

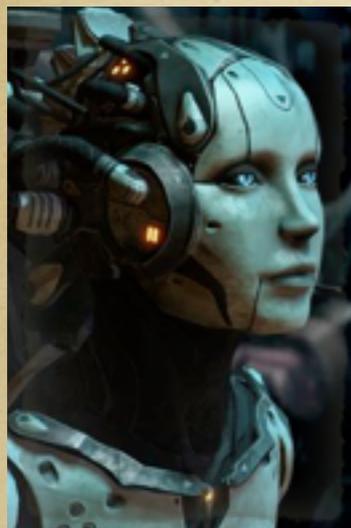
# Knowledge Base

# Robot in a room...



I can recognize everything  
in the room (proudly)

Bring me a cup of hot water



Well, I can tell you  
“where is the cup?”

➤ Recognize everything, but can do nothing

# What is missing?

Bring me a cup of hot water

- find a cup
- realize a cup has *containable* affordance

# Affordance

- A cup
- grasp
- filled in water
- pour

# Attribute

- A cup
- brittle
- made of glass,  
plastic
- has a handle

# What is missing?

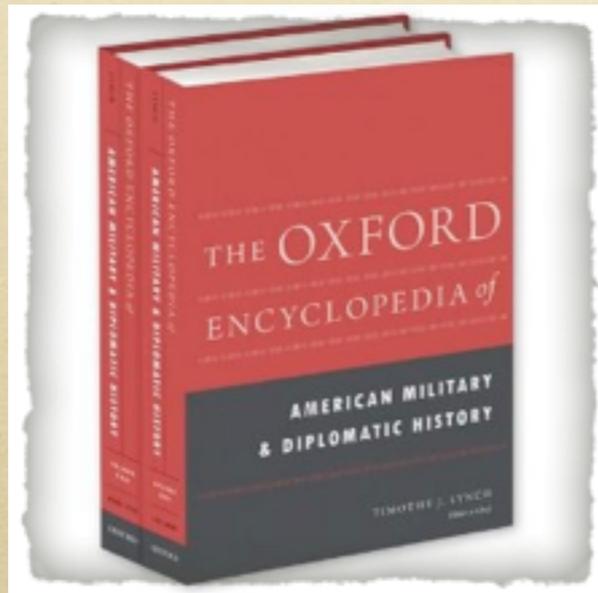
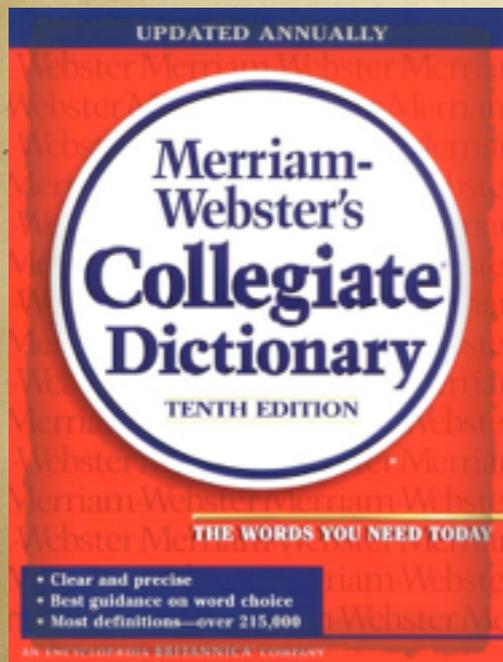
Bring me a cup of hot water

- find a cup
- realize a cup has *containable* affordance
- cup is empty
- find tap, fill in water
- find microwave
- heat it up

**The Common Knowledge**



# The Common Knowledge



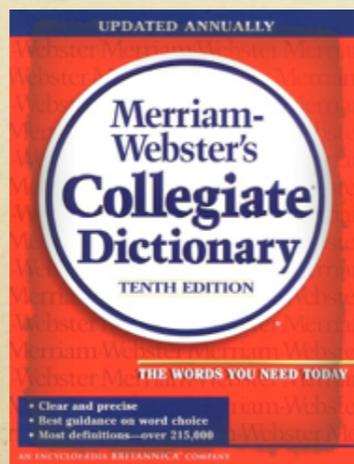
Structured



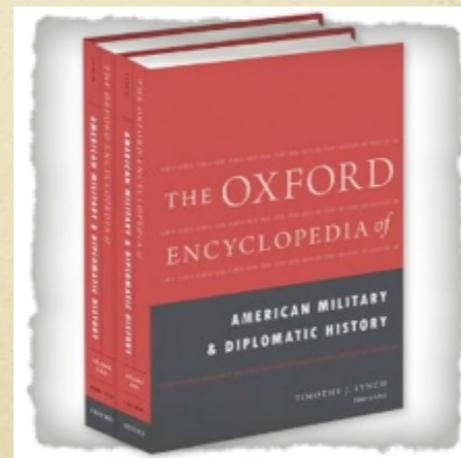
**WordNet**  
A lexical database for English



Specific



General

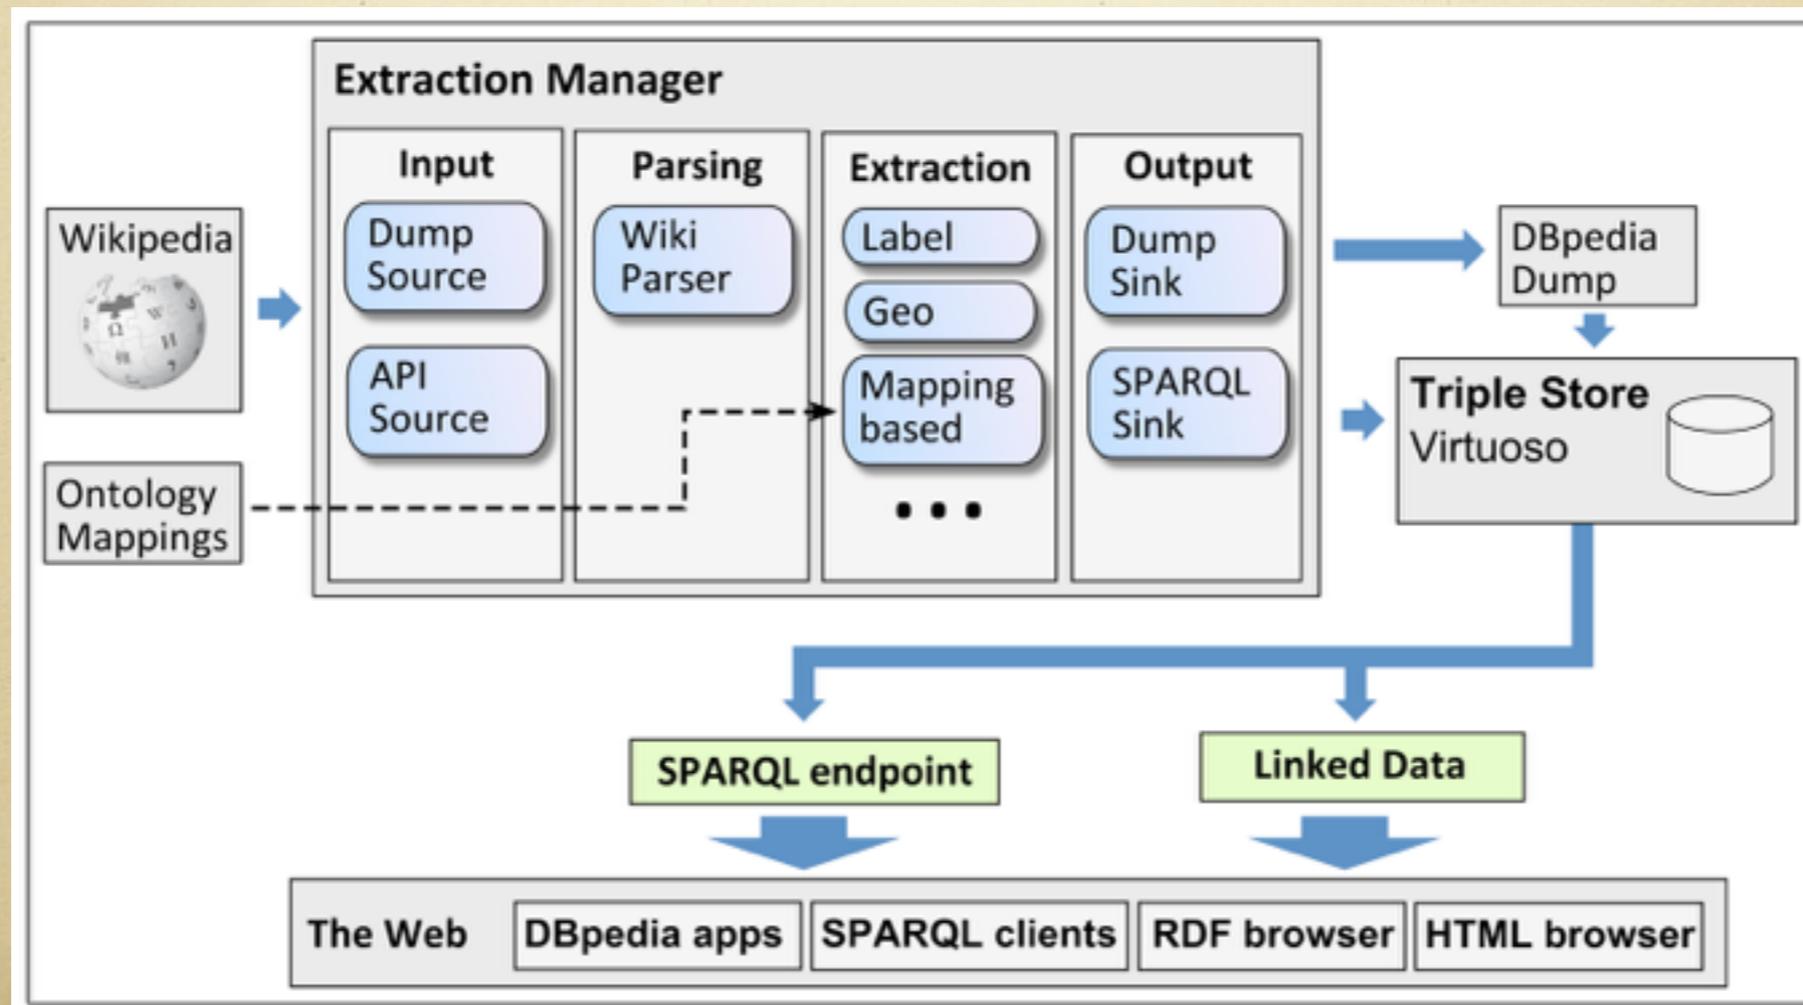


**WIKIPEDIA**  
The Free Encyclopedia

Casual format

# DBpedia

- DBpedia is a crowd-sourced community effort to extract structured information from Wikipedia





# Resource Description Framework

- A general method for conceptual description or modeling of information that is implemented in web resources.
- Make statements about web resources in the form of subject-predicate-object expression.

There is a Person identified by <http://www.w3.org/People/EM/contact#me>, whose name is Eric Miller, whose email address is e.miller123(at)example (changed for security purposes), and whose title is Dr.

- Subject: "<http://www.w3.org/People/EM/contact#me>"
- The objects are:
  - "Eric Miller" (with a predicate "whose name is"),
  - [mailto:e.miller123\(at\)example](mailto:e.miller123(at)example) (with a predicate "whose email address is"),  
and
  - "Dr." (with a predicate "whose title is").
- The predicates also have URIs. For example, the URI for each predicate:
  - "whose name is" is <http://www.w3.org/2000/10/swap/pim/contact#fullName>,
  - "whose email address is" is <http://www.w3.org/2000/10/swap/pim/contact#mailbox>,
  - "whose title is" is <http://www.w3.org/2000/10/swap/pim/contact#personalTitle>.

• RDF triples can be expressed:

- <http://www.w3.org/People/EM/contact#me>, <http://www.w3.org/2000/10/swap/pim/contact#fullName>, "Eric Miller"
- <http://www.w3.org/People/EM/contact#me>, <http://www.w3.org/2000/10/swap/pim/contact#mailbox>, [mailto:e.miller123\(at\)example](mailto:e.miller123(at)example)
- <http://www.w3.org/People/EM/contact#me>, <http://www.w3.org/2000/10/swap/pim/contact#personalTitle>, "Dr."
- <http://www.w3.org/People/EM/contact#me>, <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>, <http://www.w3.org/2000/10/swap/pim/contact#Person>

# DBpedia

➤ Revolutionize Wikipedia Search

➤ “Tell me all the episodes of Game of Thrones”  
rank them by released date.

```
SELECT *  
WHERE  
{  
  ?e <http://dbpedia.org/ontology/series>          <http://dbpedia.org/resource/Game_of_Thrones> .  
  ?e <http://dbpedia.org/ontology/releaseDate>     ?date  
  ?e <http://dbpedia.org/ontology/episodeNumber>  ?number  
  ?e <http://dbpedia.org/ontology/seasonNumber>   ?season  
}  
ORDER BY DESC(?date)
```

