

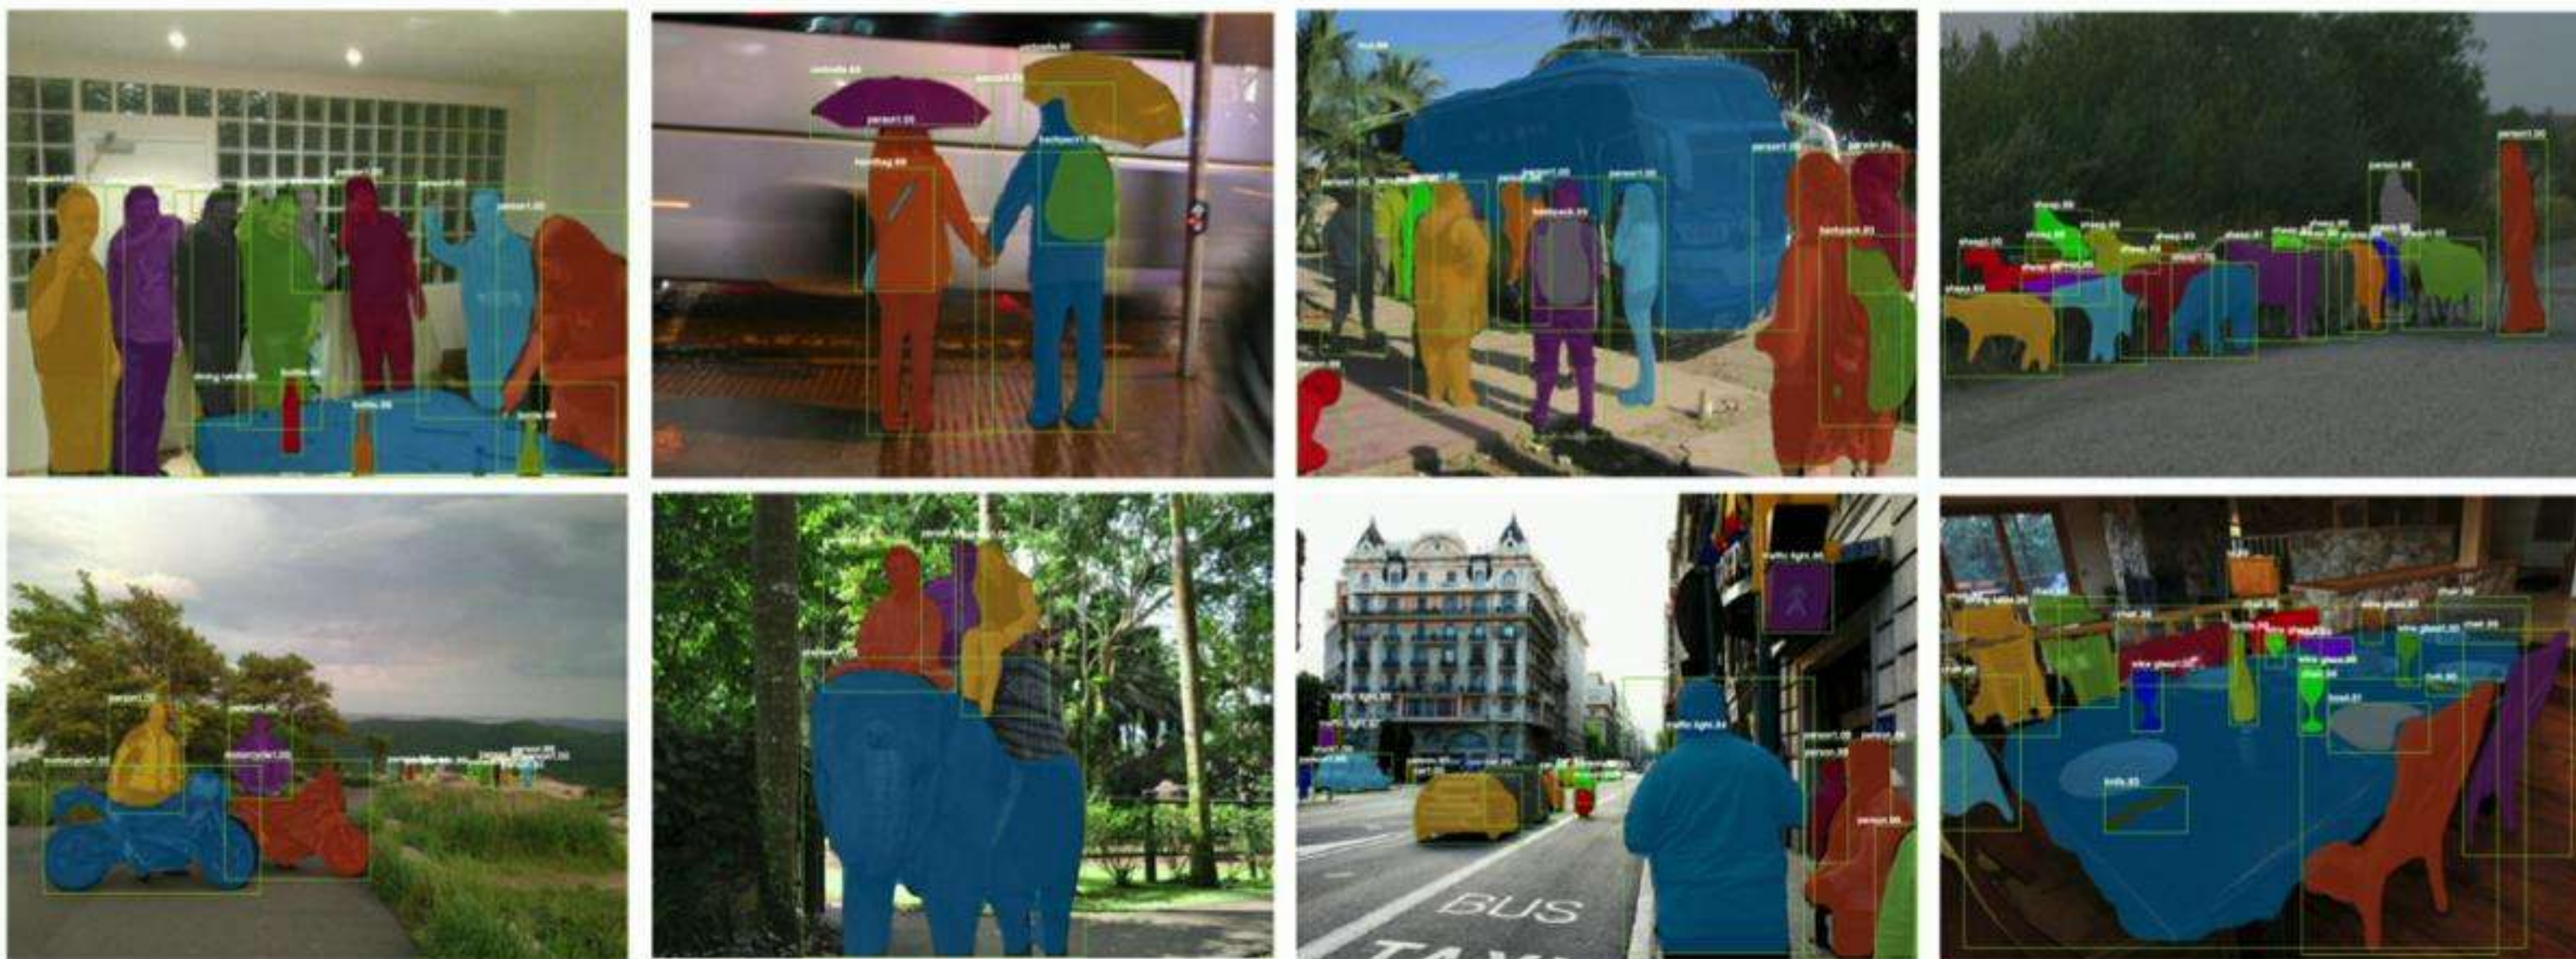
Structured Visual Understanding and Interaction with Human and Environment



Jianwei Yang
09/12/2019

The world around us is highly structured

Images are highly structured



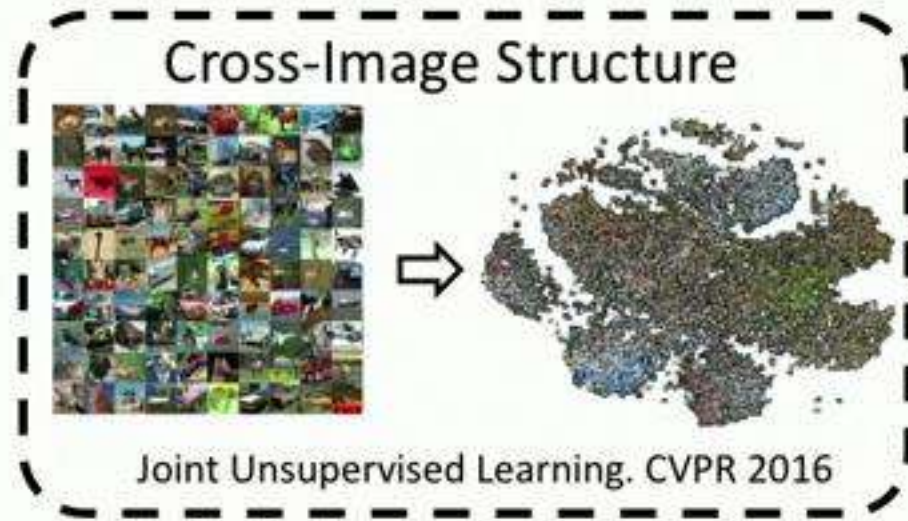
Microsoft COCO: Common Objects in Context. Lin et al. 2014

Images are highly structured



COCO-Stuff: Thing and Stuff Classes in Context. Caesar et al. 2018

My researches:



My researches:

Cross-Image Structure

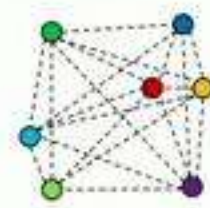


Joint Unsupervised Learning. CVPR 2016

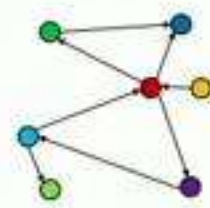
Per-Image Structure



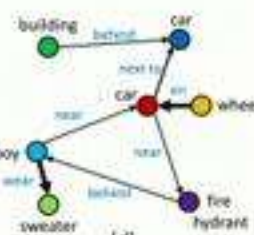
(a)



(b)



(c)



(d)

Graph R-CNN for Scene Graph Generation. ECCV 2018

My researches:

Cross-Image Structure

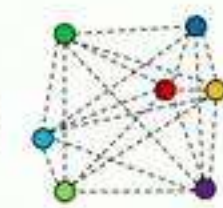


Joint Unsupervised Learning. CVPR 2016

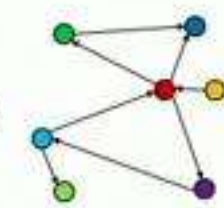
Per-Image Structure



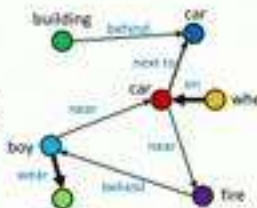
(a)



(b)



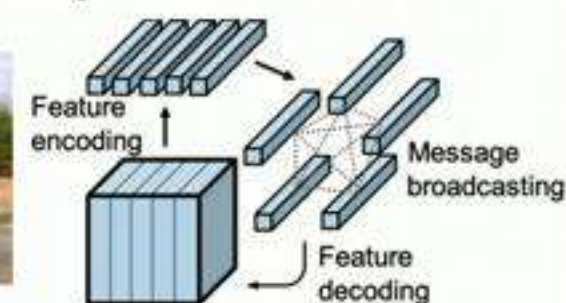
(c)



(d)

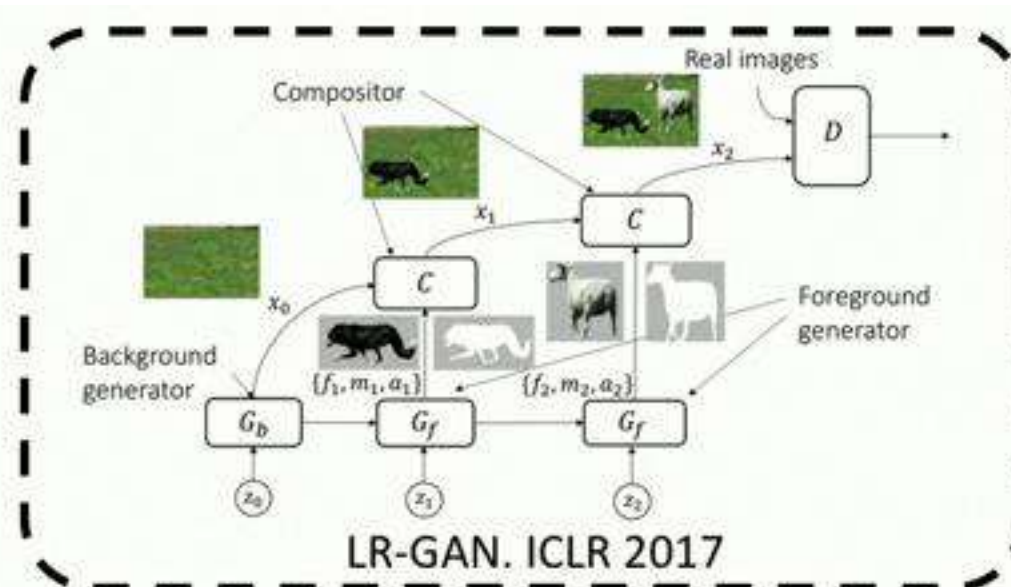
Graph R-CNN for Scene Graph Generation. ECCV 2018

Per-Object Structure



Neuron Communication Networks. NeurIPS 2019

My researches:



Cross-Image Structure

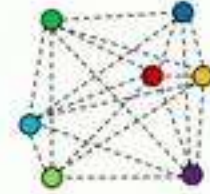


Joint Unsupervised Learning. CVPR 2016

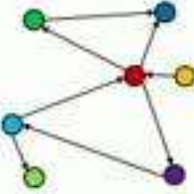
Per-Image Structure



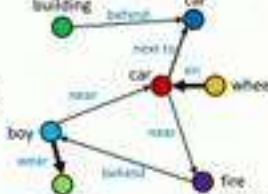
(a)



(b)



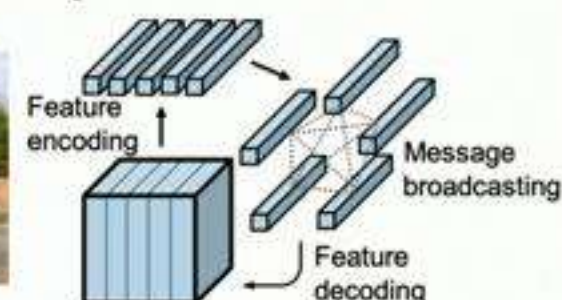
(c)



(d)

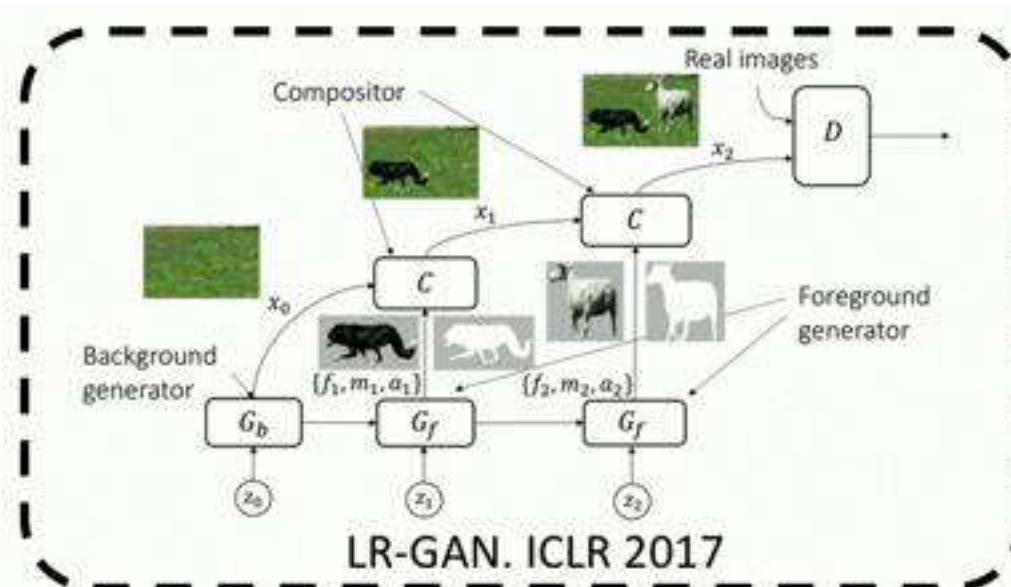
Graph R-CNN for Scene Graph Generation. ECCV 2018

Per-Object Structure



Neuron Communication Networks. NeurIPS 2019

My researches:

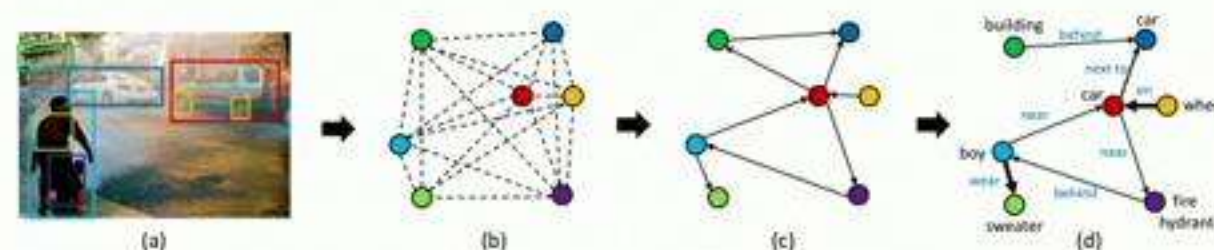


Cross-Image Structure



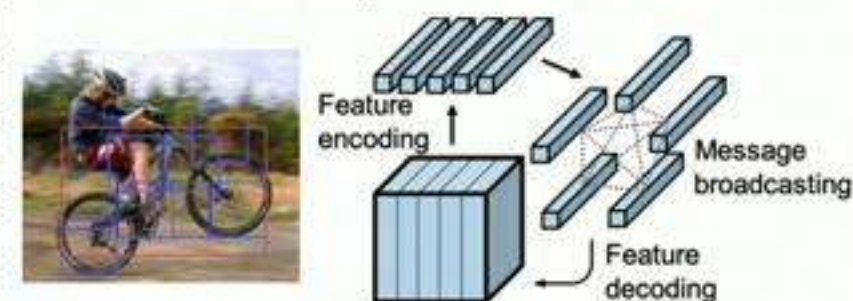
Joint Unsupervised Learning. CVPR 2016

Per-Image Structure

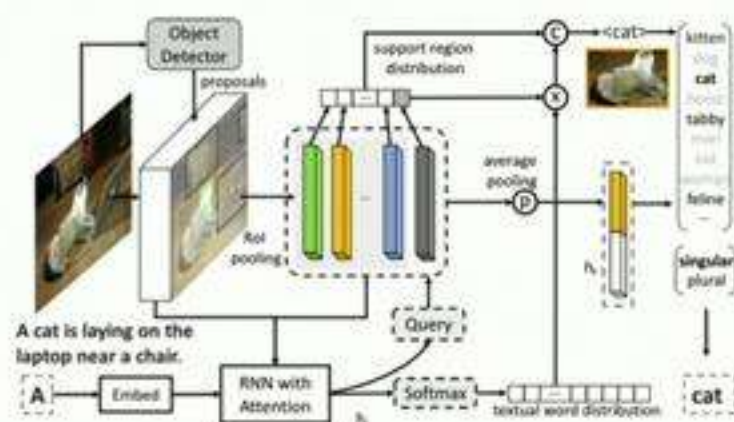


Graph R-CNN for Scene Graph Generation. ECCV 2018

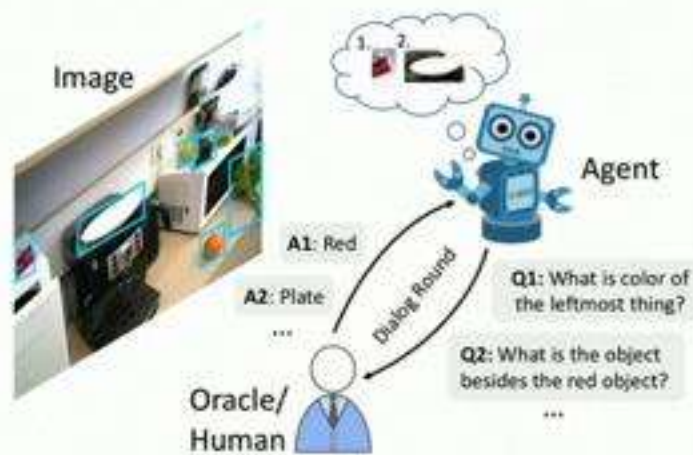
Per-Object Structure



Neuron Communication Networks. NeurIPS 2019



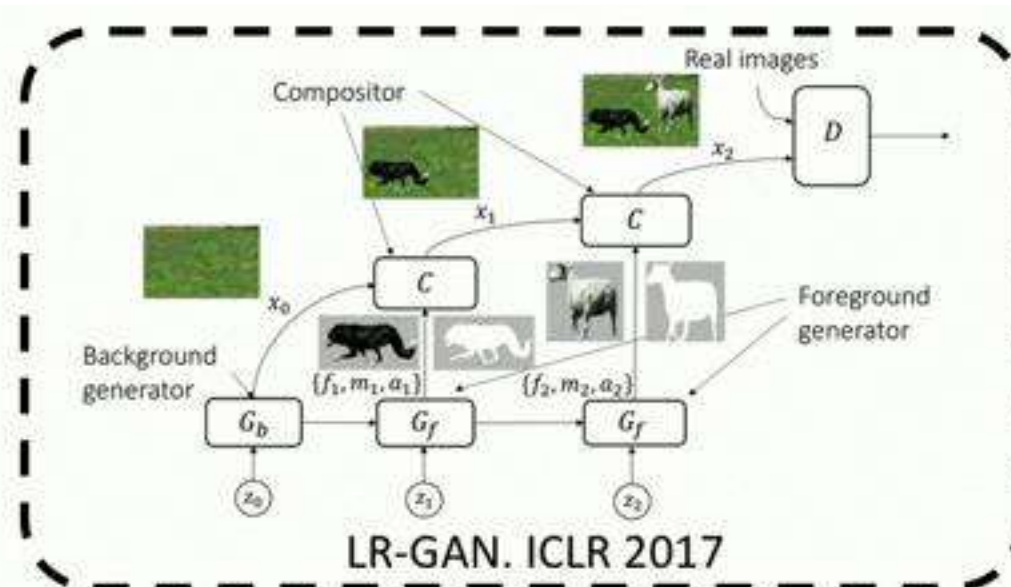
Neural Baby Talk. CVPR 2018



Visual Curiosity. CoRL 2018

Language

My researches:

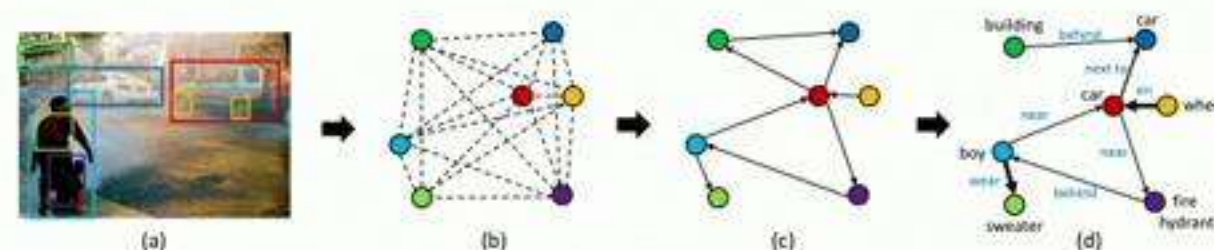


Cross-Image Structure



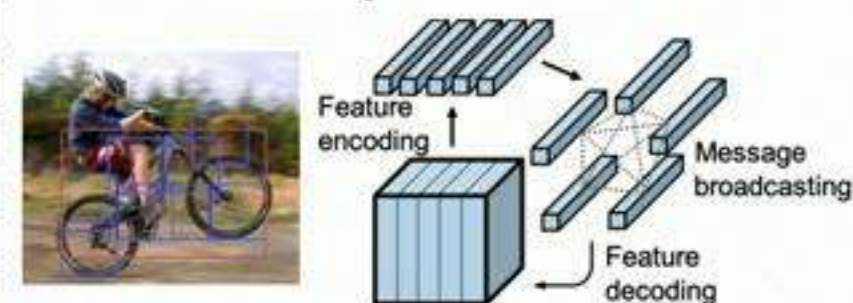
Joint Unsupervised Learning. CVPR 2016

Per-Image Structure

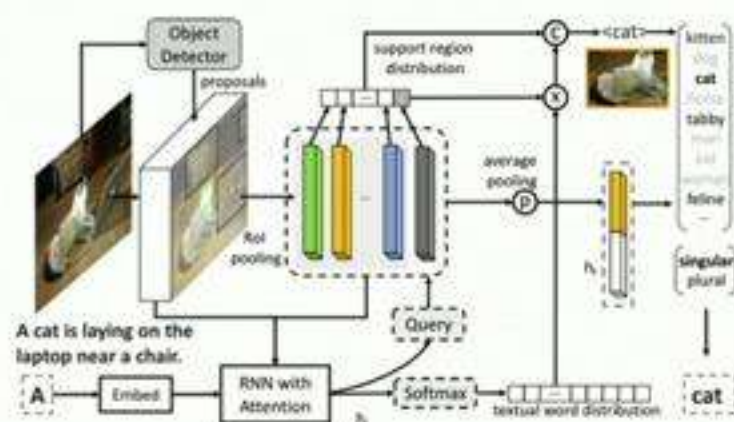


Graph R-CNN for Scene Graph Generation. ECCV 2018

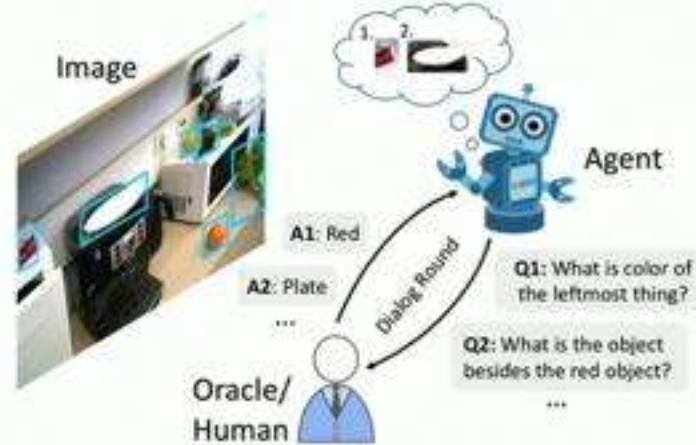
Per-Object Structure



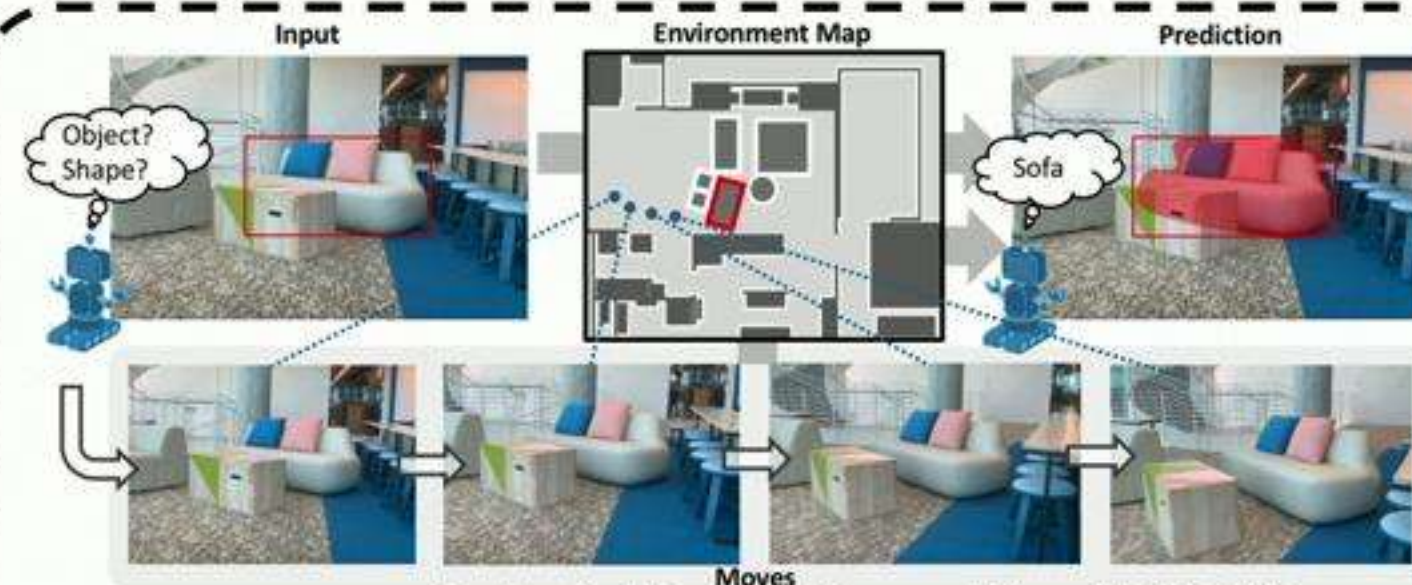
Neuron Communication Networks. NeurIPS 2019



Neural Baby Talk. CVPR 2018



Visual Curiosity. CoRL 2018

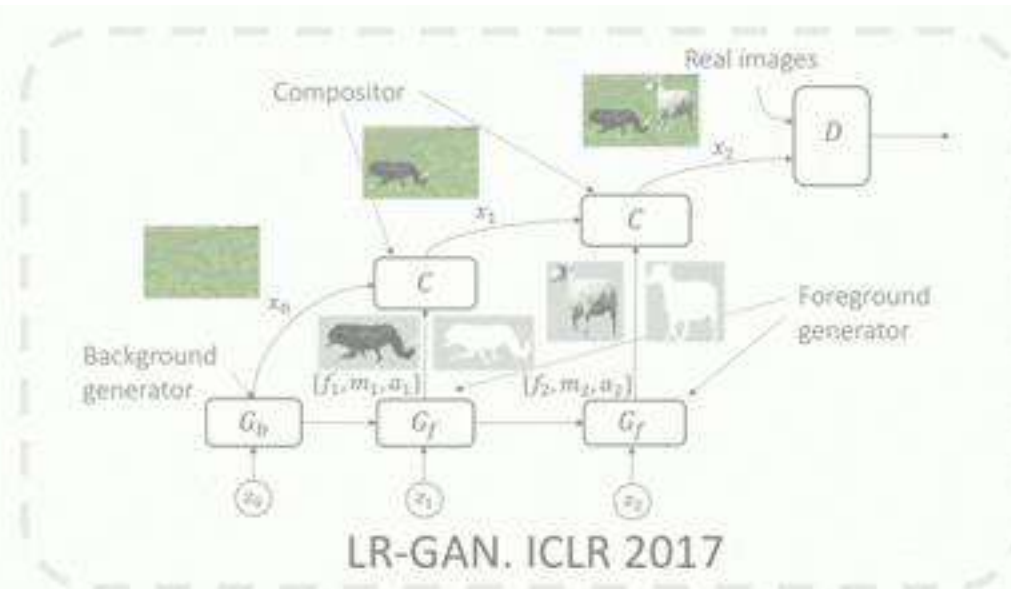


Embodied Amodal Recognition. ICCV 2019

Language

Embodiment

My researches:

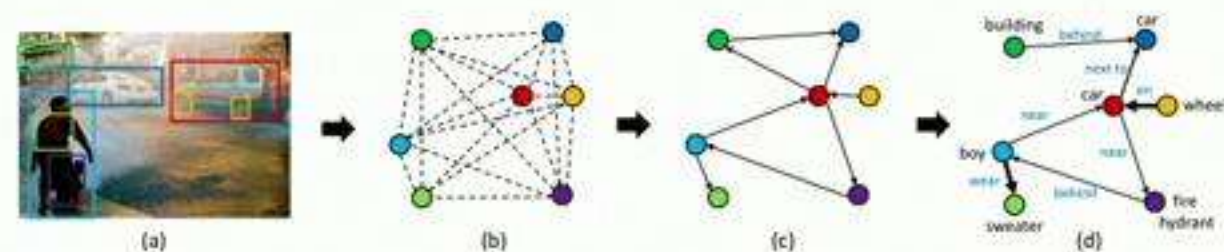


Cross-Image Structure



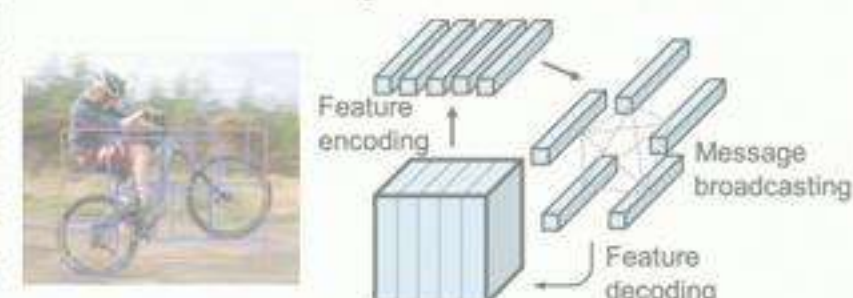
Joint Unsupervised Learning. CVPR 2016

Per-Image Structure

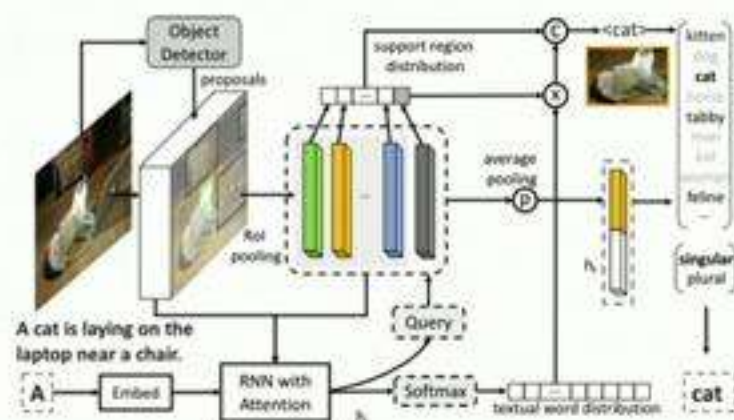


Graph R-CNN for Scene Graph Generation. ECCV 2018

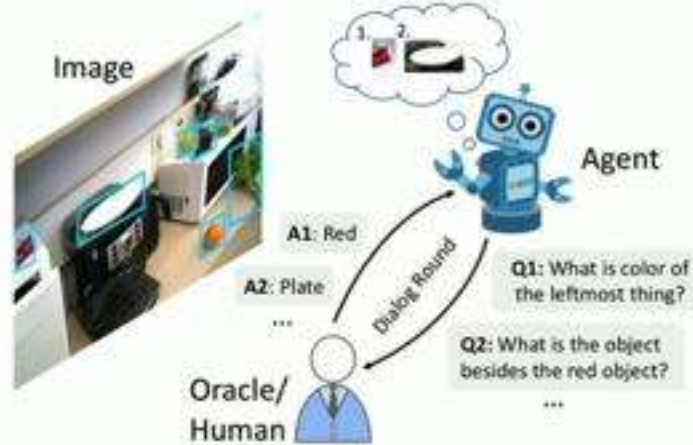
Per-Object Structure



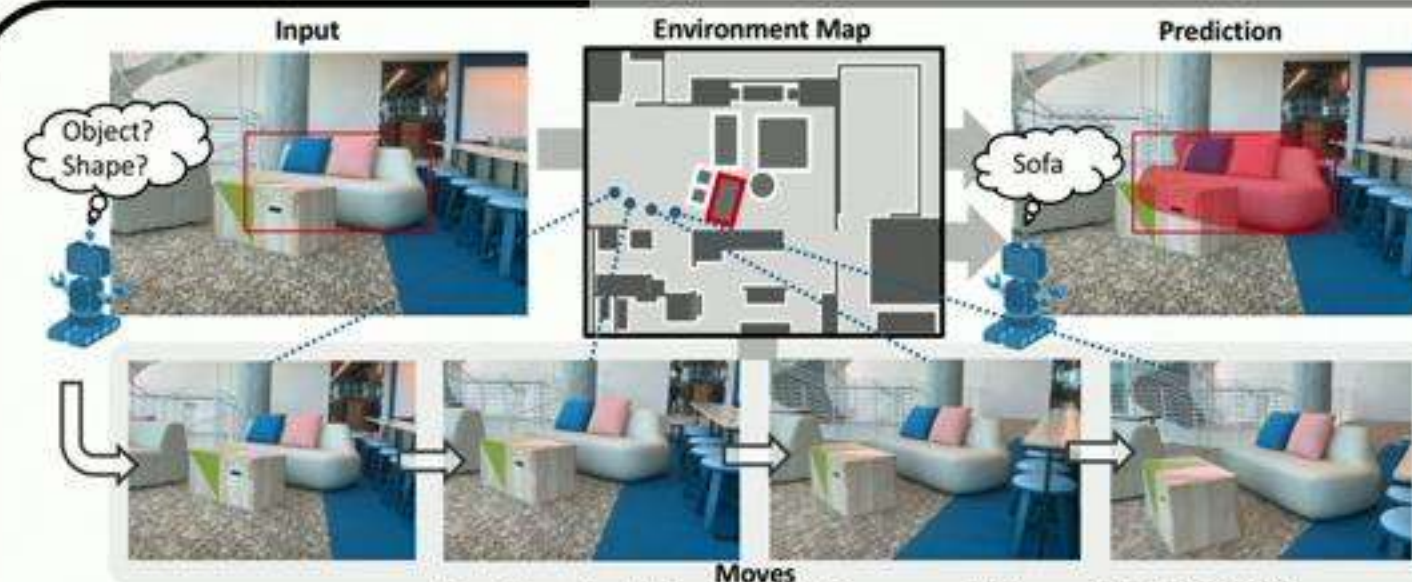
Neuron Communication Networks. NeurIPS 2019



Neural Baby Talk. CVPR 2018

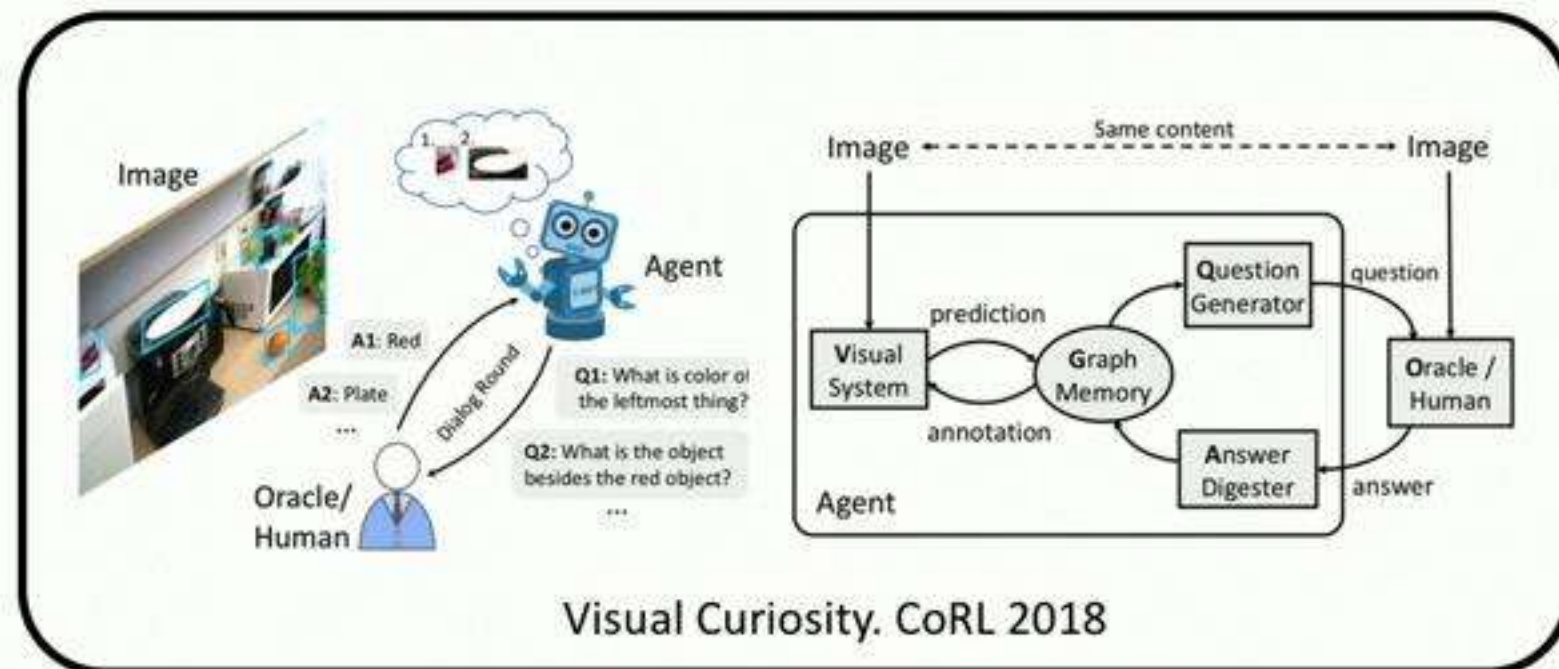


Visual Curiosity. CoRL 2018

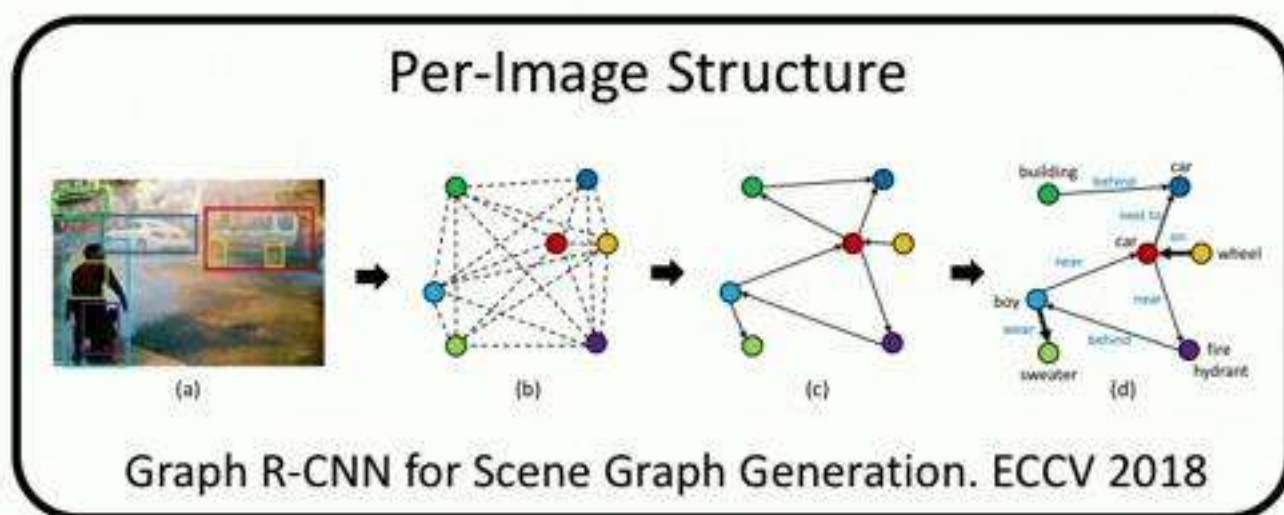


Embodied Amodal Recognition. ICCV 2019

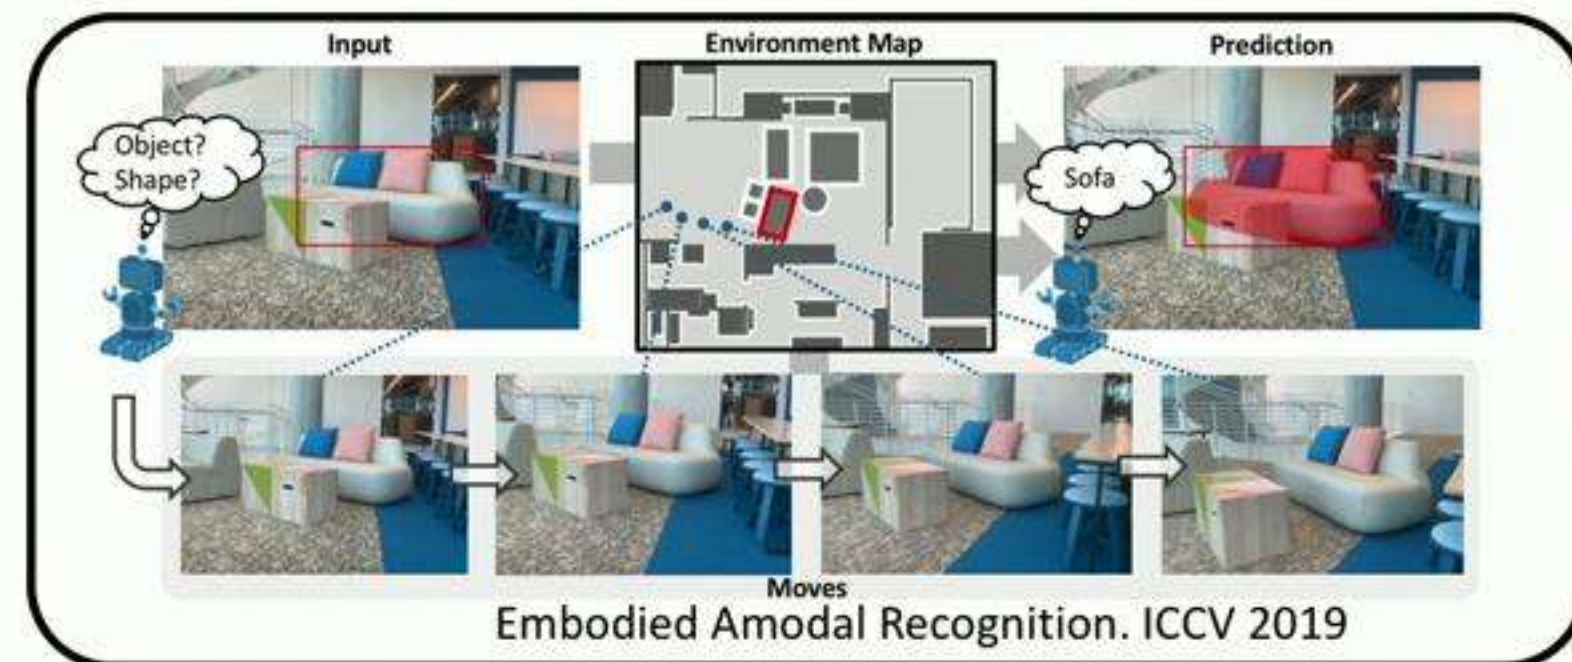
In this talk



Interact with Human



Structured Visual Understanding



Interact with Environment

Structured Visual Understanding

Graph R-CNN for Scene Graph Generation. ECCV 2018

What is scene graph?

Image as a single label



Image as an object set

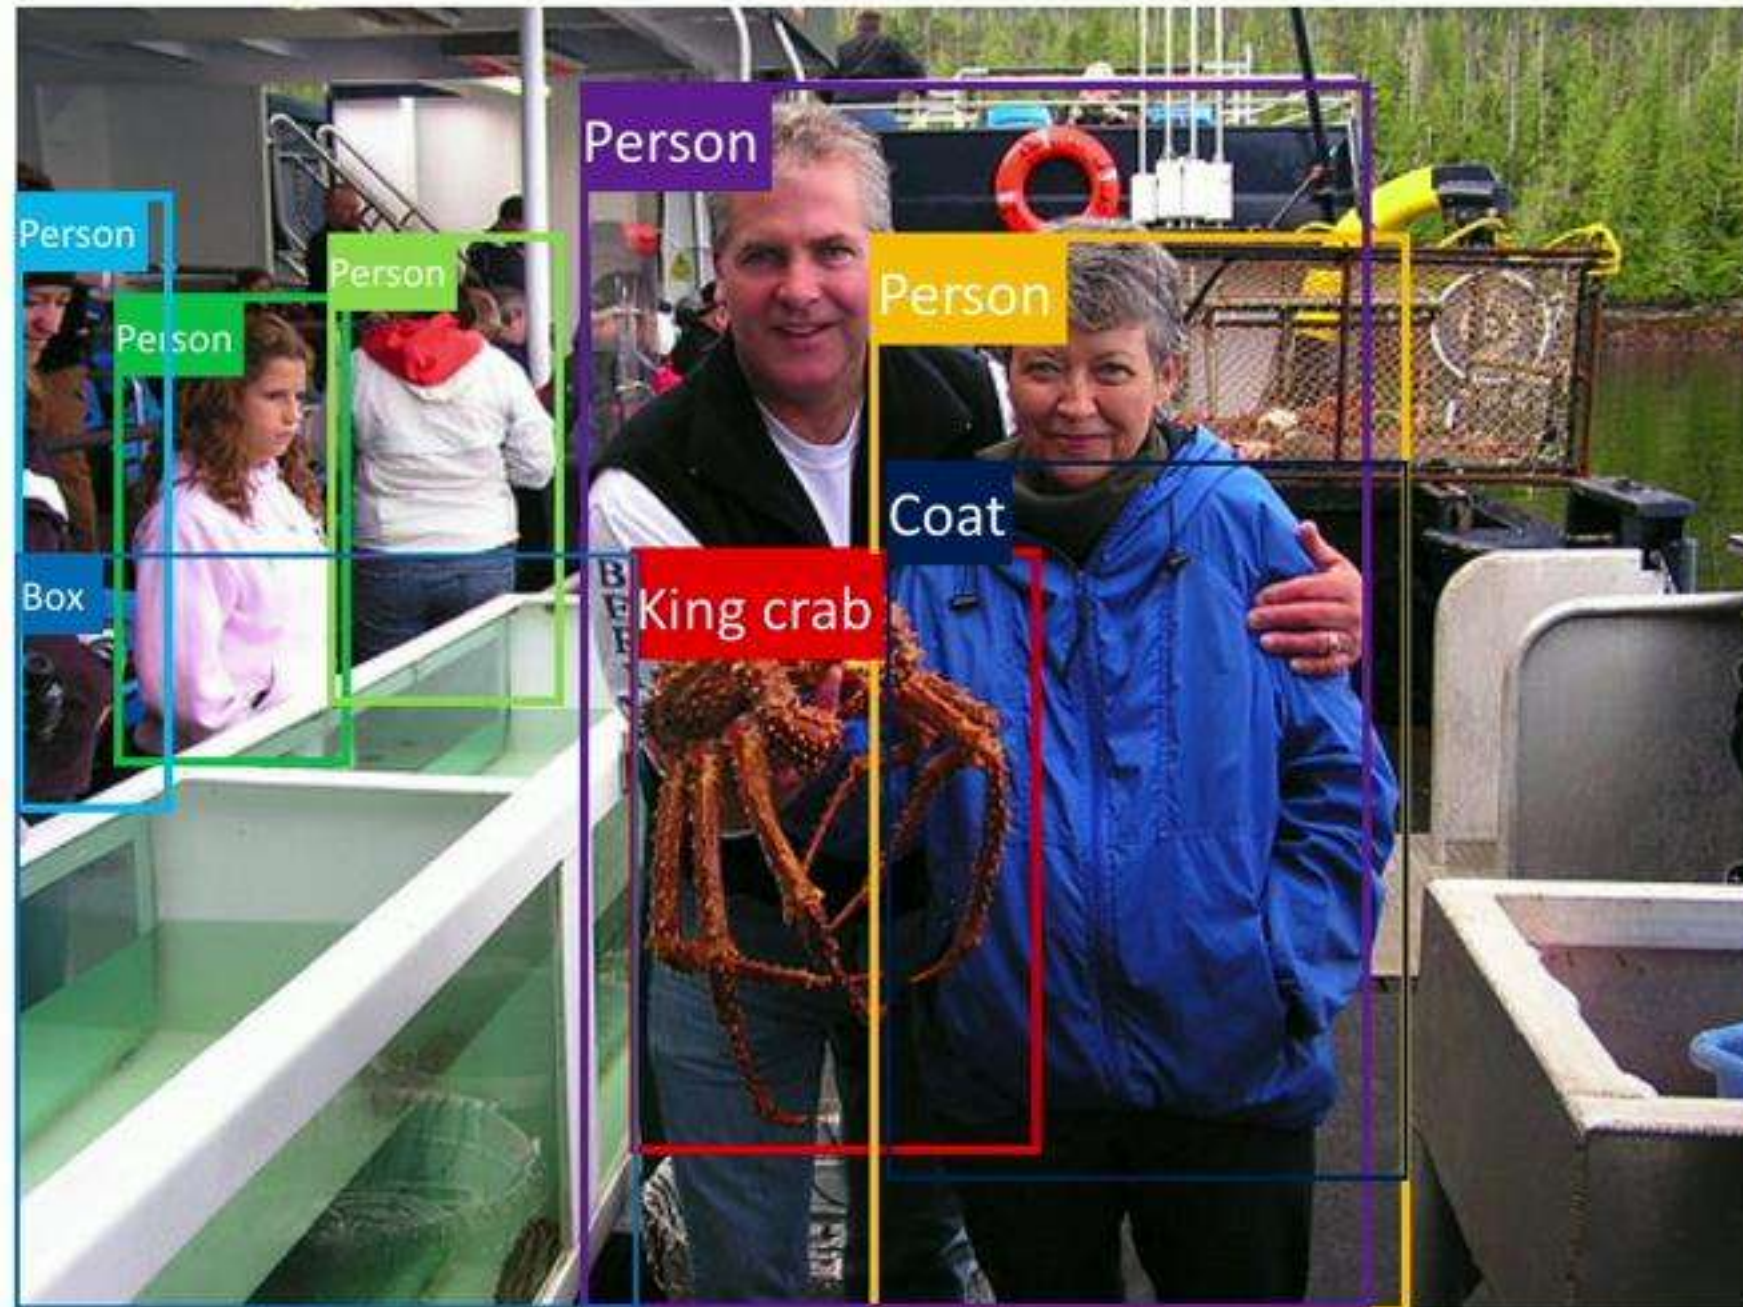
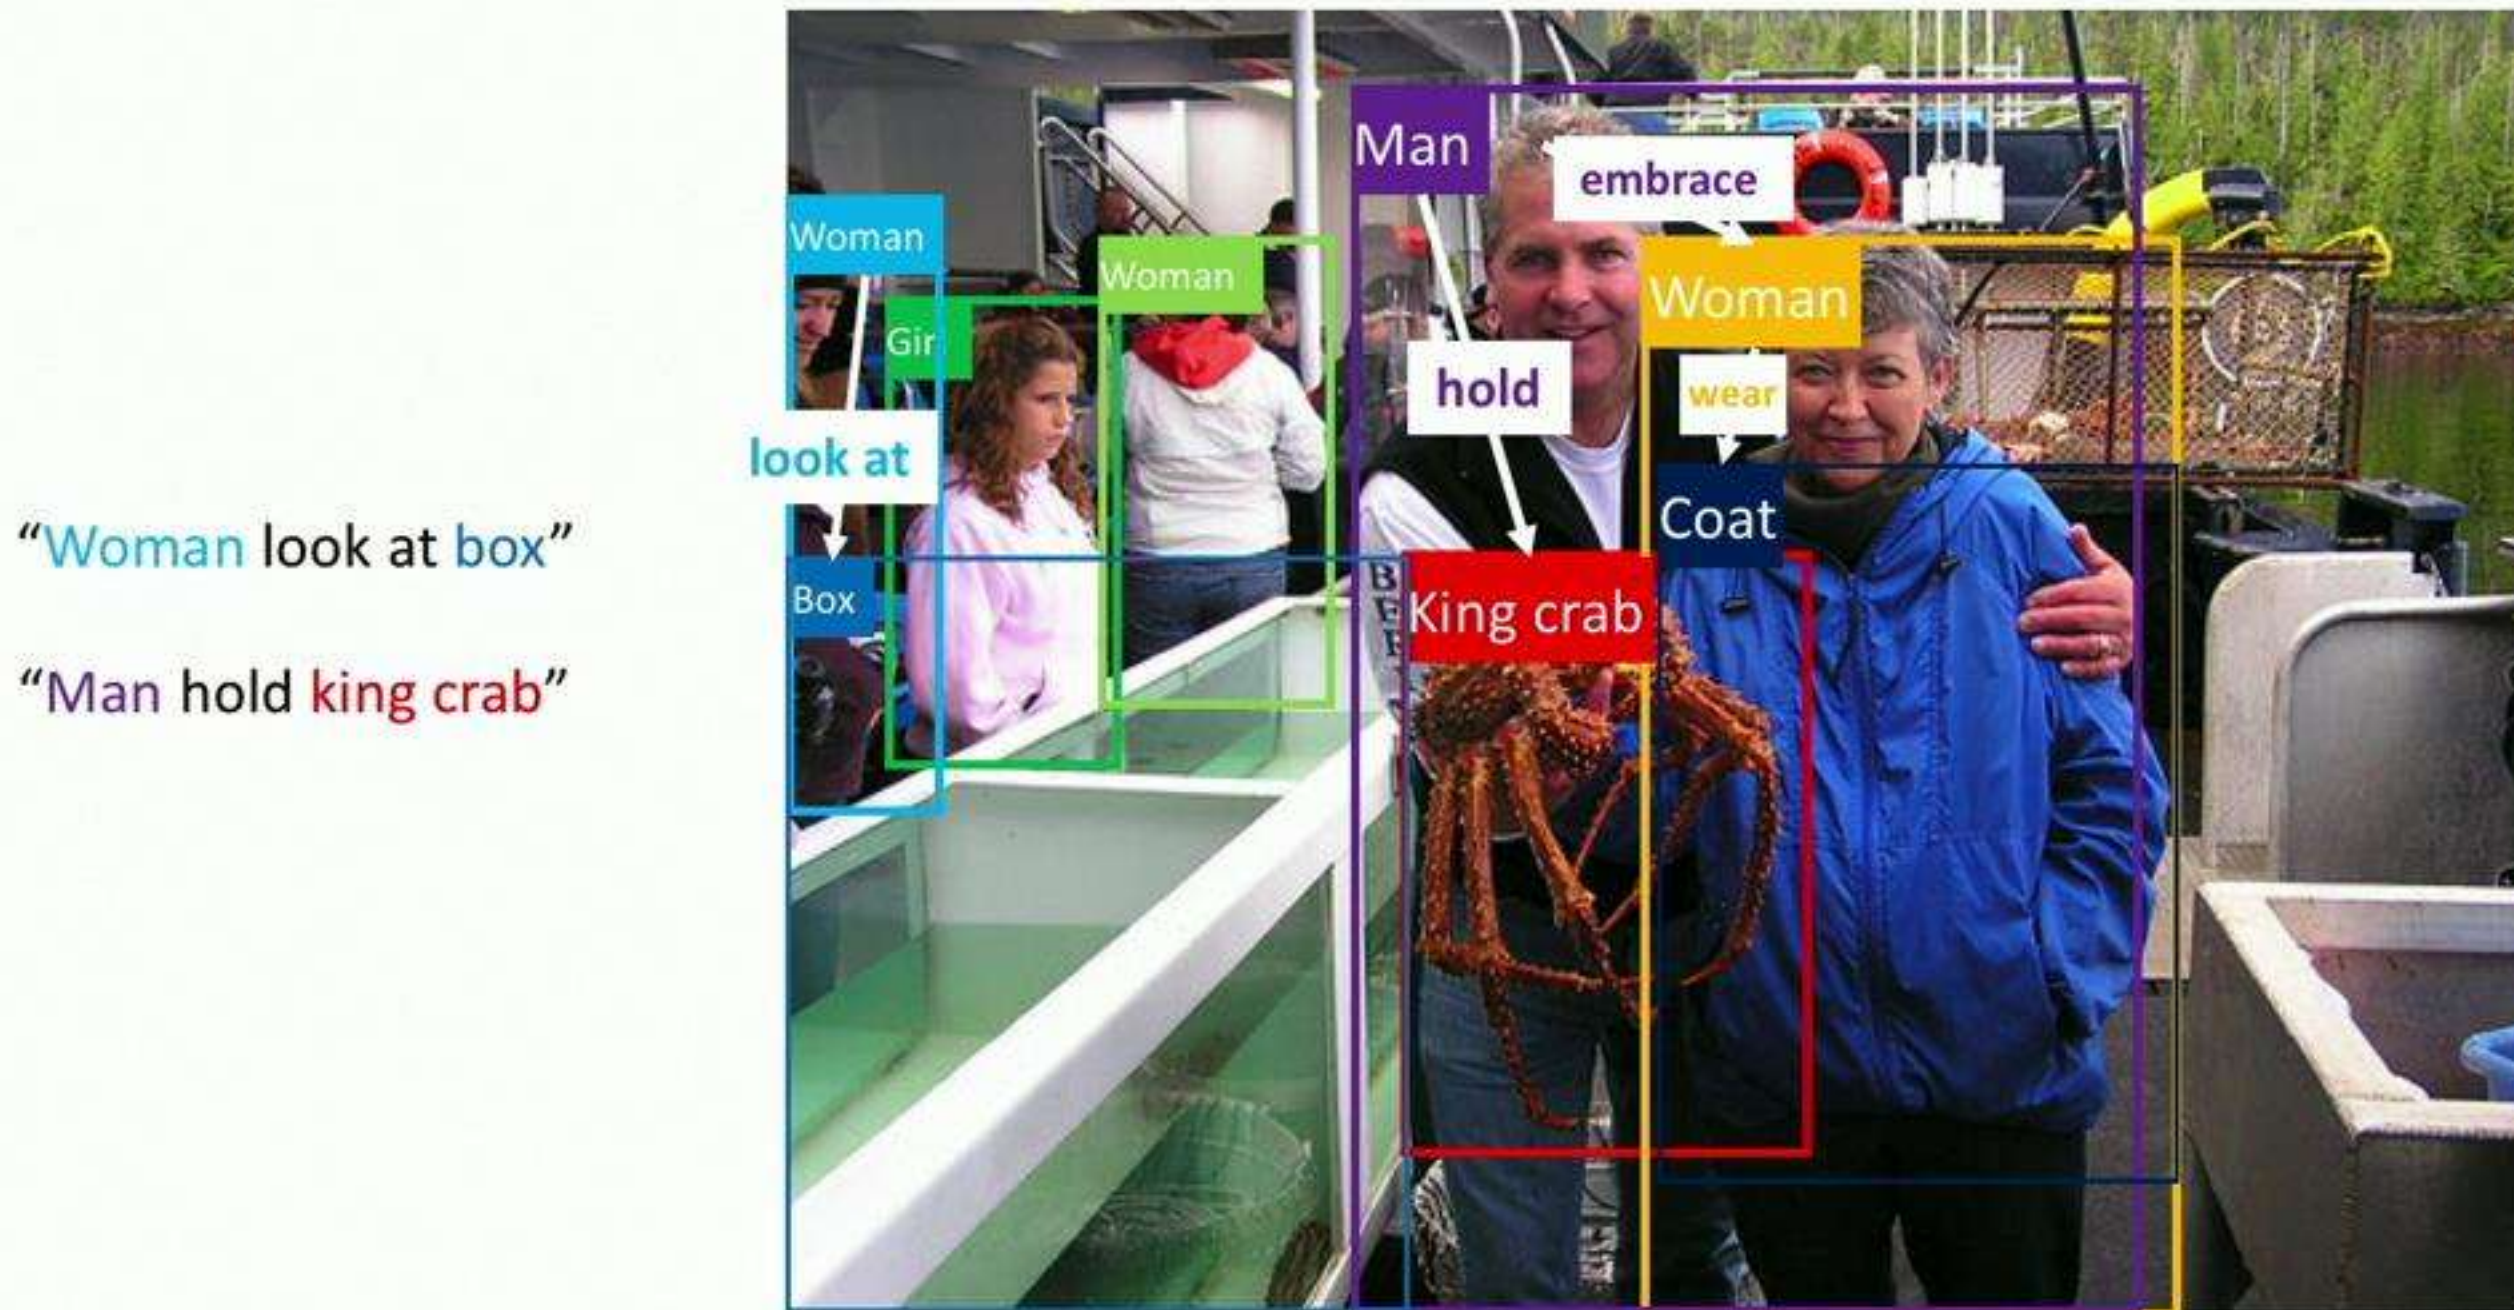
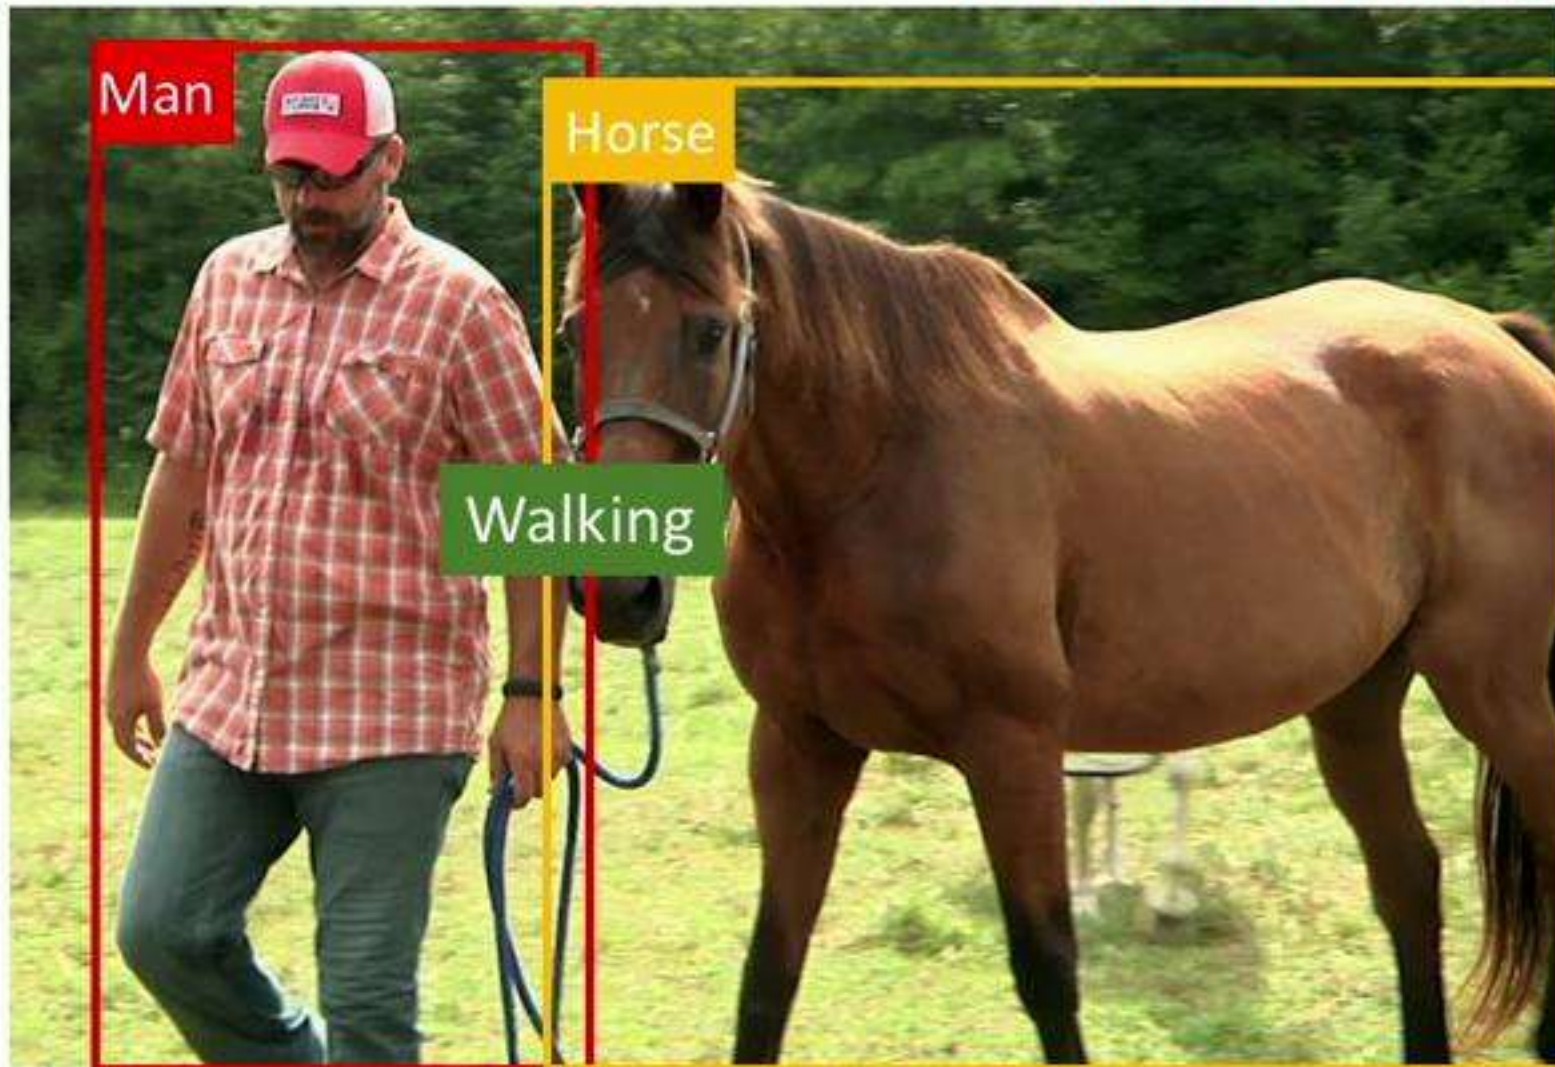


Image as a scene graph



Why we need scene graph?

