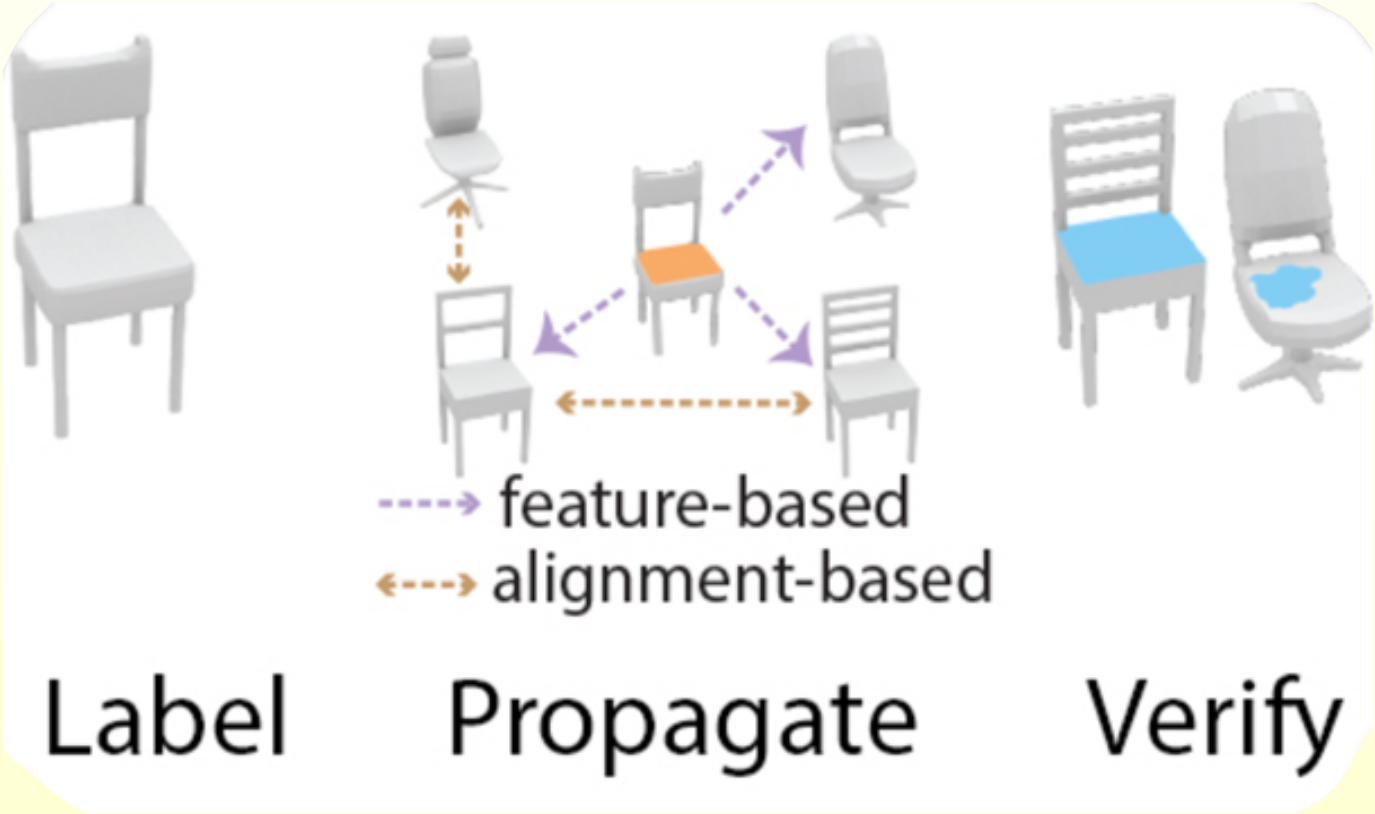

A Hybrid Part Annotation Approach

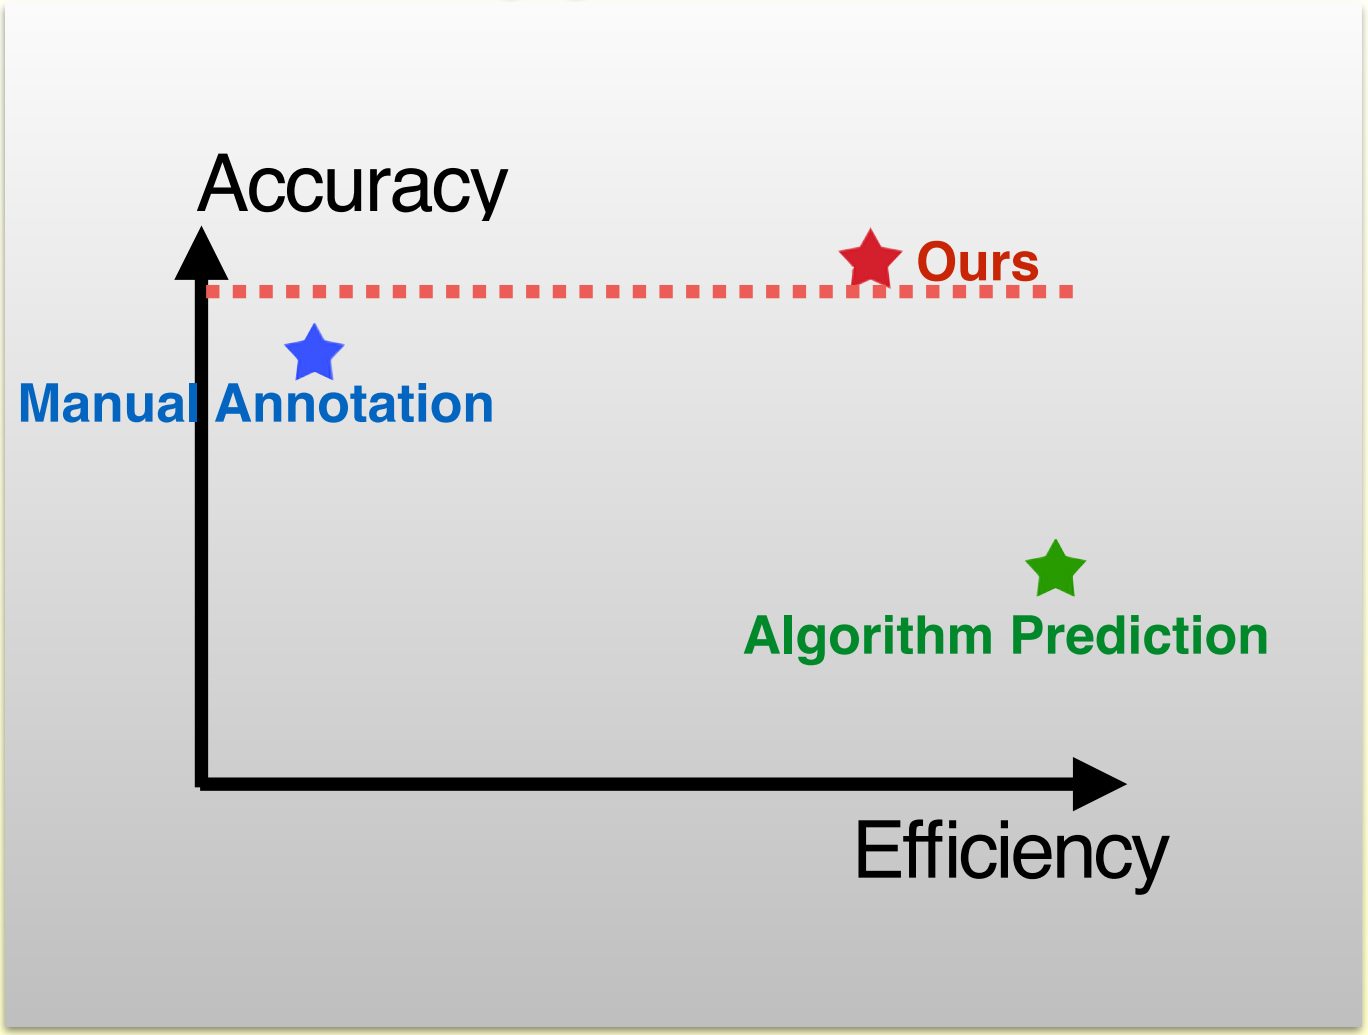

Utility Function

$$\text{Utility Function:} \quad \frac{\#(\textit{ACCURATE} \text{ labeling})}{\text{worker time}}$$

Utility function

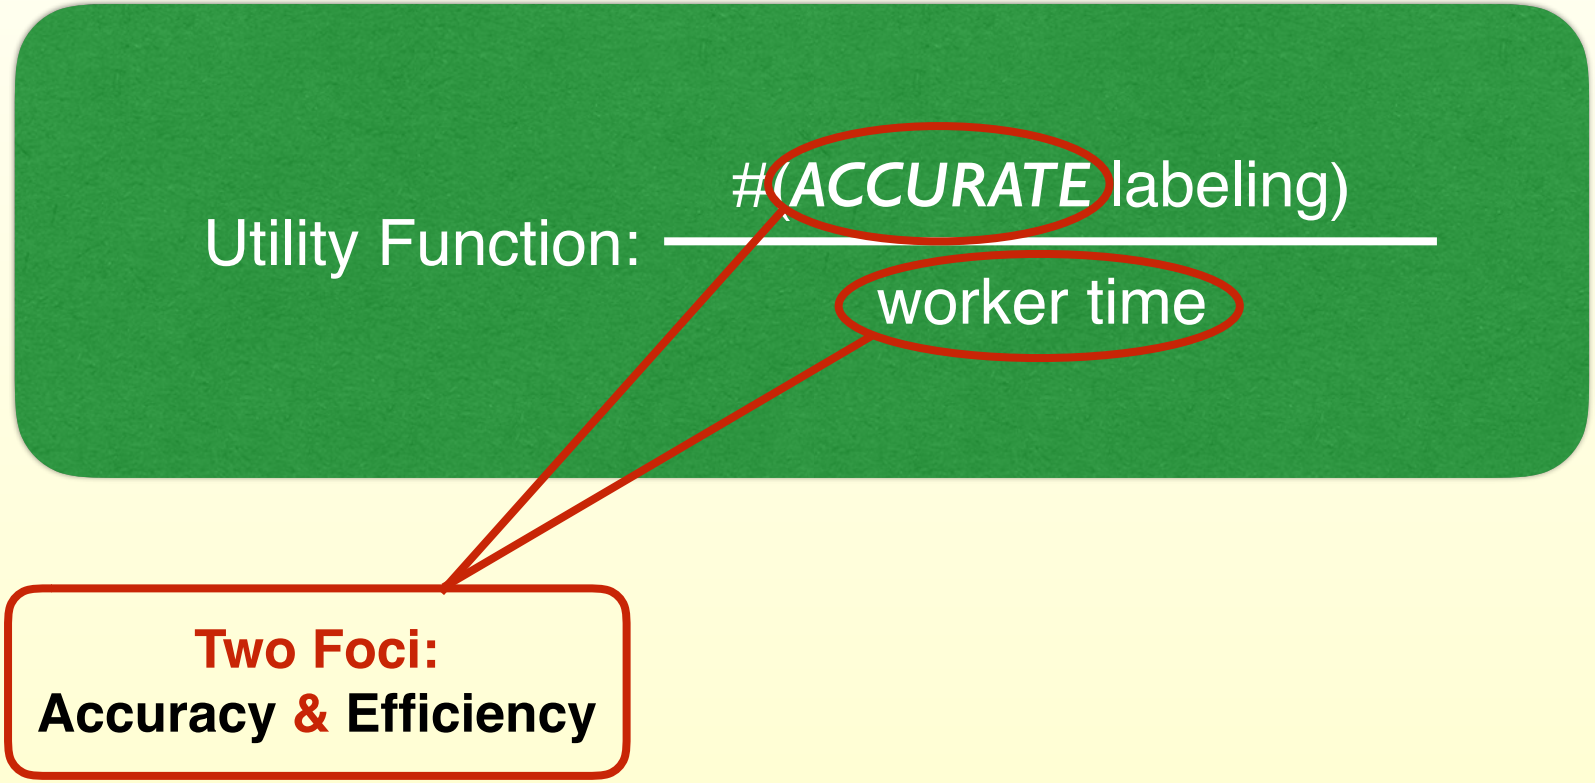

Pipeline

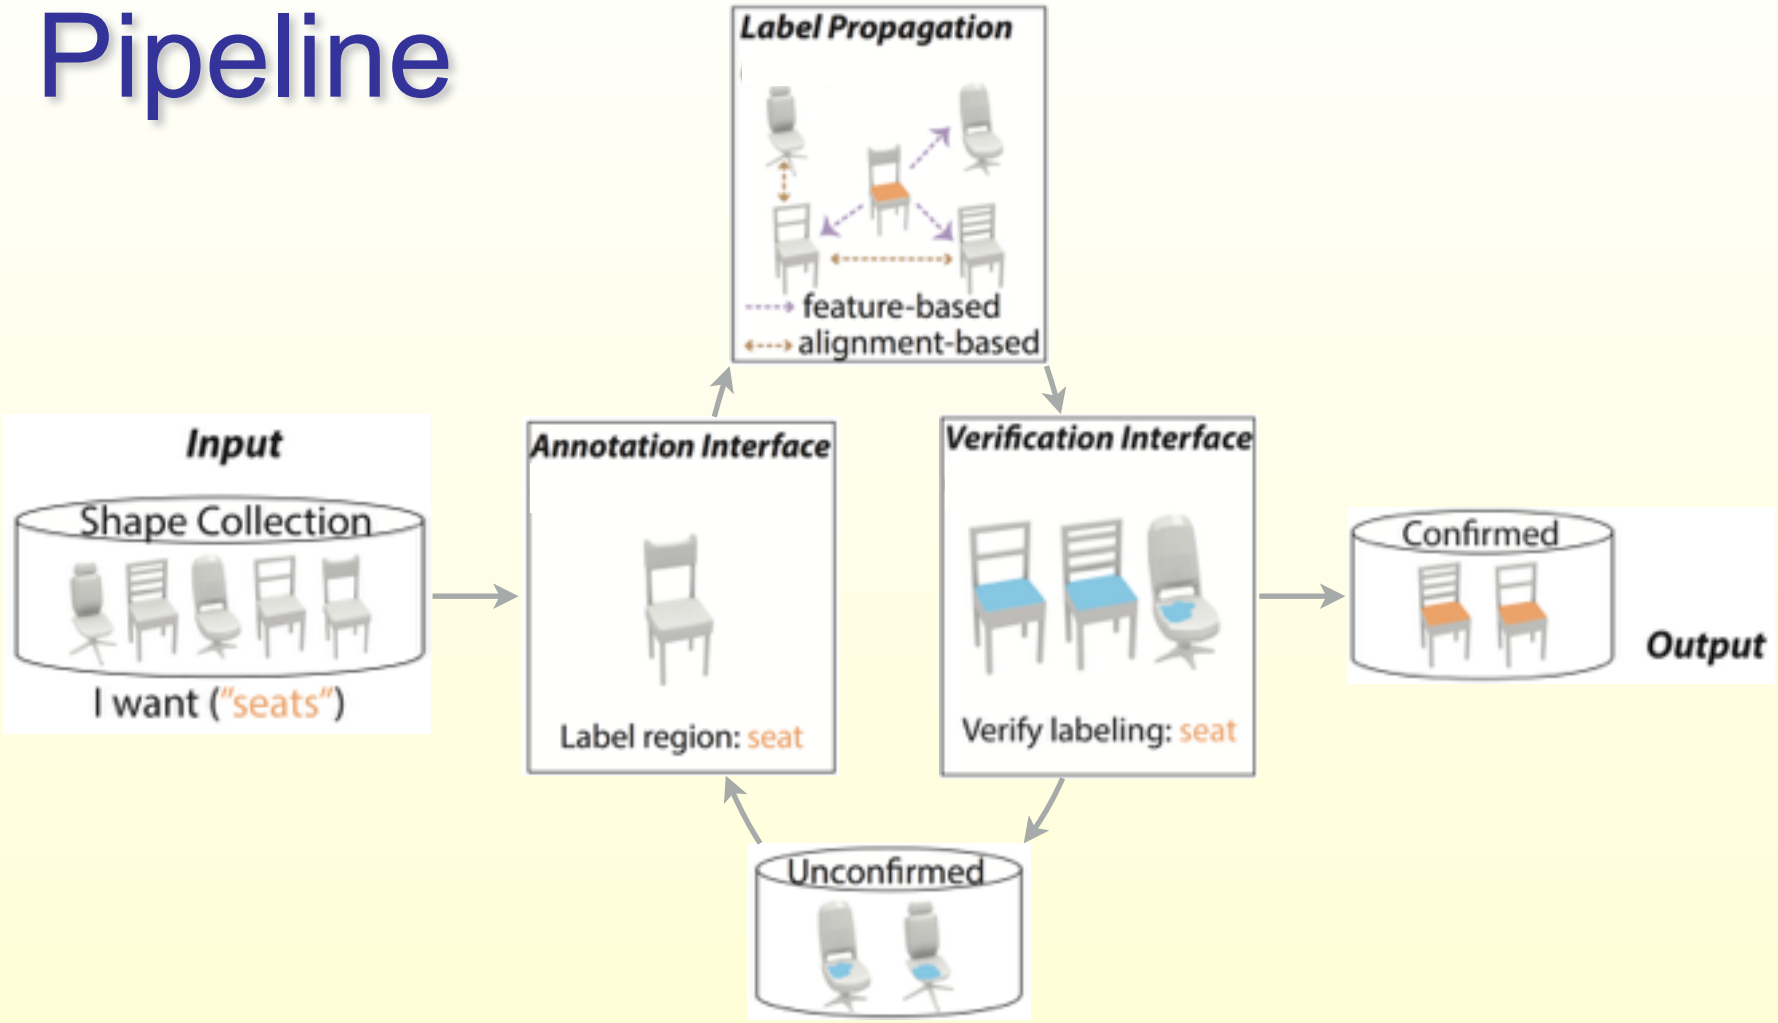
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

Pipeline

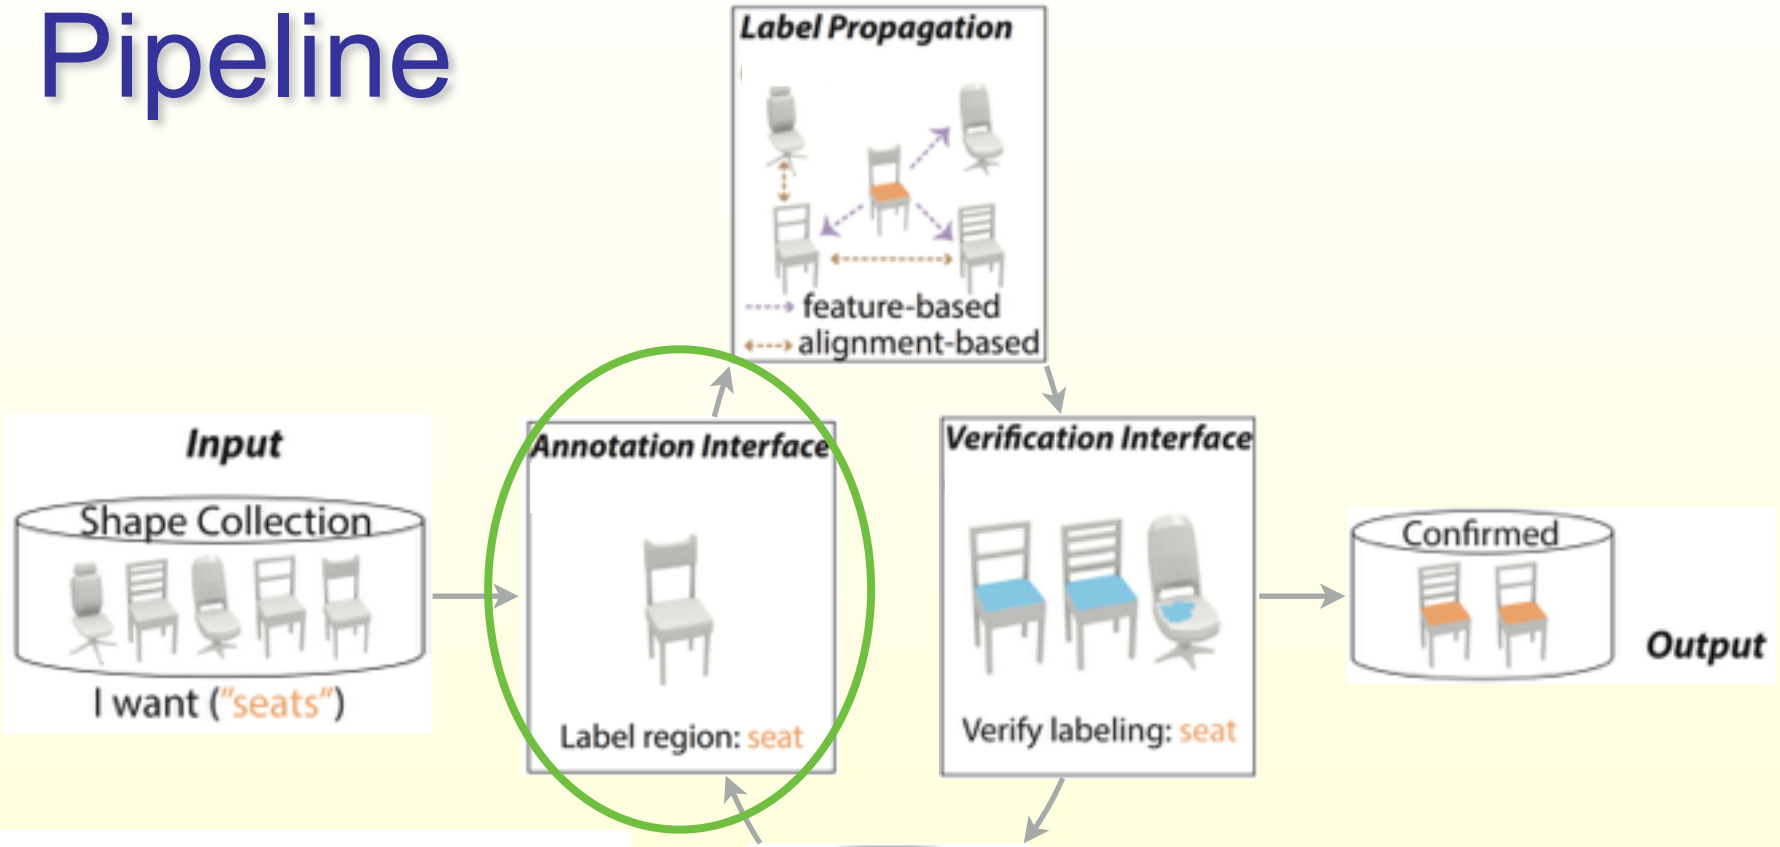
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

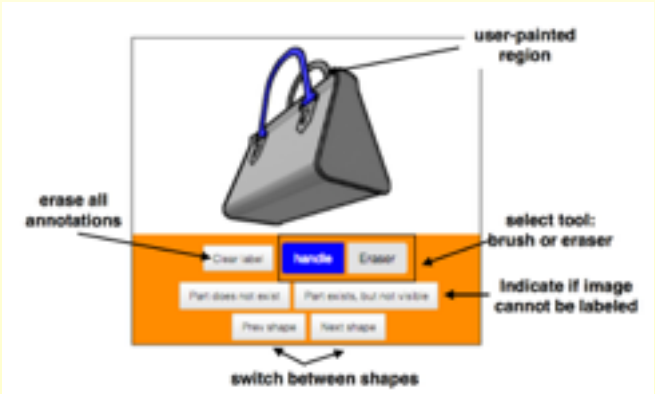
user-painted region
erase all annotations
select tool: brush or eraser
Indicate if image cannot be labeled
switch between shapes

Pipeline

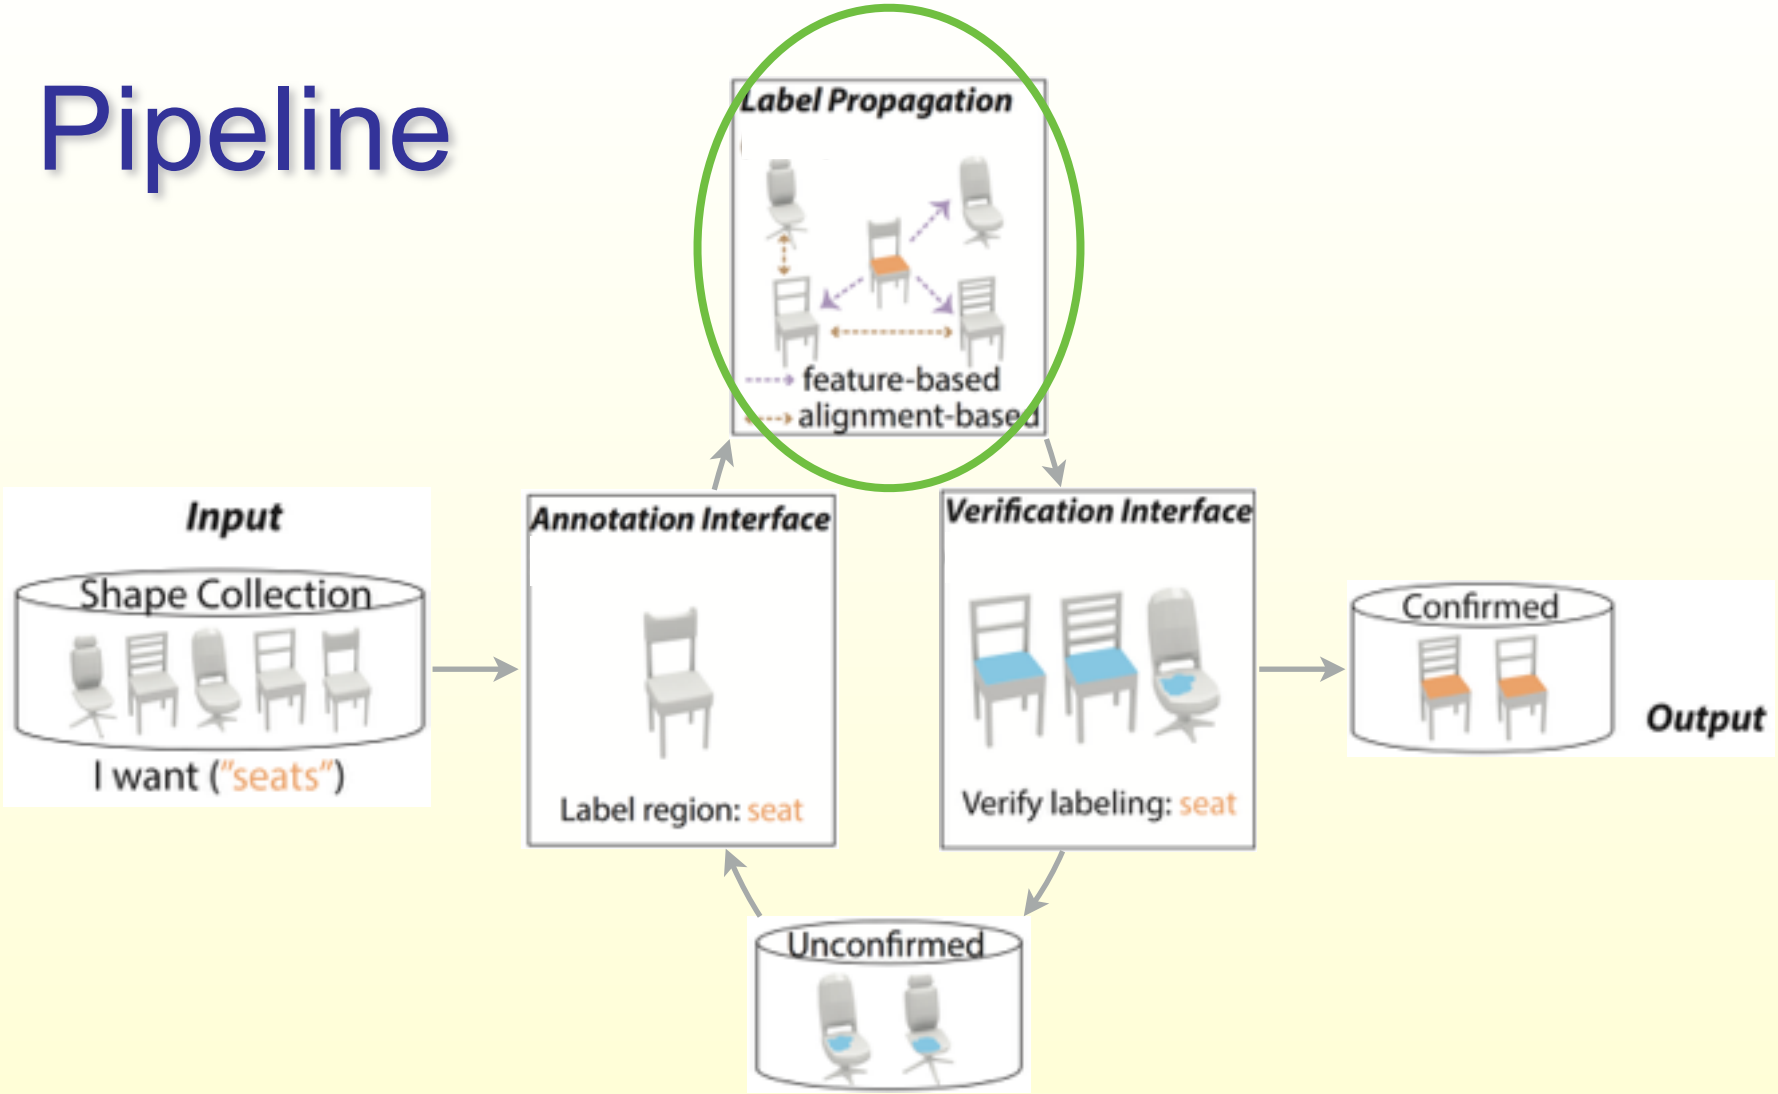
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

Label Propagation

$$E(f) = \sum_{p \in S} \psi_f(p) + \sum_{p_1, p_2 \in S \times S} \lambda_1 \psi_c(p_1, p_2) + \sum_{p_1, p_2 \in S} \lambda_2 \psi_s(p_1, p_2)$$

data term

correspondence constraint

smoothness constraint

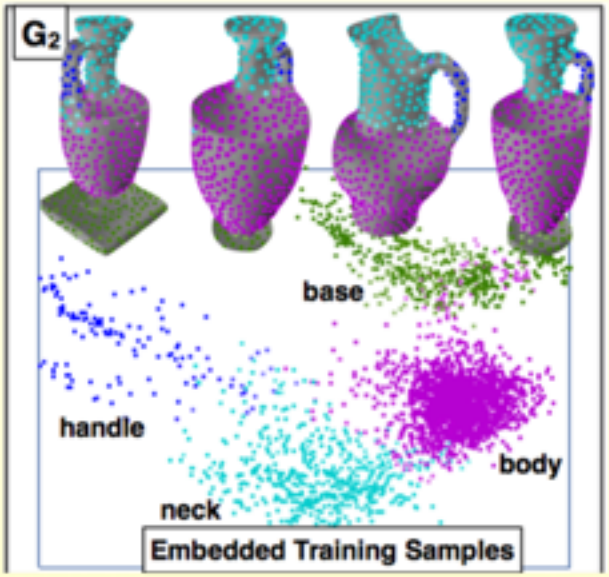

Huang et al. 2013

Pipeline

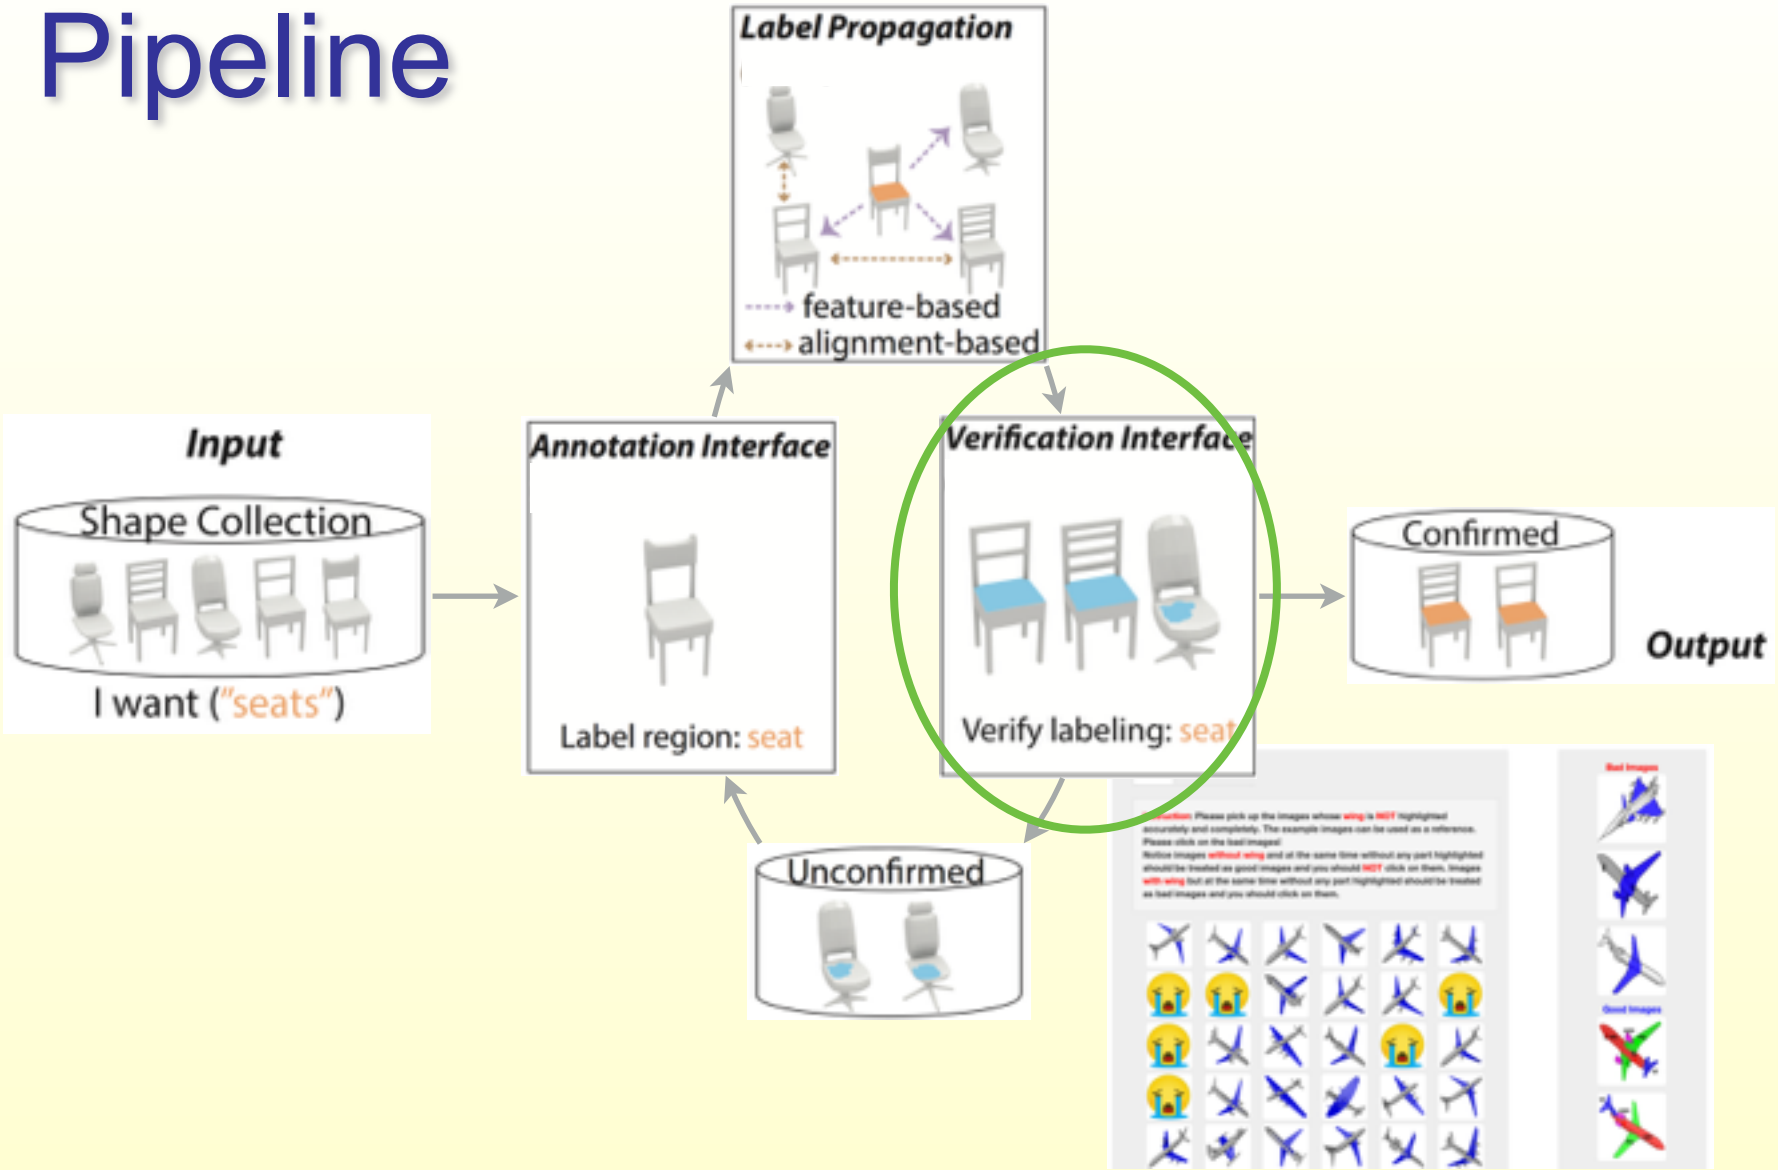
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