e	date	number	season
<a href="http://dbpedia.org/resource/Oathkeeper">http://dbpedia.org/resource/Oathkeeper</a>	2014-04-27+02:00	4	4
<a href="http://dbpedia.org/resource/Breaker_of_Chains">http://dbpedia.org/resource/Breaker of Chains</a>	2014-04-20+02:00	3	4
<a href="http://dbpedia.org/resource/The_Lion_and_the_Rose">http://dbpedia.org/resource/The Lion and the Rose</a>	2014-04-13+02:00	2	4
<a href="http://dbpedia.org/resource/Two_Swords_(Game_of_Thrones)">http://dbpedia.org/resource/Two Swords (Game of Thrones)</a>	2014-04-06+02:00	1	4
<a href="http://dbpedia.org/resource/Mhysa">http://dbpedia.org/resource/Mhysa</a>	2013-06-09+02:00	10	3
<a href="http://dbpedia.org/resource/The_Rains_of_Castamere">http://dbpedia.org/resource/The Rains of Castamere</a>	2013-06-02+02:00	9	3
<a href="http://dbpedia.org/resource/Second_Sons">http://dbpedia.org/resource/Second Sons</a>	2013-05-19+02:00	8	3
<a href="http://dbpedia.org/resource/The_Bear_and_the_Maiden_Fair">http://dbpedia.org/resource/The Bear and the Maiden Fair</a>	2013-05-12+02:00	7	3
<a href="http://dbpedia.org/resource/The_Climb_(Game_of_Thrones)">http://dbpedia.org/resource/The Climb (Game of Thrones)</a>	2013-05-05+02:00	6	3
<a href="http://dbpedia.org/resource/Kissed_by_Fire">http://dbpedia.org/resource/Kissed by Fire</a>	2013-04-28+02:00	5	3
<a href="http://dbpedia.org/resource/And_Now_His_Watch_Is_Ended">http://dbpedia.org/resource/And Now His Watch Is Ended</a>	2013-04-21+02:00	4	3
<a href="http://dbpedia.org/resource/Walk_of_Punishment">http://dbpedia.org/resource/Walk of Punishment</a>	2013-04-14+02:00	3	3
<a href="http://dbpedia.org/resource/Dark_Wings,_Dark_Words">http://dbpedia.org/resource/Dark Wings, Dark Words</a>	2013-04-07+02:00	2	3
<a href="http://dbpedia.org/resource/Valar_Dohaeris">http://dbpedia.org/resource/Valar Dohaeris</a>	2013-03-31+02:00	1	3
<a href="http://dbpedia.org/resource/Valar_Morghulis">http://dbpedia.org/resource/Valar Morghulis</a>	2012-06-03+02:00	10	2
<a href="http://dbpedia.org/resource/Blackwater_(Game_of_Thrones)">http://dbpedia.org/resource/Blackwater (Game of Thrones)</a>	2012-05-27+02:00	9	2
<a href="http://dbpedia.org/resource/The_Prince_of_Winterfell">http://dbpedia.org/resource/The Prince of Winterfell</a>	2012-05-20+02:00	8	2
<a href="http://dbpedia.org/resource/A_Man_Without_Honor">http://dbpedia.org/resource/A Man Without Honor</a>	2012-05-13+02:00	7	2
<a href="http://dbpedia.org/resource/The_Old_Gods_and_the_New">http://dbpedia.org/resource/The Old Gods and the New</a>	2012-05-06+02:00	6	2
<a href="http://dbpedia.org/resource/The_Ghost_of_Harrenhal">http://dbpedia.org/resource/The Ghost of Harrenhal</a>	2012-04-29+02:00	5	2
<a href="http://dbpedia.org/resource/Garden_of_Bones">http://dbpedia.org/resource/Garden of Bones</a>	2012-04-22+02:00	4	2
<a href="http://dbpedia.org/resource/What_Is_Dead_May_Never_Die">http://dbpedia.org/resource/What Is Dead May Never Die</a>	2012-04-15+02:00	3	2
<a href="http://dbpedia.org/resource/The_Night_Lands">http://dbpedia.org/resource/The Night Lands</a>	2012-04-08+02:00	2	2
<a href="http://dbpedia.org/resource/The_North_Remembers">http://dbpedia.org/resource/The North Remembers</a>	2012-04-01+02:00	1	2
<a href="http://dbpedia.org/resource/Fire_and_Blood_(Game_of_Thrones)">http://dbpedia.org/resource/Fire and Blood (Game of Thrones)</a>	2011-06-19+02:00	10	1
<a href="http://dbpedia.org/resource/Baelor">http://dbpedia.org/resource/Baelor</a>	2011-06-12+02:00	9	1
<a href="http://dbpedia.org/resource/The_Pointy_End">http://dbpedia.org/resource/The Pointy End</a>	2011-06-05+02:00	8	1
<a href="http://dbpedia.org/resource/You_Win_or_You_Die">http://dbpedia.org/resource/You Win or You Die</a>	2011-05-29+02:00	7	1
<a href="http://dbpedia.org/resource/A_Golden_Crown">http://dbpedia.org/resource/A Golden Crown</a>	2011-05-22+02:00	6	1
<a href="http://dbpedia.org/resource/The_Wolf_and_the_Lion">http://dbpedia.org/resource/The Wolf and the Lion</a>	2011-05-15+02:00	5	1
<a href="http://dbpedia.org/resource/Cripples,_Bastards,_and_Broken_Things">http://dbpedia.org/resource/Cripples, Bastards, and Broken Things</a>	2011-05-08+02:00	4	1
<a href="http://dbpedia.org/resource/Lord_Snow">http://dbpedia.org/resource/Lord Snow</a>	2011-05-01+02:00	3	1
<a href="http://dbpedia.org/resource/The_Kingsroad">http://dbpedia.org/resource/The Kingsroad</a>	2011-04-24+02:00	2	1
<a href="http://dbpedia.org/resource/Winter_Is_Coming">http://dbpedia.org/resource/Winter Is Coming</a>	2011-04-17+02:00	1	1

# DBpedia

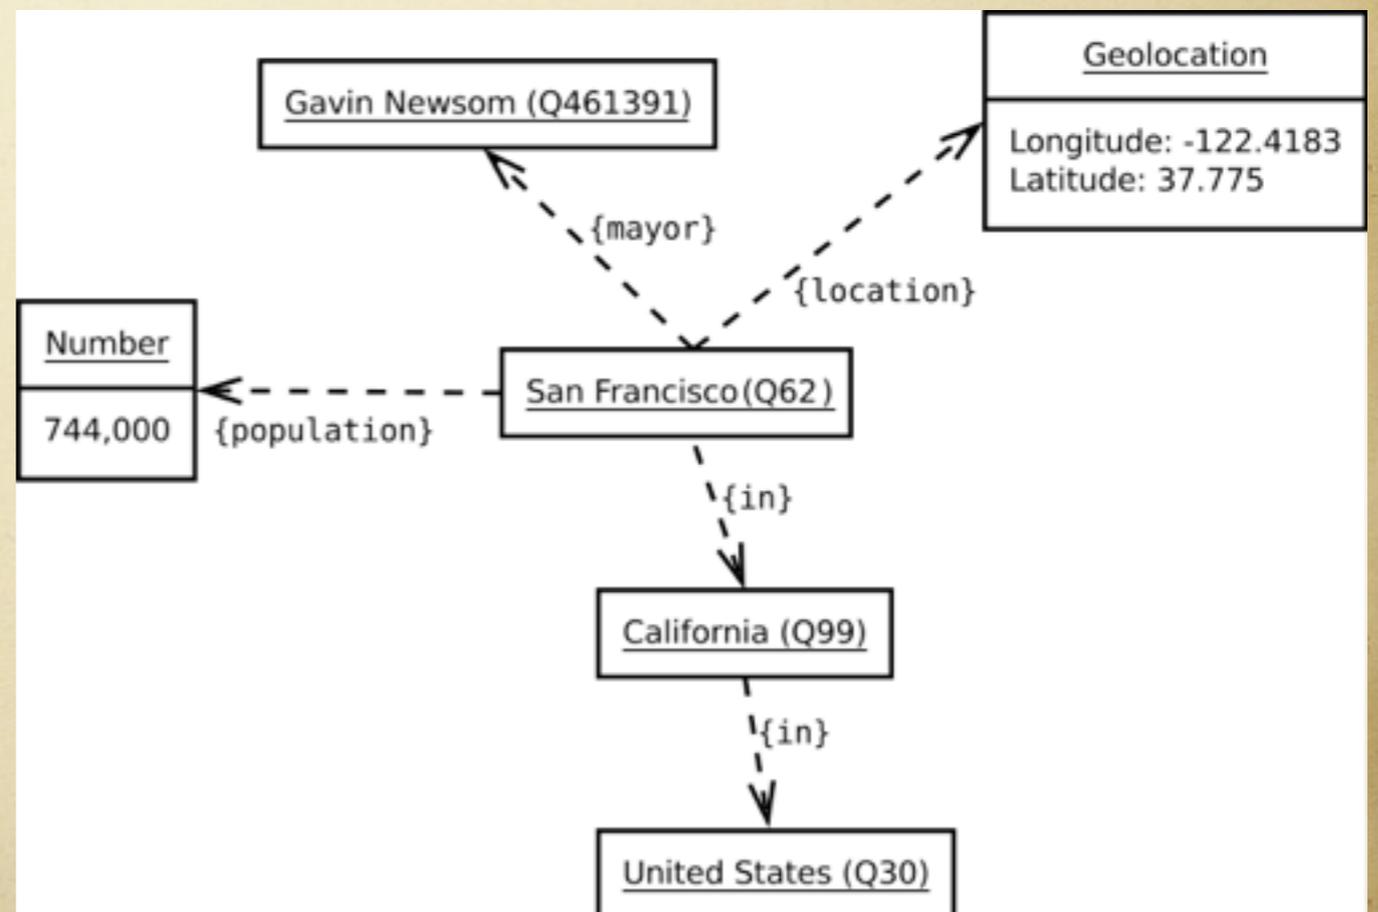
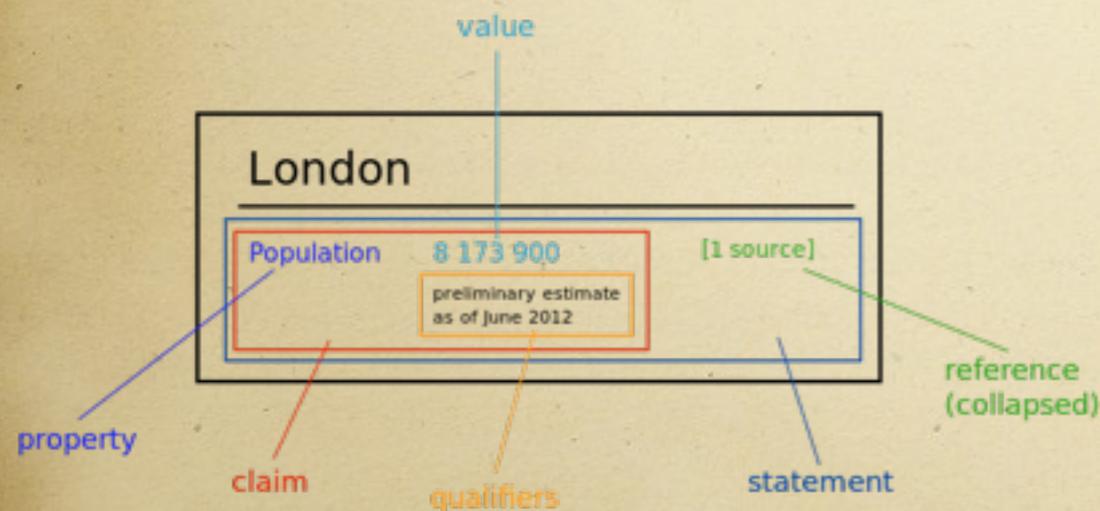
- A lot of other applications
- <http://wiki.dbpedia.org/Applications>
- Available in multiple languages
- Downloadable

# Knowledge Base

- Source of knowledge: internet, human input
- Structure: Graph = Node + Edge
  - RDF: subject-predicate-object
  - Node: entity
  - Edge: relation

# WikiData

- Very similar as DBpedia
  - link to more source
  - act as knowledge base for Wikimedia



# Wait, wait...

➤ Knowledge base, structured data organized in graph, DBpedia, Wikidata, Freebase.

➤ But...

Bring me a cup of hot water

➤ **Need low level knowledge**

- find a cup
- a cup has *containable* affordance
- cup is empty
- find tap, fill in water
- find microwave
- heat it up

# ConceptNet

- A semantic network containing lots of things computers should know about the world.

ConceptNet contains everyday basic knowledge:

learn — *MotivatedByGoal* → knowledge

*You would learn because you want knowledge.*

Cultural knowledge:

saxophone — *UsedFor* → jazz

*A saxophone is used for jazz.*

a cup has *containable* affordance

And scientific knowledge:

semantic role — *HasContext* → linguistics

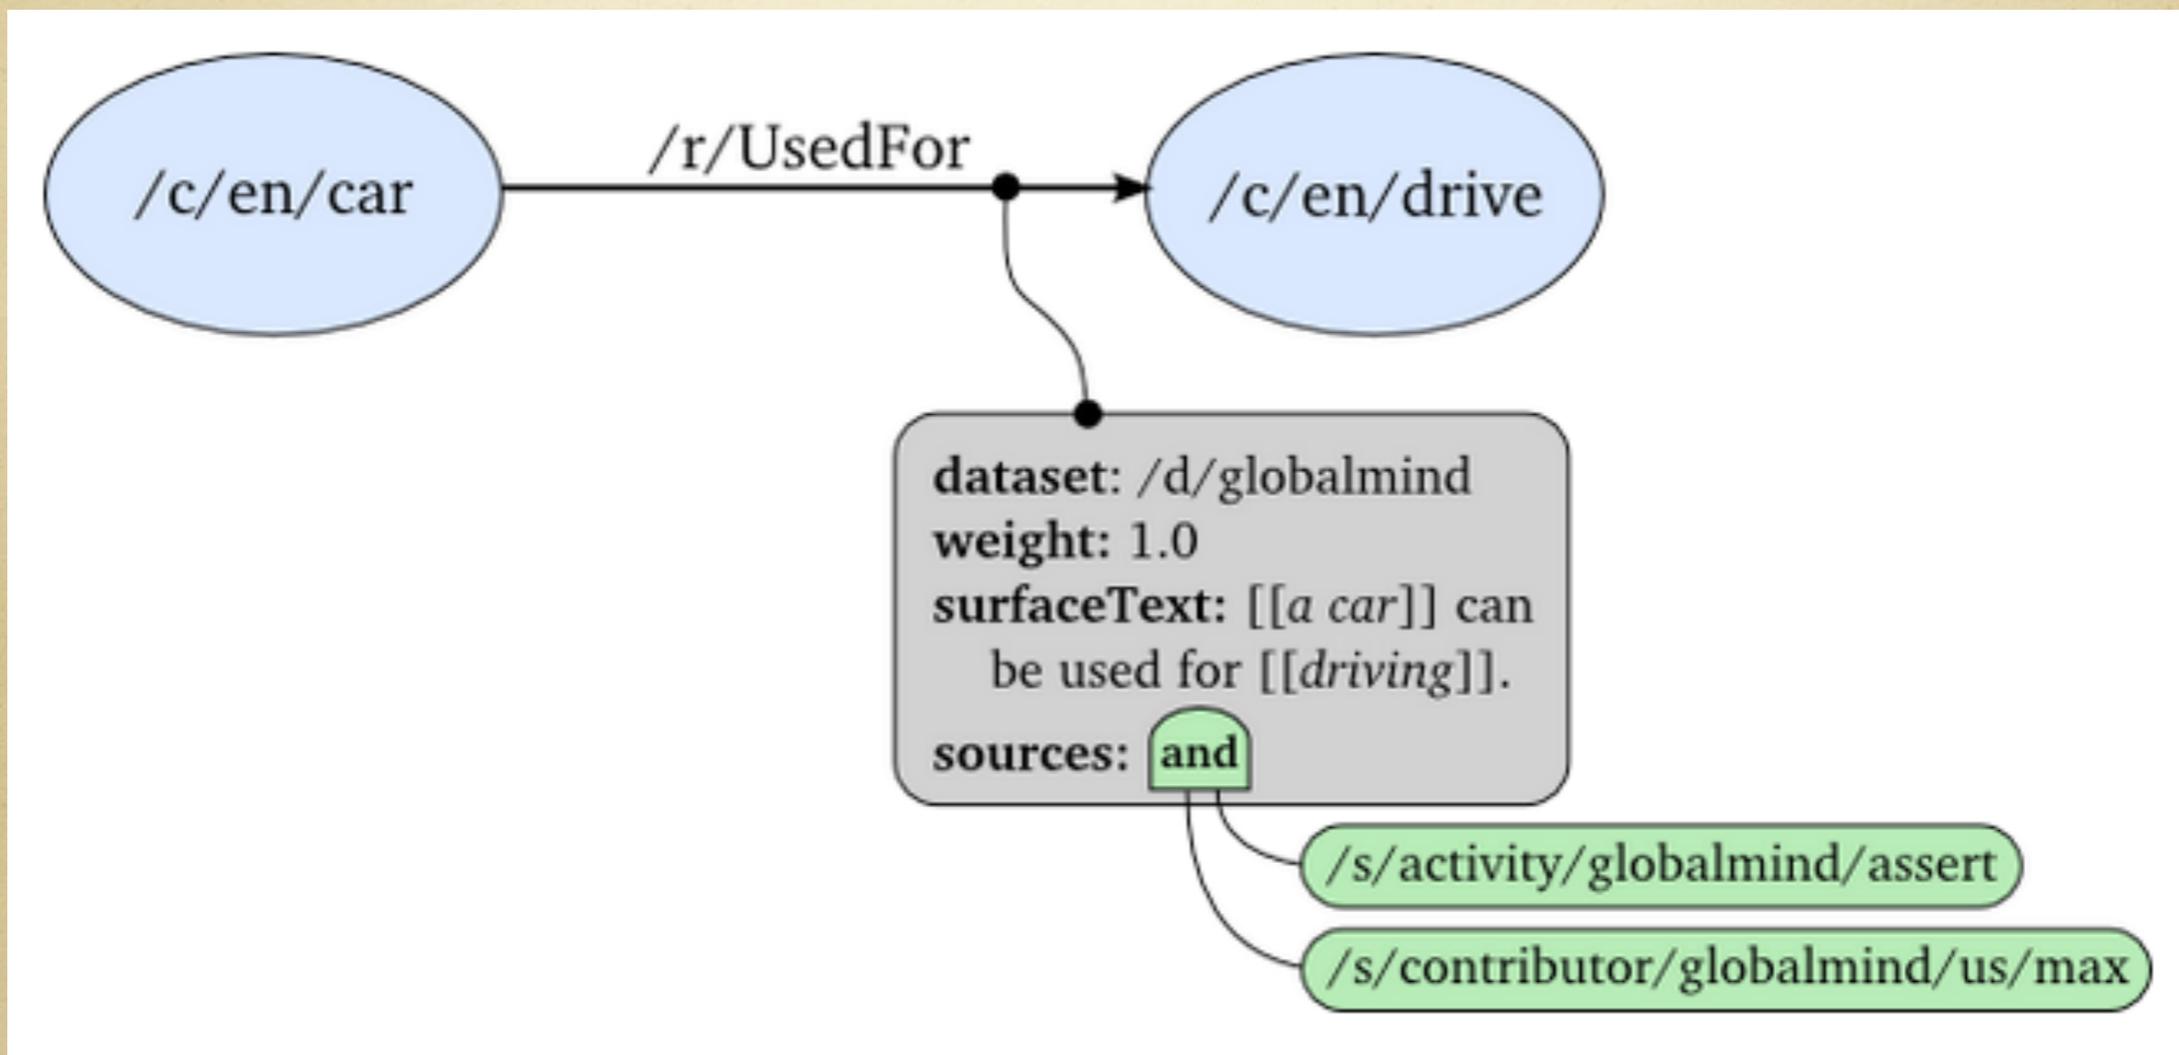
*"Semantic role" is a term in linguistics.*

It would not adequately represent human knowledge if it didn't contain other languages besides English, as well:

本 — *MadeOf* → 紙

本は紙でできている。(A book is made of paper.)