Distinguish images more accurately



[1] Image Retrieval using Scene Graphs. Johnson et al. CVPR 2015

Left: <https://cals.ncsu.edu/wp-content/uploads/2016/08/horse-1500x931.png>

Rigth: https://www.videoblocks.com/video/the-man-in-hat-feed-a-brown-horse-with-flowers-on-the-meadow-supmox_3xj0tvkb67

Why we need scene graph?

Describe images more grounding



[1]. Auto-Encoding Scene Graphs for Image Captioning. Yang et al. arXiv 2018

[2]. Exploring Visual Relationship for Image Captioning. Yao et al. ECCV 2018

Left: <https://cals.ncsu.edu/wp-content/uploads/2016/08/horse-1500x931.png>

Rigth: https://www.videoblocks.com/video/the-man-in-hat-feed-a-brown-horse-with-flowers-on-the-meadow-supmox_3xj0tvkb67

Why we need scene graph?

Answer question more precisely



[1] Graph-Structured Representations for Visual Question Answering. Teney et al. CVPR 2017

[2] Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding. Yi et al. Neurips 2018

Left: <https://cals.ncsu.edu/wp-content/uploads/2016/08/horse-1500x931.png>

Rigth: https://www.videoblocks.com/video/the-man-in-hat-feed-a-brown-horse-with-flowers-on-the-meadow-supmox_3xj0tvkb67

Why we need scene graph?

Generate questions more grounding



[1] Visual Curiosity: Learning to Ask Questions to Learn Visual Recognition. Yang et al. CoRL 2018

Left: <https://cals.ncsu.edu/wp-content/uploads/2016/08/horse-1500x931.png>

Rigth: https://www.videoblocks.com/video/the-man-in-hat-feed-a-brown-horse-with-flowers-on-the-meadow-supmox_3xj0tvkb67

Why we need scene graph?

Answer question more precisely



[1] Graph-Structured Representations for Visual Question Answering. Teney et al. CVPR 2017

[2] Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding. Yi et al. Neurips 2018

Left: <https://cals.ncsu.edu/wp-content/uploads/2016/08/horse-1500x931.png>

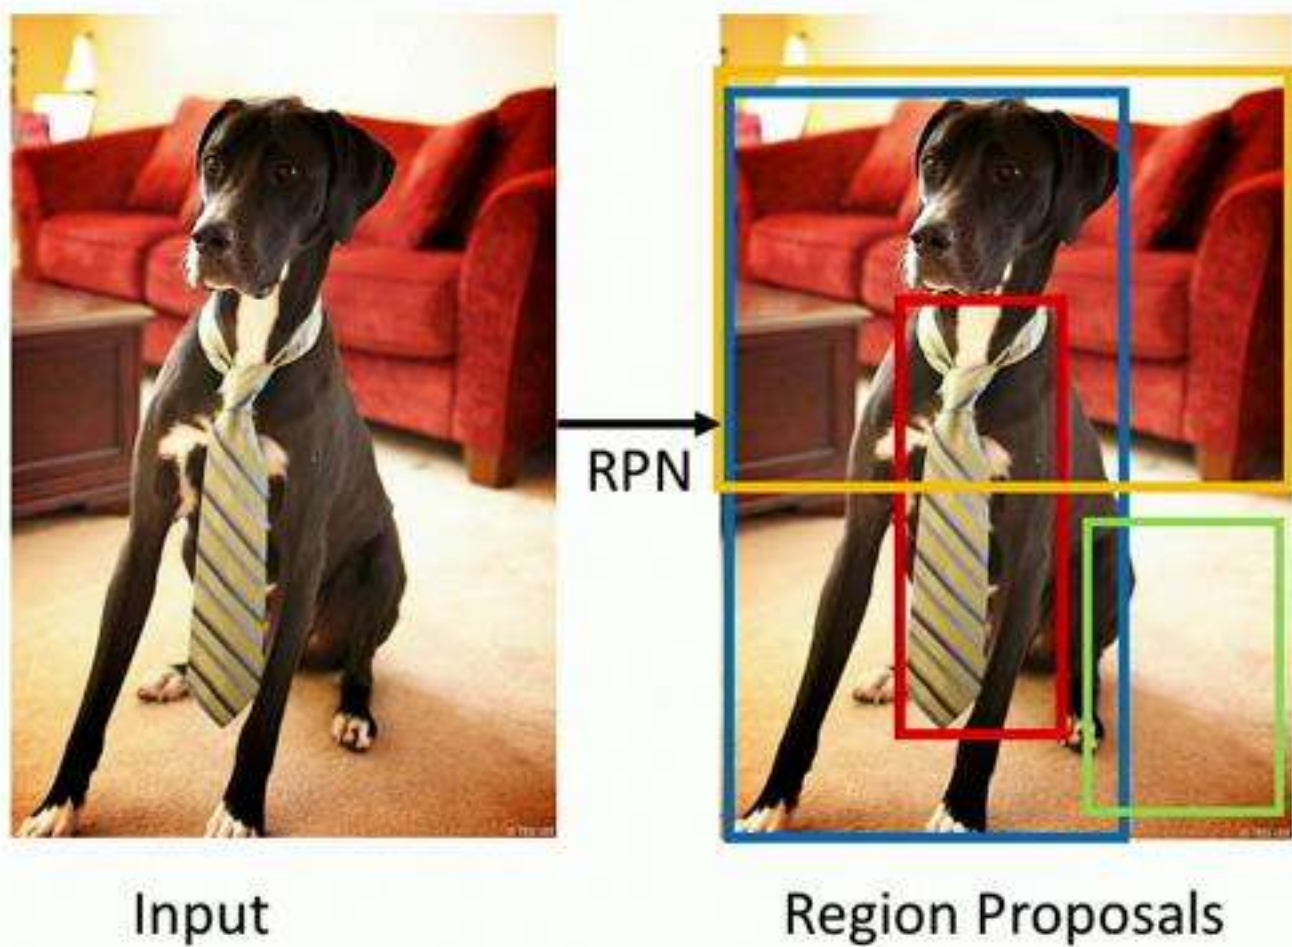
Rigth: https://www.videoblocks.com/video/the-man-in-hat-feed-a-brown-horse-with-flowers-on-the-meadow-supmox_3xj0tvkb67

General Pipeline

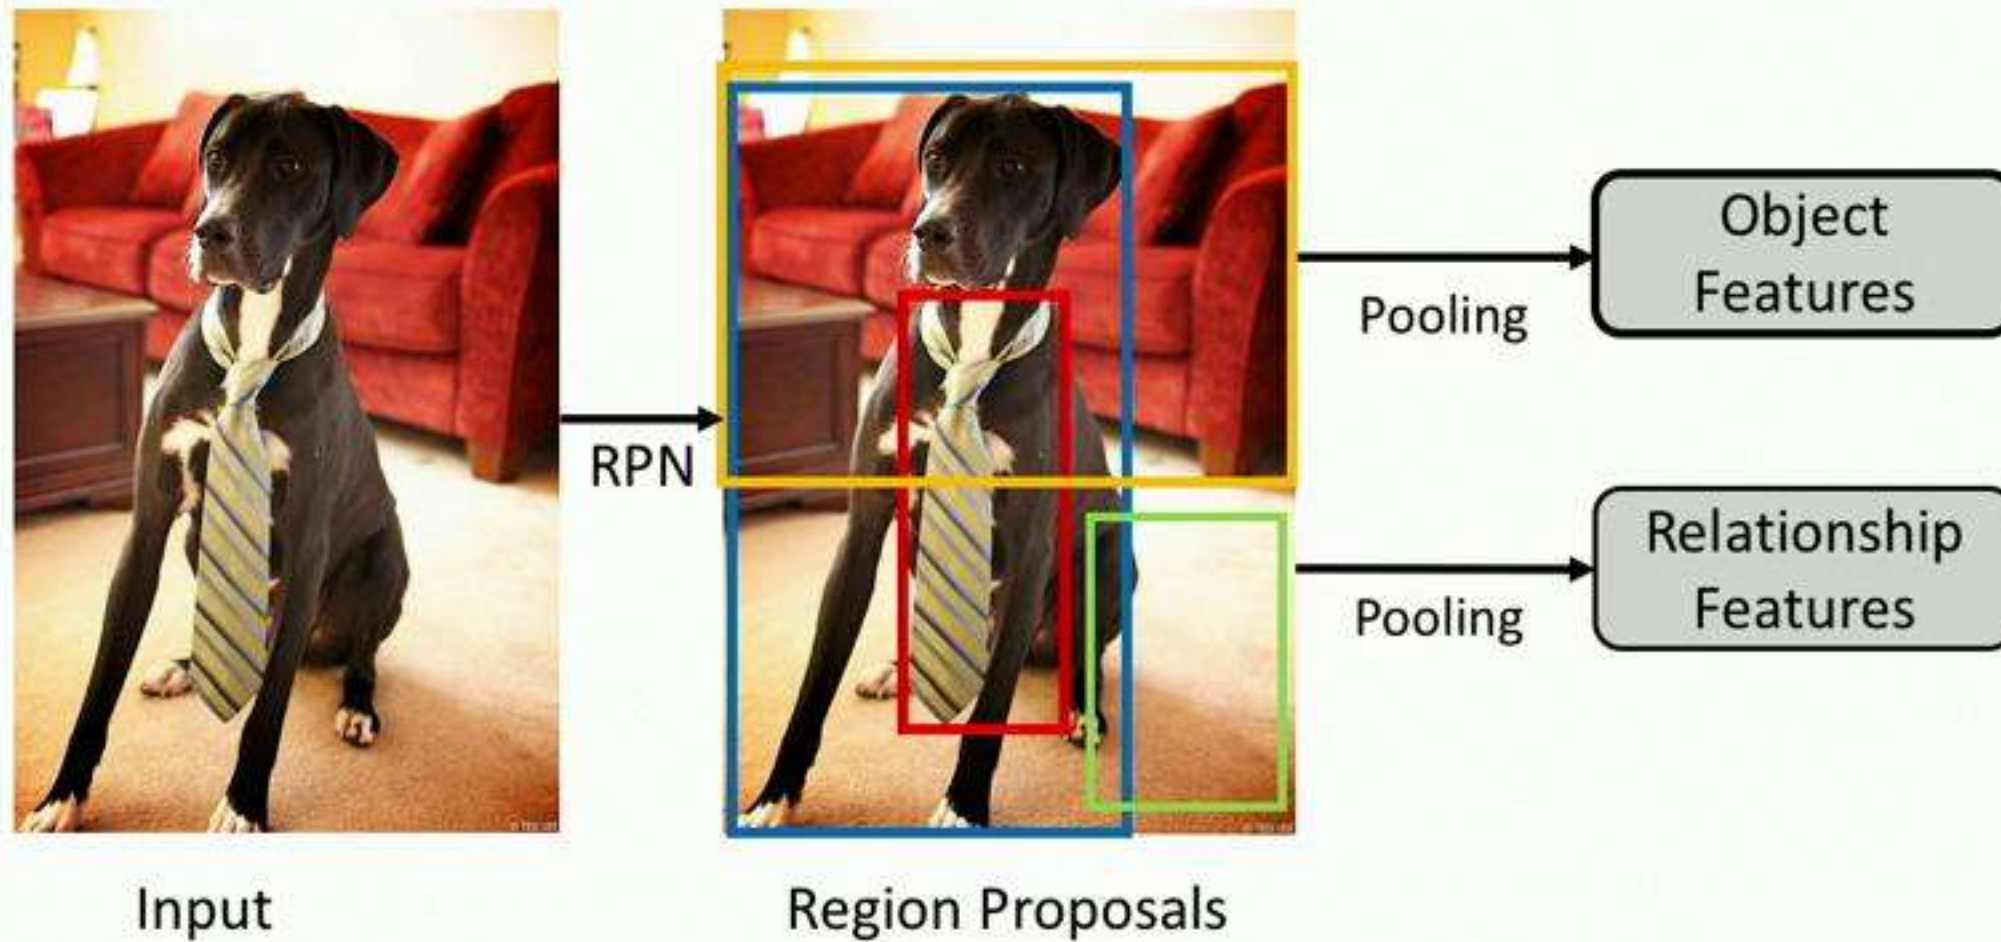


Input

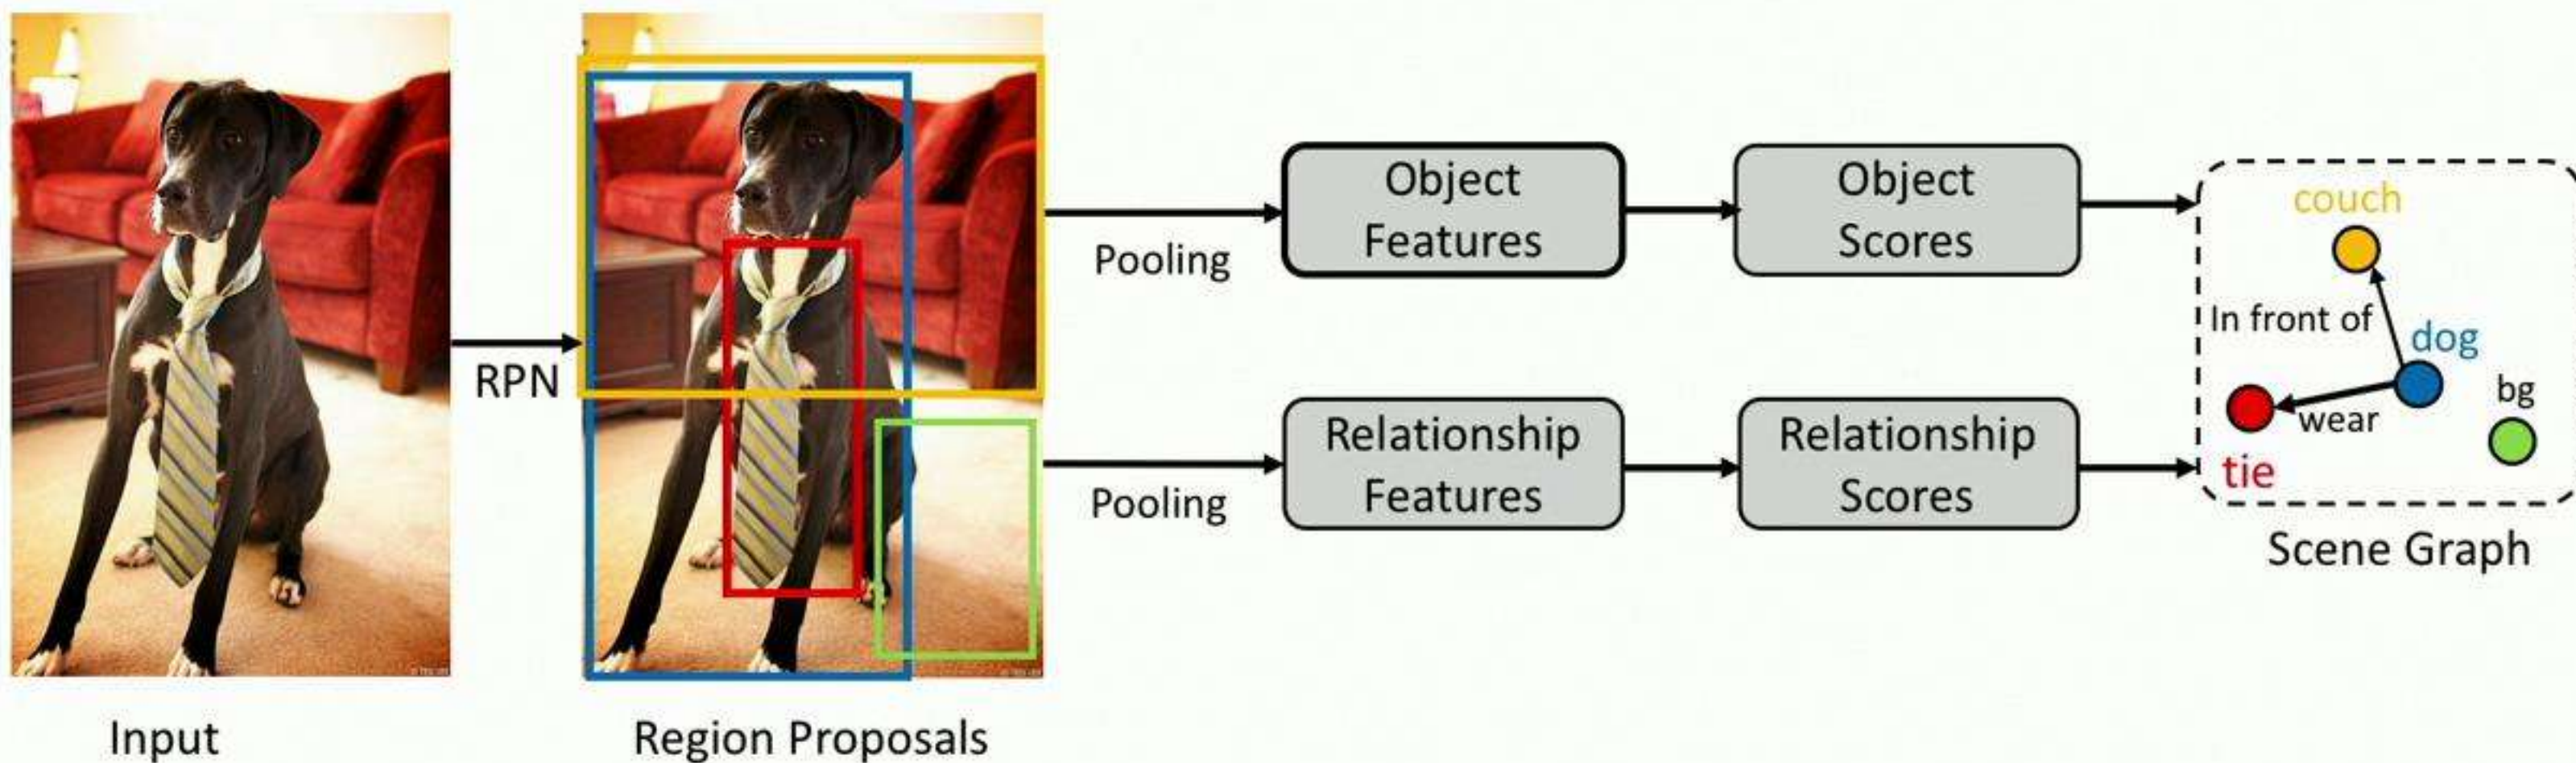
General Pipeline



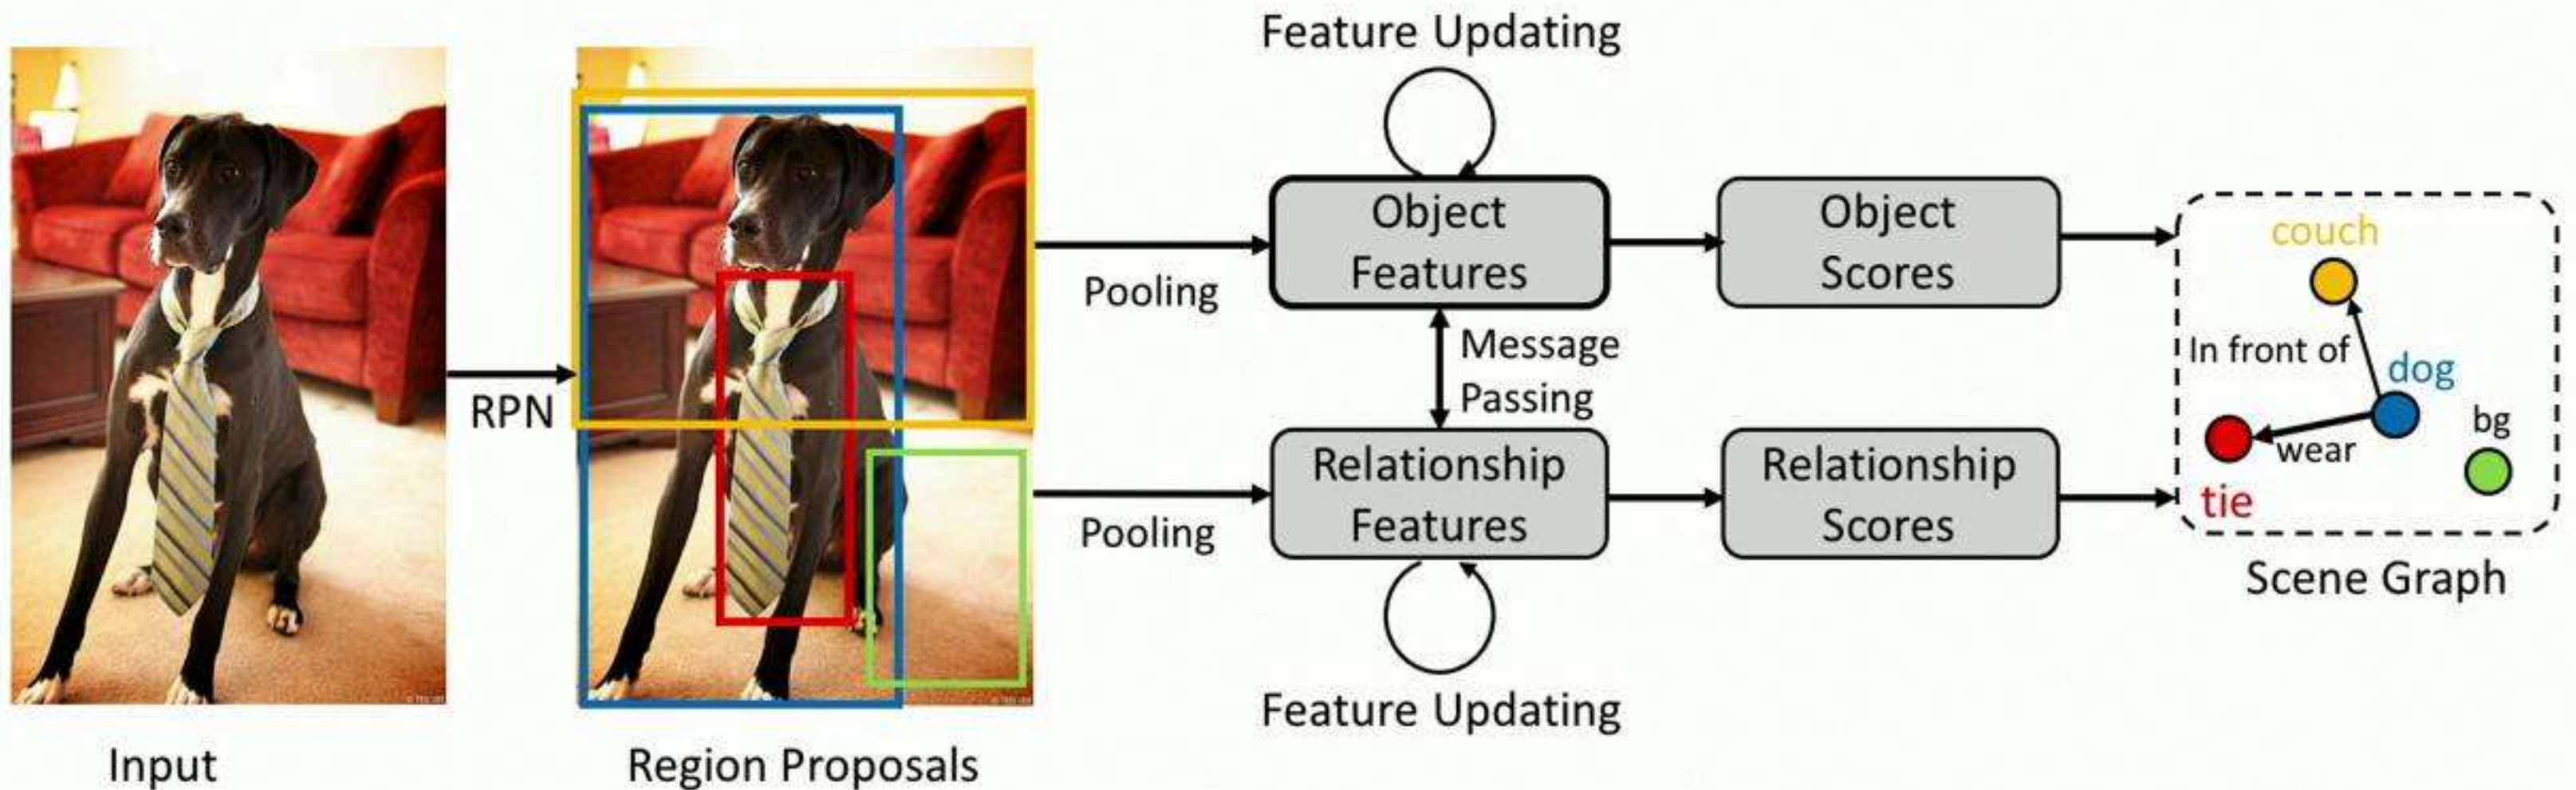
General Pipeline



General Pipeline

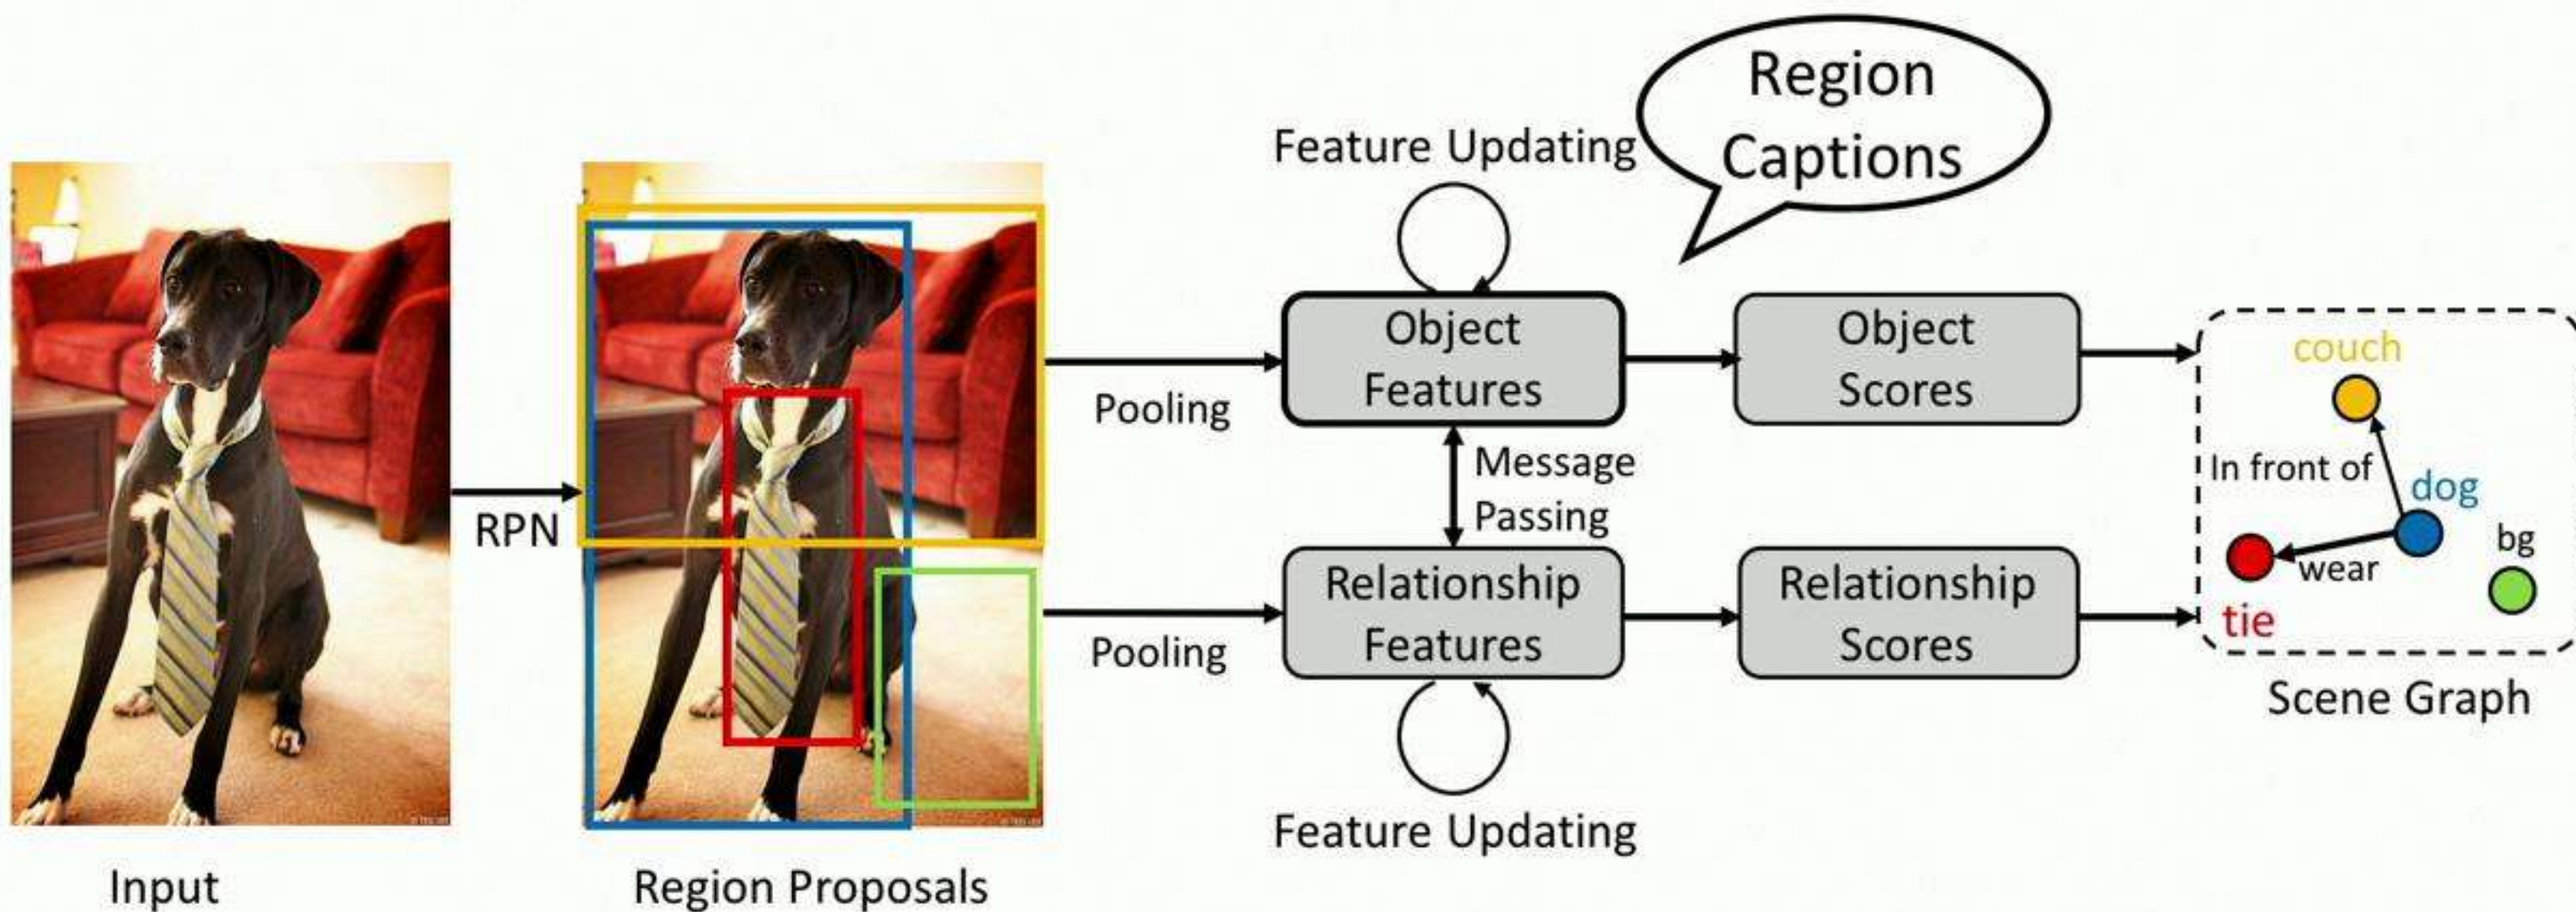


IMP Model



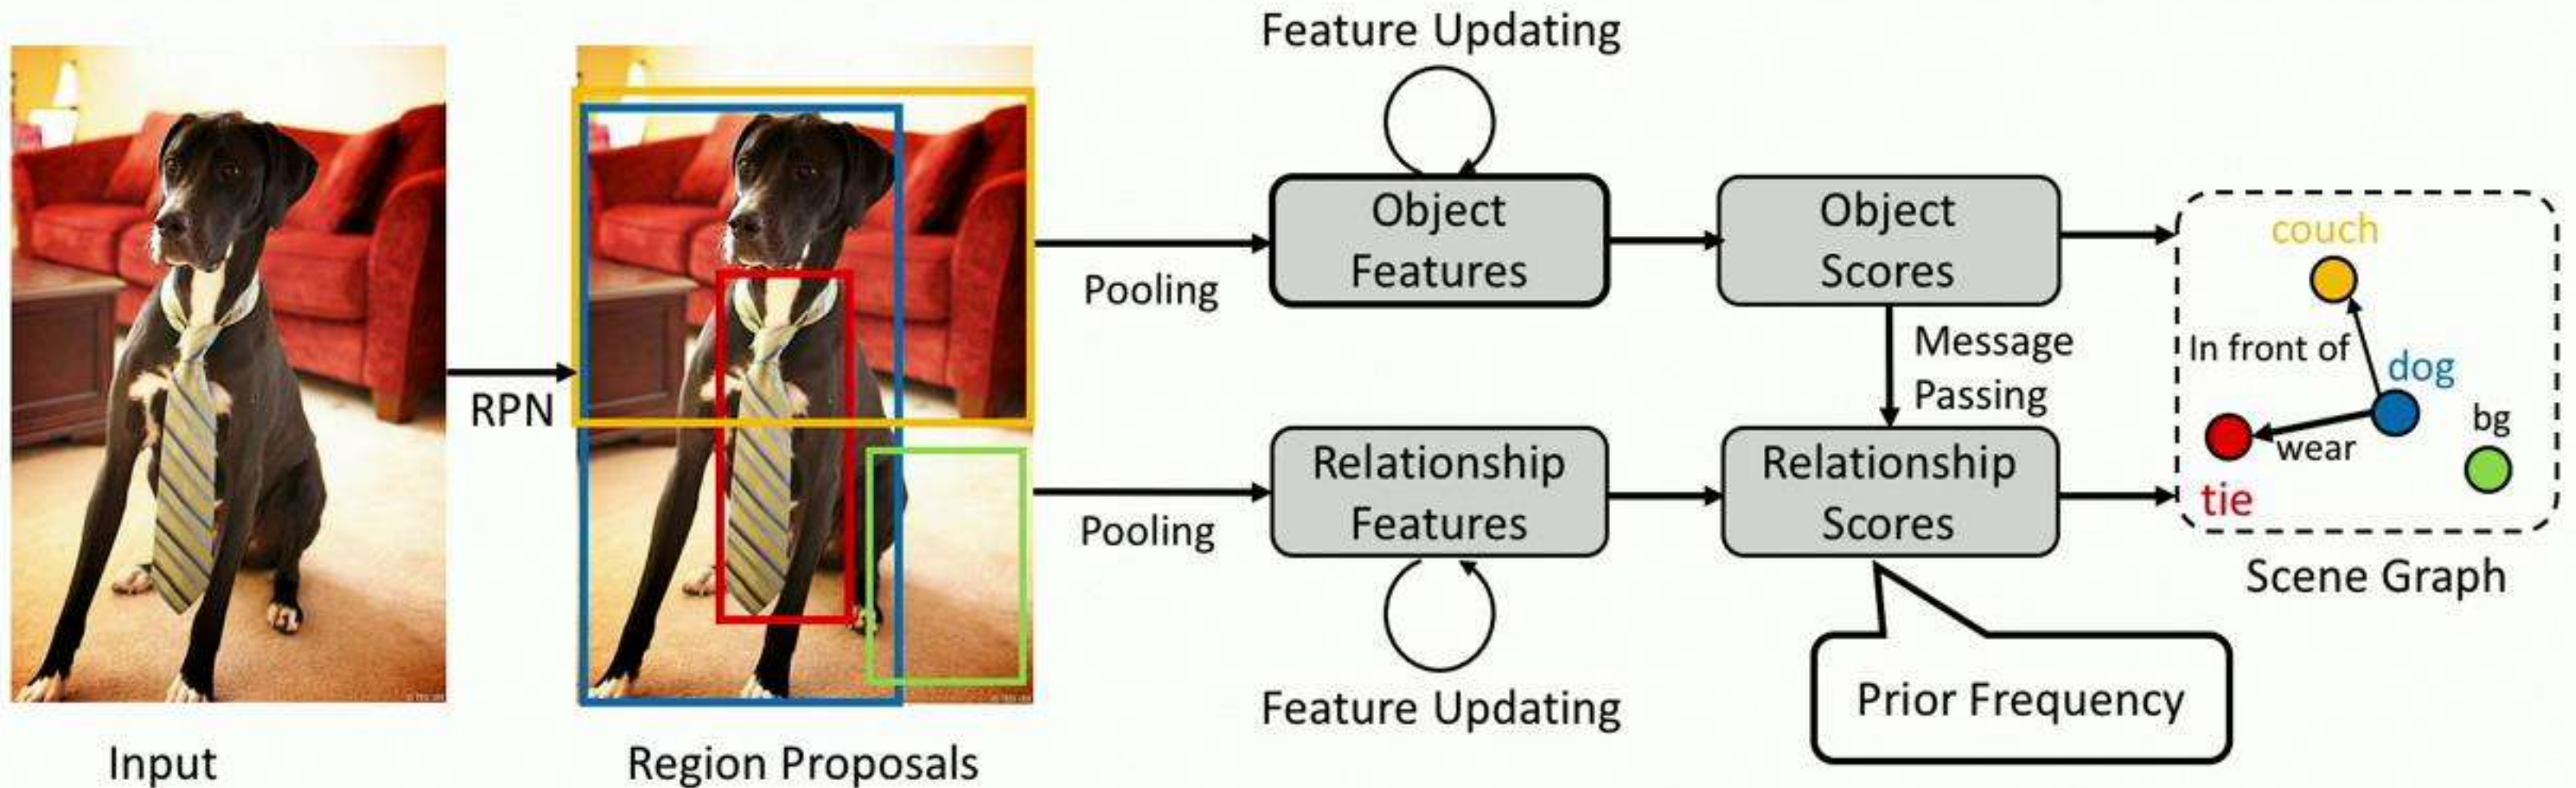
Scene Graph Generation by Iterative Message Passing. Xu et al. CVPR 2017

MSDN Model



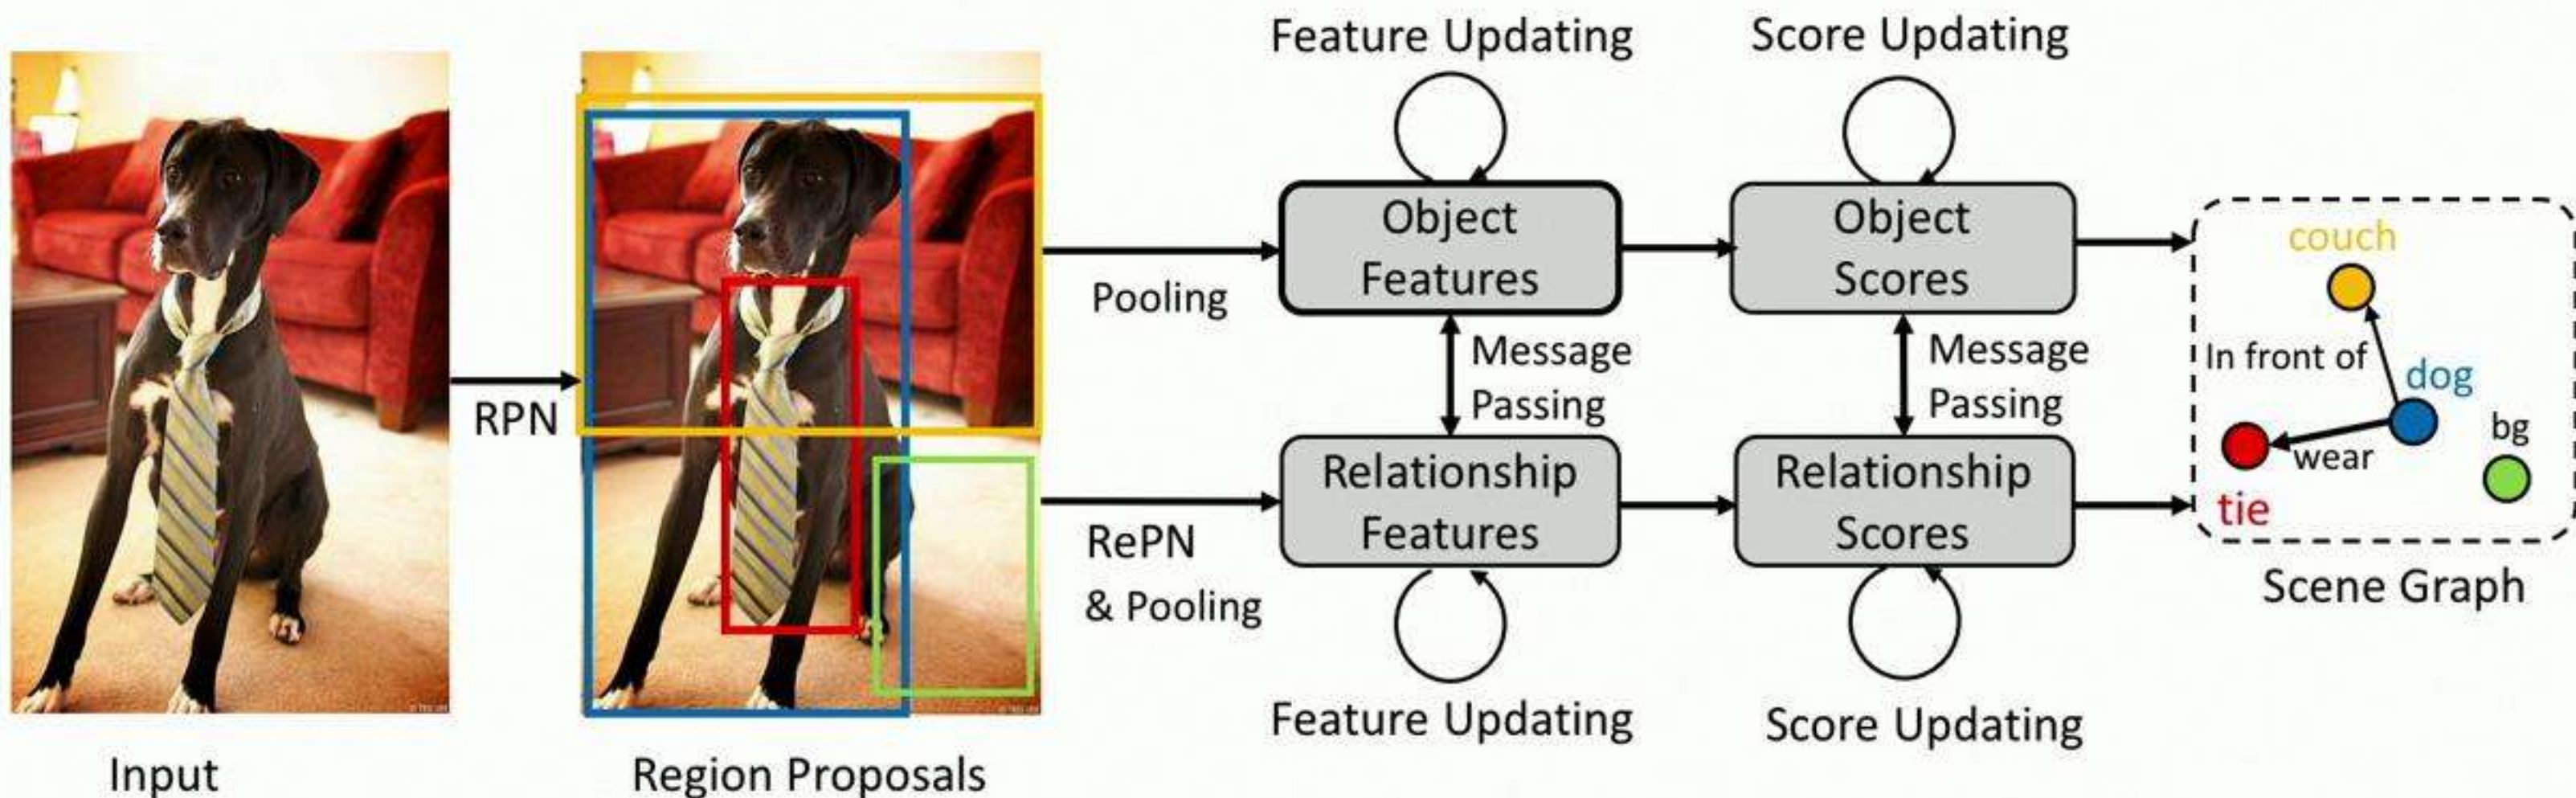
Scene Graph Generation from Objects, Phrases and Region Captions. Li et al. ICCV 2017

Neural Motif Network



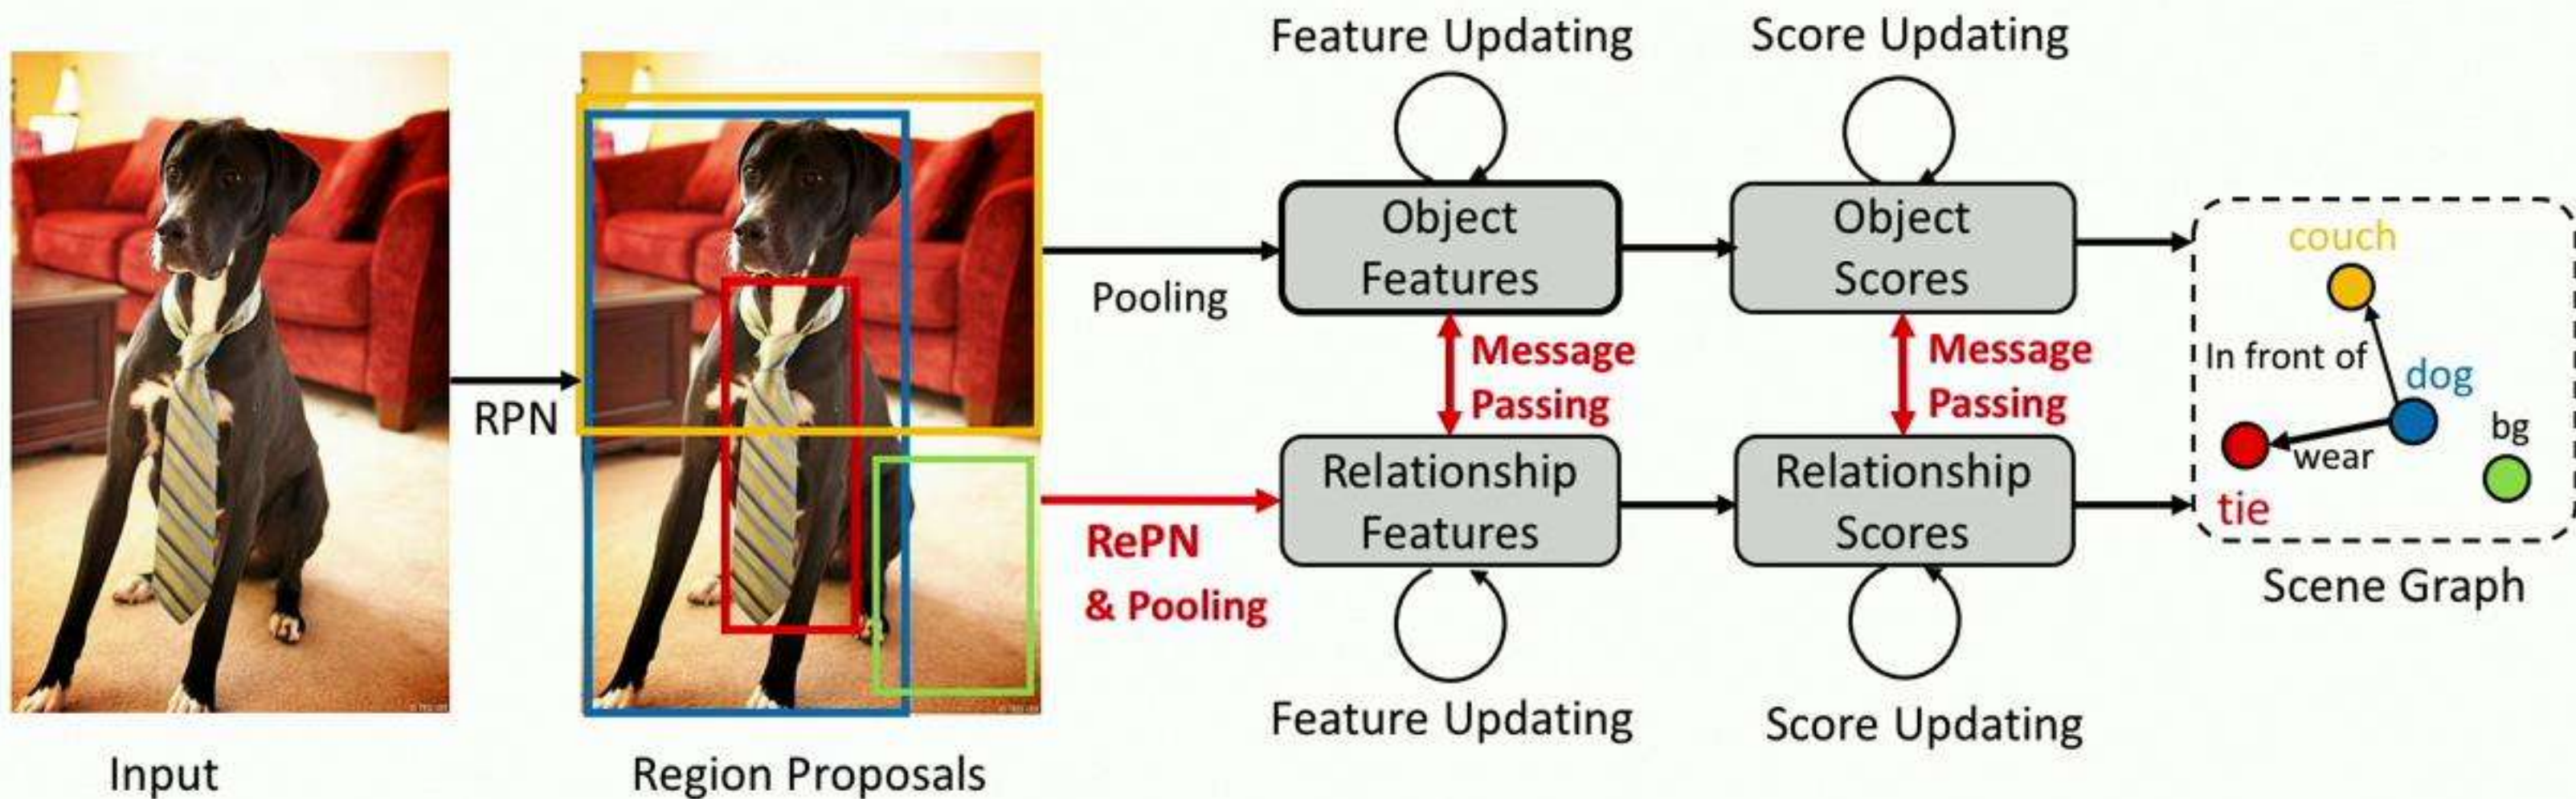
Neural Motifs: Scene Graph Parsing with Global Context. Zellers et al. CVPR 2018

Our model: Graph R-CNN

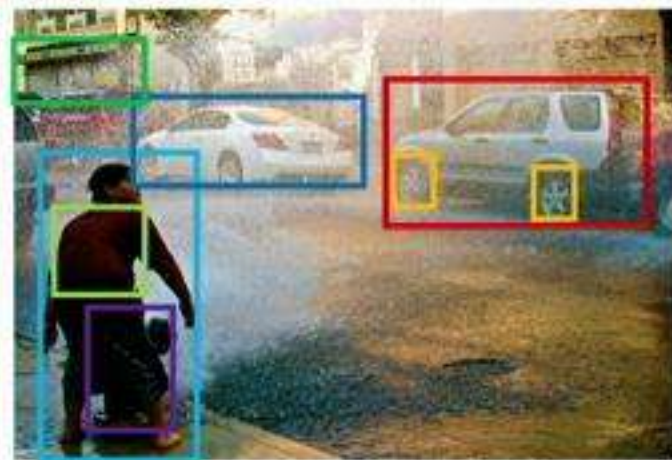


Jianwei Yang*, Jiasen Lu*, Stefan Lee, Dhruv Batra, Devi Parikh. Graph R-CNN for Scene Graph Generation. ECCV 2018.