Pipeline

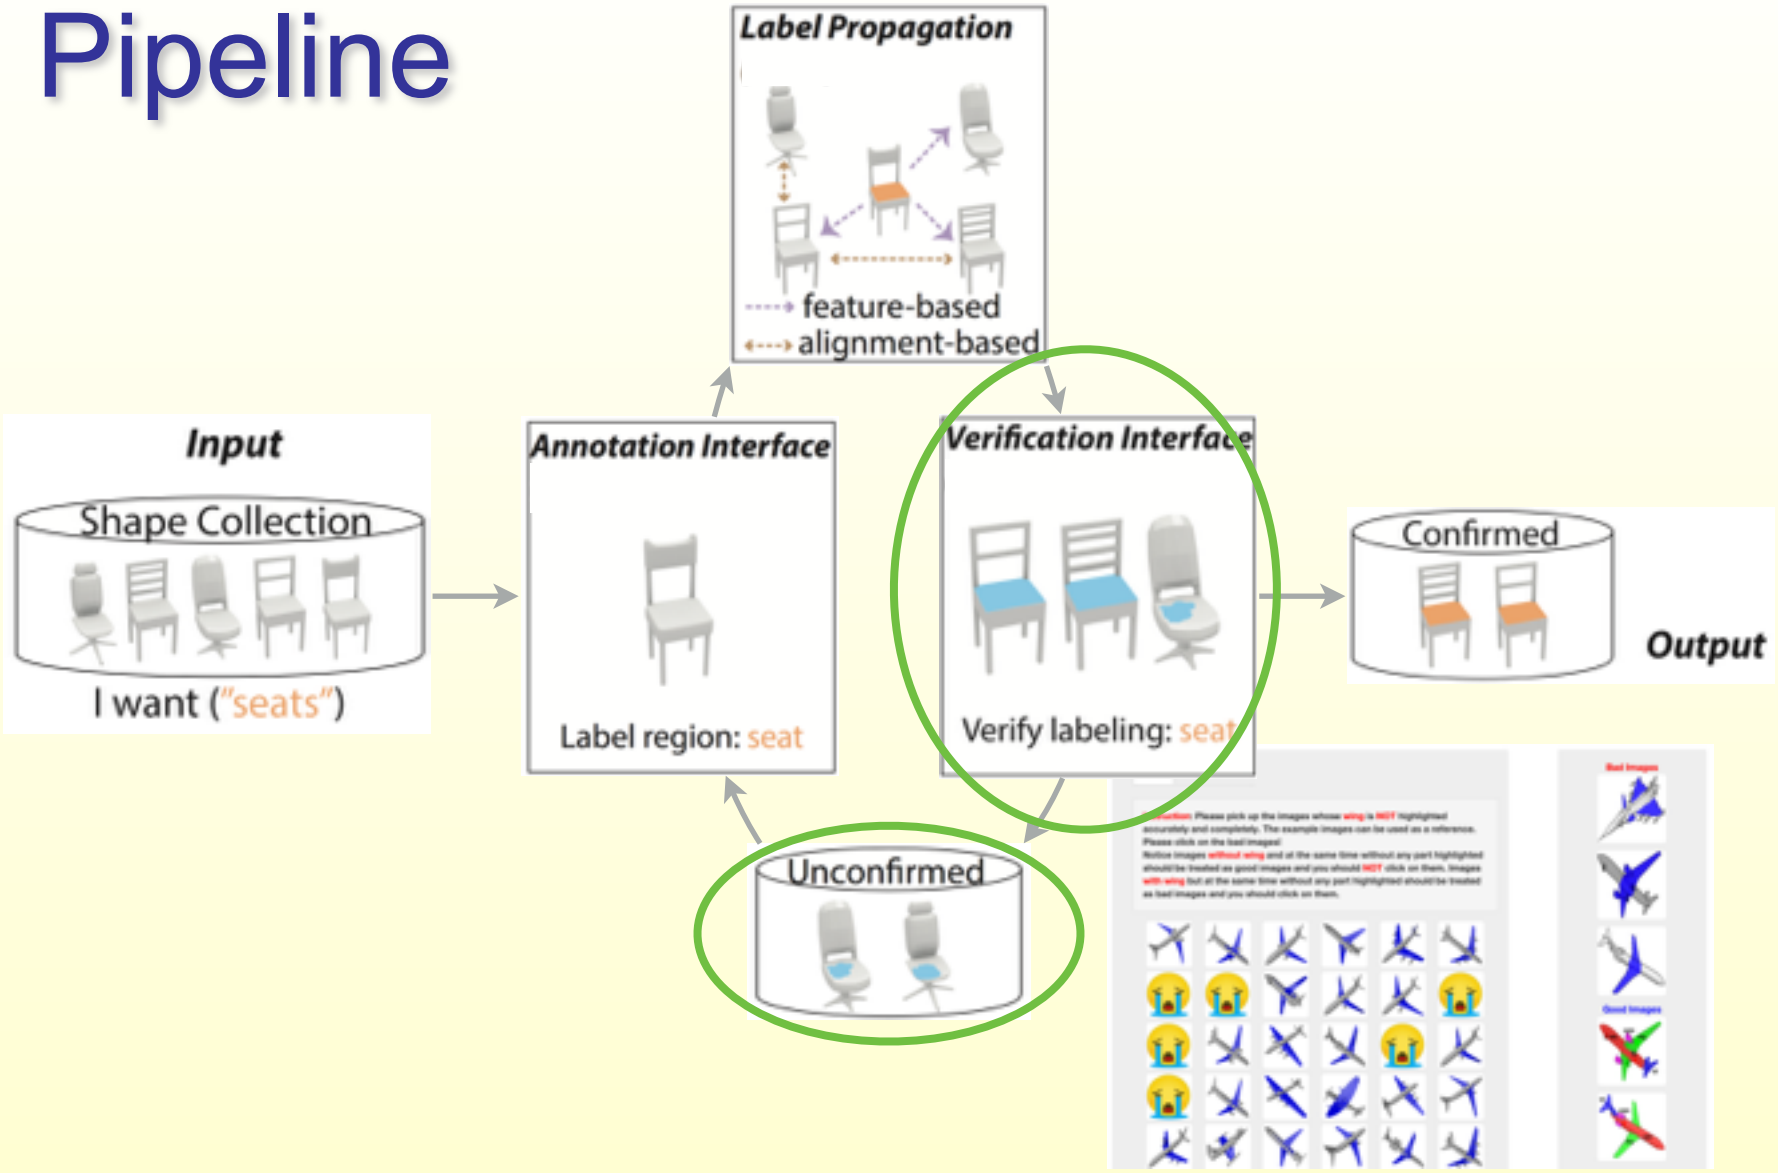
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

Pipeline

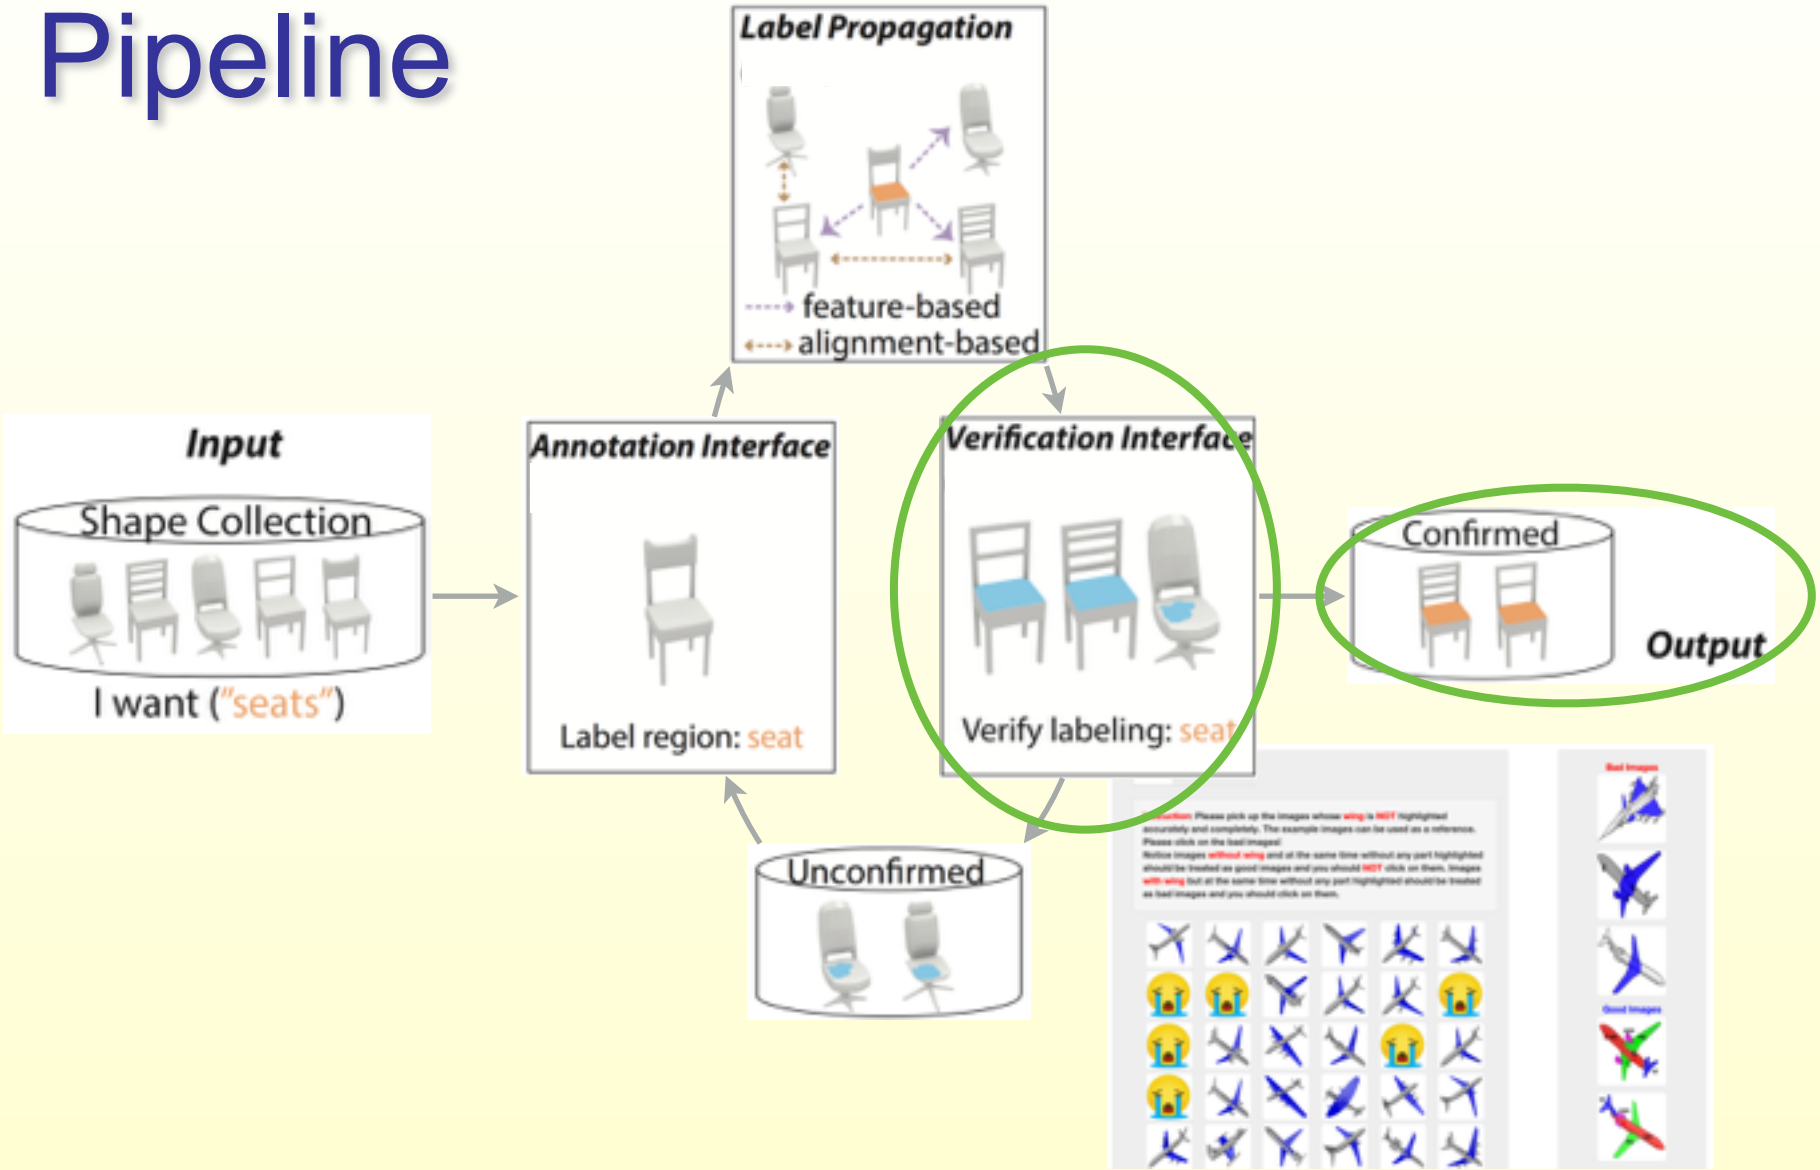
Label Propagation
feature-based
alignment-based
Input
Shape Collection
I want ("seats")
Annotation Interface
Label region: seat
Verification Interface
Verify labeling: seat
Confirmed
Output
Unconfirmed

Optimization

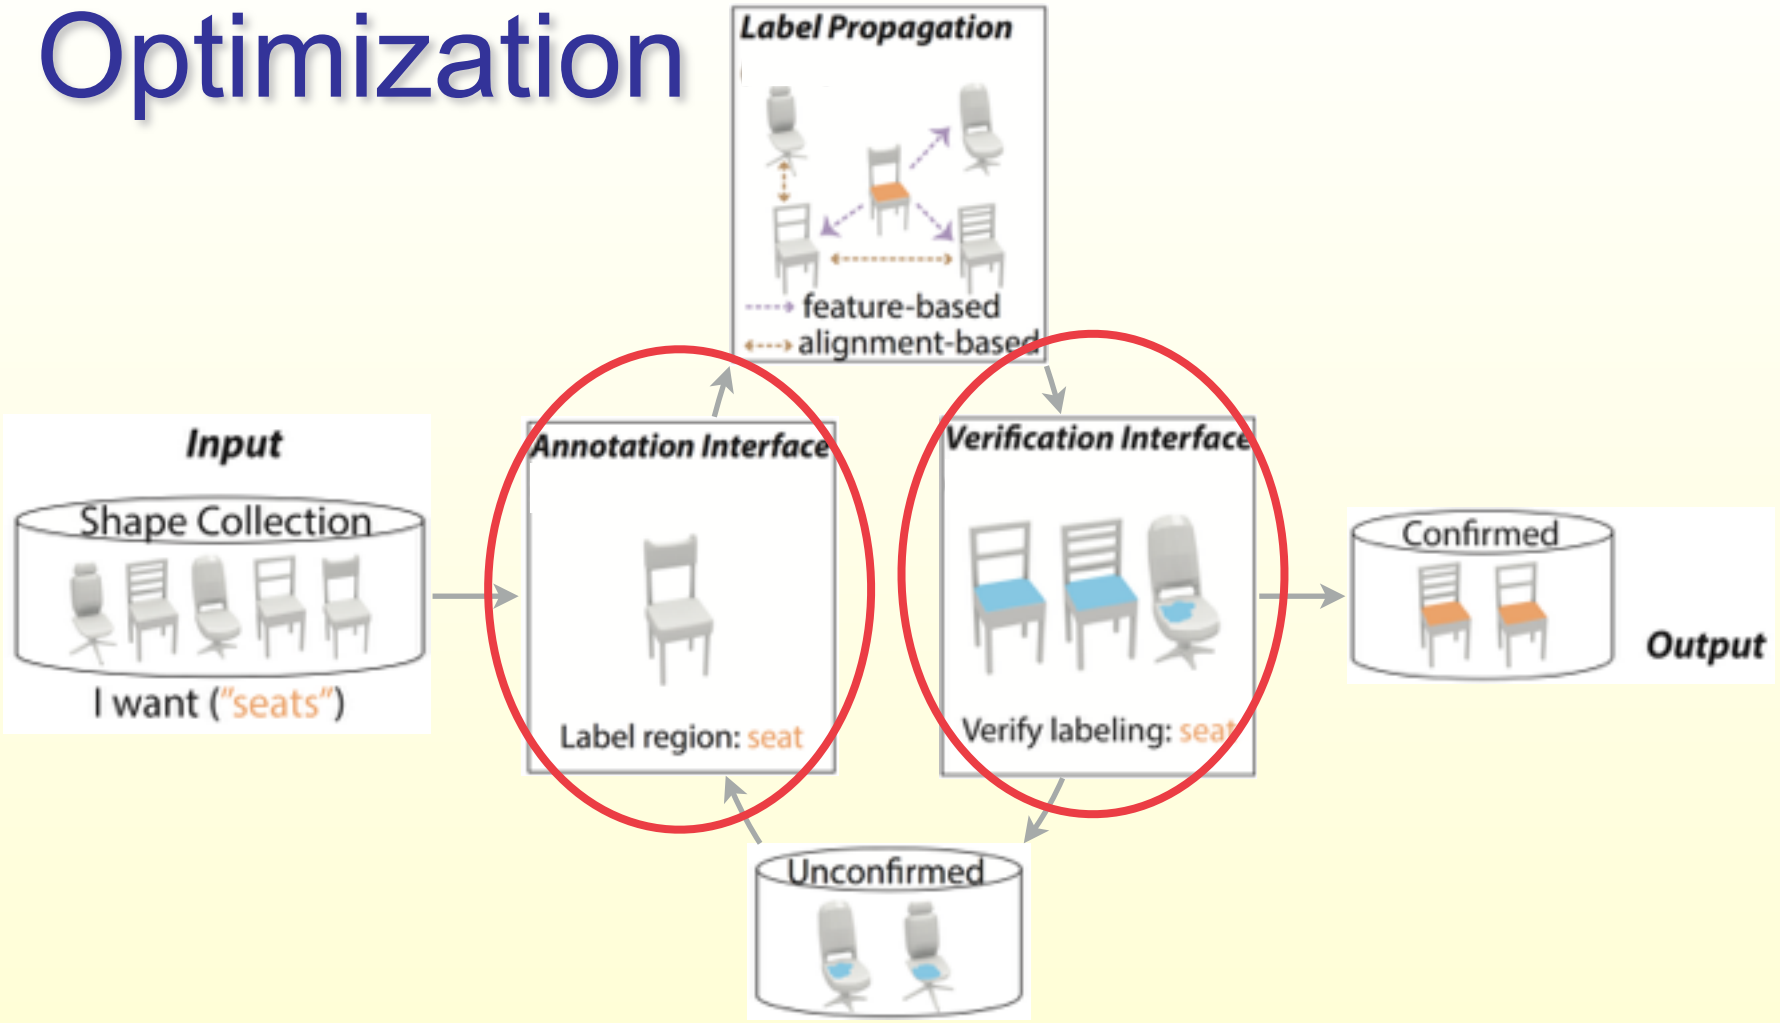

Optimization

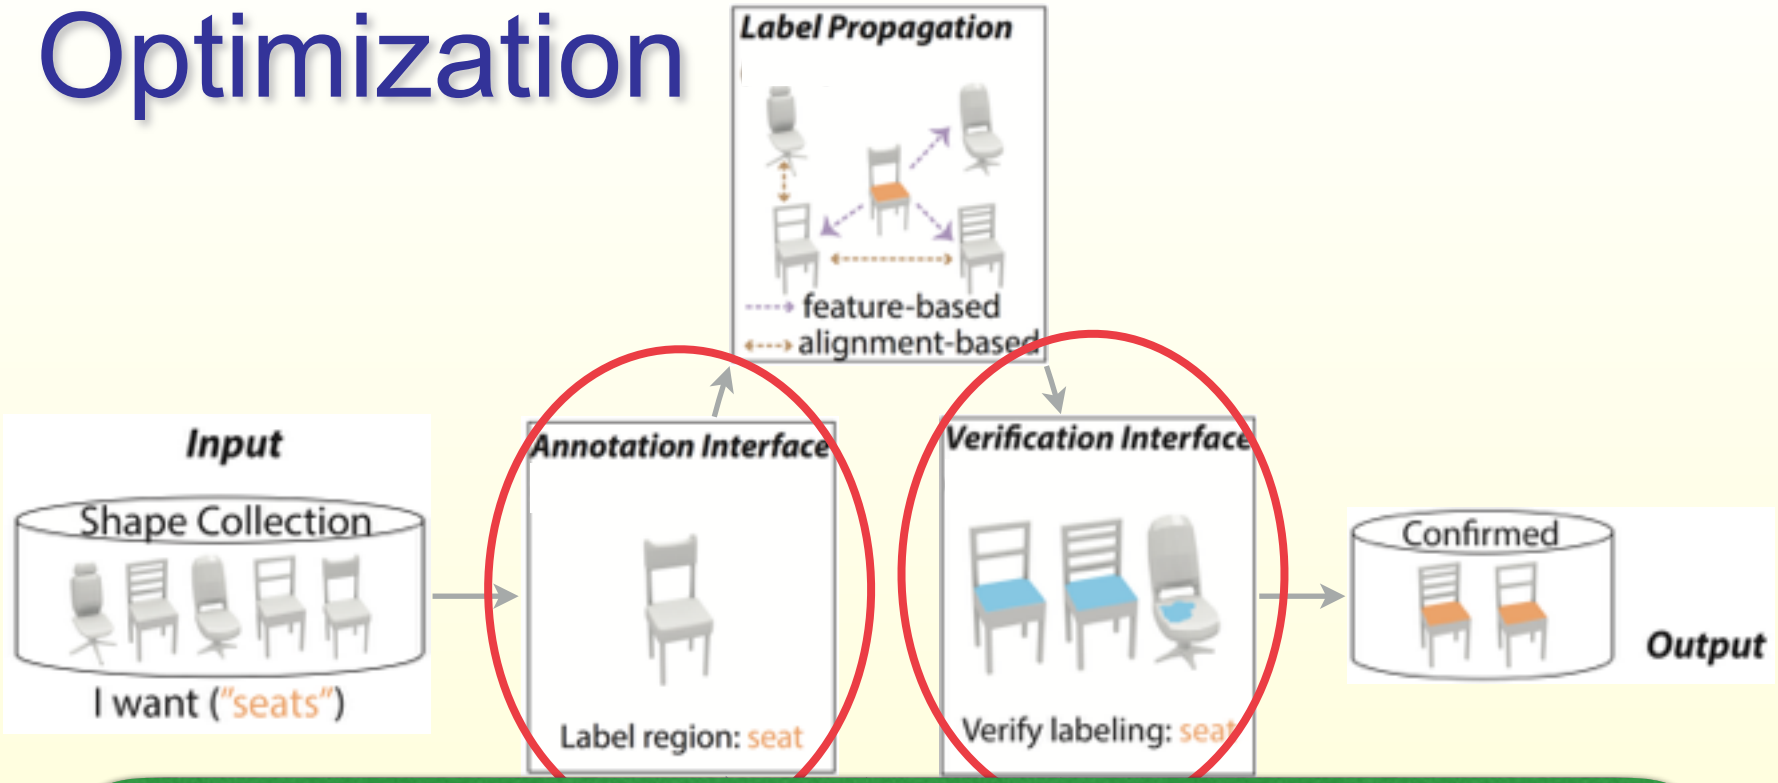

How to issue ***Annotation*** and ***Verification*** tasks

to optimize the ***Utility Function***

Optimization

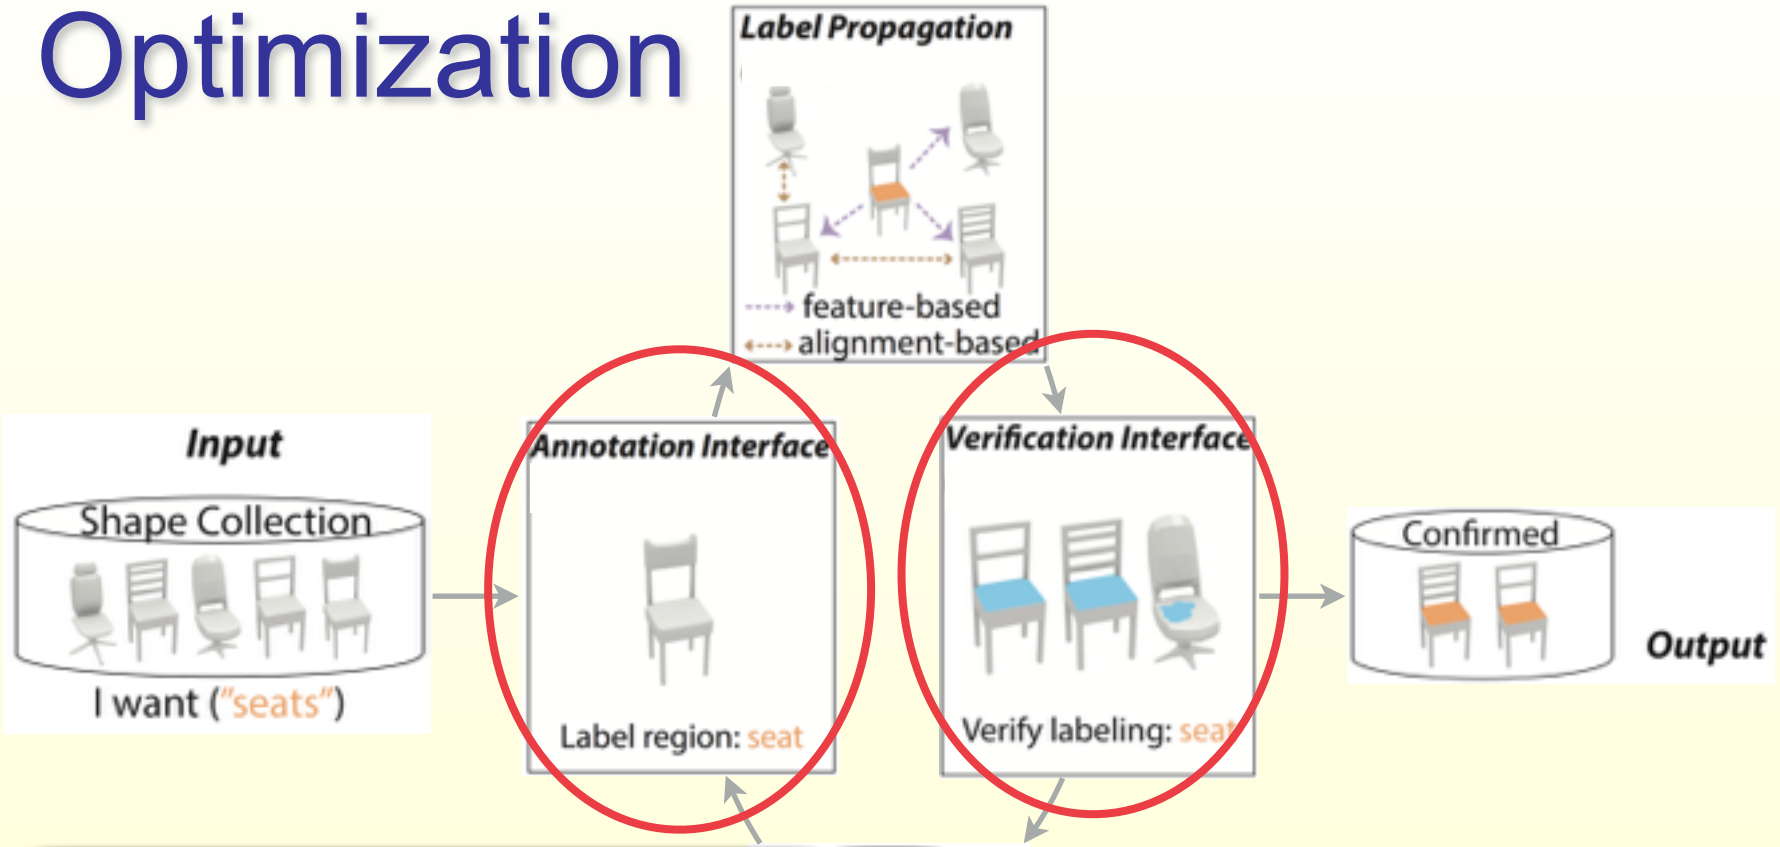

Utility Function: $\frac{\#(\textit{ACCURATE}\text{ labeling})}{\text{worker time}}$