# ConceptNet



# ConceptNet

- Free to download
- Provide API to:
  - Retrieve the data for particular nodes and edges
  - Query for edges with given properties
  - Measure and query the semantic distance between nodes

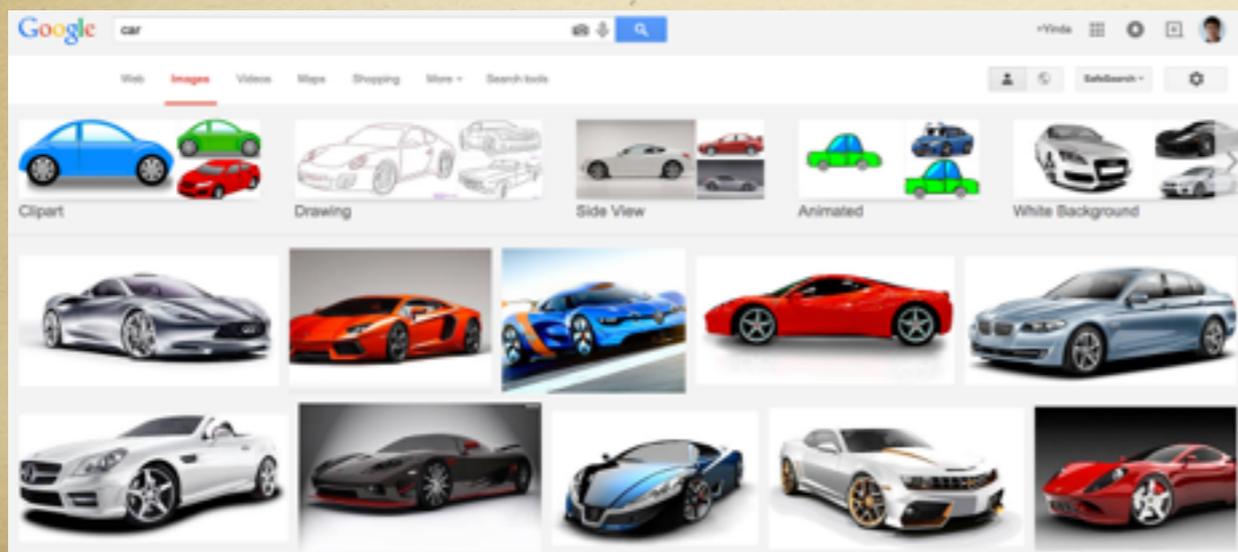
# So far...

- There are lexical knowledge base for both high-level and low-level knowledge ready online.
- To connect the knowledge with computer vision, we need visual knowledge base.
- Not as explicit as language
  - “A car can be used for driving”



# Never Ending Image Learner

- Learn from image searching engine (the weak association between image and text)
- what a car looks like?
- know that sheep are white



# Never Ending Image Learner

- NEIL is a computer program
- Run 24h per day, 7 days per week
- Automatically extract visual knowledge from internet data

➤ Learn to see



➤ Learn common sense

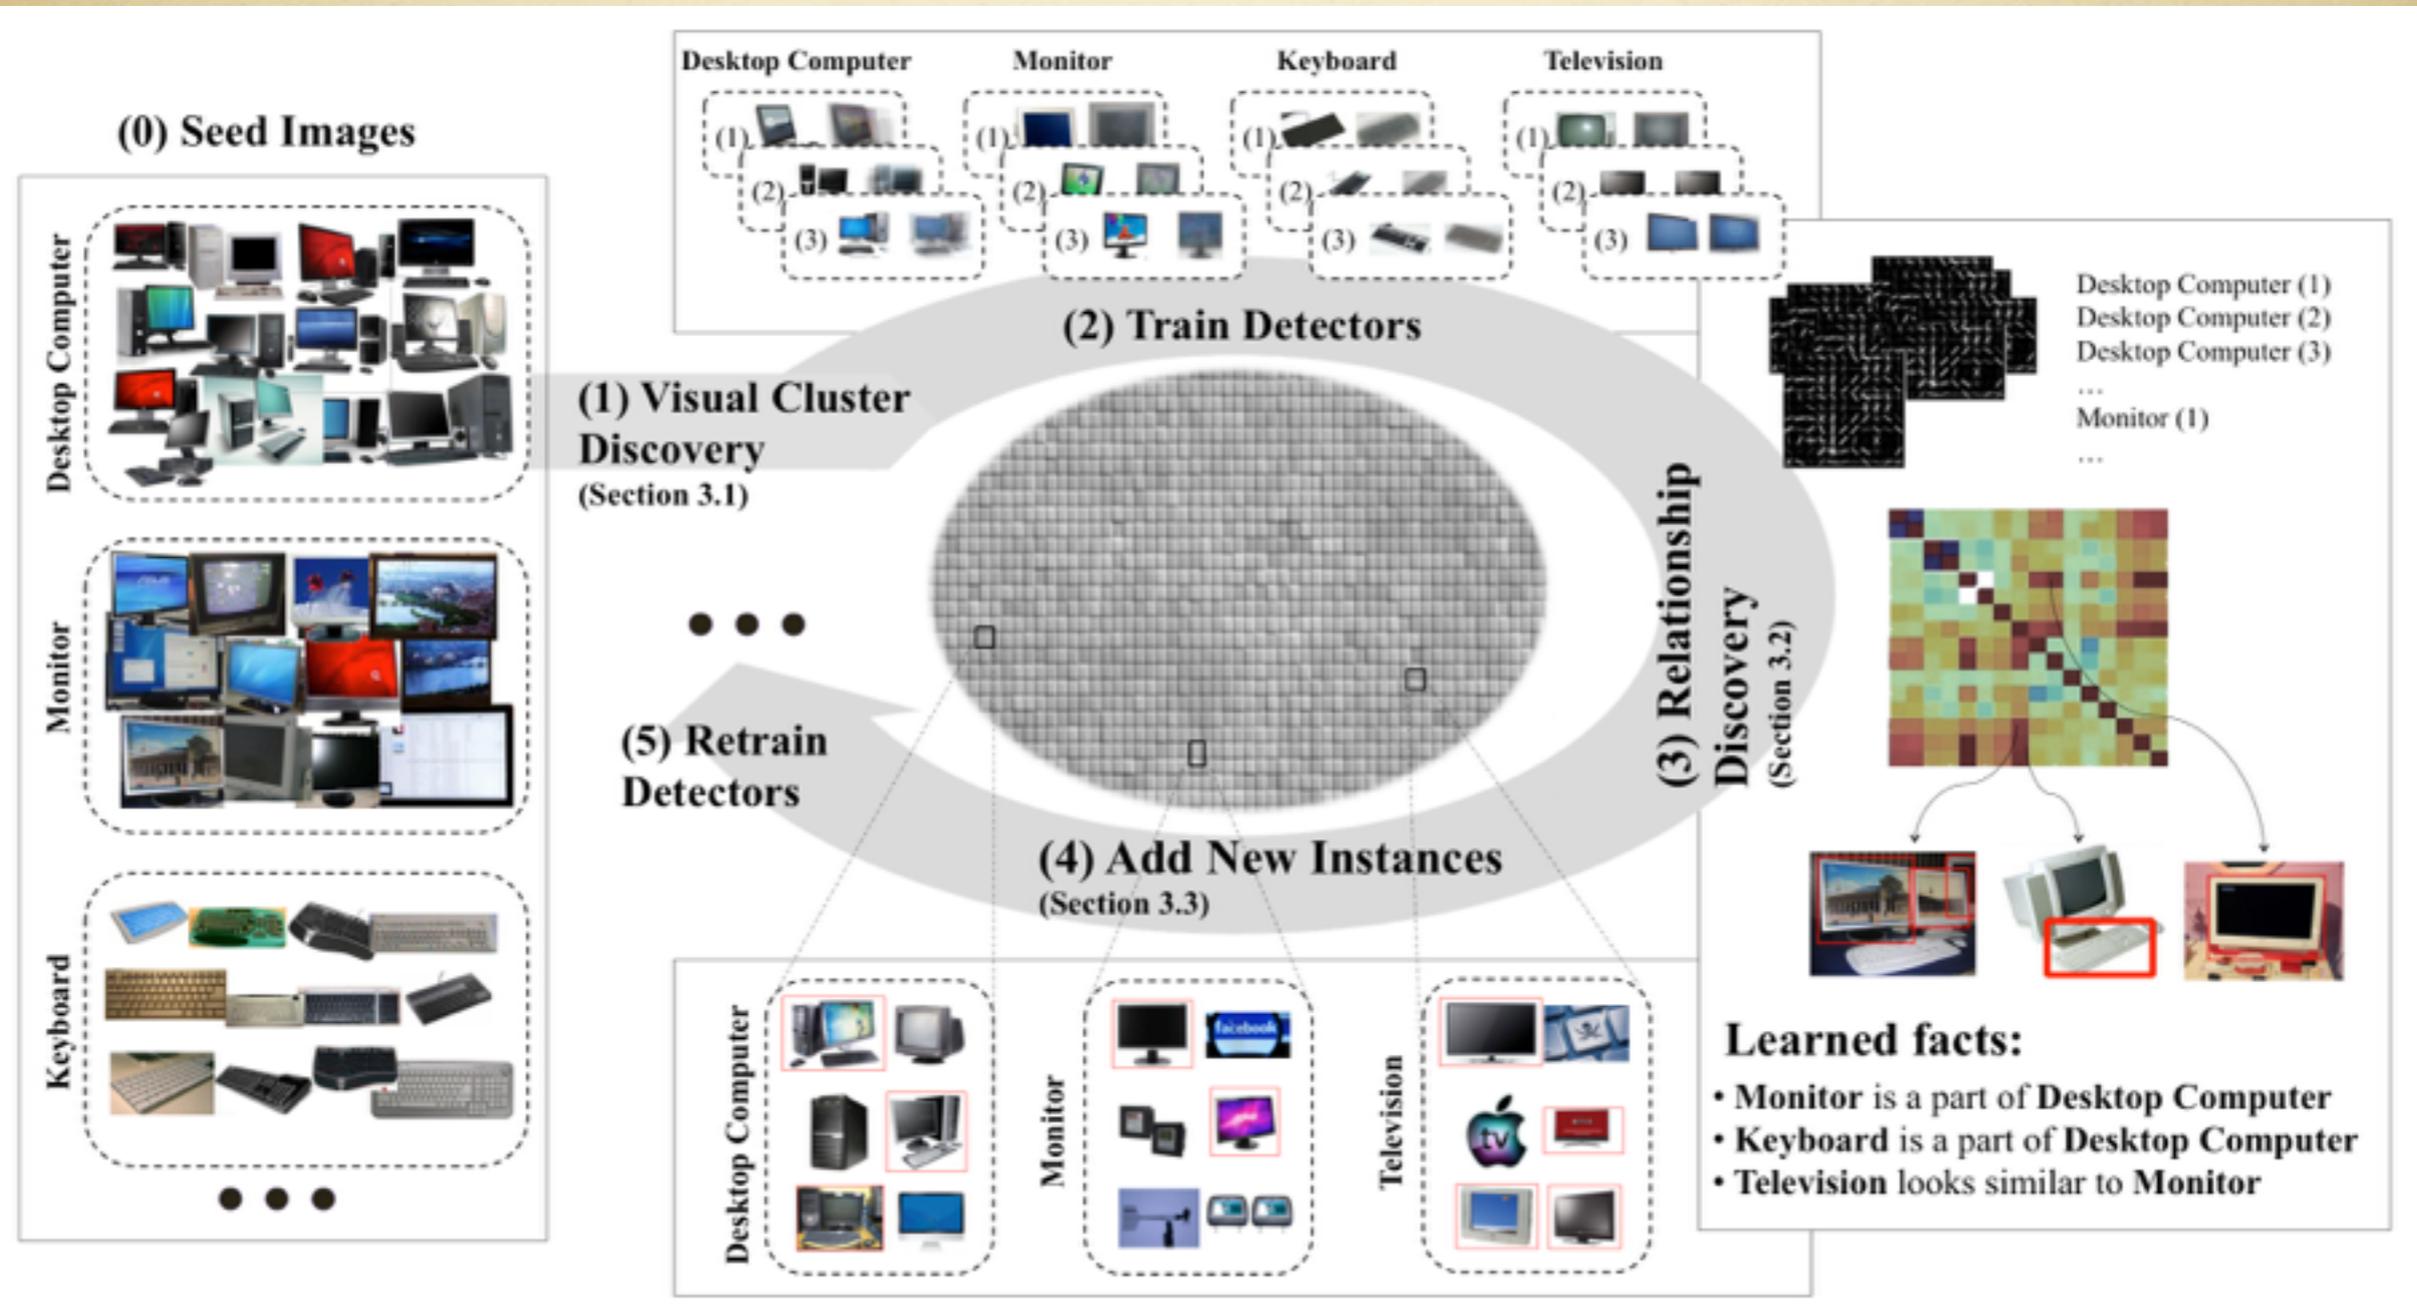
Airbus\_330 can be a kind of / look similar to Airplane.

Car can have a part Wheel.

Bamboo\_forest can be / can have Vertical\_lines.

Trading\_floor can be / can have Crowded.