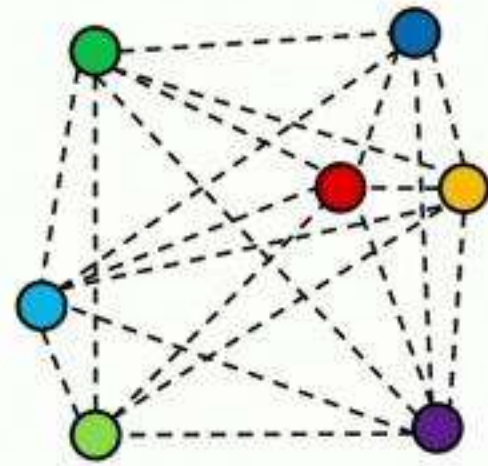
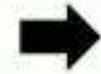
Our model: Graph R-CNN



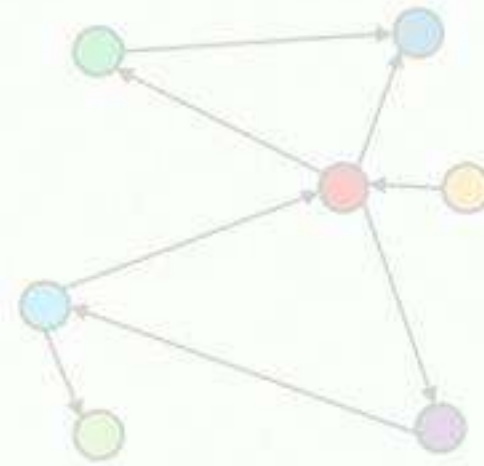
Motivations



(a)



(b)



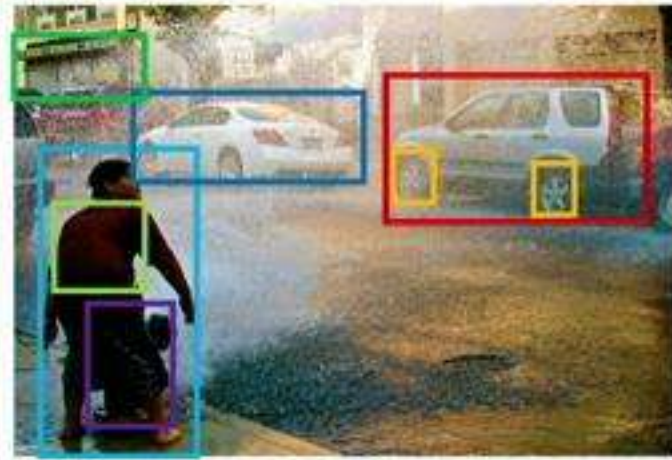
(c)



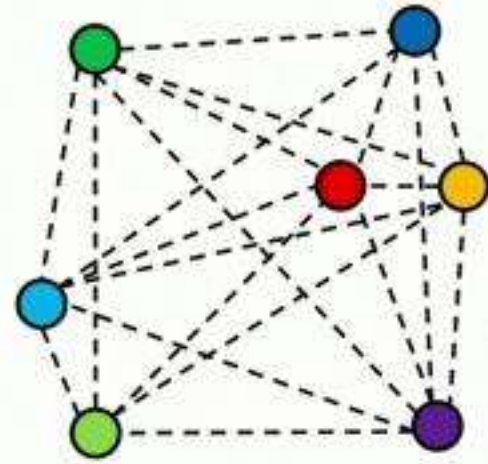
(d)

1. Objects in a scene usually have relationships with others;

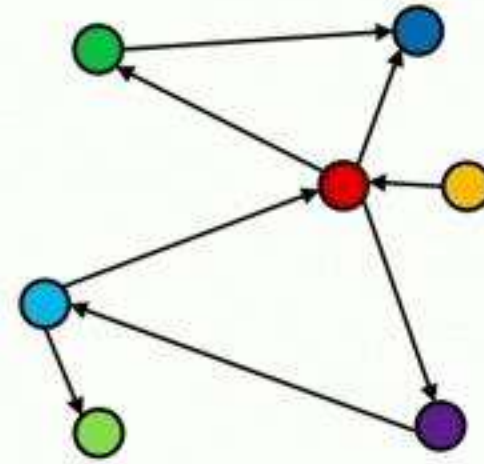
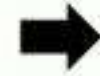
Motivations



(a)



(b)



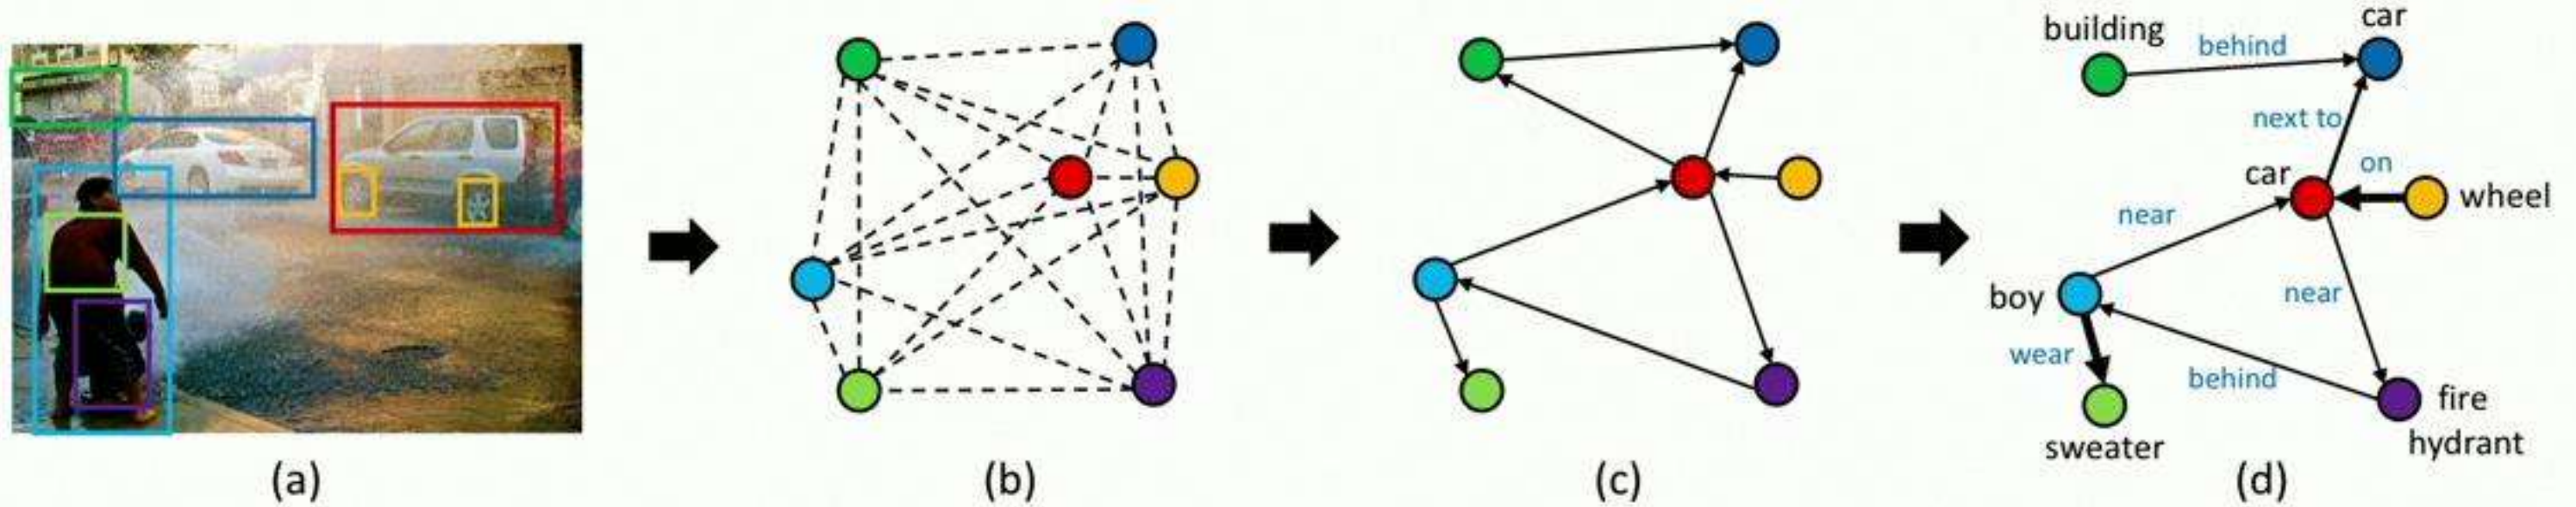
(c)



(d)

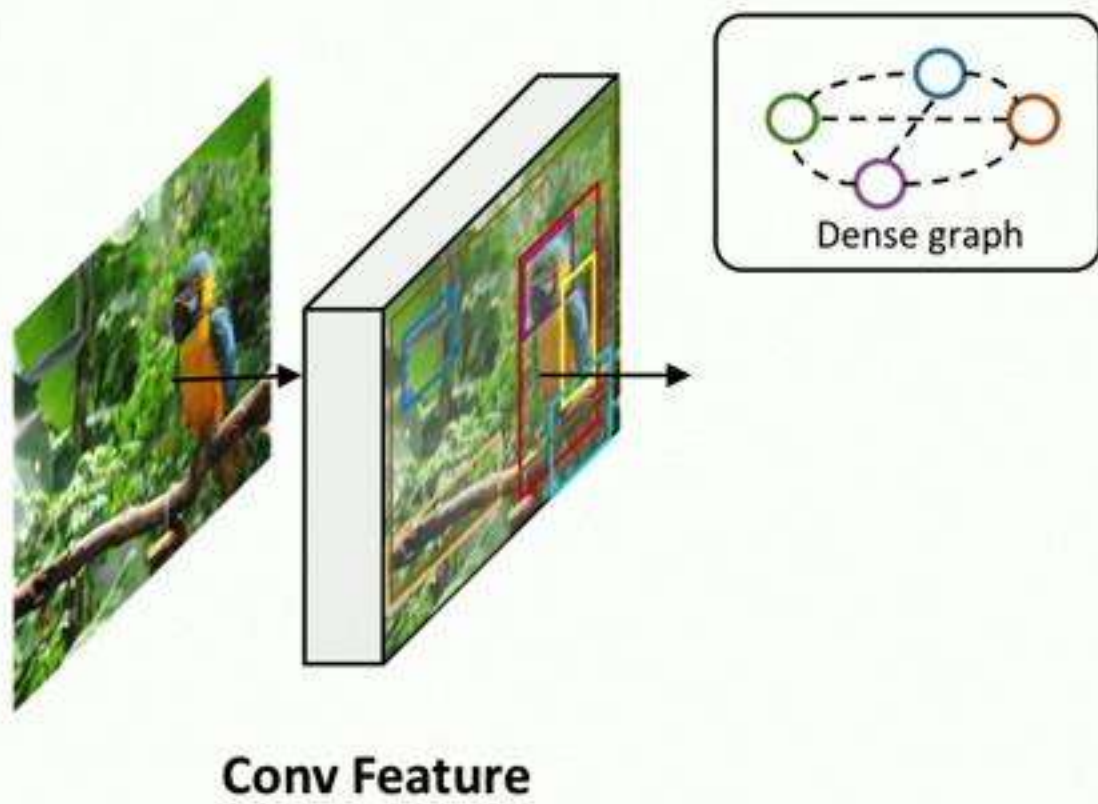
1. Objects in a scene usually have relationships with others;
2. Not all object pairs have relationships, the scene graph is usually sparse,

Motivations

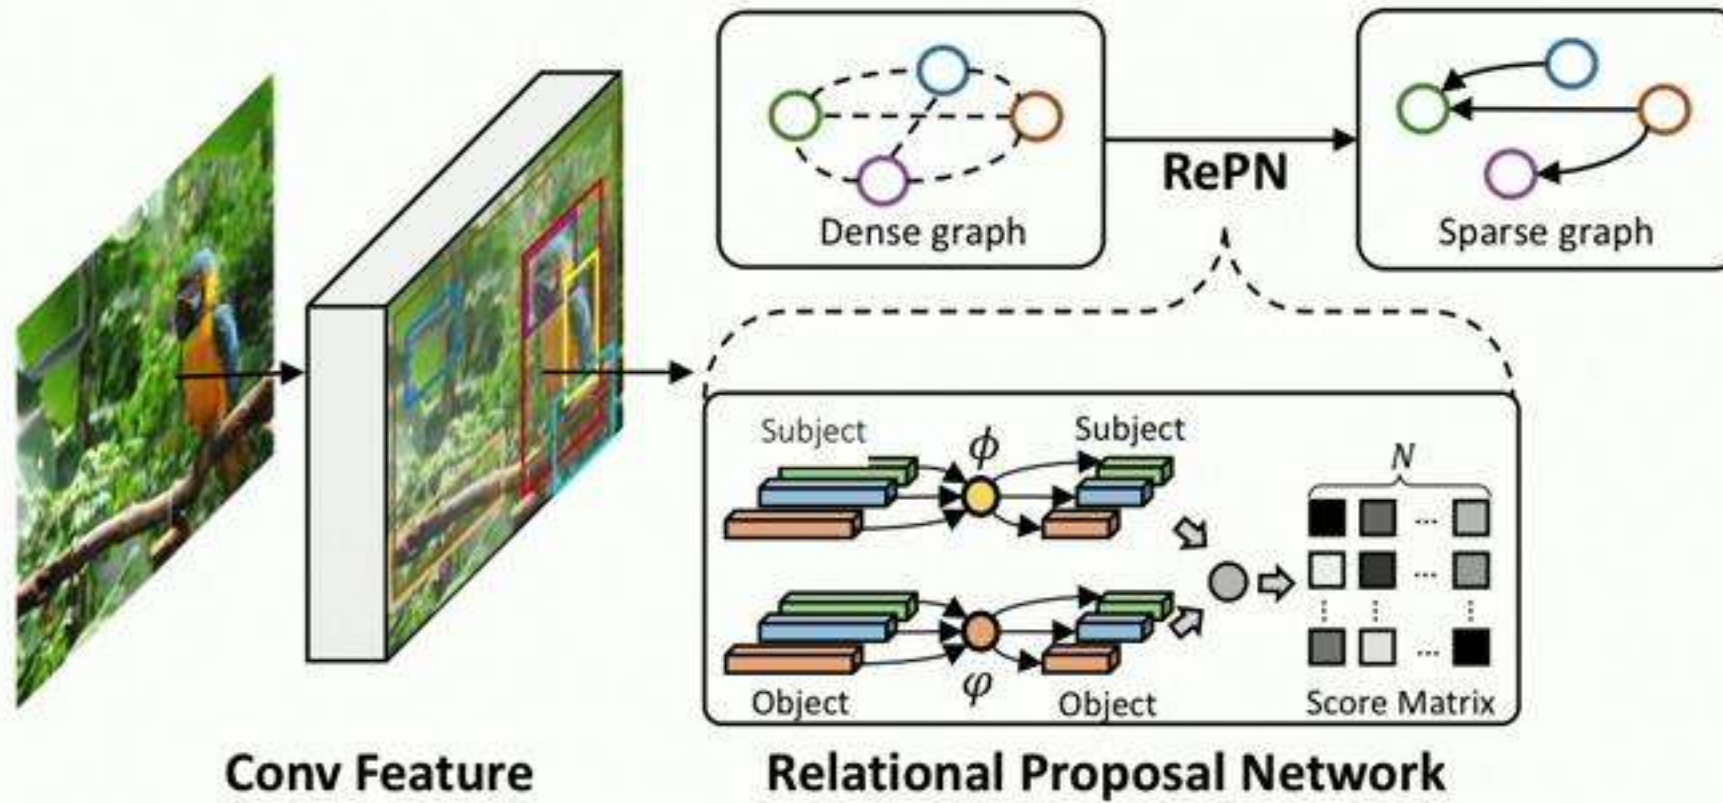


1. Objects in a scene usually have relationships with others;
2. Not all object pairs have relationships, the scene graph is usually sparse,
3. Existence of relationships highly depends on the object categories, and type of relationships highly depends on the context.

Framework

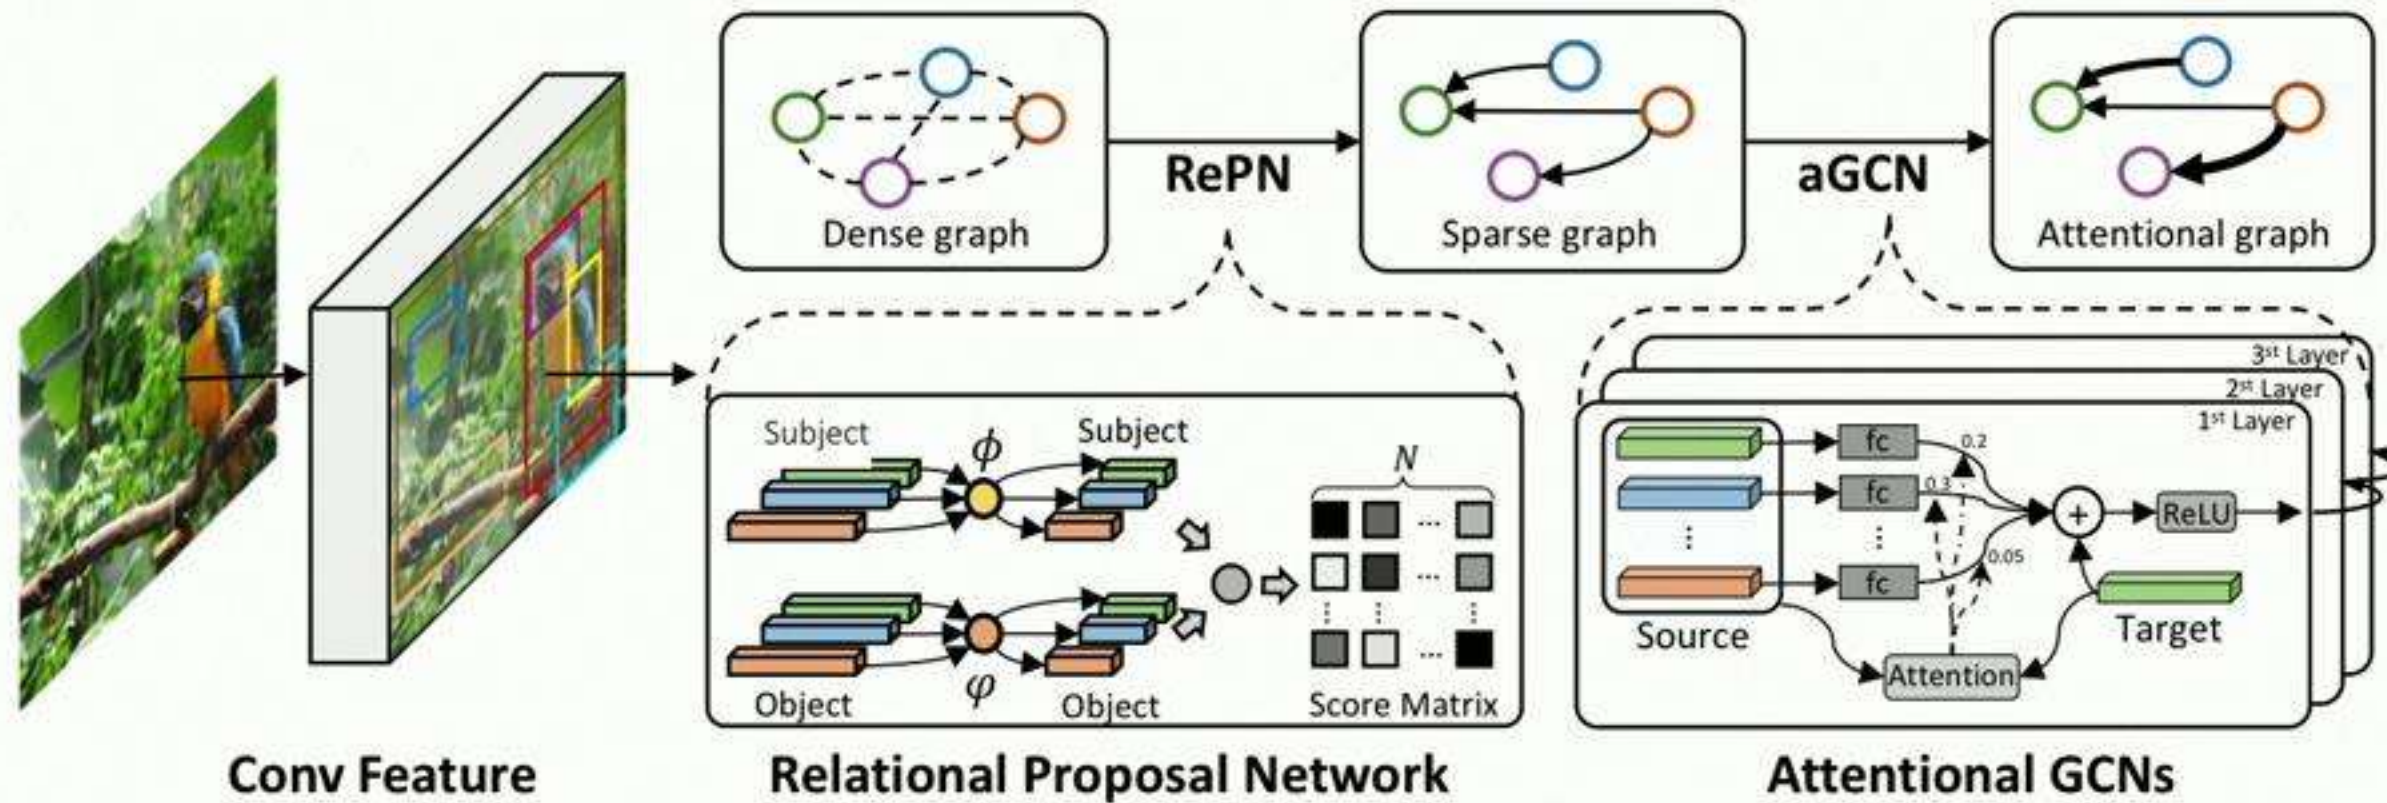


Framework



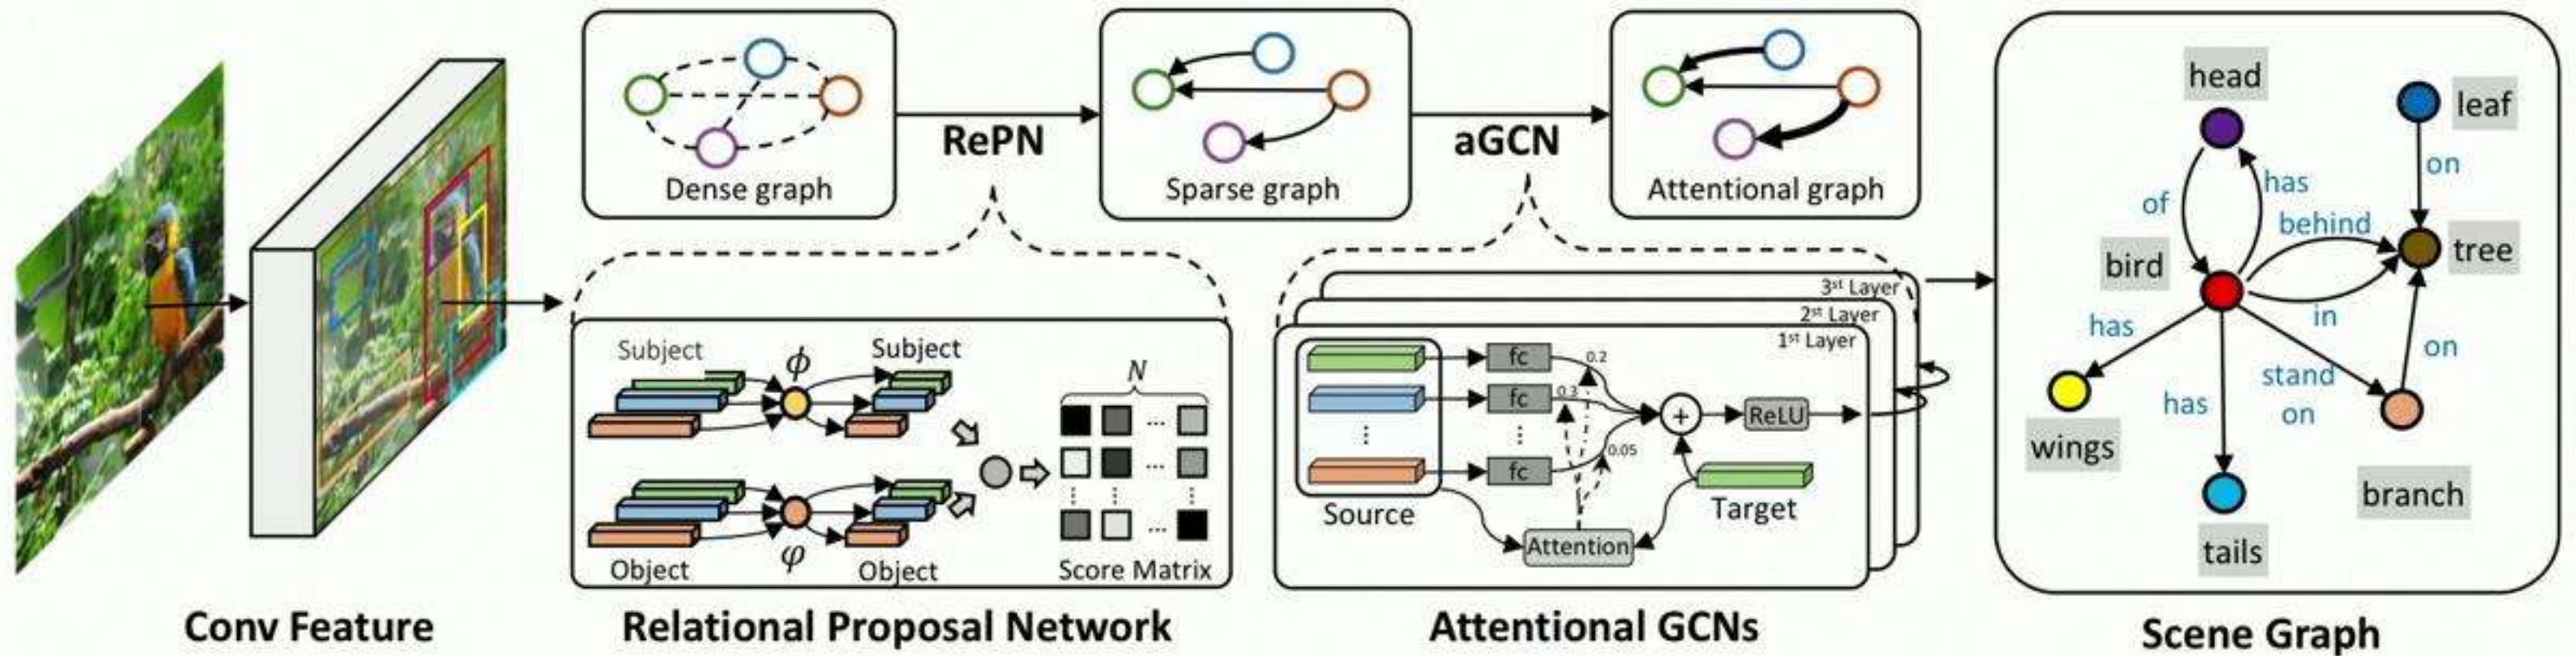
1. Relation proposal network (RePN) to learn to prune the densely connected scene graph;

Framework



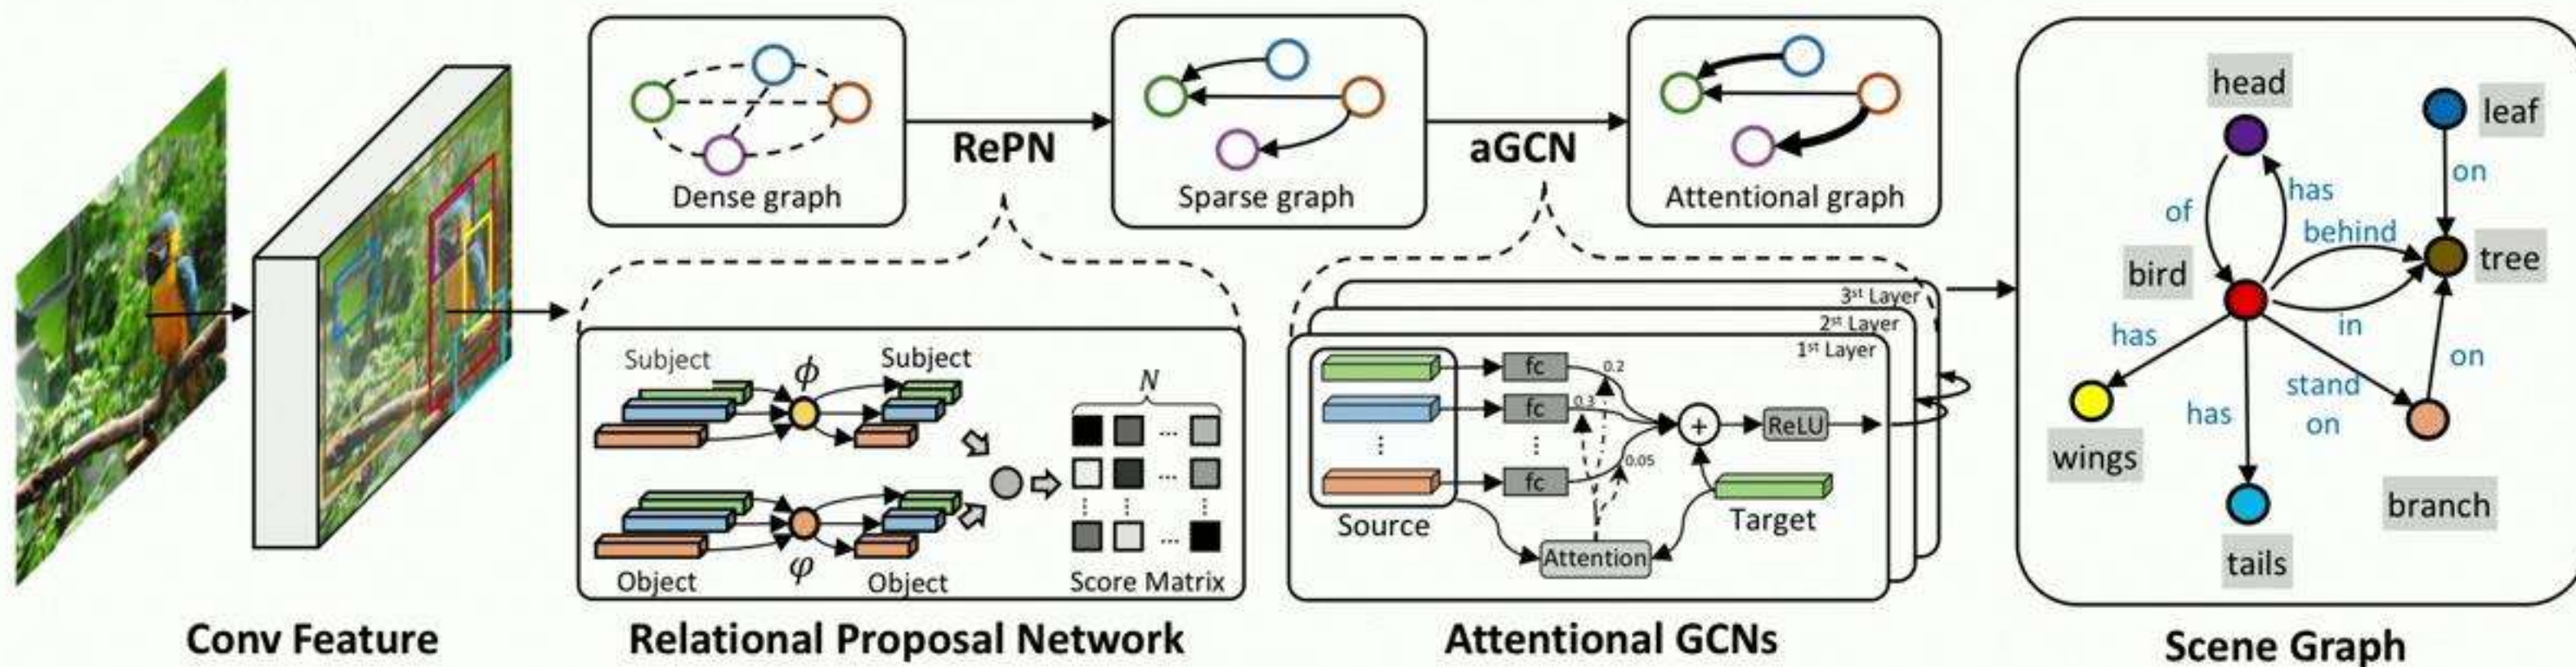
1. Relation proposal network (RePN) to learn to prune the densely connected scene graph;
2. Attentional graph convolutional networks (aGCN) to incorporate the contextual information.

Framework



1. Relation proposal network (RePN) to learn to prune the densely connected scene graph,
2. Attentional graph convolutional networks (aGCN) to incorporate the contextual information.

Framework



I — Input Image; S : Scene graph
 V — Scene graph vertices (object)
 E — Scene graph edges (relationship)
 O — Scene graph object labels
 R — Scene graph relationship labels

Region Proposal

$$P(V|I)$$

$$P(E|V, I)$$

Graph Labeling

$$P(R, O|V, E, I)$$

$$= P(S|I)$$

Relation Proposal

Relation Proposal Network

Inspired by Region Proposal Network^[1]:

Step 1: Compute Relationship-ness between subject and object:



Subj. and obj. rep.

Kernel functions^[2]

$$R(m, n) = f([x_m^o, x_n^o]) = \langle \phi(x_m^o), \phi(x_n^o) \rangle$$

Here, we use object prediction scores as the representation.

$$R(p, q) = f([x_p^o, x_q^o]) = \langle \phi(x_p^o), \phi(x_q^o) \rangle$$

Step 2: NMS for object pairs based on pair-wise IoU:

$$IoU(\{r_m^o, r_n^o\}, \{r_p^o, r_q^o\}) = \frac{I(r_m^o, r_p^o) + I(r_n^o, r_q^o)}{U(r_m^o, r_p^o) + U(r_n^o, r_q^o)}$$

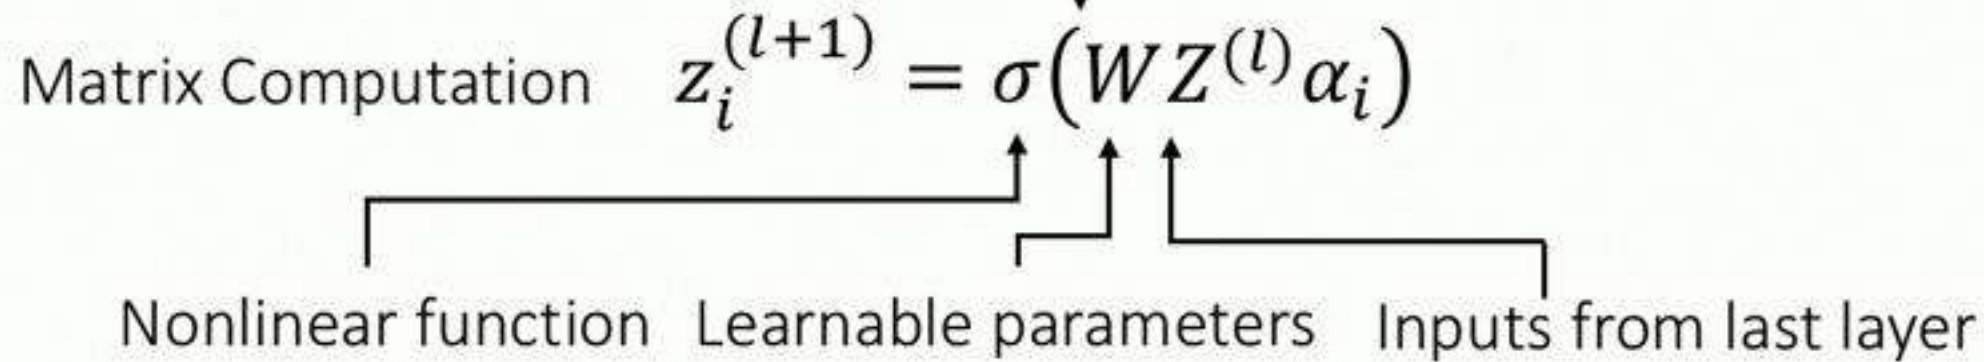
[1]. Faster R-CNN. Ren et al. Neurips 2016.

[2]. Non-local Networks. Want et al. CVPR 2018.

Attentional GCN

GCN layer with residual connection^[1]:

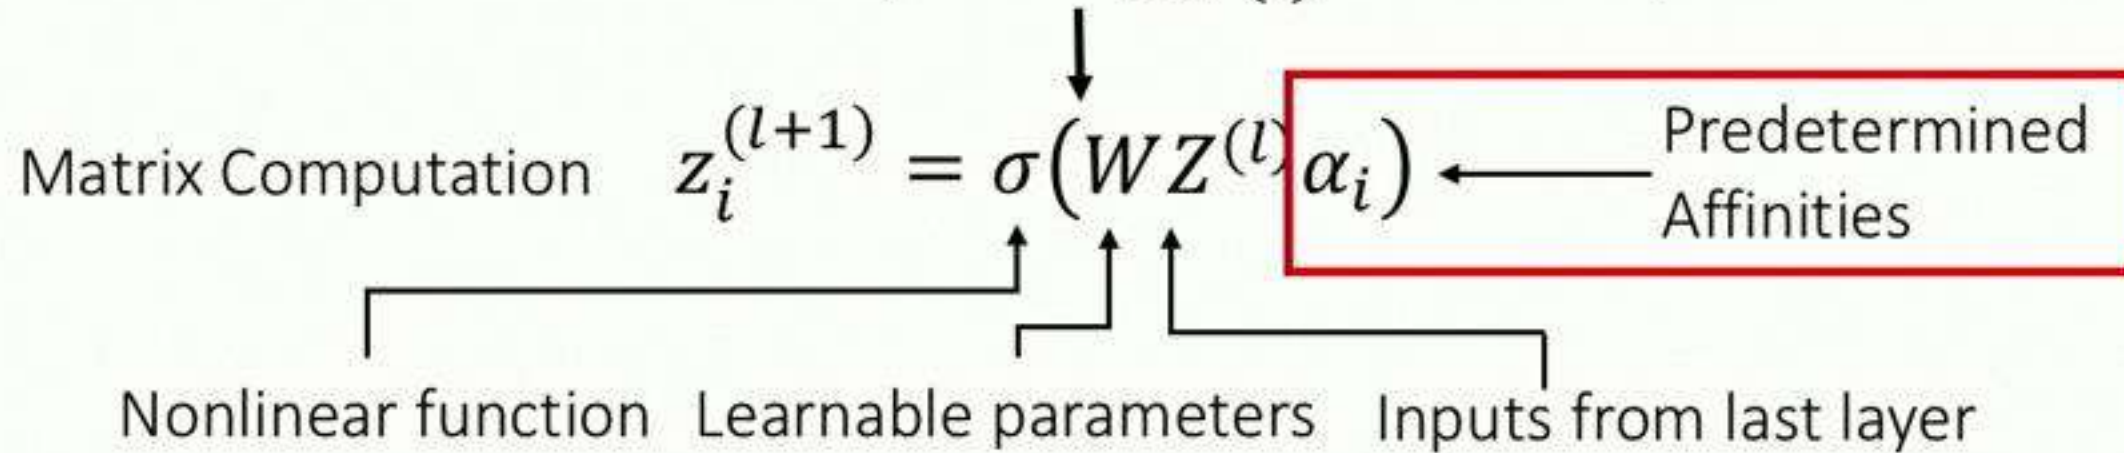
$$z_i^{(l+1)} = \sigma \left(z_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} W z_j^{(l)} \right)$$



Attentional GCN

GCN layer with residual connection^[1]:

$$z_i^{(l+1)} = \sigma \left(z_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} W z_j^{(l)} \right)$$



Attentional GCN

GCN layer with residual connection^[1]:

$$z_i^{(l+1)} = \sigma \left(z_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} W z_j^{(l)} \right)$$

Matrix Computation $z_i^{(l+1)} = \sigma(W Z^{(l)} \alpha_i)$

Nonlinear function Learnable parameters Inputs from last layer

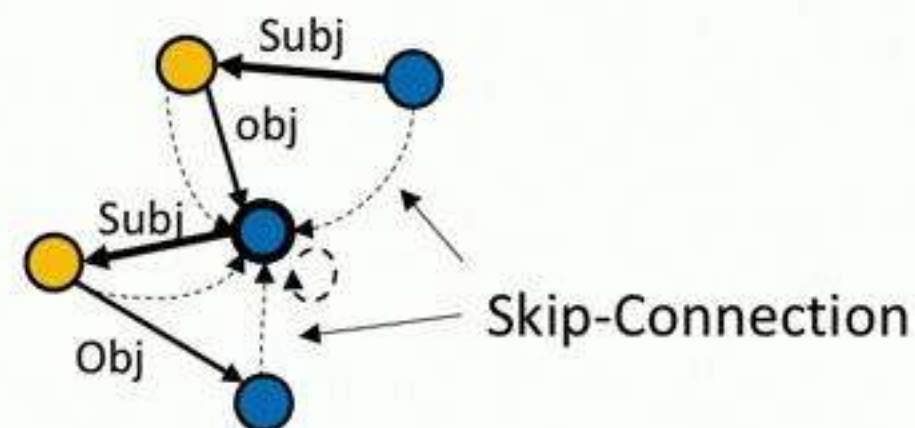
Learning the affinities!

$$u_{ij} = w_h^T \sigma \left(W_a \begin{bmatrix} z_i^{(l)} \\ z_j^{(l)} \end{bmatrix} \right)$$
$$\alpha_i = \text{softmax}(u_i)$$

Attentional GCNs (aGCN) on scene graph:

Update object representations:

$$z_i^o = \sigma \left(W^{\text{skip}} Z^o \alpha^{rs} + W^{sr} Z^r \alpha^{sr} + W^{or} Z^r \alpha^{or} \right)$$



[1]. Semi-Supervised Classification with Graph Convolutional Networks. Kipf et al. ICLR 2017

[2]. Graph Attention Networks. Veličković et al. ICLR 2018

Attentional GCN

GCN layer with residual connection:

$$z_i^{(l+1)} = \sigma \left(z_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} W z_j^{(l)} \right)$$

Matrix Computation $z_i^{(l+1)} = \sigma(W Z^{(l)} \alpha_i)$

Nonlinear function Learnable parameters Inputs from last layer

Learning the affinities!

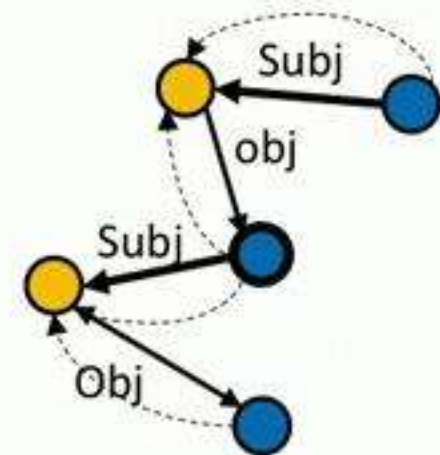
$$u_{ij} = w_h^T \sigma \left(W_a \begin{bmatrix} z_i^{(l)} \\ z_j^{(l)} \end{bmatrix} \right)$$

$$\alpha_i = \text{softmax}(u_i)$$

Attentional GCNs (aGCN) on scene graph:

Update predicate representations:

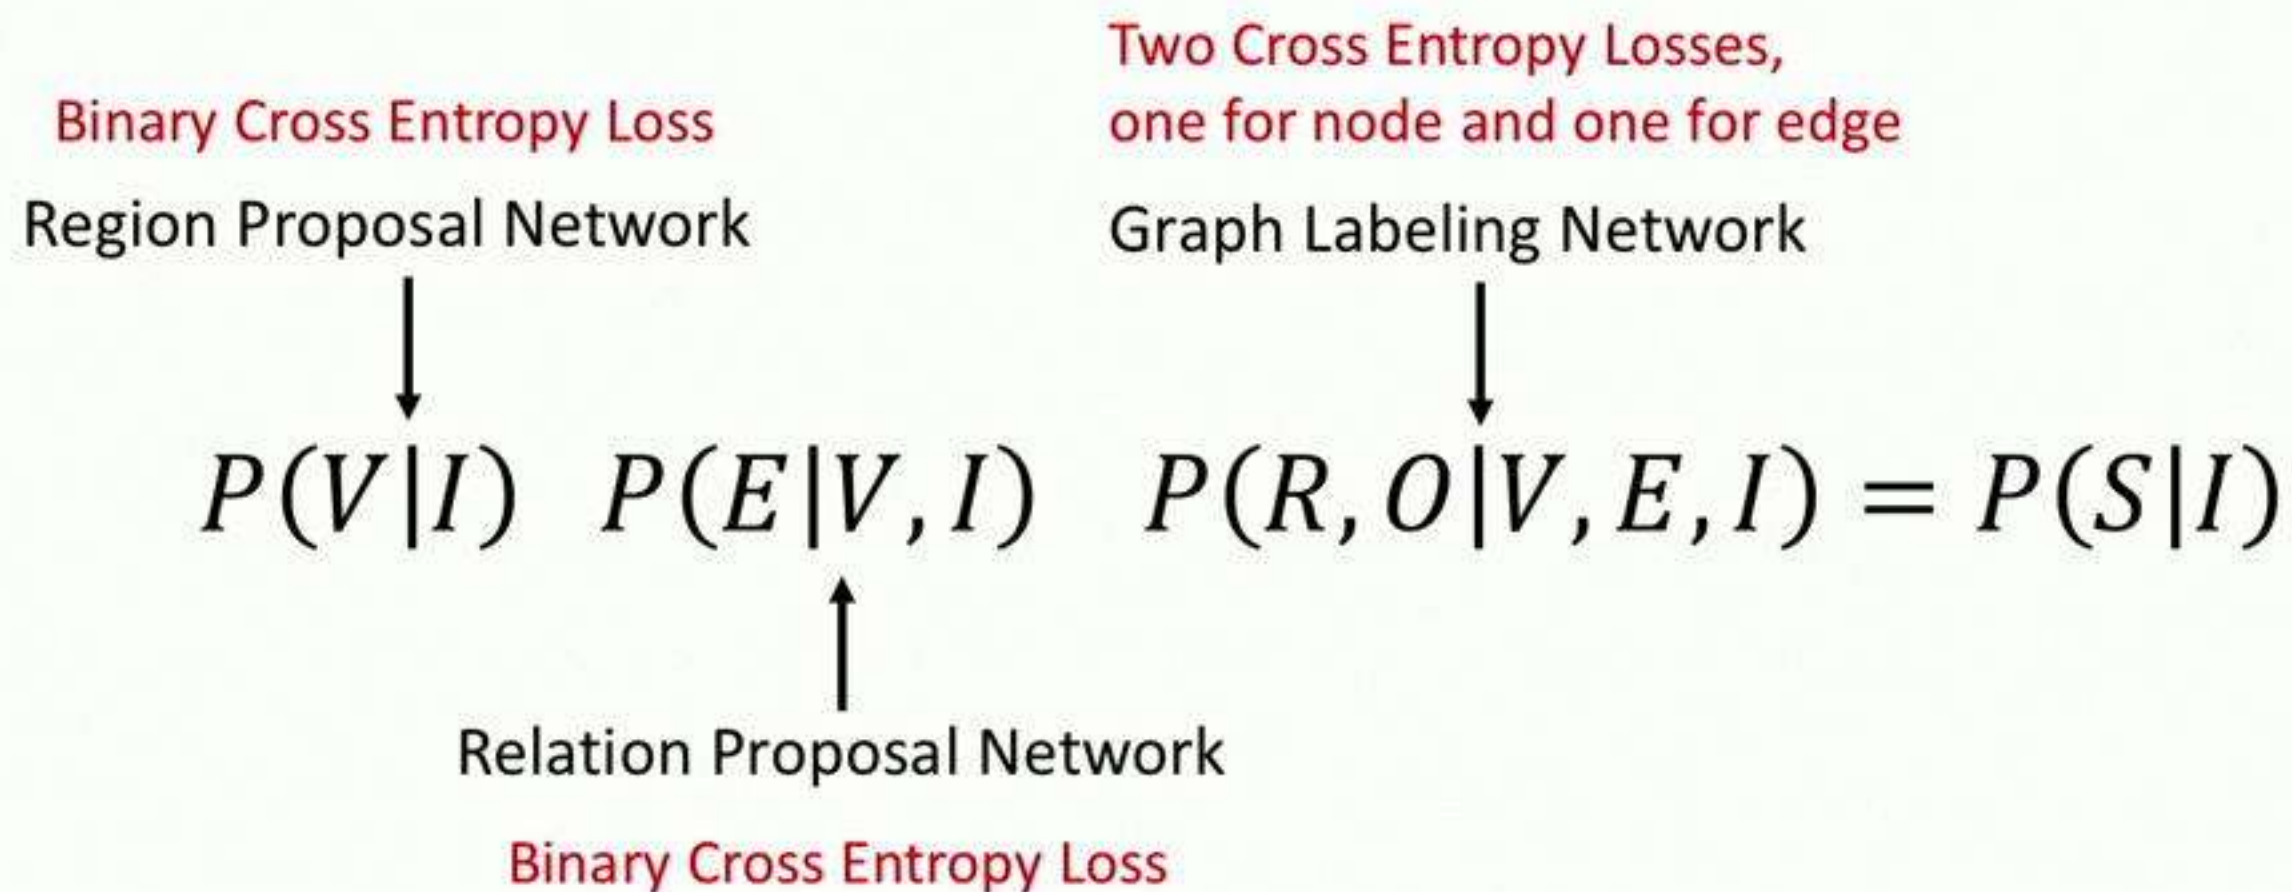
$$z_i^r = \sigma(z_i^r + W^{rs} Z^o \alpha^{rs} + W^{ro} Z^o \alpha^{ro})$$



[1]. Semi-Supervised Classification with Graph Convolutional Networks. Kipf et al. ICLR 2017

[2]. Graph Attention Networks. Veličković et al. ICLR 2018

Training



Metrics

Assume there are N objects extracted from an image, then $N * (N - 1)$ edges

Step 1: Take maximum for object scores and predicate scores, excluding background class.

Step 2: Compute relationship scores: $Rel(i, j) = Subj(i) * Obj(j) * Pred(i, j)$

Step 3: Sort the relationship triplets in a descending order:

Step 4: Compute the triplet recalls (Recall@50, Recall@100) based on the ground-truth

$$\textbf{SGGen: } Recall = \frac{C(T_{pred} \text{ and } T_{gt})}{N(T_{gt})} \quad \text{IoU} > 0.5$$

PhrCls: all object locations are known

PredCls: all object locations and labels are known

Experiments

Table. Implementation Details.

Dataset	Backbone	#objects	#predicates	Metrics
Visual Genome Train: 75,651 Test: 32,422	VGG-16 Faster R-CNN ^[1]	150	50	PredCls,SGCls, SGGen, mAP

[1] A Faster Implementation of Faster R-CNN. Yang and Lu et al.

Metrics

Assume there are N objects extracted from an image, then $N * (N - 1)$ edges

Step 1: Take maximum for object scores and predicate scores, excluding background class.

Step 2: Compute relationship scores: $Rel(i, j) = Subj(i) * Obj(j) * Pred(i, j)$

Step 3: Sort the relationship triplets in a descending order:

Step 4: Compute the triplet recalls (Recall@50, Recall@100) based on the ground-truth

$$\textbf{SGGen: } Recall = \frac{C(T_{pred} \text{ and } T_{gt})}{N(T_{gt})} \quad \text{IoU} > 0.5$$

PhrCls: all object locations are known

PredCls: all object locations and labels are known

Experiments

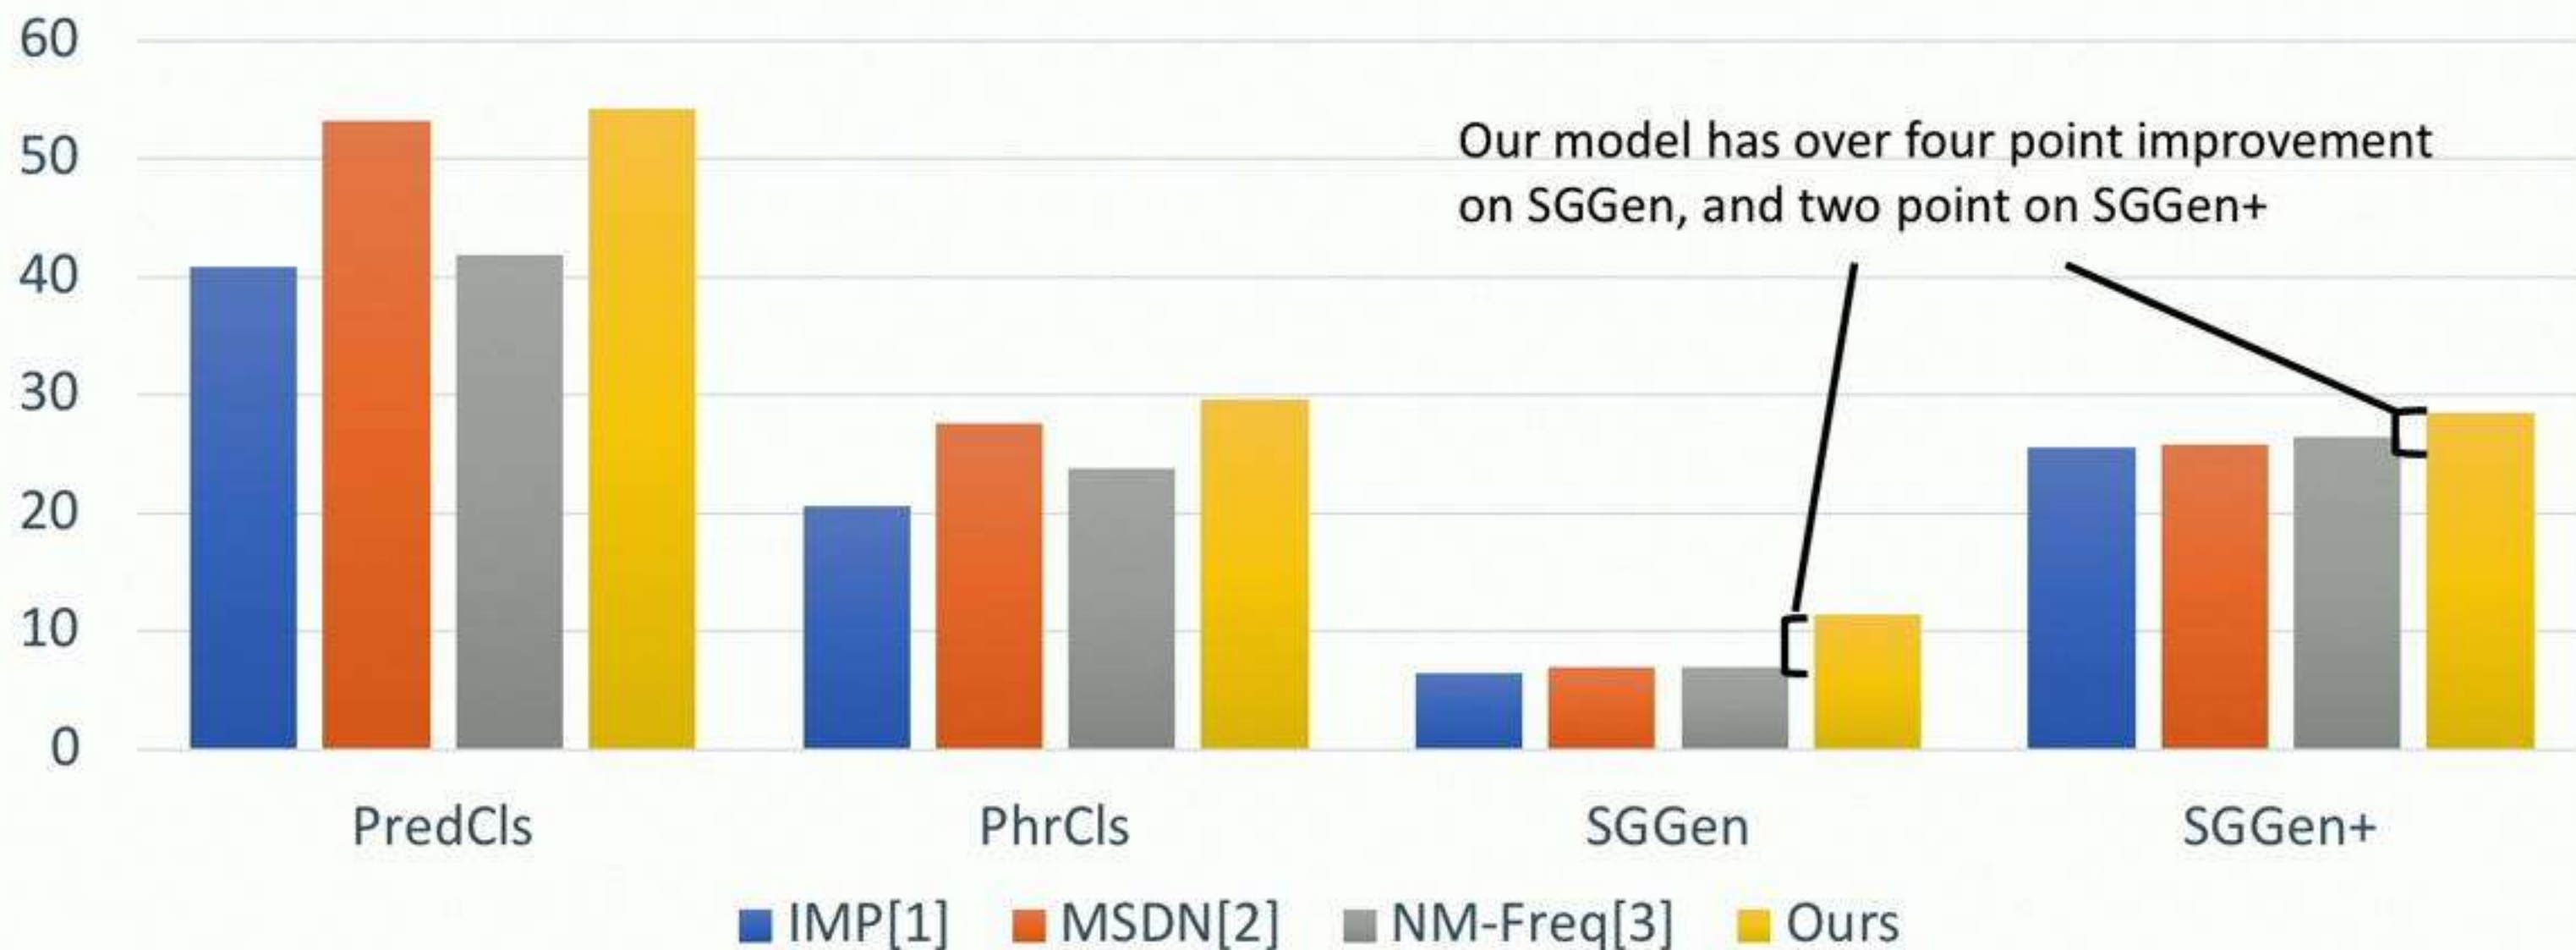
Table. Implementation Details.

Dataset	Backbone	#objects	#predicates	Metrics
Visual Genome Train: 75,651 Test: 32,422	VGG-16 Faster R-CNN ^[1]	150	50	PredCls,SGCls, SGGen, mAP

[1] A Faster Implementation of Faster R-CNN. Yang and Lu et al.