Optimization

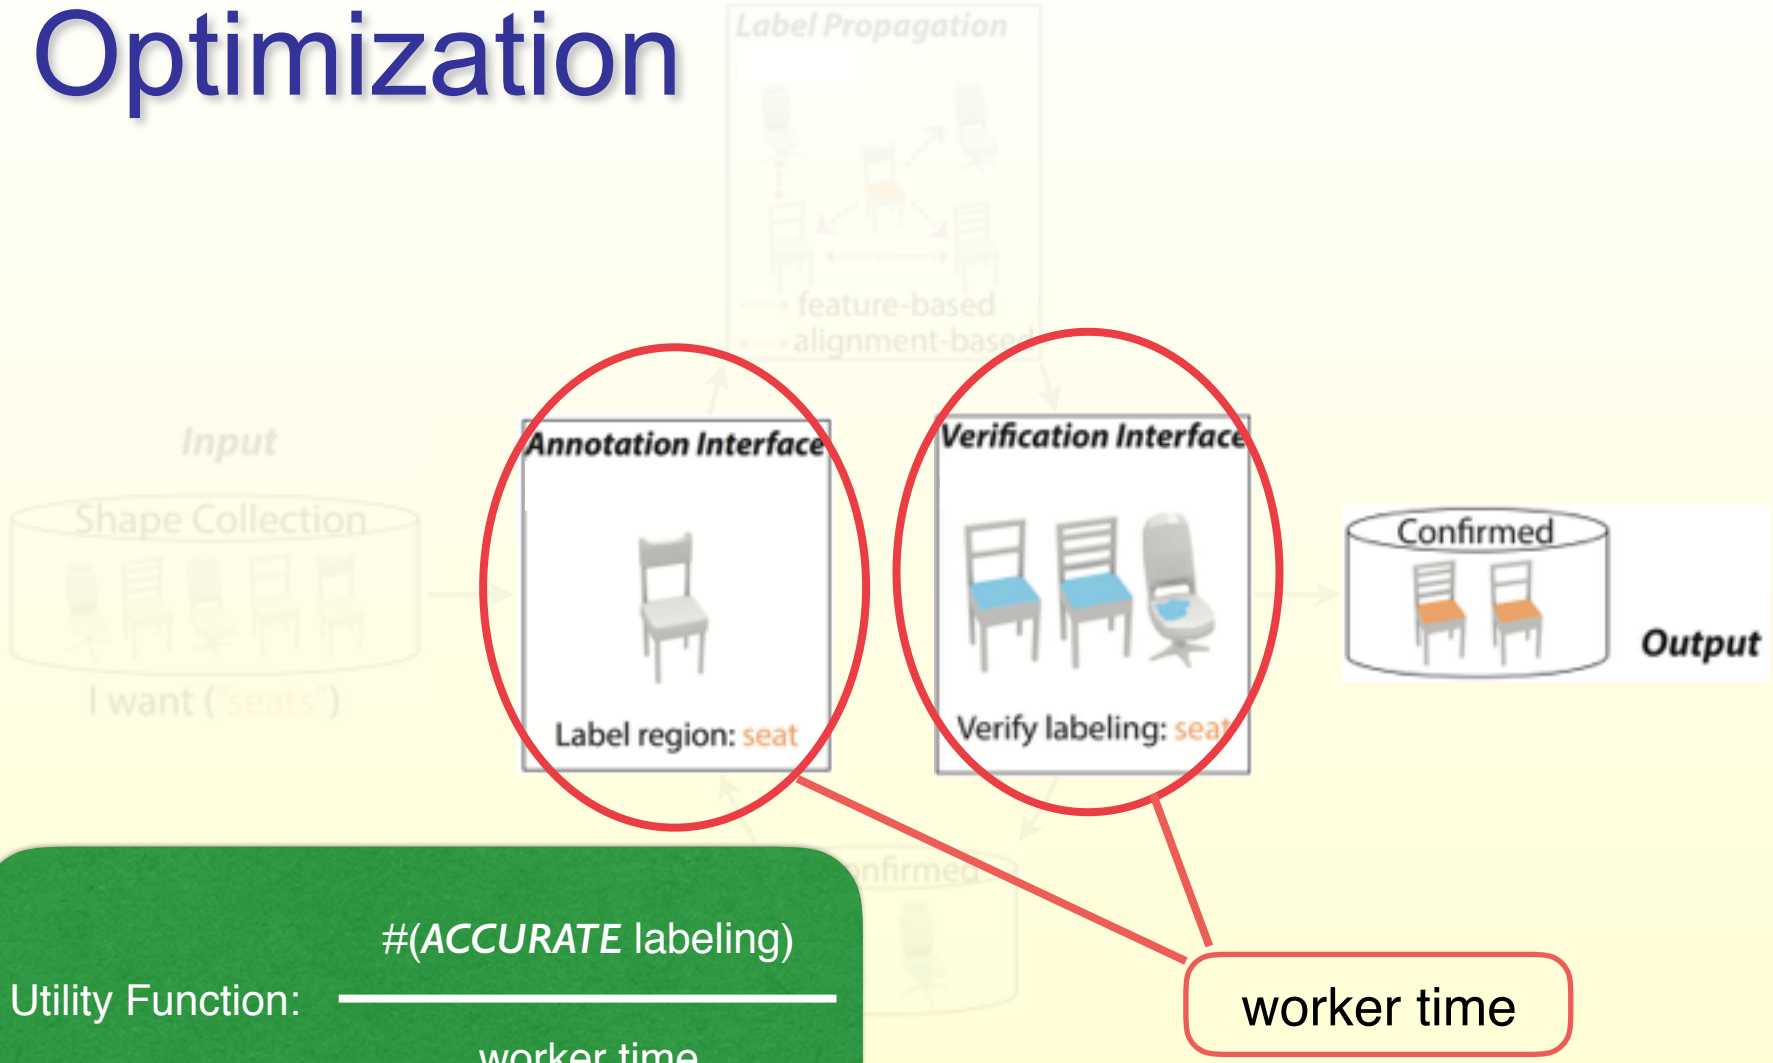

Optimization

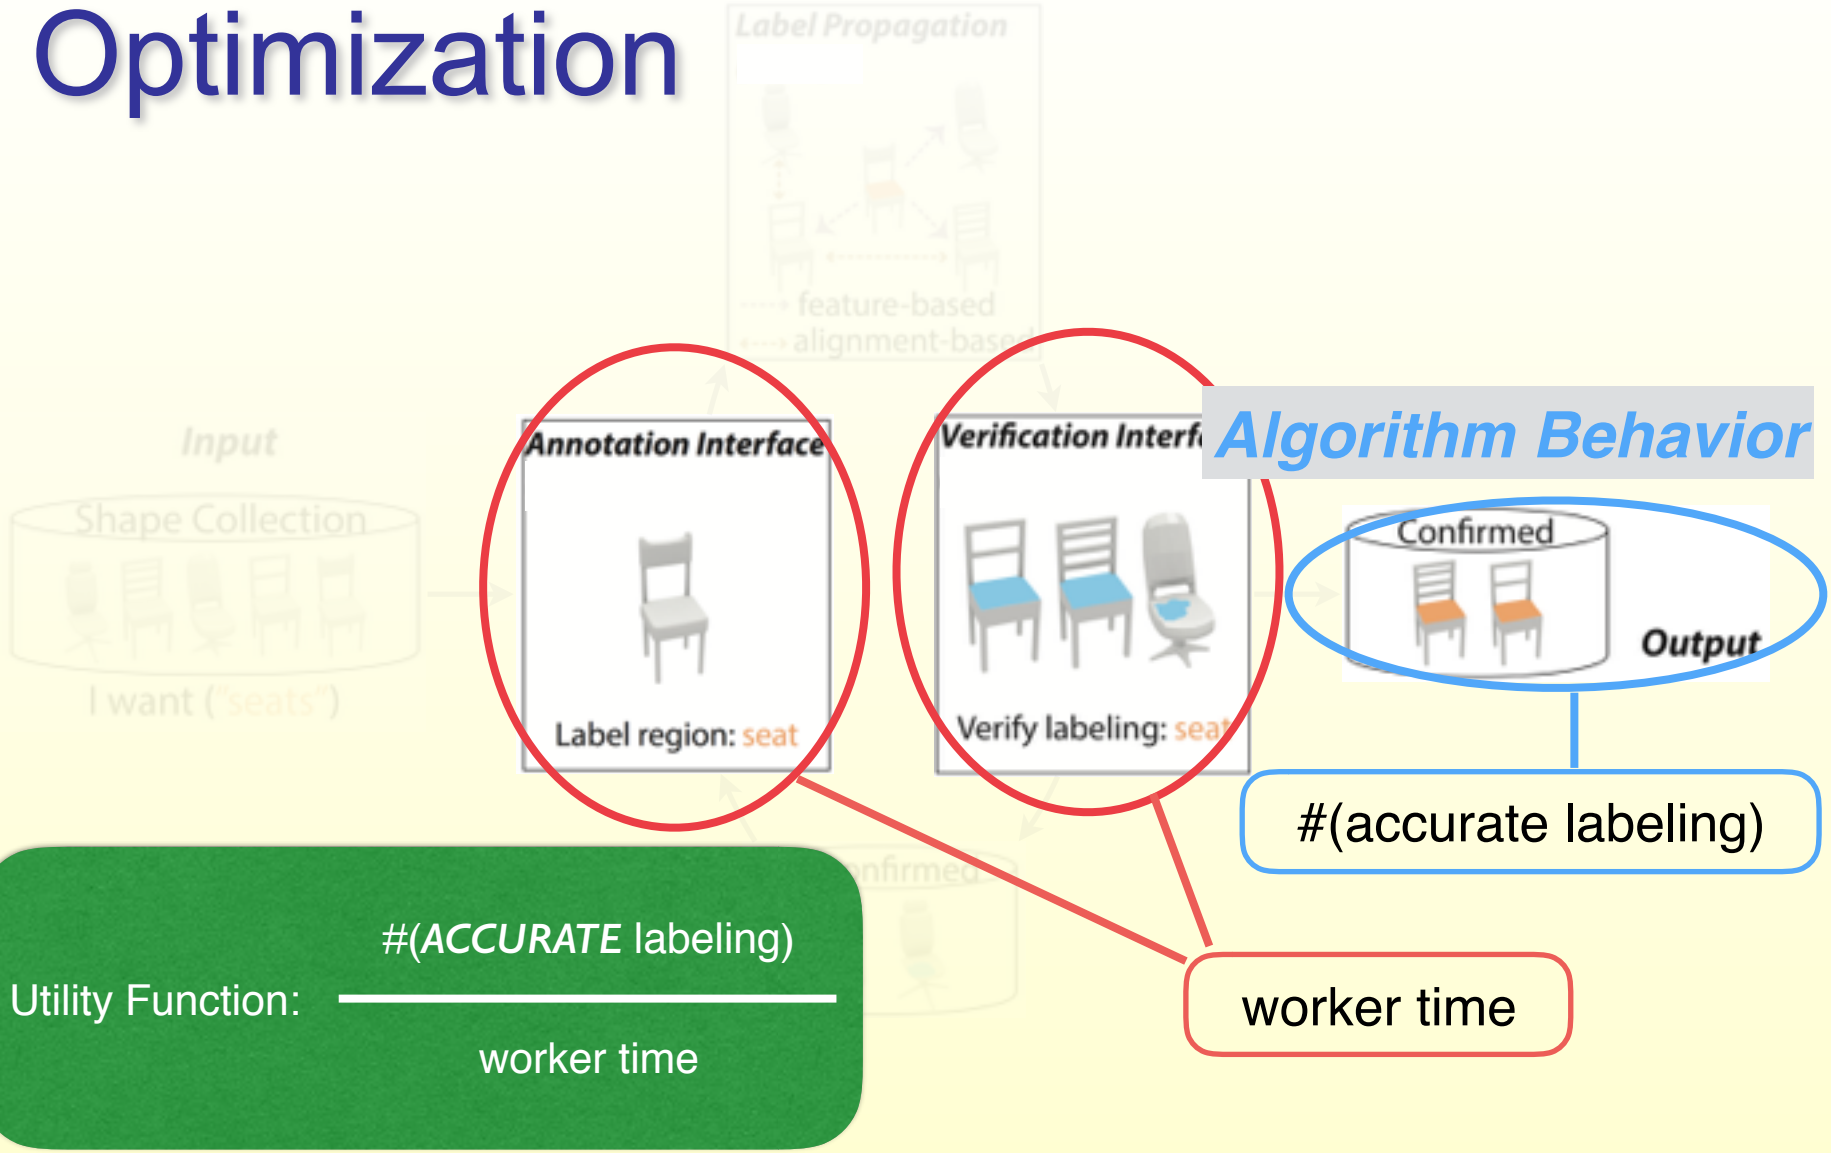

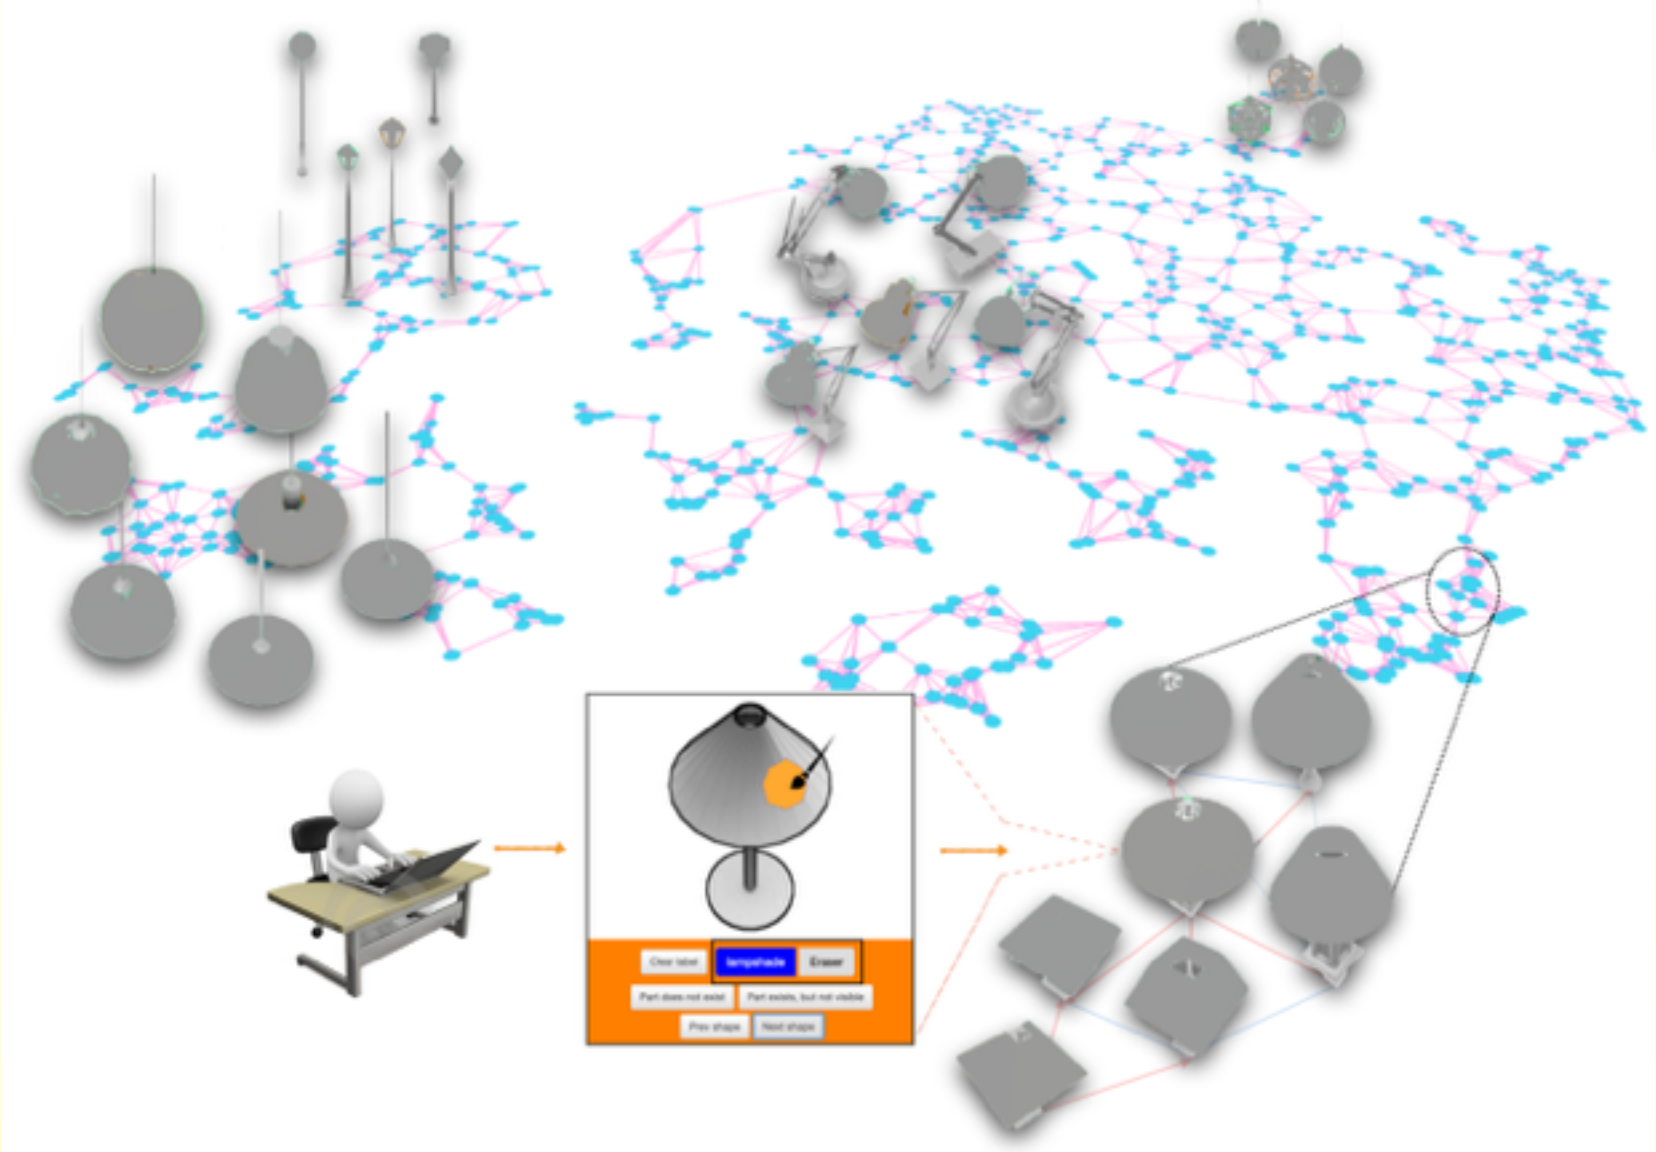
Clear label
lampshade
Eraser
Part does not exist
Part exists, but not visible
Prev shape
Next shape

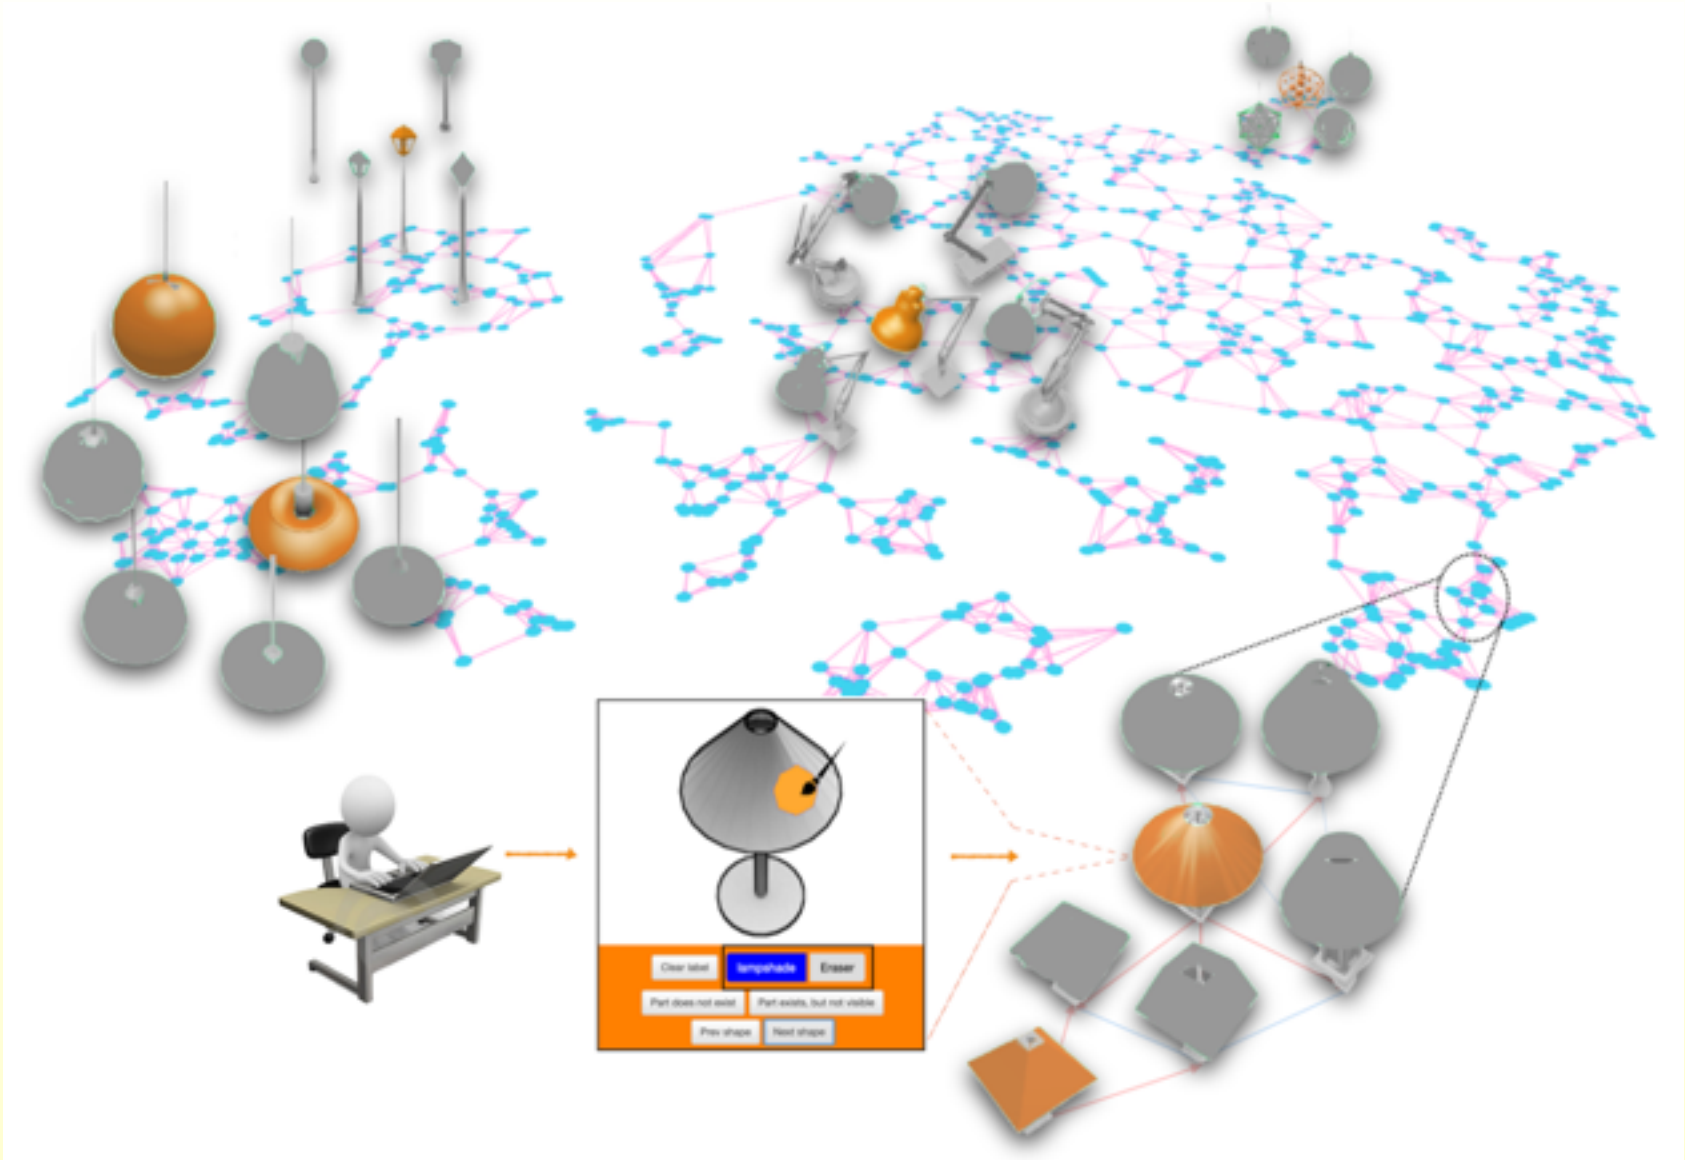
Clear label
lampshade
Eraser
Part does not exist
Part exists, but not visible
Prev shape
Next shape

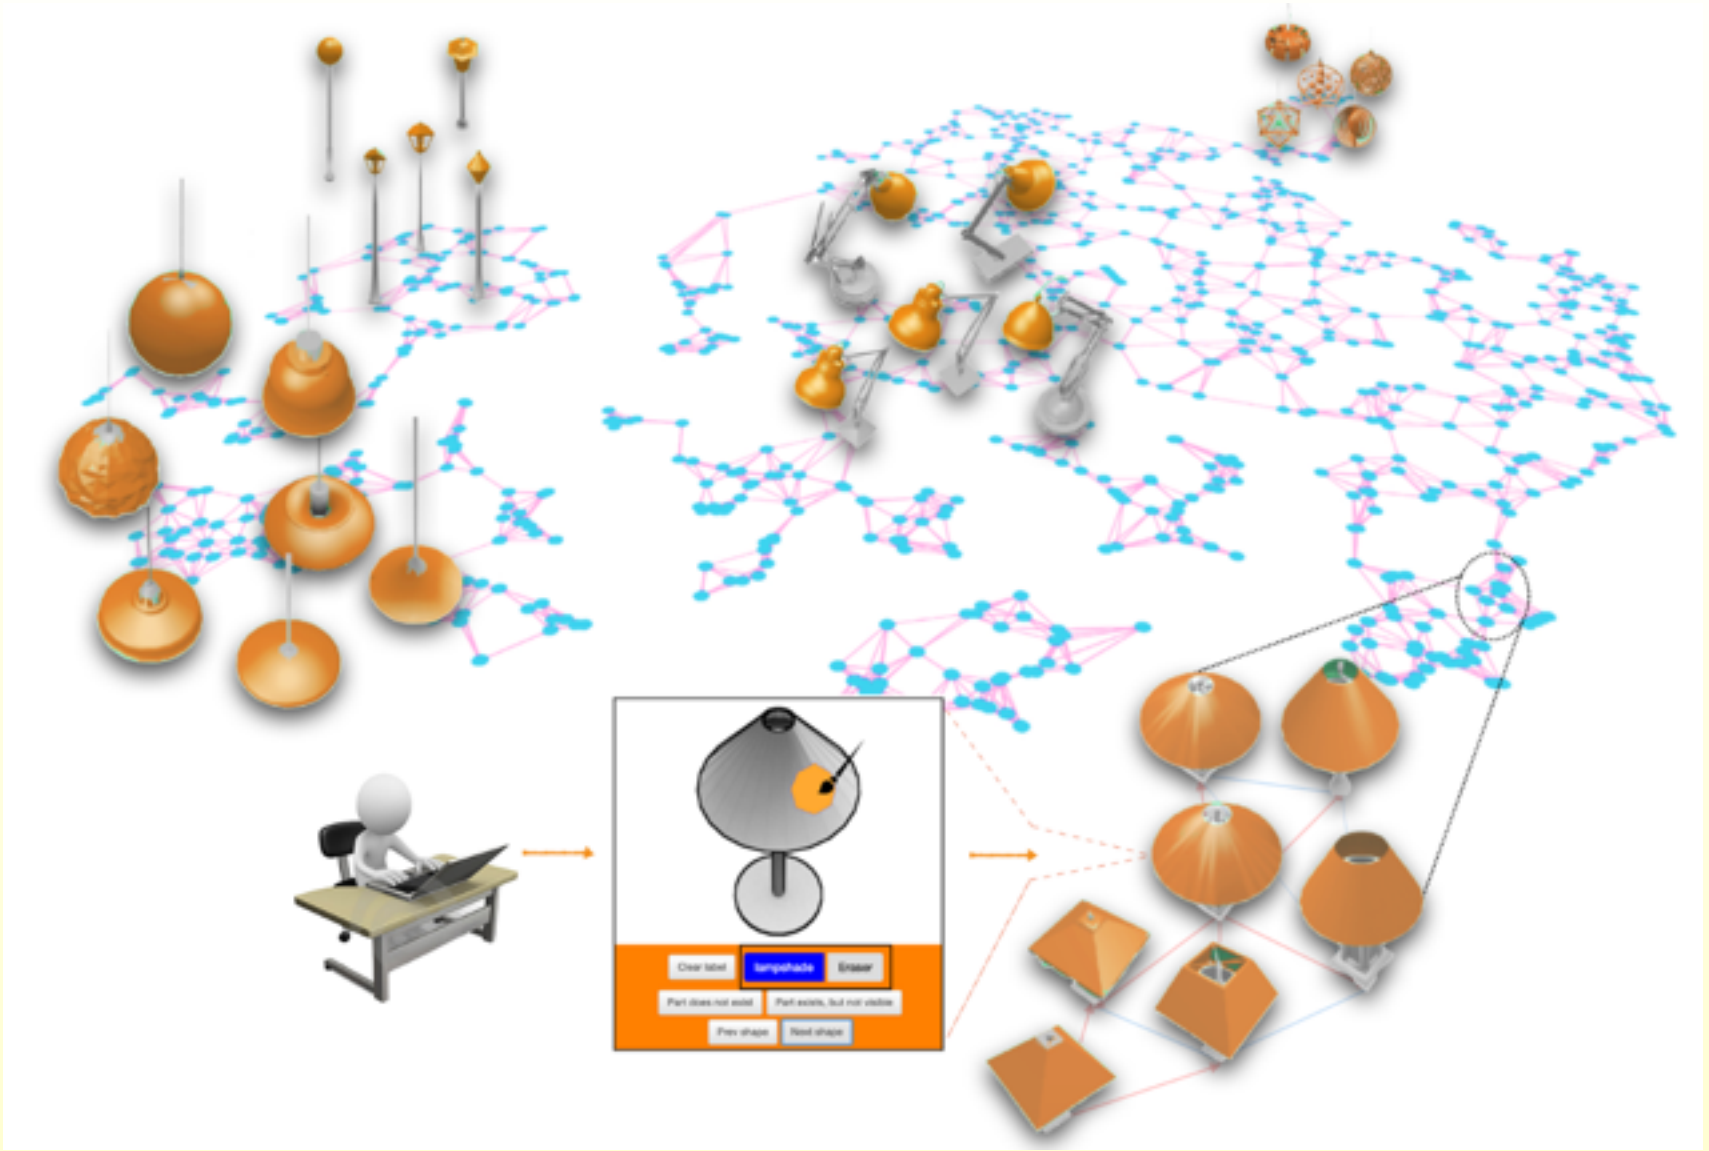
Clear label
lampshade
Eraser
Part does not exist
Part exists, but not visible
Prev shape
Next shape

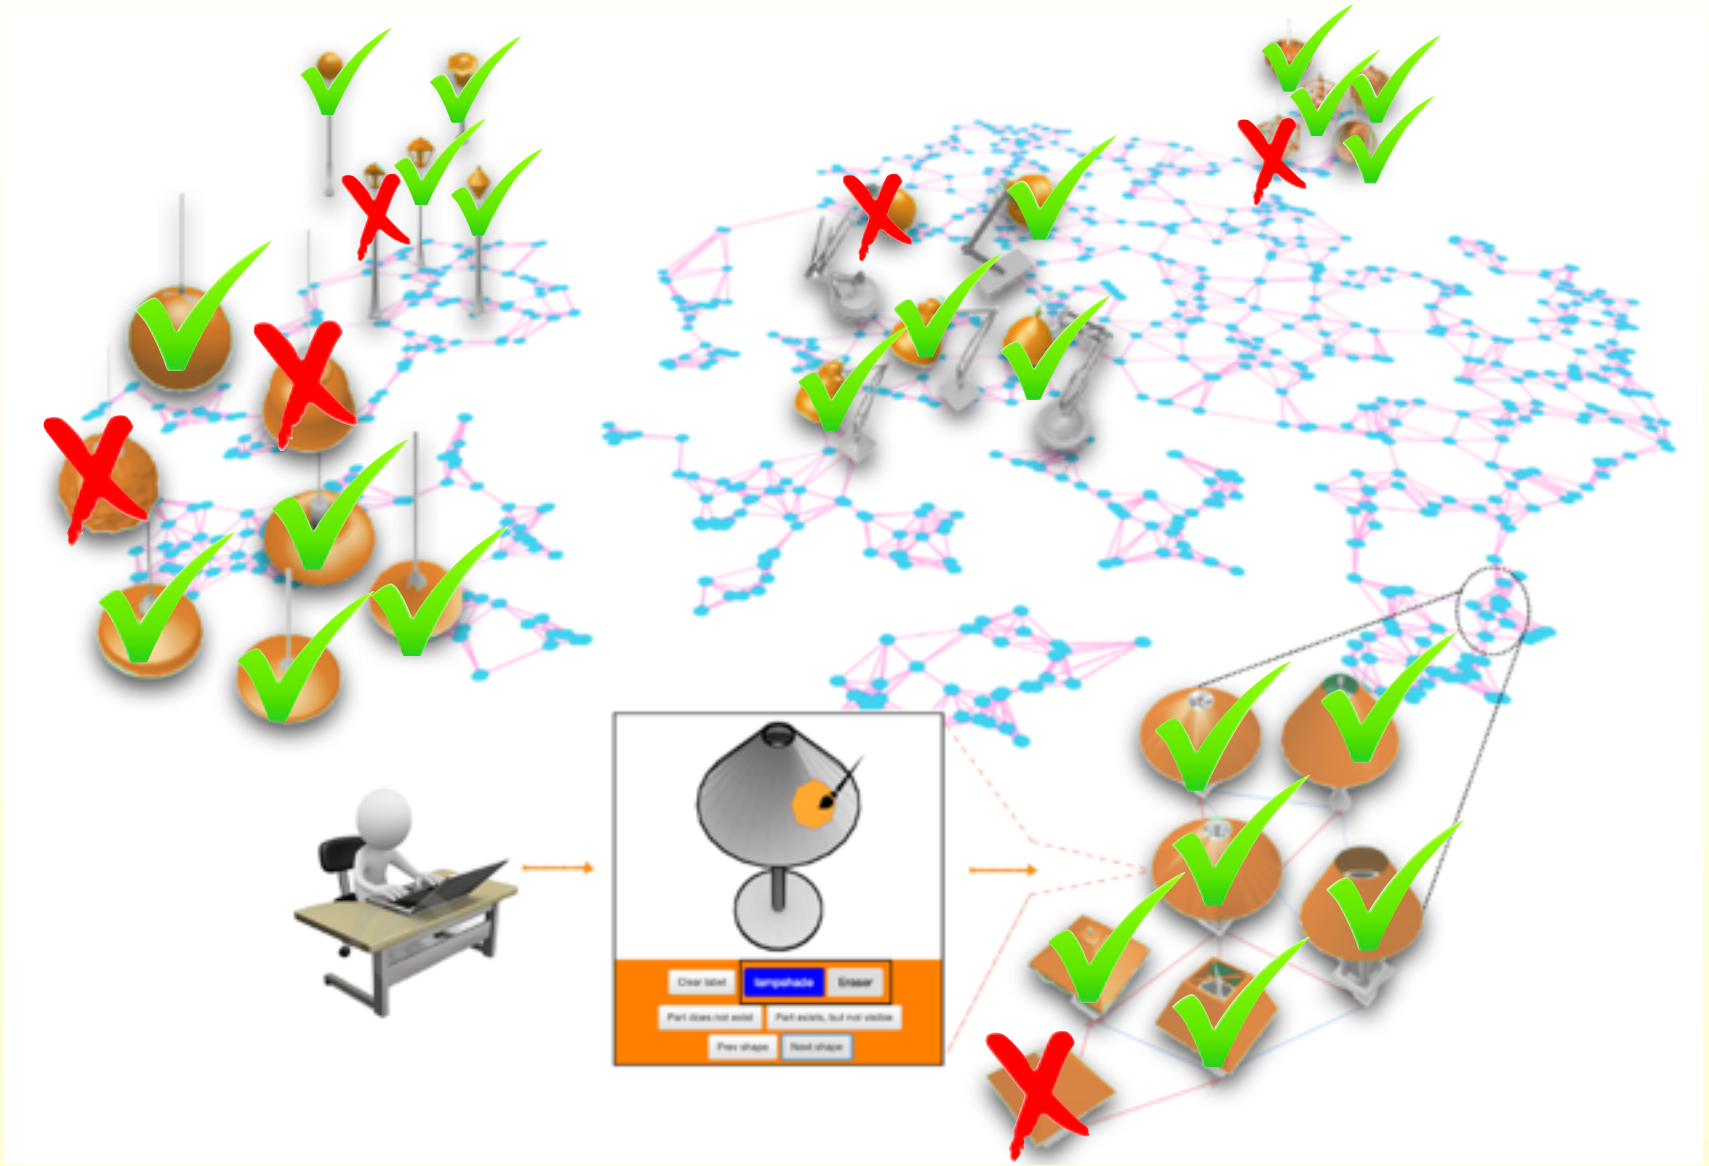
Clear label
lampshade
Eraser
Part does not exist
Part exists, but not visible
Prev shape
Next shape