# Never Ending Image Learner



# Never Ending Image Learner

- Seeding Classifier via Google Image Search
  - scene, attribute classifier; object, attribute detector.
  - Directly train scene and attribute classifier on downloaded images.
  - However, fail for object and attribute detector
    - Outlier, Polysemy, Visual diversity, Localization

# Never Ending Image Learner

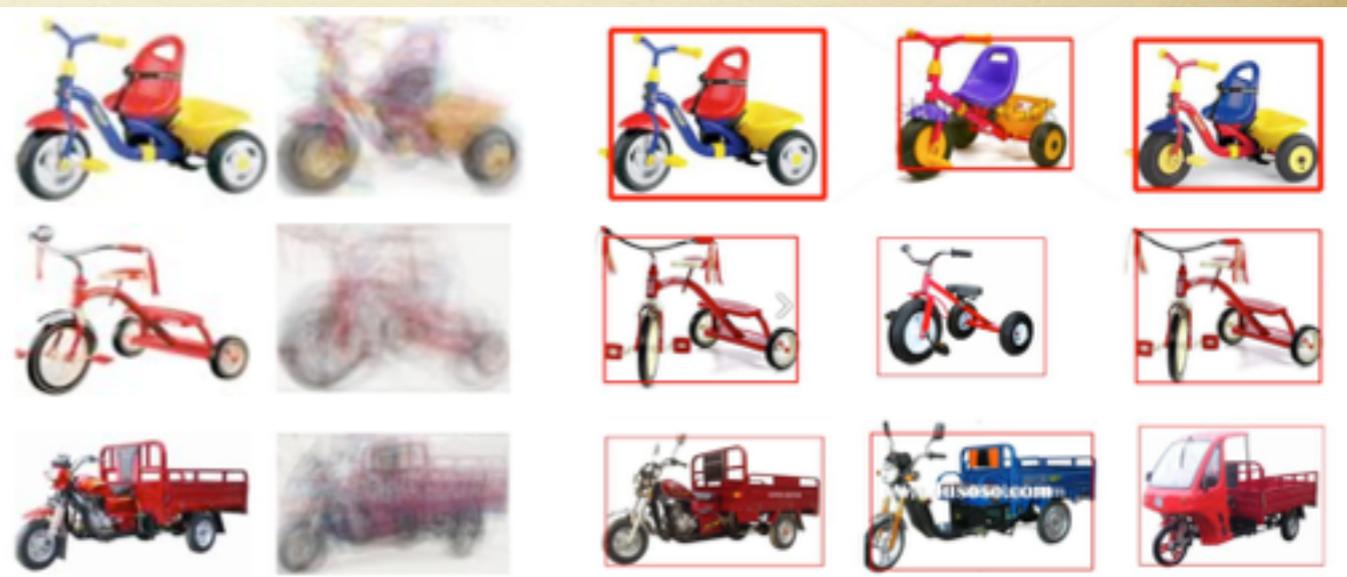
- Seeding Classifier via Google Image Search
- Train exemplar-LDA for each image
- Run detection on all images
- Get top K windows with high scores from multiple detectors
- Clustering with ELDA score vector
- Train classifier for each cluster

# Never Ending Image Learner

➤ Seeding Classifier via Google Image Search



(a) Google Image Search for "tricycle"



(b) Sub-category Discovery

# Never Ending Image Learner

➤ Extract Relationships

➤ Object-Object Relationships:

➤ Partonomy: Eye is a part of Baby.

➤ Taxonomy: BMW 320 is a kind of Car.

➤ Similarity: Swan looks similar to Goose.

# Never Ending Image Learner

- Extract Relationships
- Build co-occurrence matrix
- Get co-occurred object pairs
- Learn relationship in terms of mean and variance of relative position, aspect ratio, score, size.

# Never Ending Image Learner

- Object-Attribute Relationships

  - “Pizza has Round Shape”, “Sunflower is Yellow”

- Scene-Object Relationships

  - “Bus is found in Bus depot”

- Scene-Attribute Relationships

  - “Ocean is Blue”

# Never Ending Image Learner

➤ Discover new instance and retrain

$$\phi_i(x) + \sum_{i,j \in \mathcal{R}_O \cup \mathcal{R}_A} \phi_j(x_l) \psi_{i,j}(x, x_l) + \sum_{i,k \in \mathcal{R}_S} \omega_k(x)$$

object detector

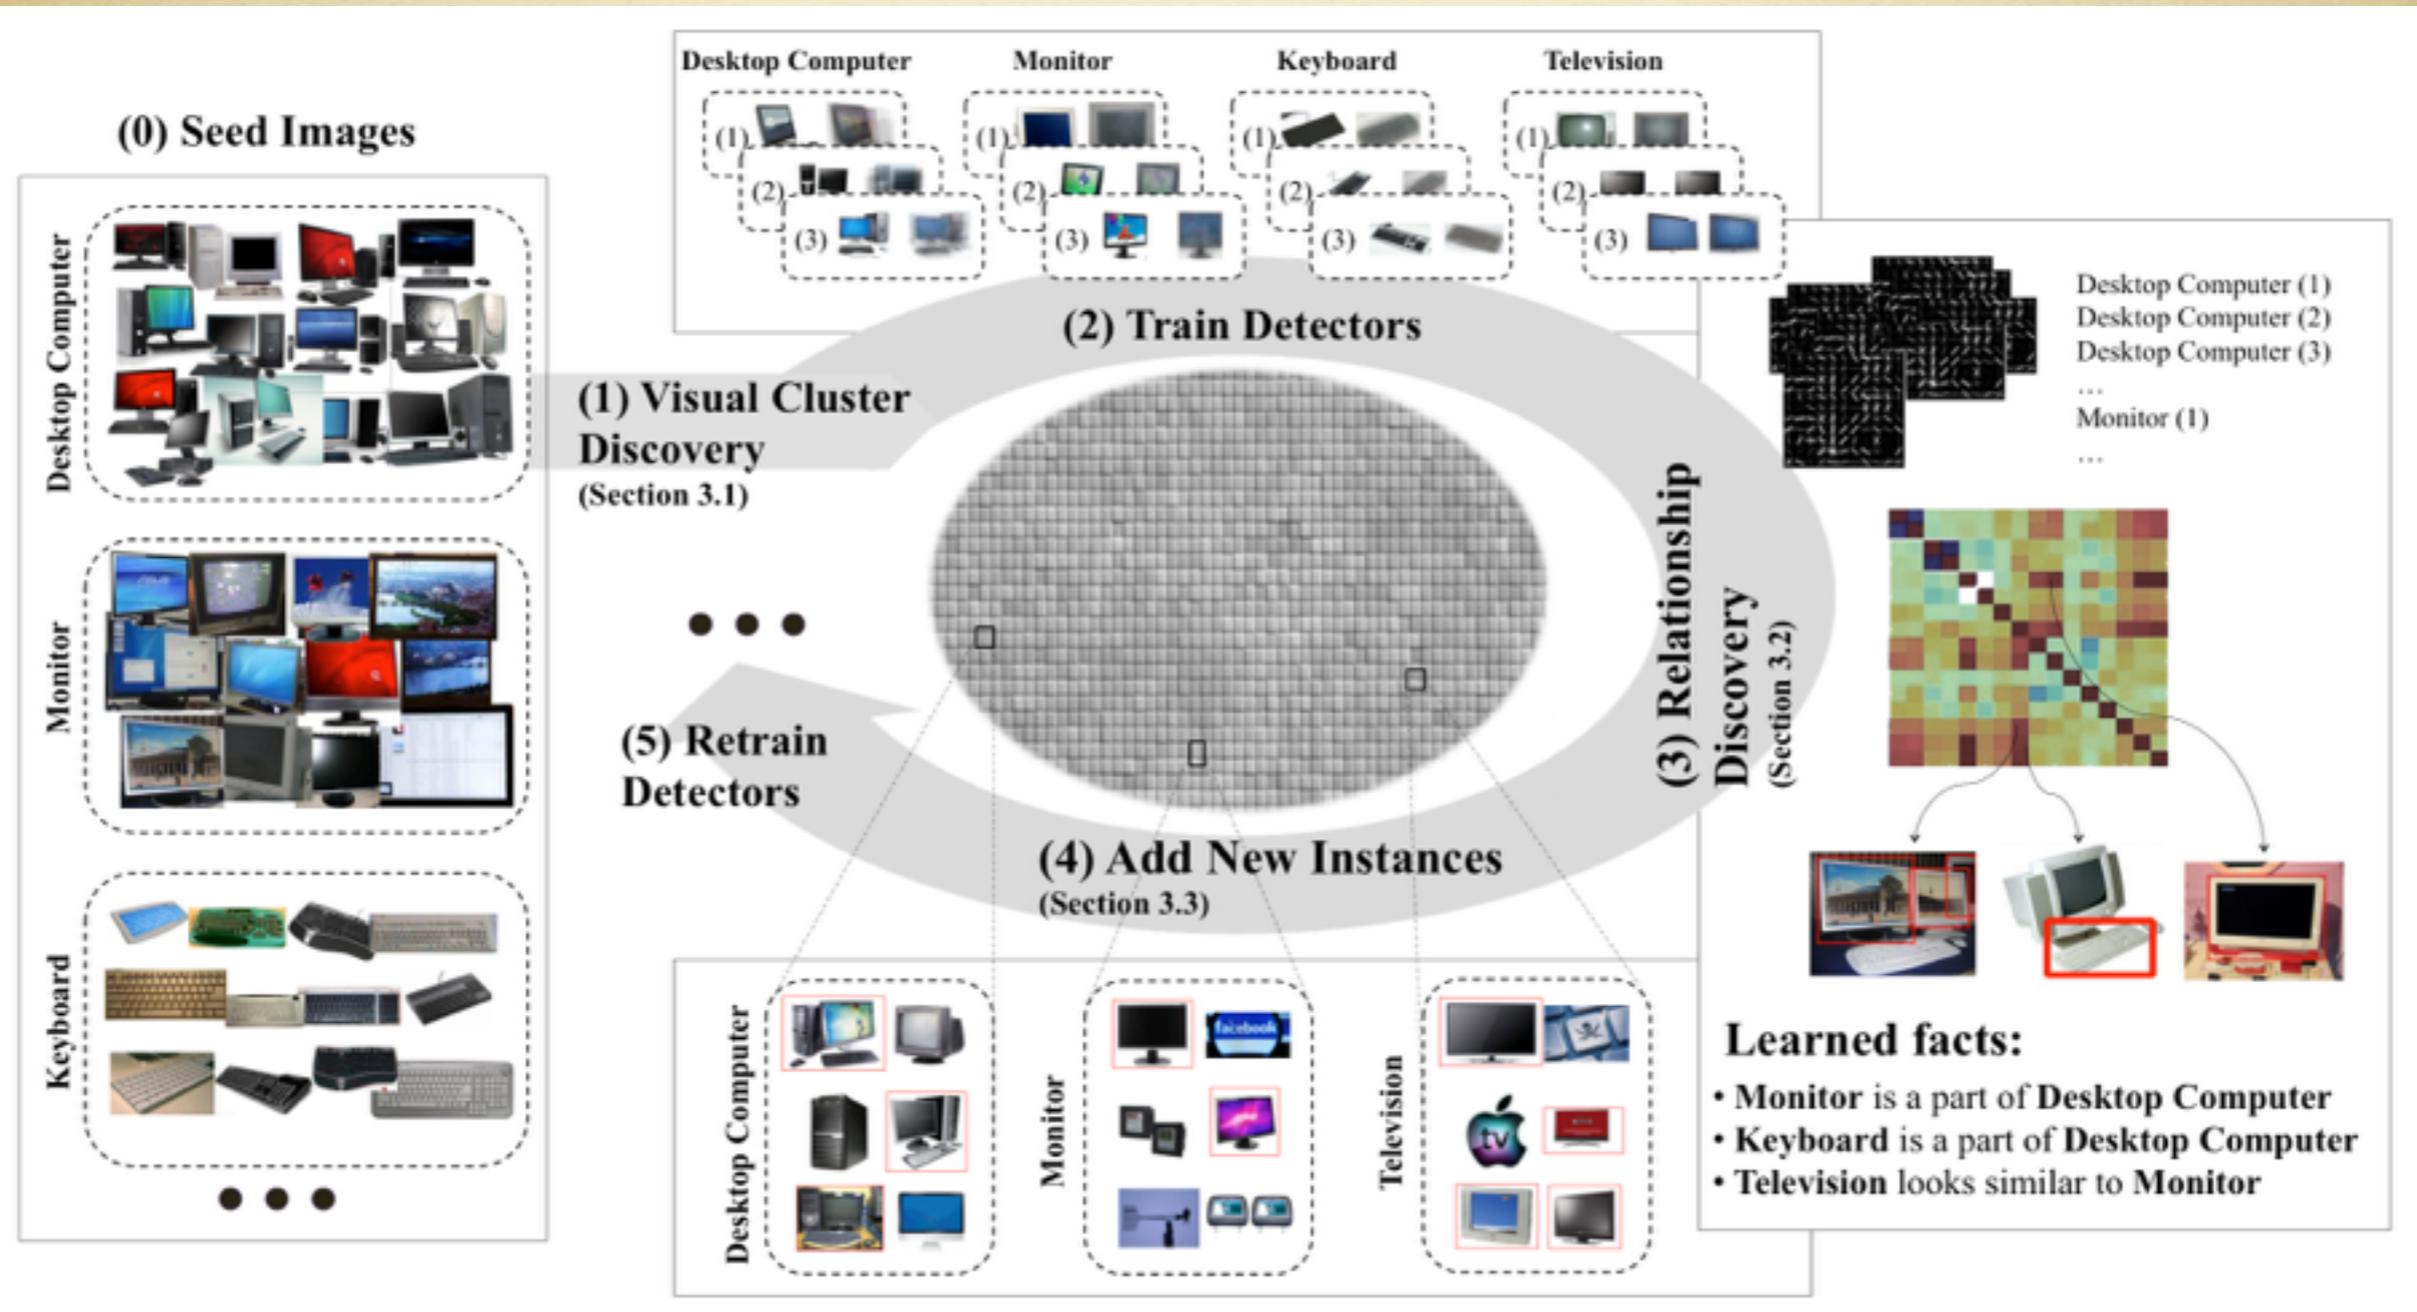
binary relationship

all related objects and attributes

scene classifier

all related scenes

# Never Ending Image Learner



# Never Ending Image Learner

➤ Bootstrapping

➤ Words: NELL (never ending language learning)

➤ Images: ImageNet, SUN, Google Image Search

Table 1. mAP performance for scene classification on 12 categories.

	mAP
Seed Classifier (15 Google Images)	0.52
Bootstrapping (without relationships)	0.54
NEIL Scene Classifiers	0.57
NEIL (Classifiers + Relationships)	<b>0.62</b>

Table 2. mAP performance for object detection on 15 categories.

	mAP
Latent SVM (50 Google Images)	0.34
Latent SVM (450 Google Images)	0.28
Latent SVM (450, Aspect Ratio Clustering)	0.30
Latent SVM (450, HOG-based Clustering)	0.33
Seed Detector (NEIL Clustering)	0.44
Bootstrapping (without relationships)	0.45
NEIL Detector	0.49
NEIL Detector + Relationships	<b>0.51</b>