Comparing with Previous Work

Recall@50

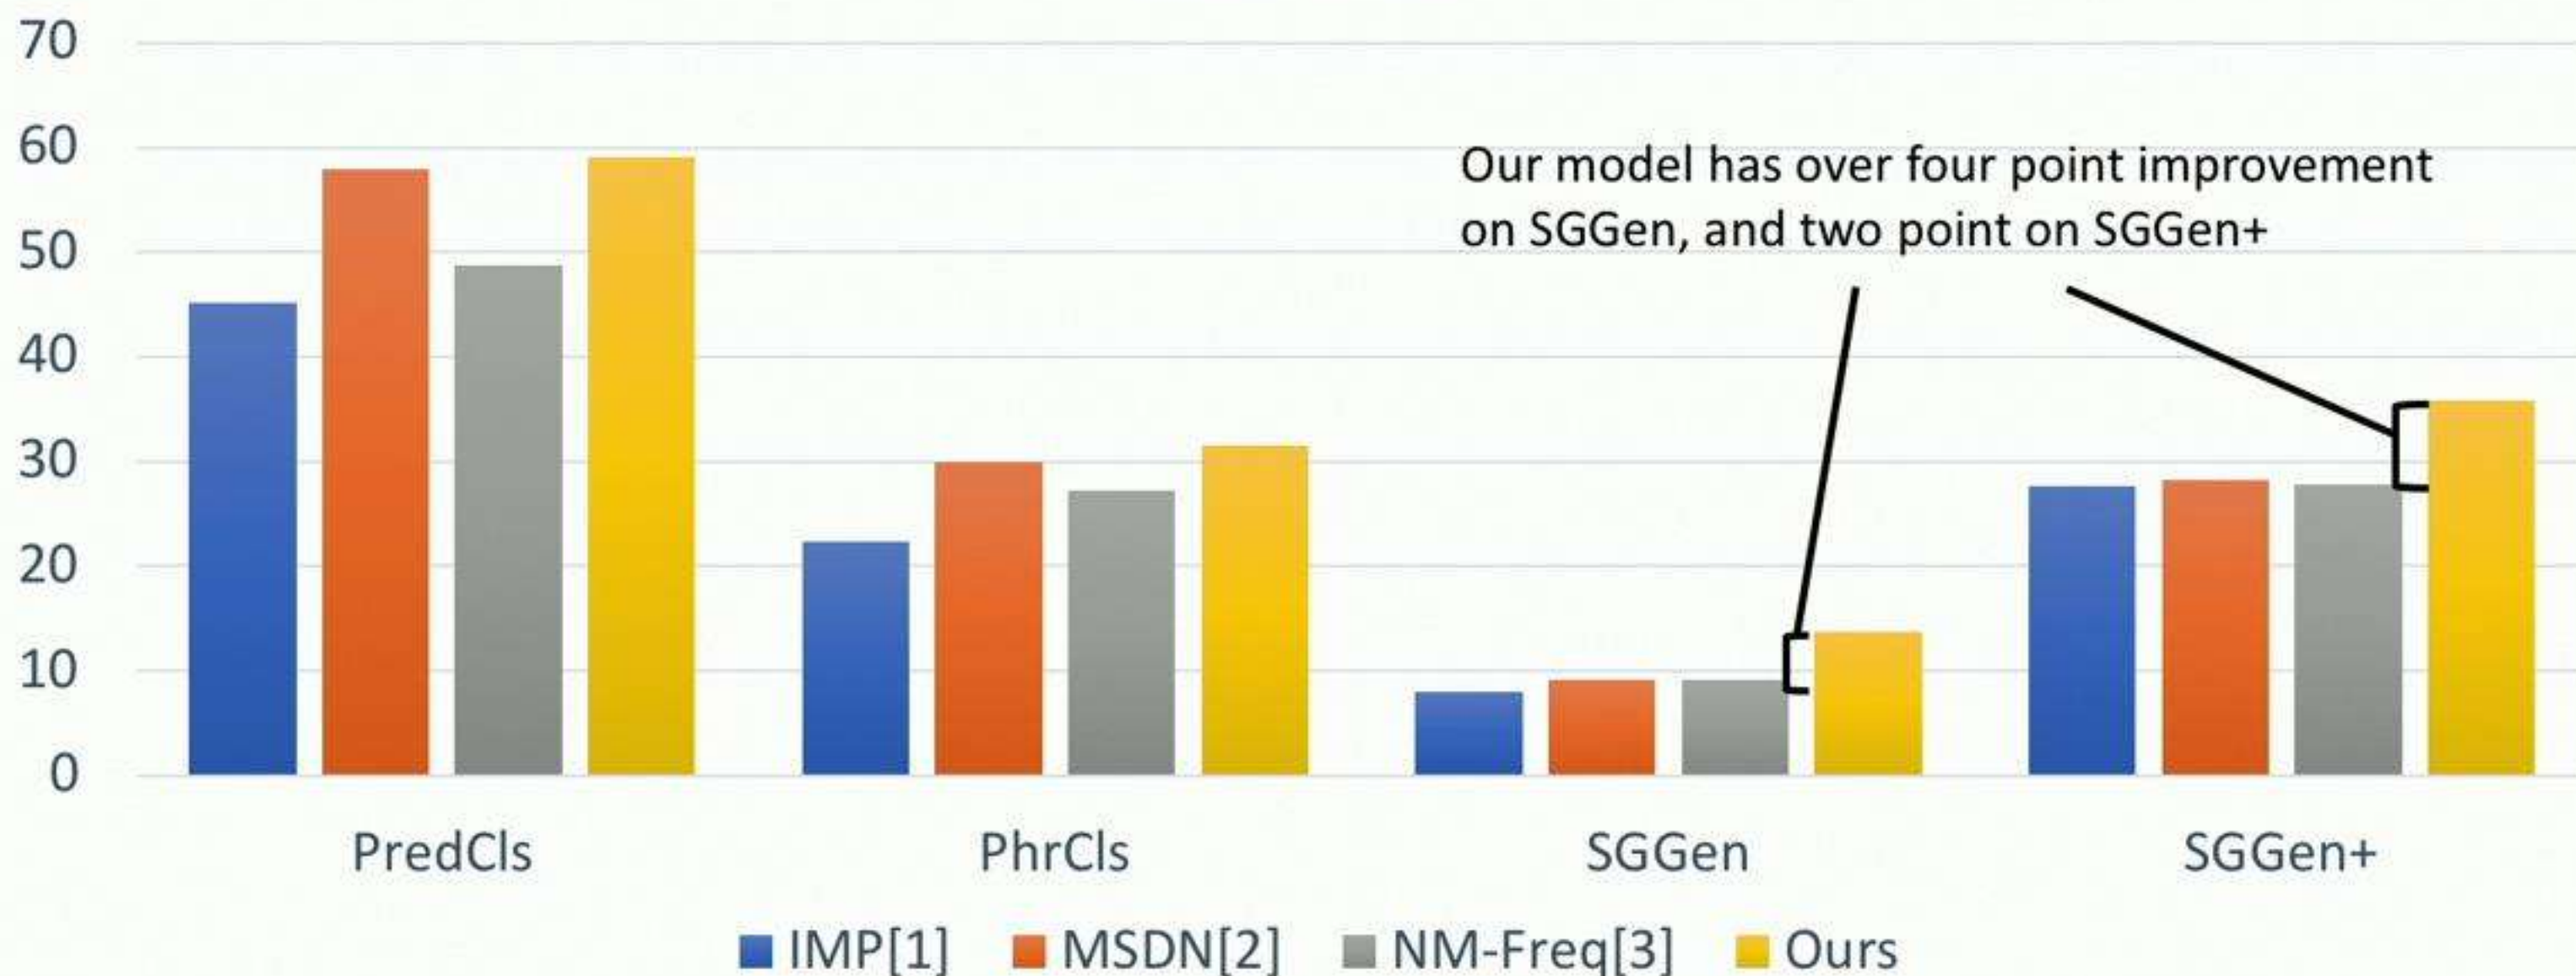


Our model has over four point improvement on SGGen, and two point on SGGen+

- [1] Scene Graph Generation by Iterative Message Passing. Xu et al. CVPR 2017
- [2] Scene Graph Generations from Objects, Phrases and Captions. Li et al. ICCV 2017
- [3] Neural Motif: Scene Graph Parsing with Global Context. Zellers et al. CVPR 2018

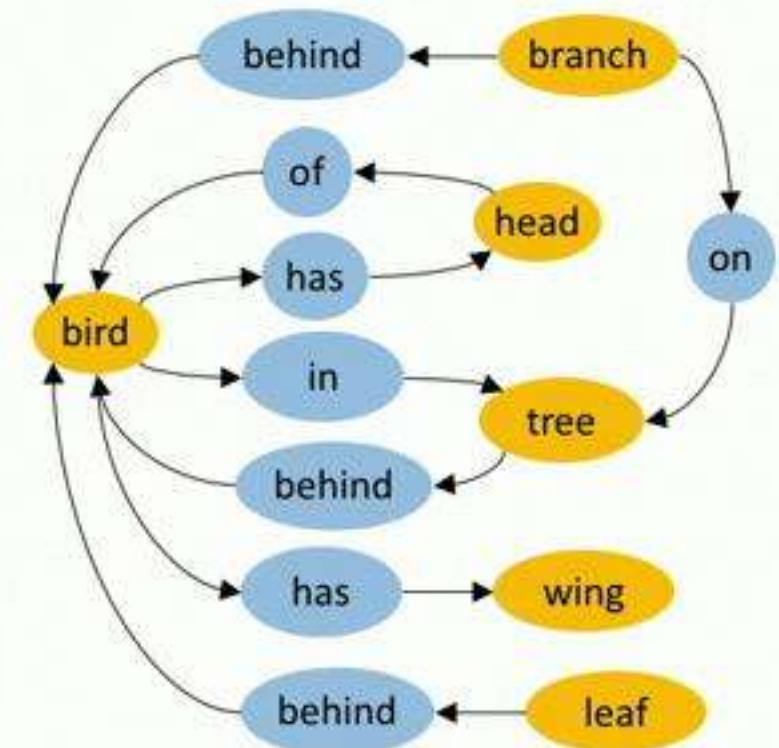
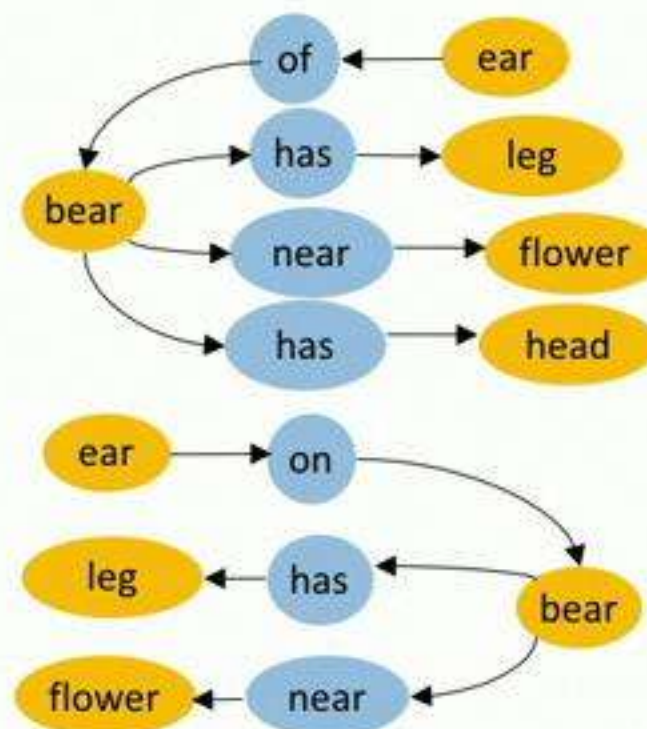
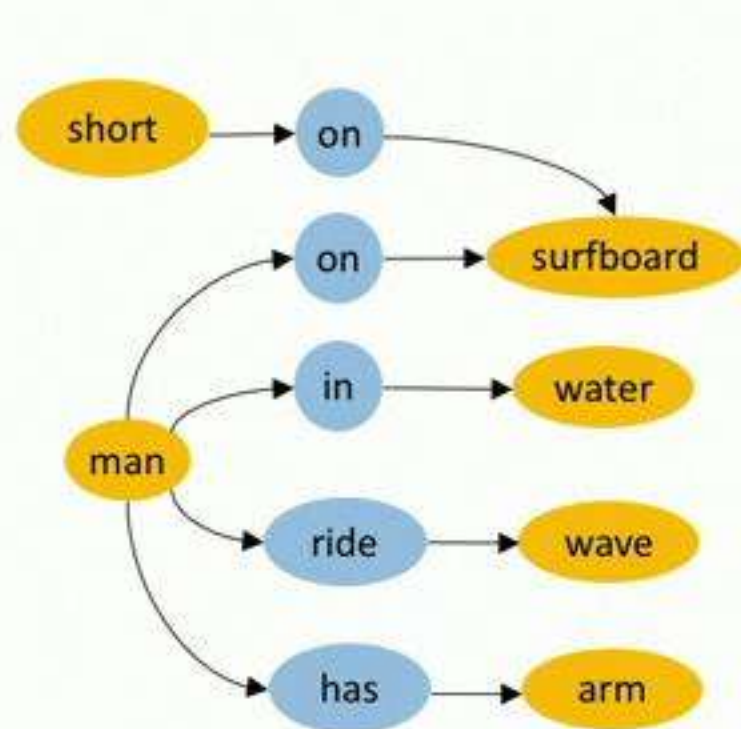
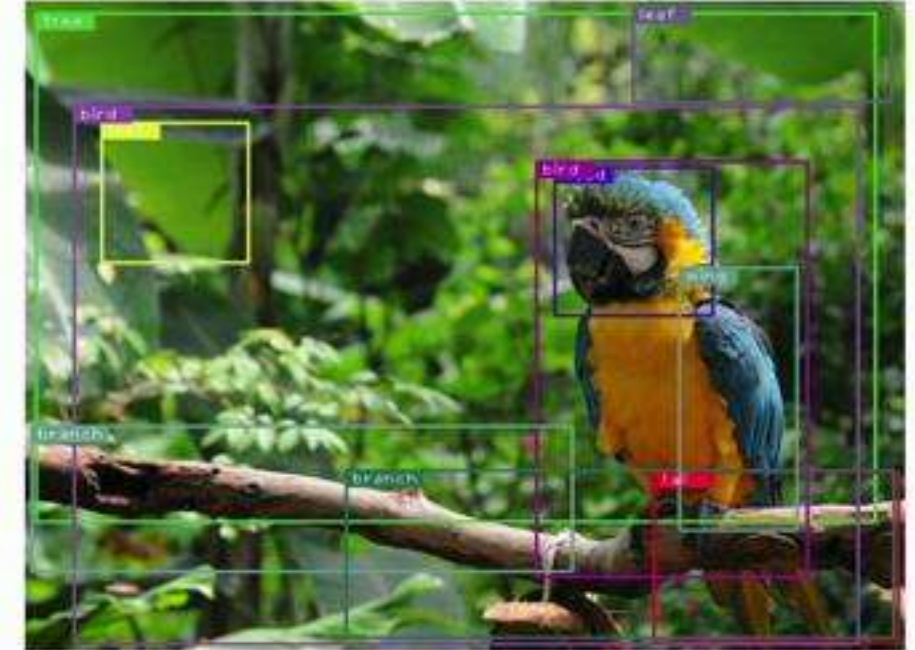
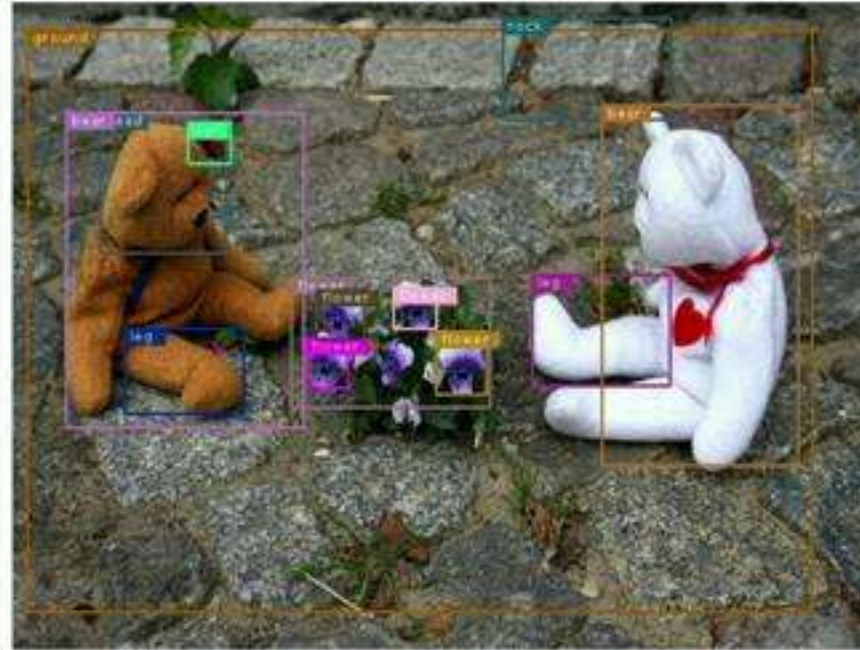
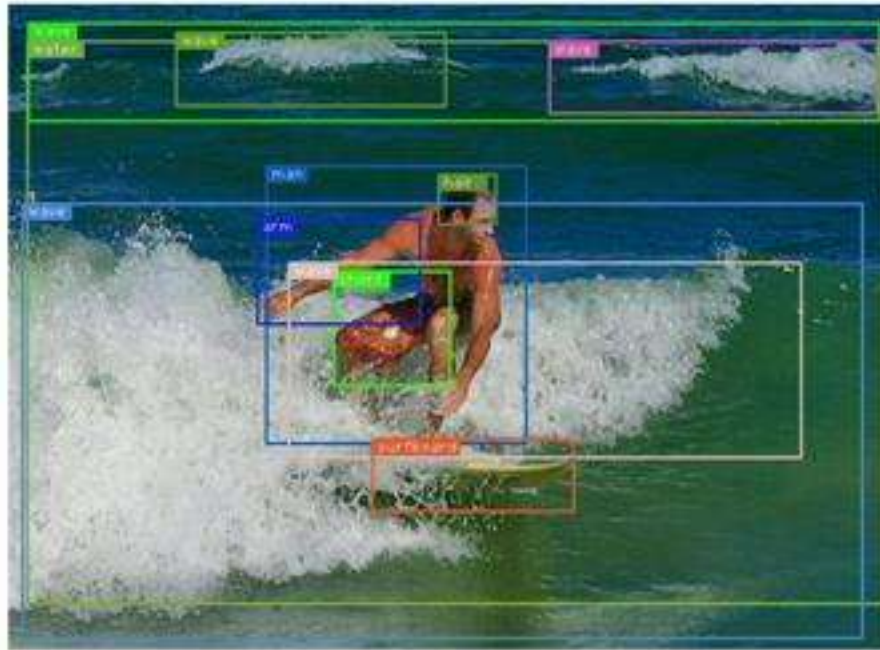
Comparing with Previous Work

Recall@100



- [1] Scene Graph Generation by Iterative Message Passing. Xu et al. CVPR 2017
- [2] Scene Graph Generations from Objects, Phrases and Captions. Li et al. ICCV 2017
- [3] Neural Motif: Scene Graph Parsing with Global Context. Zellers et al. CVPR 2018

Qualitative Results



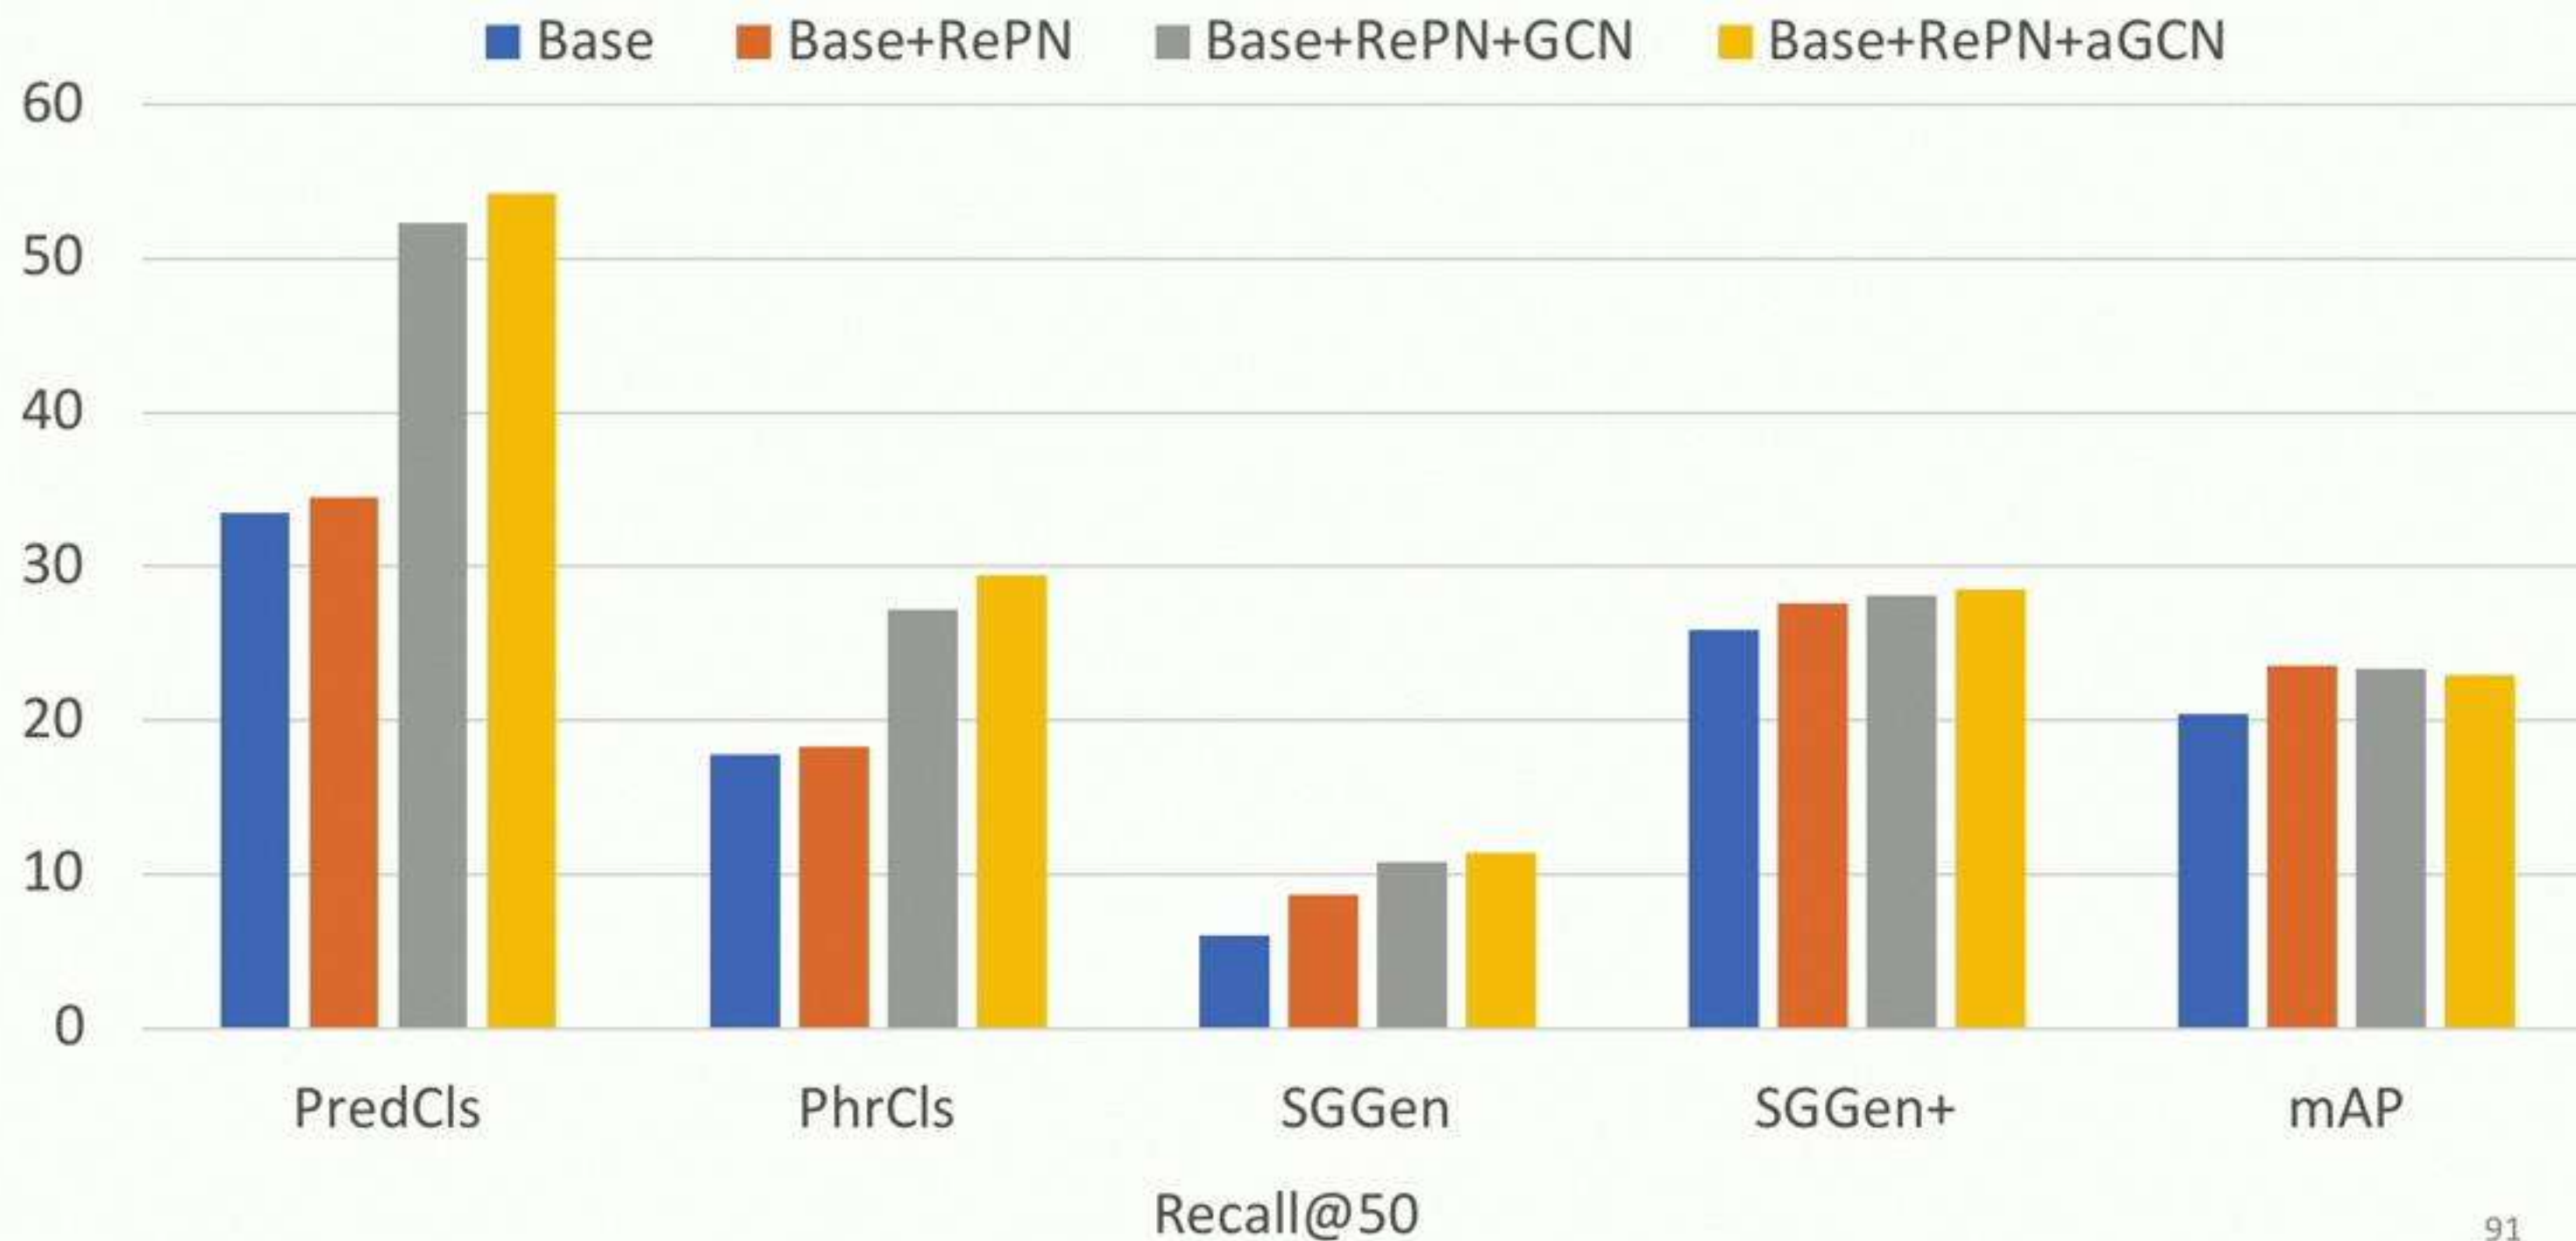
Common Sense Emerges

We extract the weights in the score-level aGCN layer, and sort it in descending order.

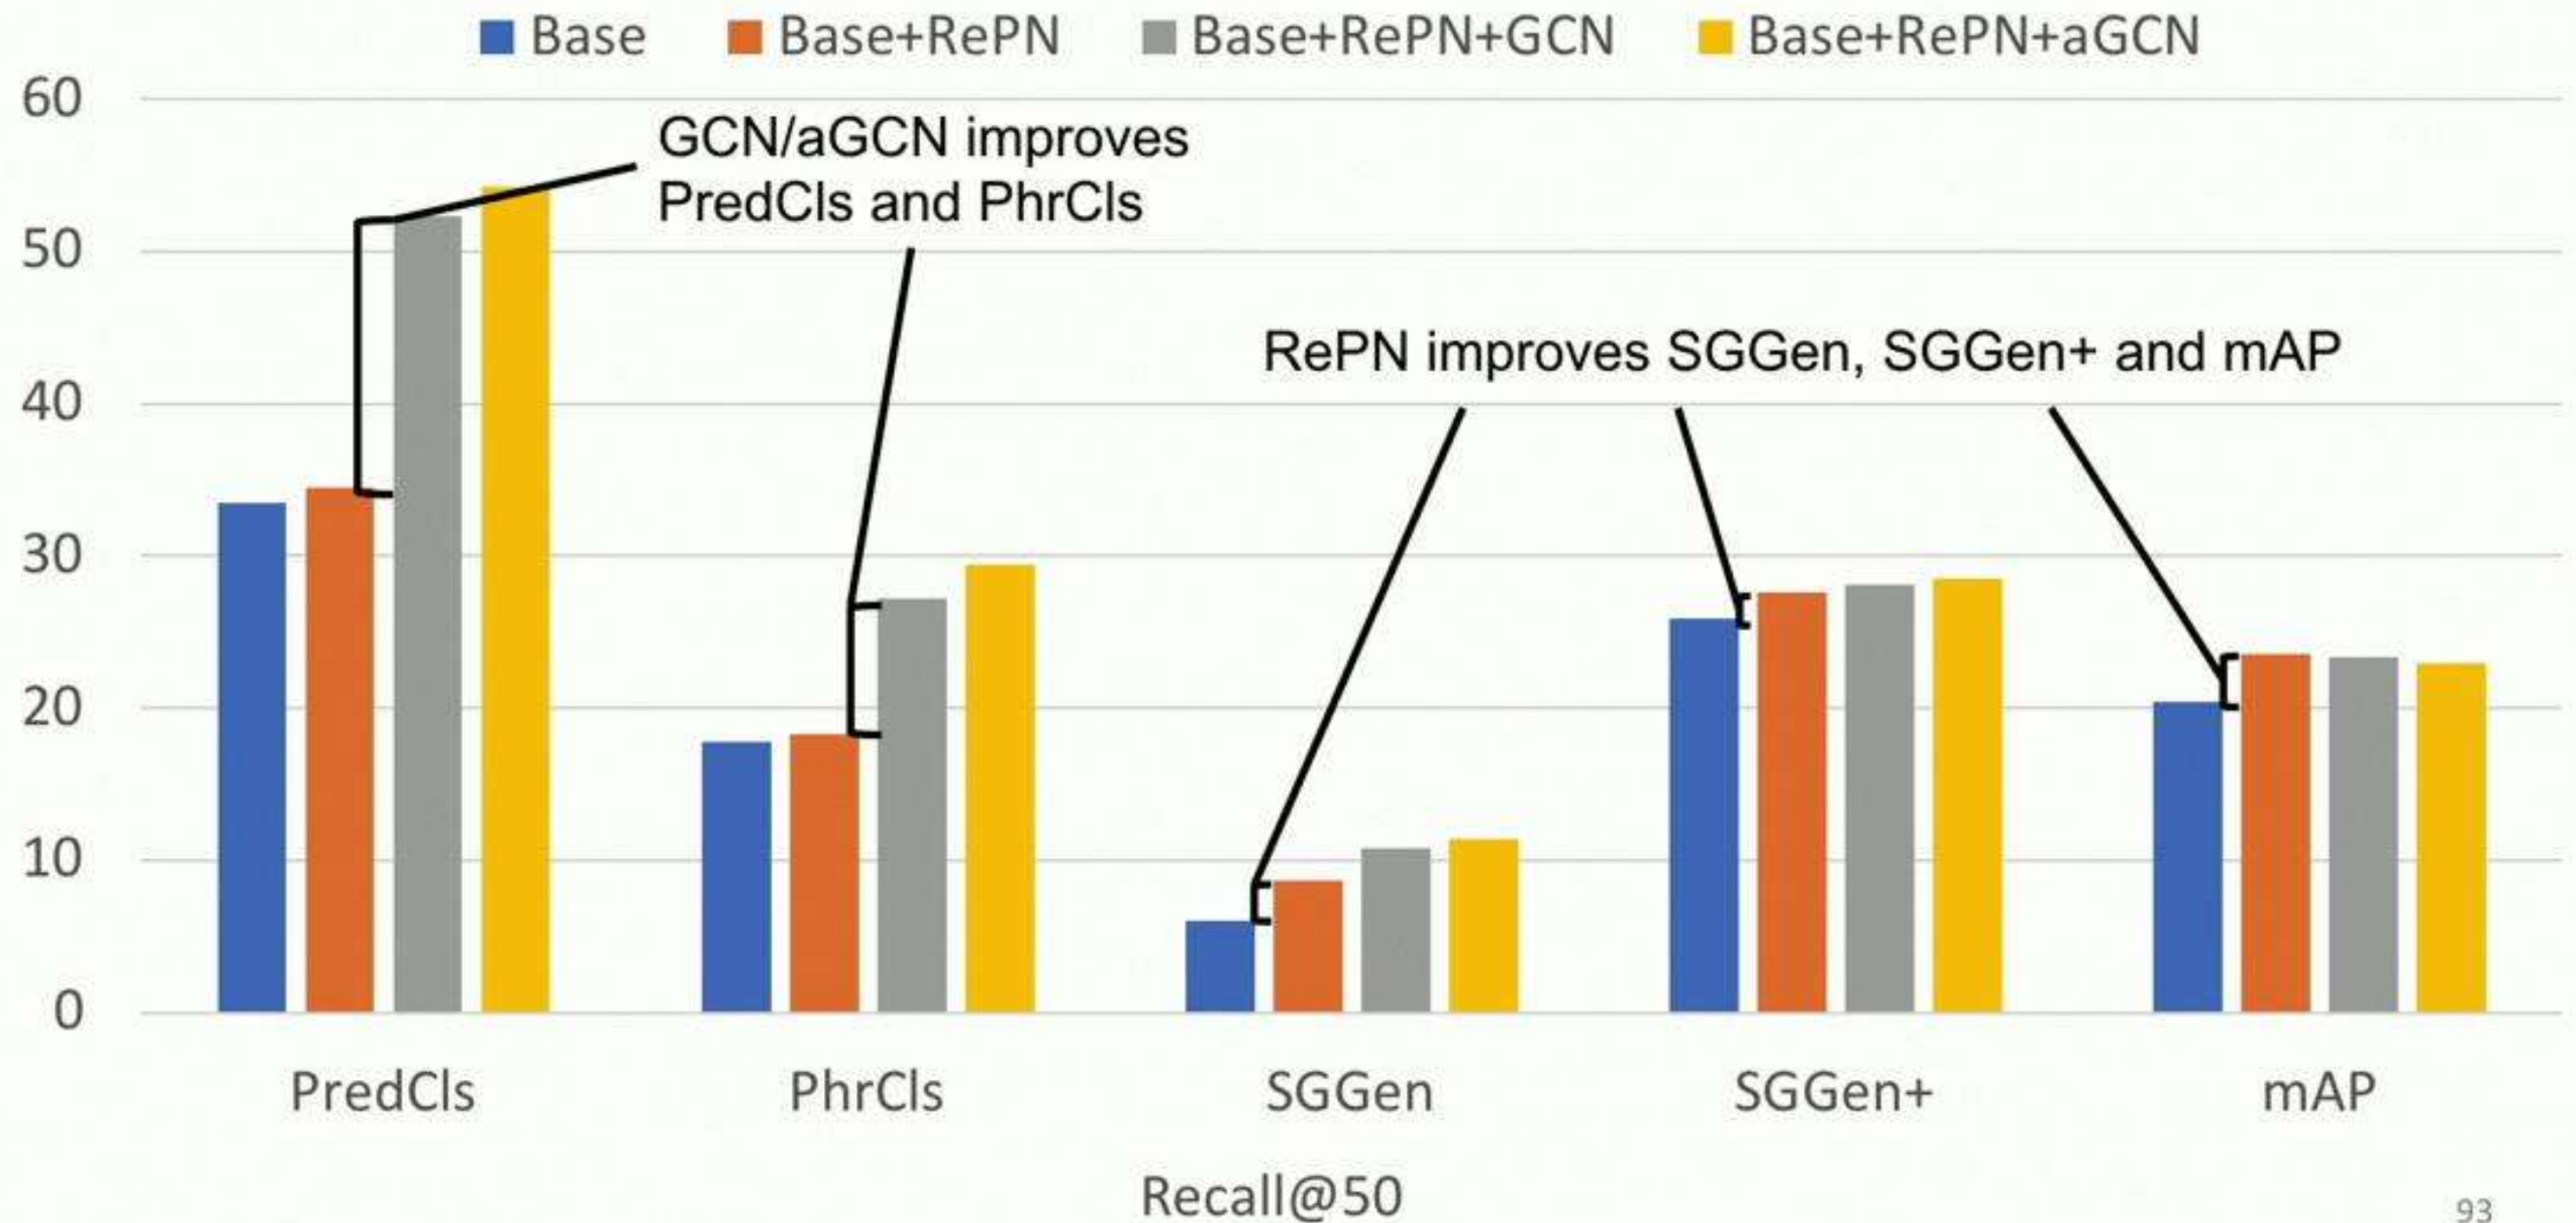
Object-Object Co-Occurrence					
Object	Top-1	Top-2	Object	Top-1	Top-2
boat	water	beach	girl	woman	hair
plane	wing	tail	cow	horse	dog
clock	building	root	sidewalk	street	bus
bottle	cup	glass	handle	plate	food
bus	truck	vehicle	snow	pole	ski

Object-Predicate Co-Occurrence					
Object	Top-1	Top-2	Object	Top-1	Top-2
hat	hold	wear	kite	watch	look at
boat	in	sit in	girl	look at	watch
umbrella	carry	hold	jacket	wear	with
track	with	on	stripe	on	has
sidewalk	at	walk on	snow	on	near