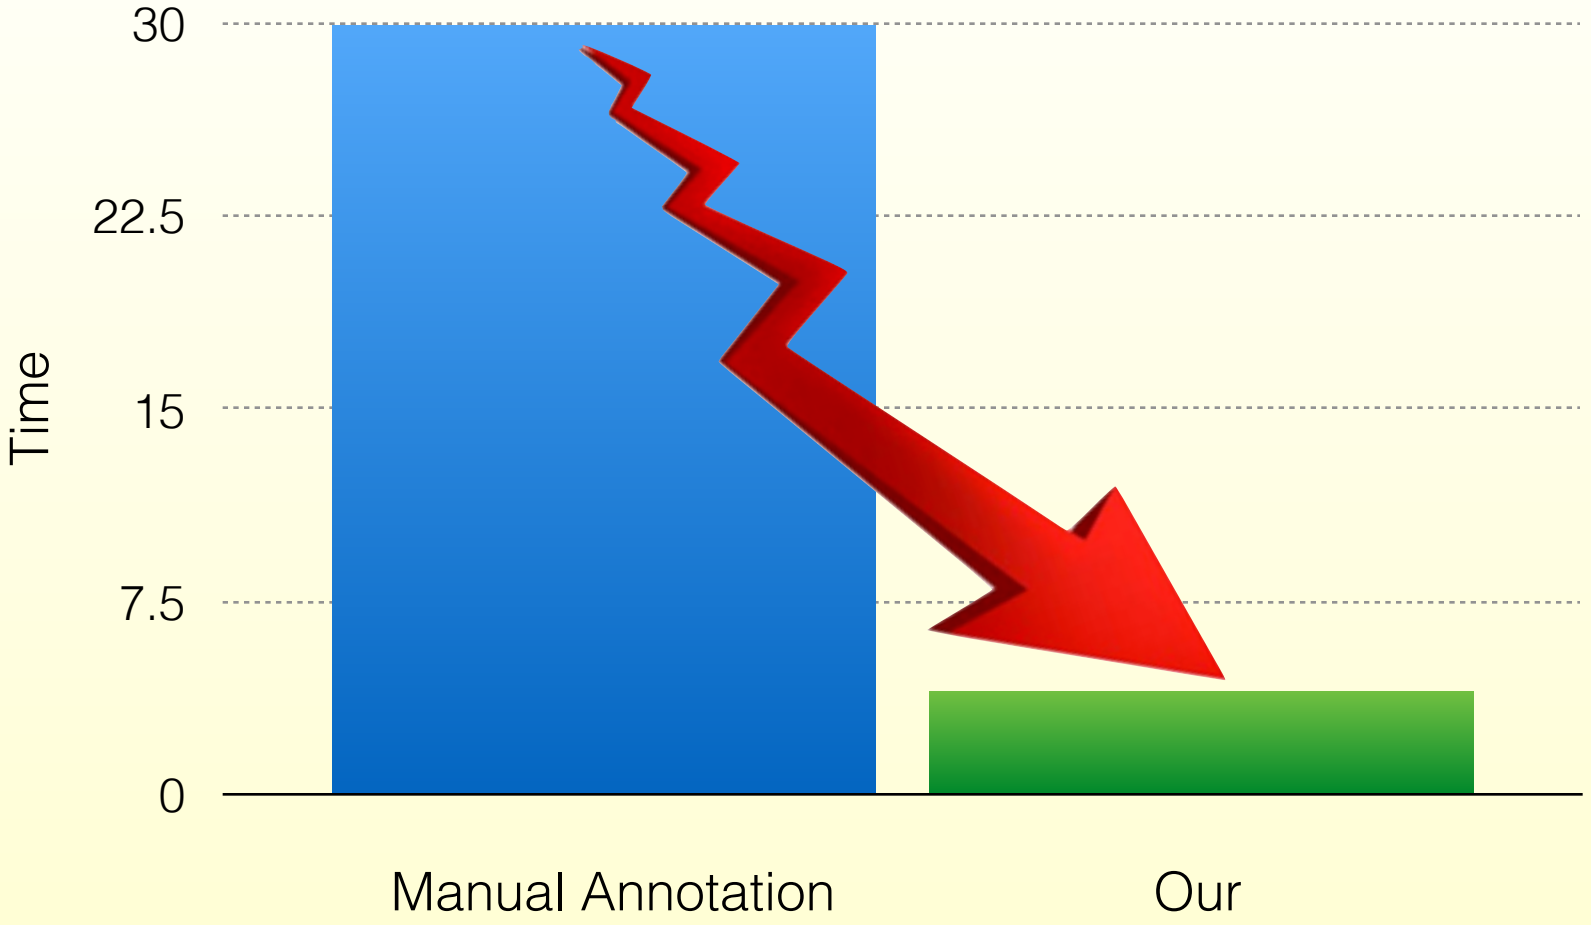
30
22.5
15
7.5
0
Time
Manual Annotation
Our

Results

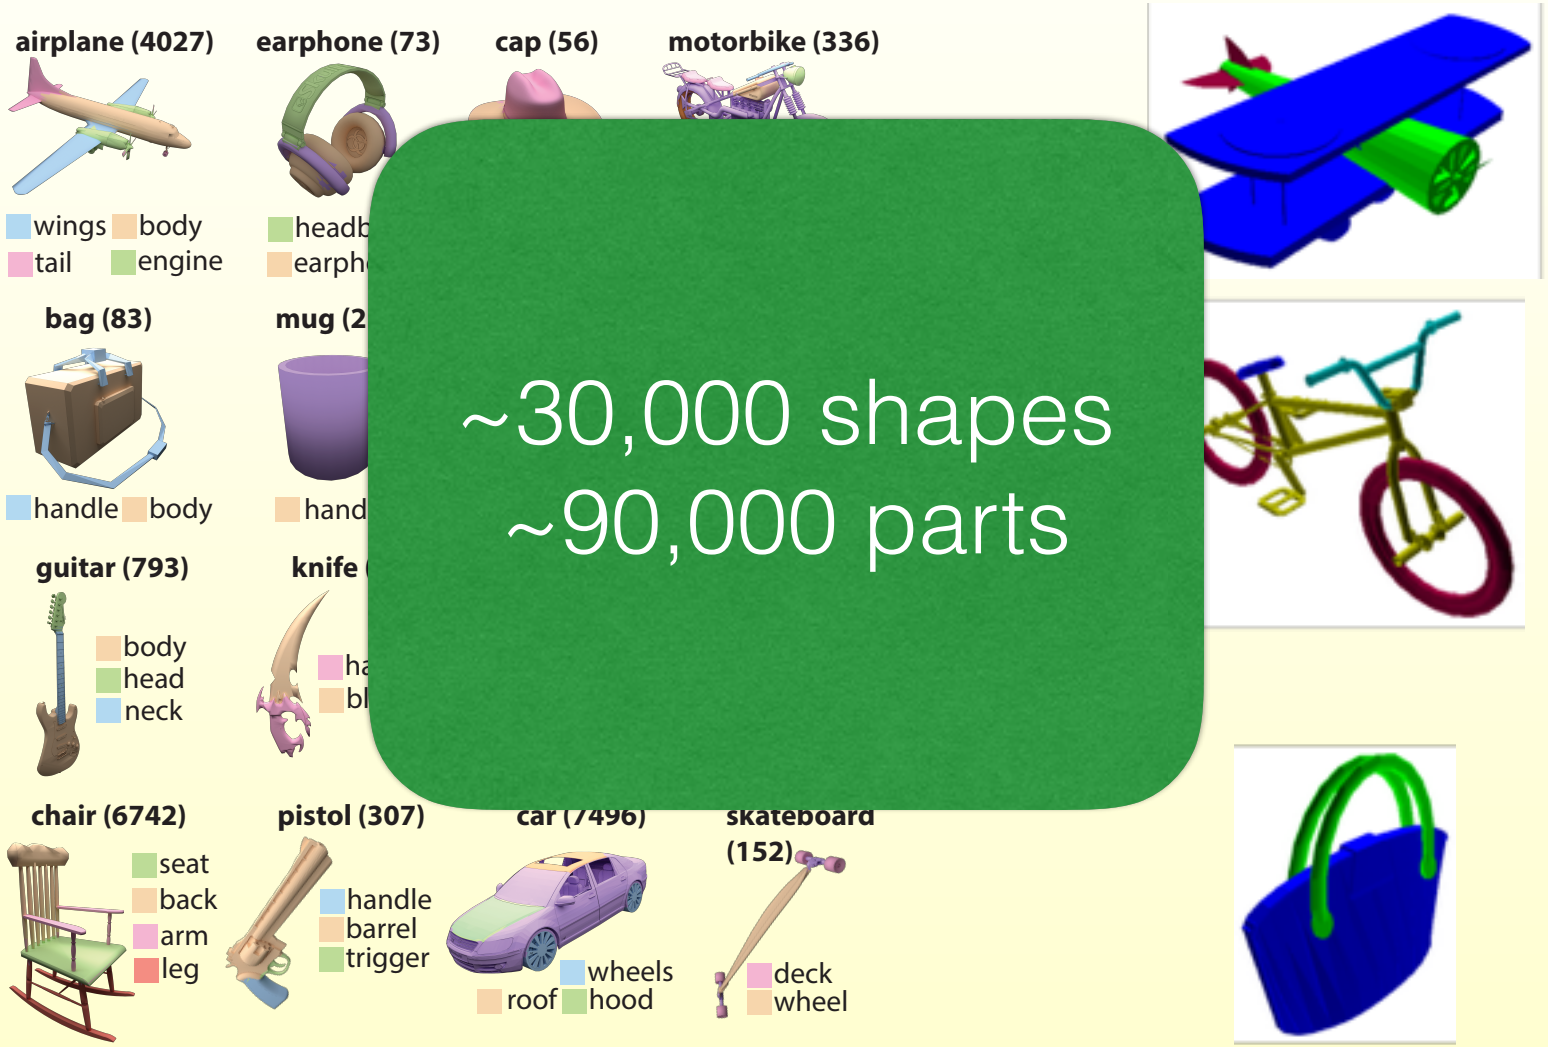
airplane (4027)
wings
body
tail
engine
earphone (73)
cap (56)
motorbike (336)
bag (83)
handle
body
guitar (793)
body
head
neck
chair (6742)
seat
back
arm
leg
pistol (307)
handle
barrel
trigger
car (7496)
wheels
roof
hood
skateboard (152)
deck
wheel
~30,000 shapes
~90,000 parts

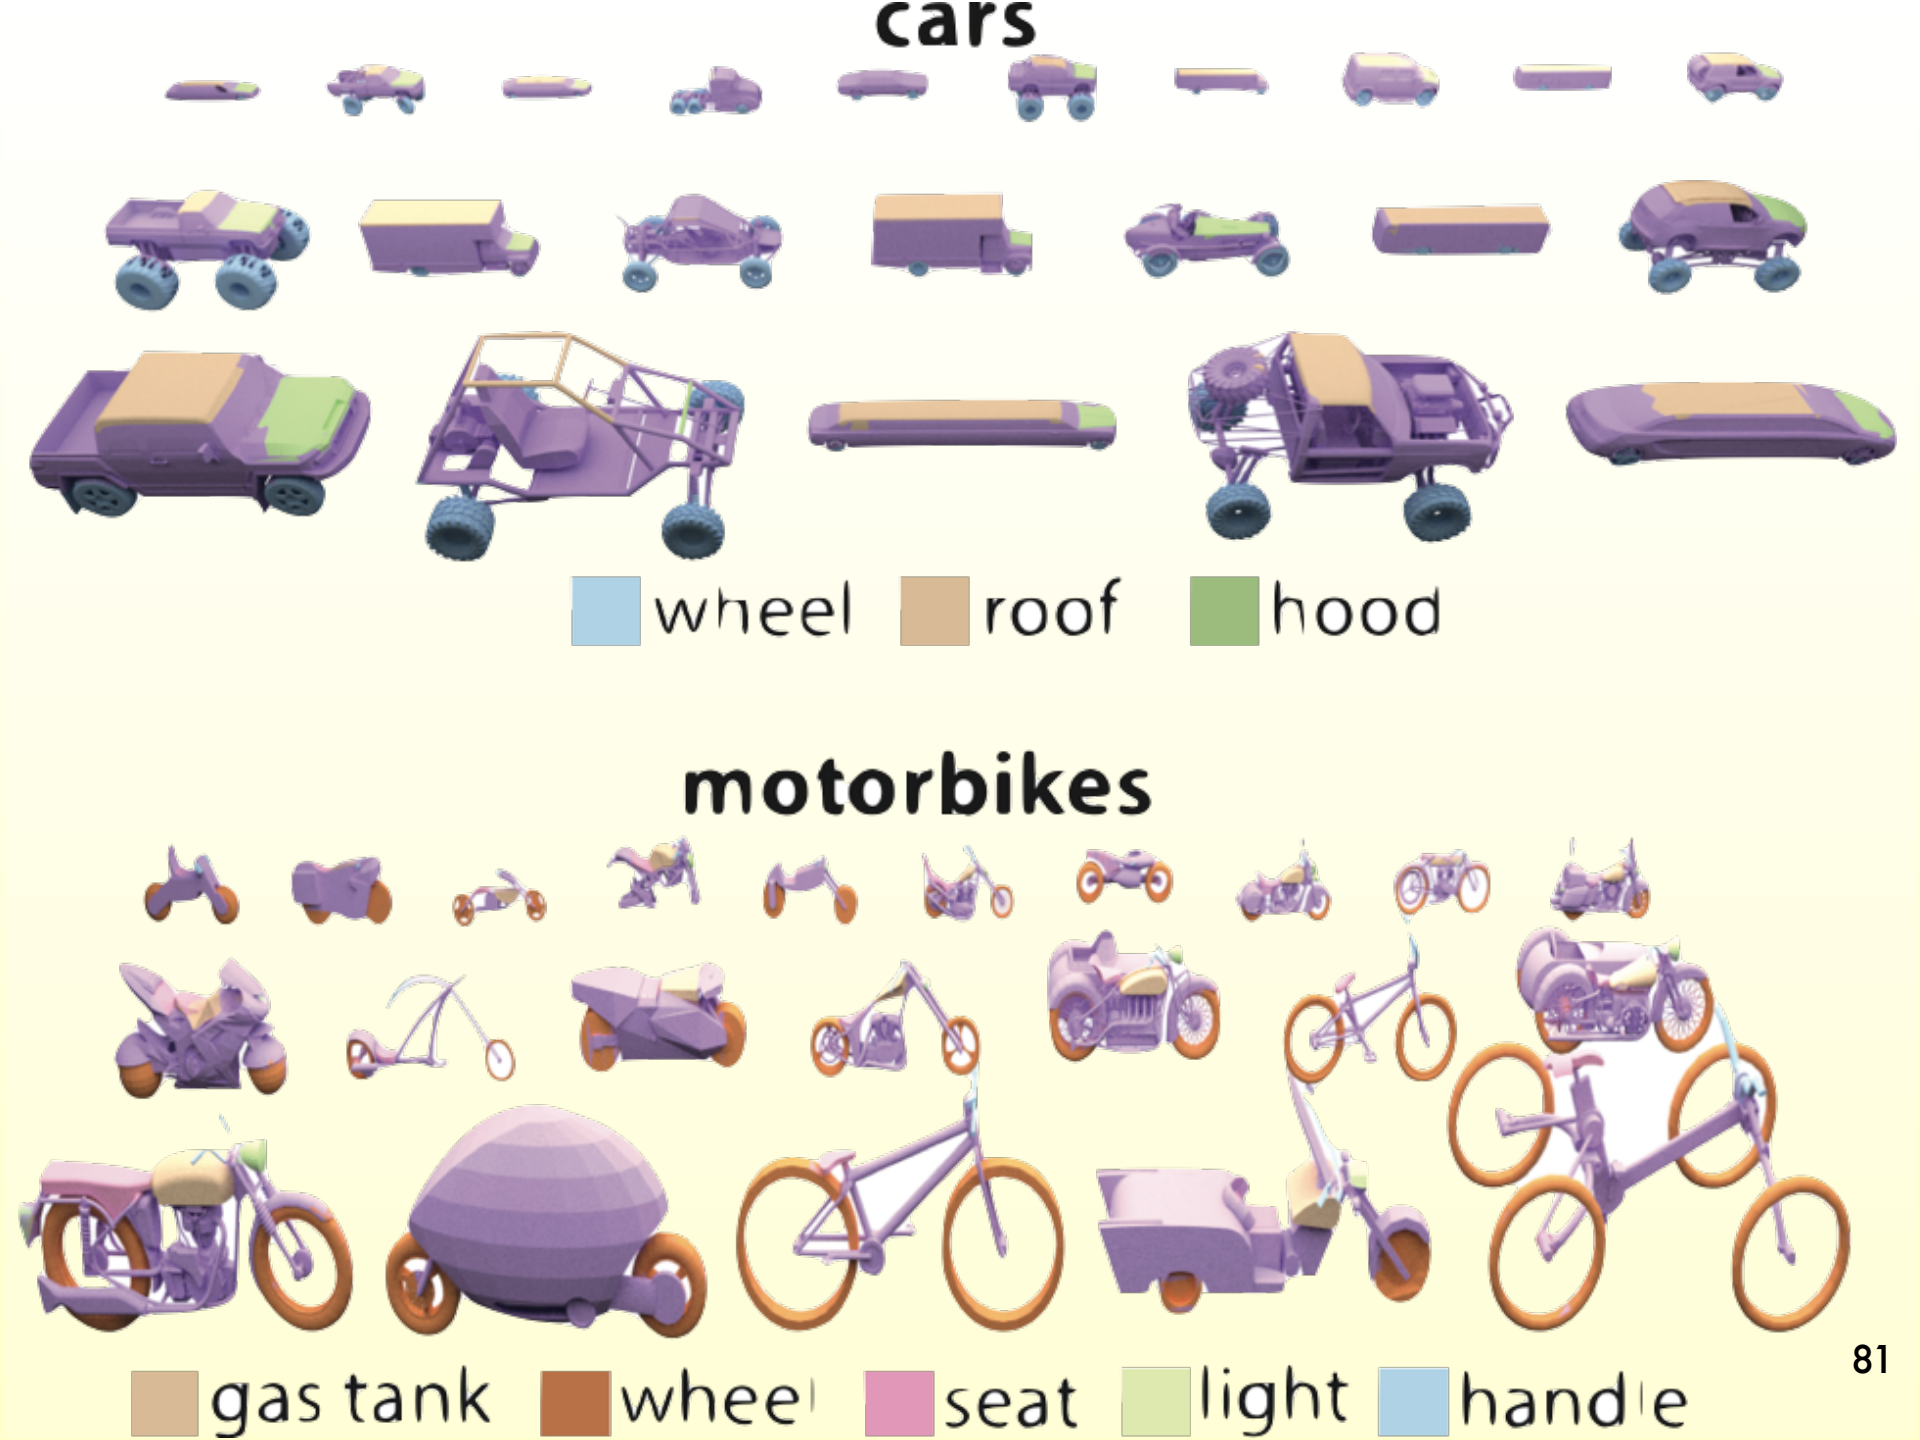
cars
wheel
roof
hood
motorbikes
gas tank
wheel
seat
light
handle

pistols

handle barrel trigger

airplanes

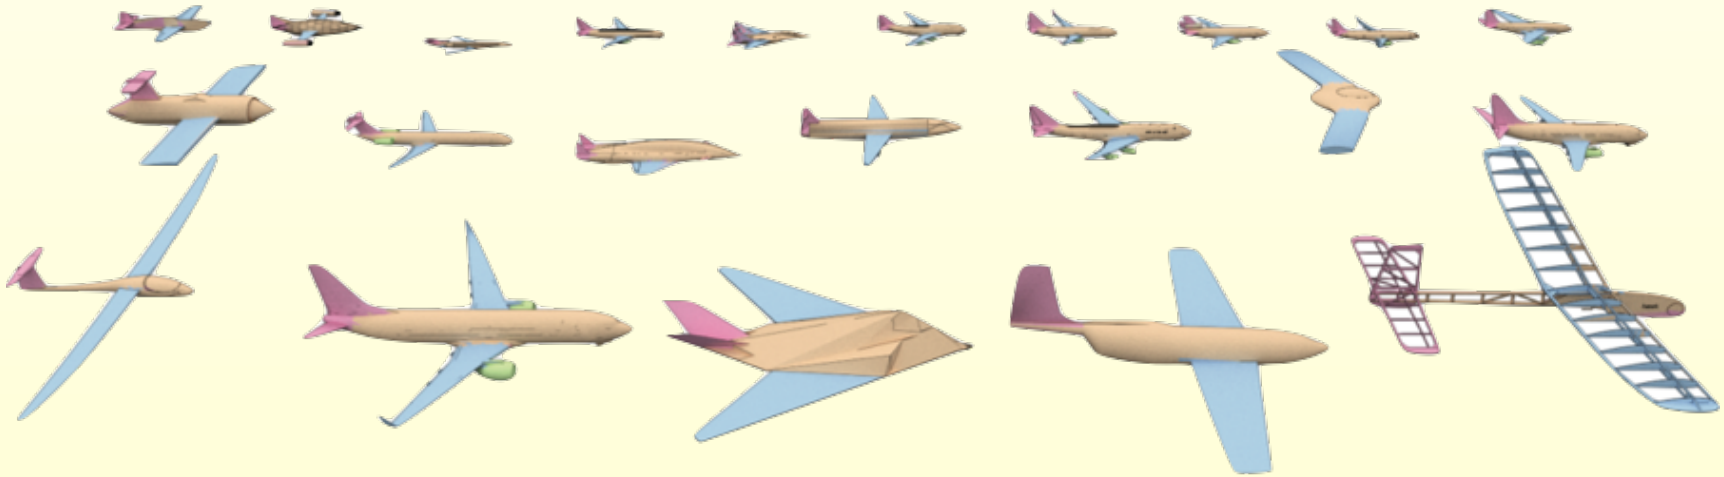

body wing engine tail

bags

handle body

lamps

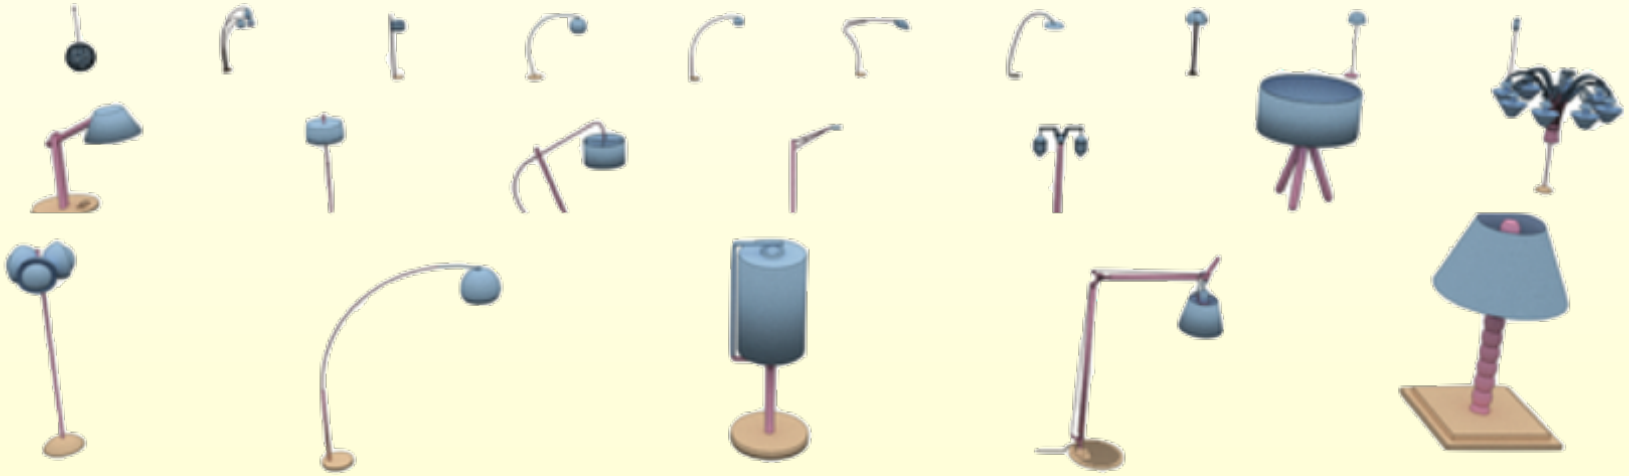

shade base tube

Take-Home Message

- The hybrid data labeling approach boosts the annotation efficiency dramatically
- The annotation and verification tasks need to be carefully balanced
- The “vertical” label prediction plus the “horizontal” label propagation is the key for success

Applications of Horizontal Networks

- A scalable active framework for part annotation in ShapeNet
- Learning hierarchical shape segmentation and labeling in a weakly supervised manner from online repositories

Problem Definition

Car

Input

Problem Definition

Car

Input

Parts

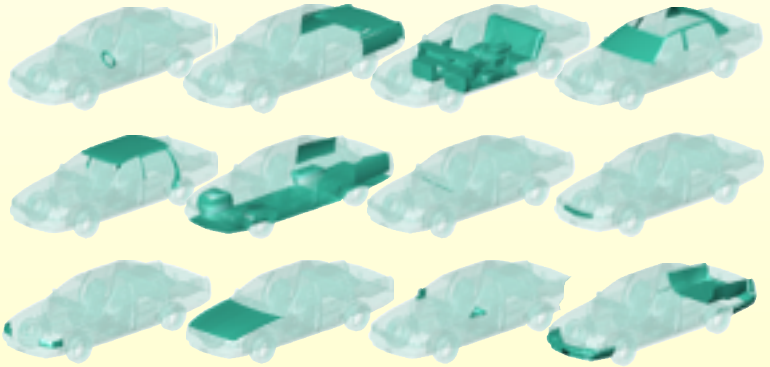

Problem Definition

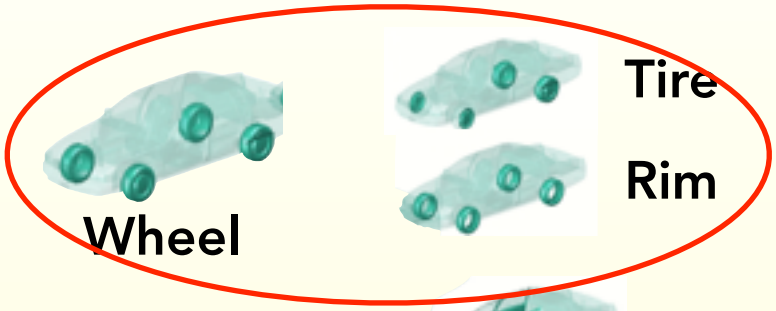

Parts Labels


Car


Door

Input


Body

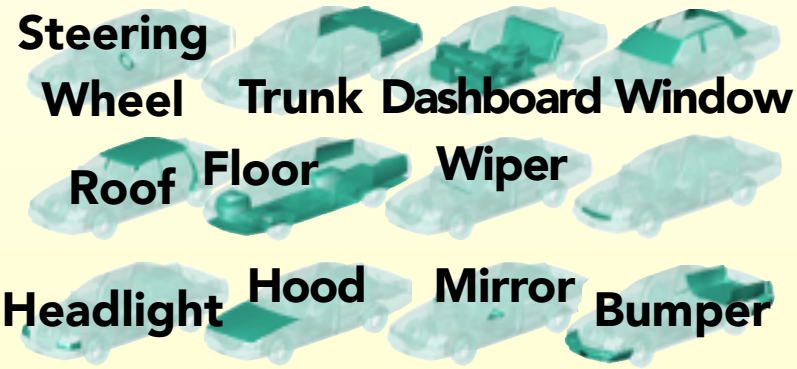

Problem Definition

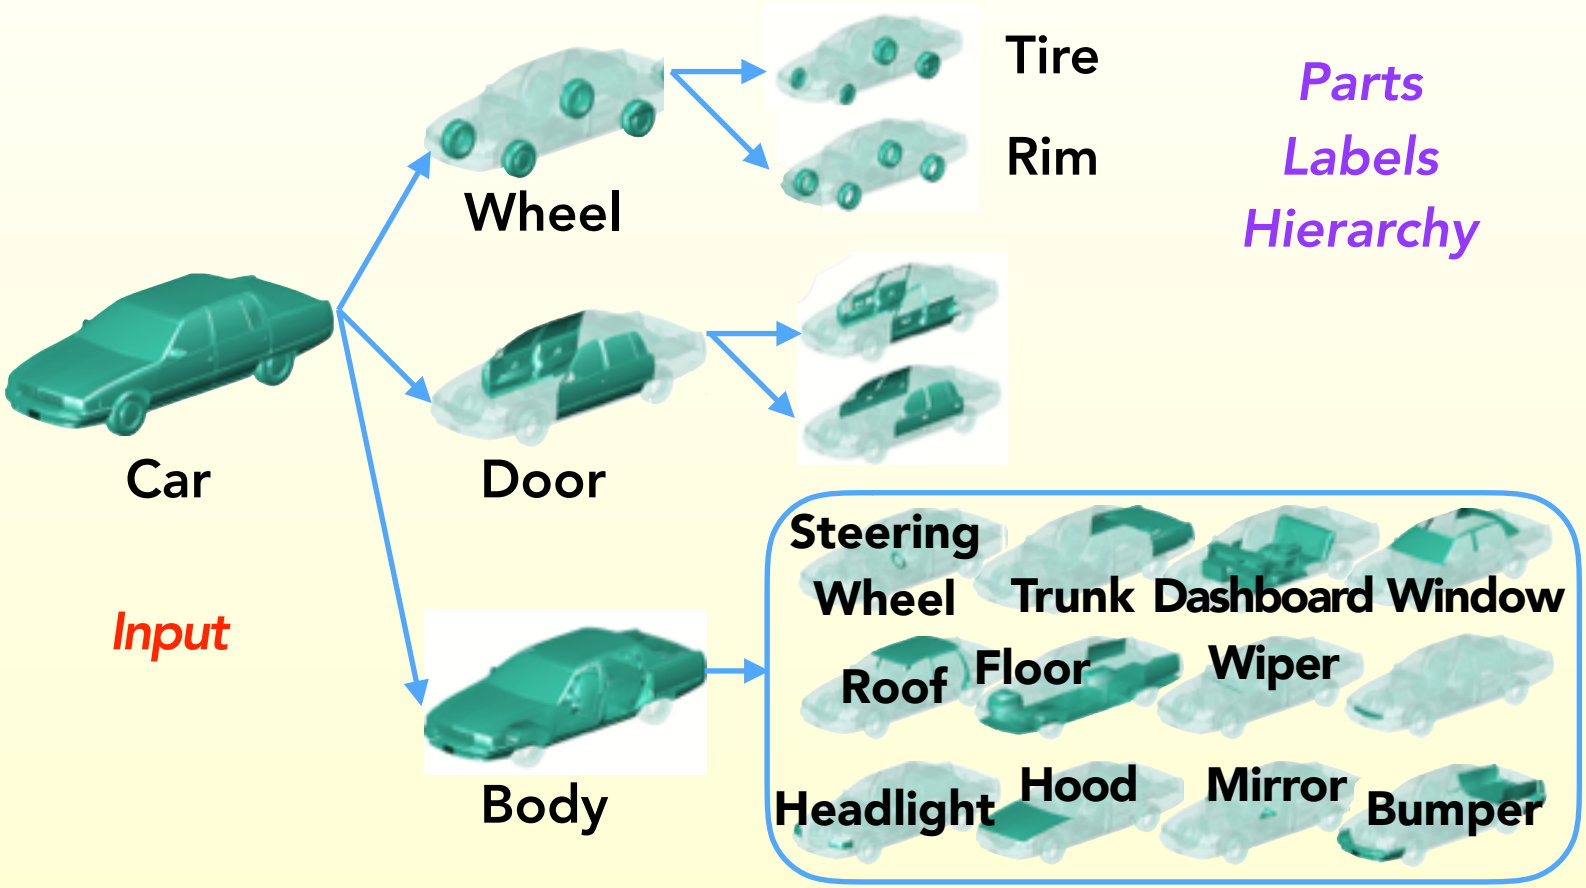

Parts
Labels
Hierarchy

Problem Definition

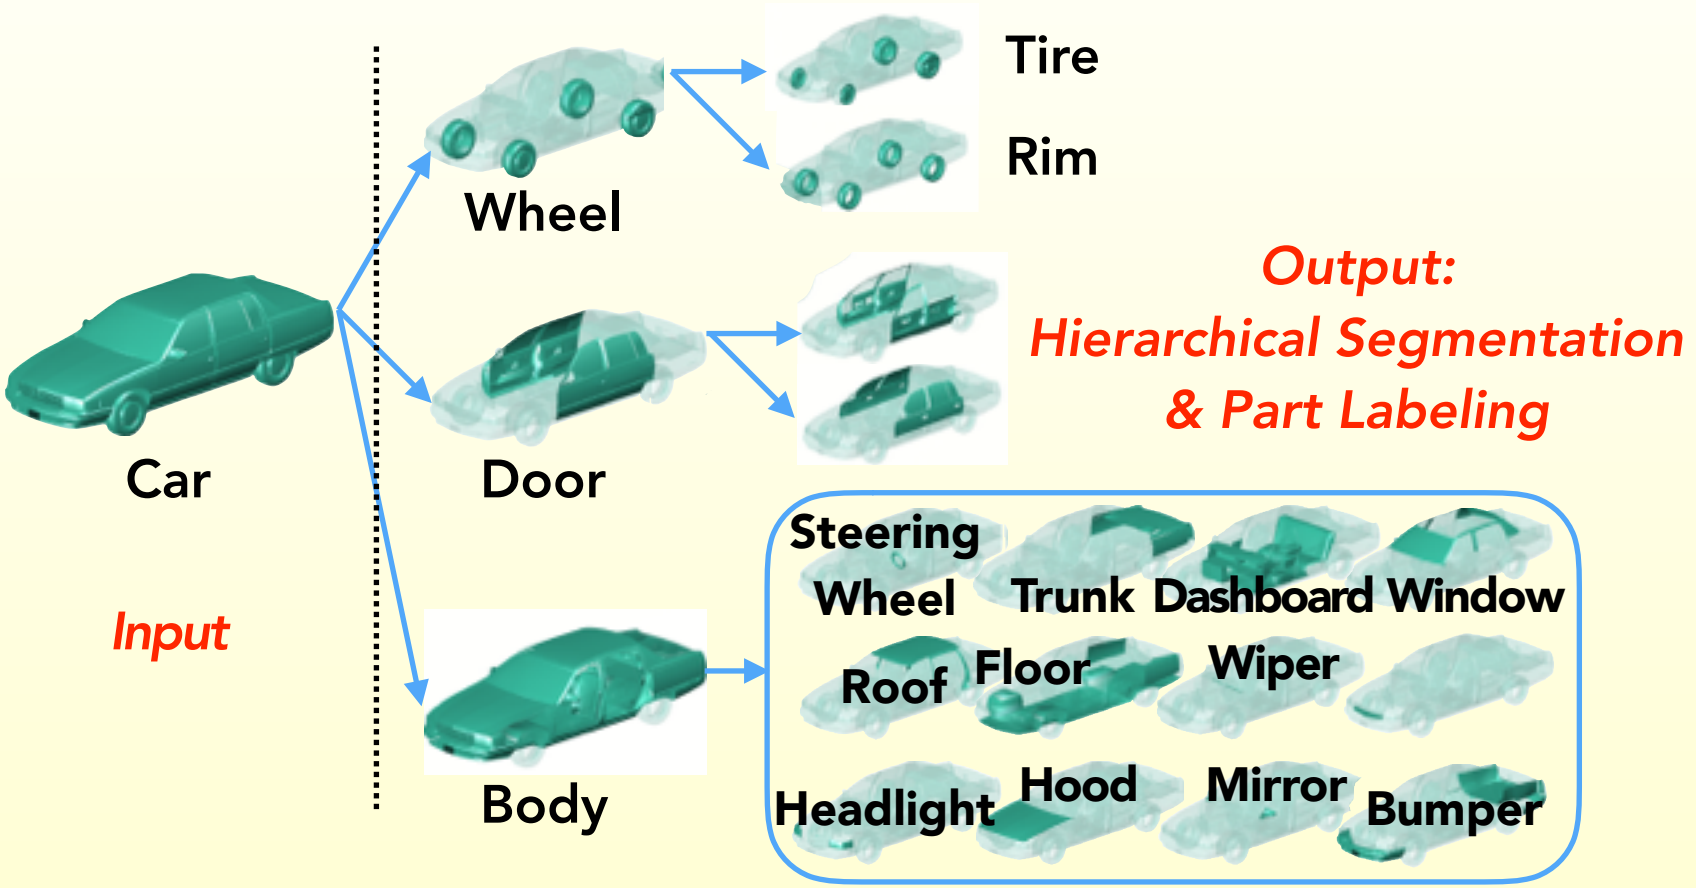

Output:
Hierarchical Segmentation & Part Labeling

How to Define Parts?