# Hey, it's about time...

- to fix the annoying problem

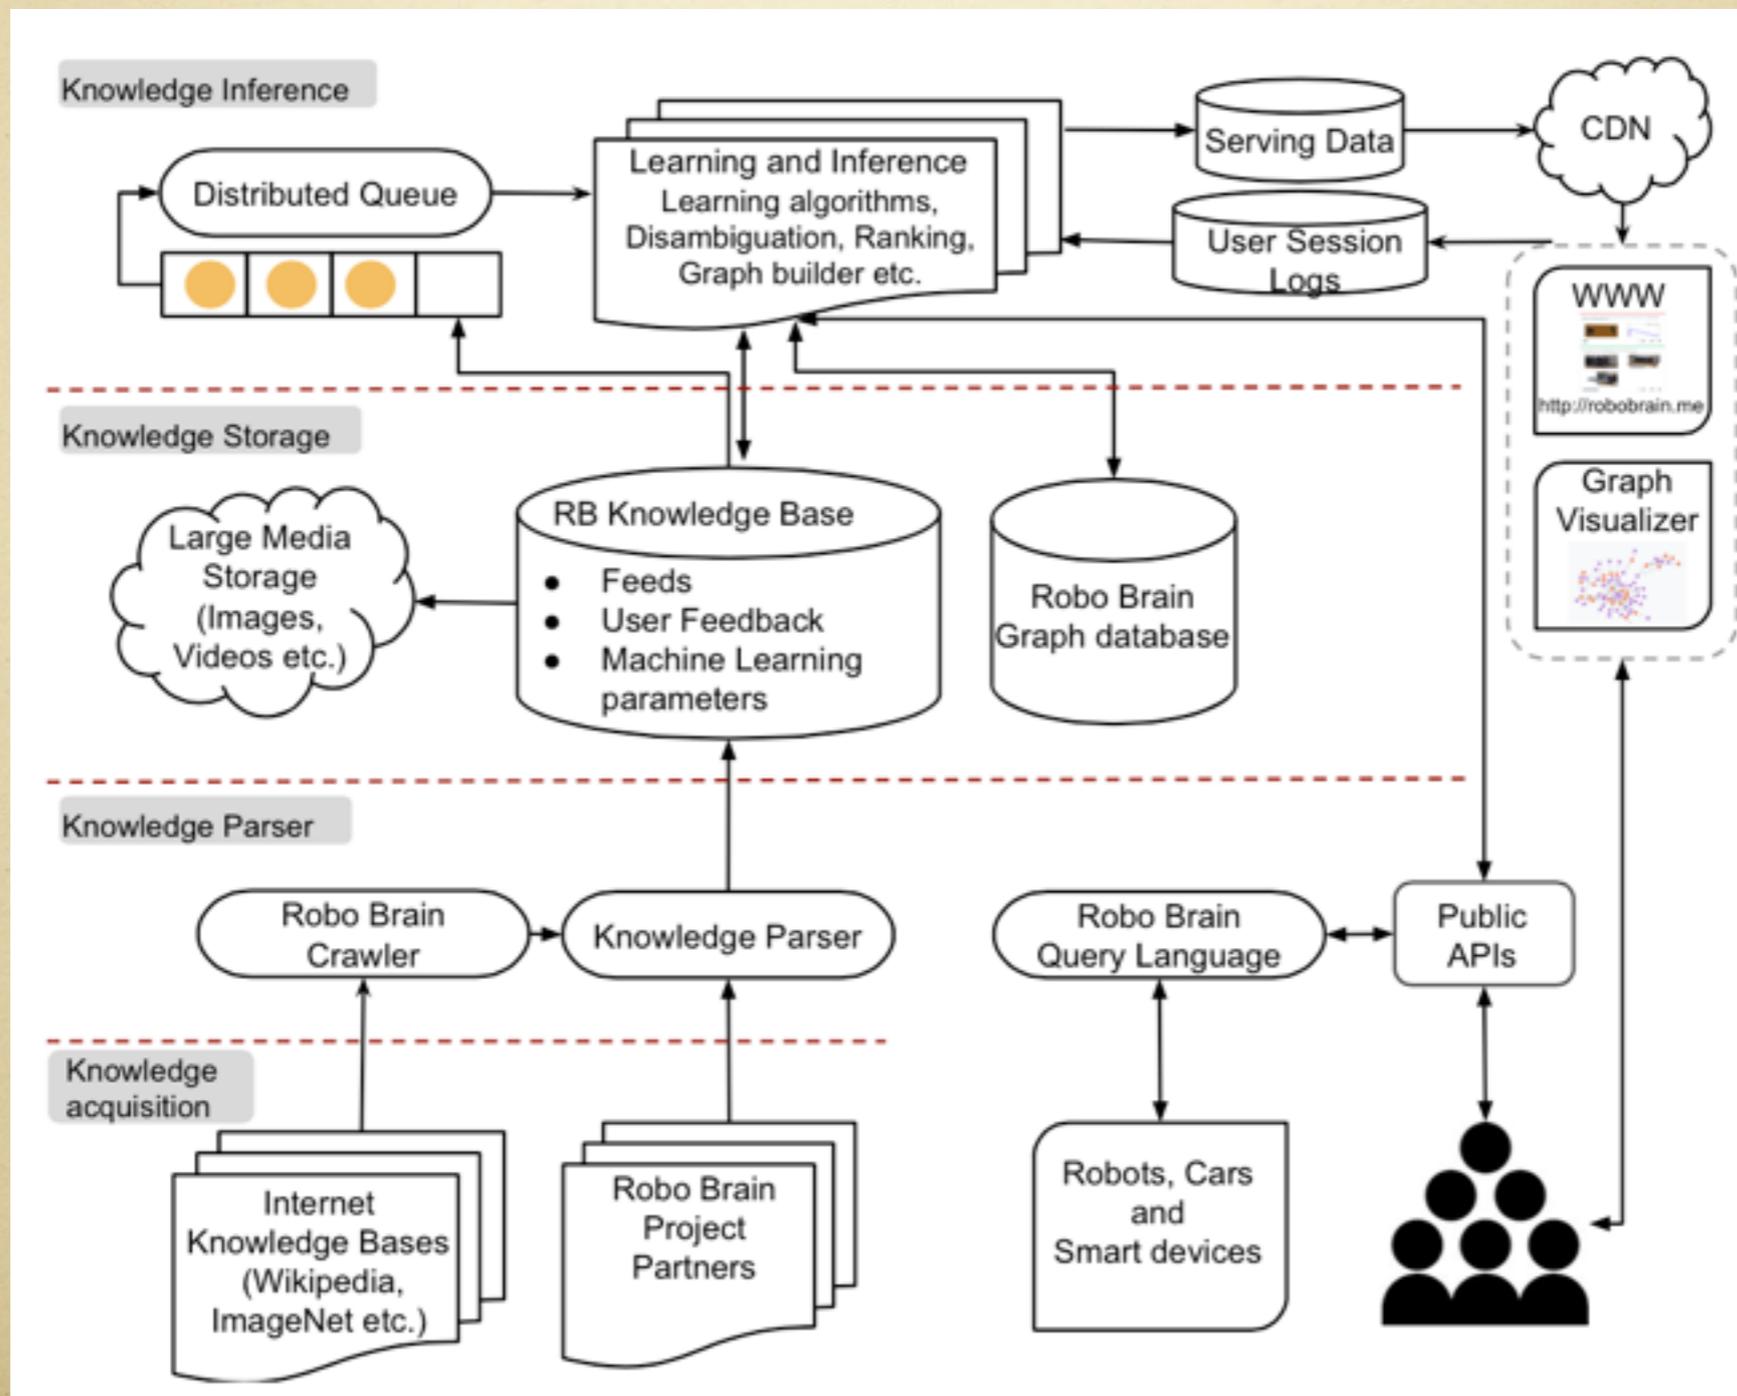
Bring me a cup of hot water

- Design a robot with knowledge base

# RoboBrain

- A large-scale knowledge engine for robot
- Build a knowledge base similar as ConceptNet
- More diverse edges
- Edges have beliefs
  - measure the confidence of learned relations
  - labelled by crowd-sourced feedback

# RoboBrain



# RoboBrain

- How to build knowledge base?
- again, graph represented in triplets

Word	An english word represented as an ASCII string.
DeepFeature	Feature function trained with a Deep Neural Network
Image	2D RGB Image.
PointCloud	3D point cloud

IsTypeOf	human <i>IsTypeOf</i> a mammal.
HasAppearance	floor <i>HasAppearance</i> as follows (this image).
CanPerformAction	human <i>CanPerformAction</i> cutting.
SpatiallyDistributedAs	location of human is <i>SpatiallyDistributedAs</i> .

- (StandingHuman, Shoe, *CanUse*)
- (Grasping, DeepFeature23, *UsesFeature*)
- (StandingHuman,  $\mathcal{N}(\mu, \Sigma)$ , *SpatiallyDistributedAs*)

# RoboBrain

➤ Knowledge acquisition

$$G = (V, E)$$

+

$$(v_1^1, v_2^1, \ell^1) \dots (v_1^N, v_2^N, \ell^N)$$

→

$$G' = (V', E')$$

Original Database

New Feeds

New Database

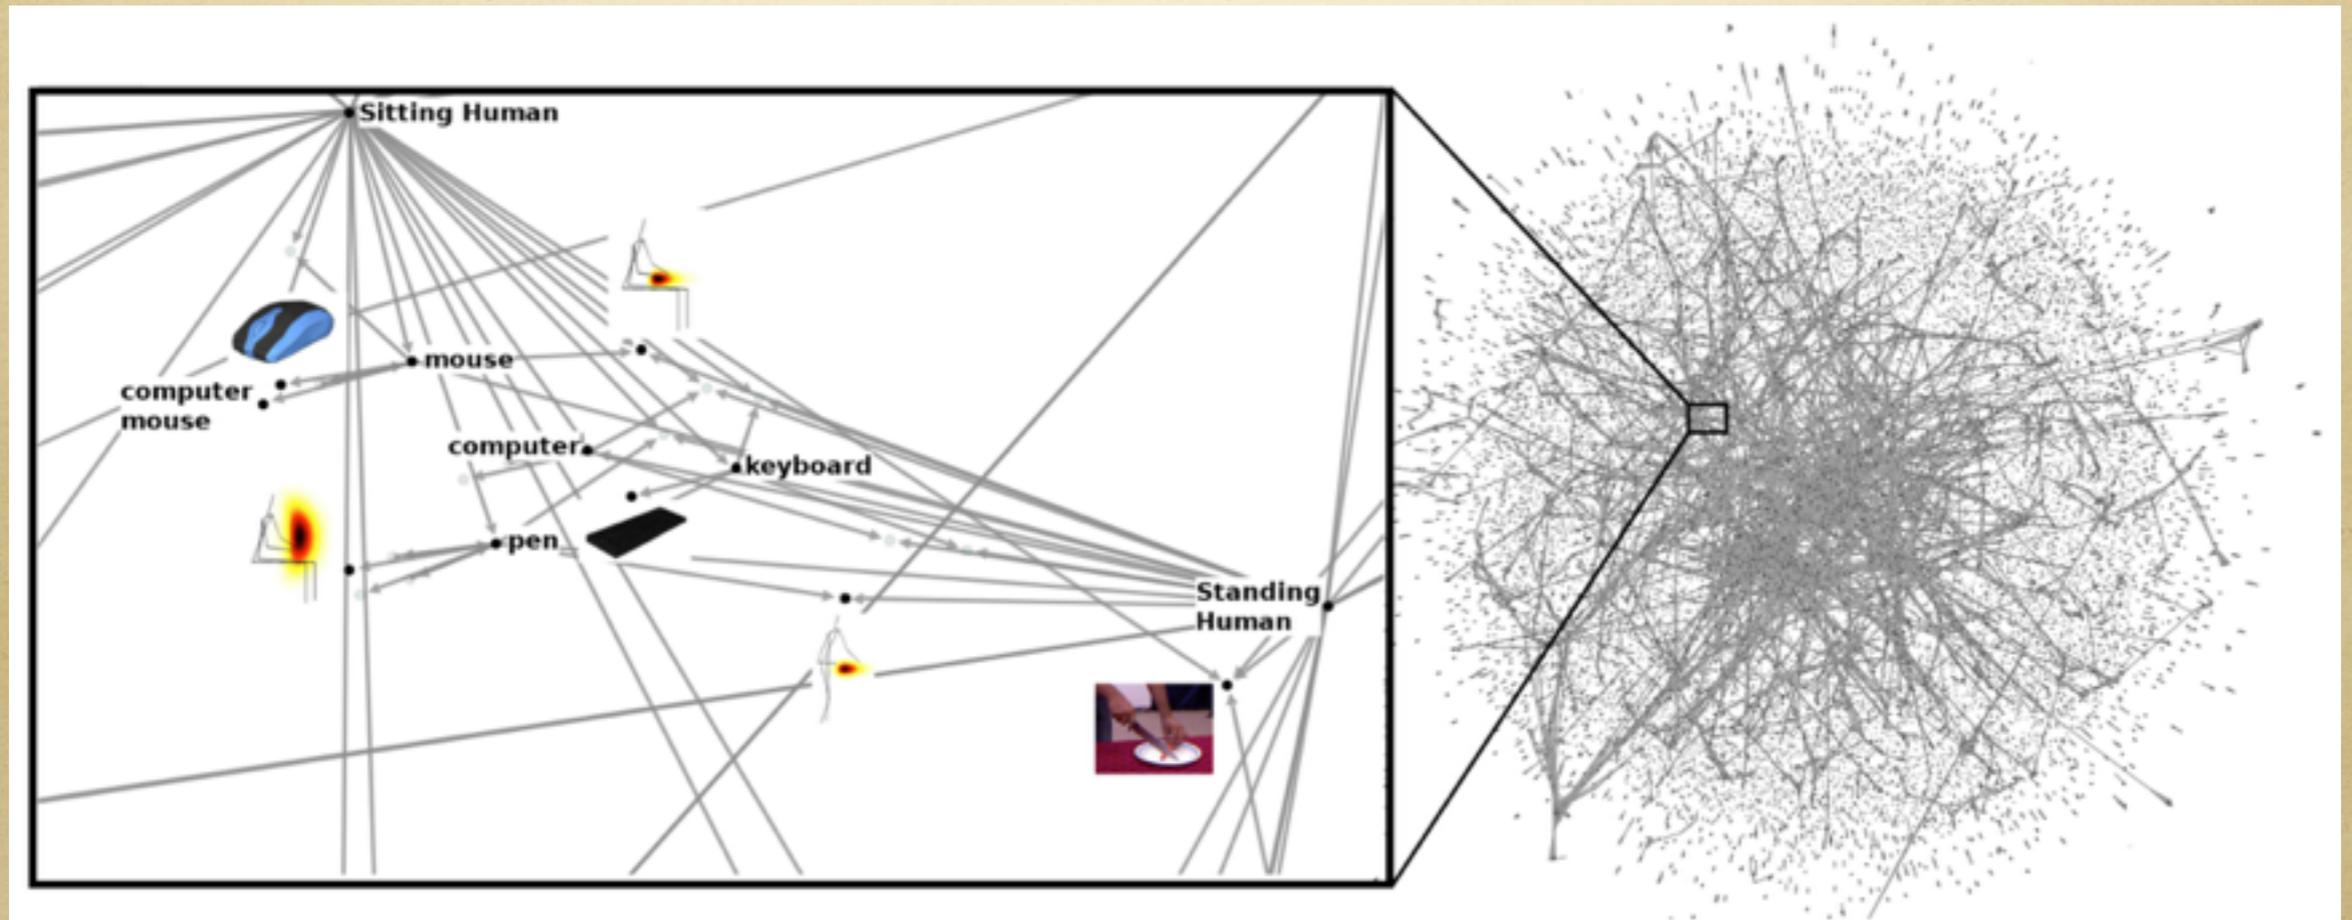
$$V' = v_1^1 \cup v_2^1 \cup \dots \cup v_1^N \cup v_2^N \cup V$$

$$E' = (v_1^1, v_2^1, \ell^1) \cup \dots \cup (v_1^N, v_2^N, \ell^N) \cup E$$



# RoboBrain

➤ Visualization of Knowledge Base



50K nodes, 100K edges

# RoboBrain

- Grounding a natural language sentence
  - “fill a cup with water”