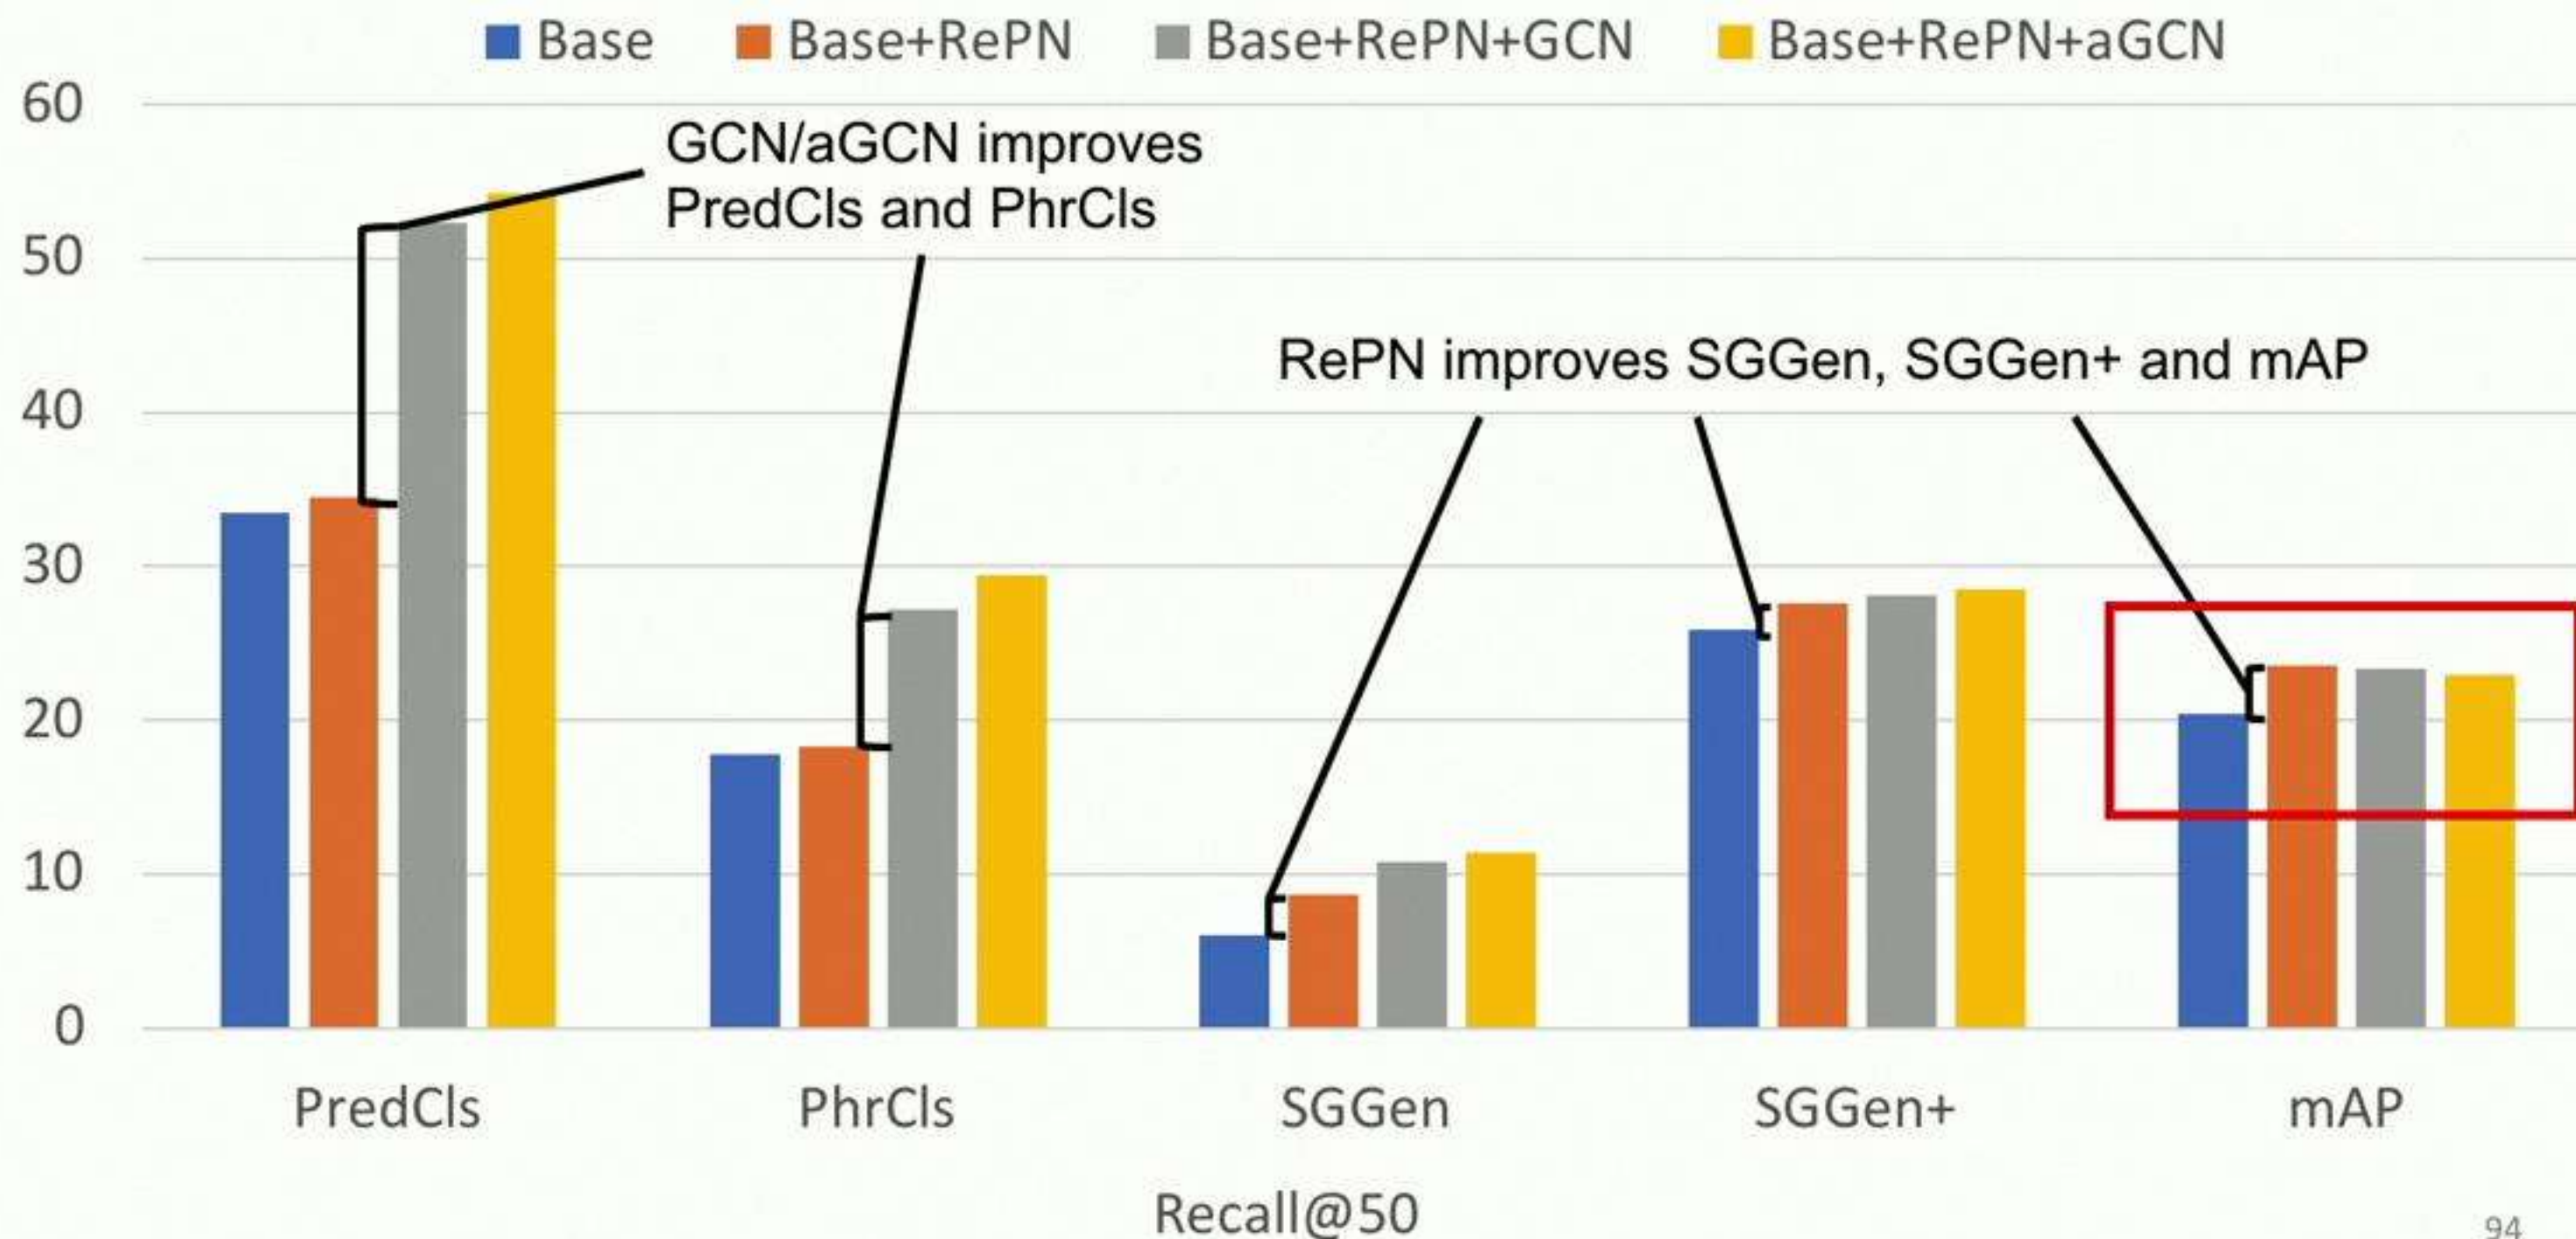
Ablation Study



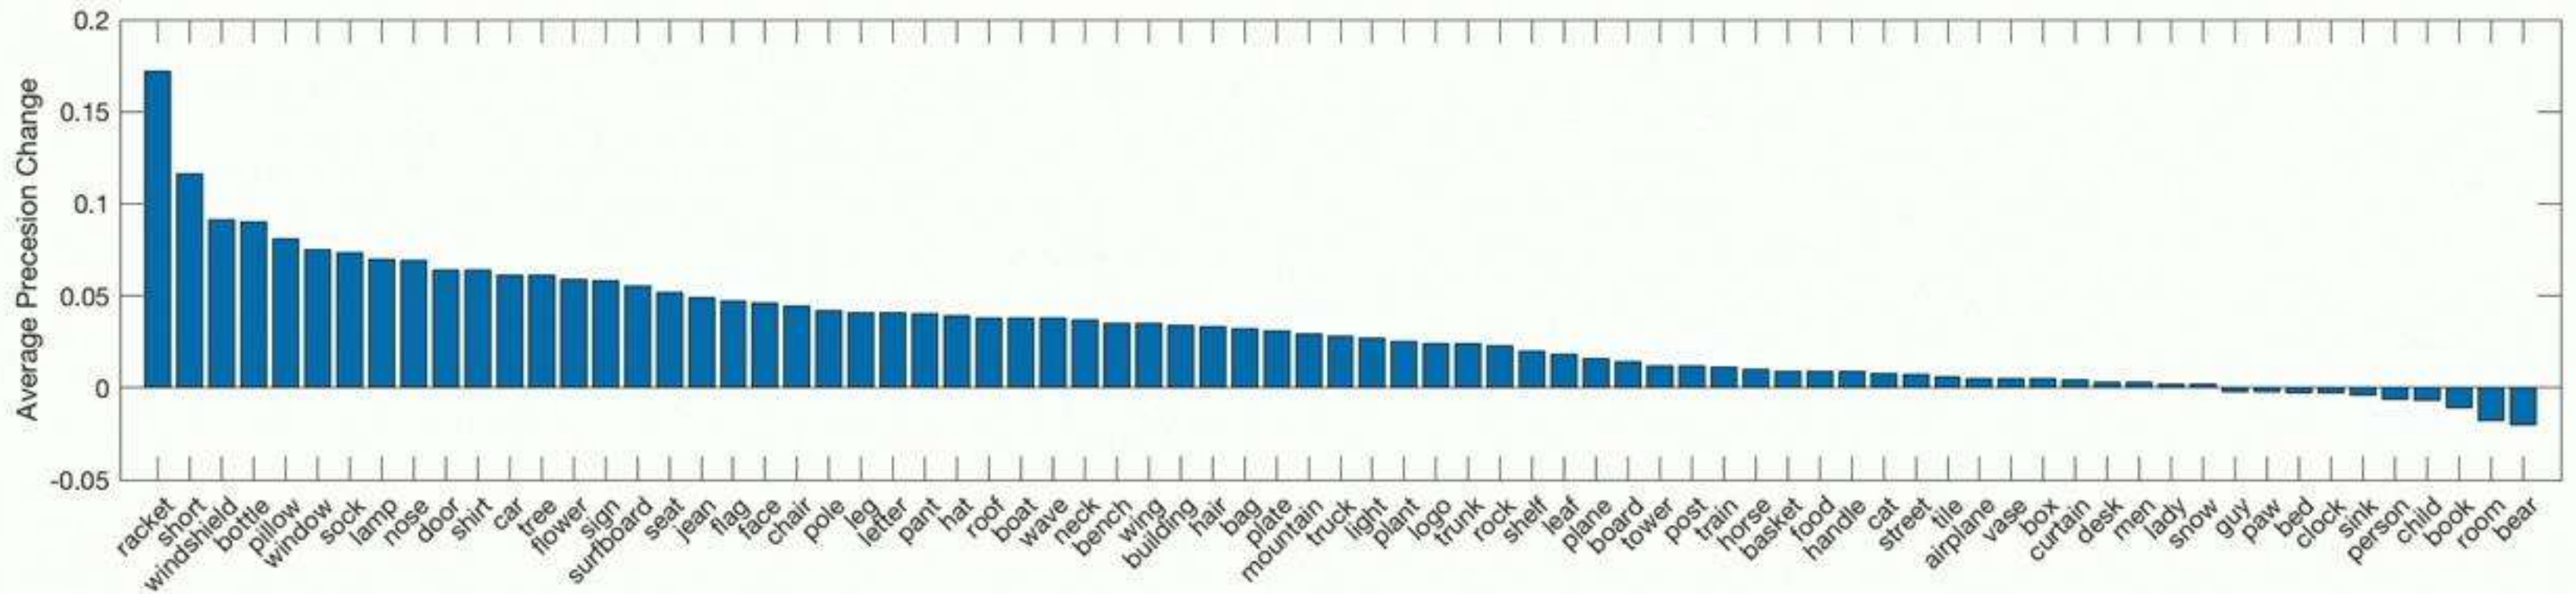
Ablation Study



Ablation Study

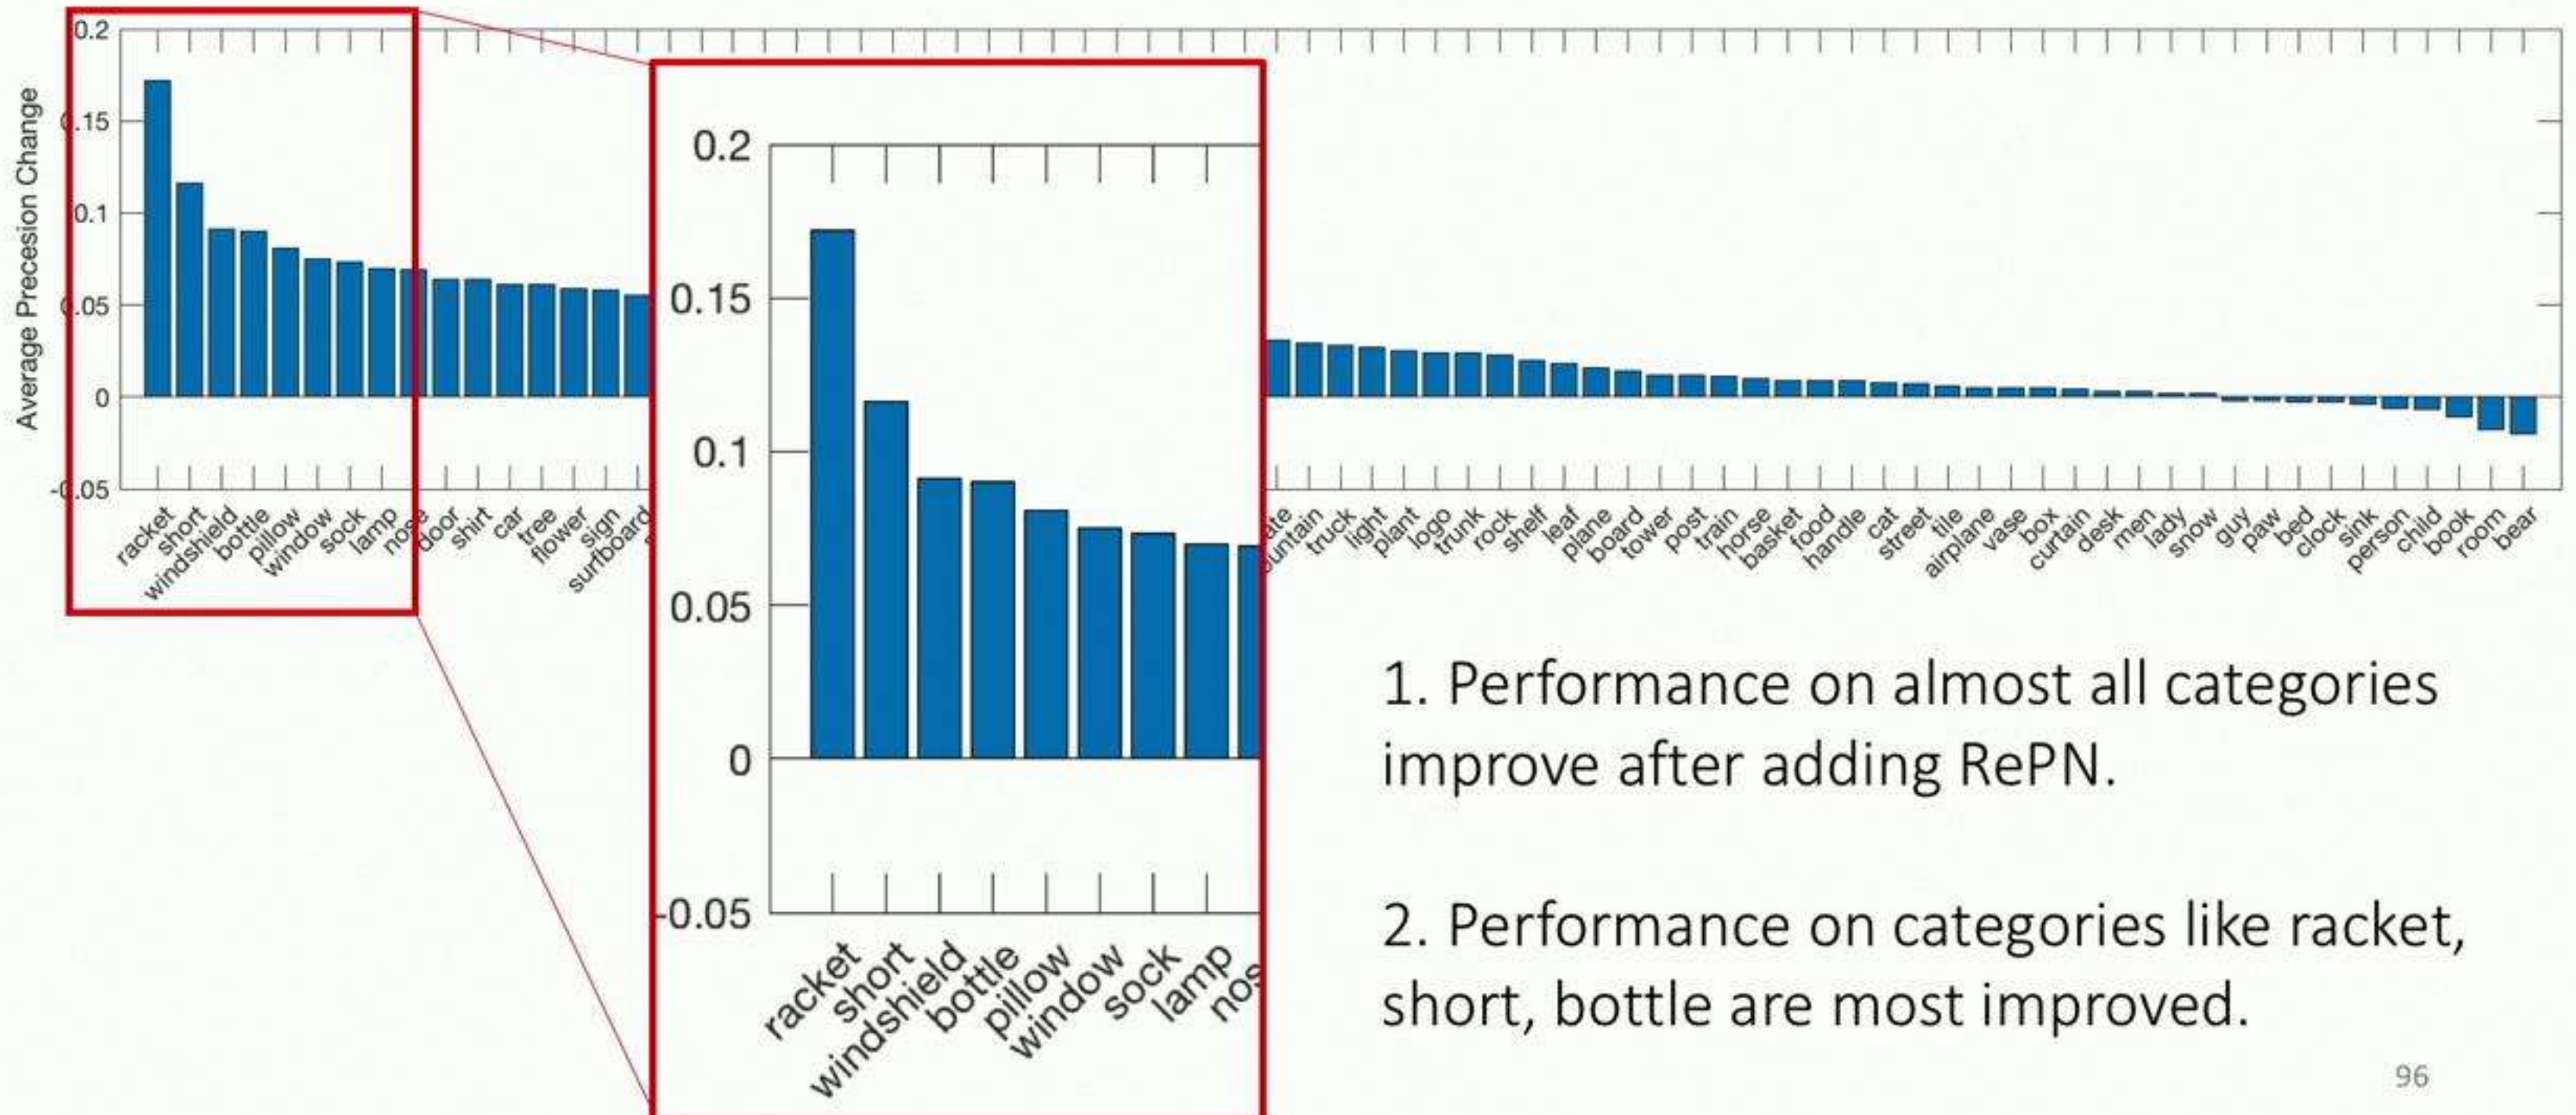


Object Detection Investigation



1. Performance on almost all categories improve after adding RePN.

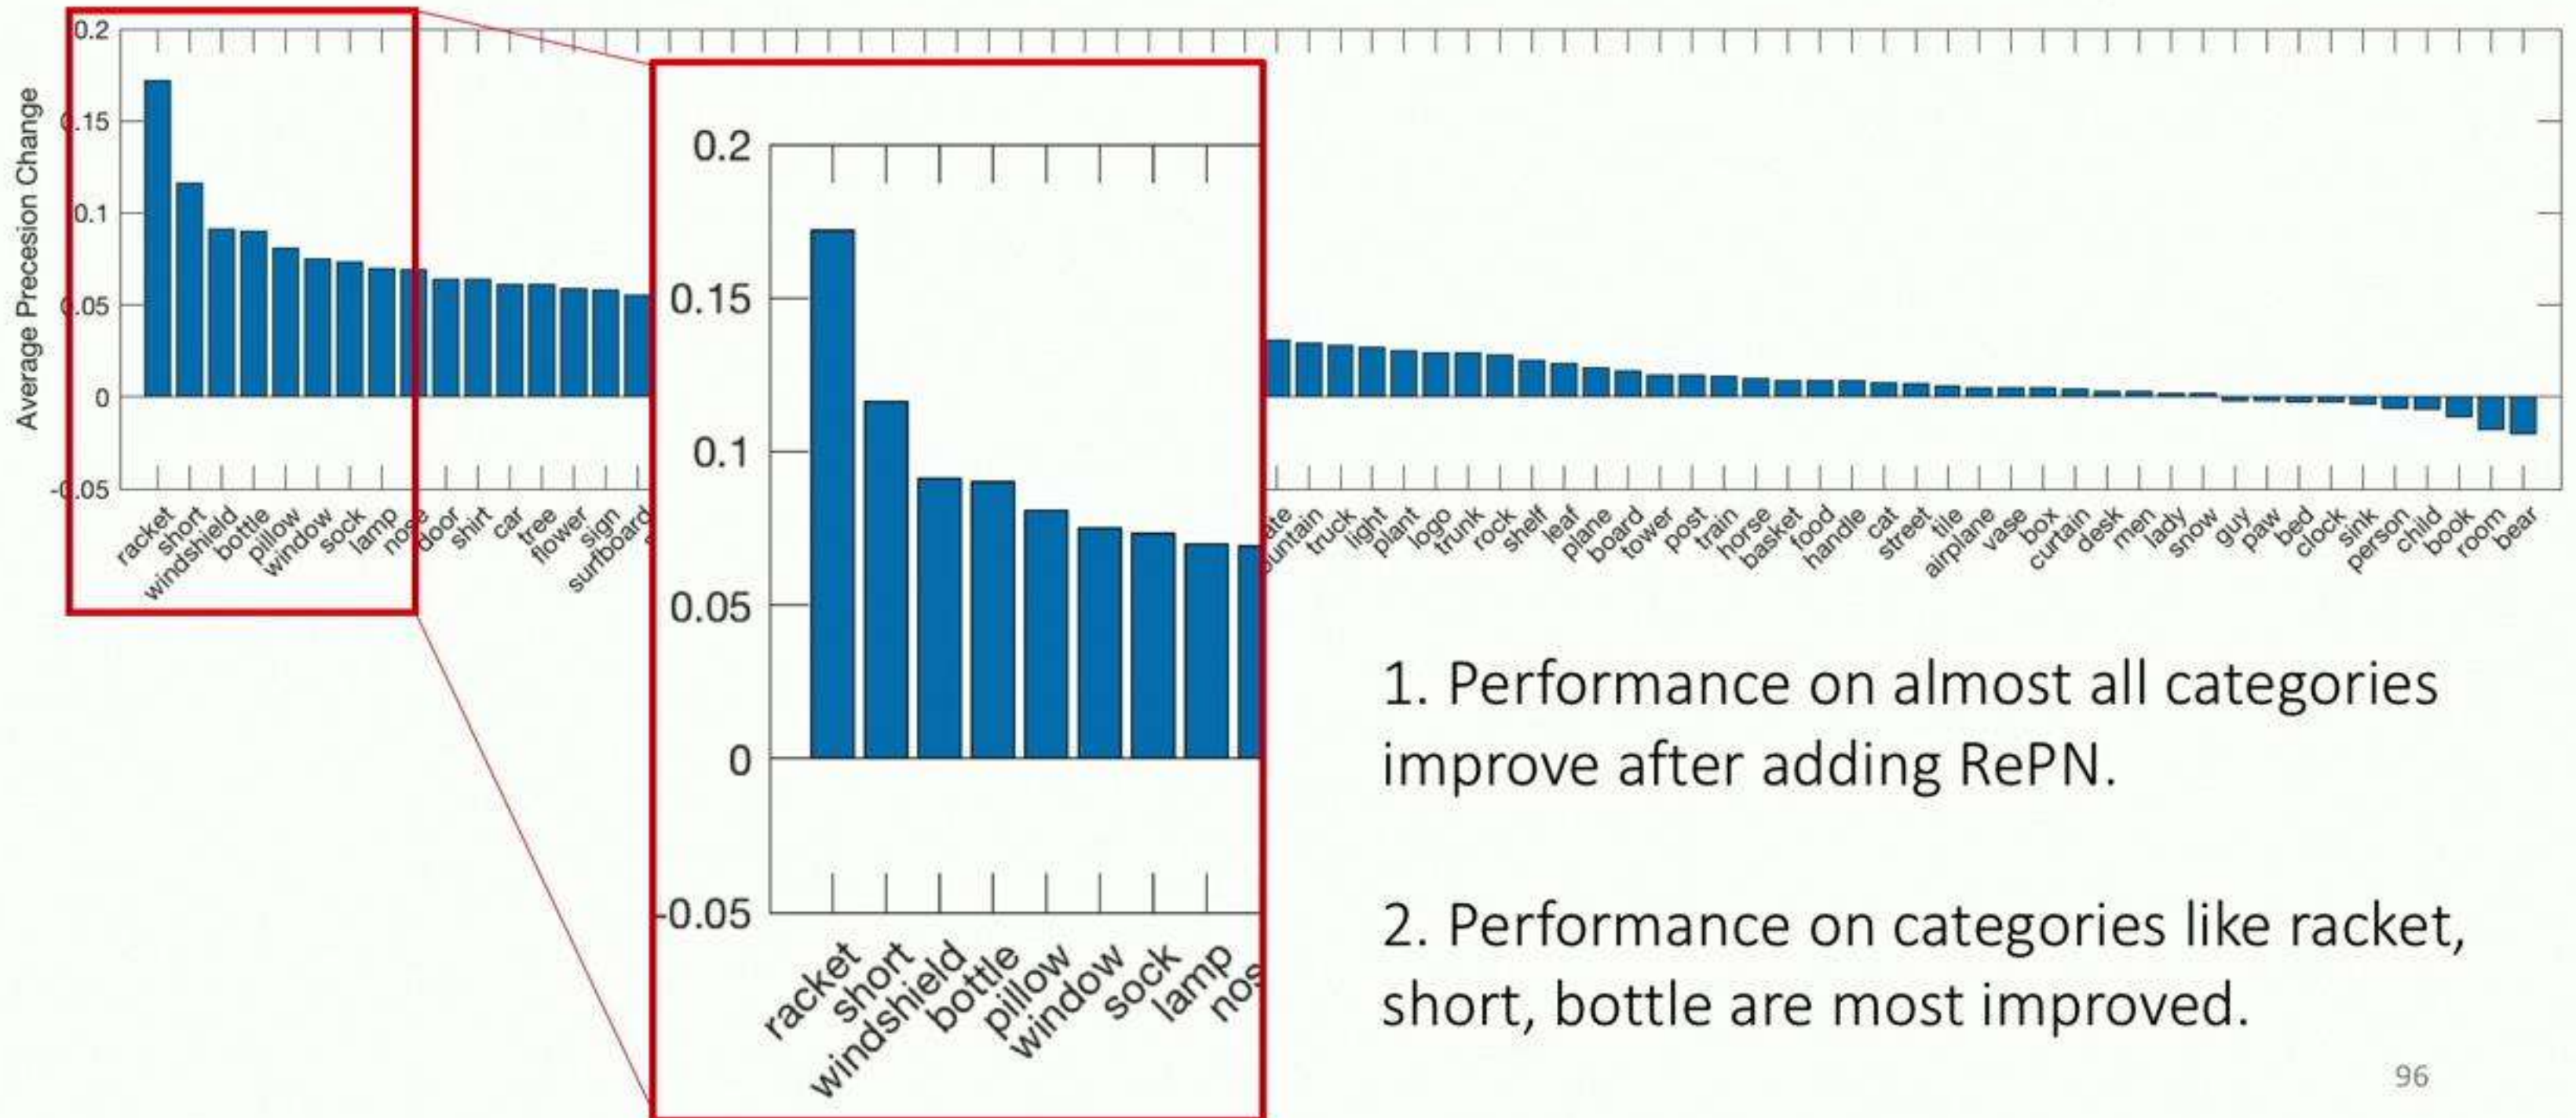
Object Detection Investigation



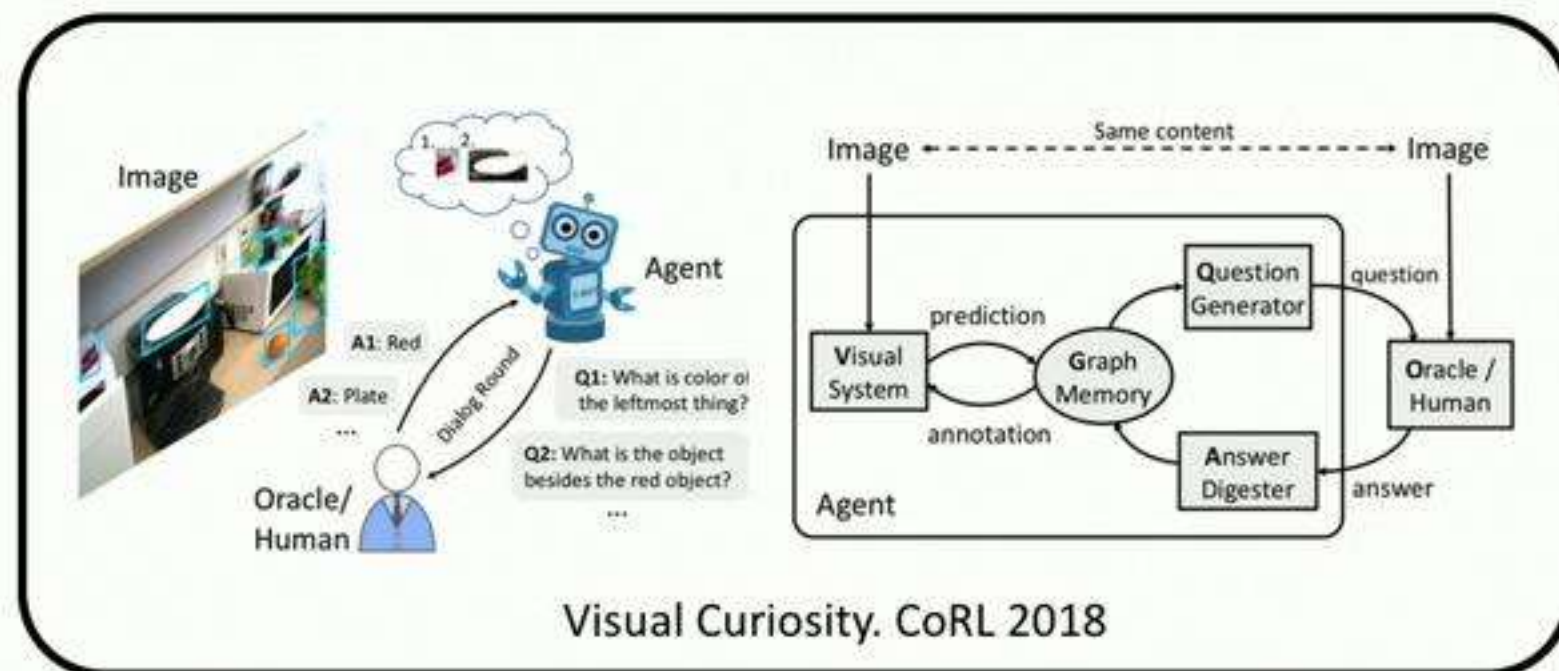
Takeaways

- Introducing a general base model for scene graph generation
- Pruning the fully-connected graph is important for scene graph generation
- Exploiting the context across objects and predicates is crucial
- Scene graph generation helps to improve object detection

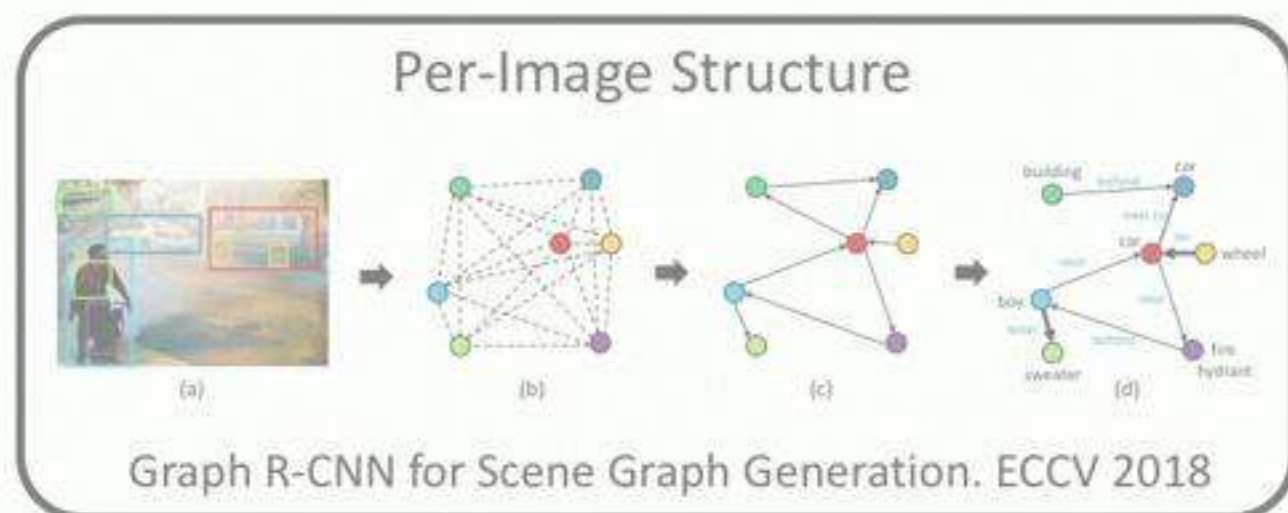
Object Detection Investigation



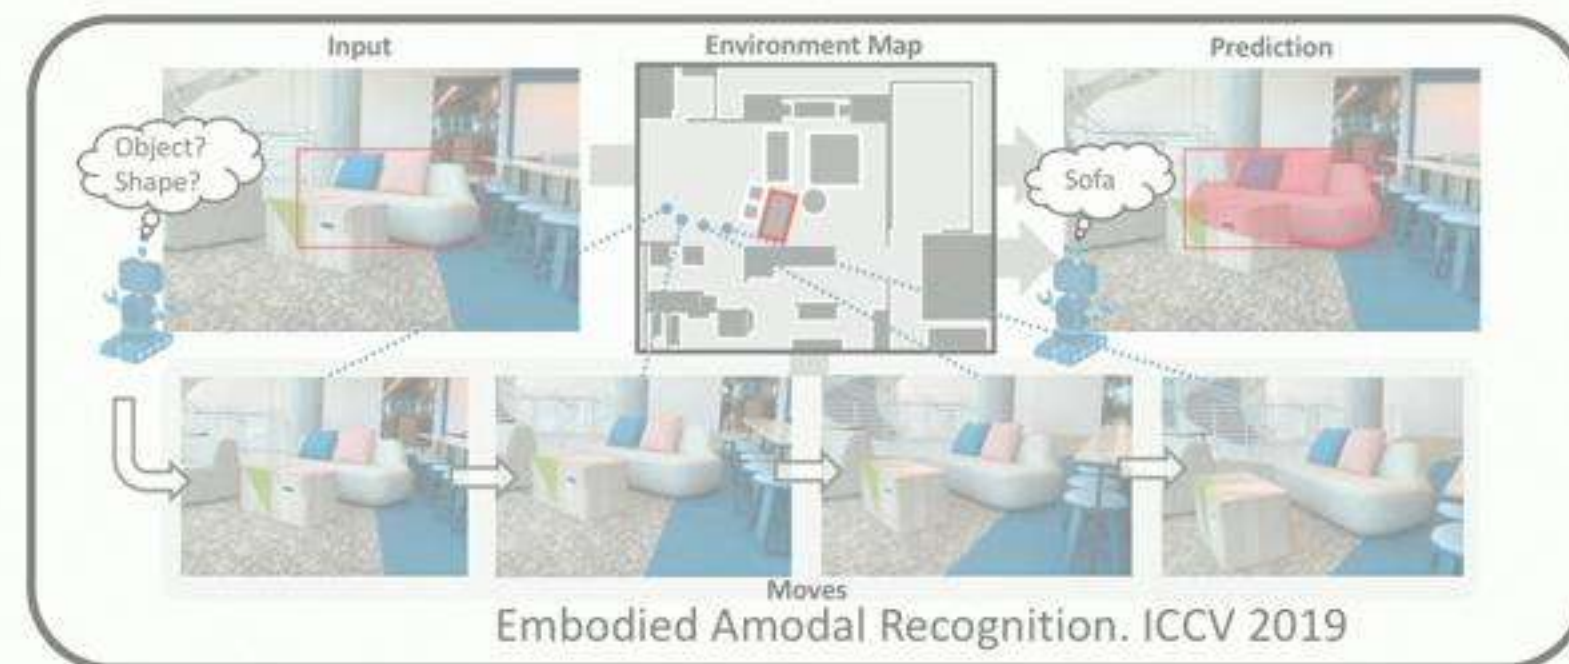
In this talk



Interact with Human



Structured Visual Understanding

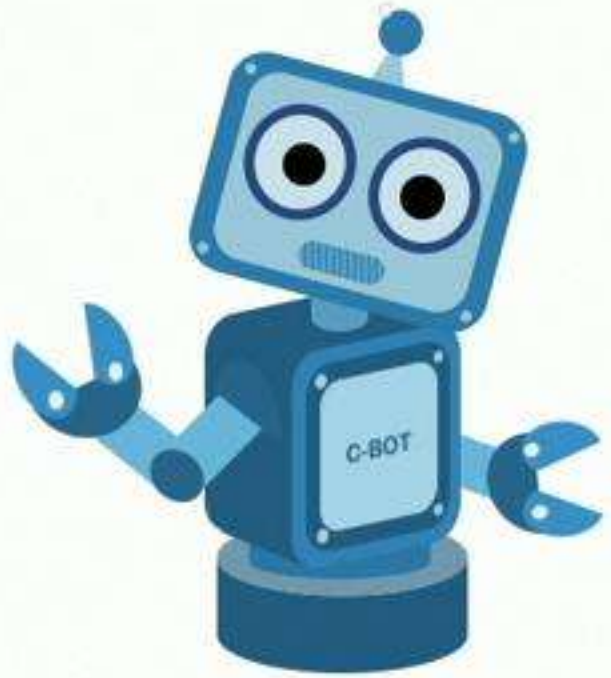


Interact with Environment

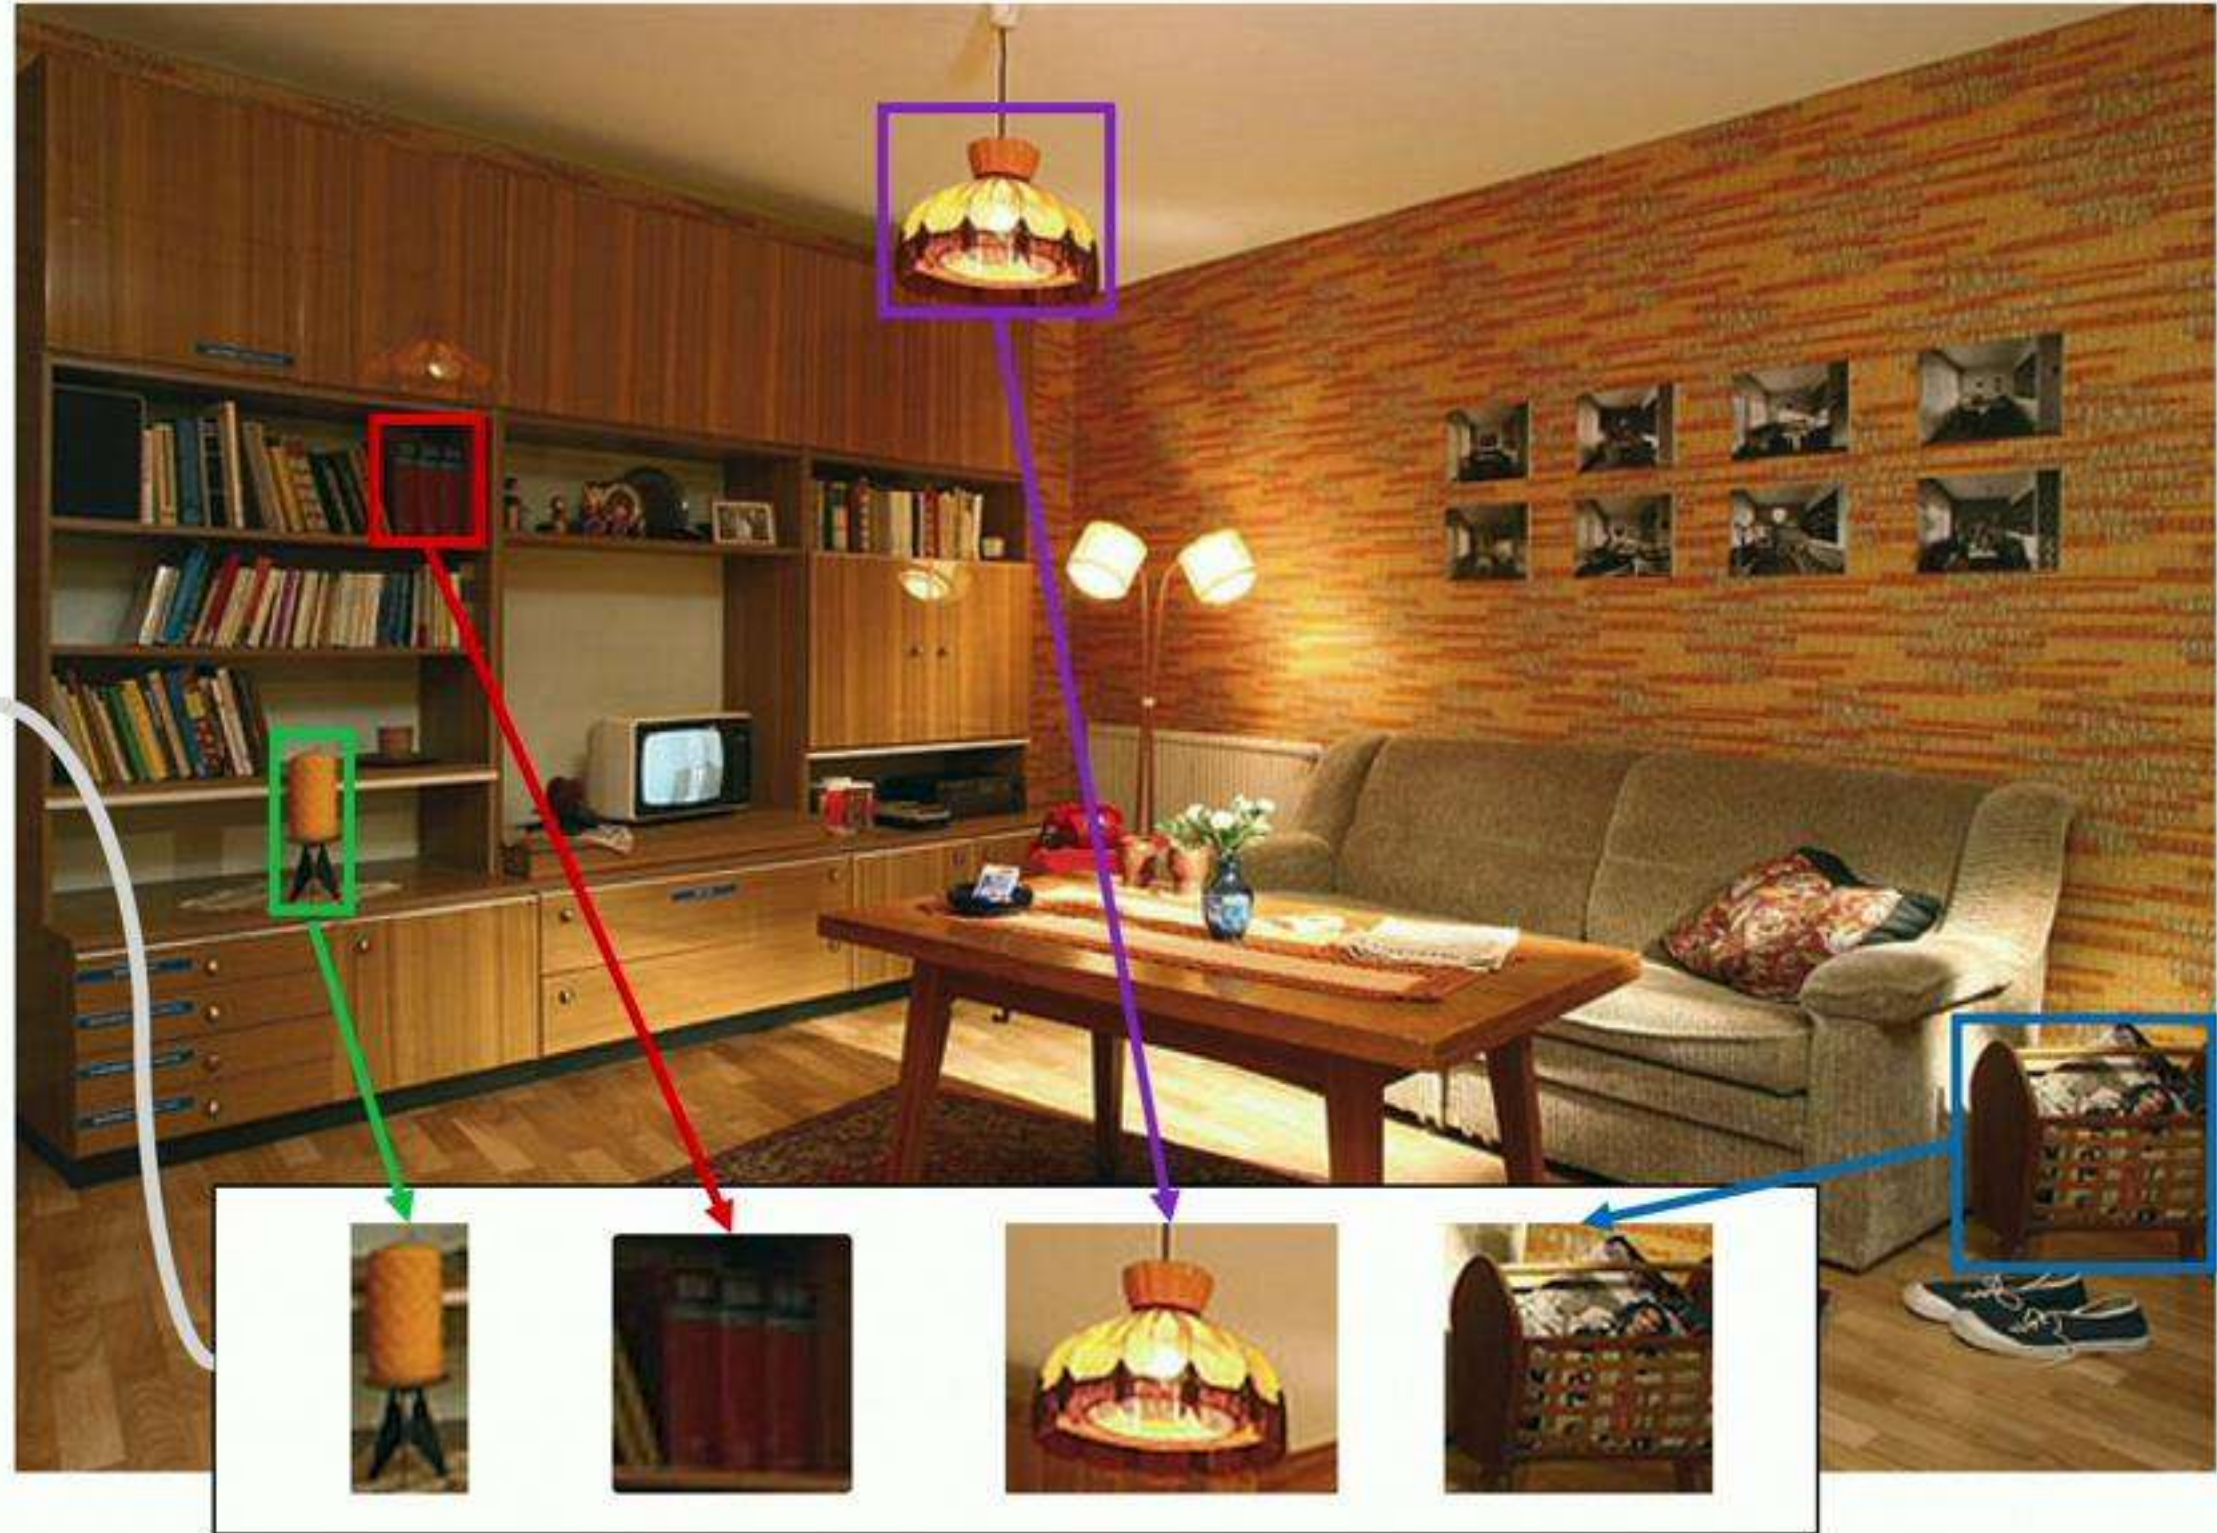
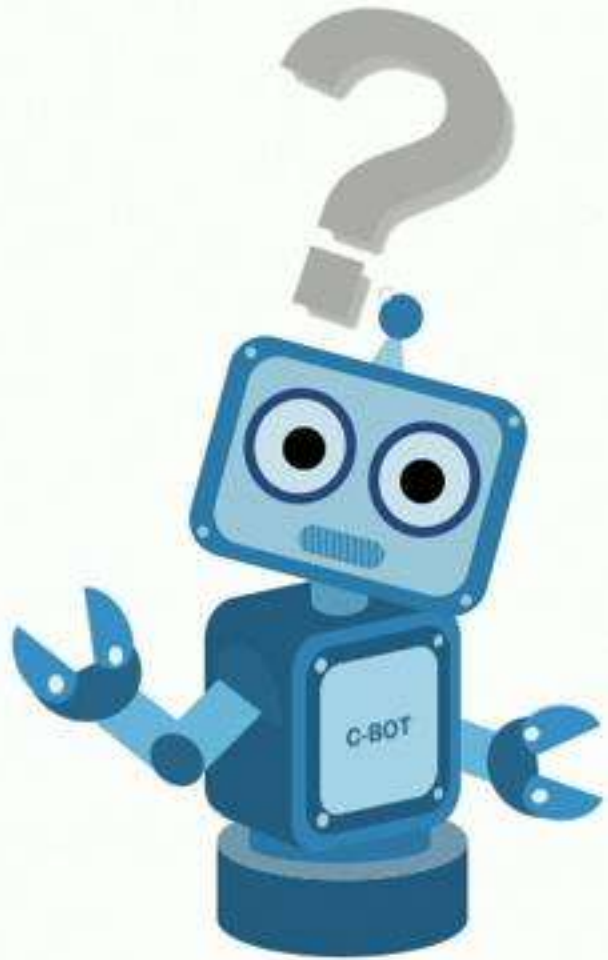
Visual Understanding by Asking Questions

Visual Curiosity: Learning to Ask Questions to Learn
Visual Recognition. CoRL 2018.

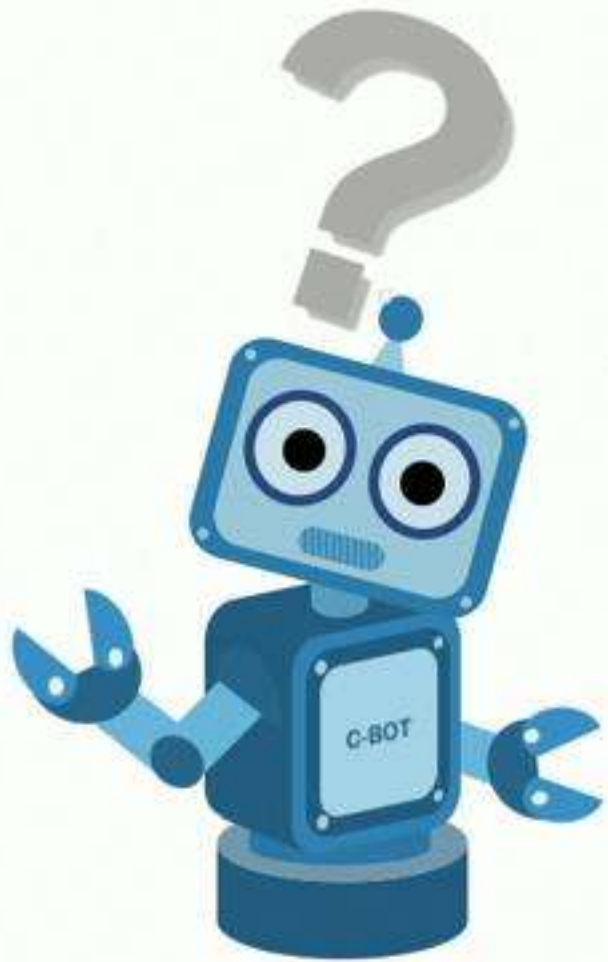
The Open-World Recognition Problem



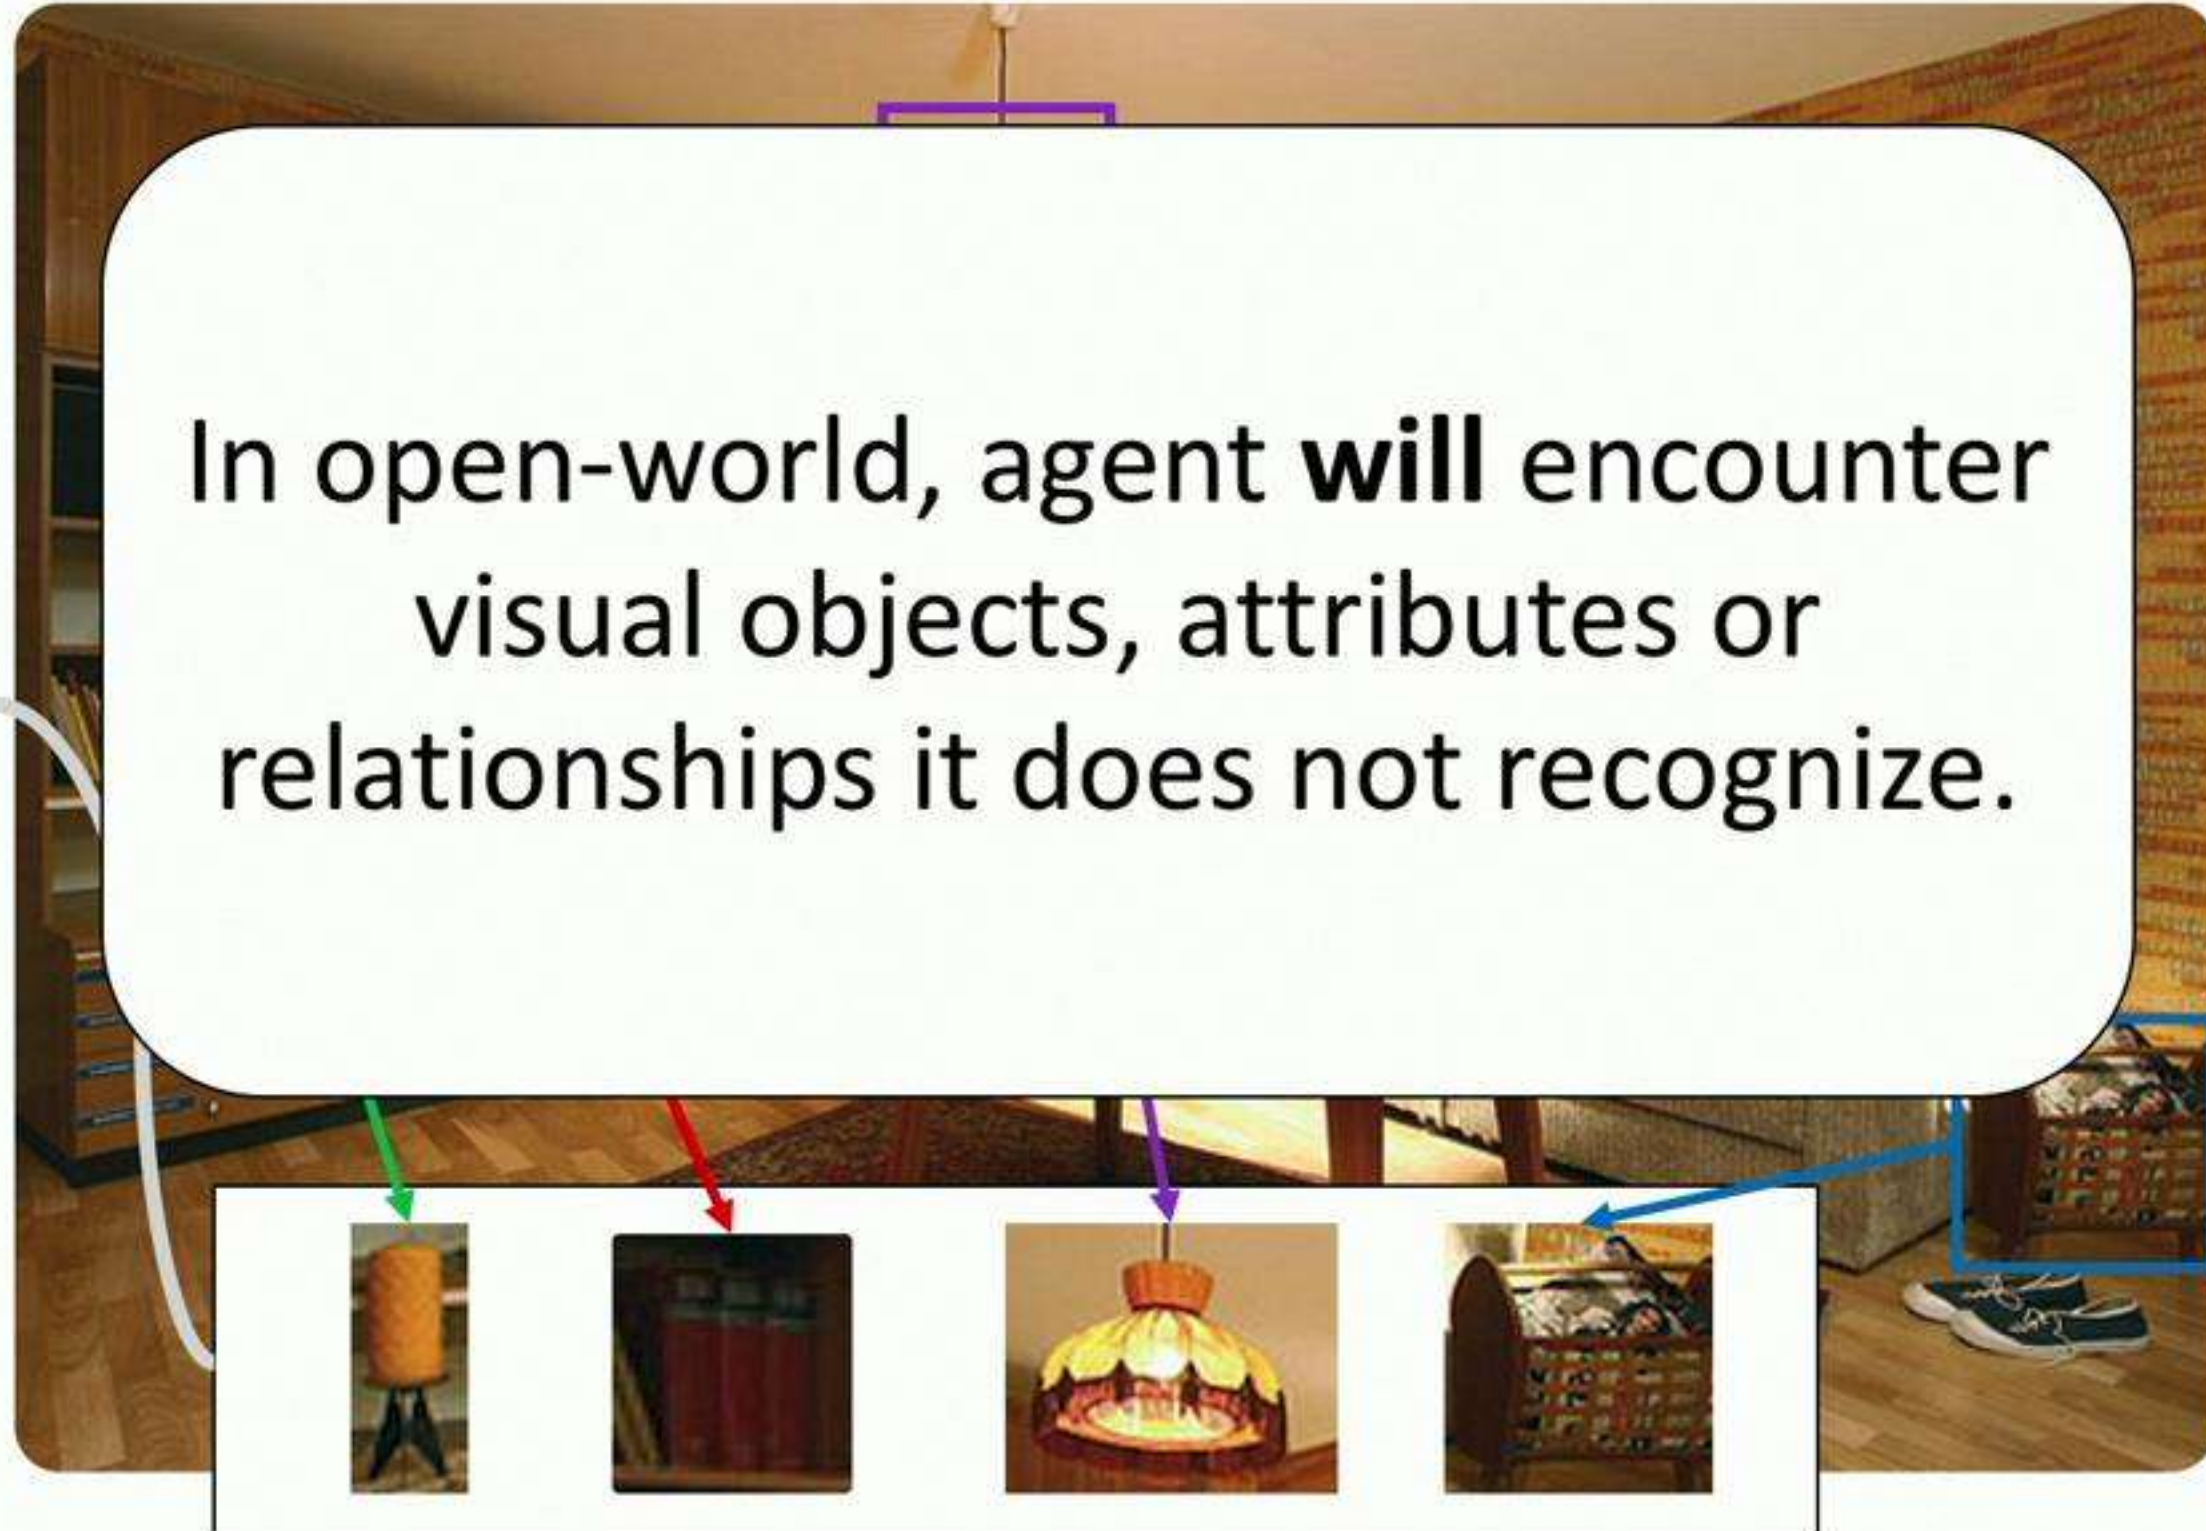
The Open-World Recognition Problem



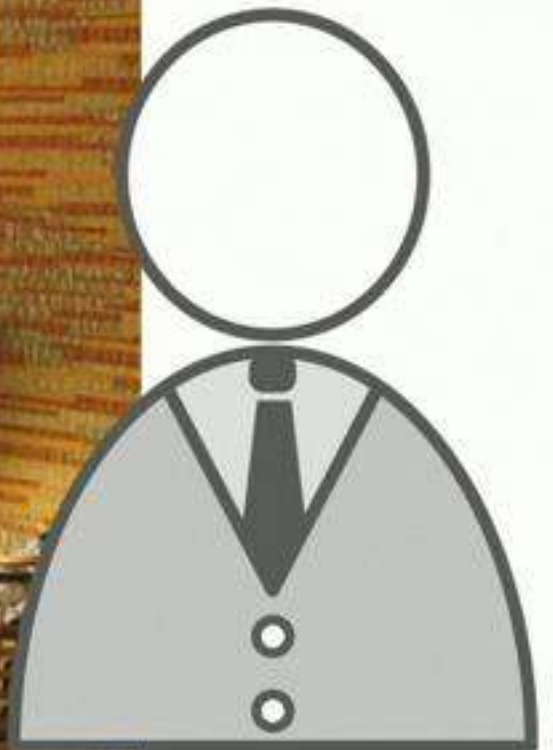
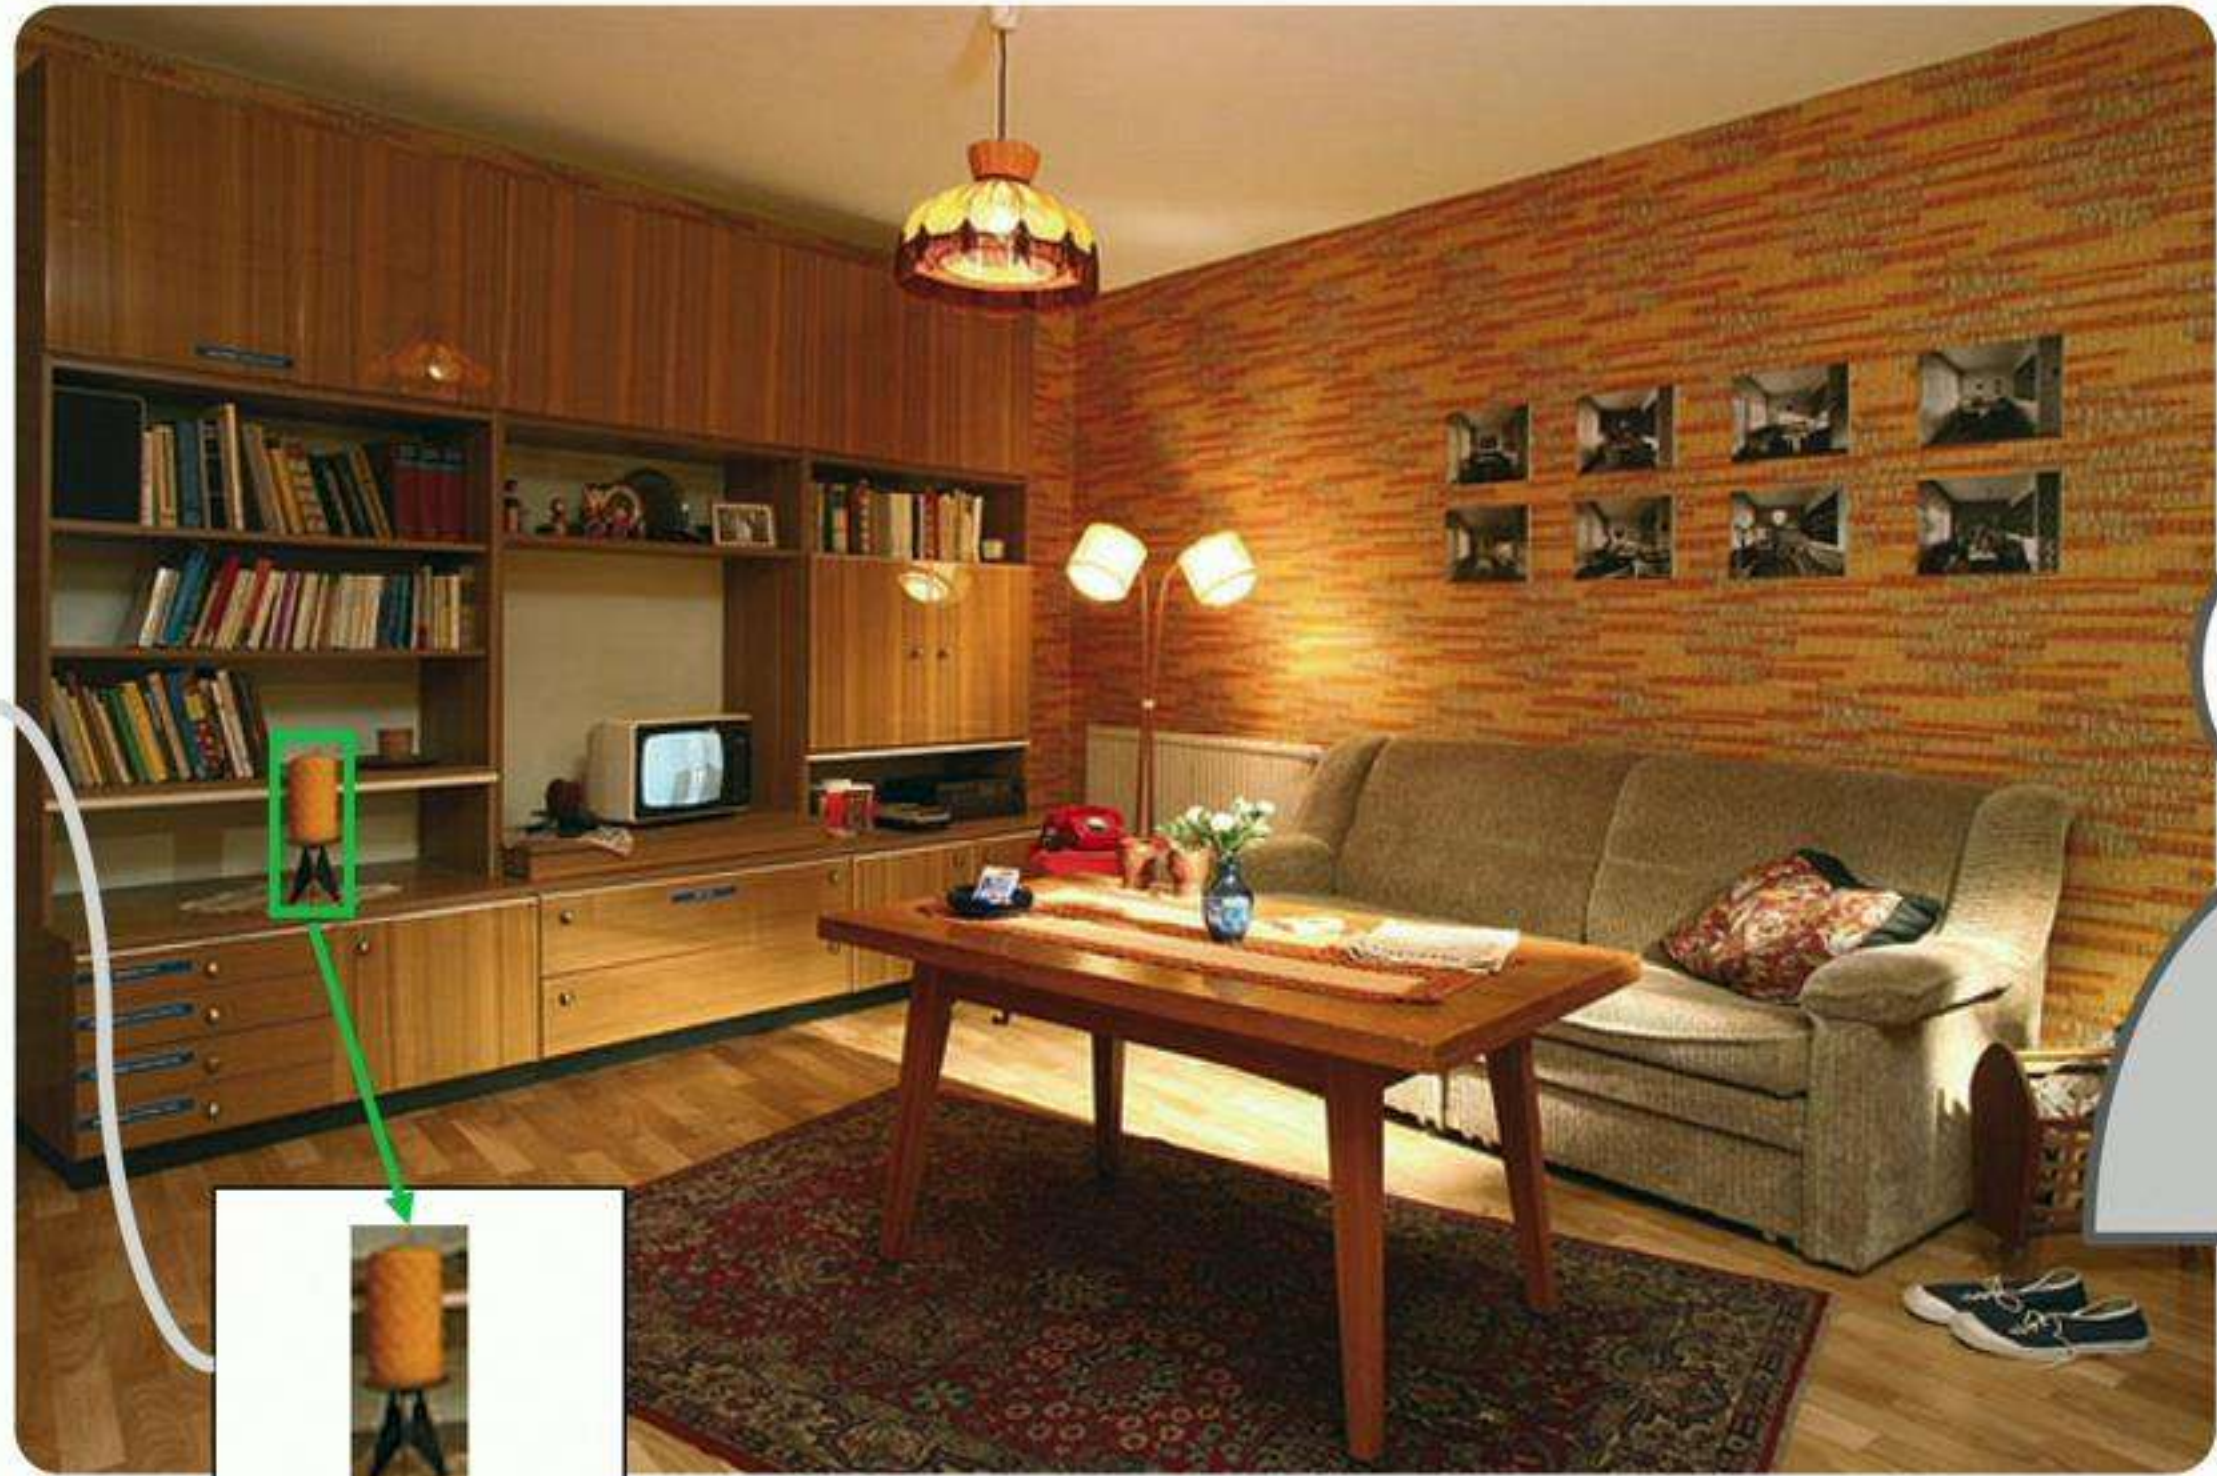
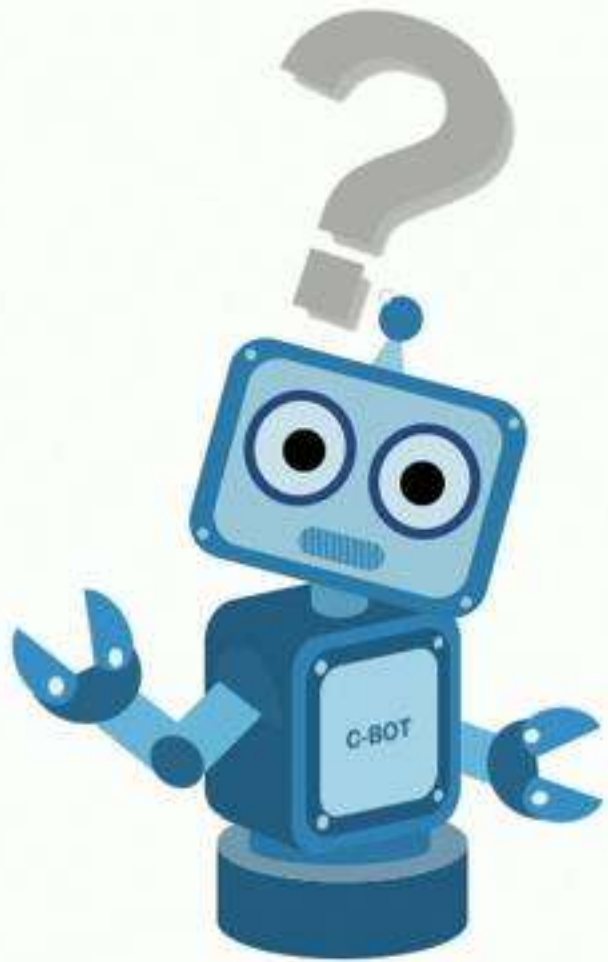
The Open-World Recognition Problem



In open-world, agent **will** encounter visual objects, attributes or relationships it does not recognize.

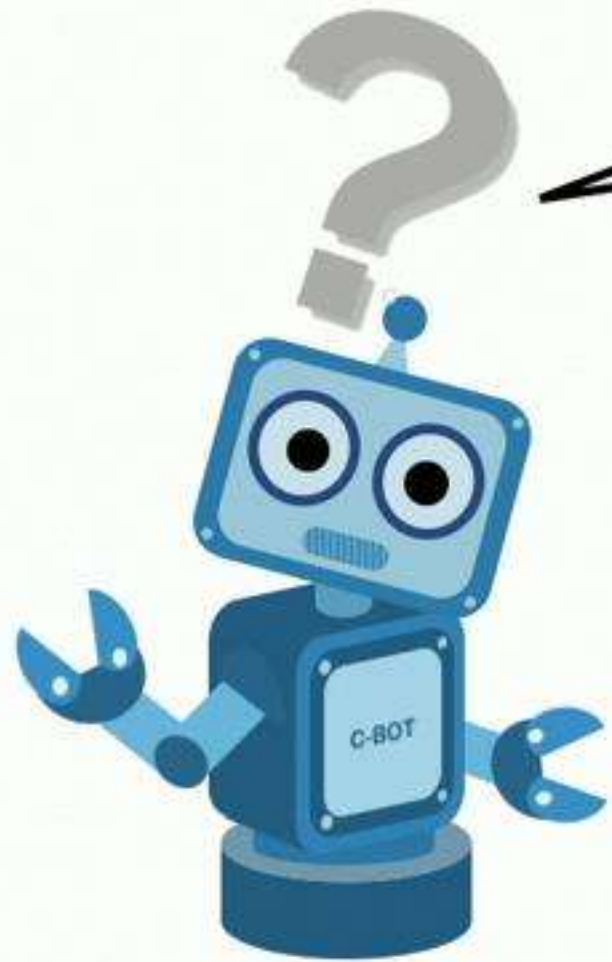


How can an agent learn about these concepts?