- Previously expert defined, flat and coarse

- Knowledge could emerge from horizontal collection wise analysis

Distill Knowledge from Object Graphs

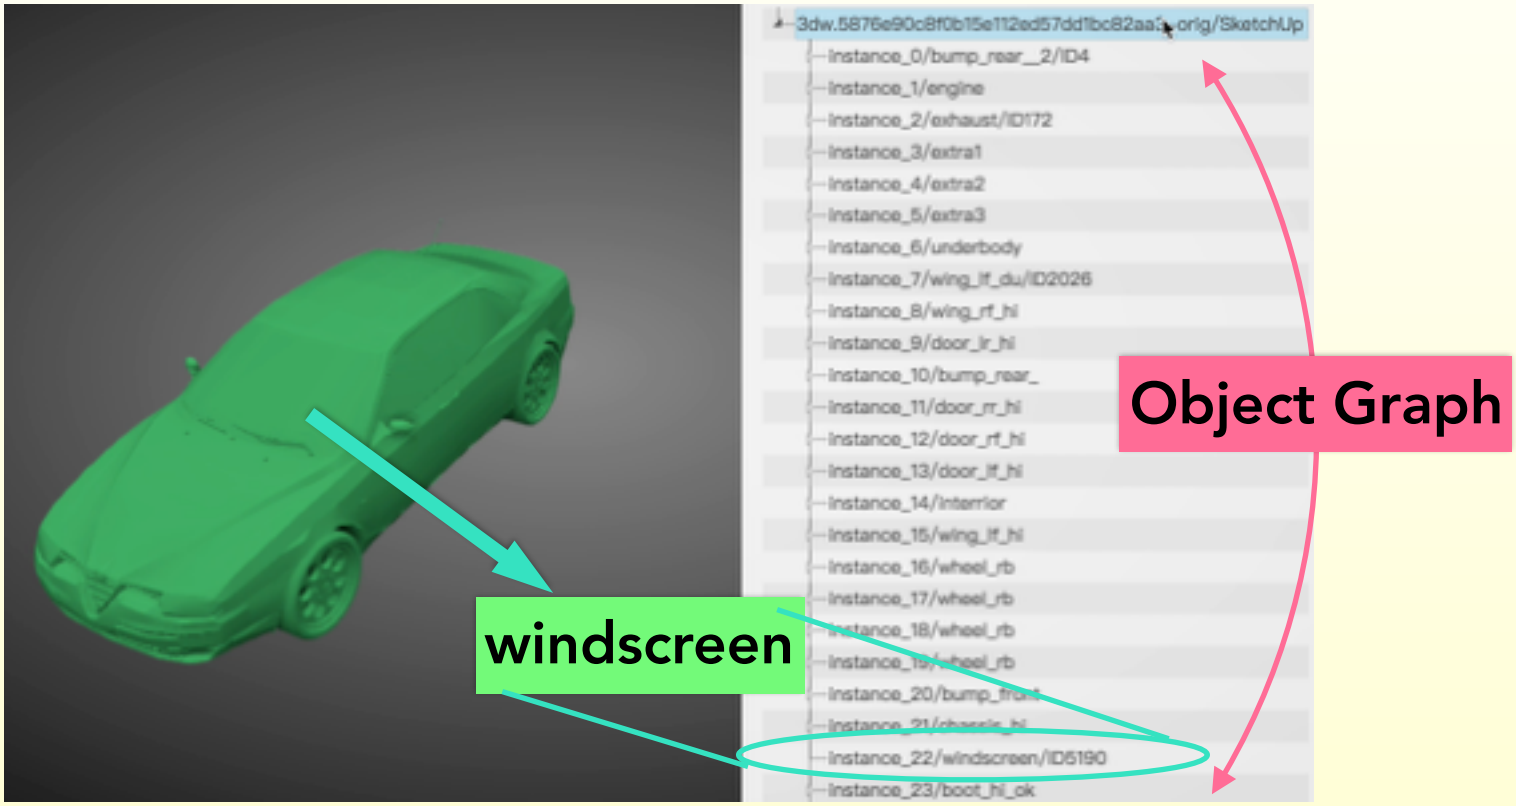

Object Graph

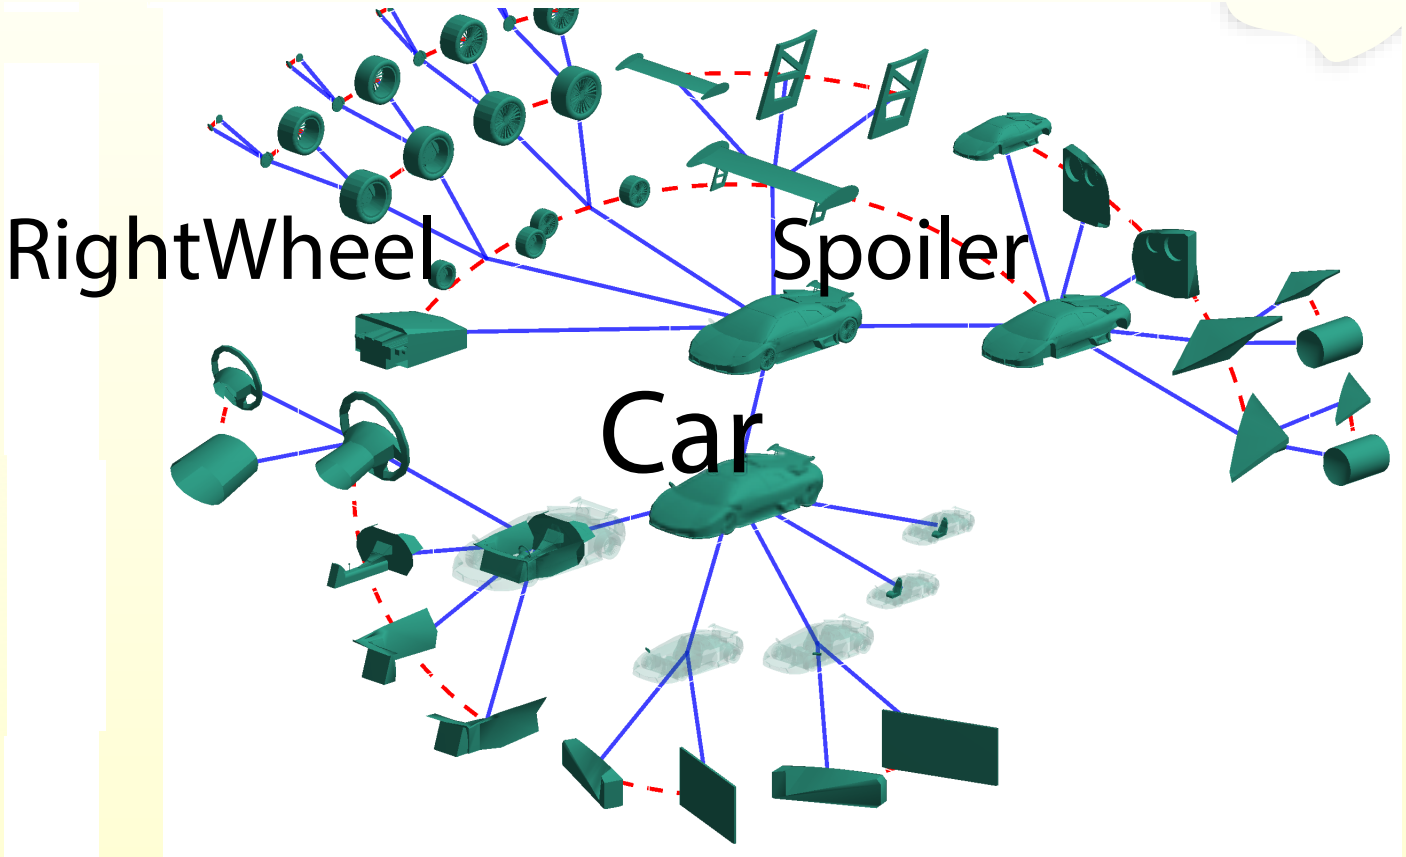

Object Graph

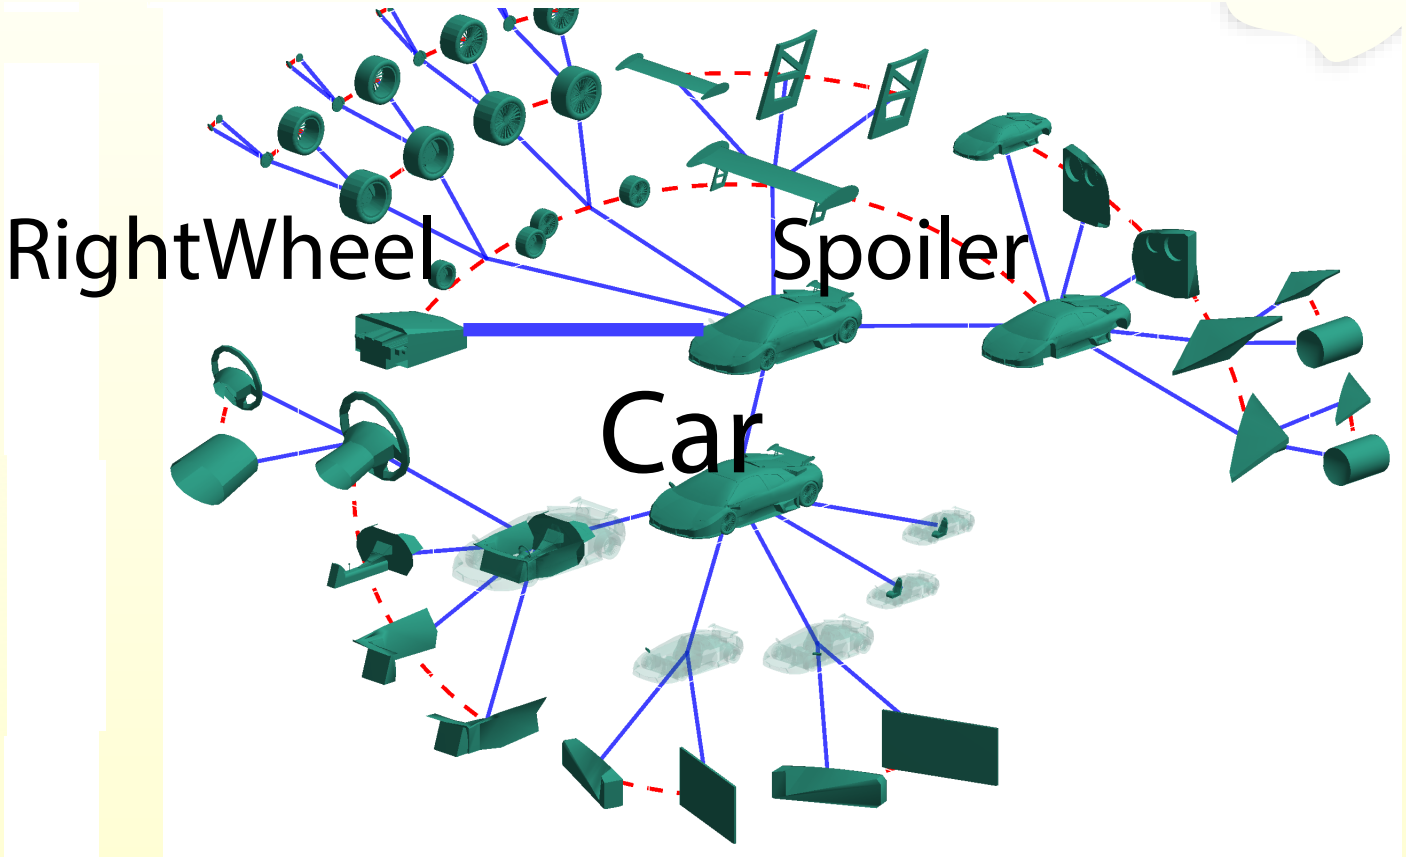

Object Graph

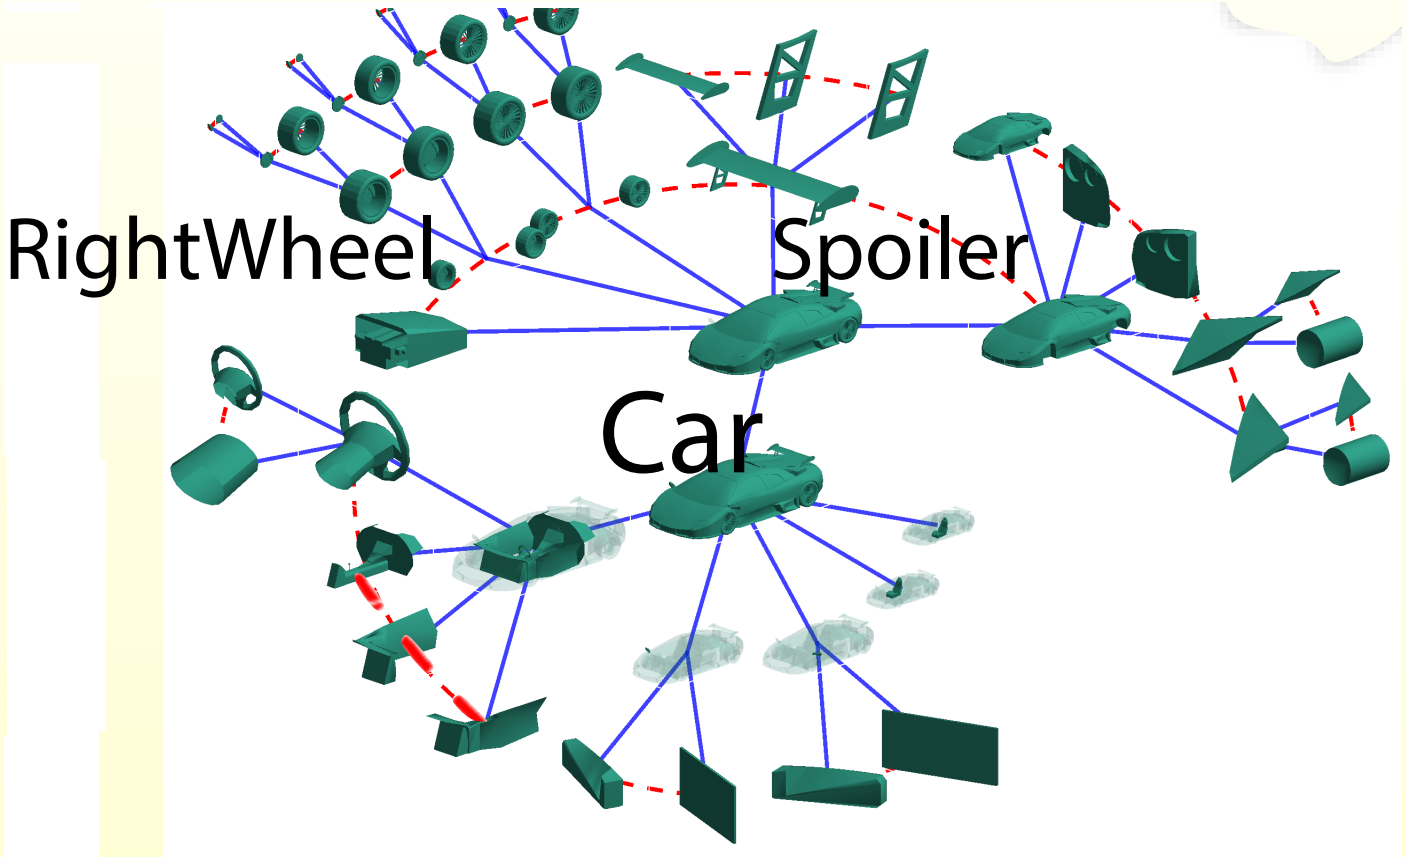

Object Graph

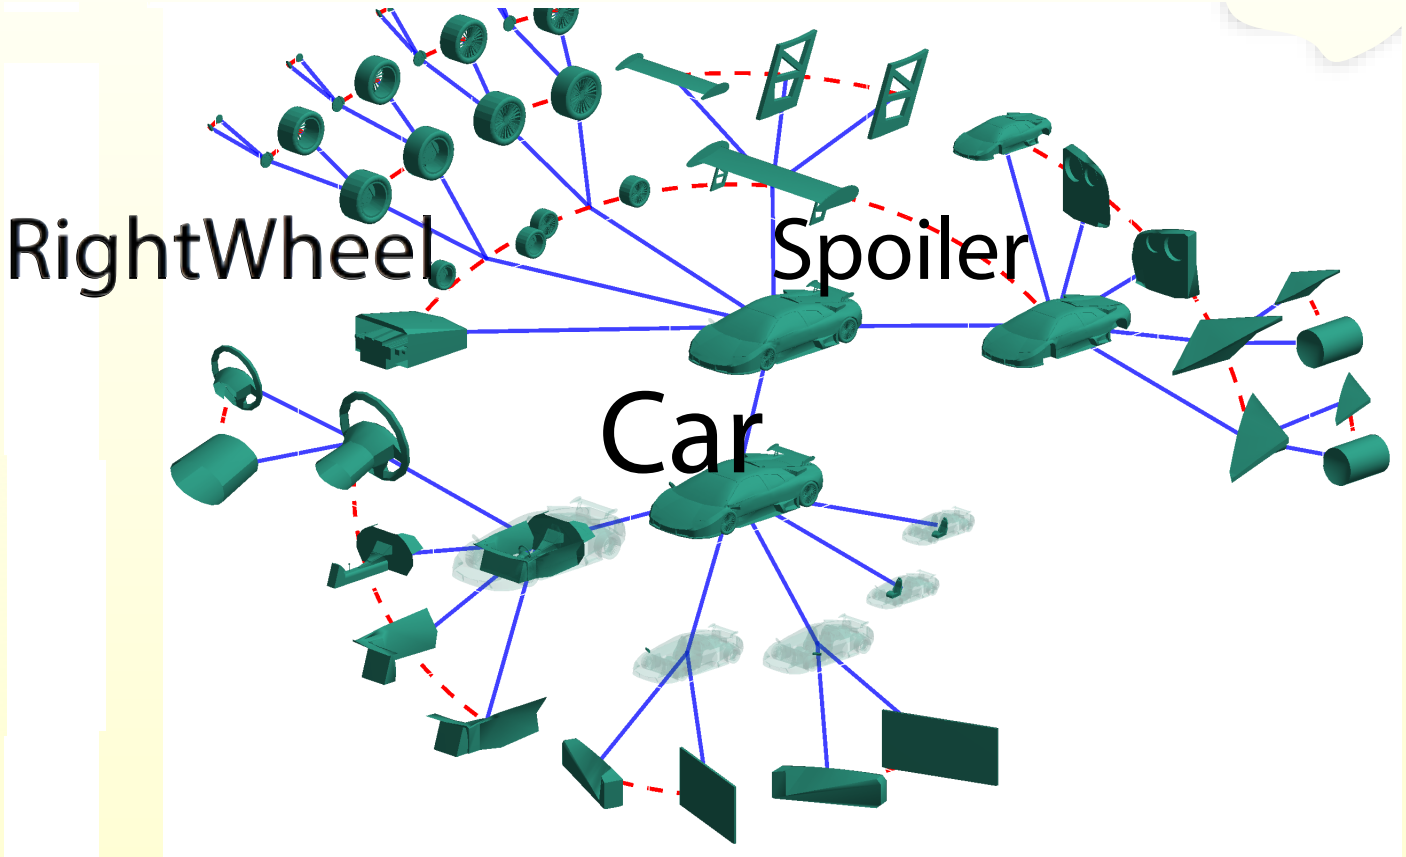

Common Structures in Object Graphs

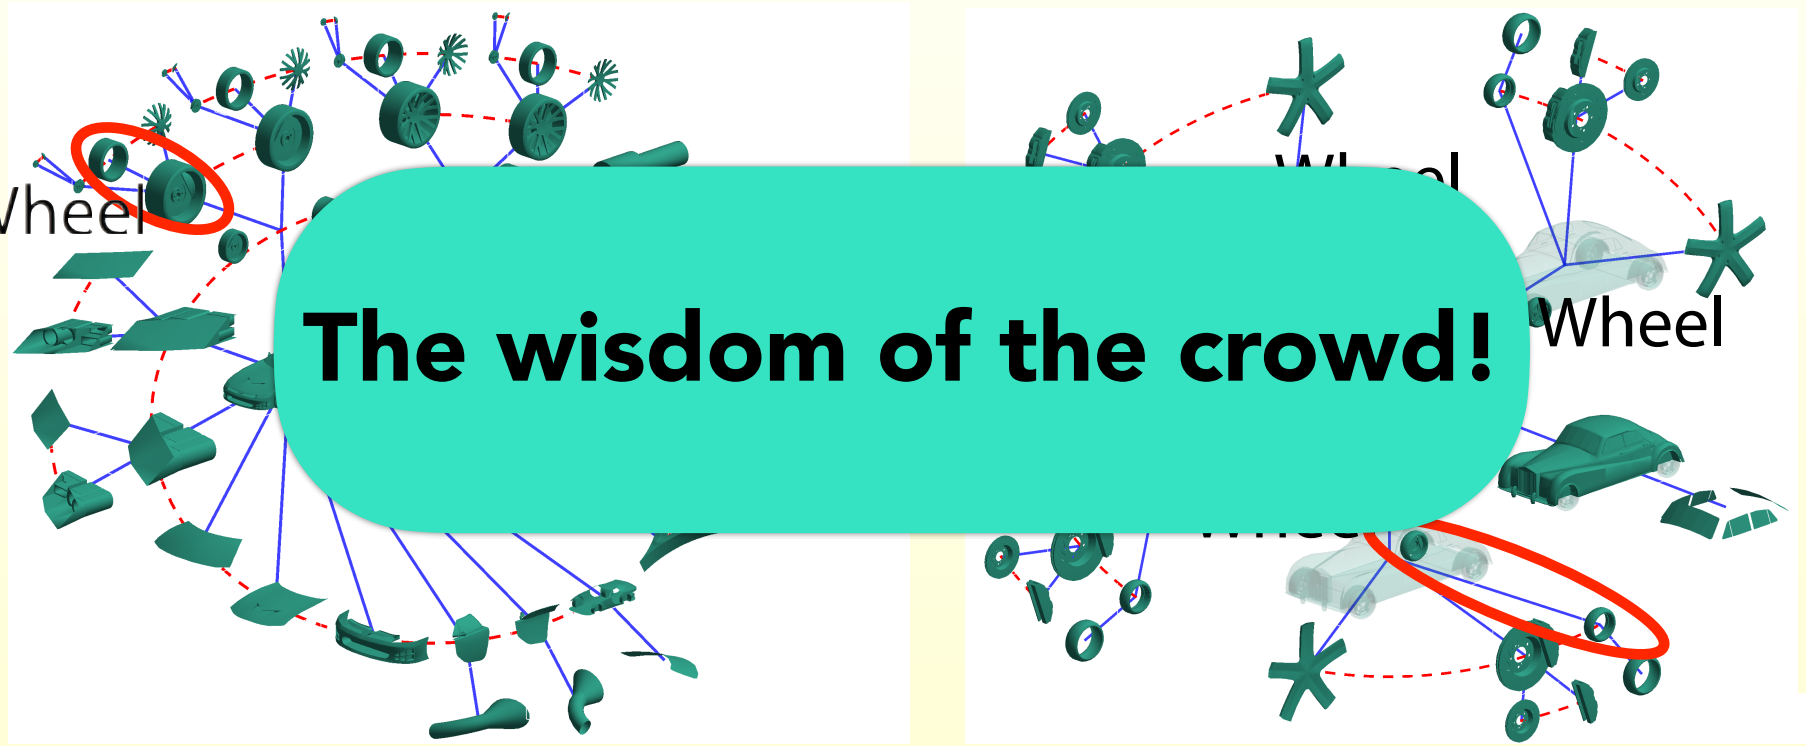

Challenges - Heterogeneous Data

Different Parts

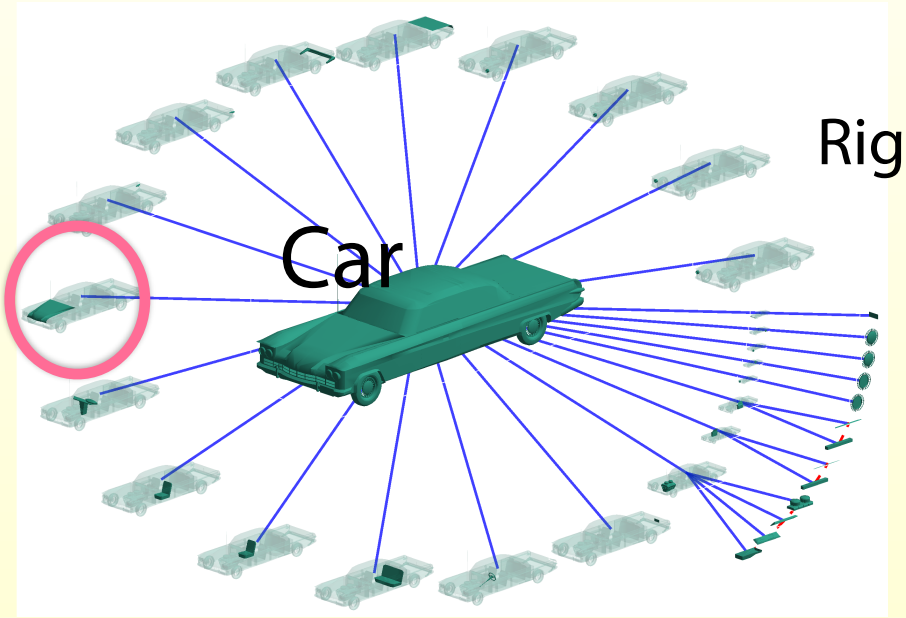

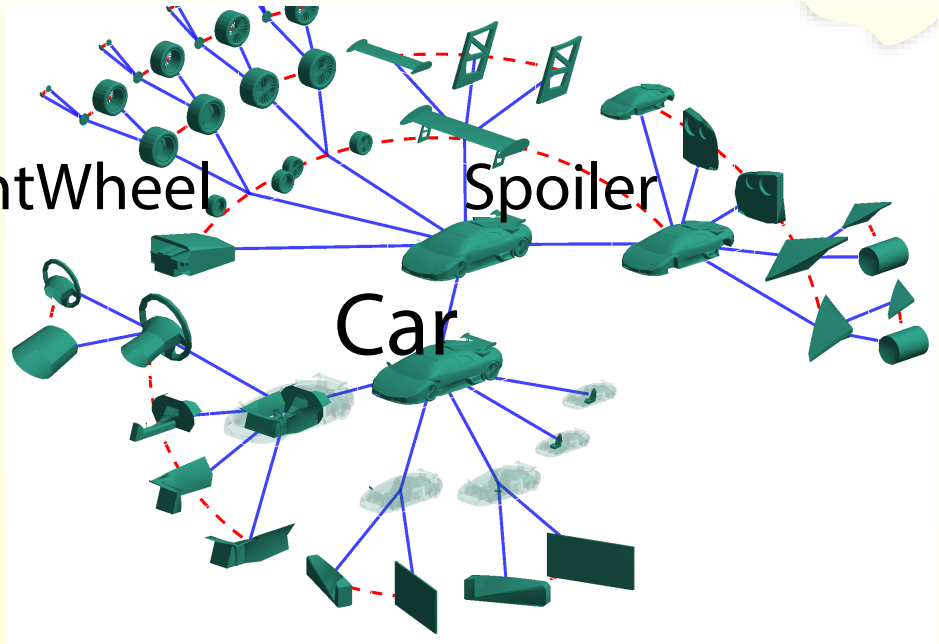

Challenges - Heterogeneous Data

Different Parts, Different Hierarchy

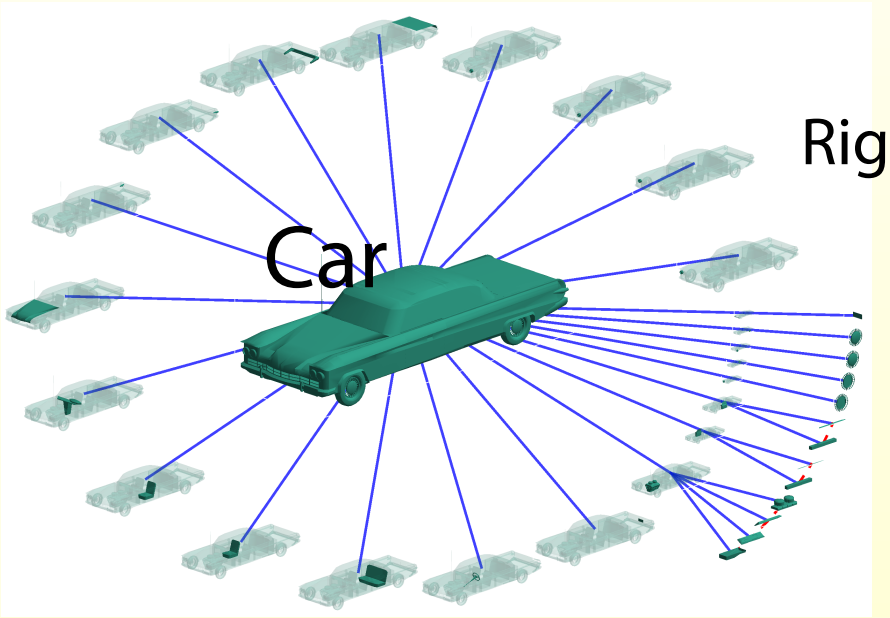

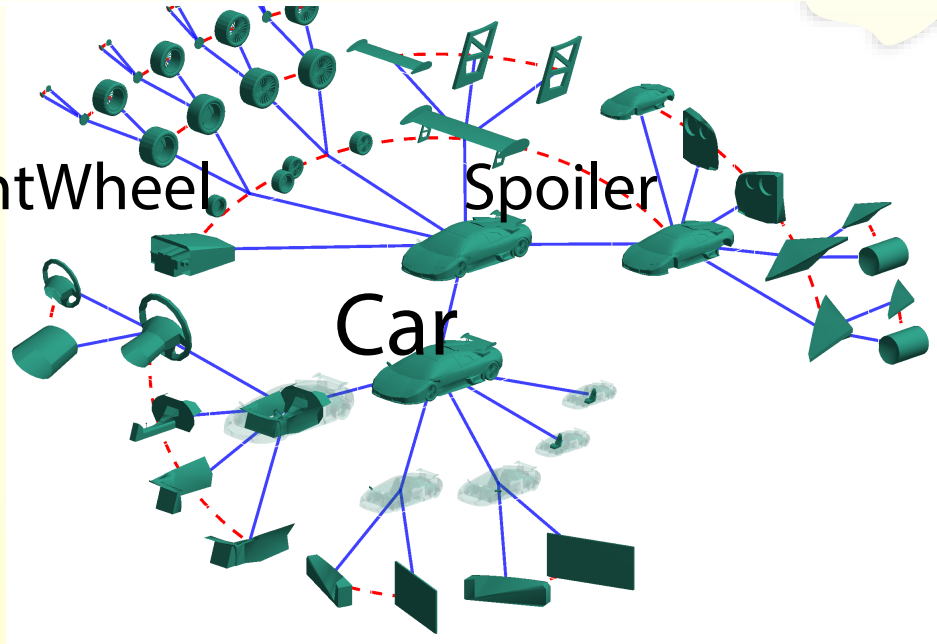

Challenges - Heterogeneous Data

Different Parts, Different Hierarchy, Sparse Tagging

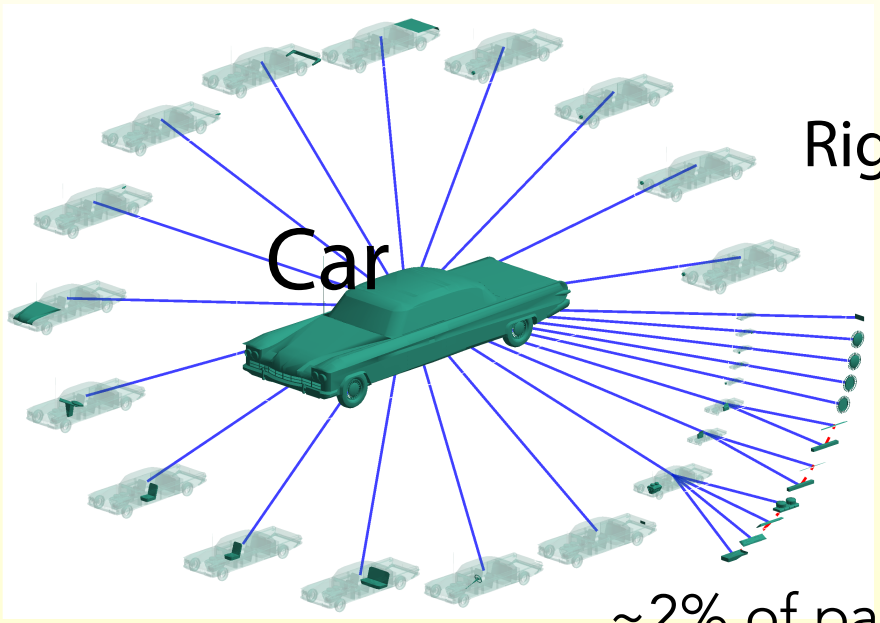

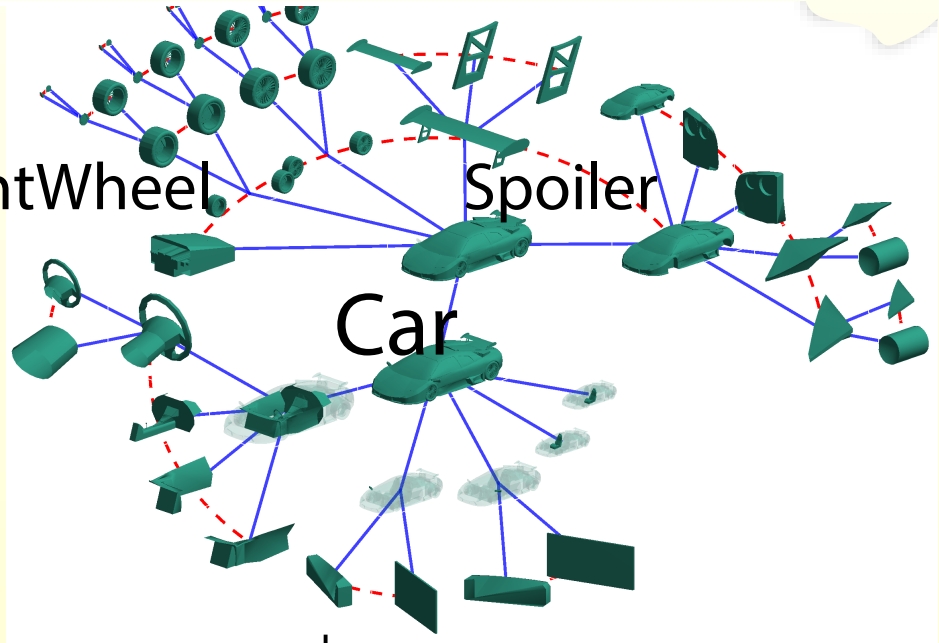

~2% of parts are tagged

Approach Overview

Training Stage

Inference Stage

Approach Overview

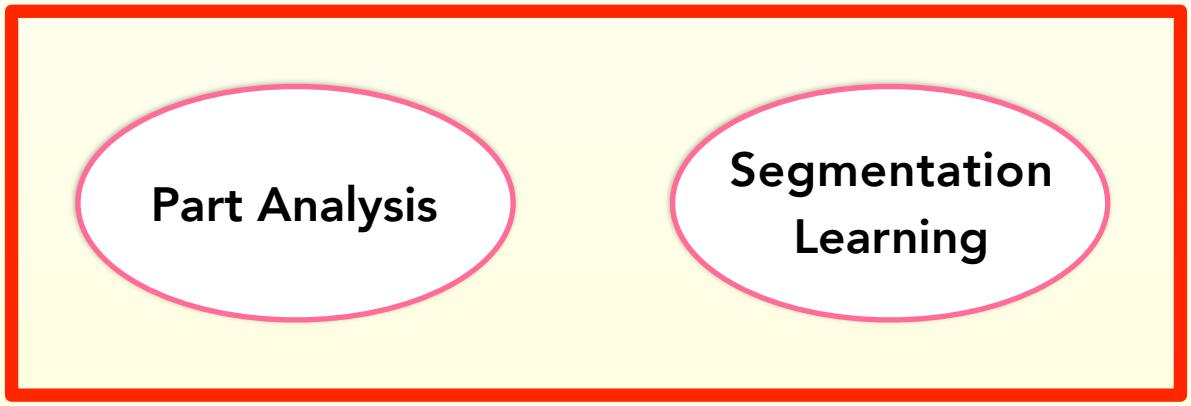

Training Stage

Approach Overview

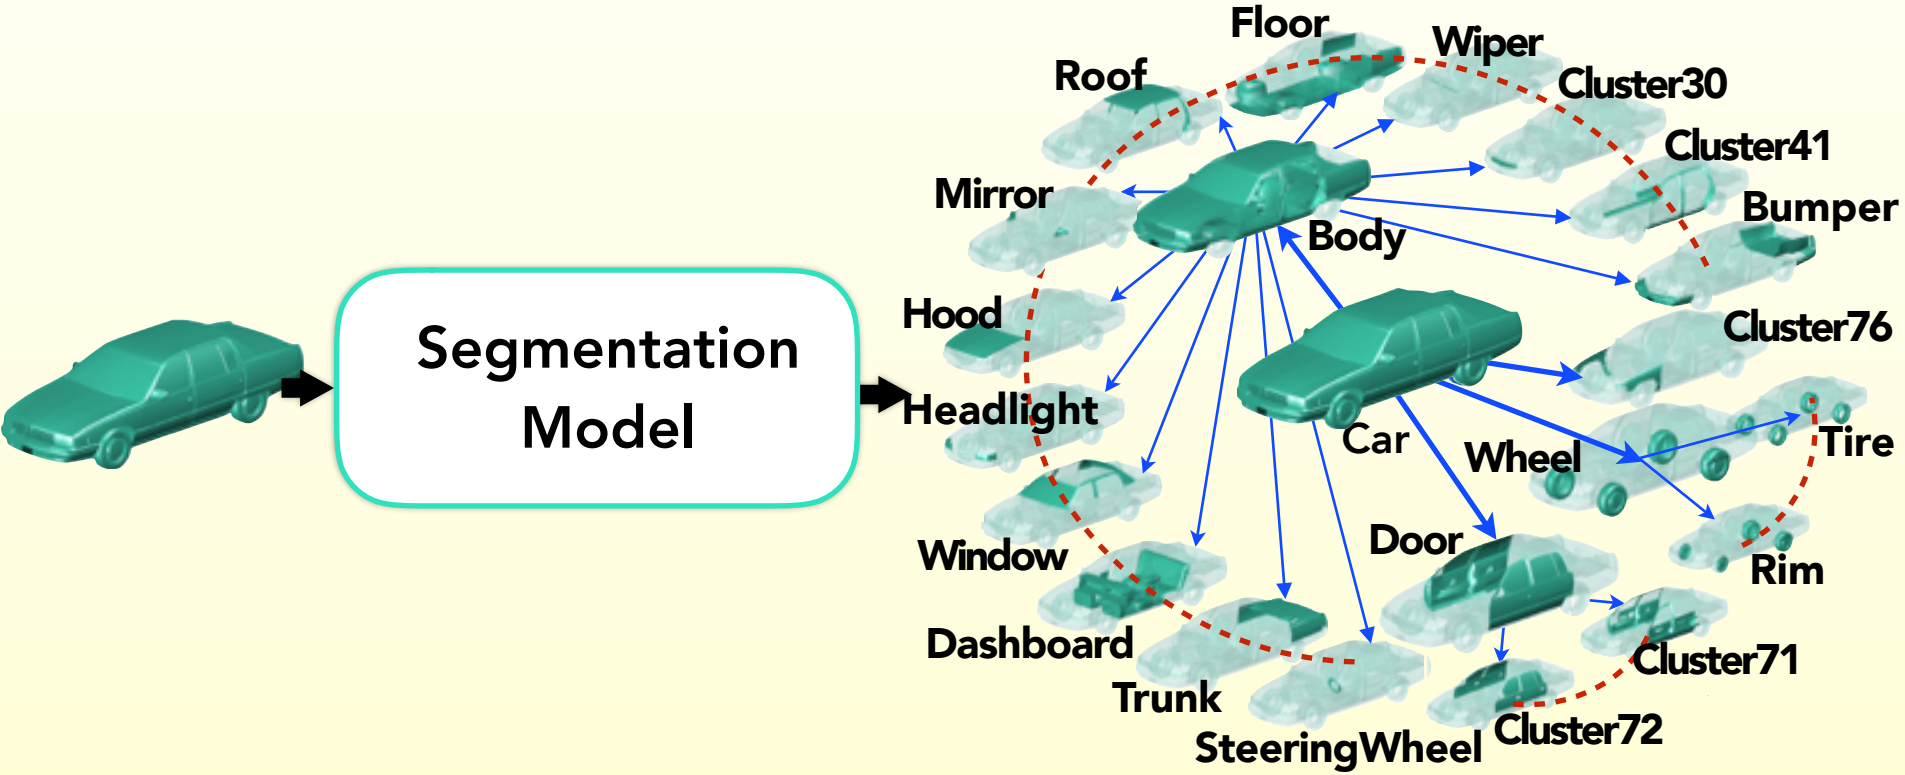

Inference Stage

Approach Overview

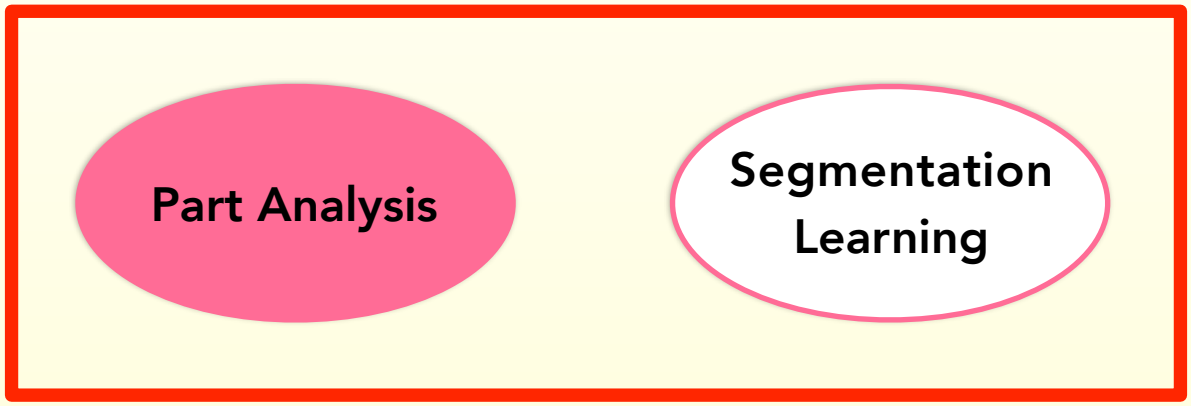

Training Stage

Part Analysis

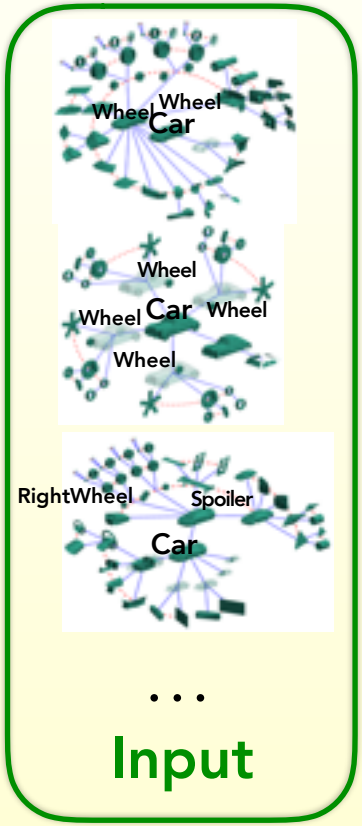

Part Analysis

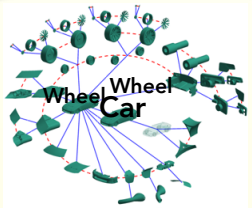

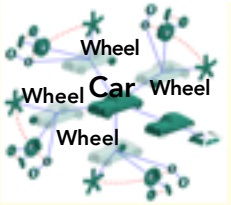

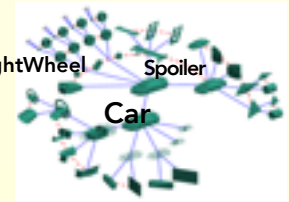

. . .