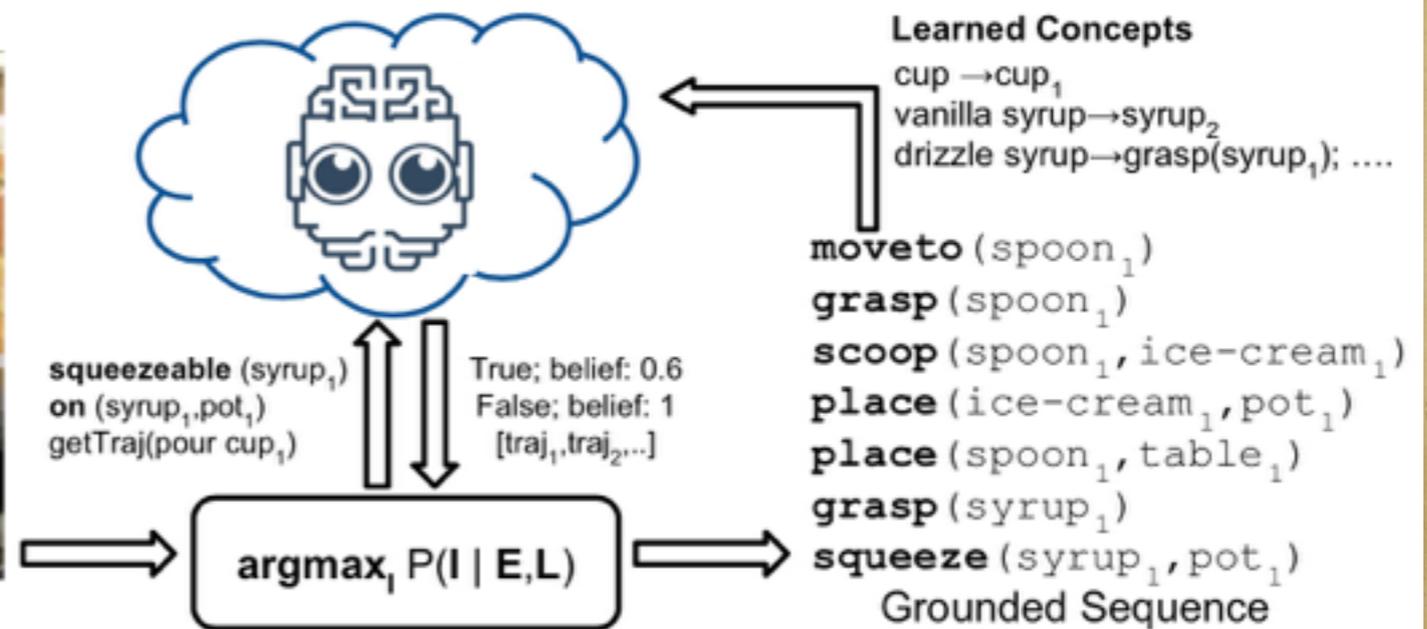
```
moveto(cup01); grasp(cup01); moveto(sink01); keep(cup01,  
on, sink01); toggle(sink_knob01); wait(); toggle(sink_knob01);
```

# RoboBrain

- Grounding a natural language sentence
- appearance, affordance, possible action, associated trajectory, manipulation feature



"add ice-cream to cup and drizzle syrup over it"



# RoboBrain

➤ Support action planning

*squeeze(syrup<sub>1</sub>,pot<sub>1</sub>)*

- *grasping pr2 syrup<sub>1</sub>*: Robot is grasping the syrup.
- *squeezeable syrup<sub>1</sub>*: Syrup bottle should be squeezeable.
- *on syrup<sub>1</sub> pot<sub>1</sub>*: Syrup bottle should be placed directly above the pot.

*squeezeable syrup<sub>1</sub>* = len fetch (u{name : 'syrup'}) →  
['HasAffordance'] → (v{name : 'squeezeable'}) > 0

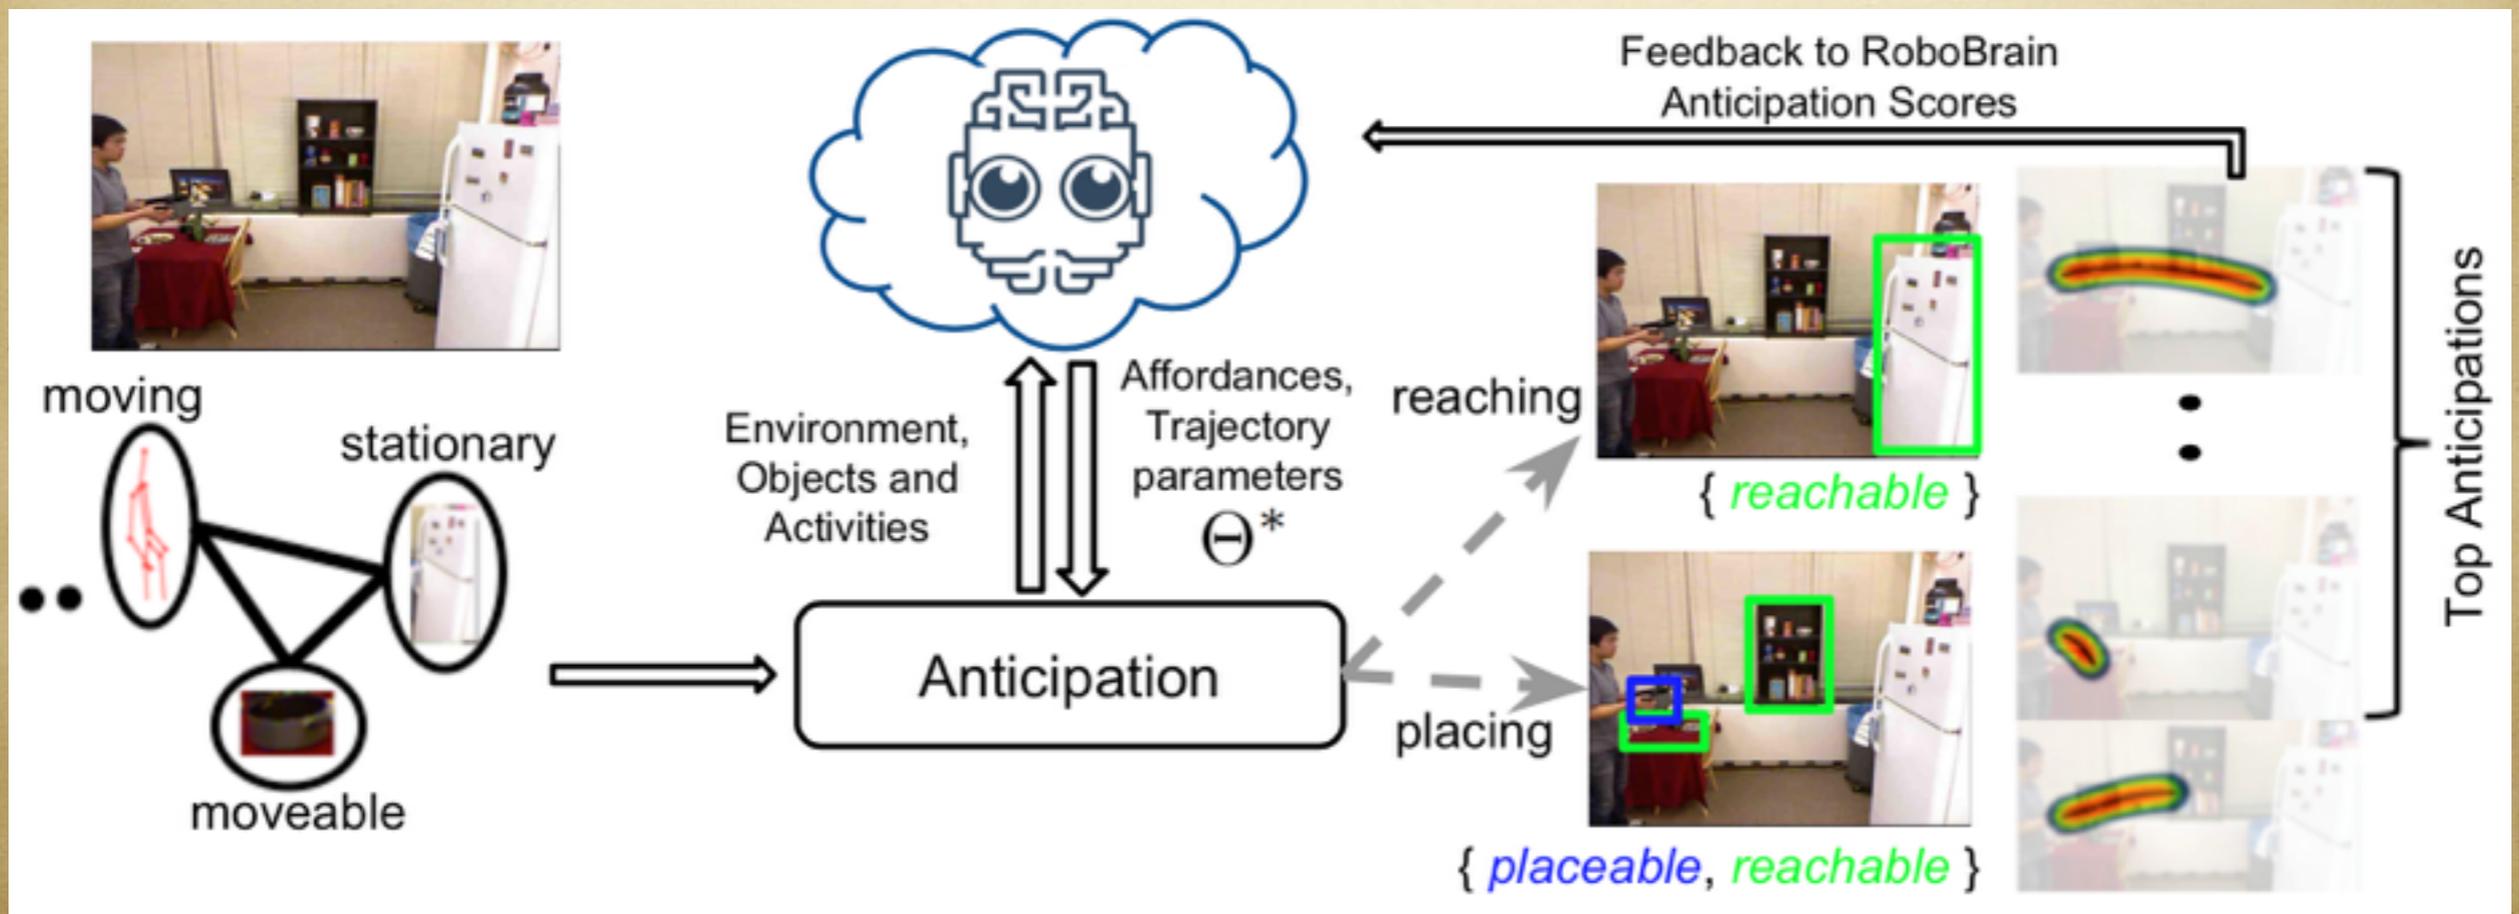
# RoboBrain

- Transfer action primitives to trajectory

```
pour cup01 = fetch ({name : 'pour'}) →  
(v{name : 'HasTrajectory'}) ← ({name : 'cup'})
```

# RoboBrain

- Other application
  - anticipating human activity



# RoboBrain

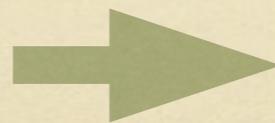
## ➤ Summary

- a knowledge base integrates knowledge about physical world that robots live in.
- share knowledge to support complicated tasks
  - natural language grounding
  - activity prediction

# Can we do more?

- So far, we know how to reuse learned knowledge.
- Can we generalize the learned knowledge to understand what we never seen before?

edible



# Zero-shot Affordance Prediction

➤ Idea

➤ affordance, attribute, human interaction are highly correlated

# Zero-shot Affordance Prediction

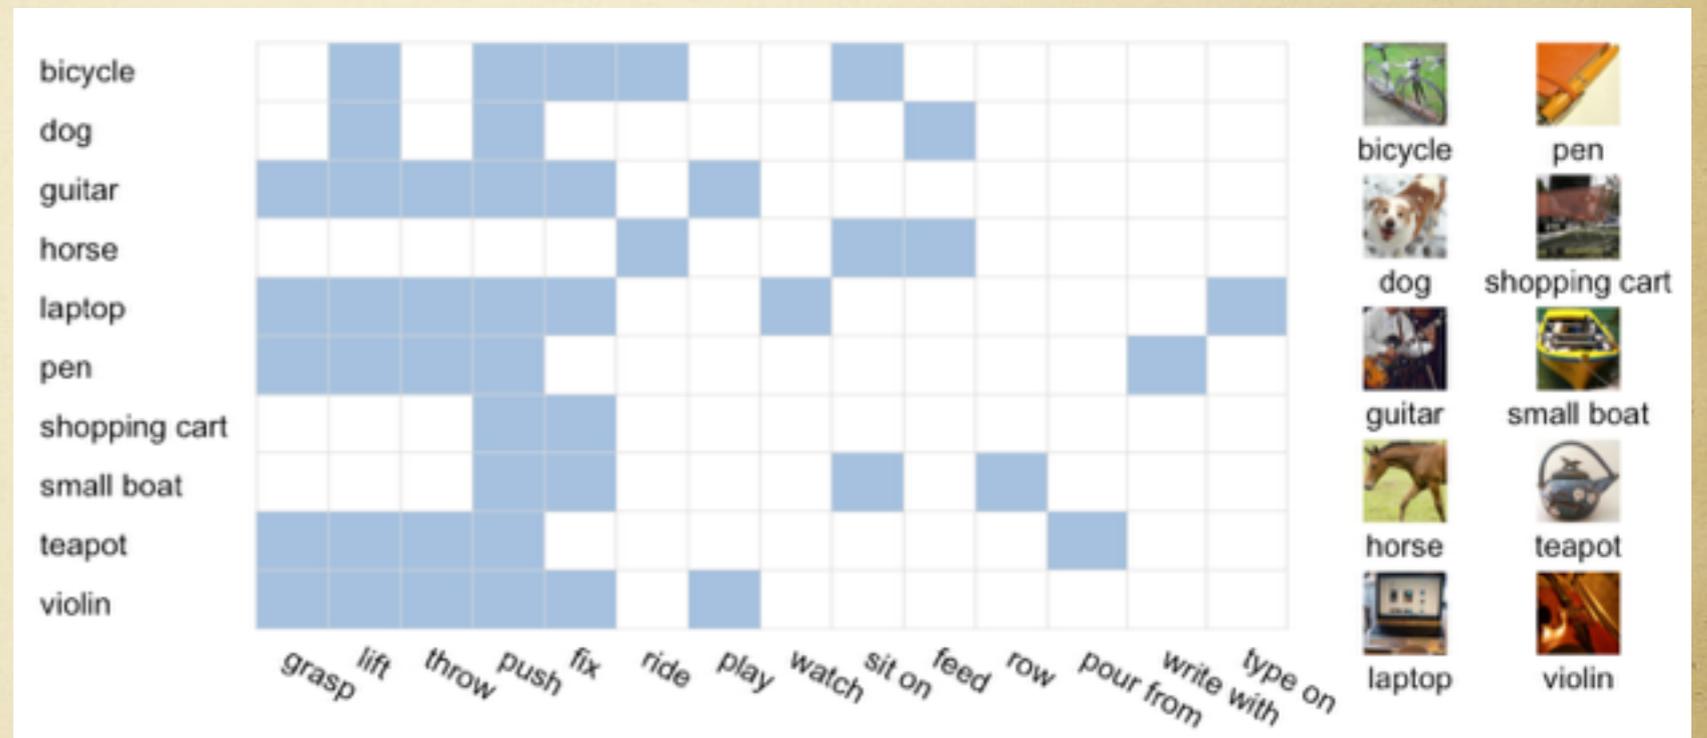
- Learning the knowledge base: choose 40 objects (Stanford 40 Action Database)
- Nodes (Entities):
  - Attribute:
    - visual: 33 per-trained classifiers, “round”, “shiny”
    - physical: weight, size, from FreeBase, Amazon
    - categorical: 22 from WordNet, “animal”, “vehicle”