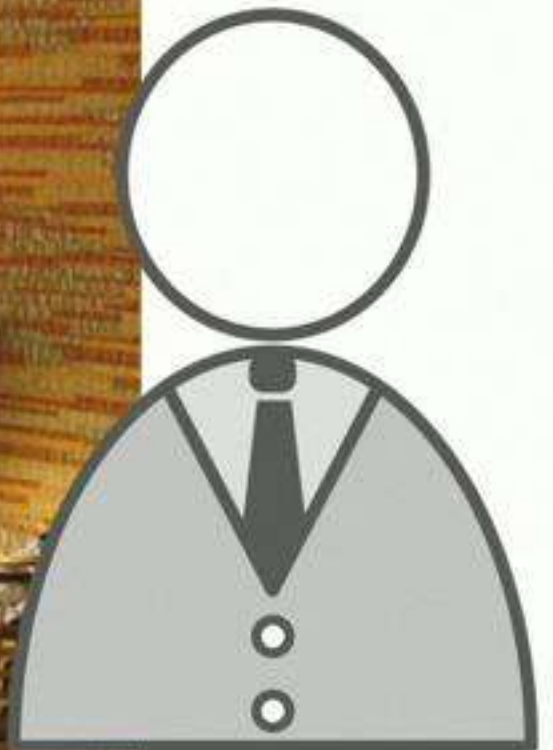
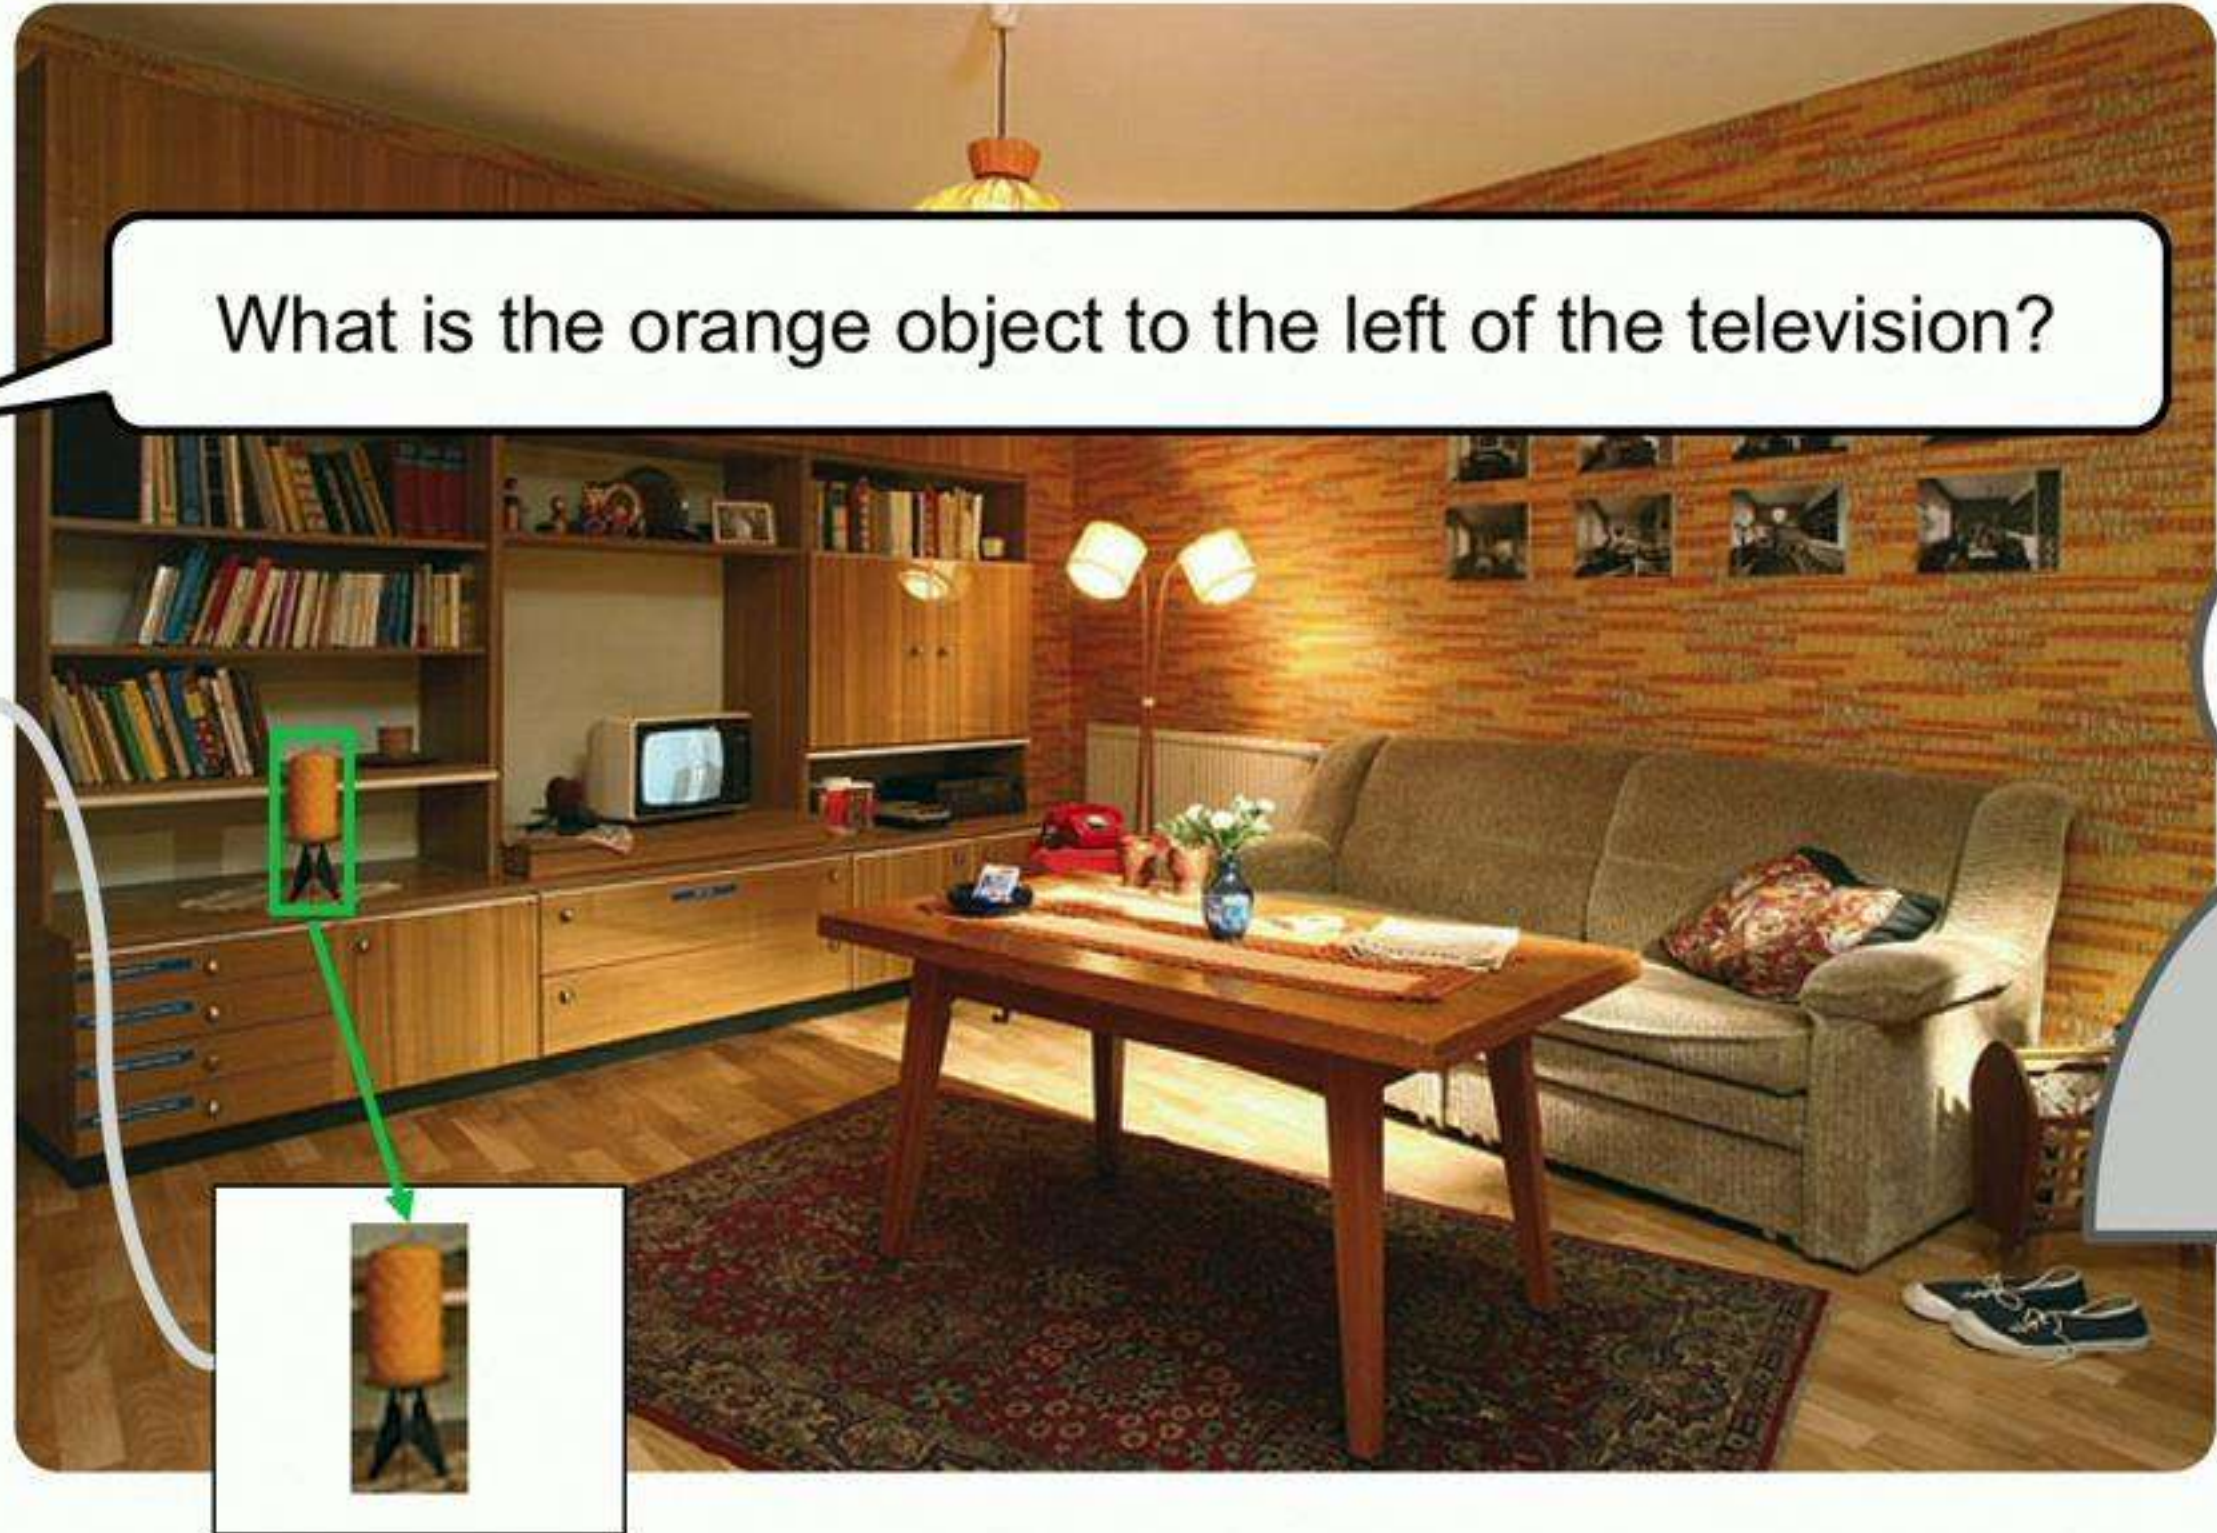


Human
(Oracle)

How can an agent learn about these concepts?

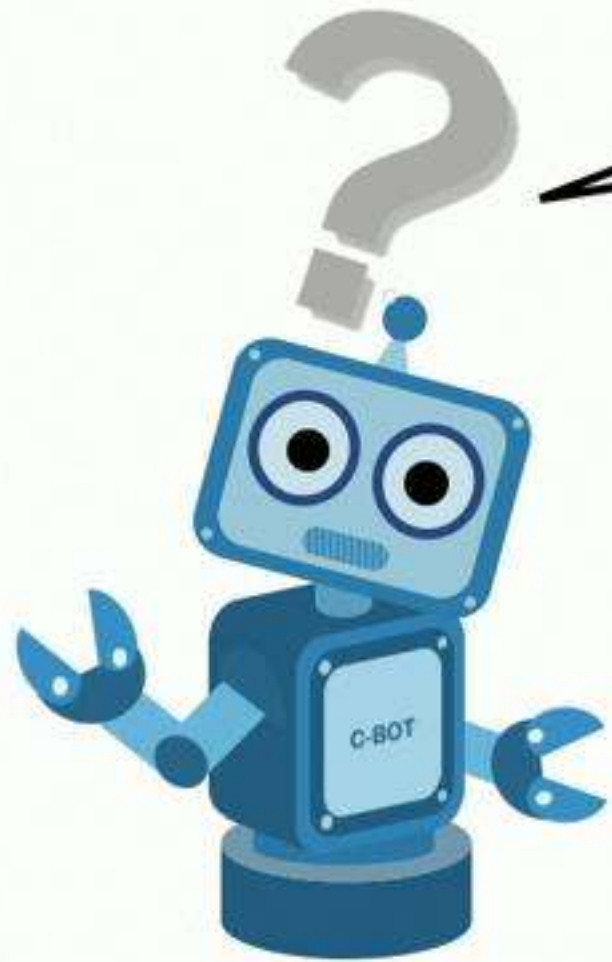


What is the orange object to the left of the television?



Human
(Oracle)

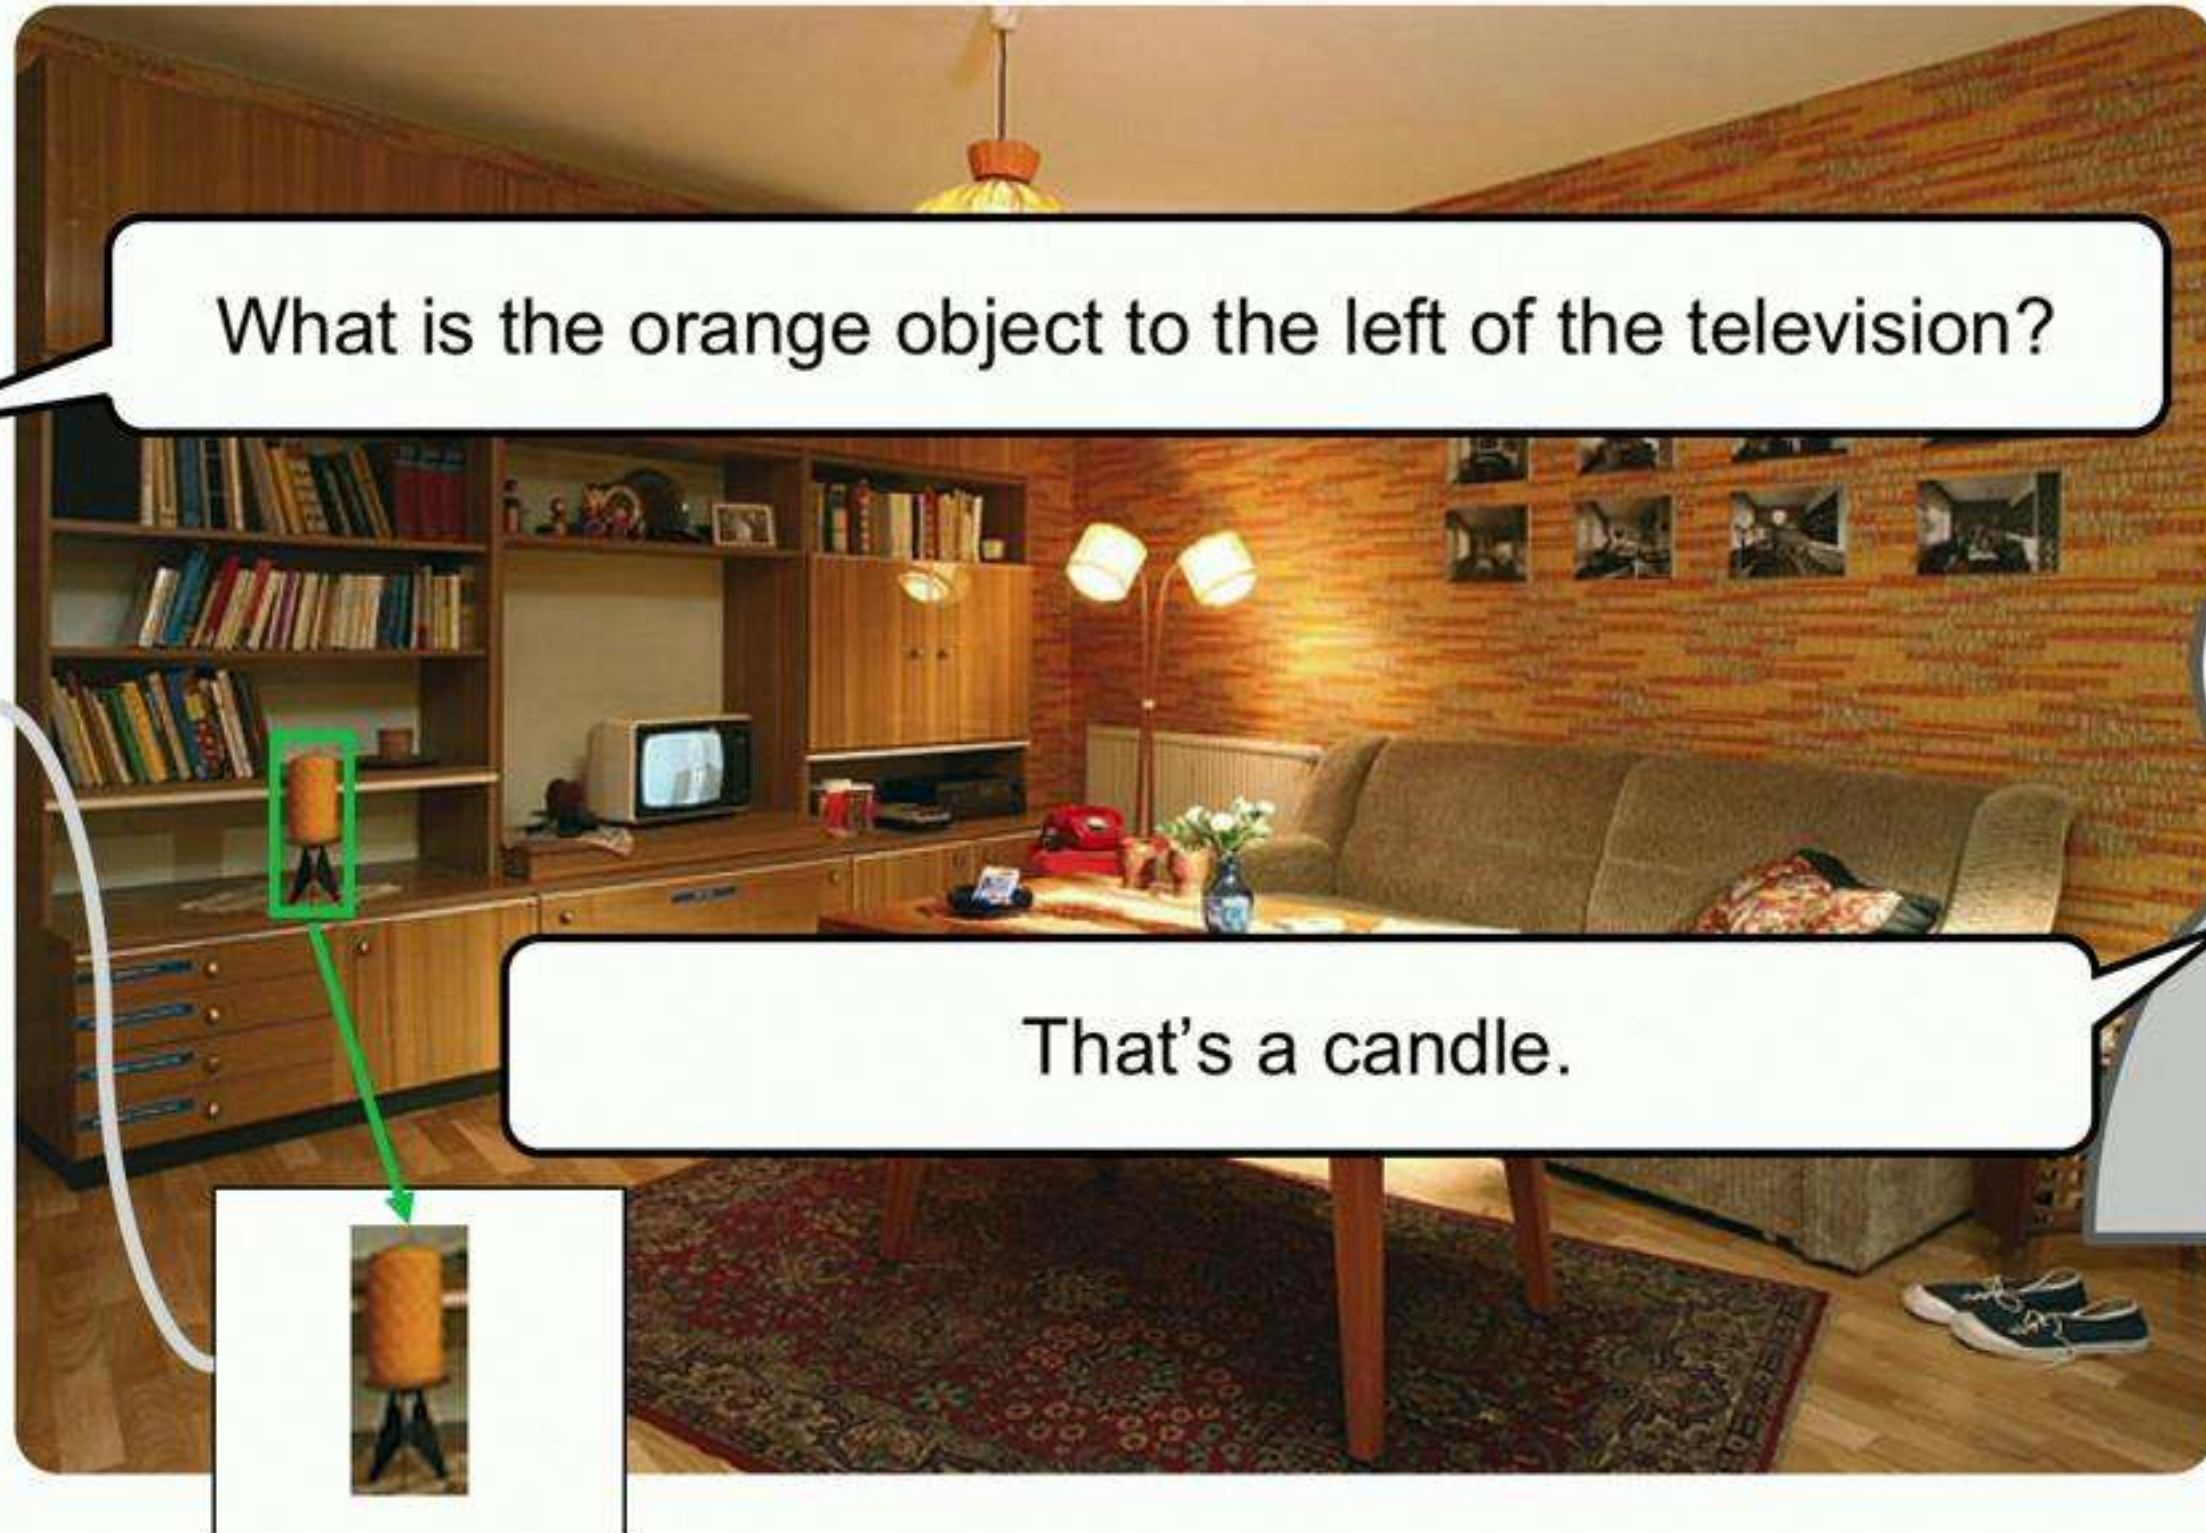
How can an agent learn about these concepts?



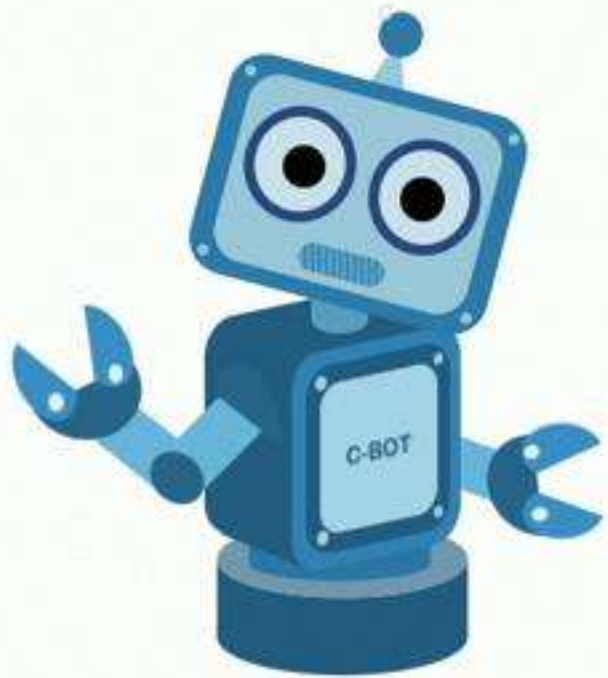
What is the orange object to the left of the television?

That's a candle.

Human
(Oracle)



Asking informative questions is challenging!



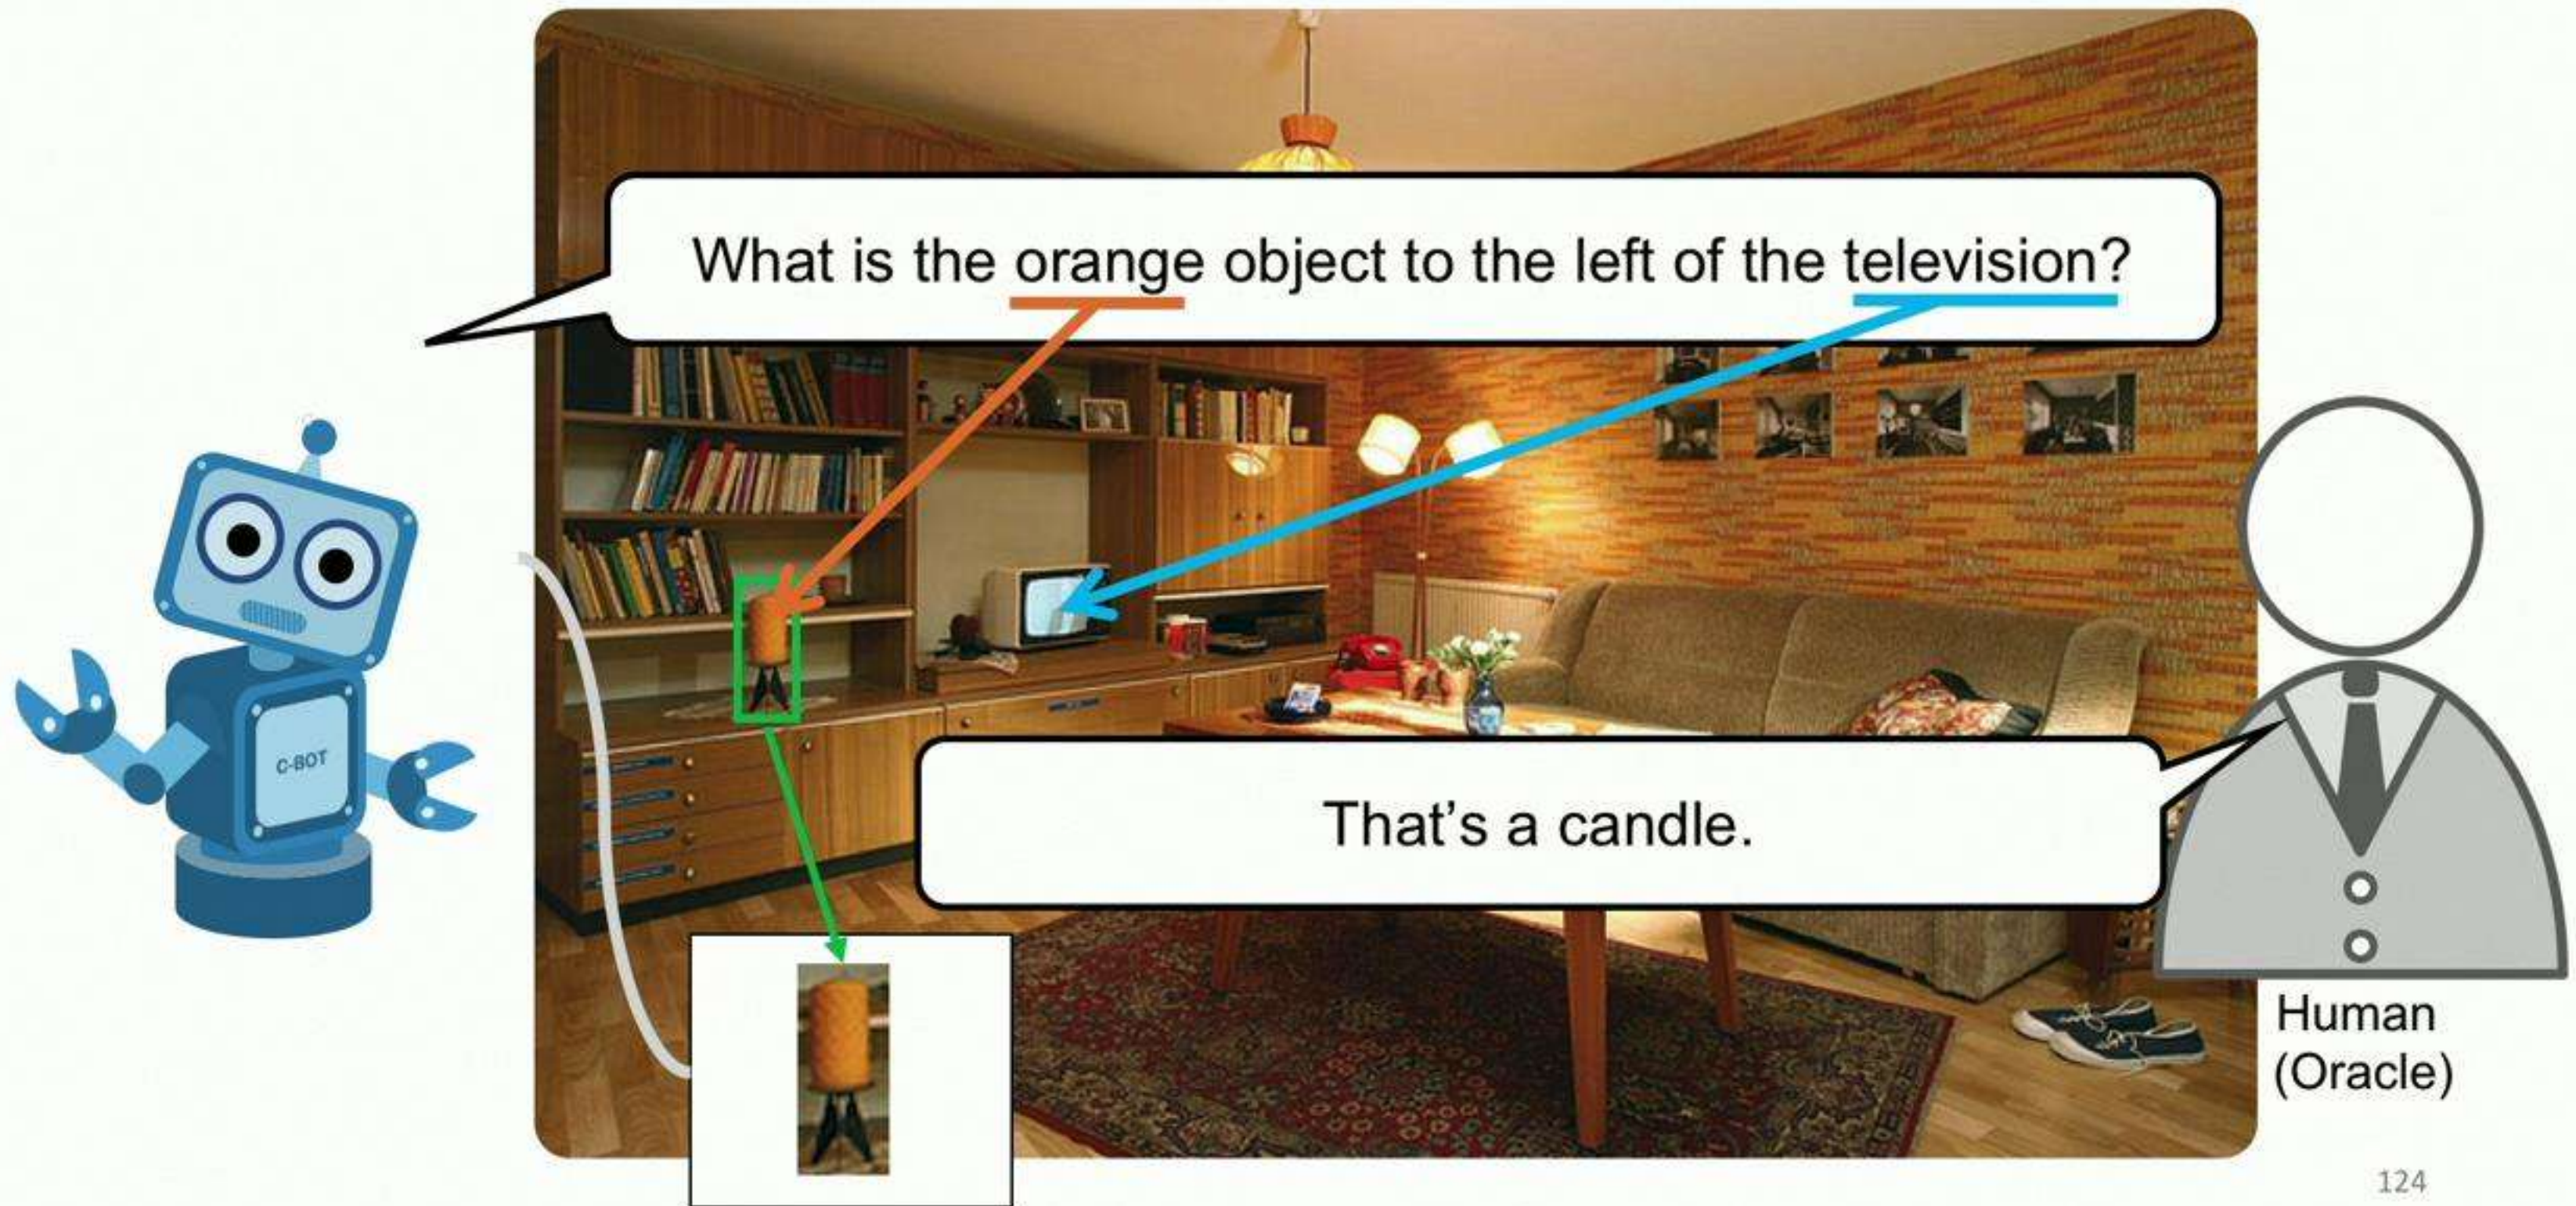
What is the orange object to the left of the television?

That's a candle.

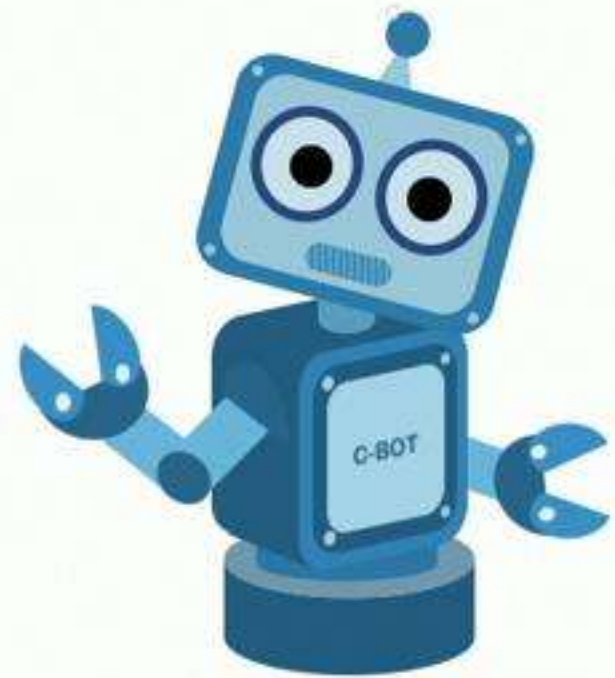
Human
(Oracle)



Asking informative questions is challenging!



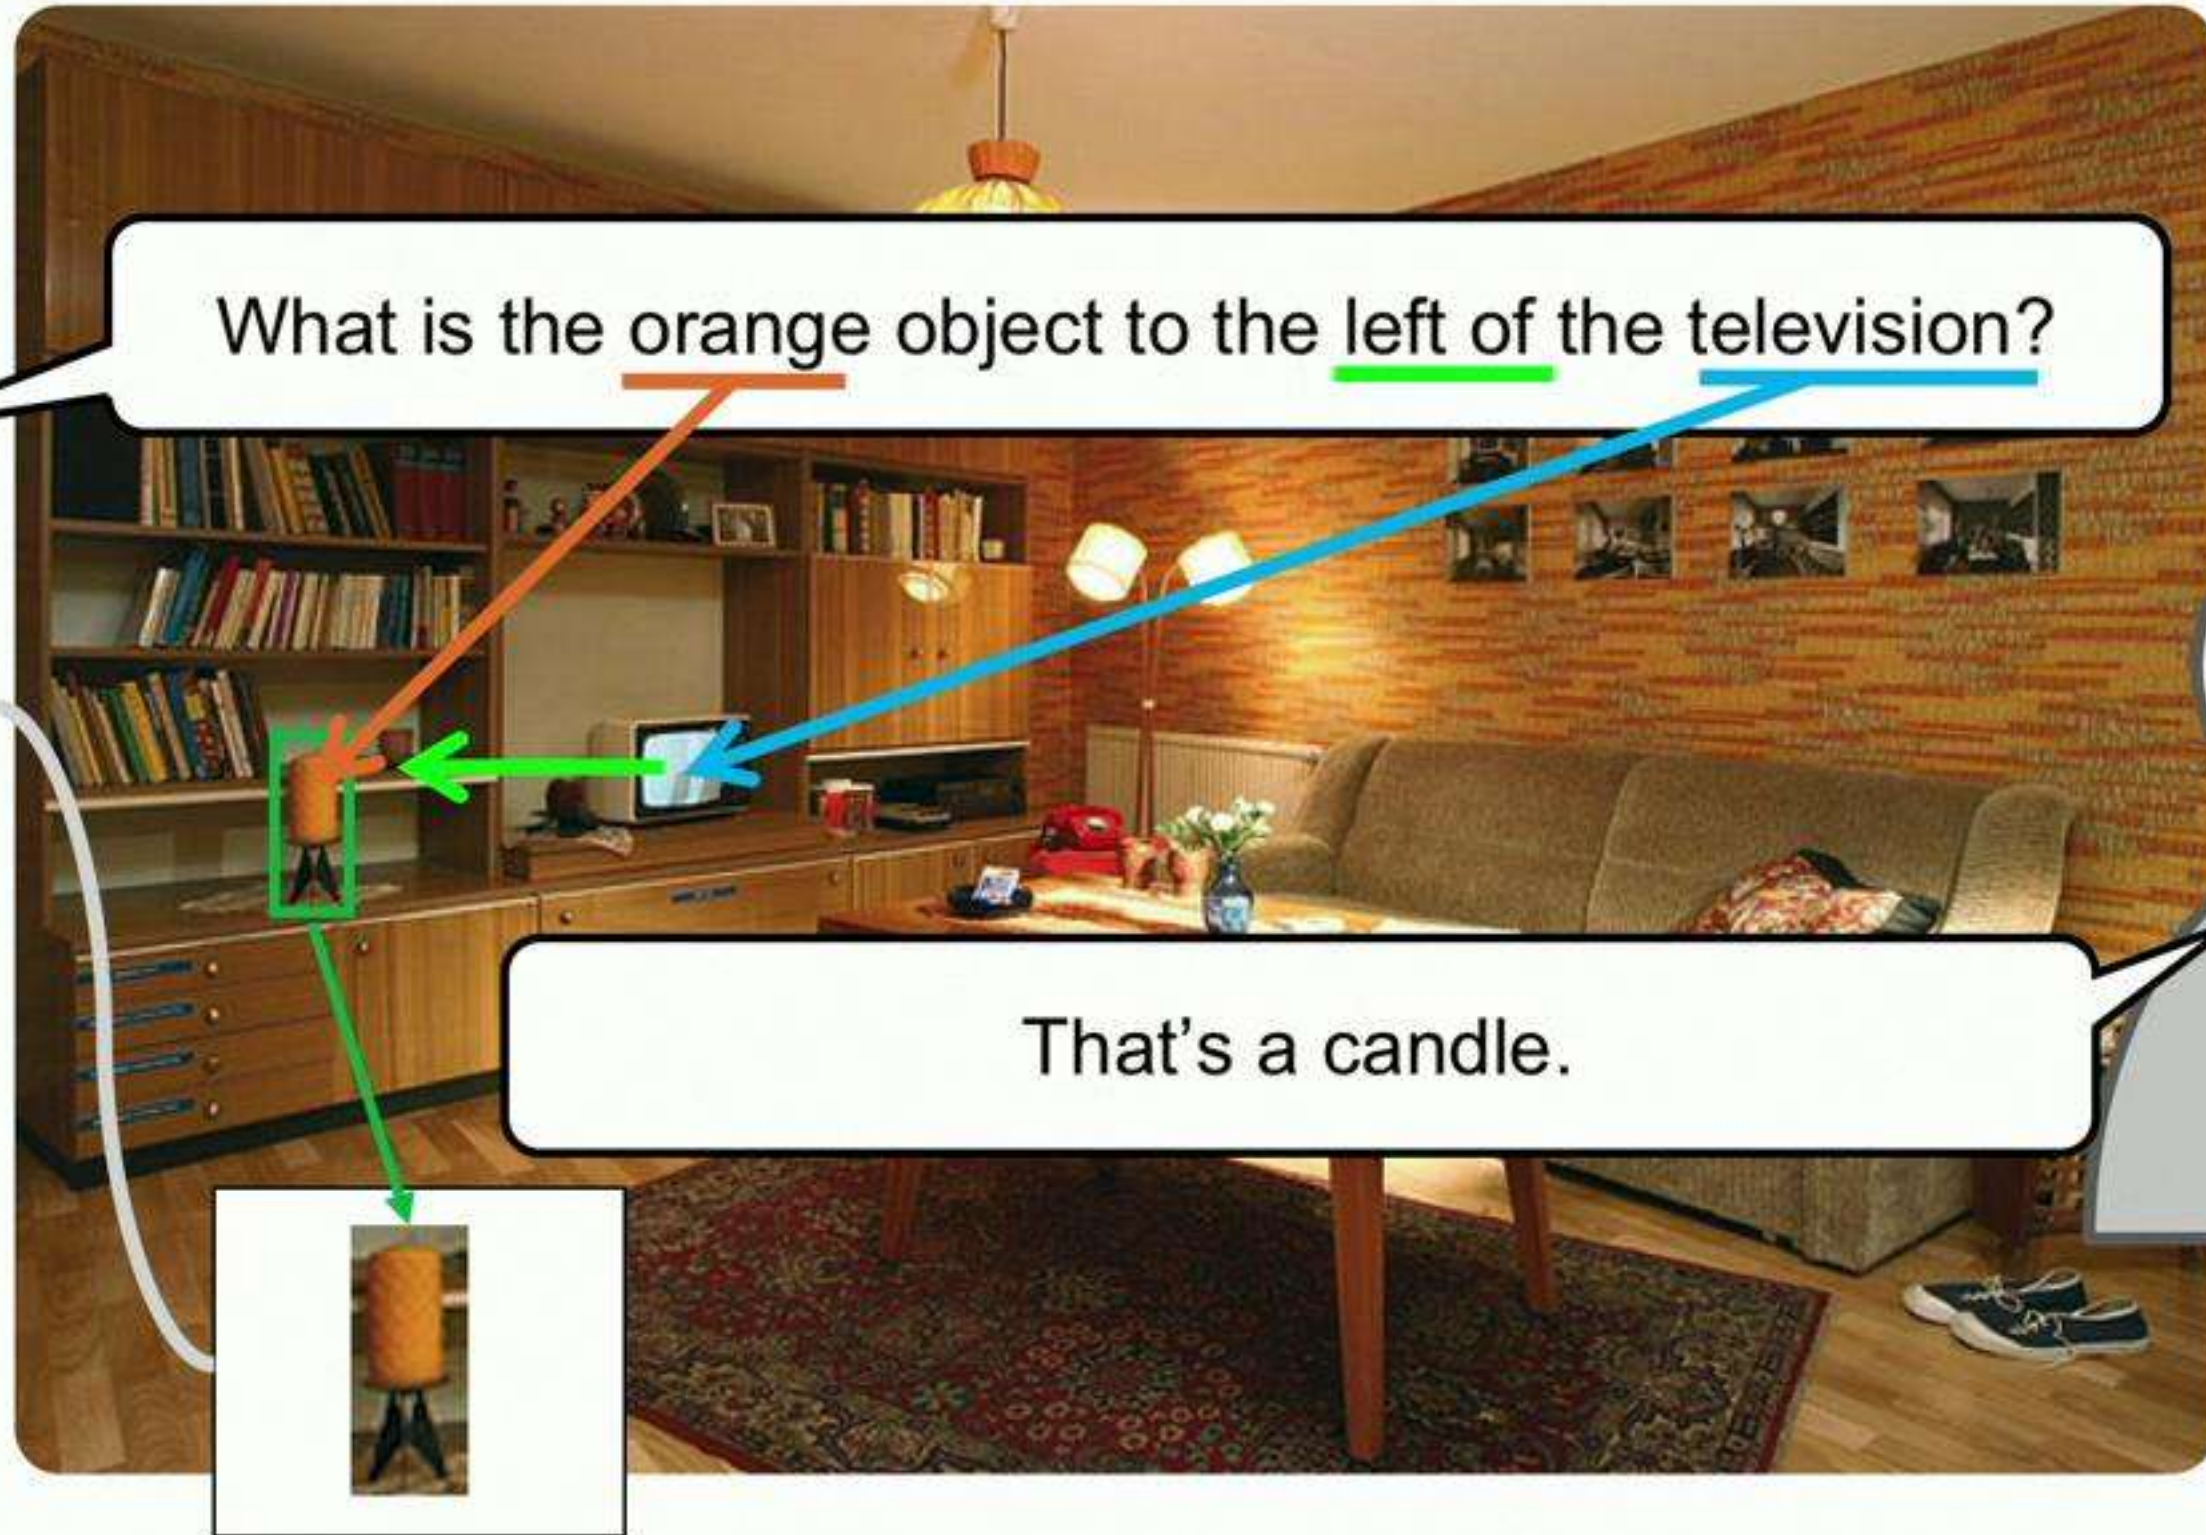
Asking informative questions is challenging!



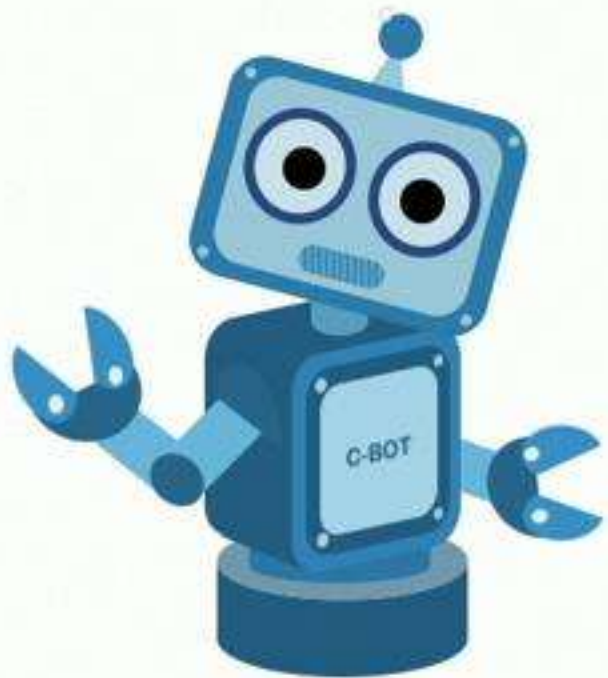
What is the orange object to the left of the television?

That's a candle.

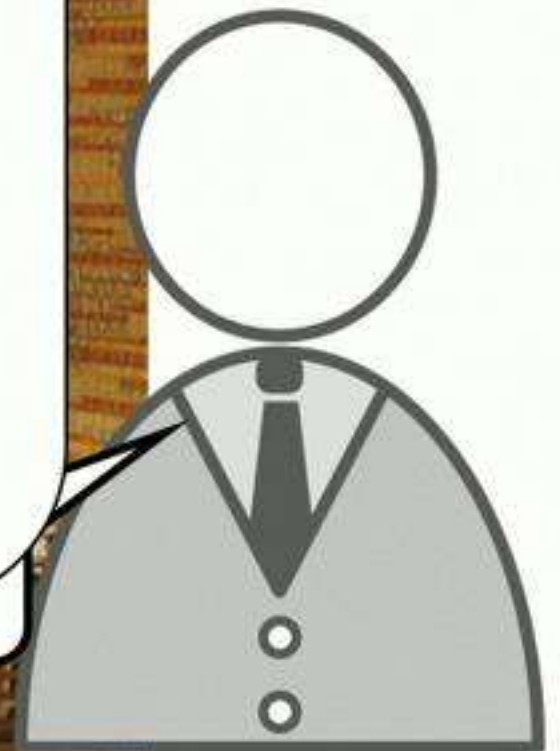
Human
(Oracle)



Asking informative questions is challenging!

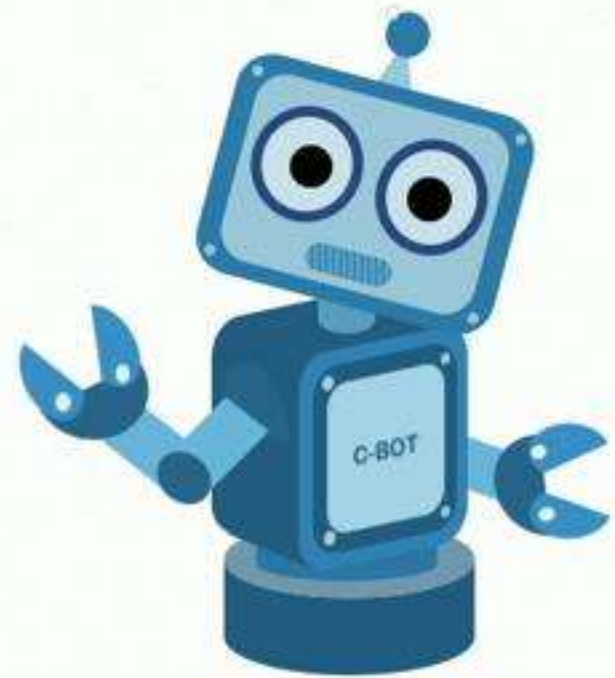


How to train an agent to **ask questions about its surroundings** to improve its visual understanding capabilities?



Human
(Oracle)

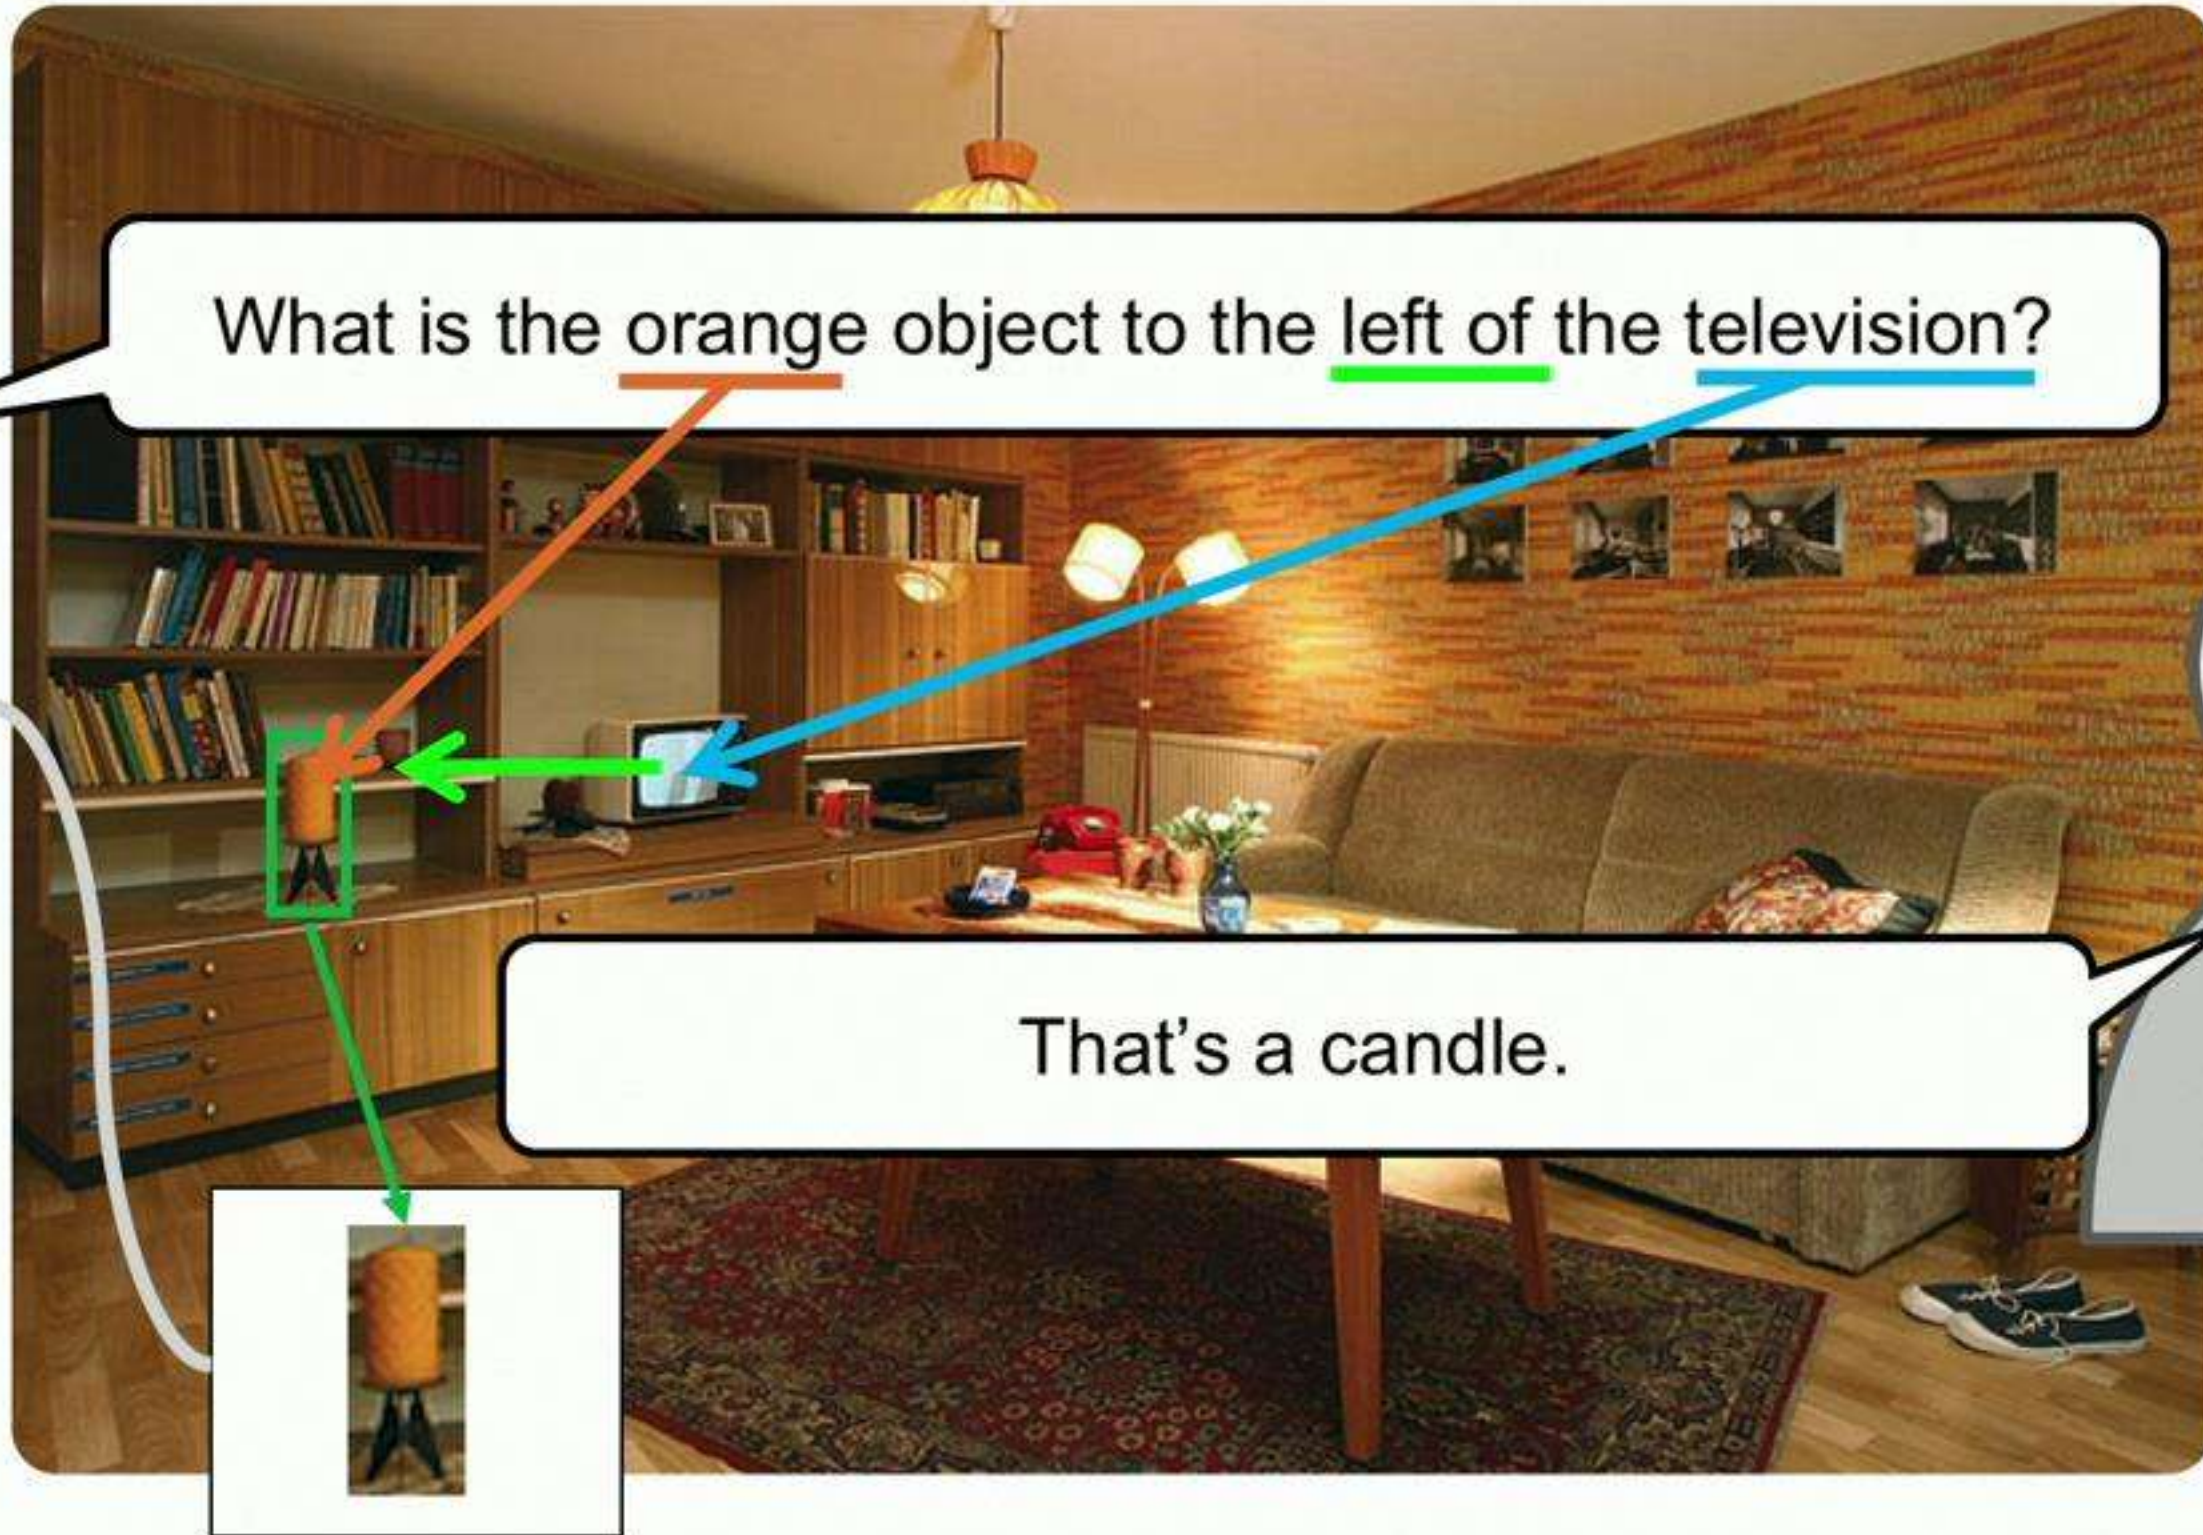
Asking informative questions is challenging!