Input

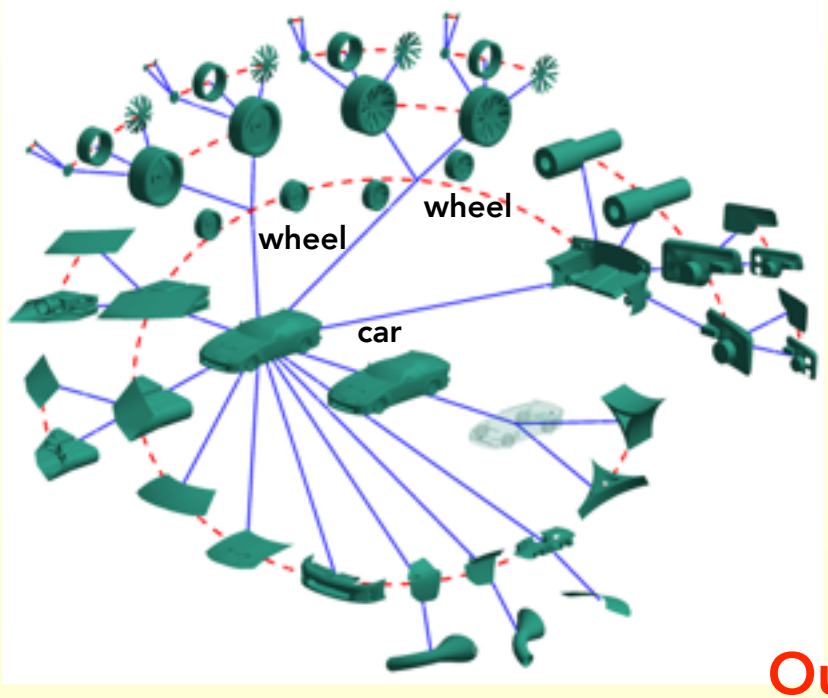

Output

Part Analysis

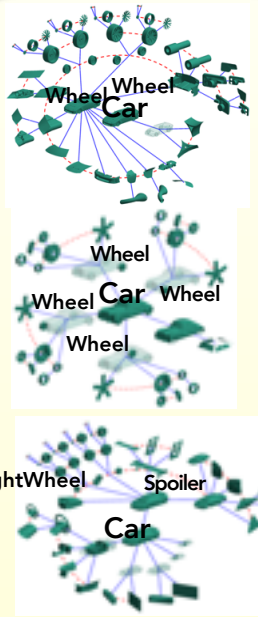

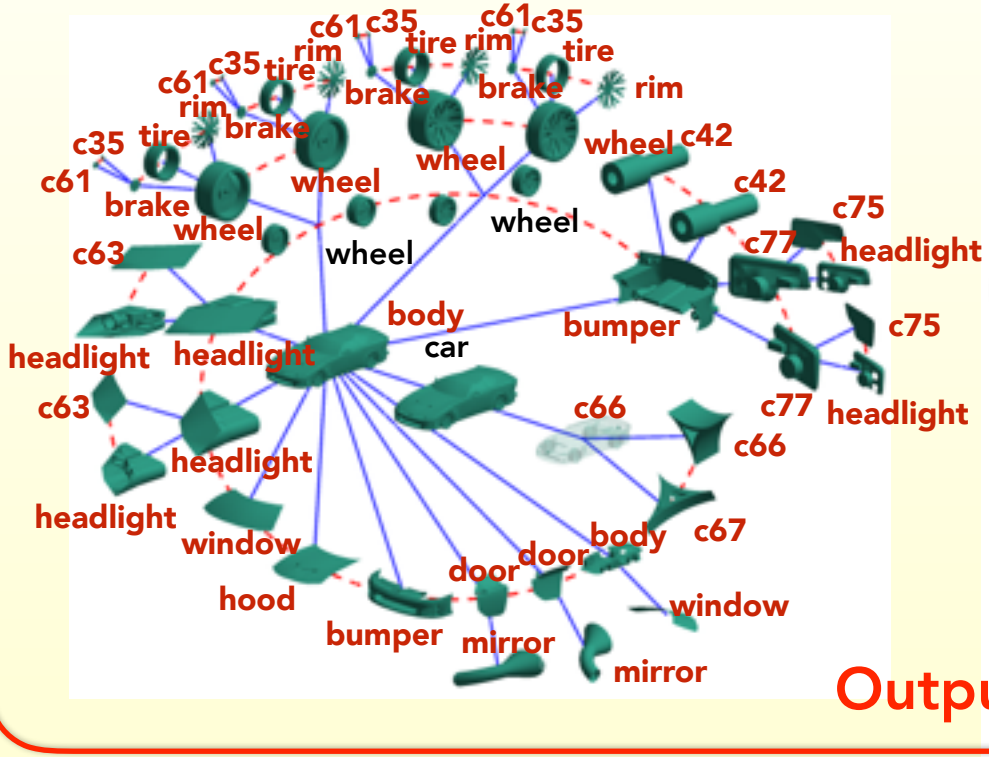

Label Propagation through Part Correspondences

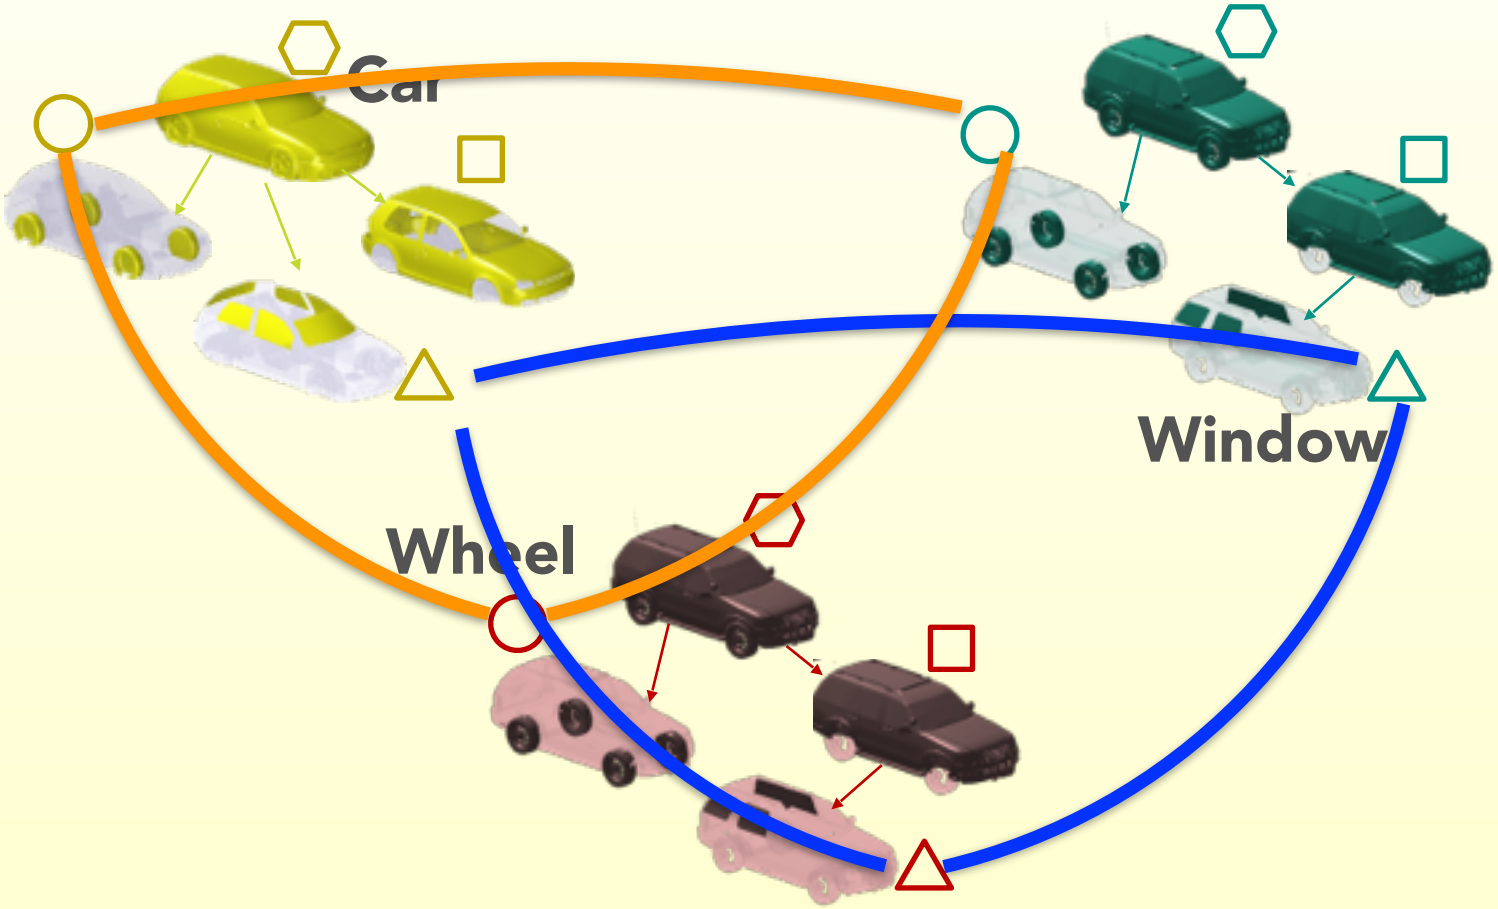

Key Idea — Semi-Supervised Clustering

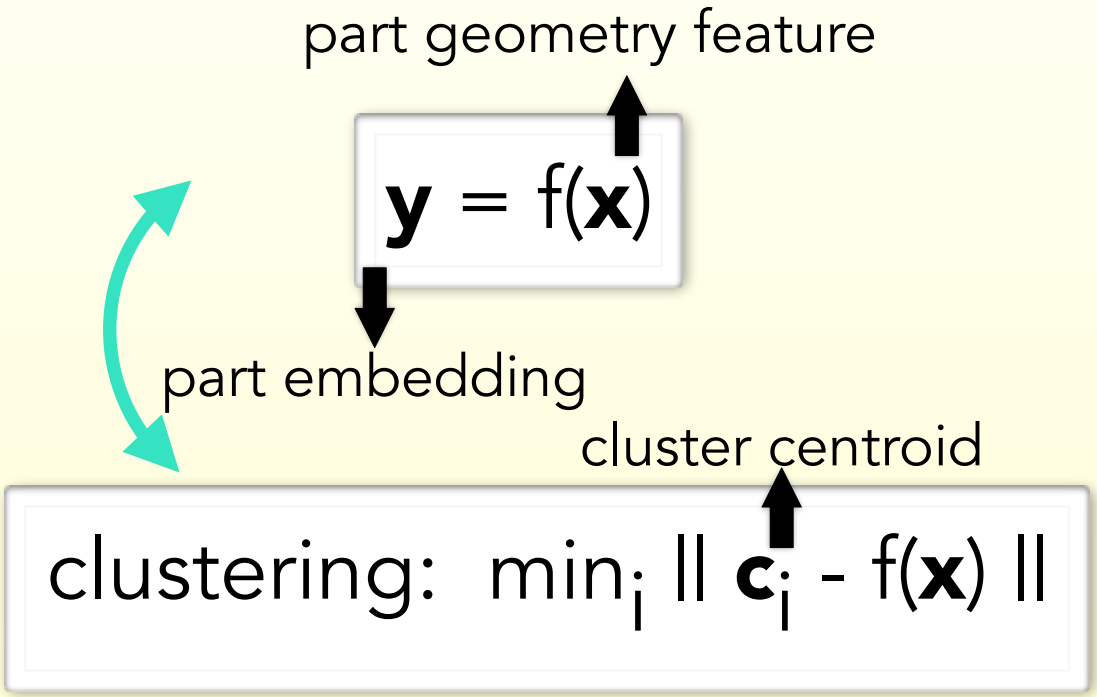

Key Idea — Semi-Supervised Clustering

part geometry feature

$\mathbf{y} = f(\mathbf{x})$

part embedding

cluster centroid

clustering: $\min_i \| \mathbf{c}_i - f(\mathbf{x}) \|$

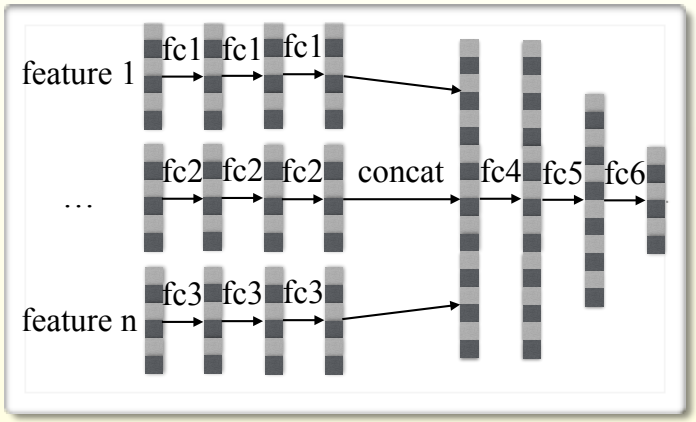

Supervision: sparse tags, inconsistent hierarchies

Key idea — Semi-Supervised Clustering

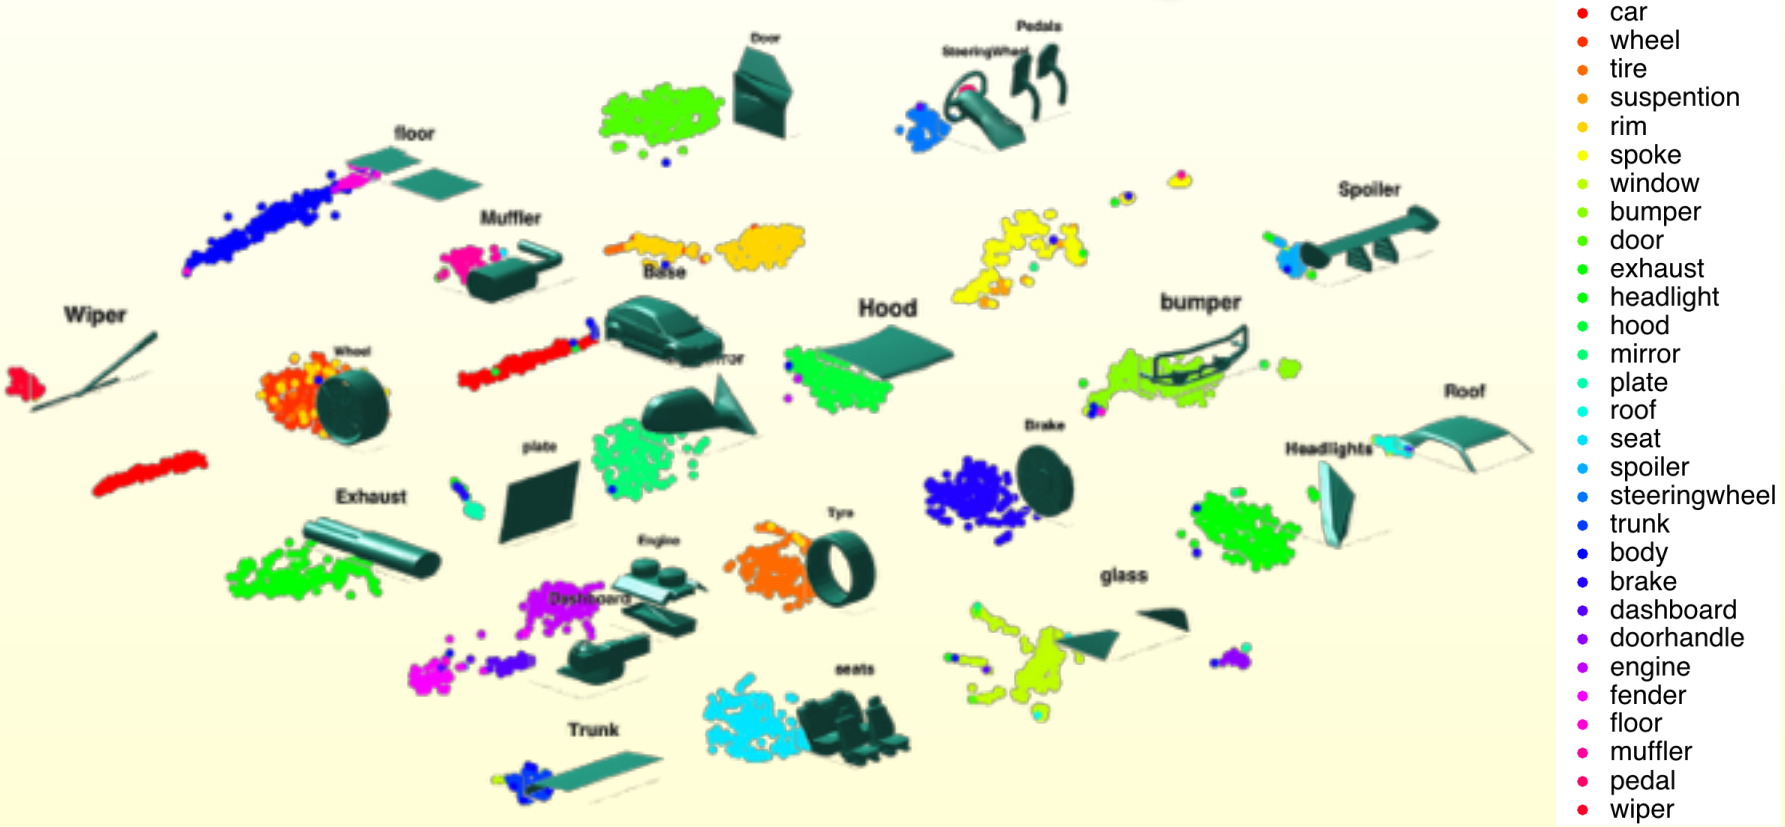

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

c.f. Basu et al. KDD 2004

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

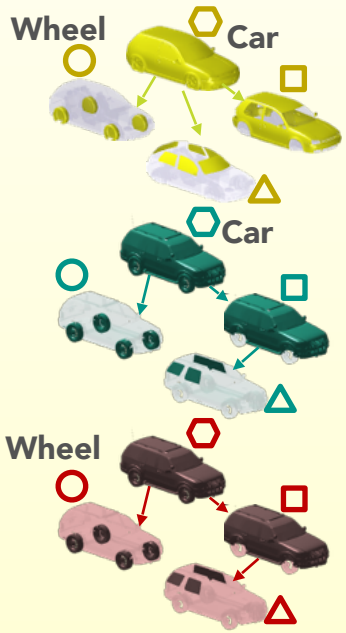

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

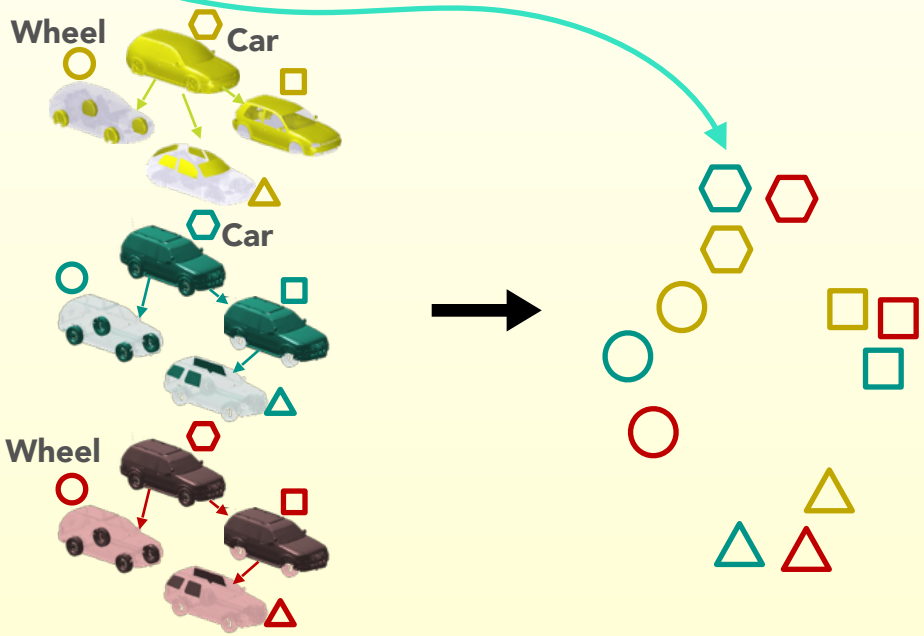

Embedding parameter

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

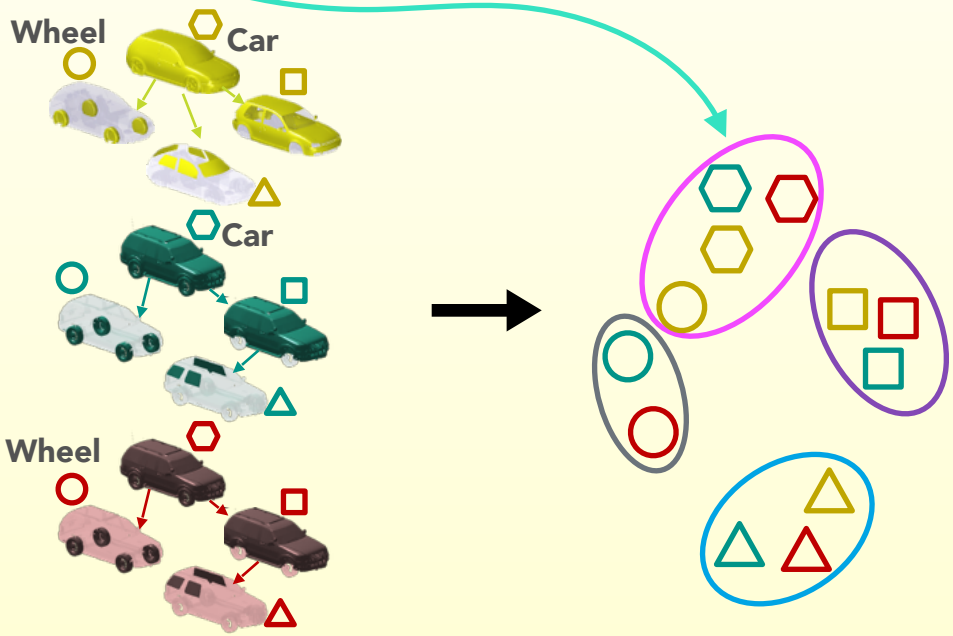

Clustering labels

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

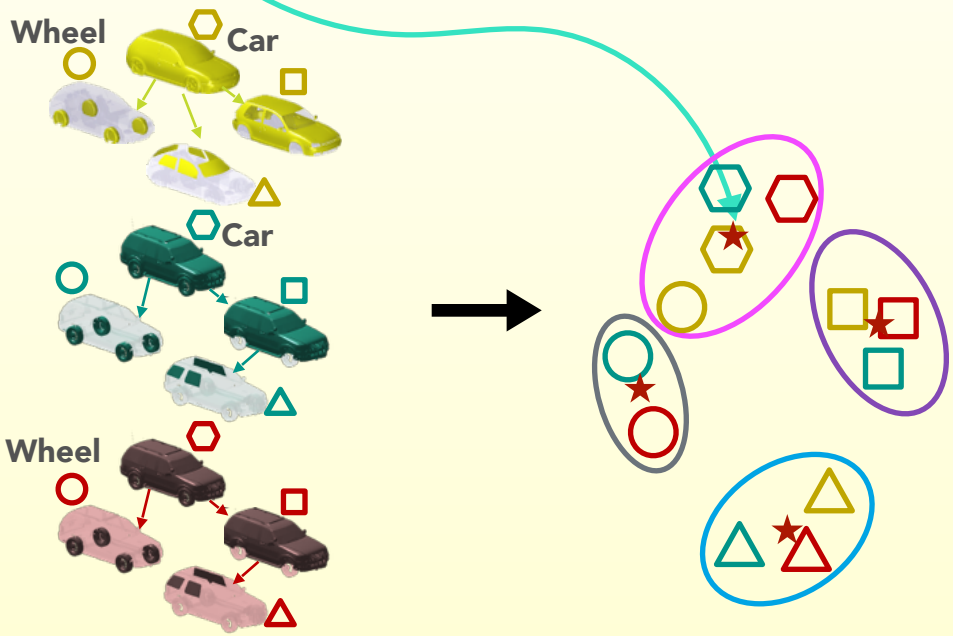

Clustering centroids

Objective Function

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

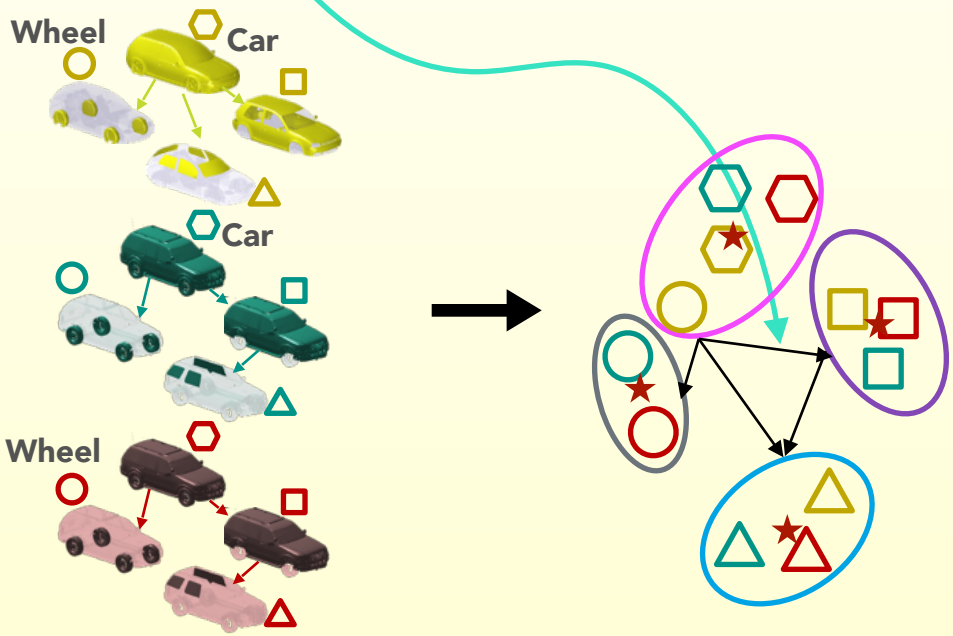

Soft hierarchy graph

Objective Function

Clustering

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

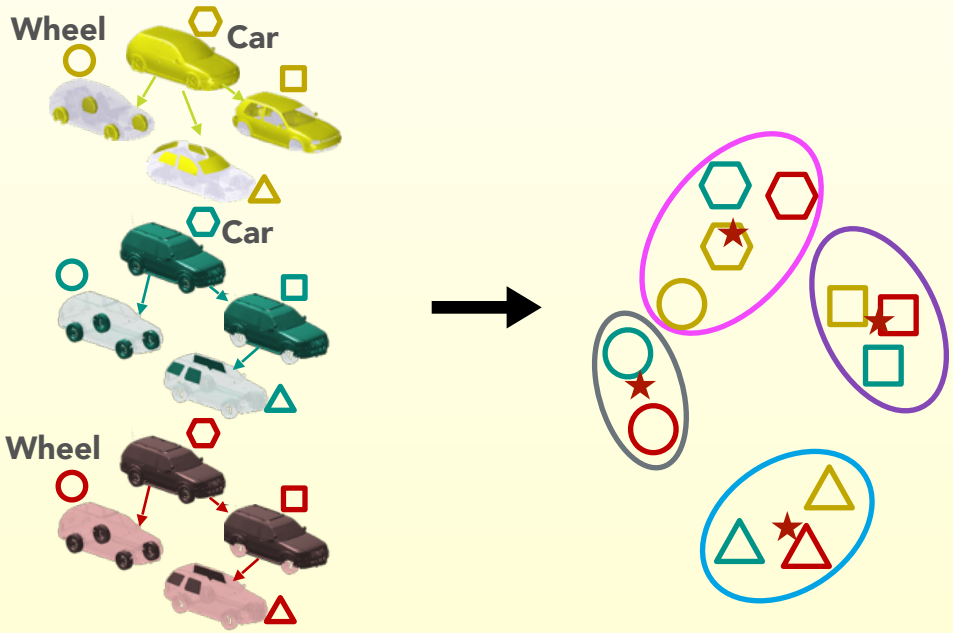

Clustering term:
Encourage parts to form clusters

Objective Function

Clustering

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

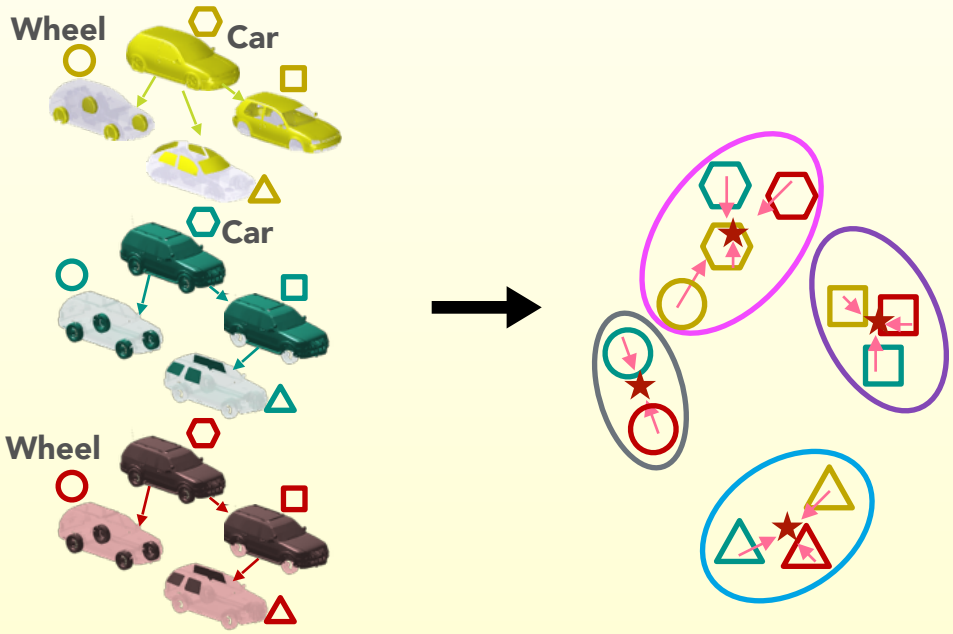

Clustering term:
Encourage parts to form clusters

Objective Function

Clustering

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

$$E_c = \sum_{\text{parts}} \sum_{\text{centroids}} p_{\text{part,centroid}} ||f(\text{part}) - \text{centroid}||$$

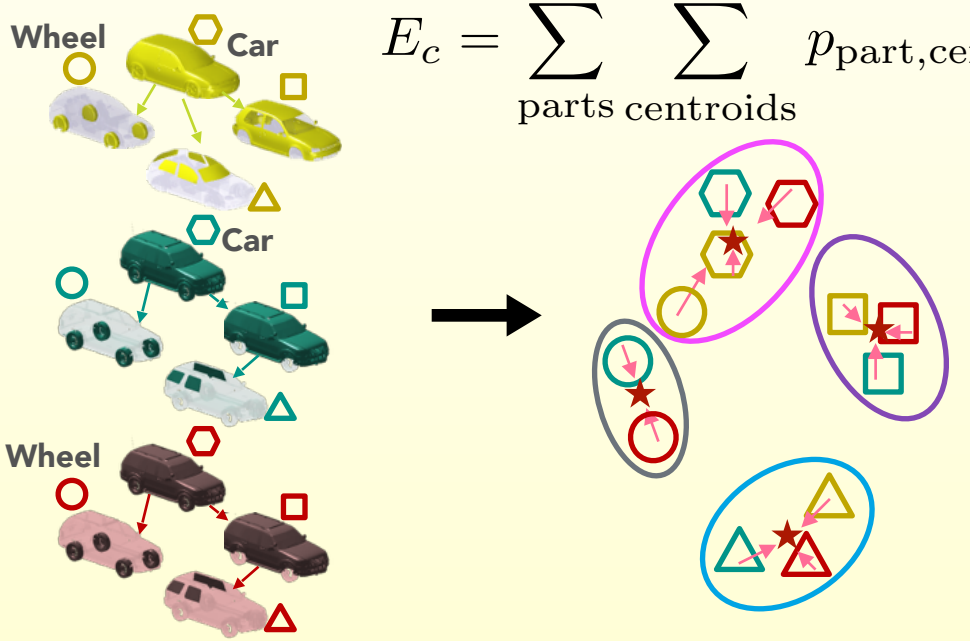

Clustering term:

Encourage parts to form clusters

Objective Function

Clustering Similarity

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

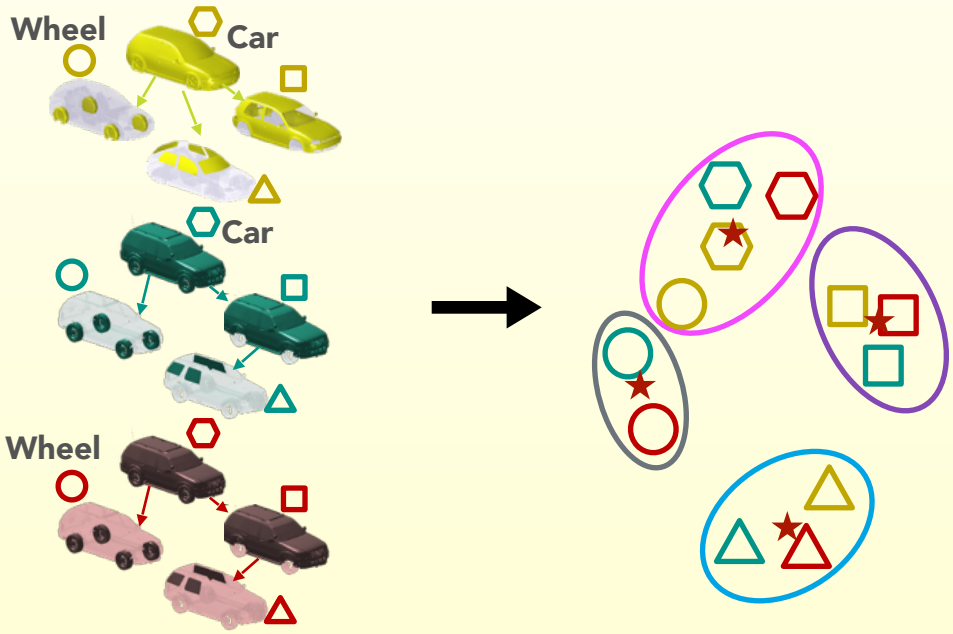