# Zero-shot Affordance Prediction

➤ Nodes

➤ Attributes

➤ Affordance



➤ choose 14 from Stanford 40 Action

➤ manual labeling for 40 objects

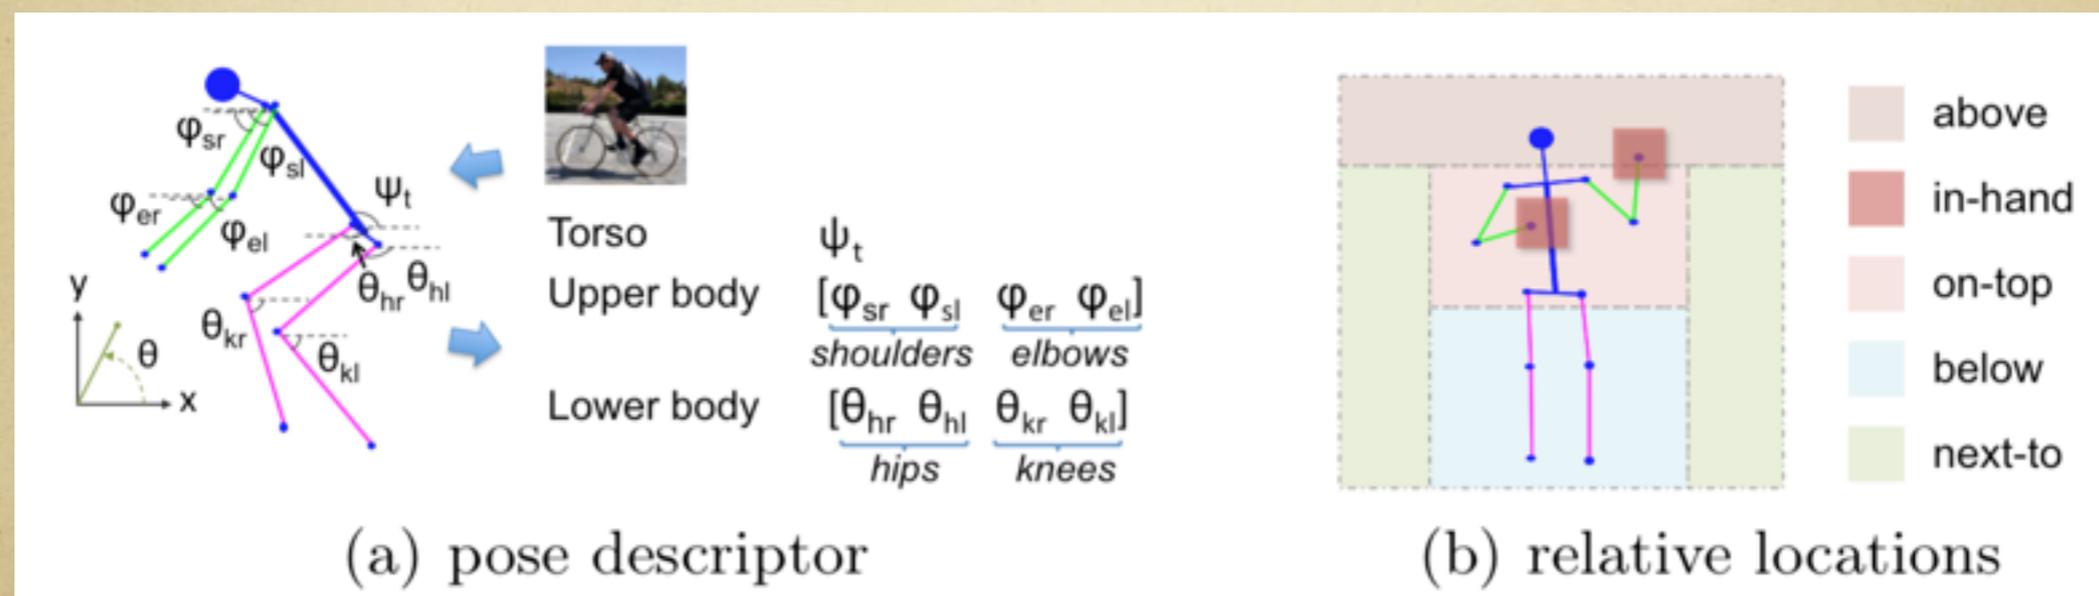
➤ on average, 4.25 per object

# Zero-shot Affordance Prediction

➤ Nodes:

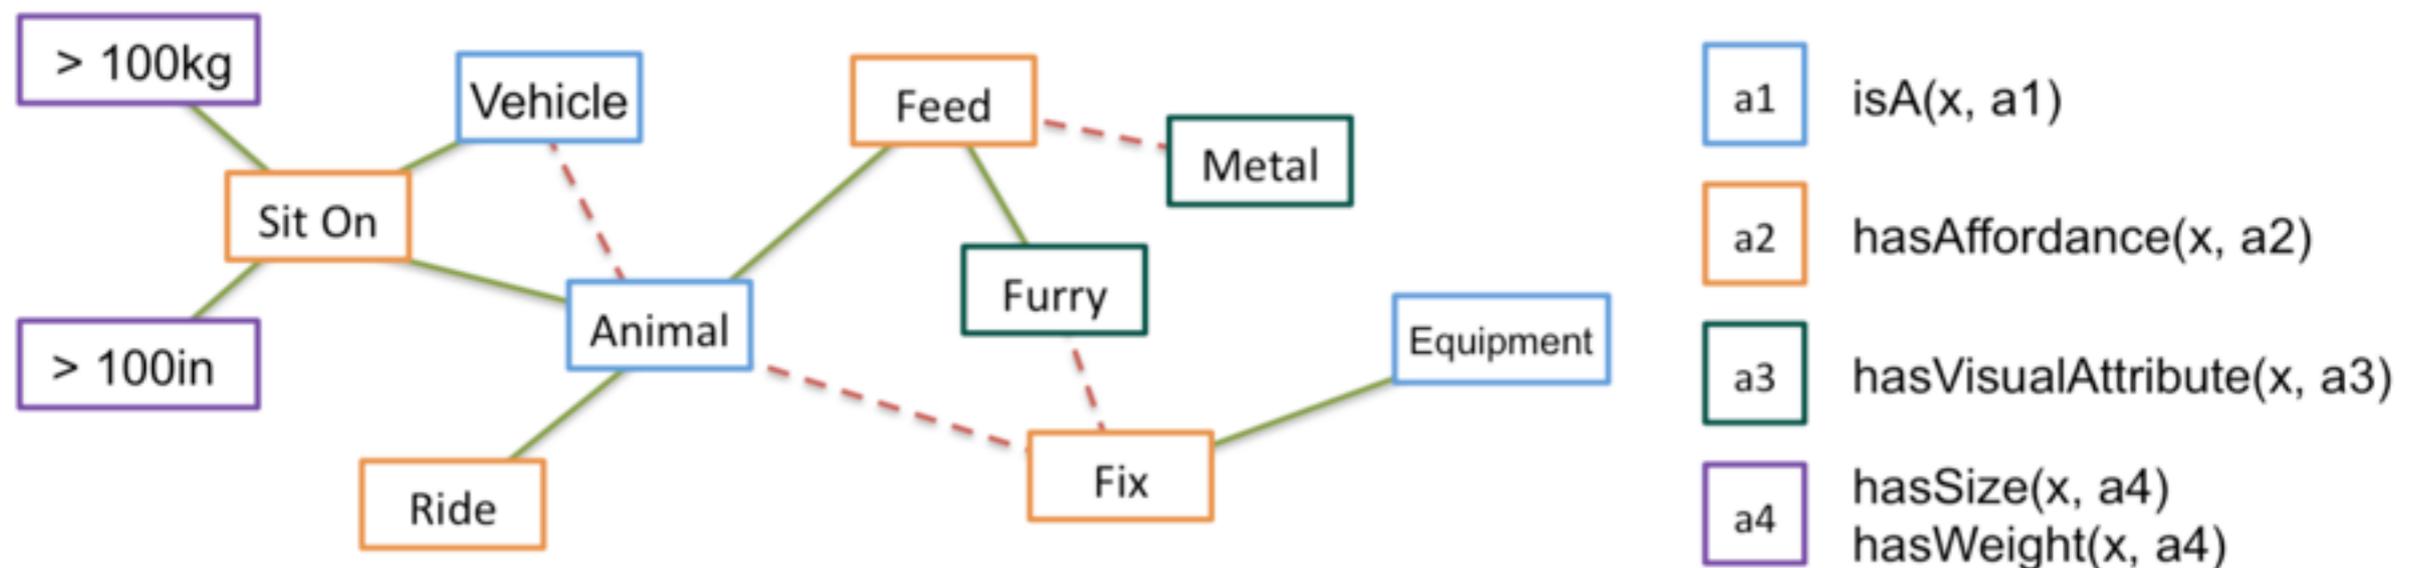
➤ Human pose: cluster centroids of descriptor.

➤ Human object relative position



# Zero-shot Affordance Prediction

- Learn a Markov Logic Network (MRF) to represent the relationships between nodes
- Use training data to build such relationships



# Zero-shot Affordance Prediction

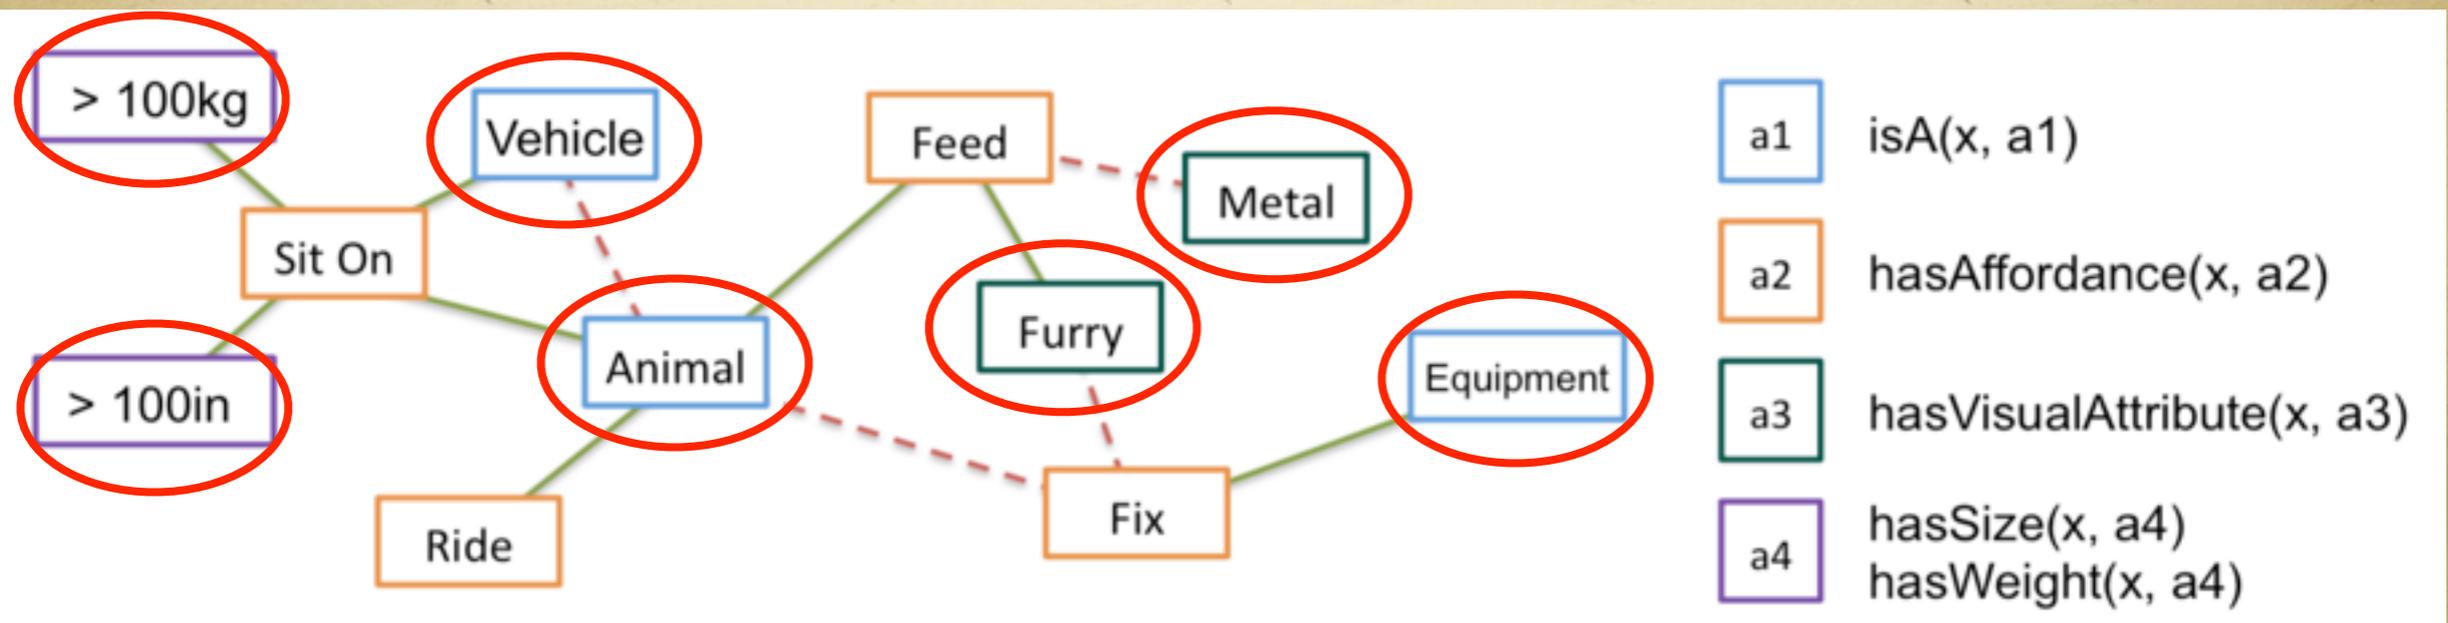
- Zero-shot prediction:
  - choose 22 objects that are semantically similar as the 40 training objects.
  - sample 50 images per objects as testing set.