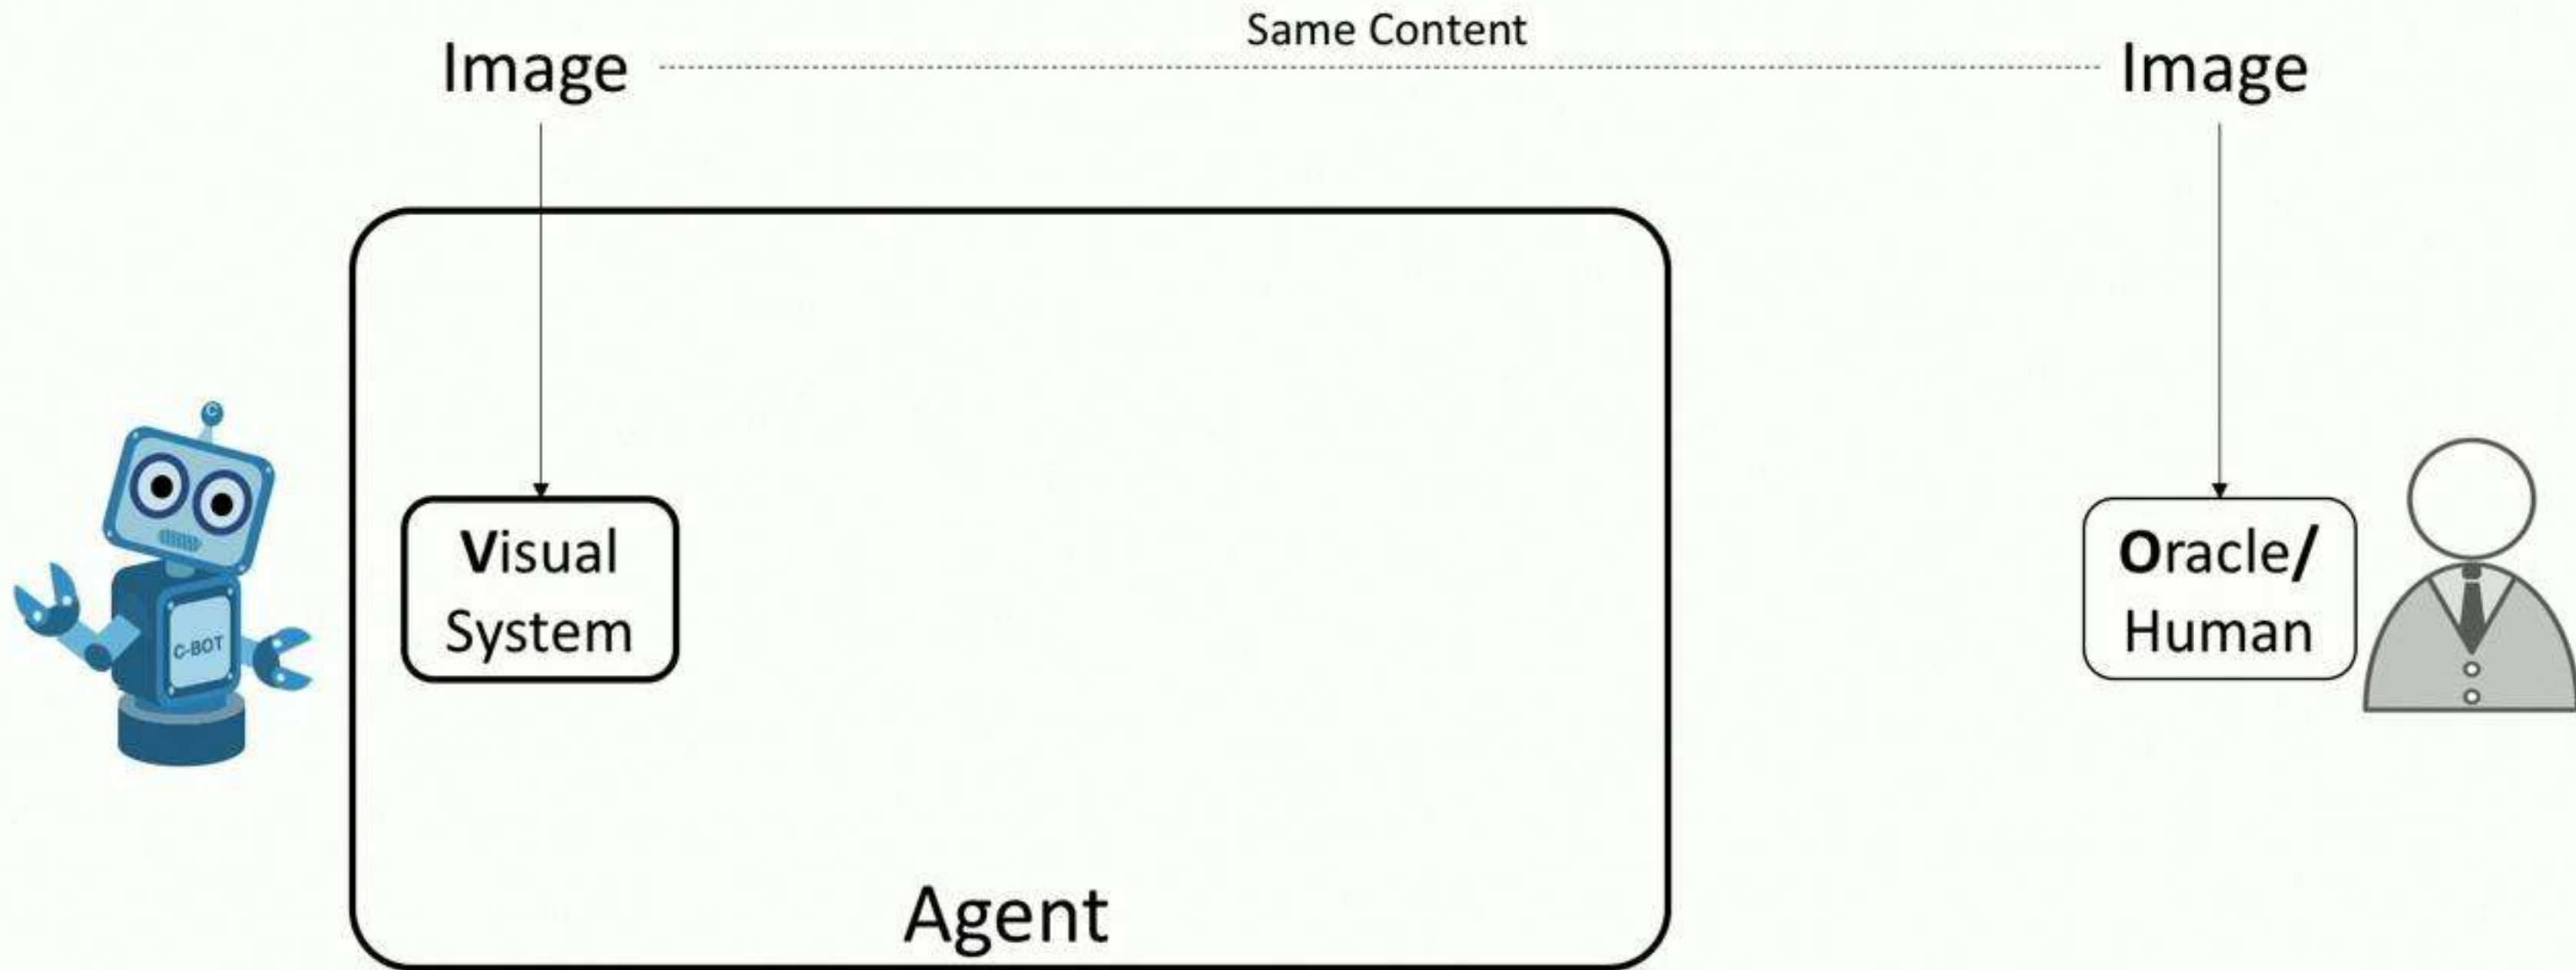
What is the orange object to the left of the television?

That's a candle.

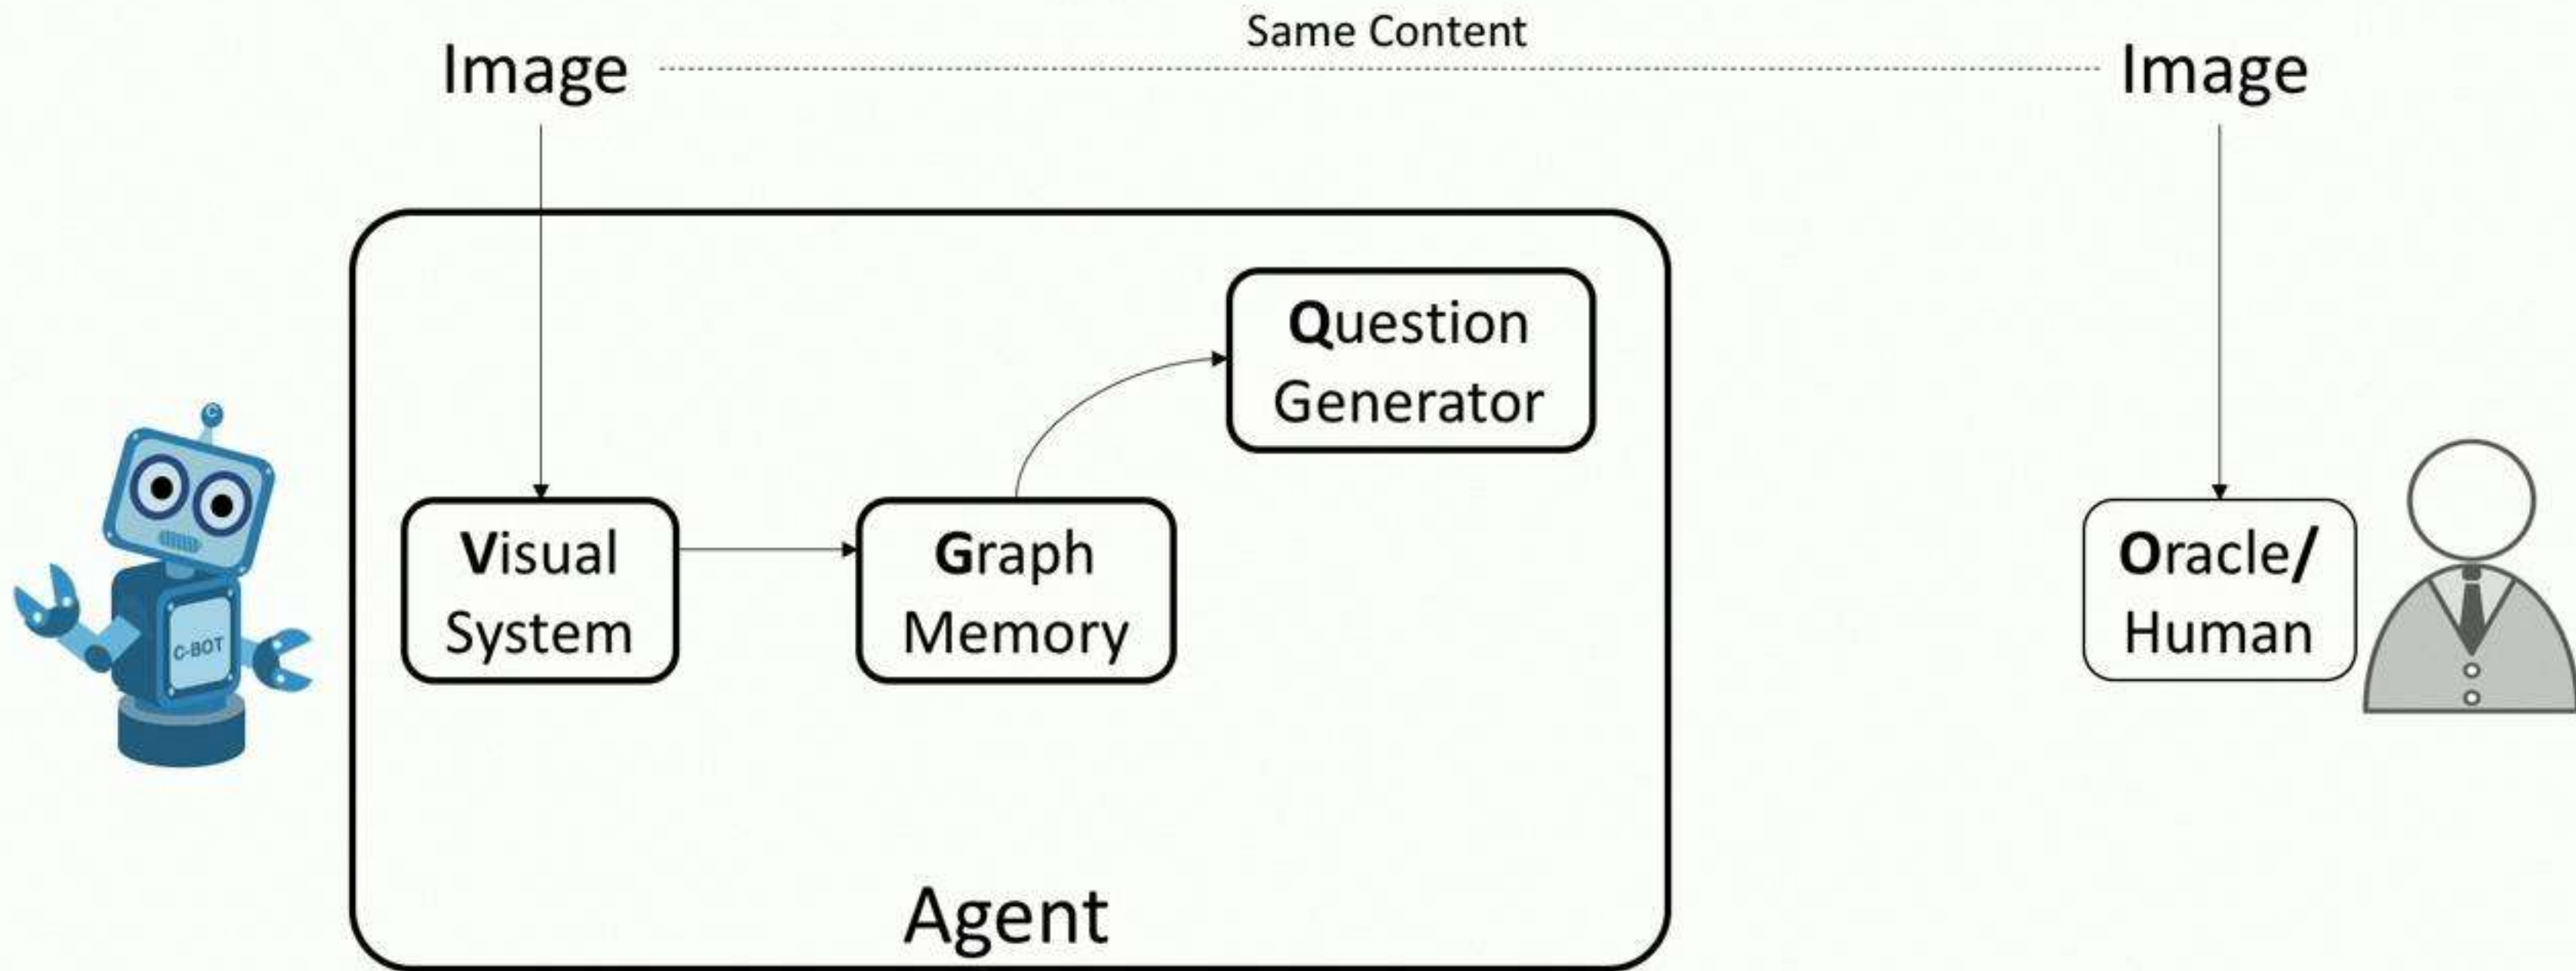
Human
(Oracle)



Agent Architecture

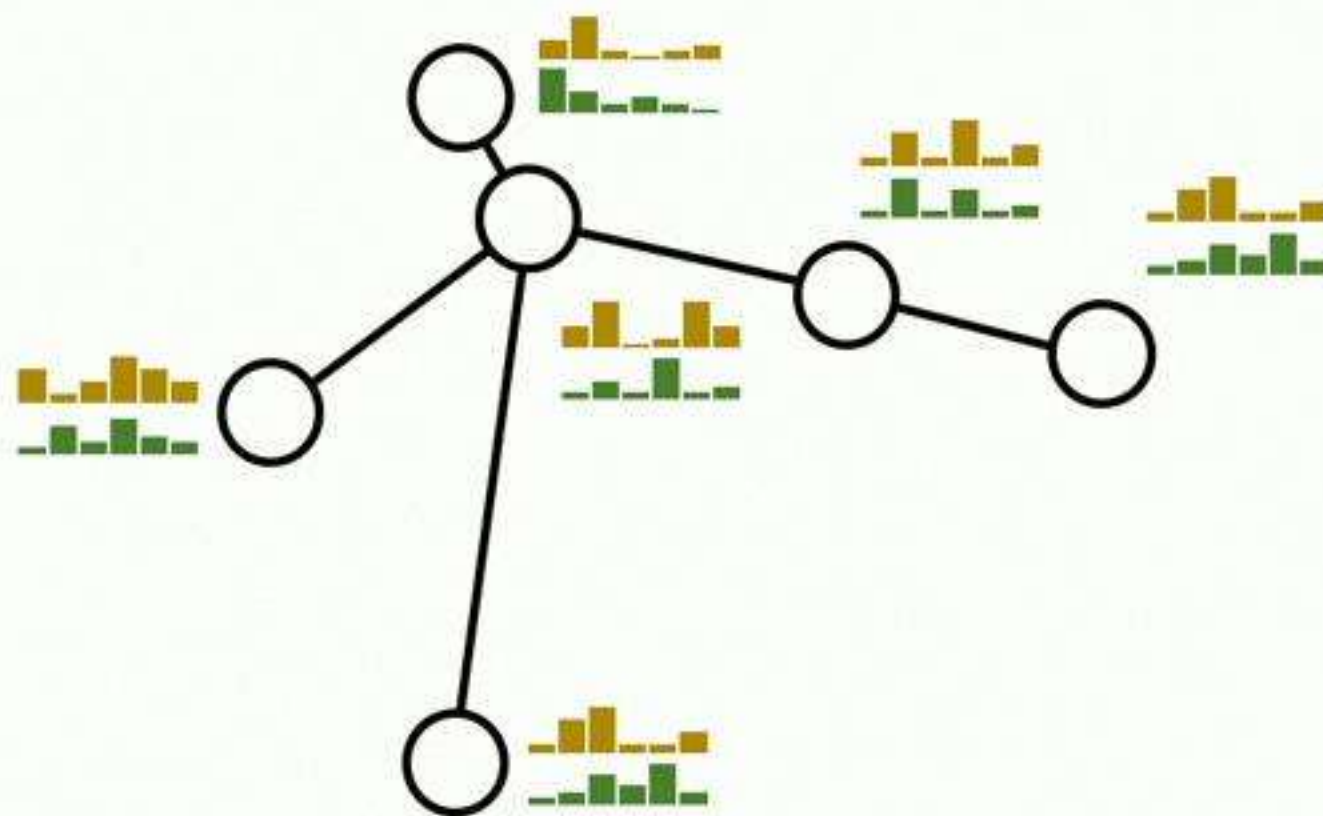
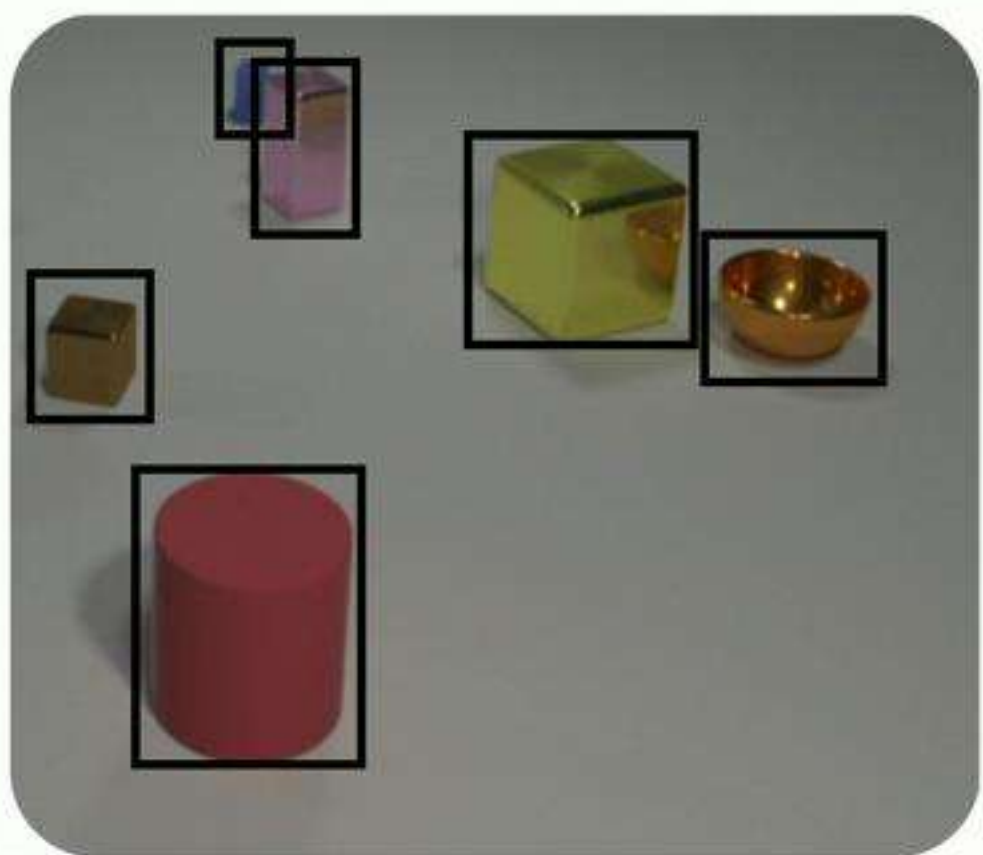


Agent Architecture



Agent Architecture

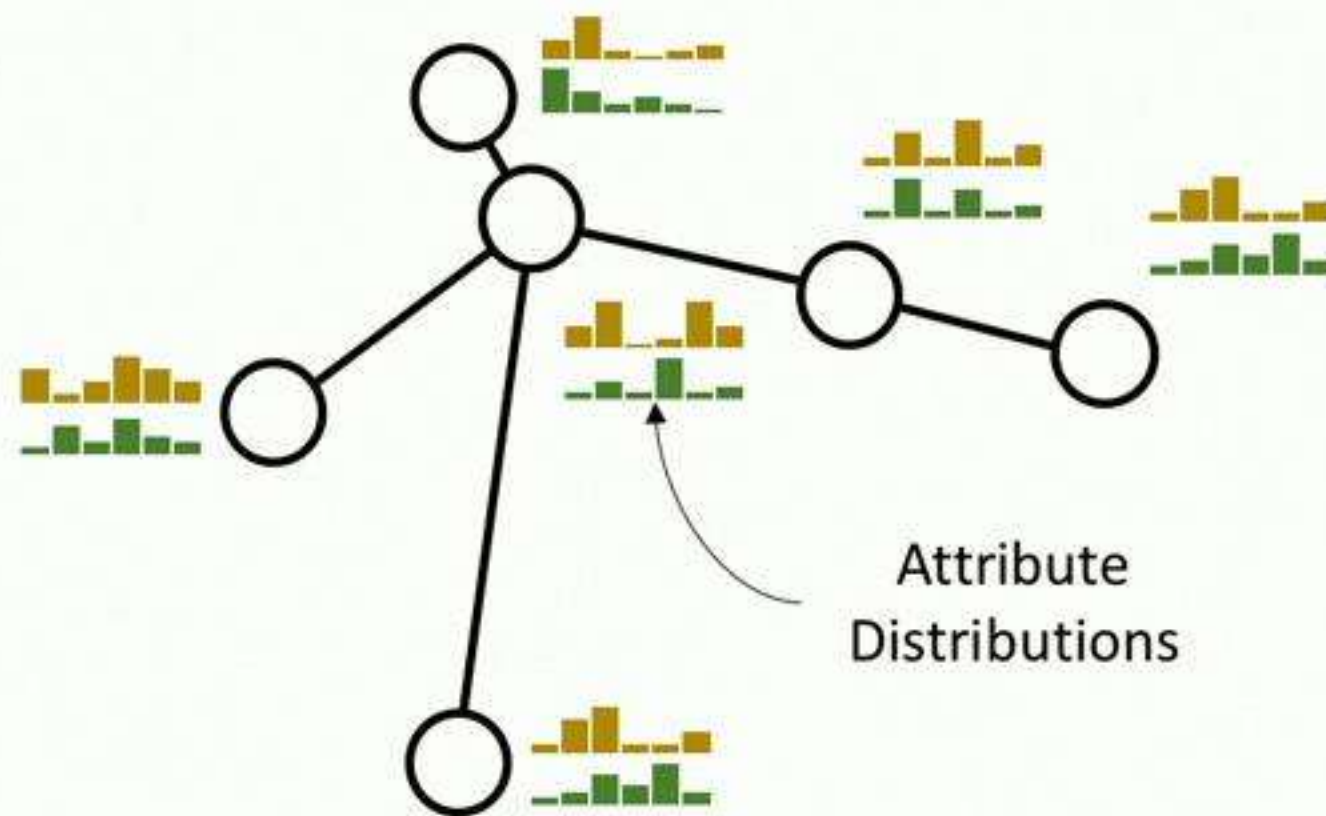
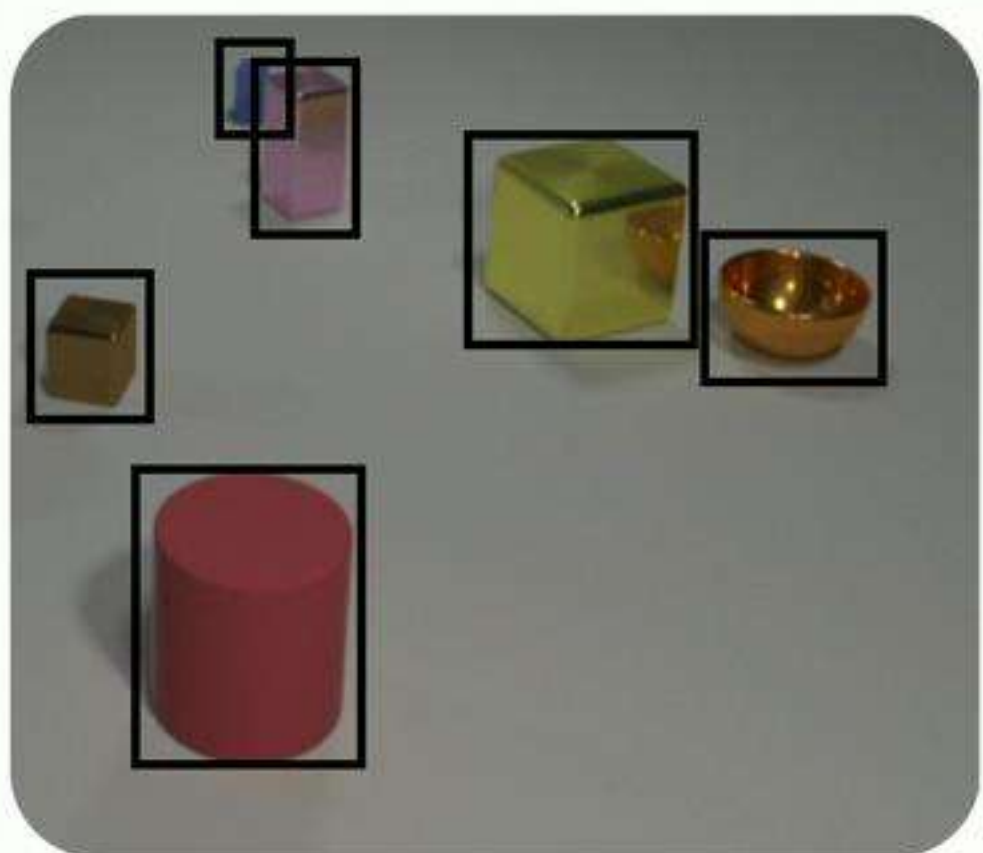
Visual System



Visual Graph

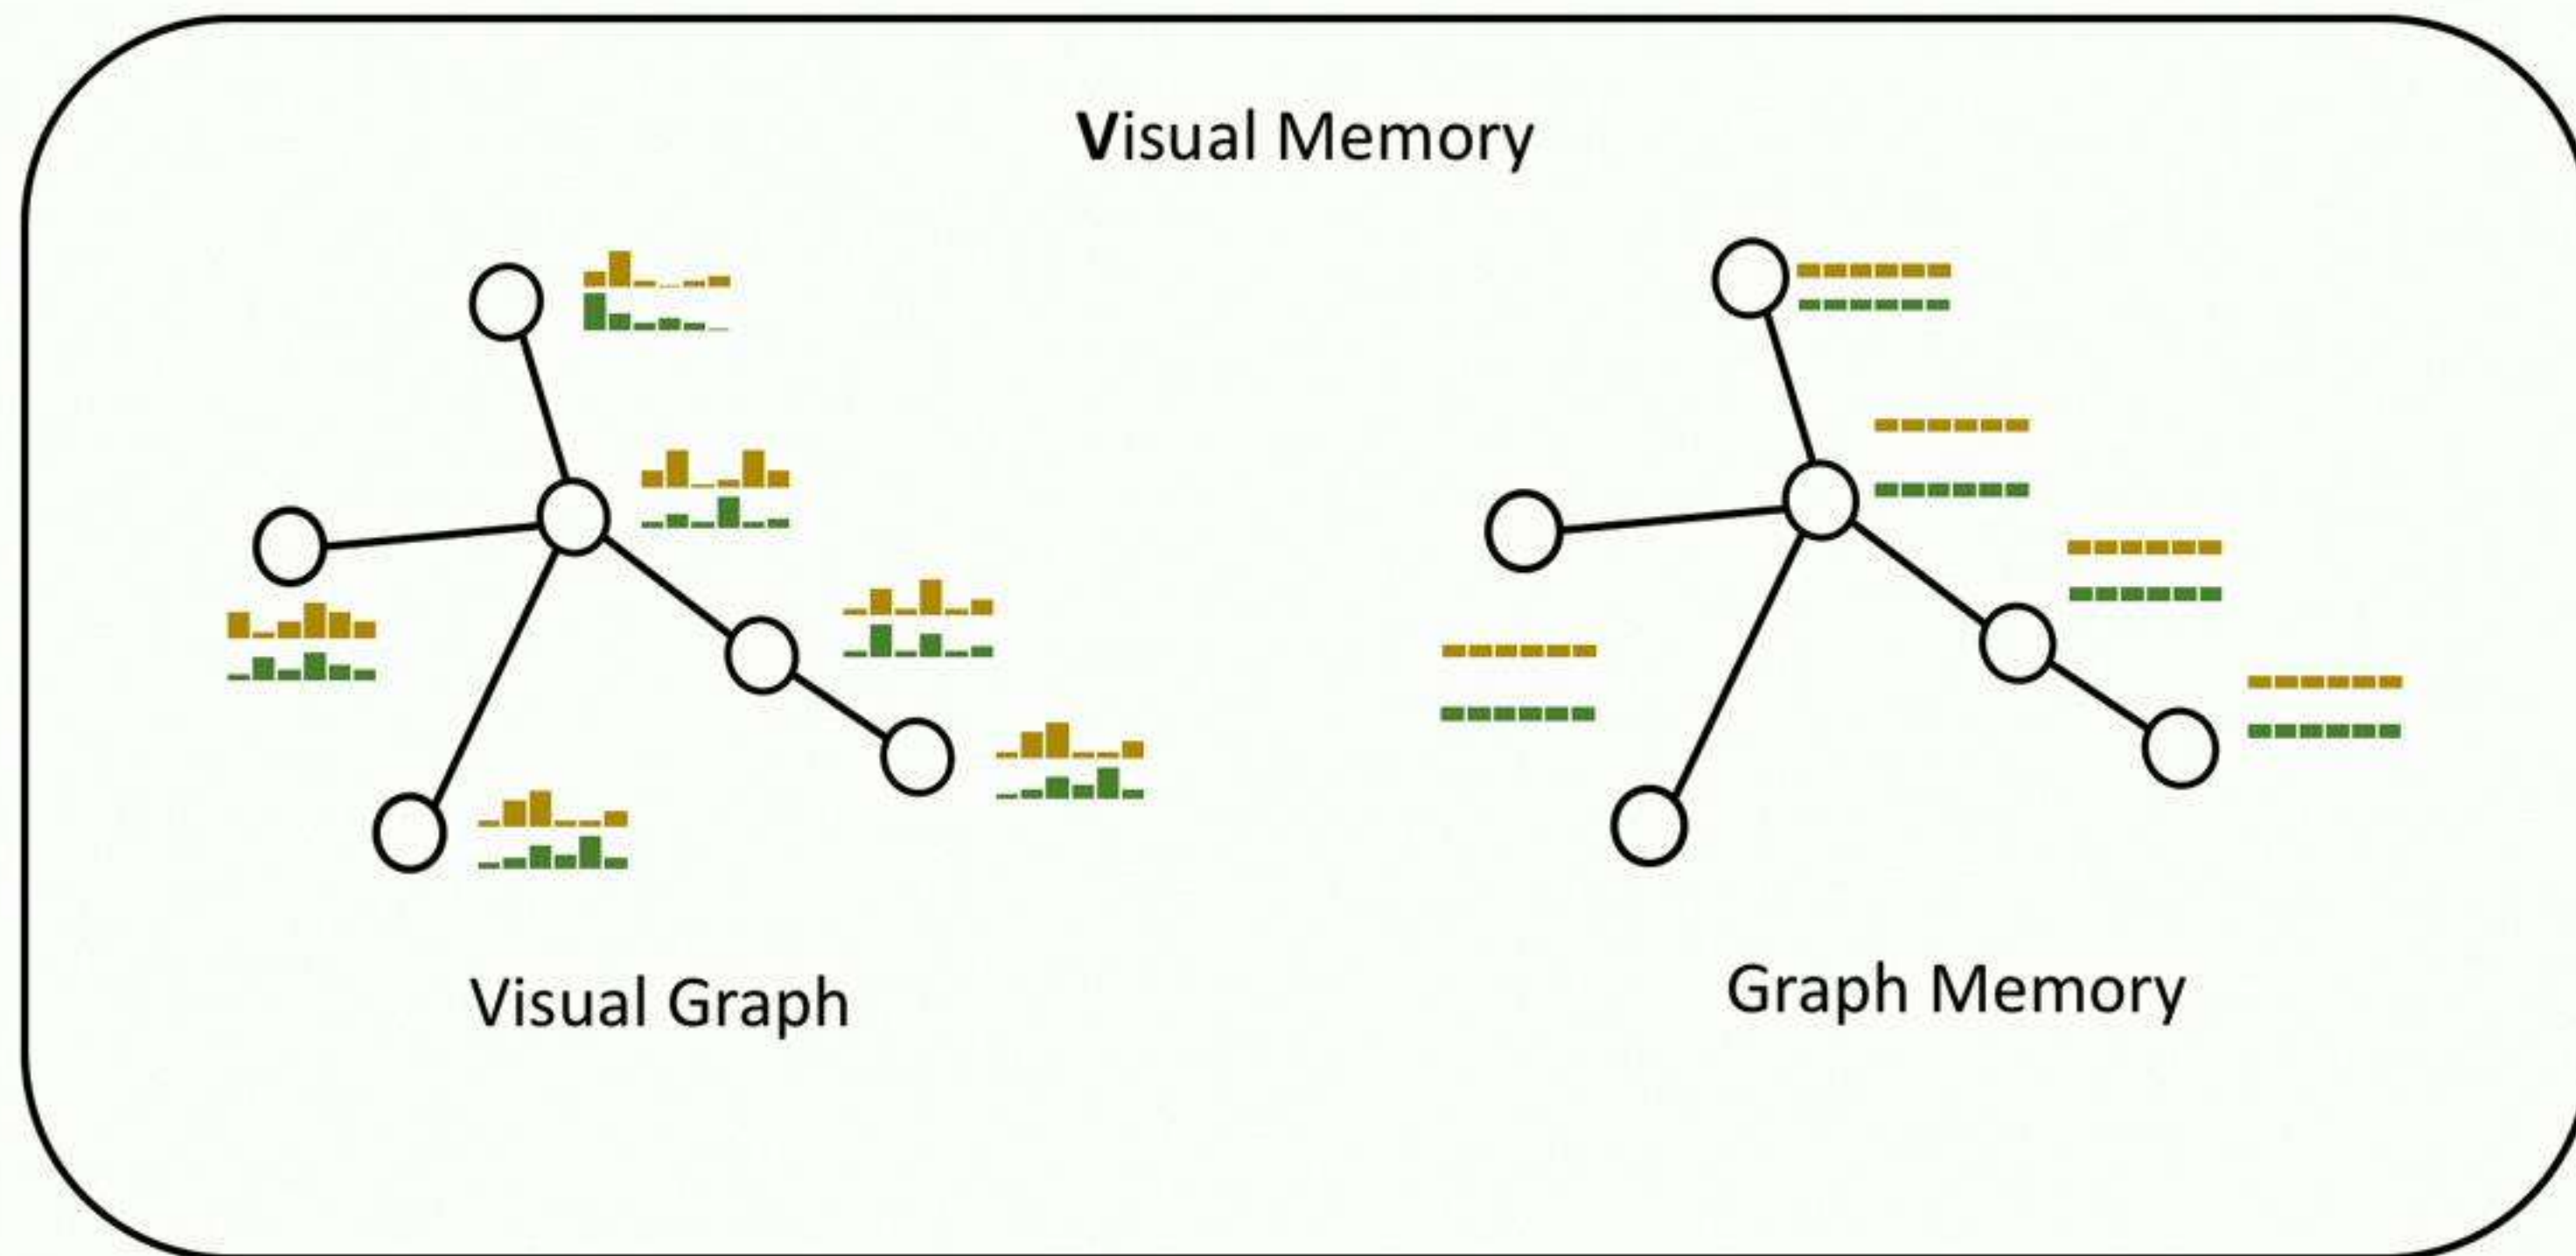
Agent Architecture

Visual System

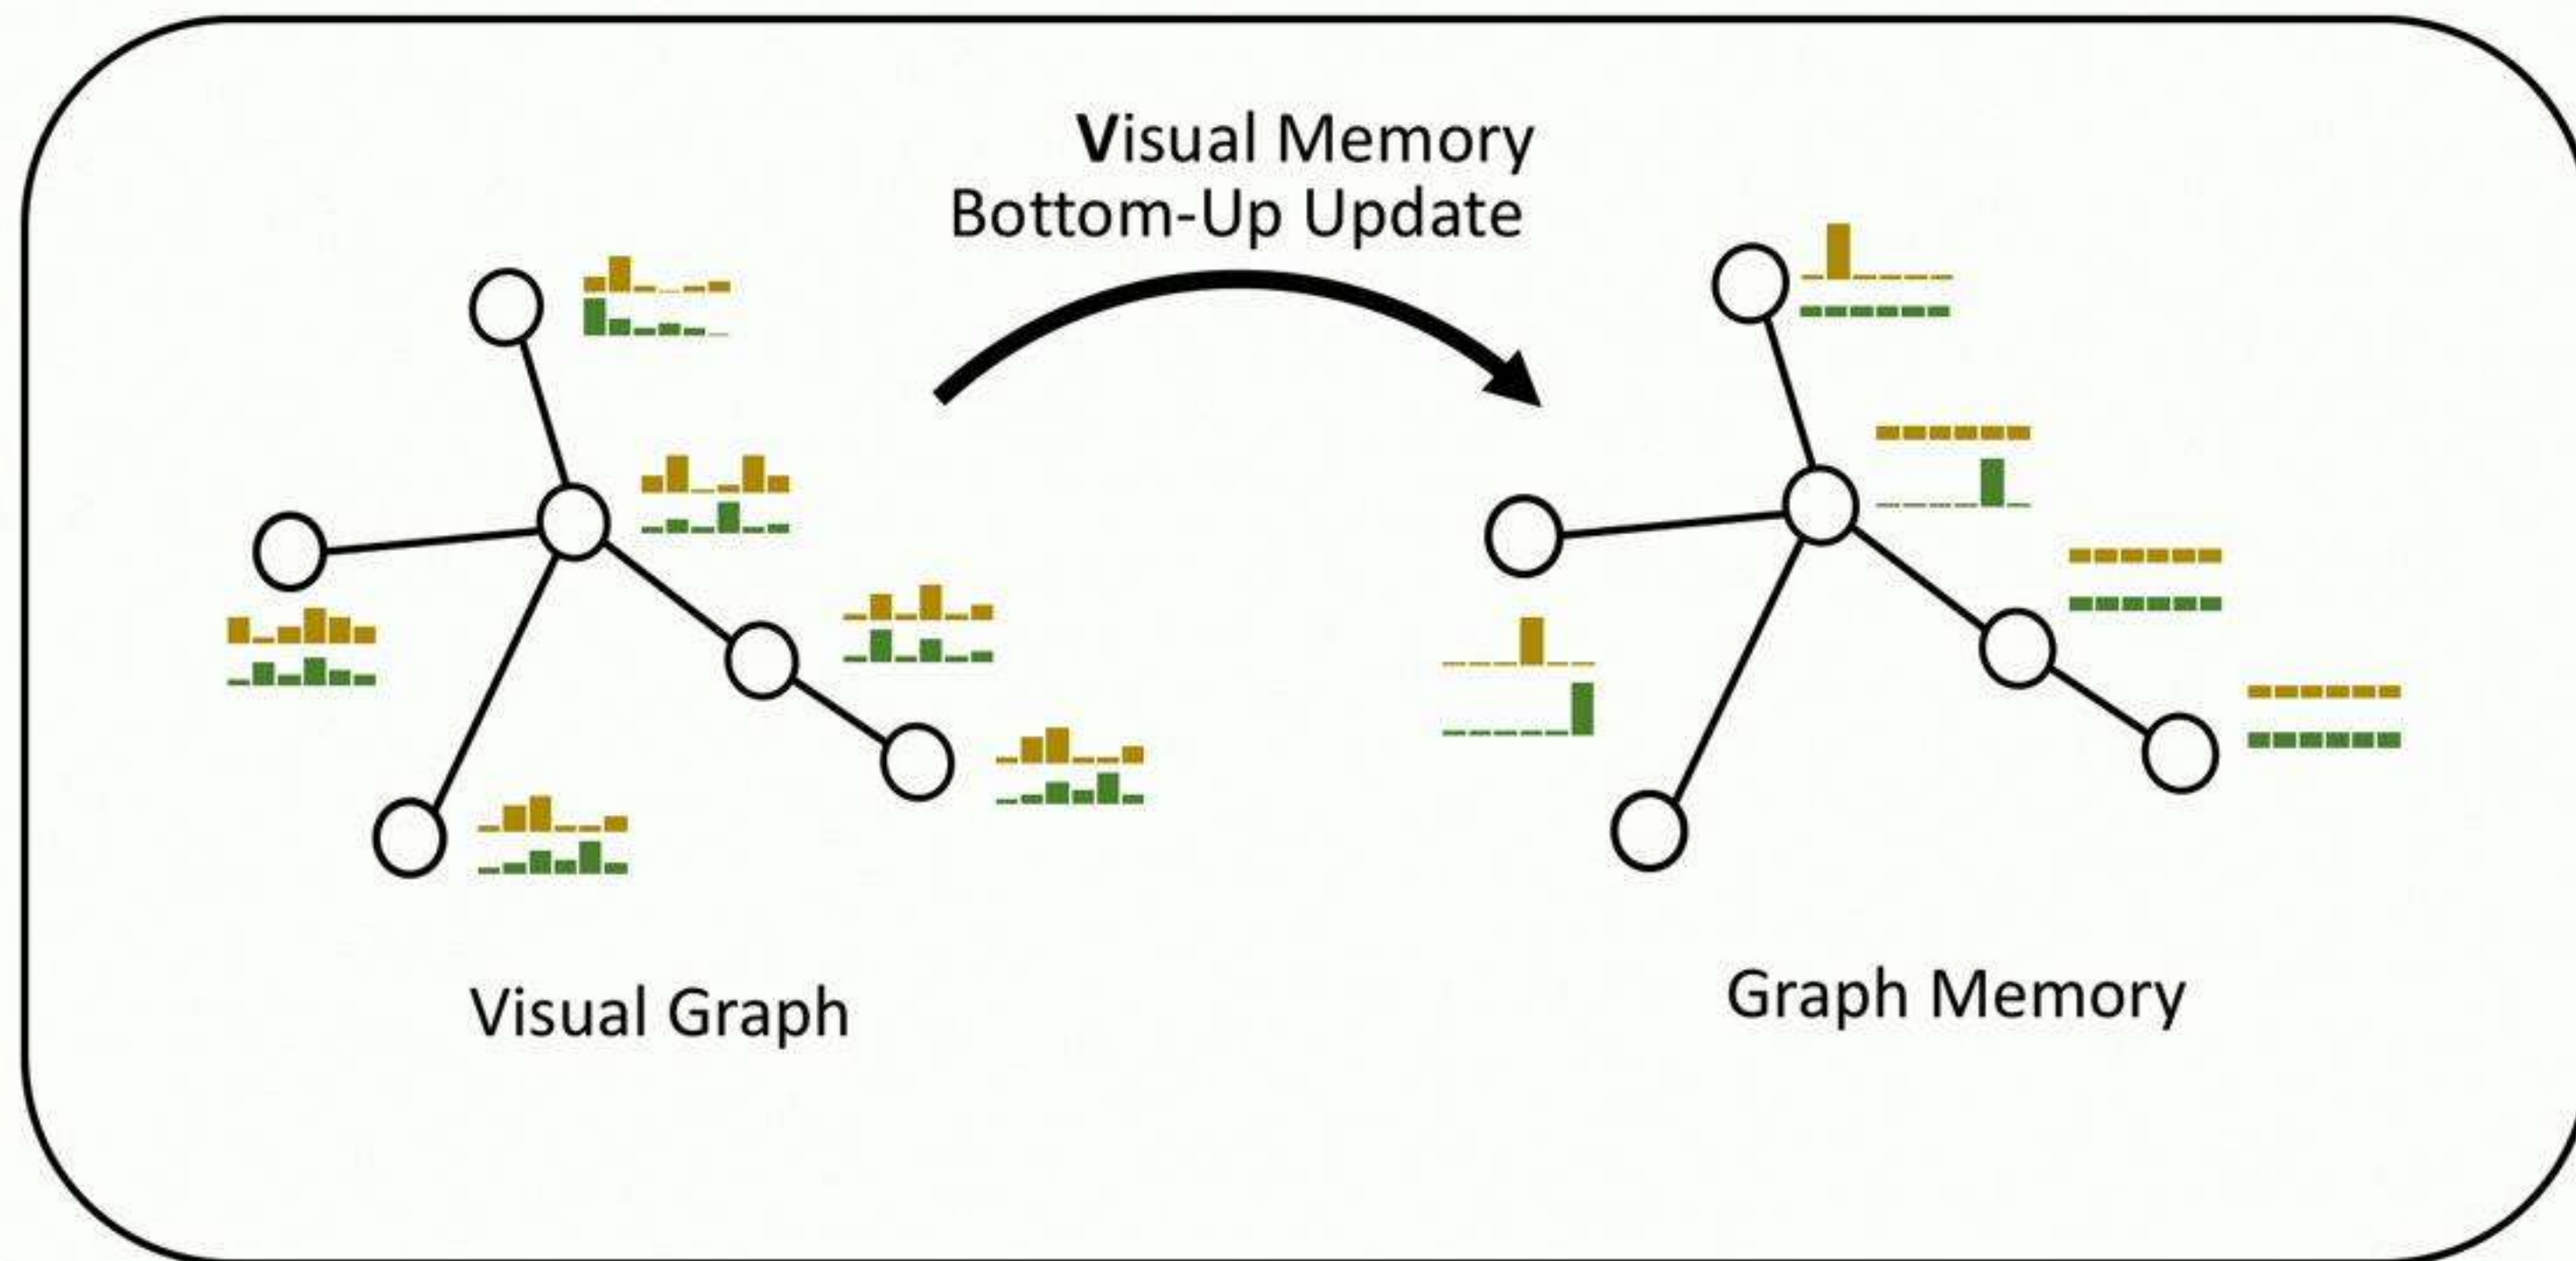


Visual Graph

Agent Architecture

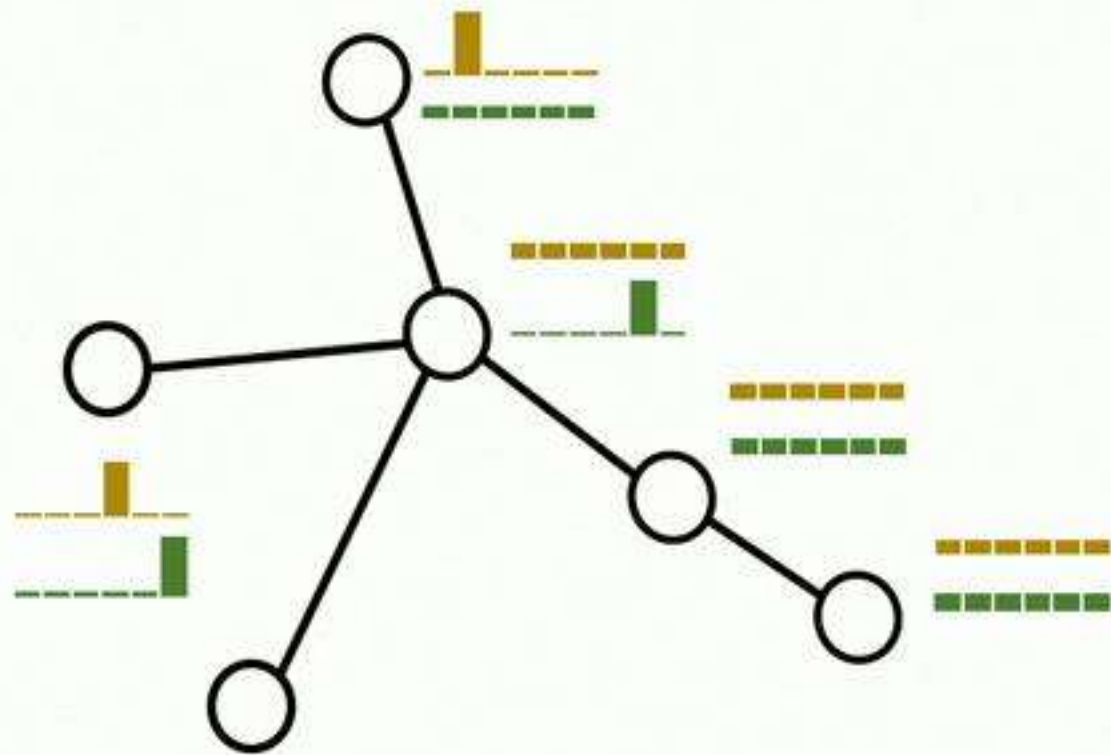


Bottom-up Update



Top-down Update

Visual Memory

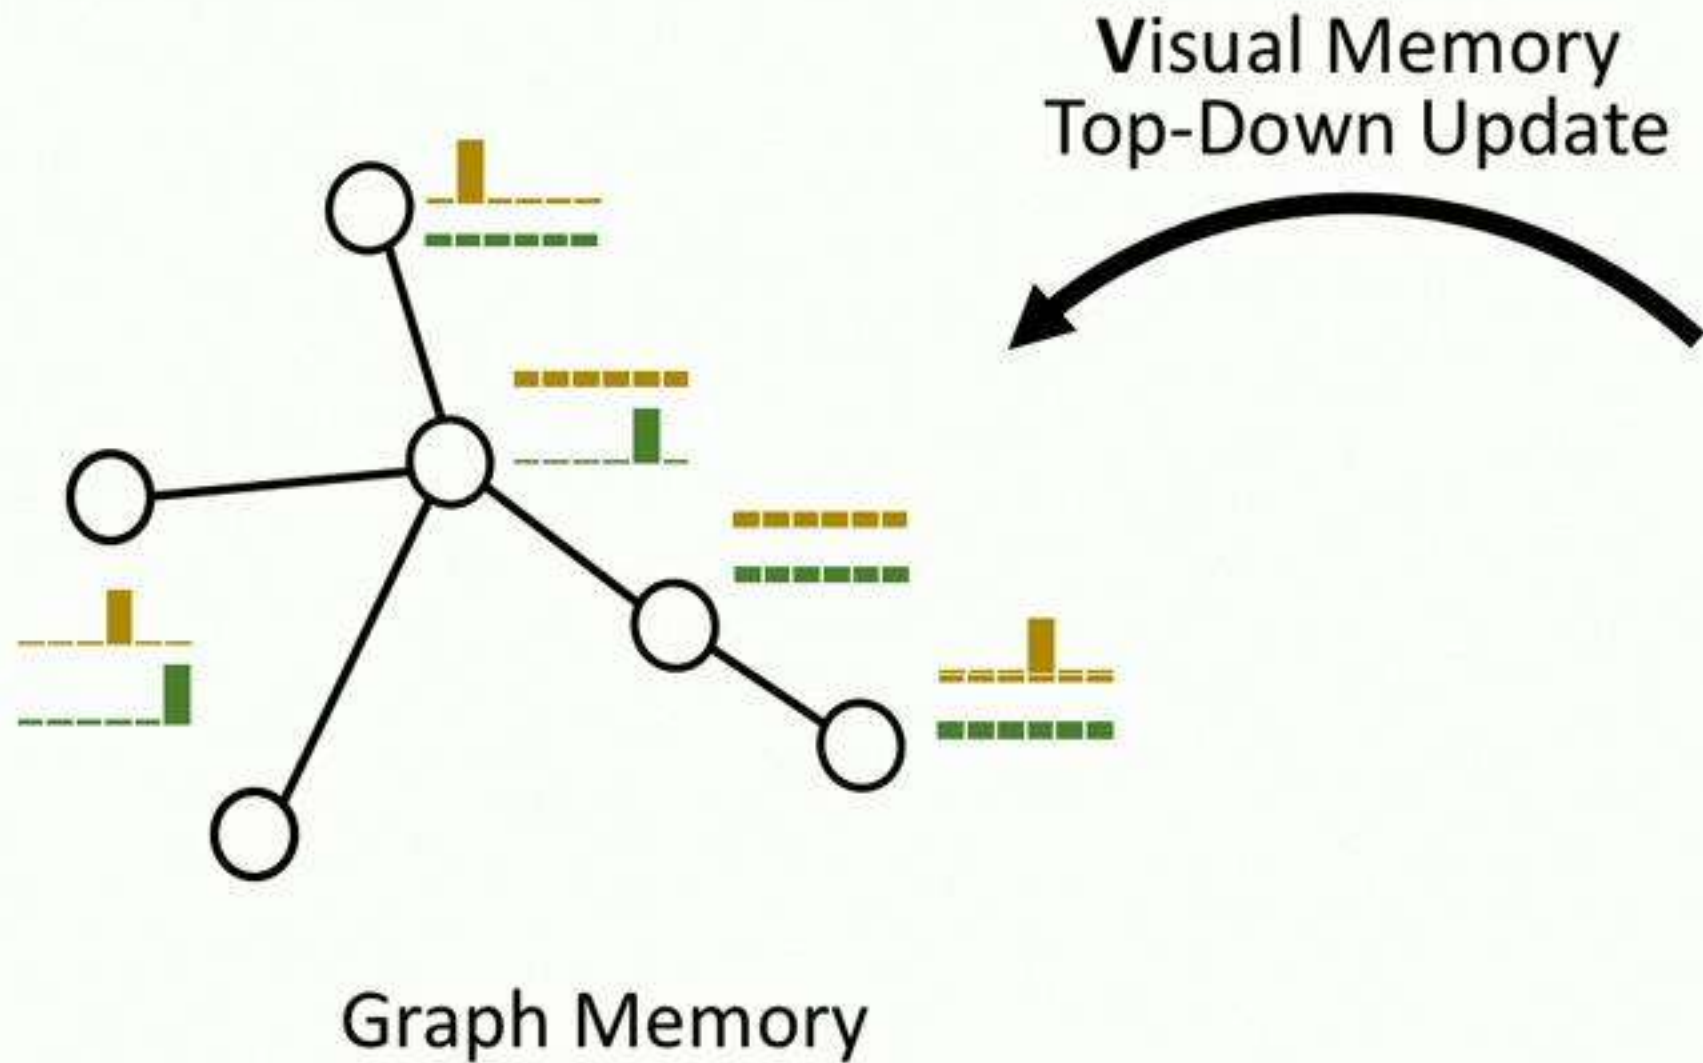


Memory Graph

Q: What is the color of
the right most object?