Similarity term:

Group parts with the same tag or geometry

Objective Function

Clustering Similarity

$$E(\theta, p, c, M) = \lambda_c E_c + \boxed{\lambda_s E_s} + \lambda_d E_d + \lambda_m E_m - H$$

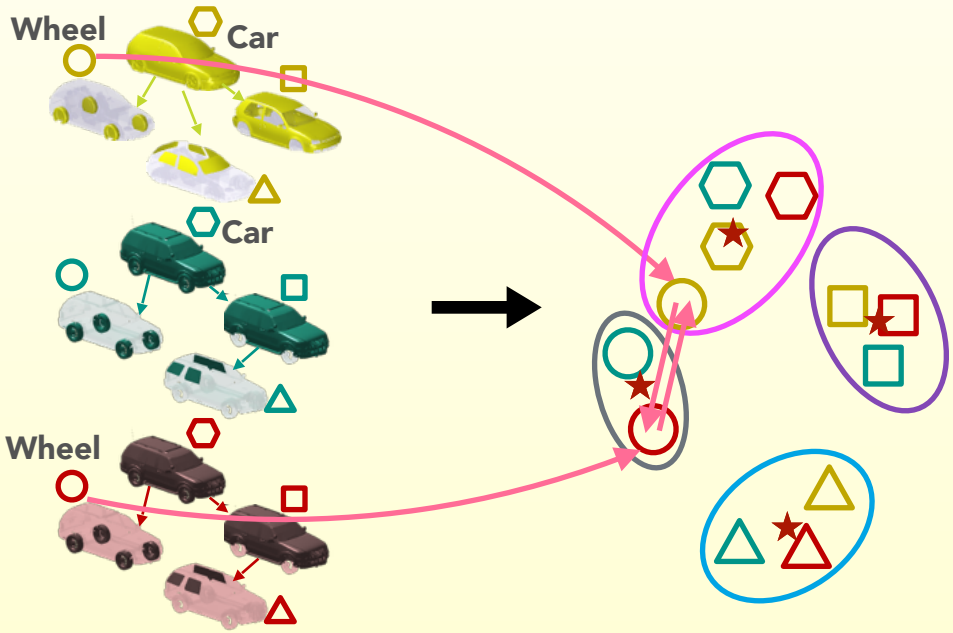

Similarity term:

Group parts with the same tag or geometry

Objective Function

Clustering Similarity

$$E(\theta, p, c, M) = \lambda_c E_c + \boxed{\lambda_s E_s} + \lambda_d E_d + \lambda_m E_m - H$$

$$E_s = \sum_{\text{part 1}} \sum_{\text{part 2}} ||f(\text{part 1}) - f(\text{part 2})|| \quad \begin{matrix}\text{iff almost identical} \\ \text{or tags are the same}\end{matrix}$$

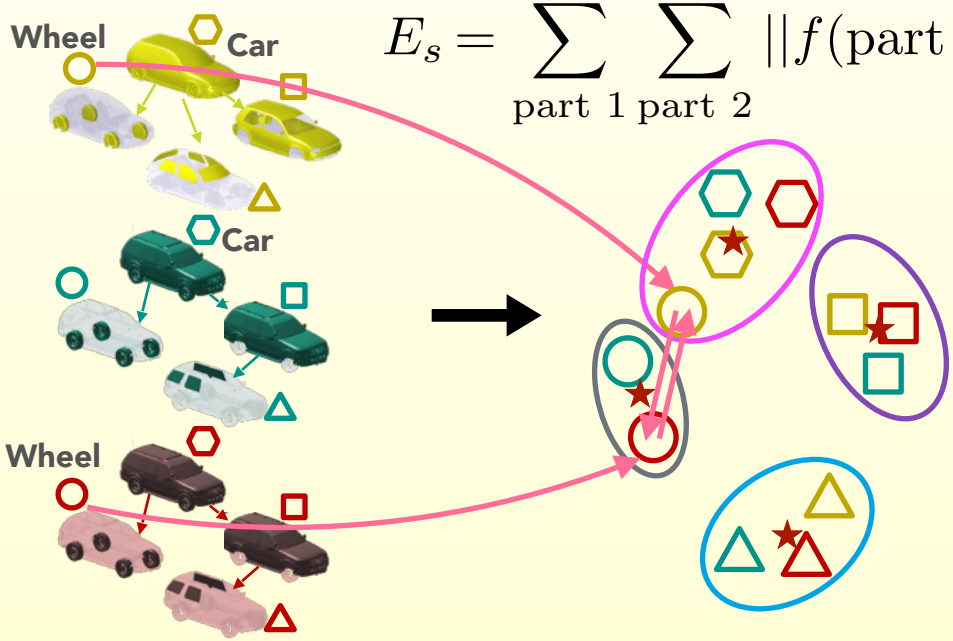

Similarity term:

Group parts with the same tag or geometry

Objective Function

Clustering Similarity Dissimilarity

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

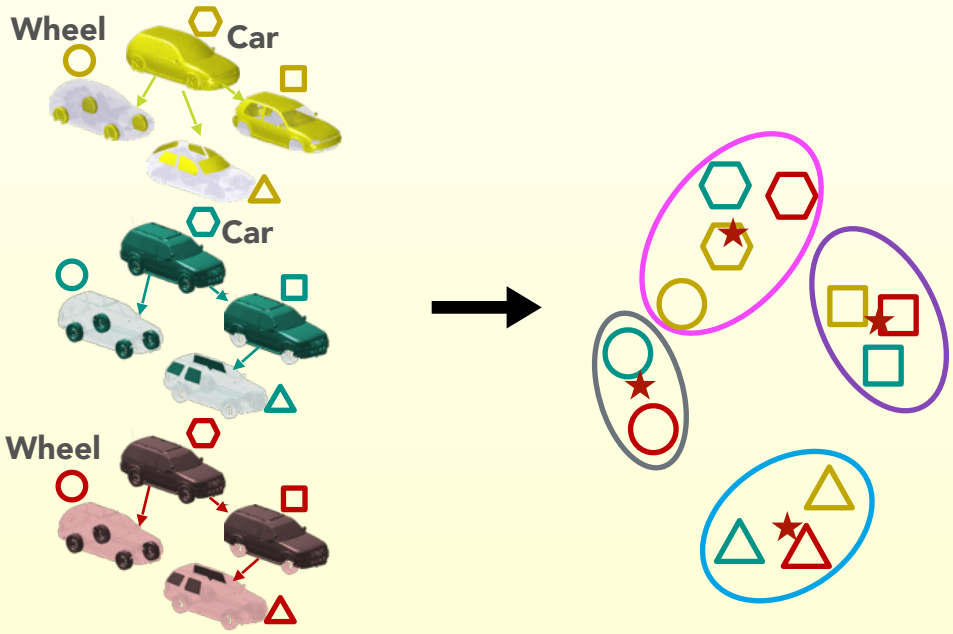

Dissimilarity term:

Separate parts which should not have the same label

Objective Function

Clustering Similarity Dissimilarity

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \boxed{\lambda_d E_d} + \lambda_m E_m - H$$

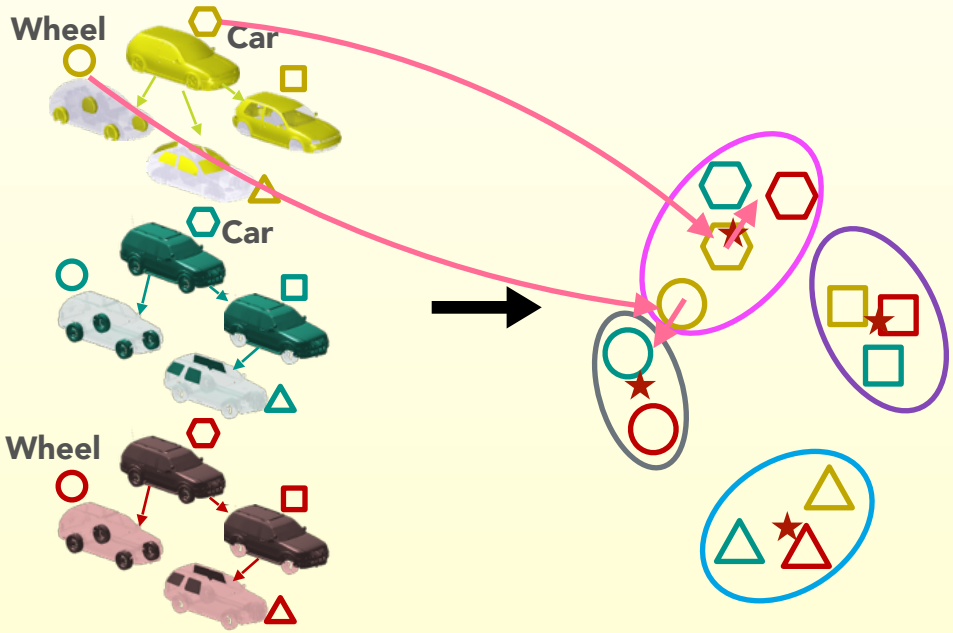

Dissimilarity term:

Separate parts which should not have the same label

Objective Function

Clustering Similarity Dissimilarity Structure

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

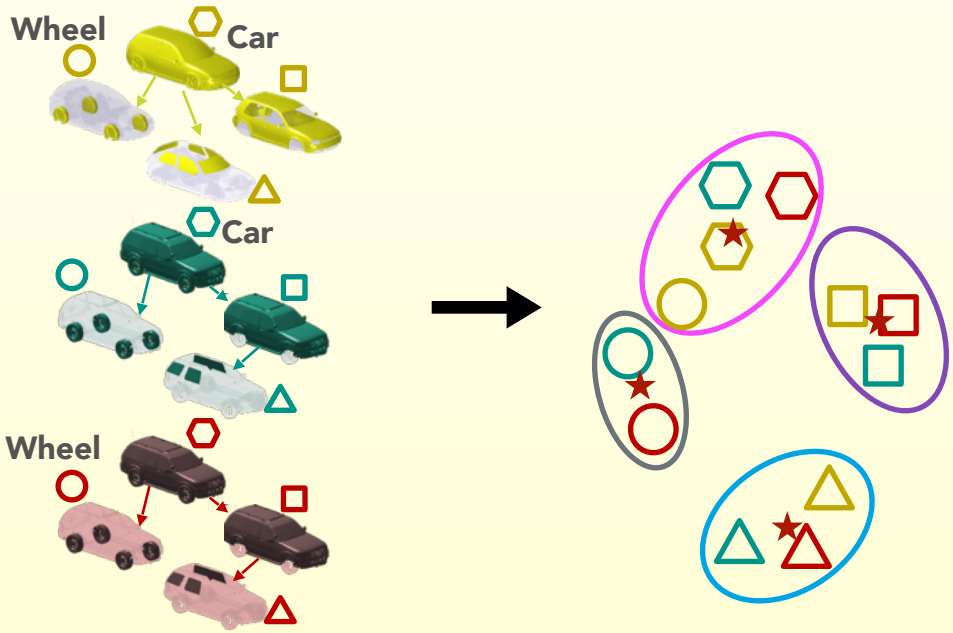

Structure term:

Encourage cluster labels to follow overall parent child relationship

Objective Function

Clustering Similarity Dissimilarity Structure

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

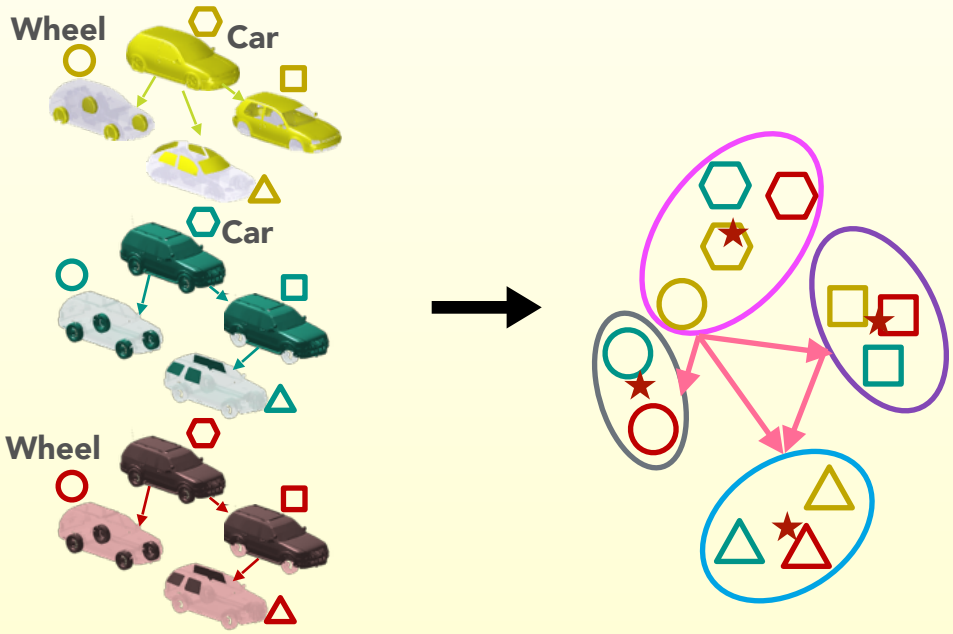

Structure term:

Encourage cluster labels to follow overall parent child relationship

Optimizing the Objective

Clustering Similarity Dissimilarity Structure

$$E(\theta, p, c, M) = \lambda_c E_c + \lambda_s E_s + \lambda_d E_d + \lambda_m E_m - H$$

SEM optimization:

E step: optimize for p

M step: SGD optimize for θ, c, M

Outputs

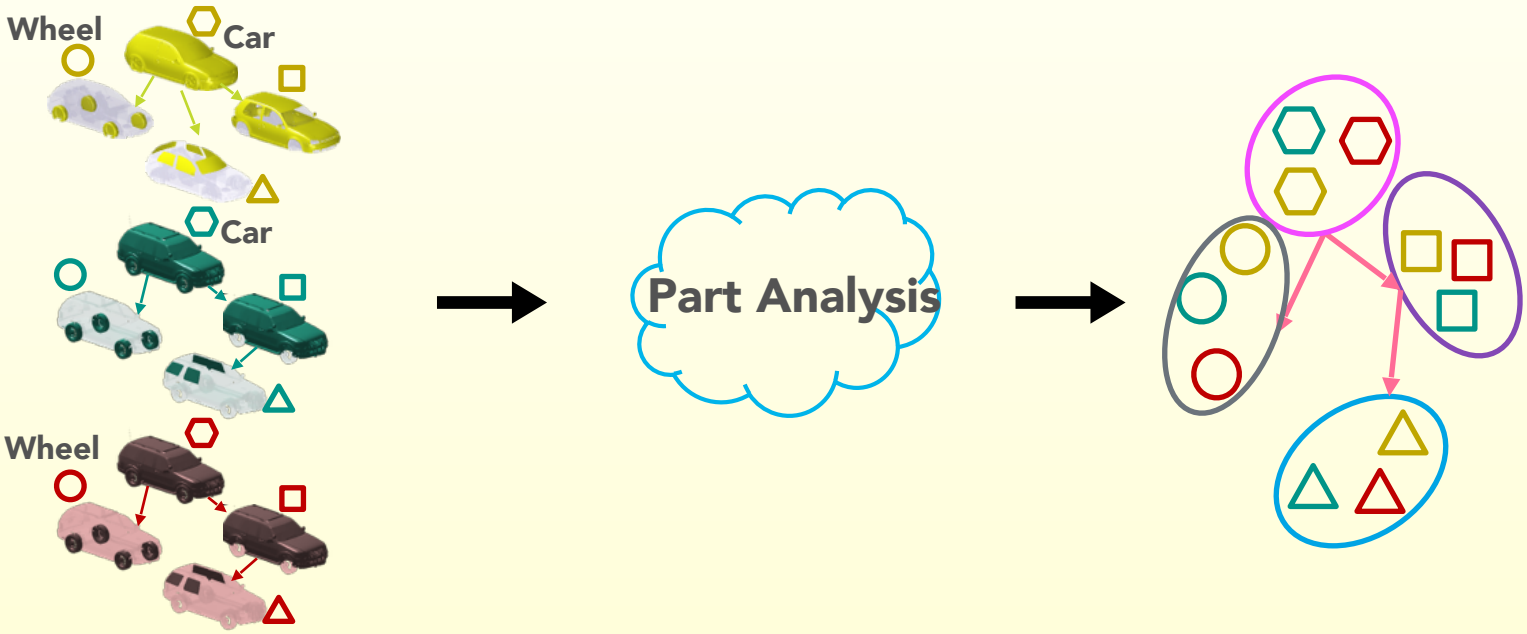

Outputs

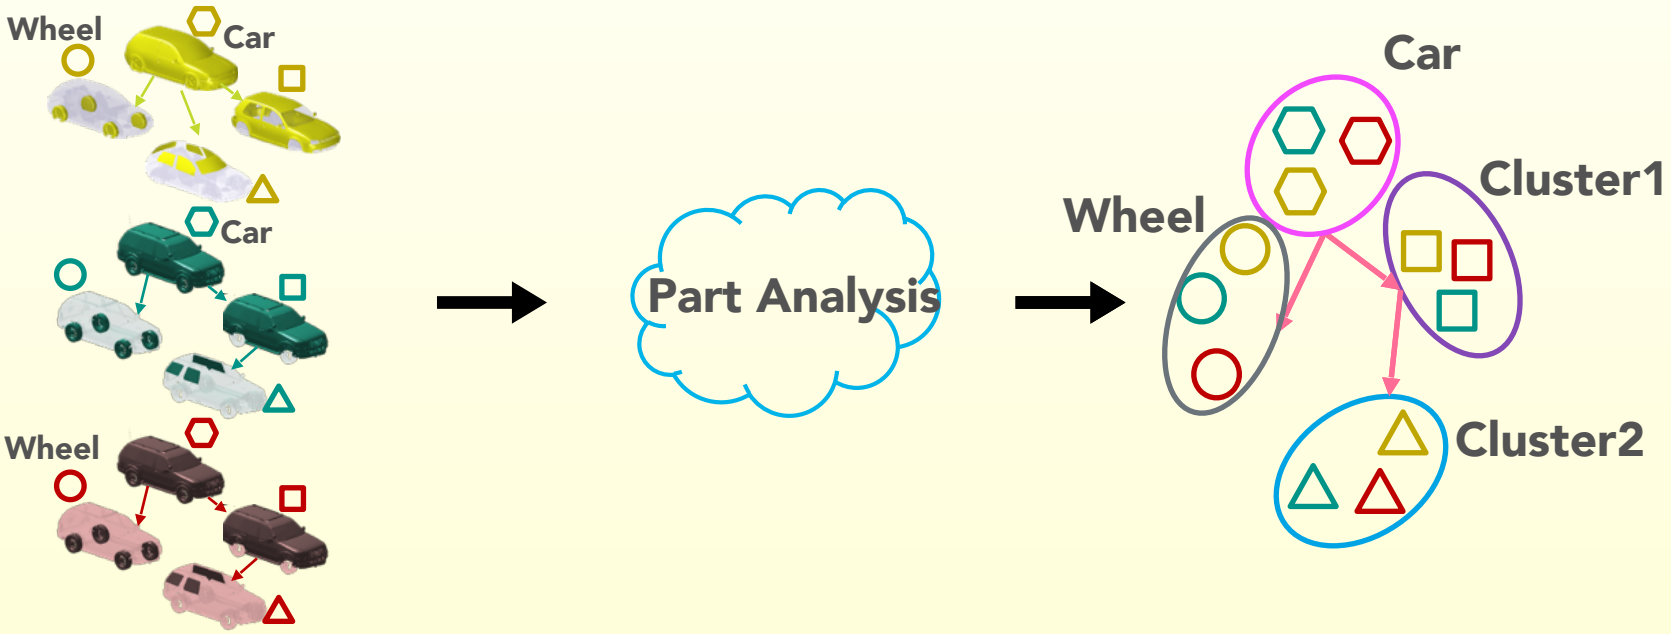
Wheel
Car
Car
Wheel
Part Analysis
Car
Cluster1
Wheel
Cluster2