# Zero-shot Affordance Prediction

- Zero-shot prediction:
  - Estimating visual attributes: run classifiers
  - Inferring:
    - Categorical attributes: learn regression from image feature and VA
    - Physical attributes: regression from image feature

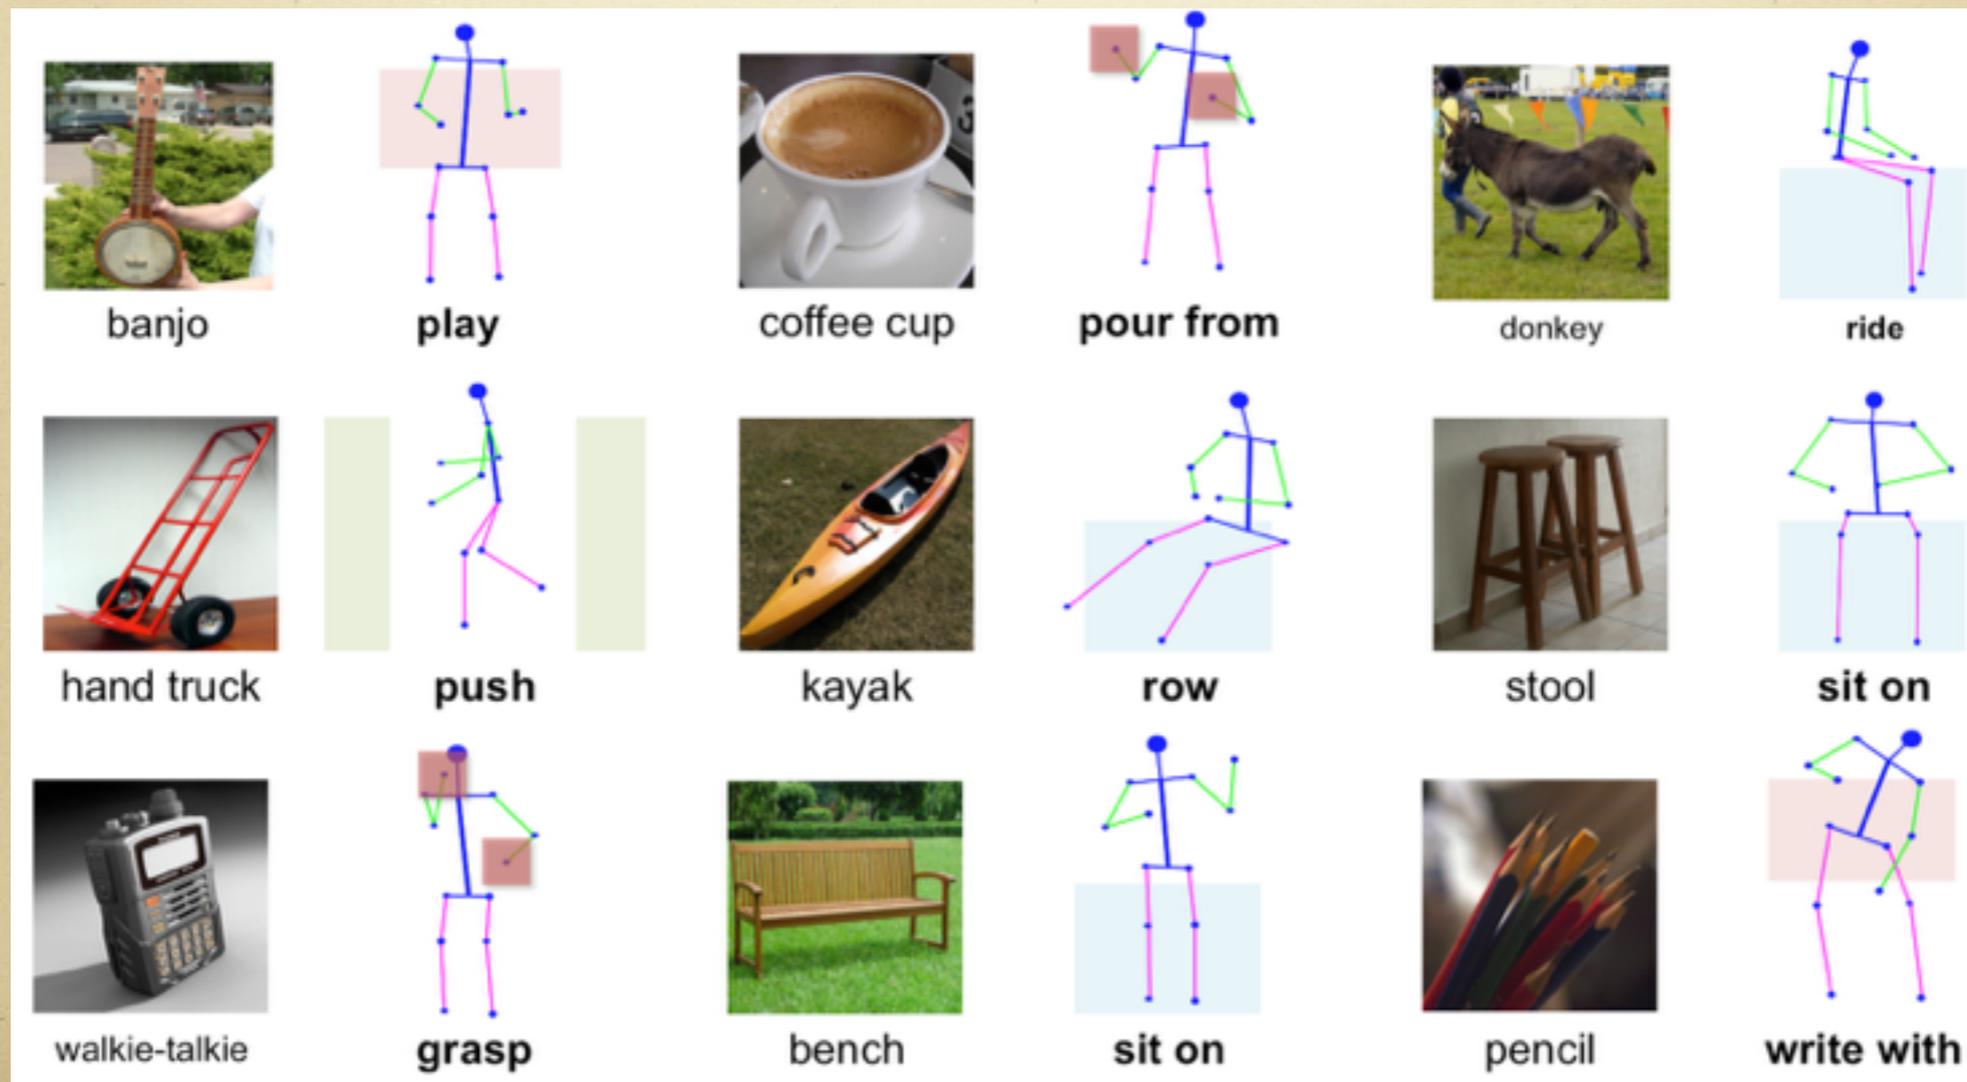
# Zero-shot Affordance Prediction

- Zero-shot prediction:
- Now, we have confidence on attribute nodes. Run belief propagation on MRF, we get confidence on affordance nodes.



# Zero-shot Affordance Prediction

➤ Zero-shot prediction:



# Zero-shot Affordance Prediction

➤ Zero-shot prediction:

**Table 1.** Performance of Zero-shot Affordance Prediction (measured in mAUC)

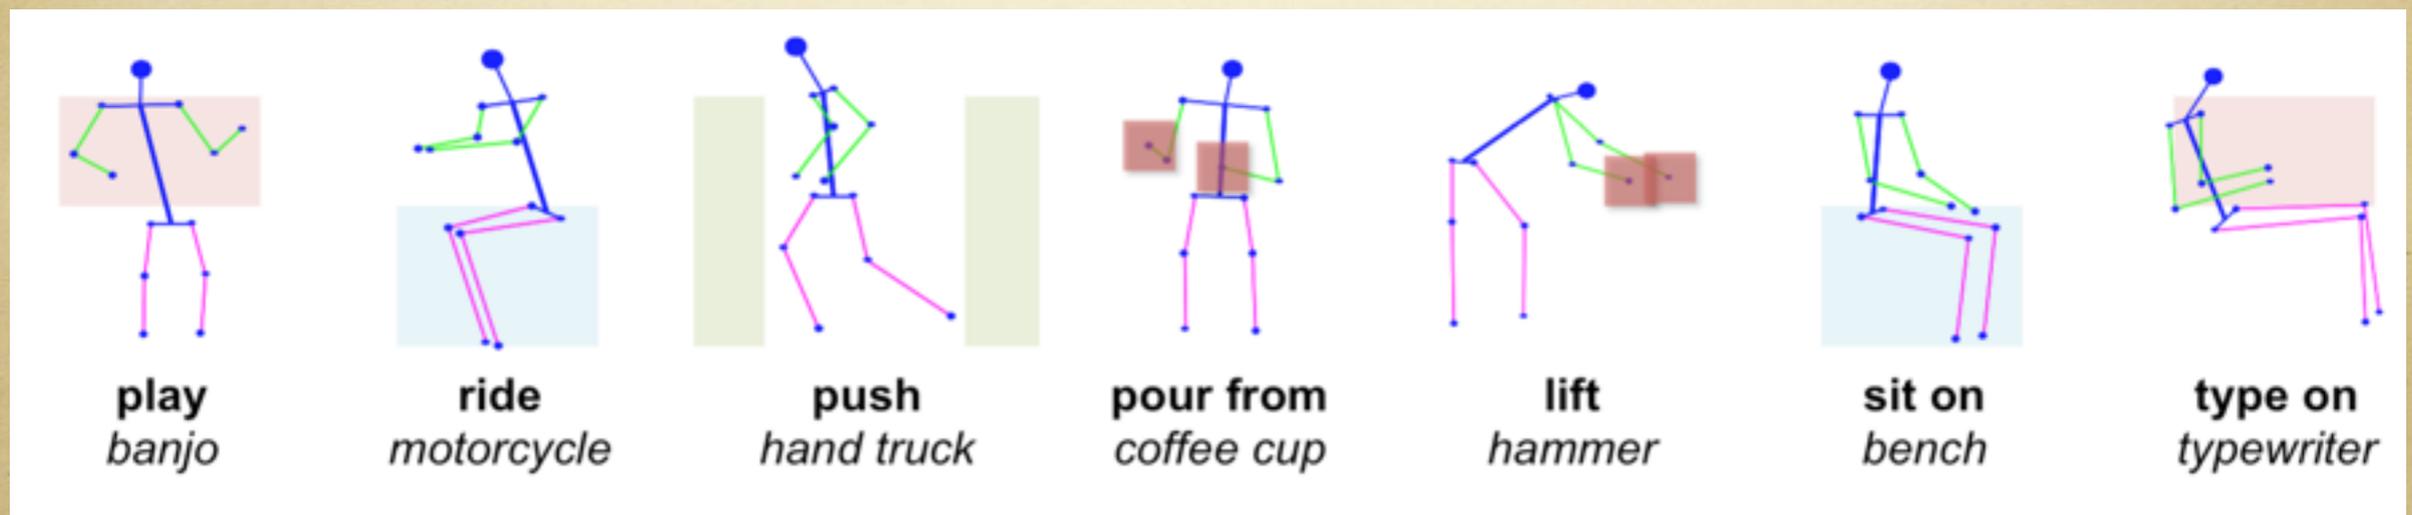
Method	L1-LR [9]	$\chi^2$ -SVM [25]	Ours
base features (BF)	0.7858	-	-
visual attributes (VA)	0.7525	0.7533	0.7432
categorical & physical (CP)	0.7919	0.7924	0.8234
combined (VA+CP)	0.8006	0.7985	<b>0.8409</b>

**Table 2.** Performance of Estimating Human Poses (in Hamming distance)

Method	nearest neighbor	attributes	affordances	attributes+affordances
Distance	0.928	1.027	0.630	<b>0.527</b>

# Zero-shot Affordance Prediction

➤ Prediction from human pose:



**Table 3.** Predicting Actions and Objects from Human-Object Interactions

Method	Action	Object
human poses	50.4%	46.2%
poses + locations	81.2%	64.5%

# Zero-shot Affordance Prediction

➤ Robust to partial observation:

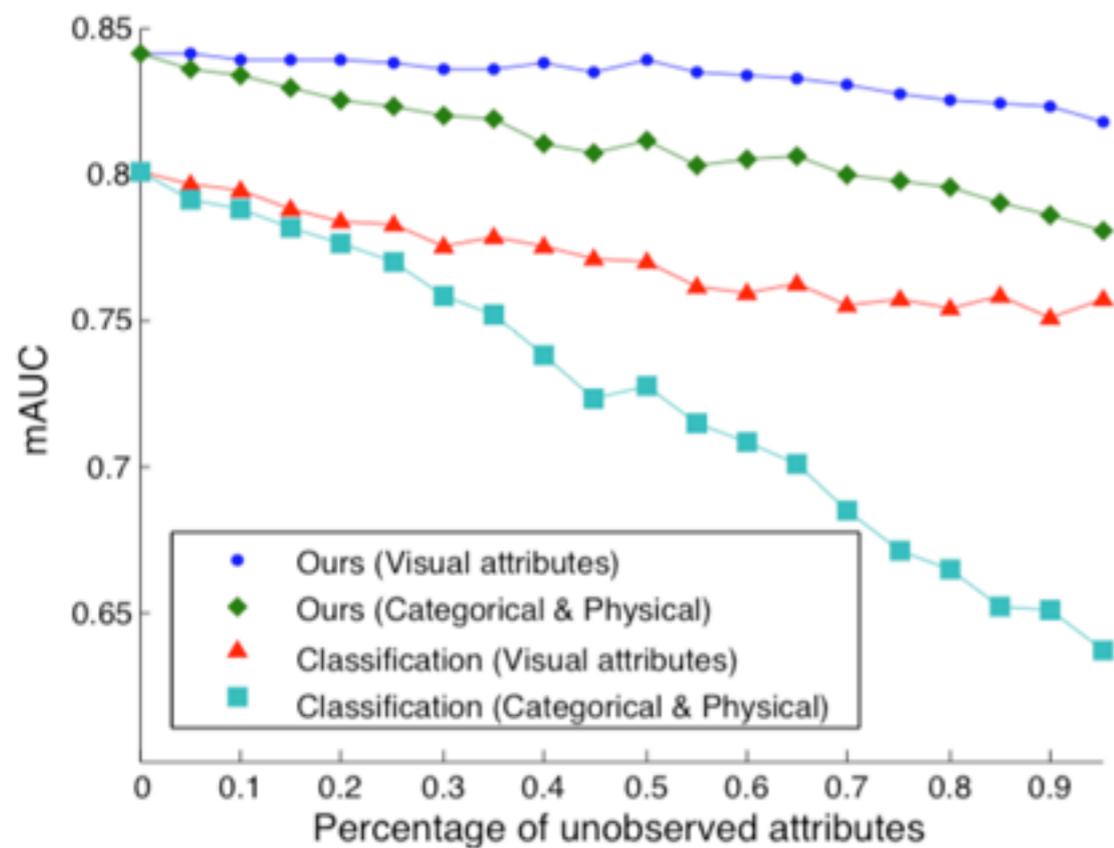


Fig. 12: Performance variations against partial observation. The  $x$ -axis denotes the percentage of unobserved evidence. The  $y$ -axis denotes the performance (mAUC). The top two curves correspond to our method. The bottom two are the classification-based method. In comparison, the knowledge base representation is more robust against partial observation.

# Zero-shot Affordance Prediction

## ➤ Question Answering:

Question	Evidence	Query	Top Answers
What do animals look like?	isA(N1, <b>Animal</b> )	hasVisualAttribute(N1, x)	hasVisualAttribute(N1, <b>Leather</b> ) hasVisualAttribute(N1, <b>Head</b> ) hasVisualAttribute(N1, <b>Tail</b> ) hasVisualAttribute(N1, <b>Furry</b> )
I saw something shiny and metallic. What is it?	hasVisualAttribute(N1, <b>Shiny</b> ) hasVisualAttribute(N1, <b>Metal</b> )	isA(N1, x)	isA(N1, <b>Instrumentality</b> ) isA(N1, <b>Device</b> ) isA(N1, <b>Container</b> ) isA(N1, <b>Computer</b> )
Here is a vehicle and it's quite heavy. What can I do with it?	isA(N1, <b>Vehicle</b> ) hasWeight(N1, W4) (> 100 kg)	hasAffordance(N1, x)	hasAffordance(N1, <b>Ride</b> ) hasAffordance(N1, <b>Row</b> ) hasAffordance(N1, <b>SitOn</b> ) hasAffordance(N1, <b>Fix</b> )
Tell me how heavy and large a wooden musical instrument is.	isA(N1, <b>Musical_instrument</b> ) hasVisualAttribute(N1, <b>Wood</b> )	hasWeight(N1, x) hasSize(N1, x)	hasSize(N1, D2) <b>(10-100 in)</b> hasWeight(N1, W2) <b>(1-10 kg)</b>

# Summary

- Online knowledge base
  - high-level: DBpedia, Wikidata
  - low-level: ConceptNet
- How to learn visual knowledge base: NEIL
- How to create KB for robot to do complicated tasks: RoboBrain
- How to generalize KB: zero-shot affordance prediction