Sample Clusters

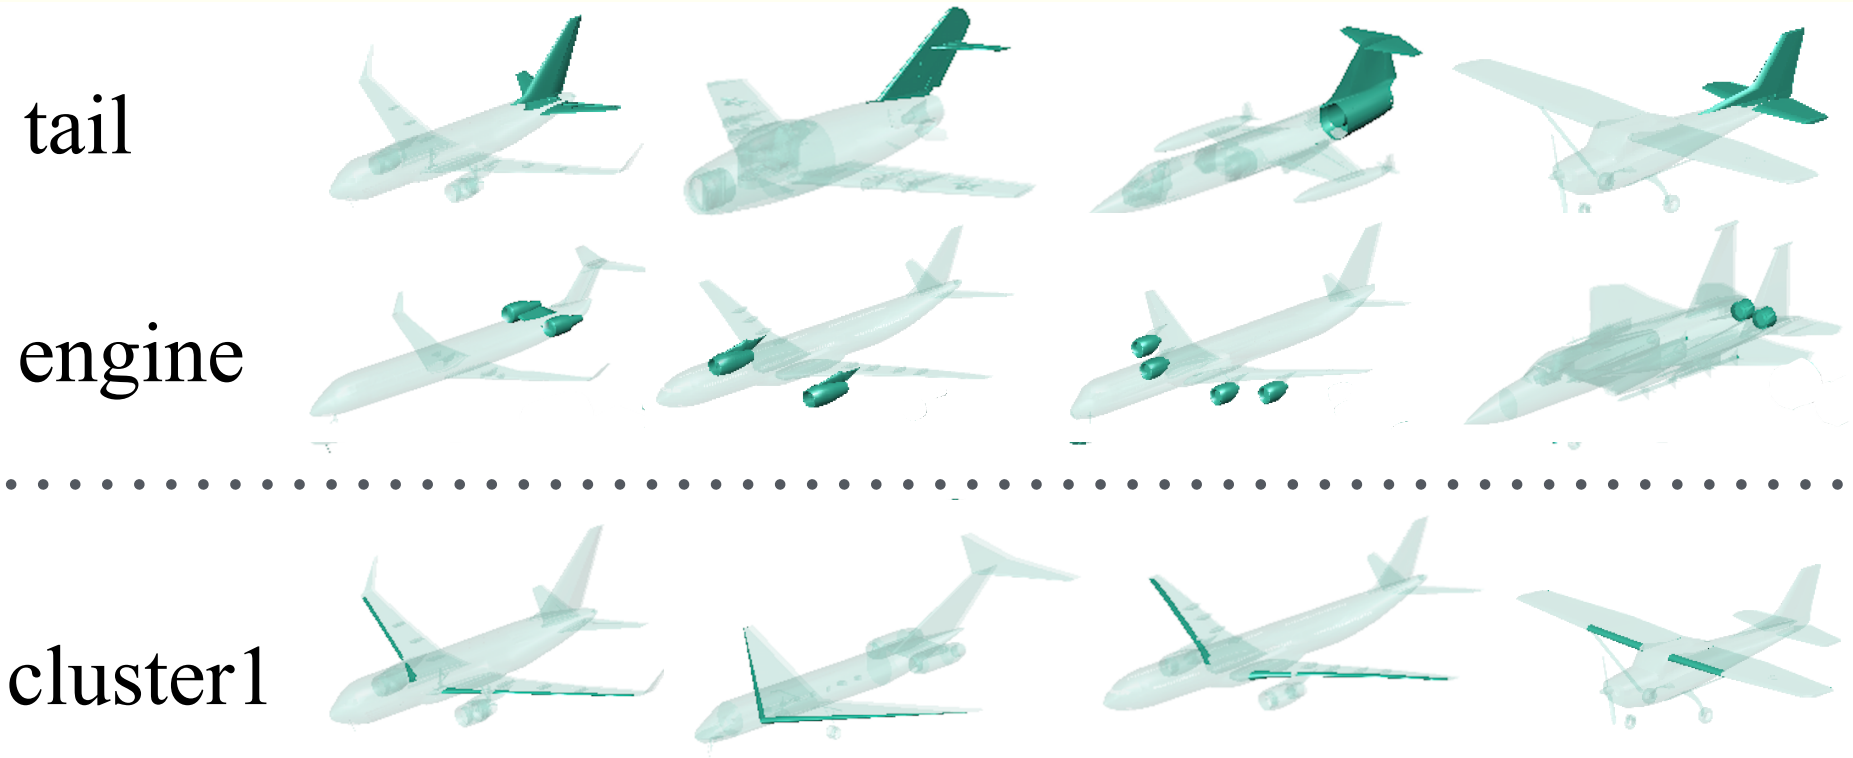

Sample Labeling

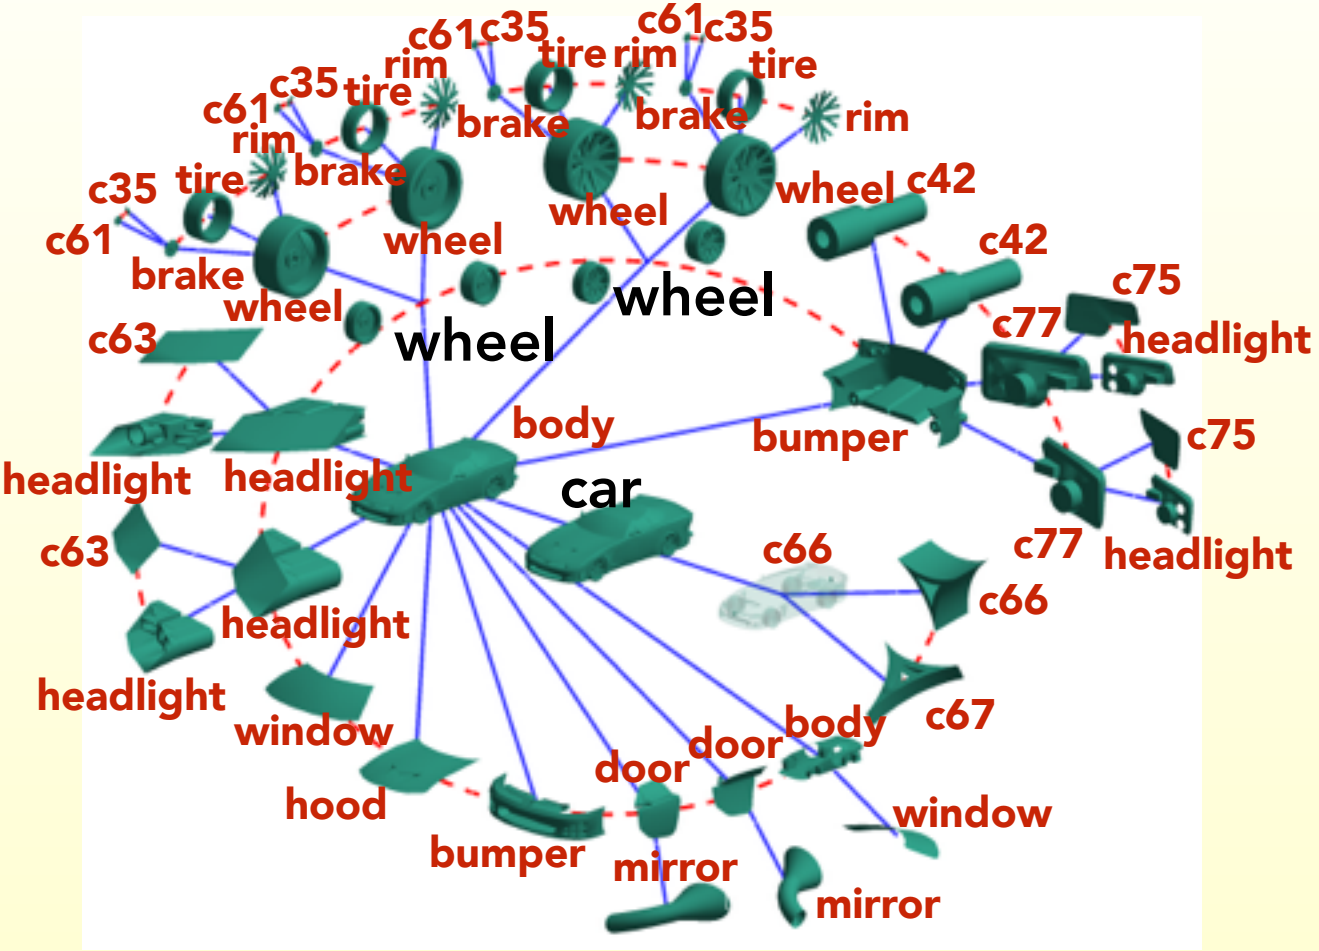

Sample Labeling

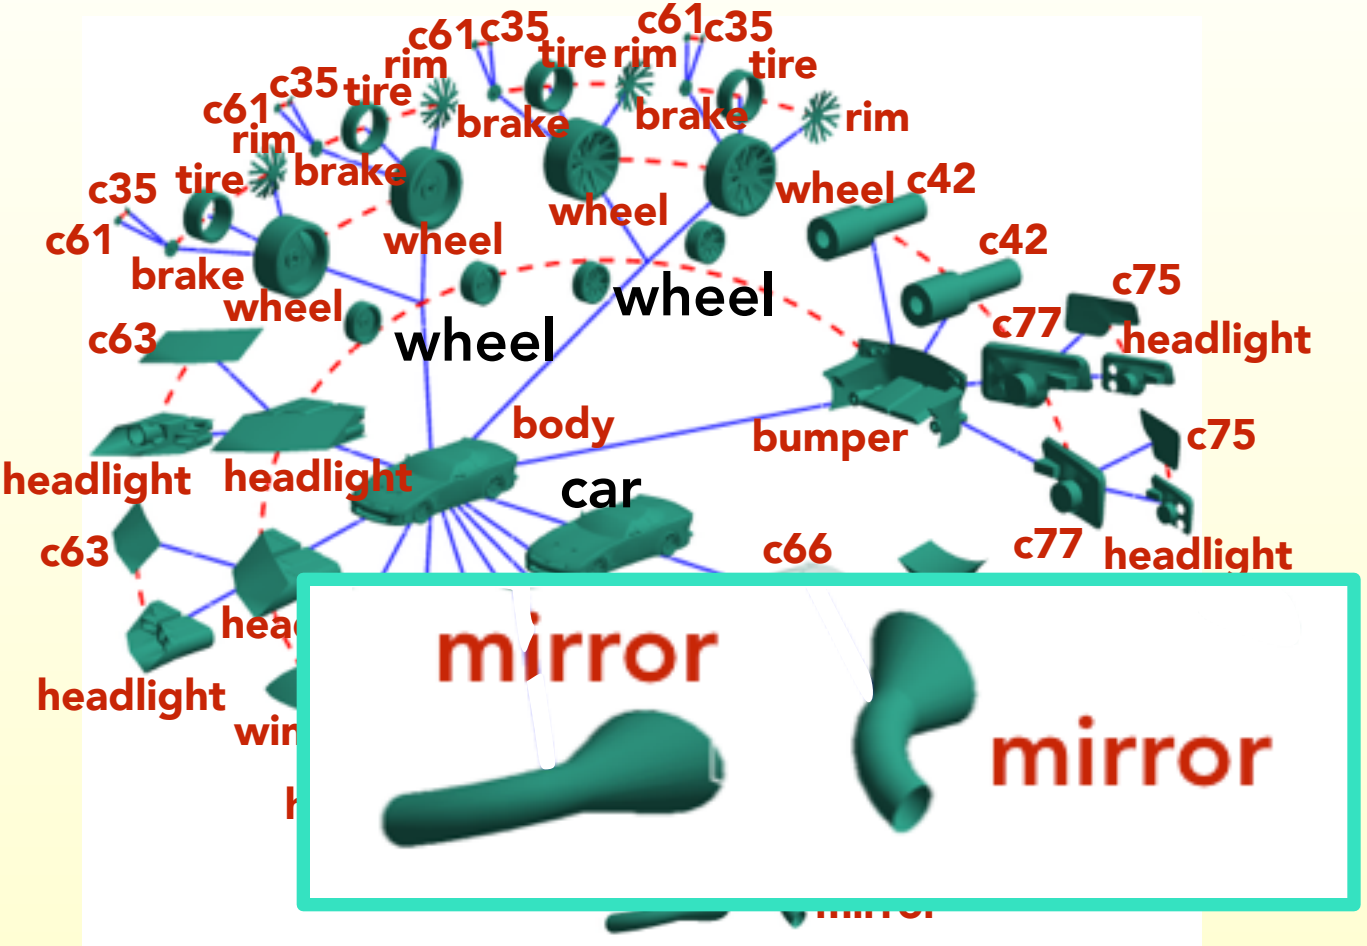

Sample Labeling

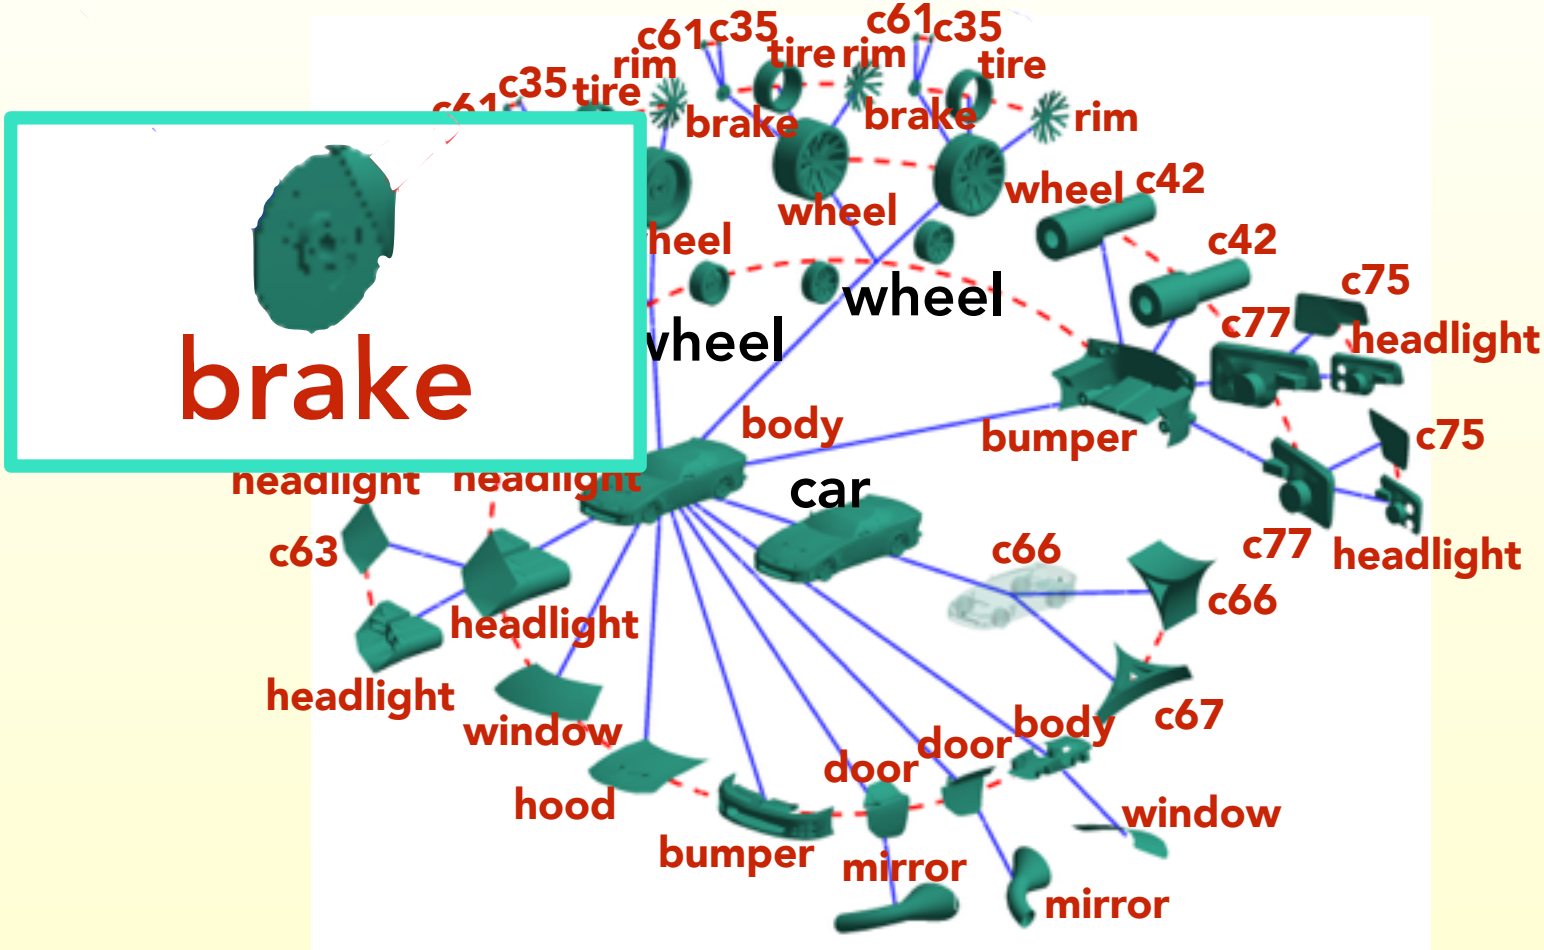

Inferred Canonical Hierarchy

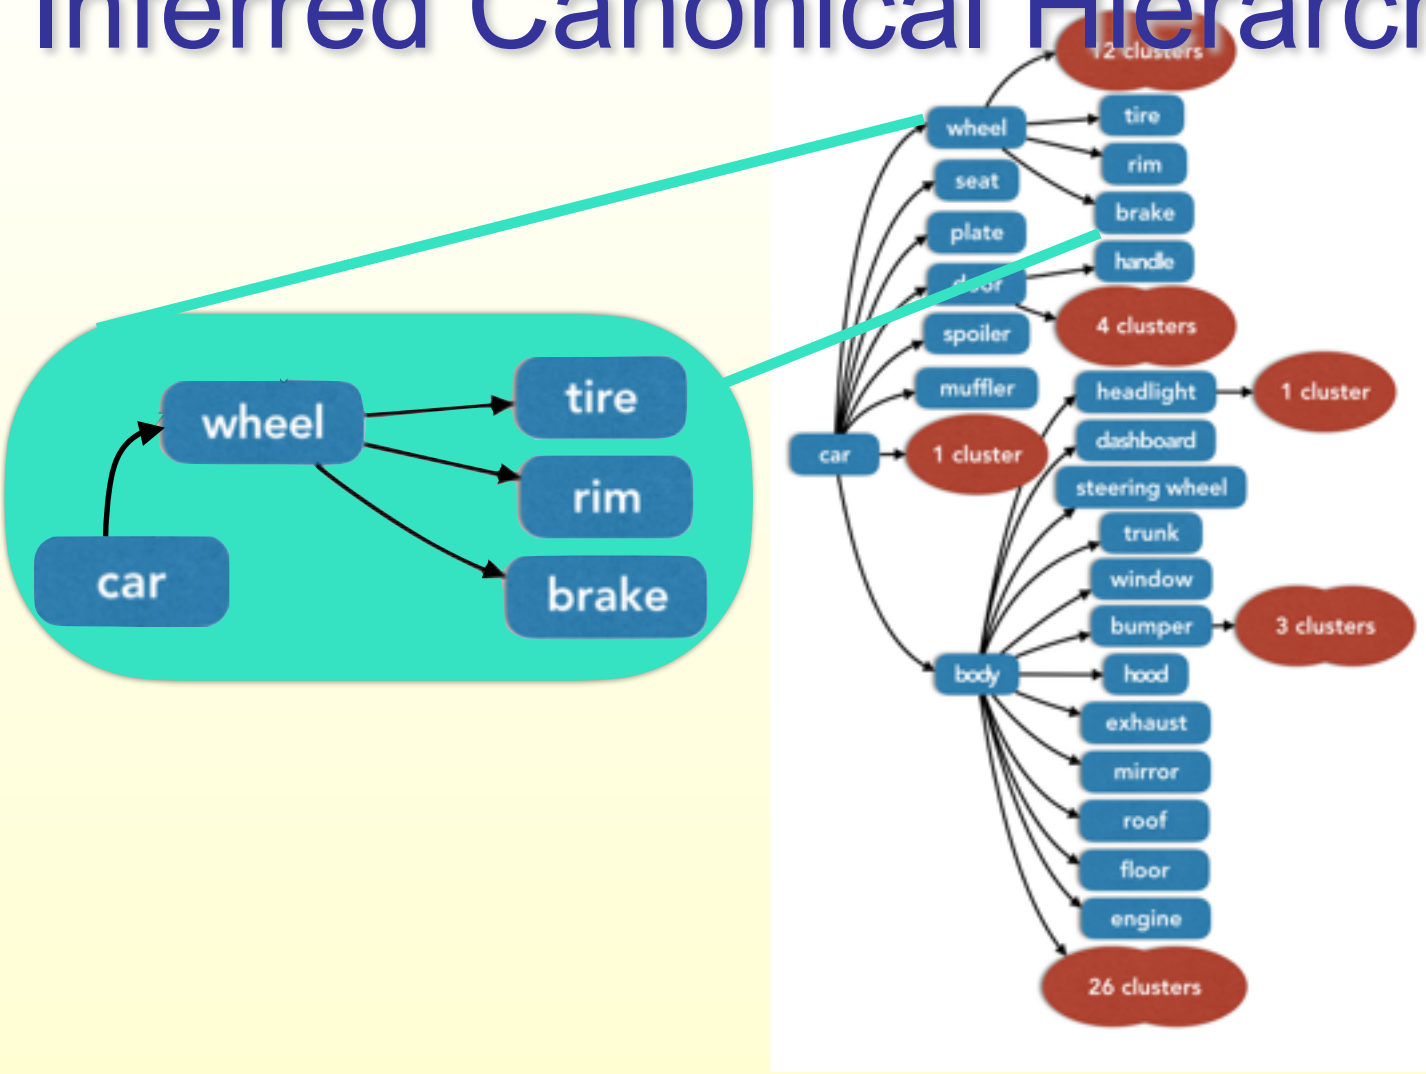

Approach Overview

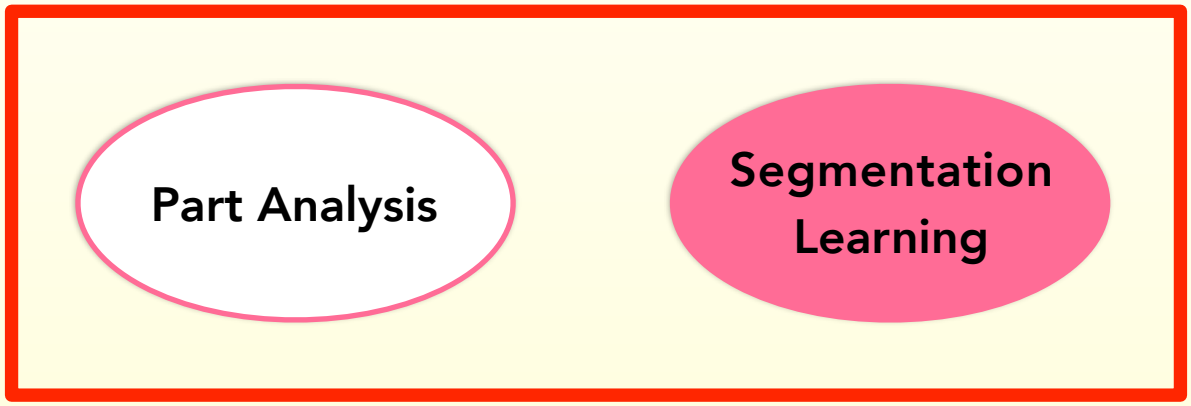

Training Stage

Learning Hierarchical Mesh Segmentation

Supervised Labeling

triangle face?

Results

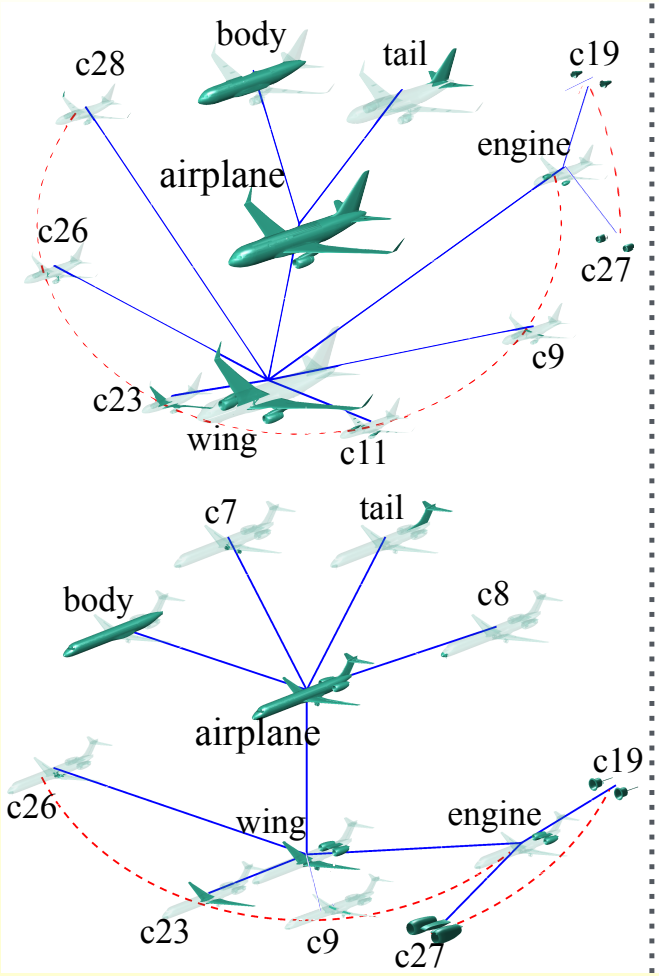
body
tail
c28
c19
engine
airplane
c26
c27
c9
c23
wing
c11
c7
tail
body
c8
airplane
c19
c26
wing
engine
c23
c9
c27

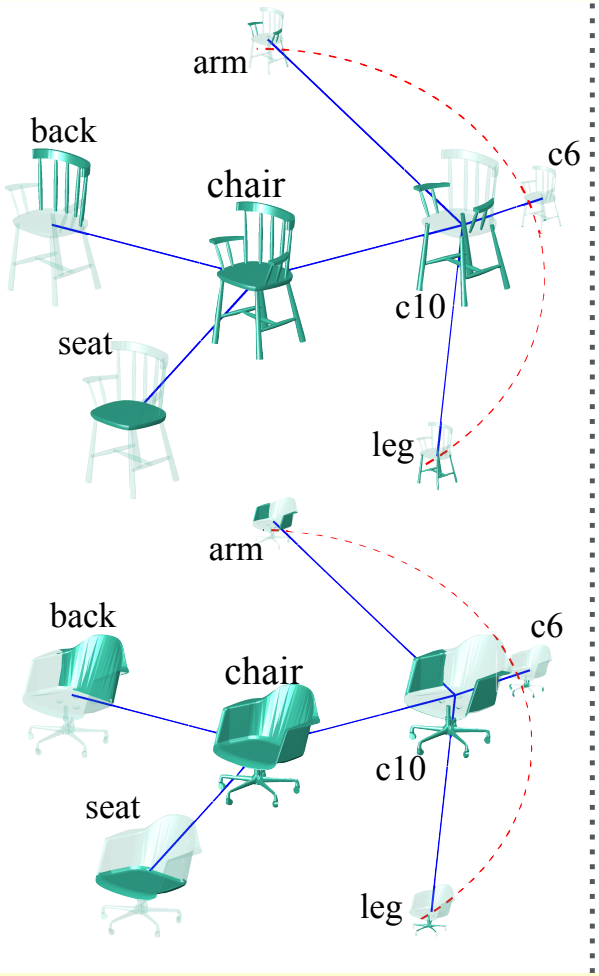
arm
back
c6
chair
c10
seat
leg
arm
back
c6
chair
c10
seat
leg

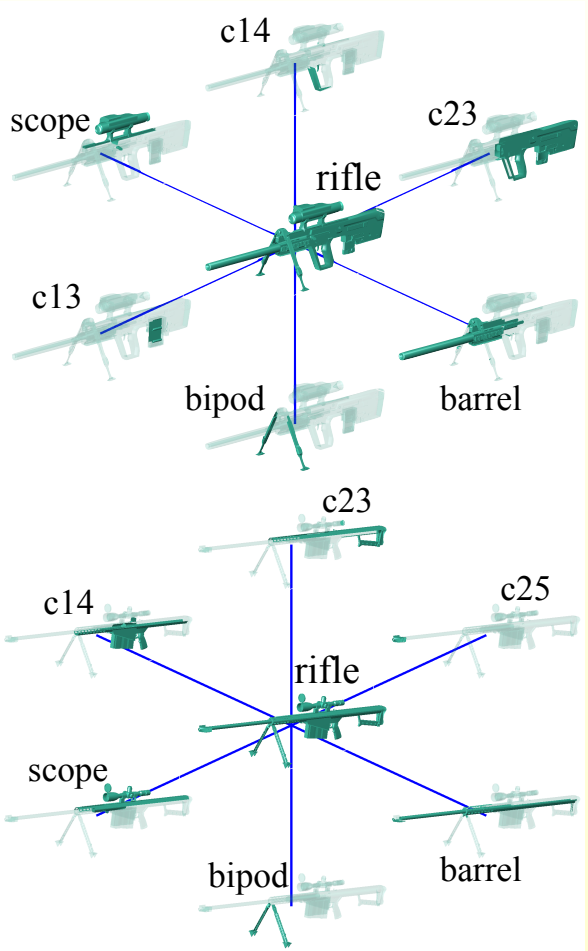
c14
scope
c23
rifle
c13
bipod
barrel
c23
c14
c25
rifle
scope
bipod
barrel

Results

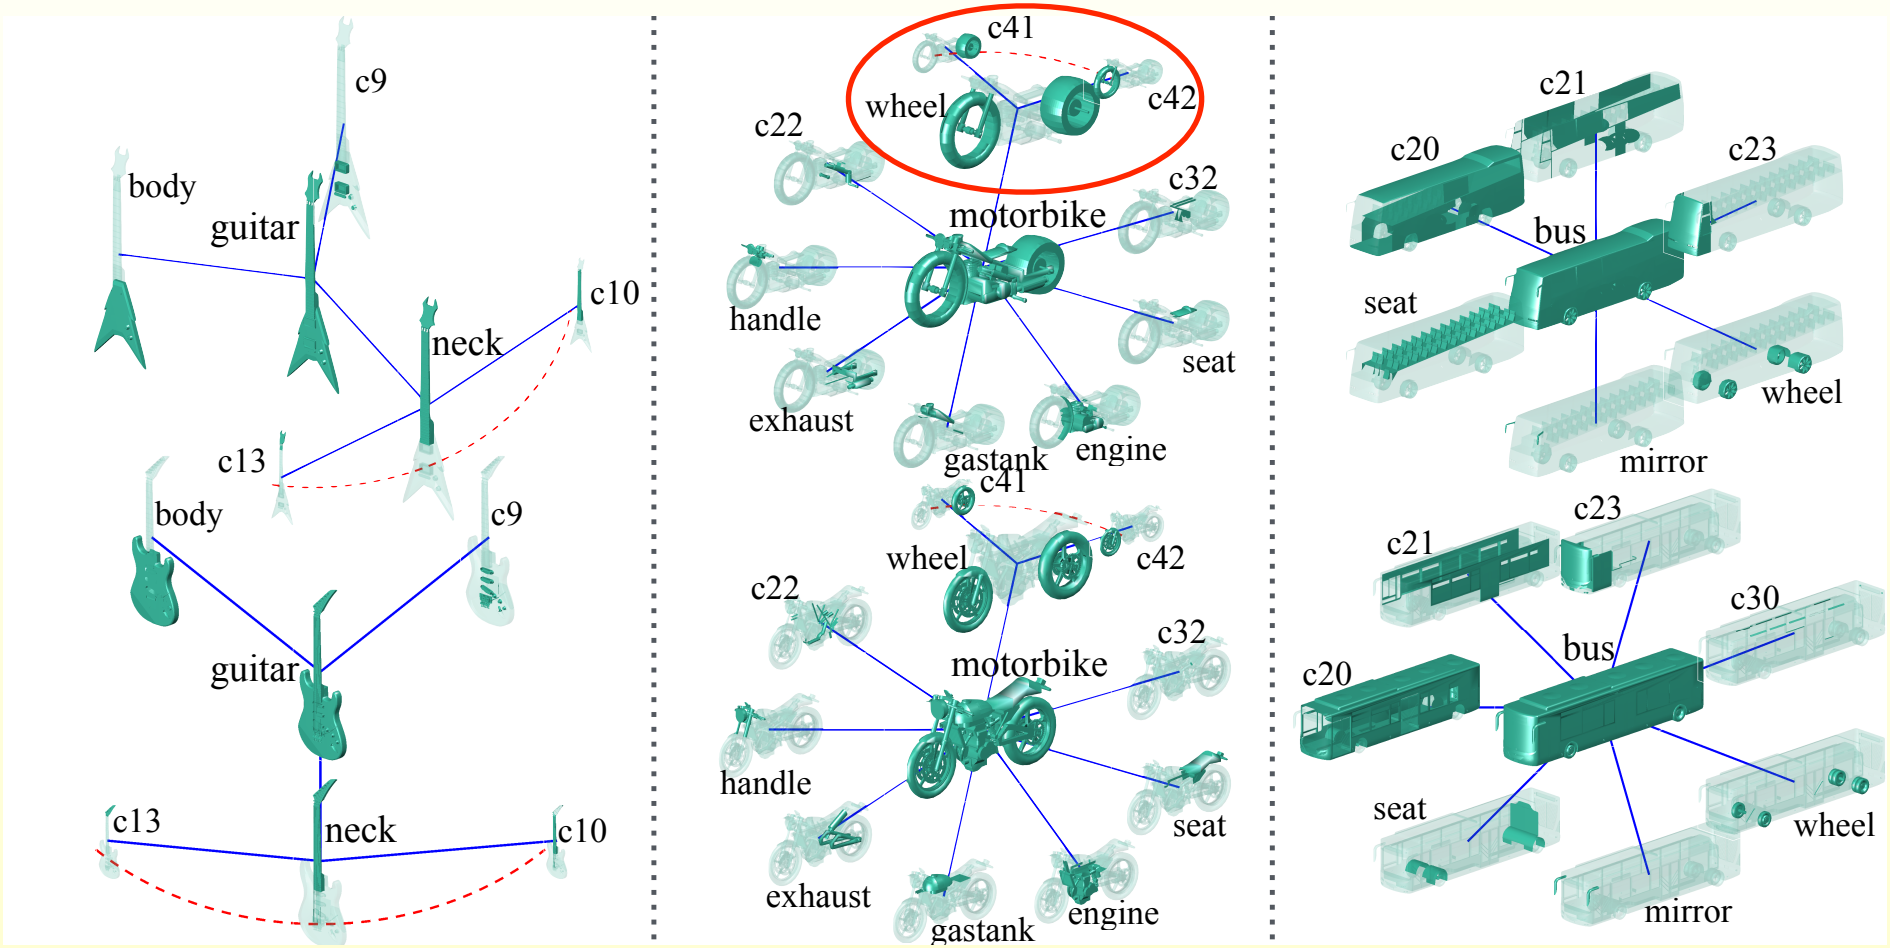
c9
body
guitar
c10
neck
c13
body
c9
guitar
c13
neck
c10
c41
wheel
c42
c22
c32
motorbike
handle
seat
exhaust
gastank
engine
c41
wheel
c42
c22
c32
motorbike
handle
seat
exhaust
gastank
engine
c21
c20
c23
bus
seat
wheel
mirror
c23
c21
c30
bus
c20
seat
wheel
mirror

Take-Home Message

- Distill wisdom from the crowd
- Knowledge could emerge while jointly analyzing a collection of shapes
- A novel method for mining massive but sparsely annotated object graphs “in the wild”

Agenda

- Goal of this lecture
- From semantic networks to data networks
  - WordNet, ImageNet and ShapeNet
- A hybrid approach for annotation acquisition
- From vertical networks to horizontal networks
  - Annotation transportation in ShapeNet

Useful Resources for Geometry Processing

- Software
  - **Blender** (free)
  - **Meshlab** (free)
  - Autodesk Maya (free for students)
  - Autodesk 3DSMax (free for students)
  - SketchUp (free)
  - …
- Geometry processing libraries
  - CGAL (C++)
  - libIGL (C++)
  - gptoolbox (MATLAB)
  - …