A: Orange

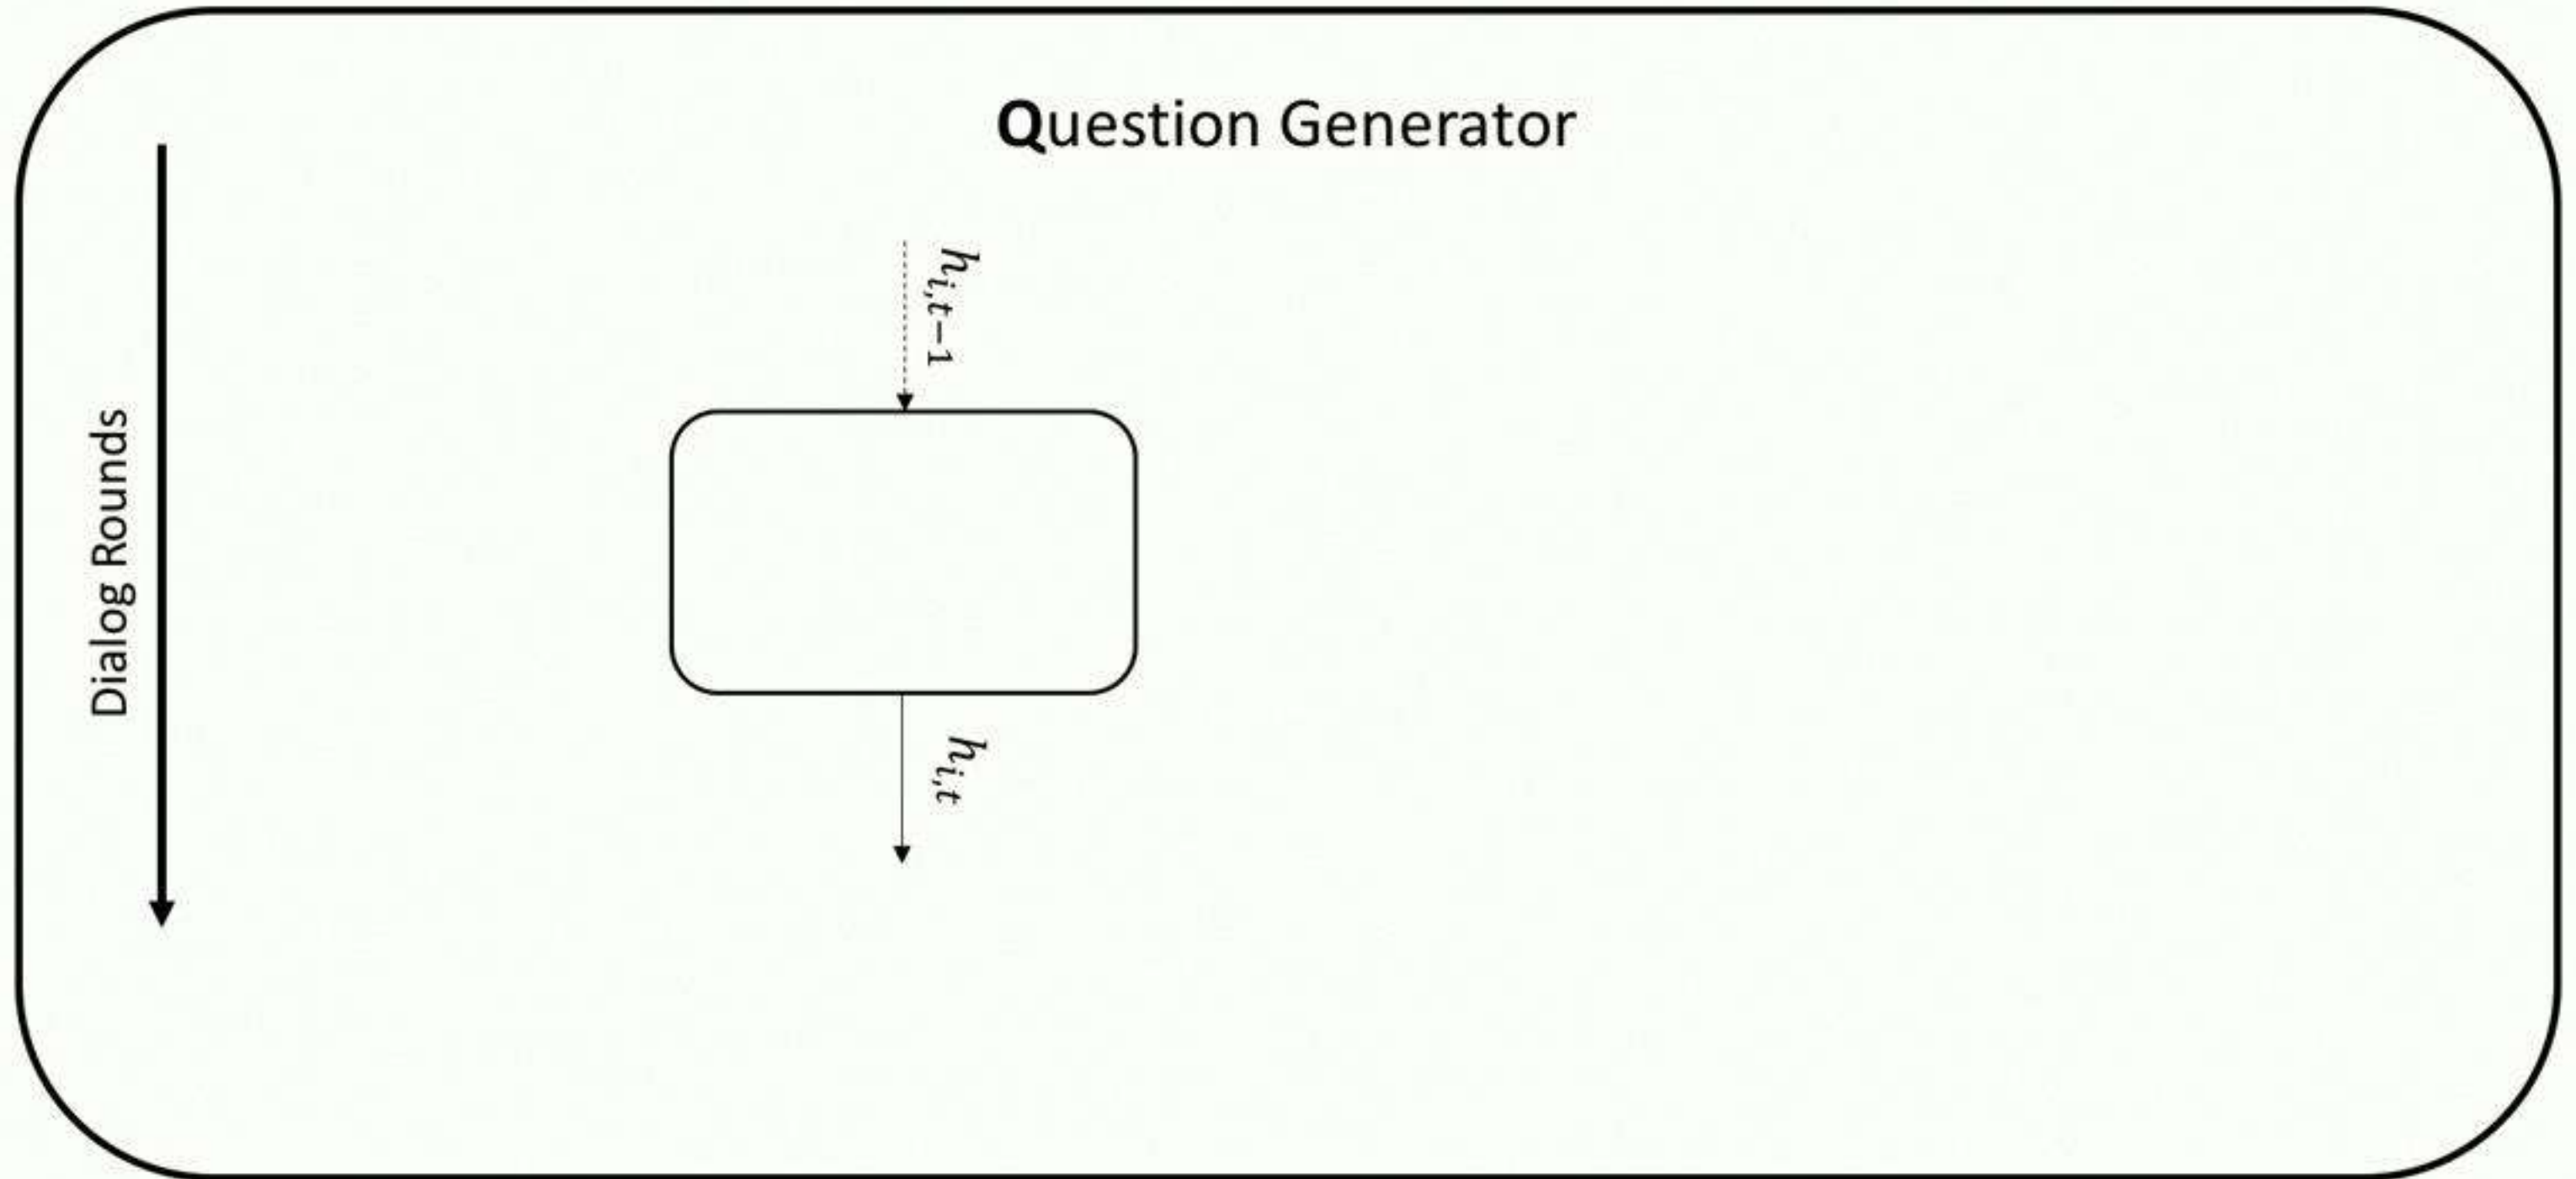
Top-down Update



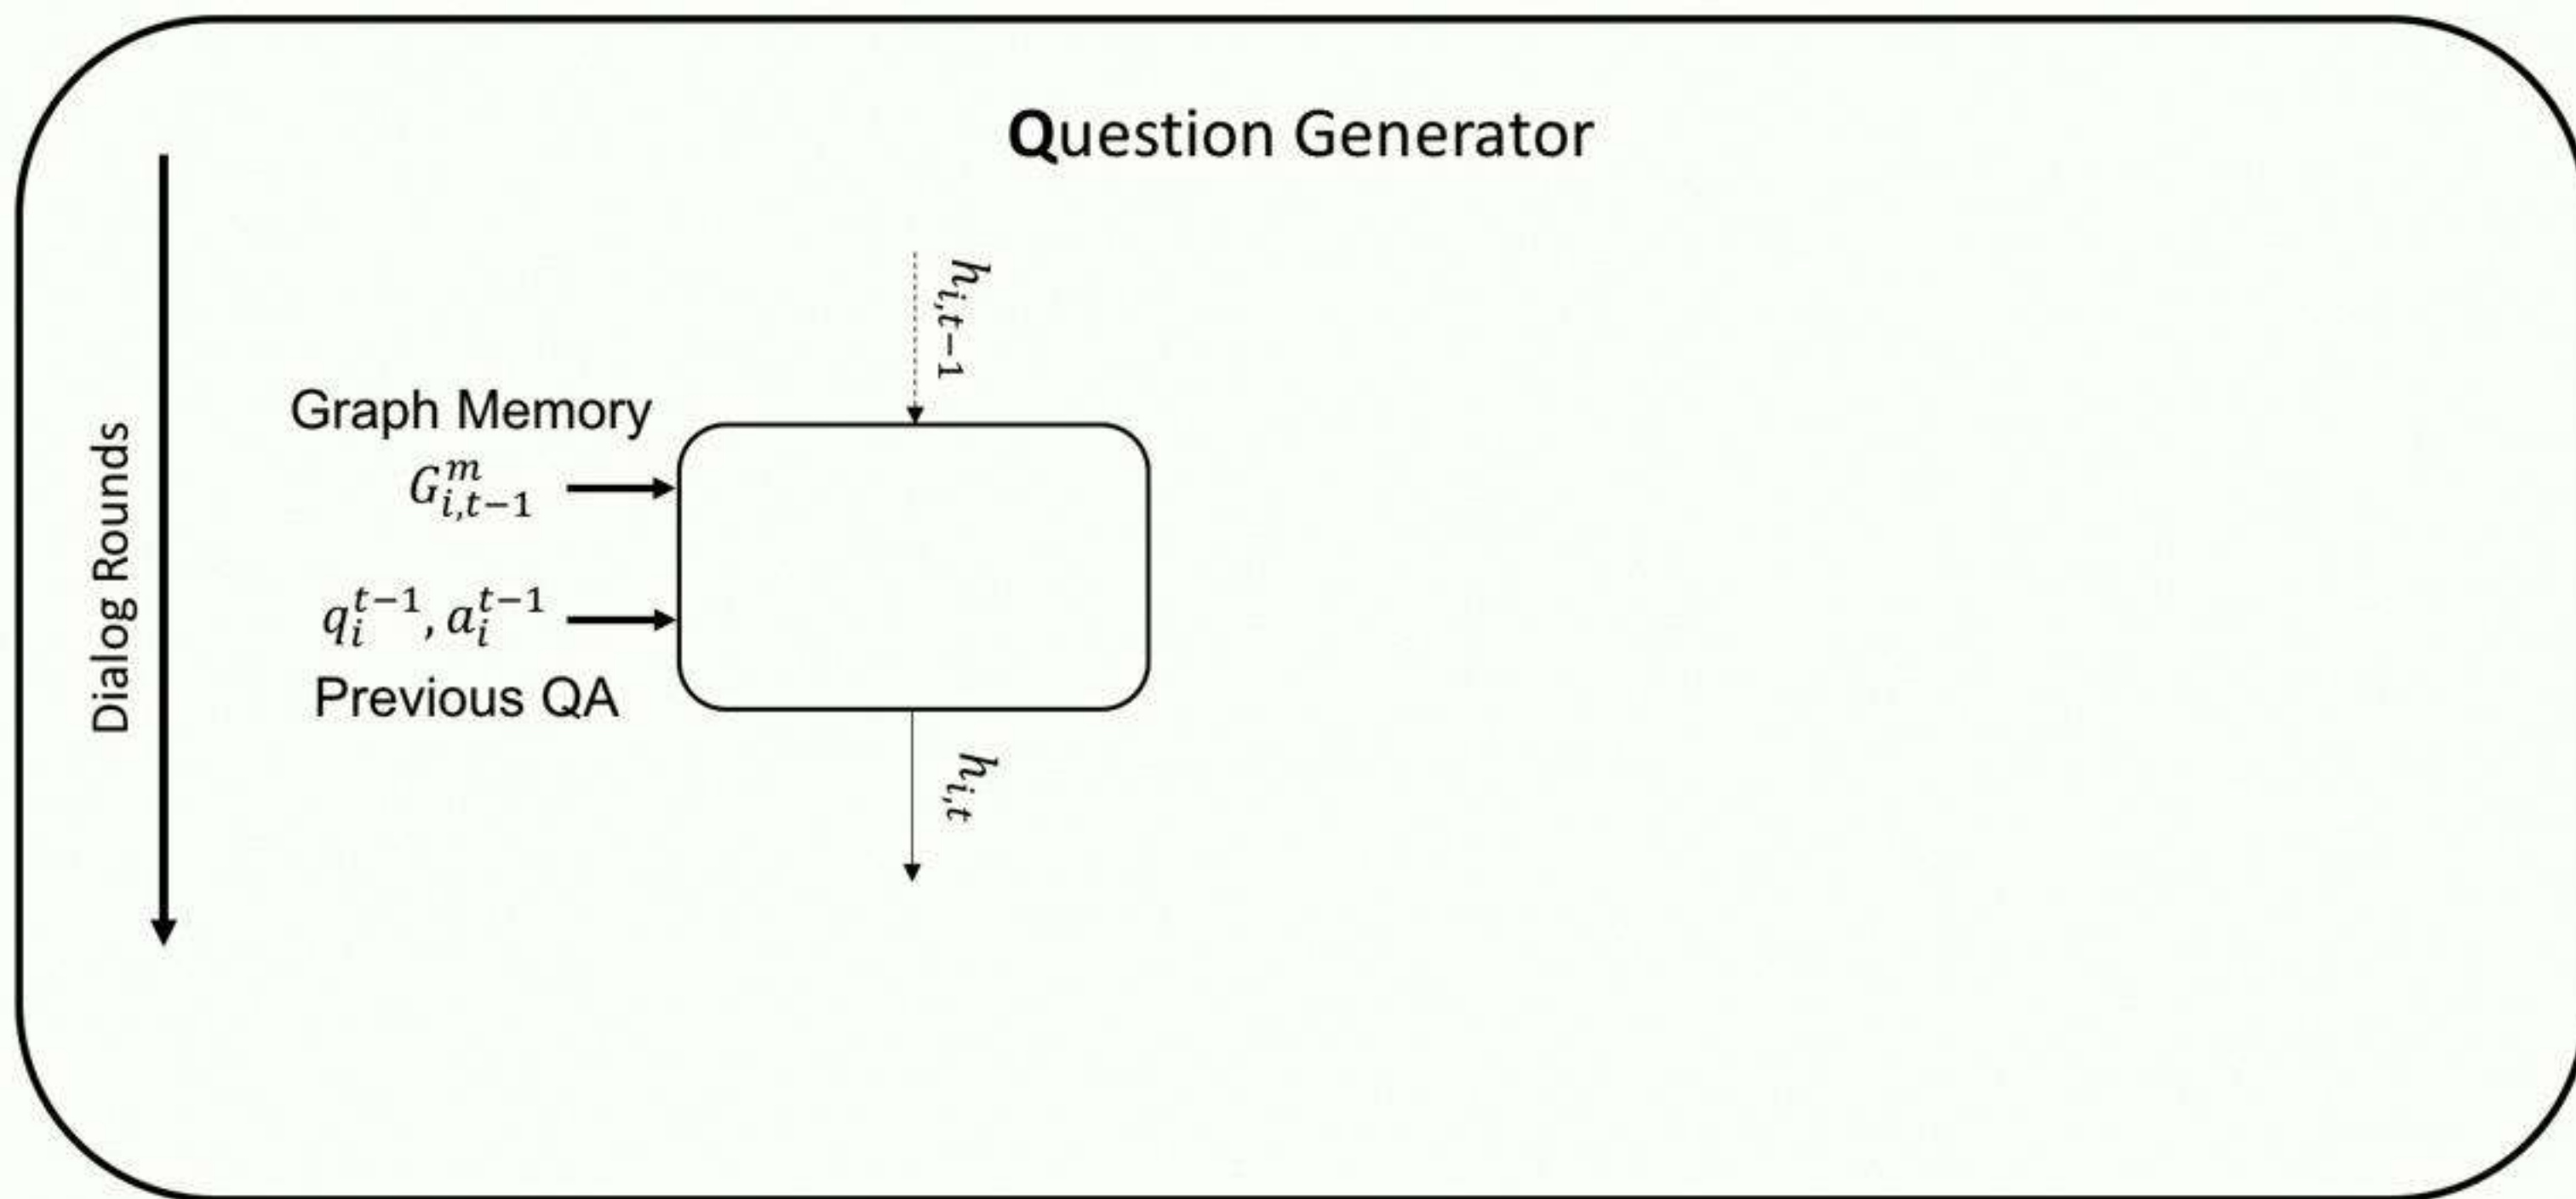
Q: What is the color of
the right most object?

A: Orange

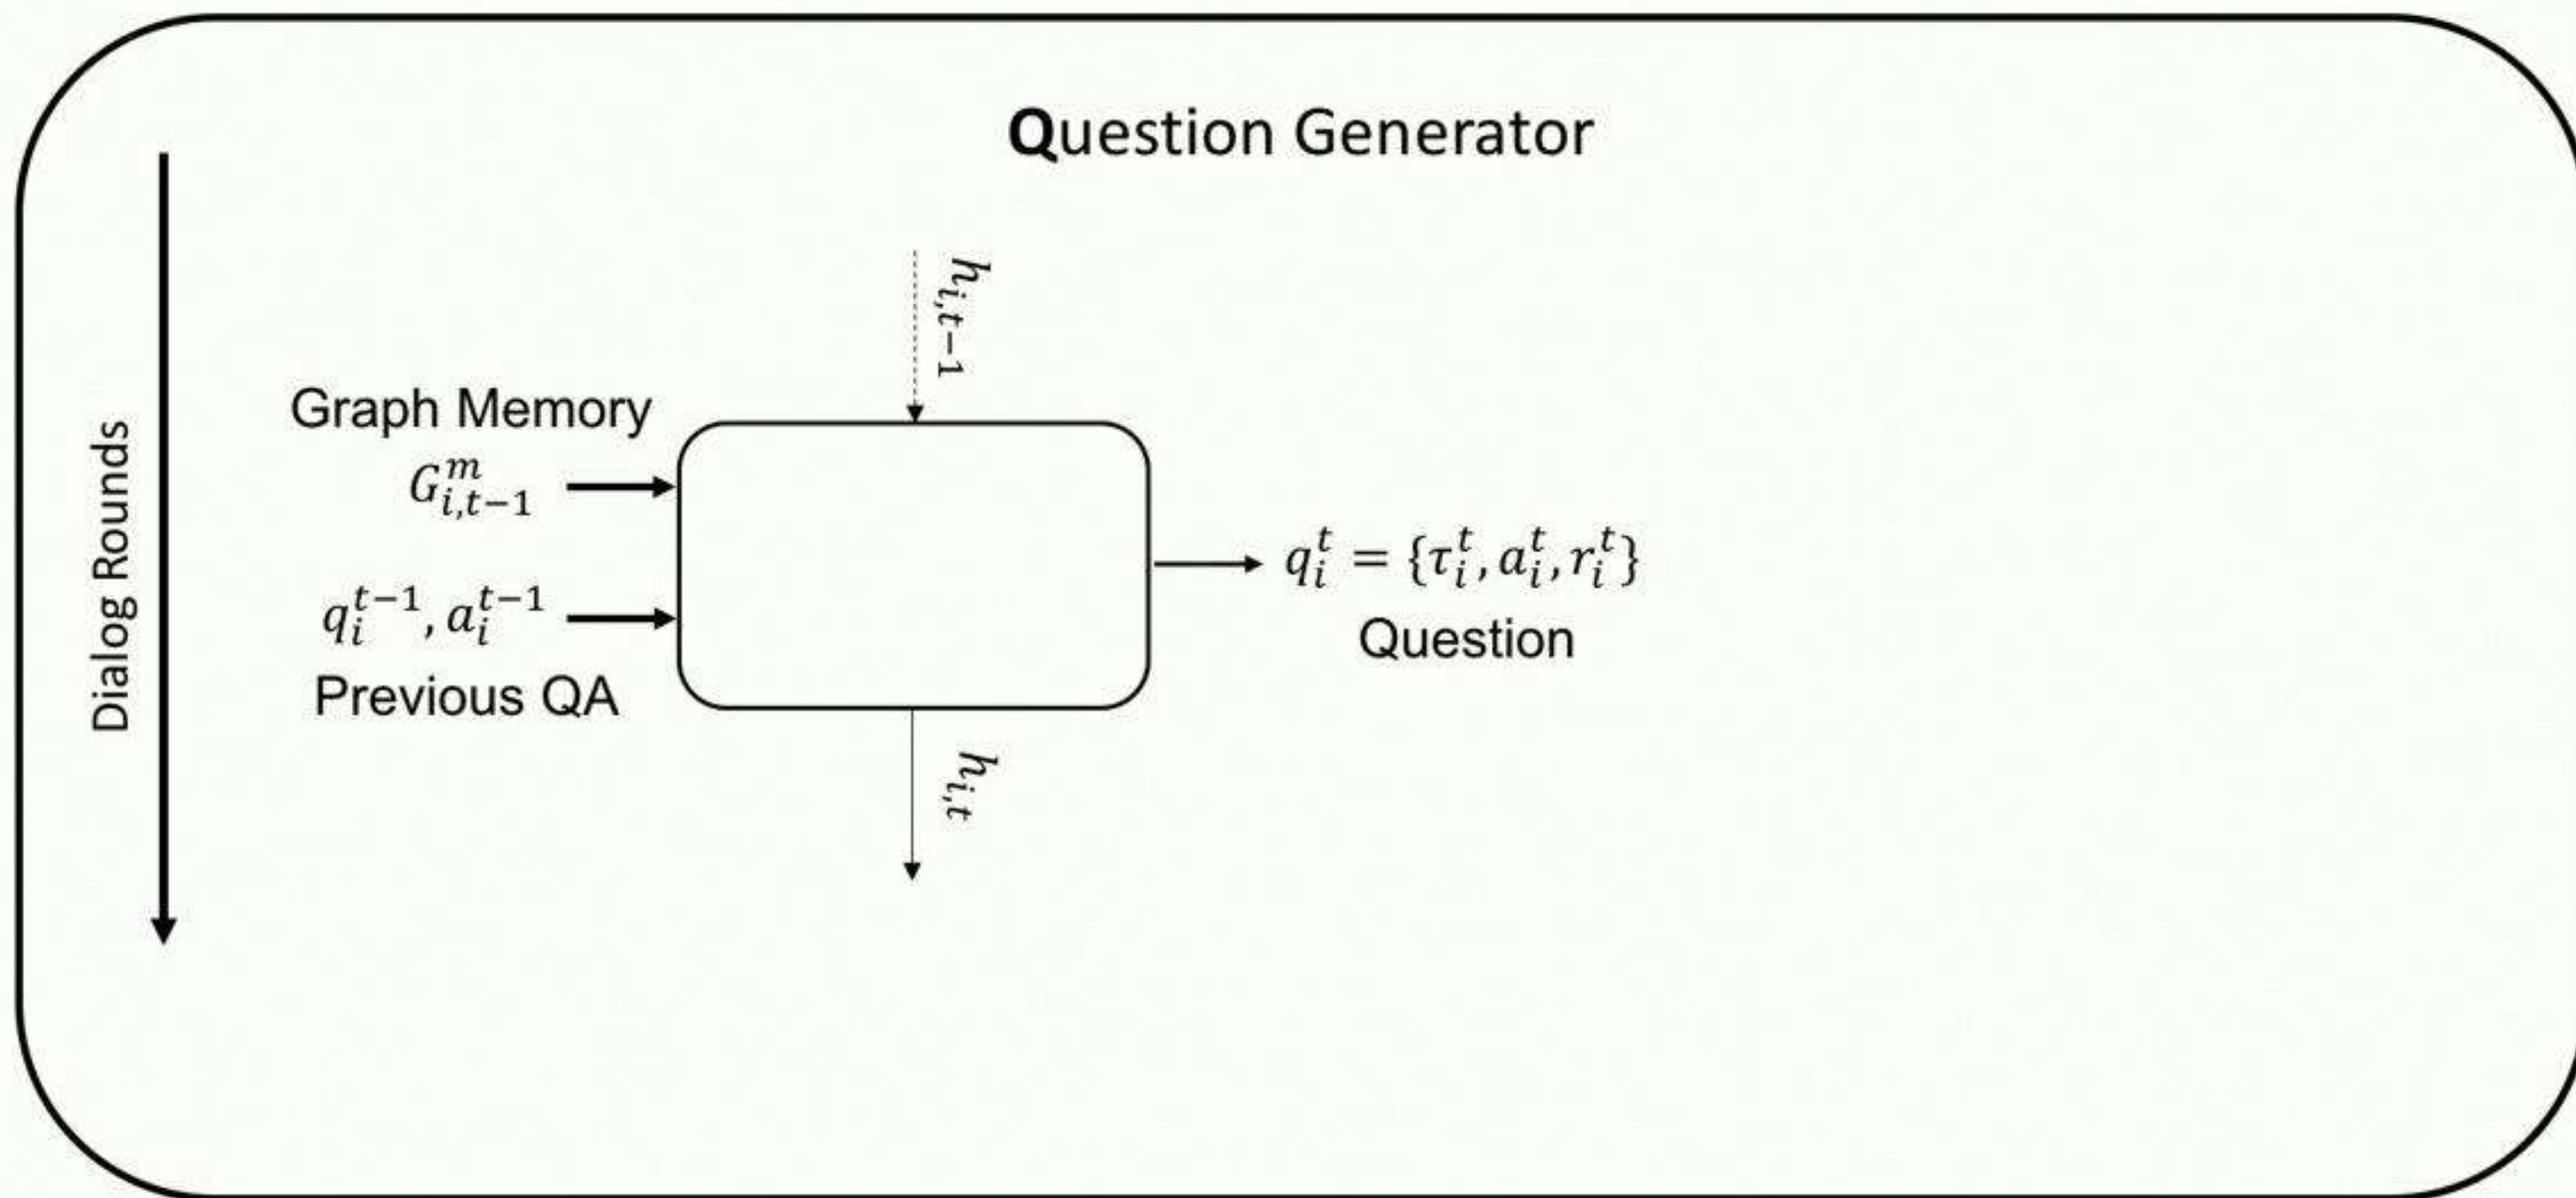
Agent Architecture



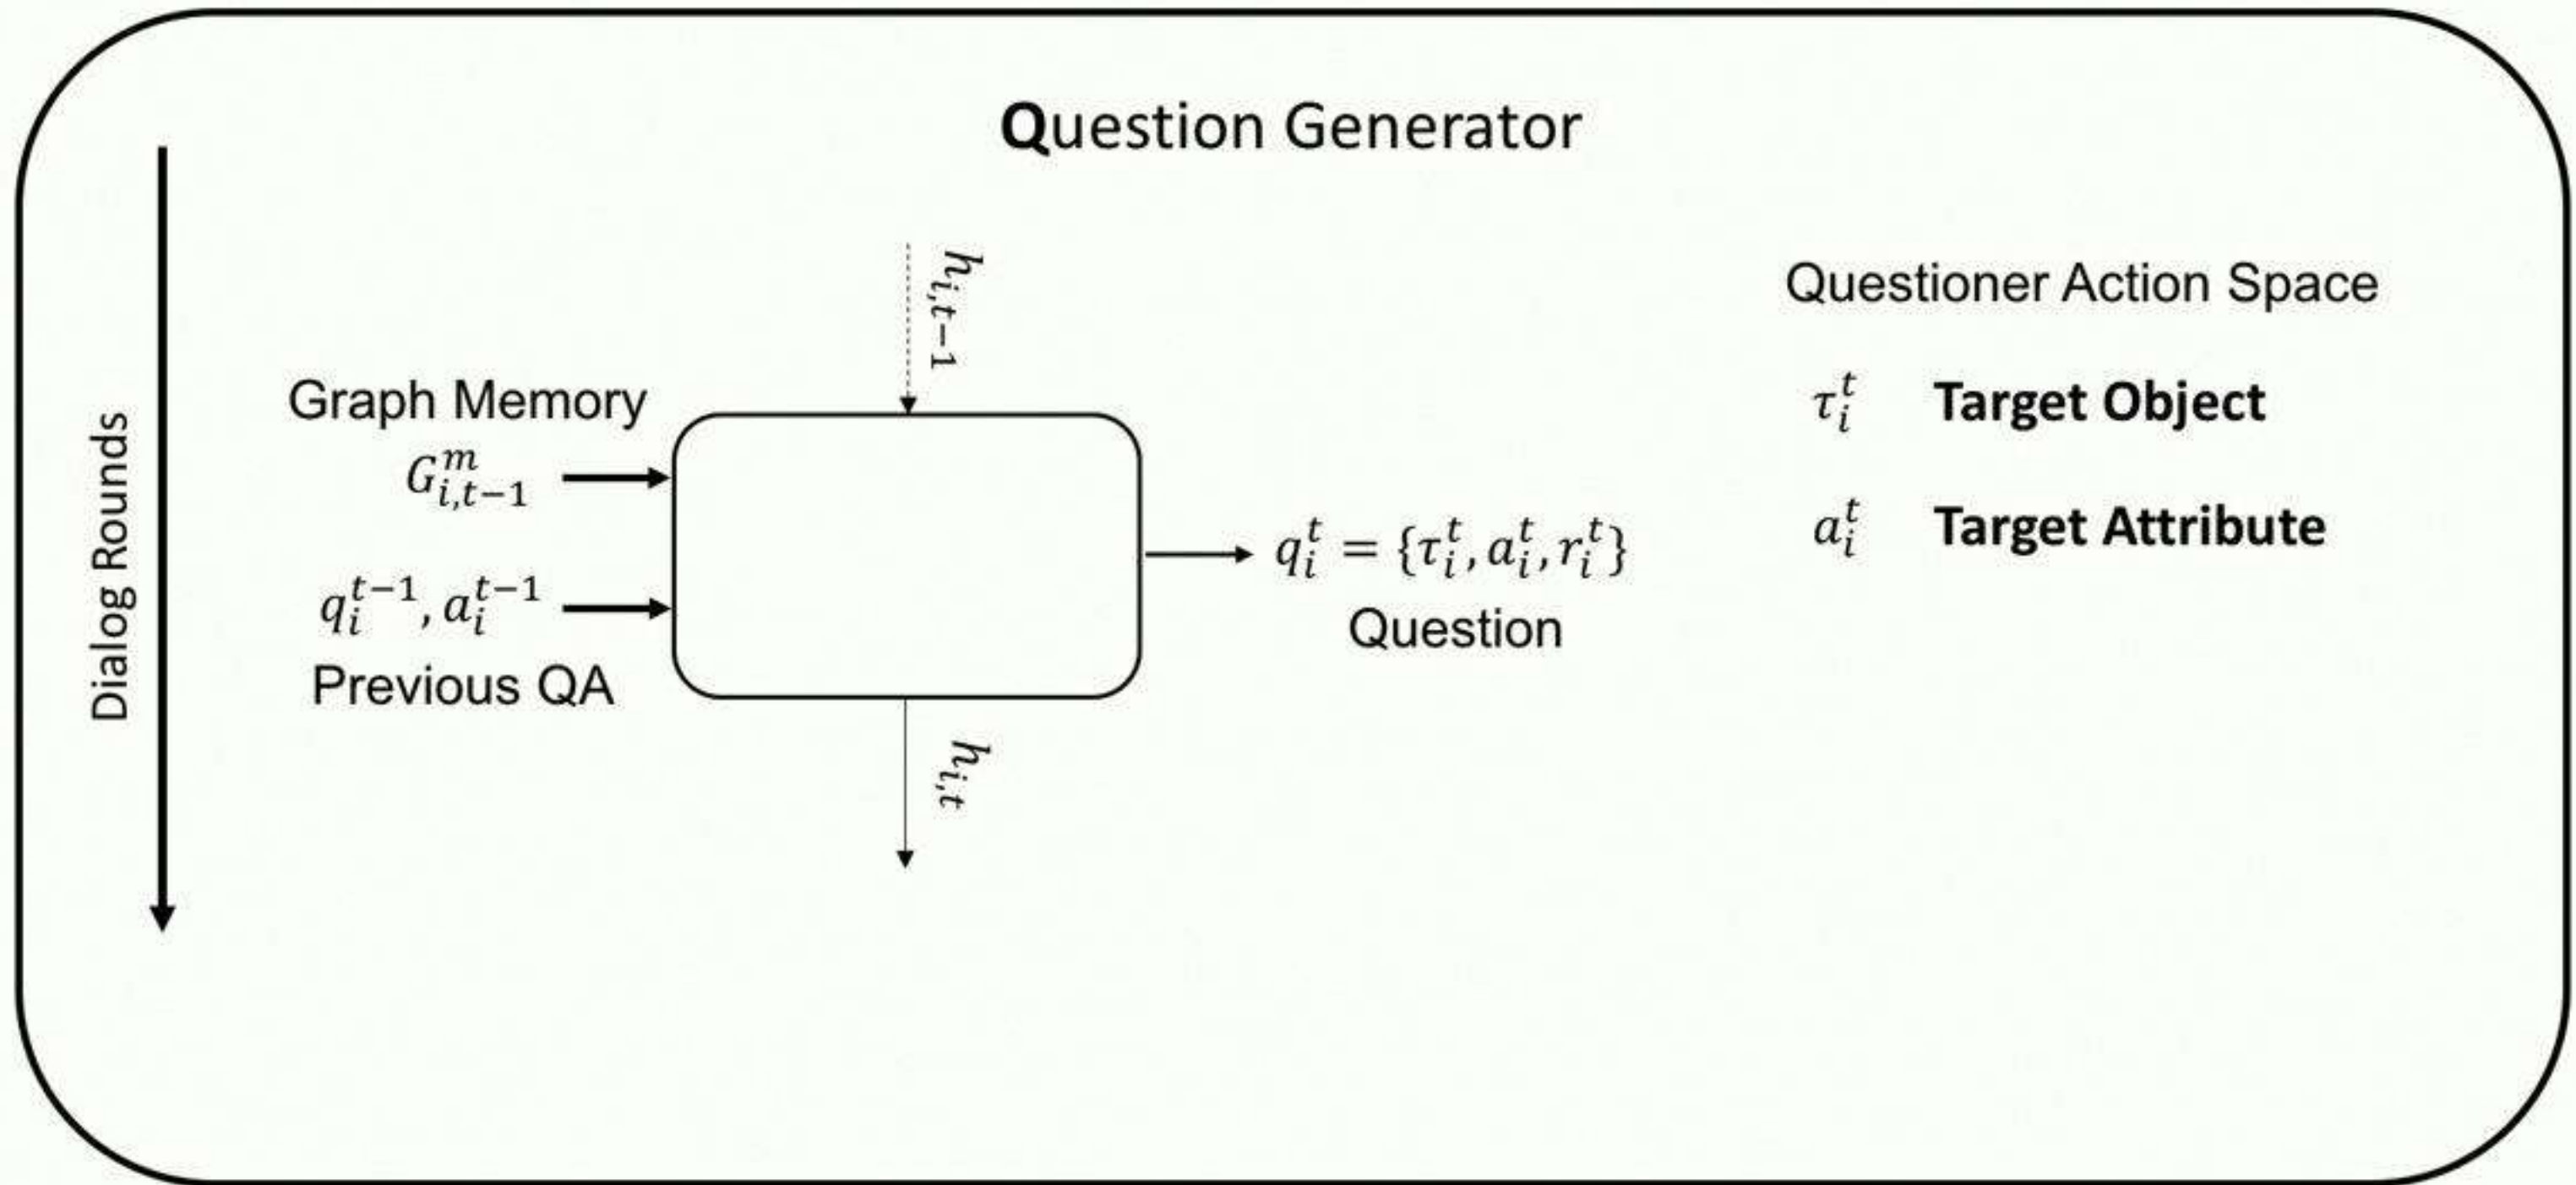
Agent Architecture



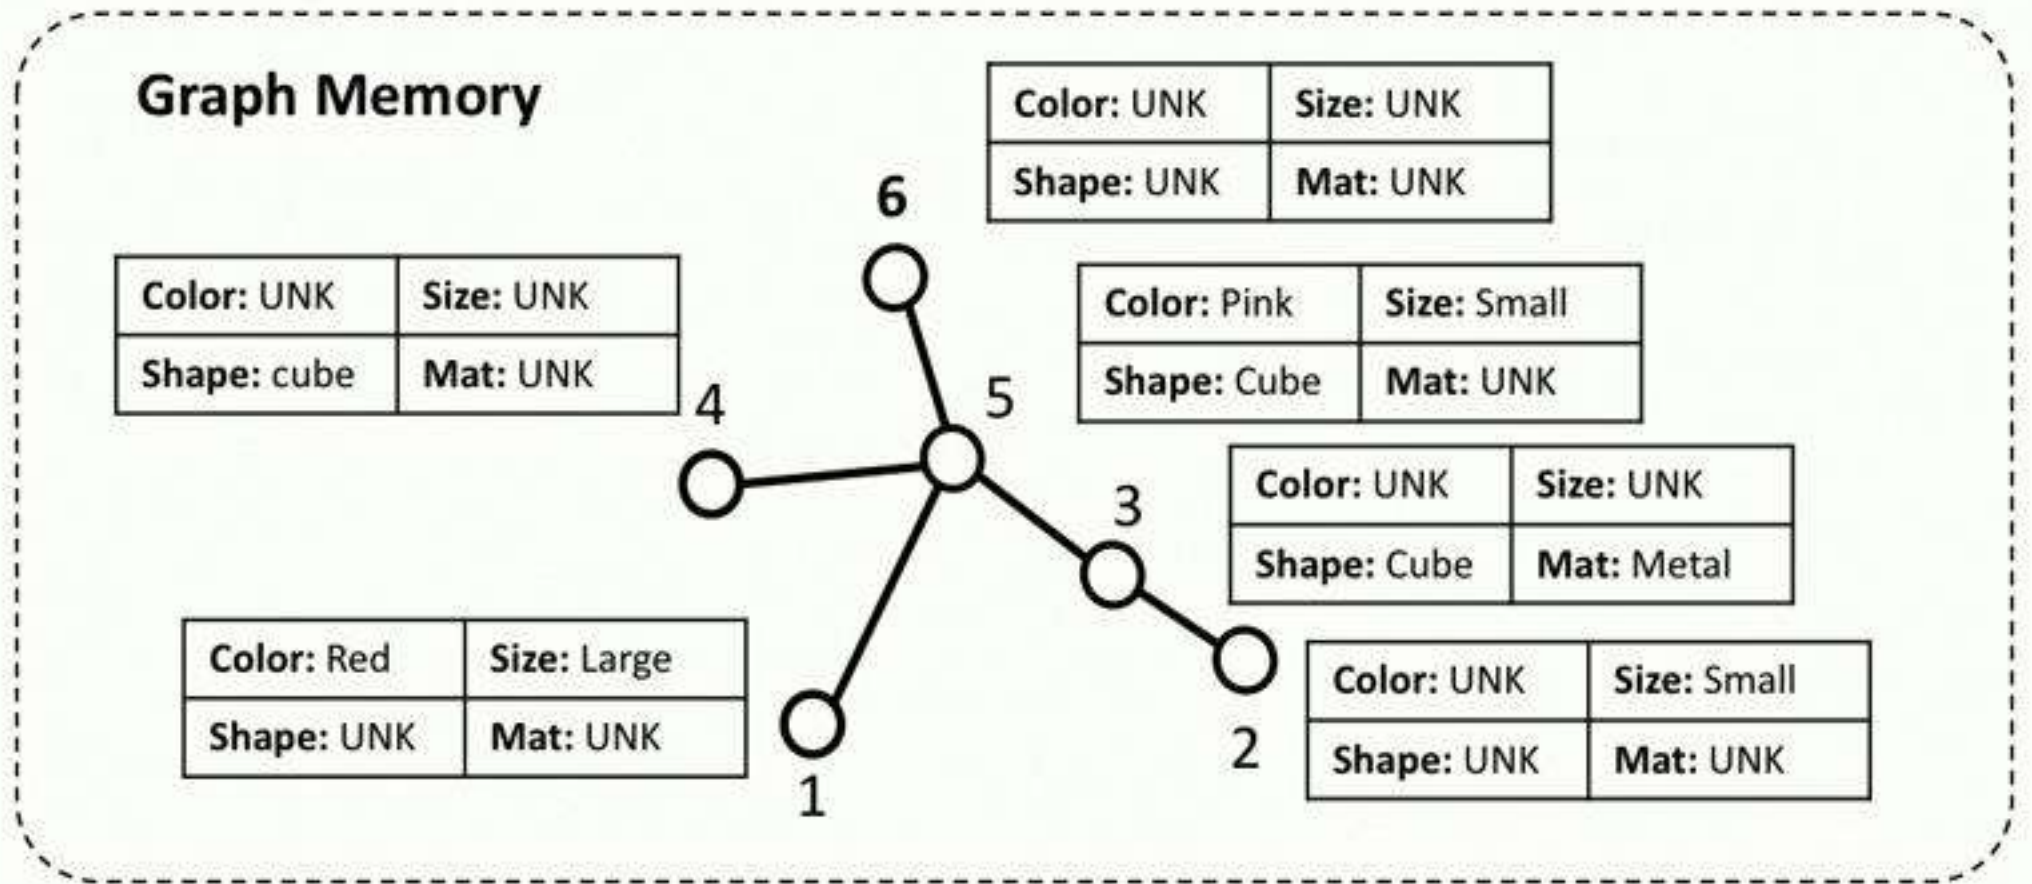
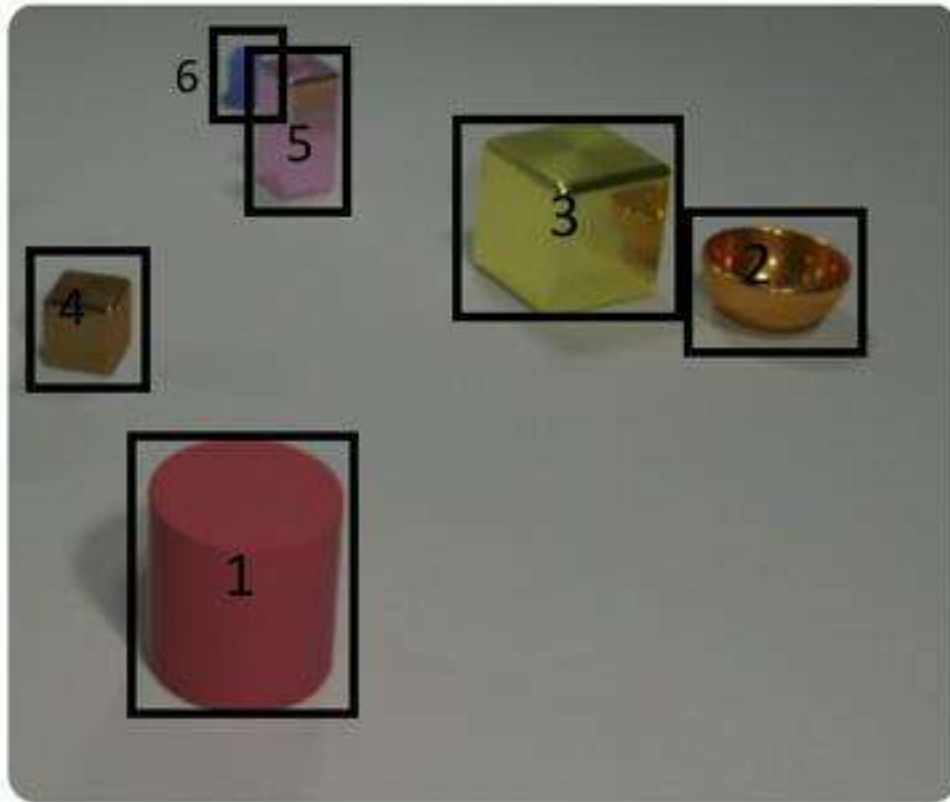
Agent Architecture



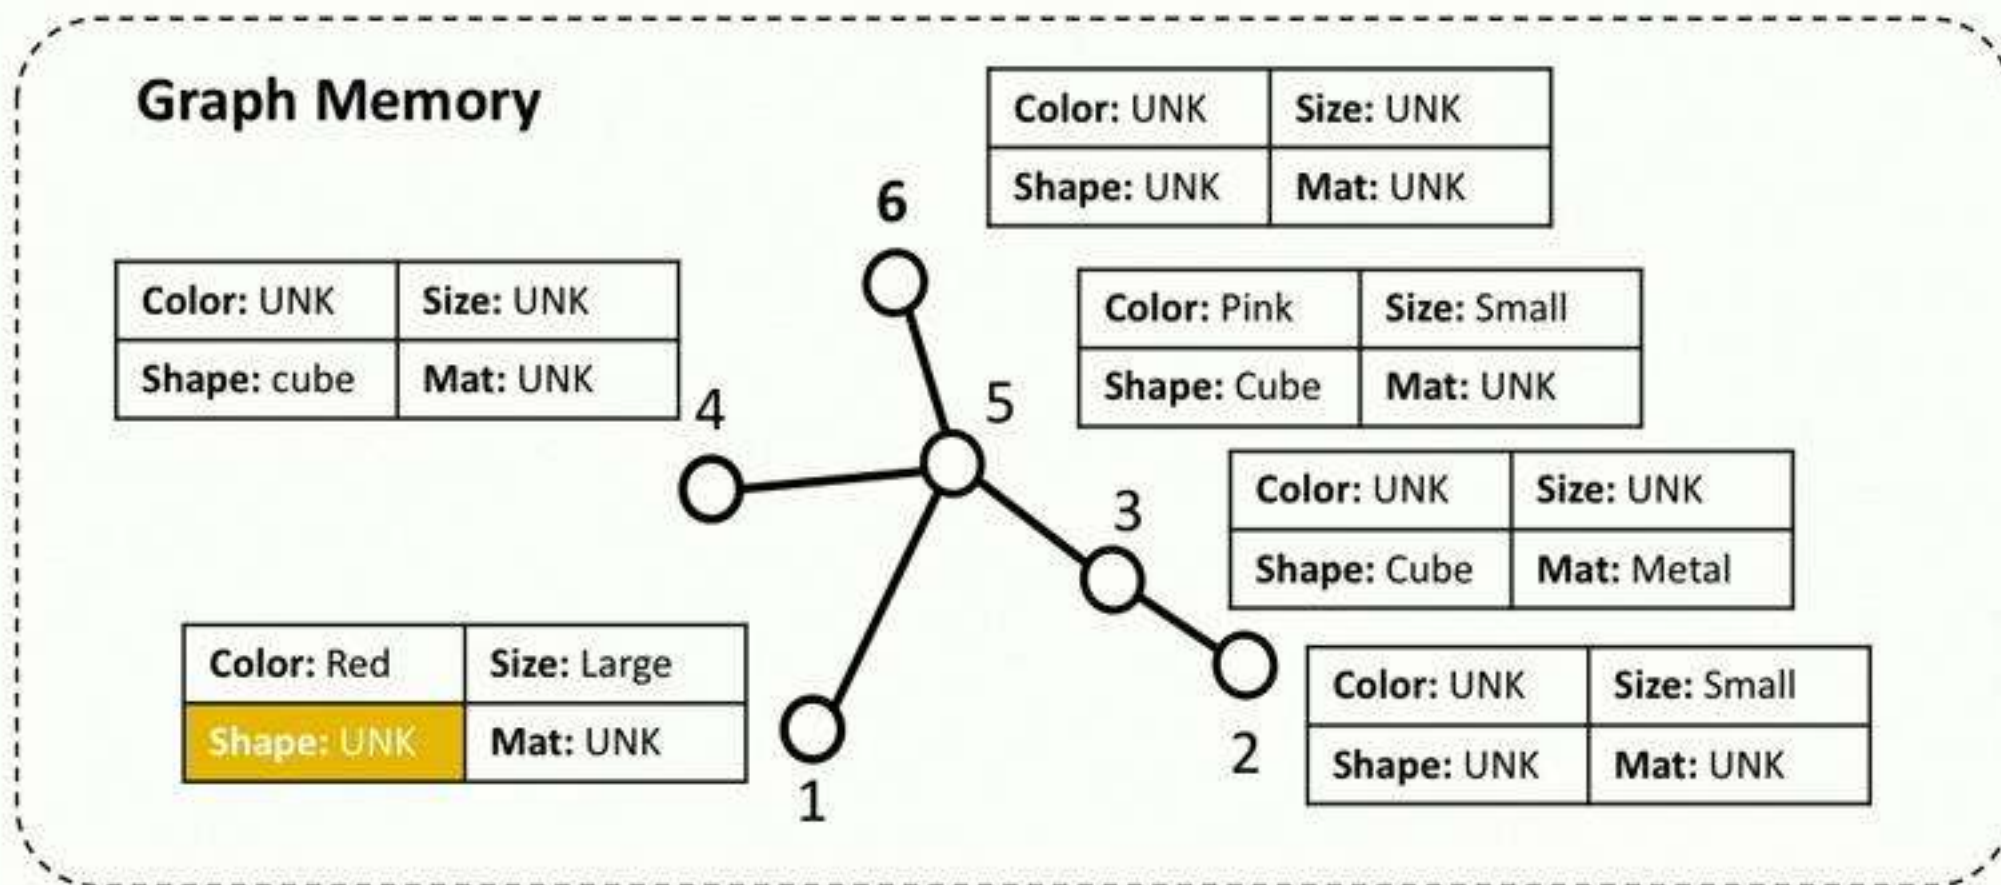
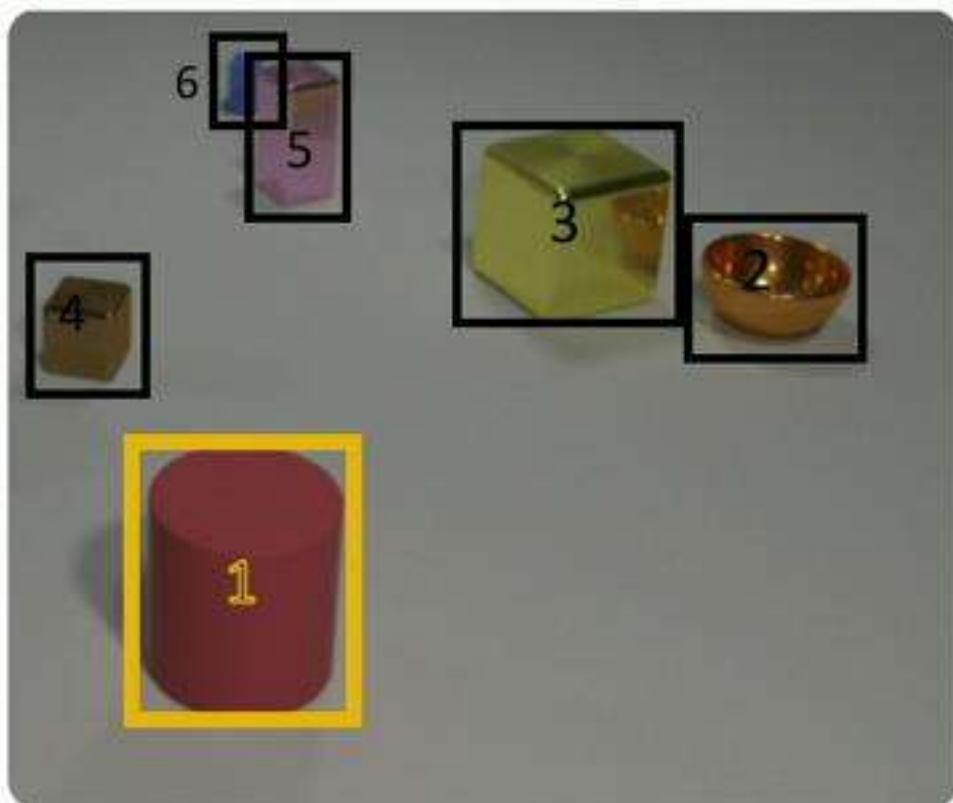
Agent Architecture



Question Template



Question Template



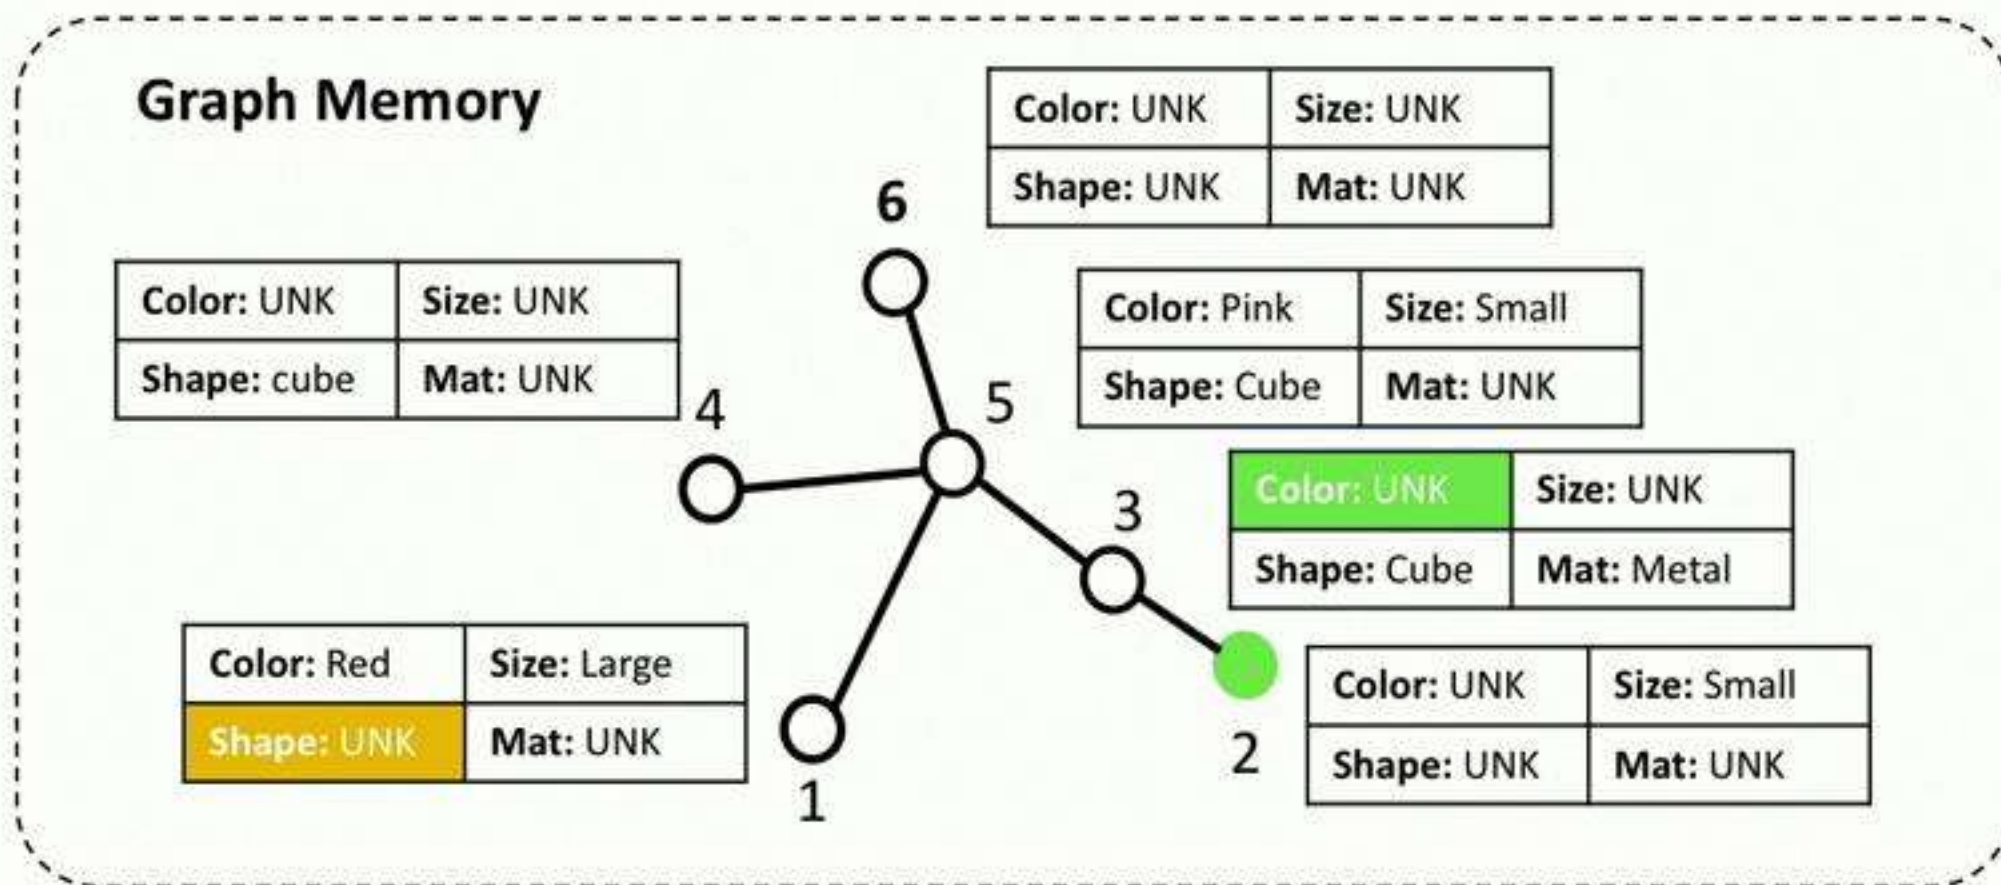
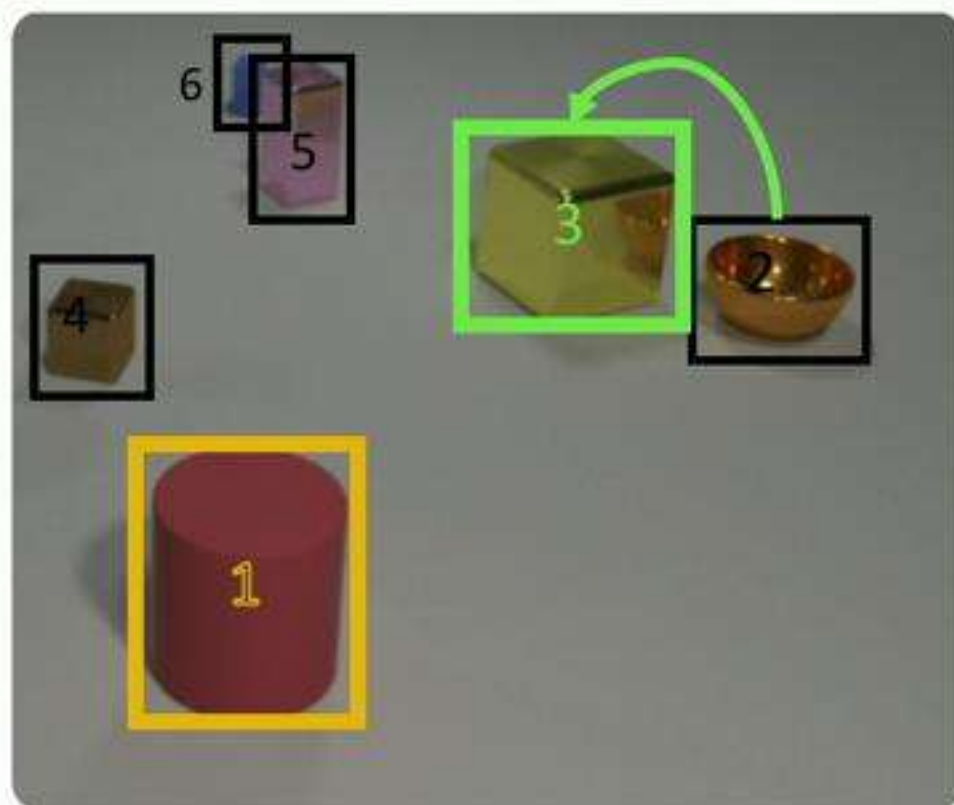
Target 1

Attribute Shape

Reference None

Question What is the shape of the front most large red object?

Question Template



Target 1

Attribute Shape

Reference None

Question What is the shape of the front most large red object?

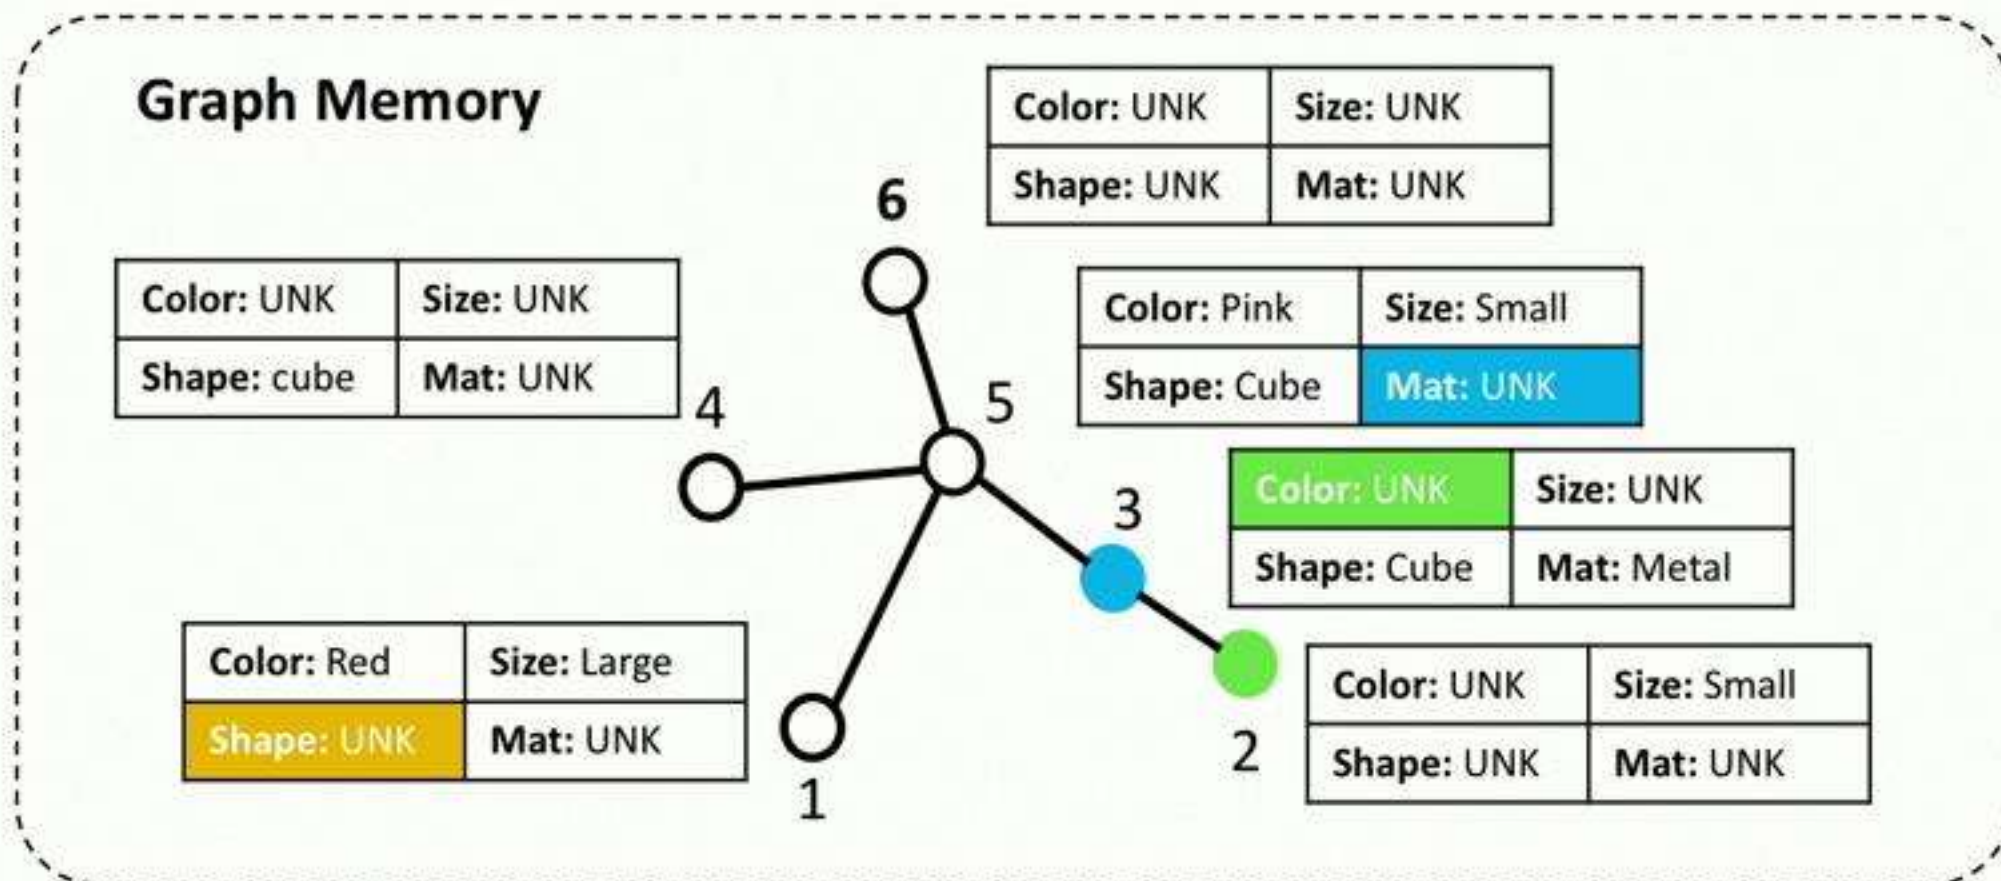
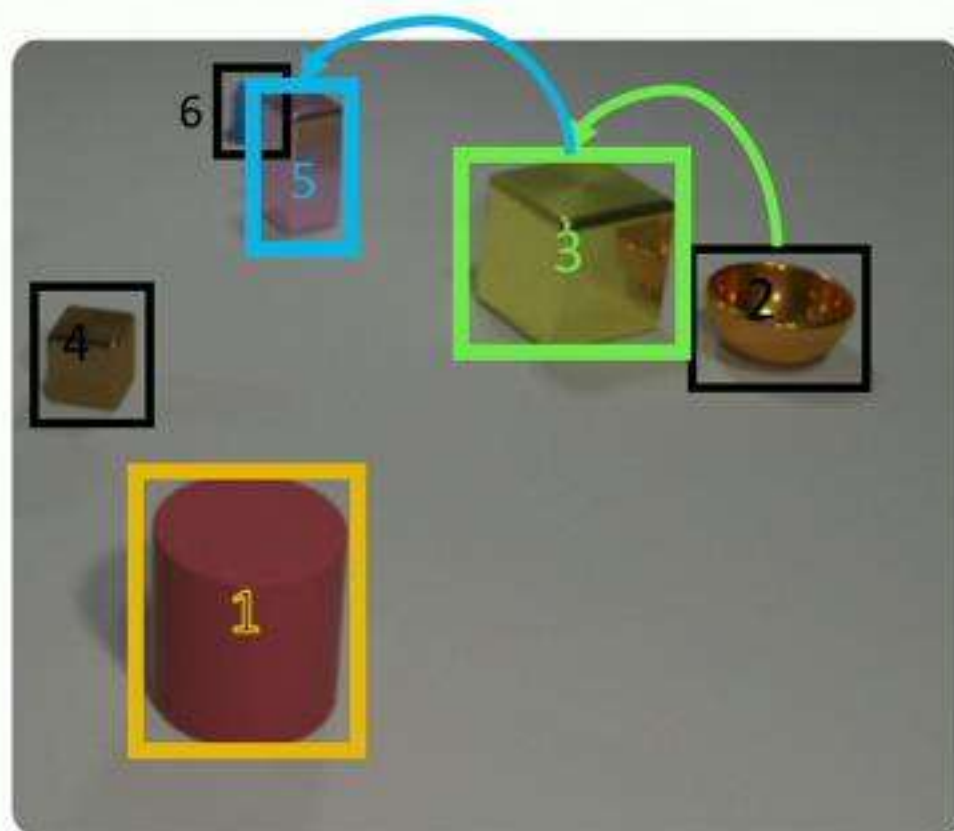
3

Color

2

What is the color of the metal cube on the left side of a small object?

Question Template



Target 1

Attribute Shape

Reference None

Question What is the shape of the front most large red object?

3

Color

2

What is the color of the metal cube on the left side of a small object?

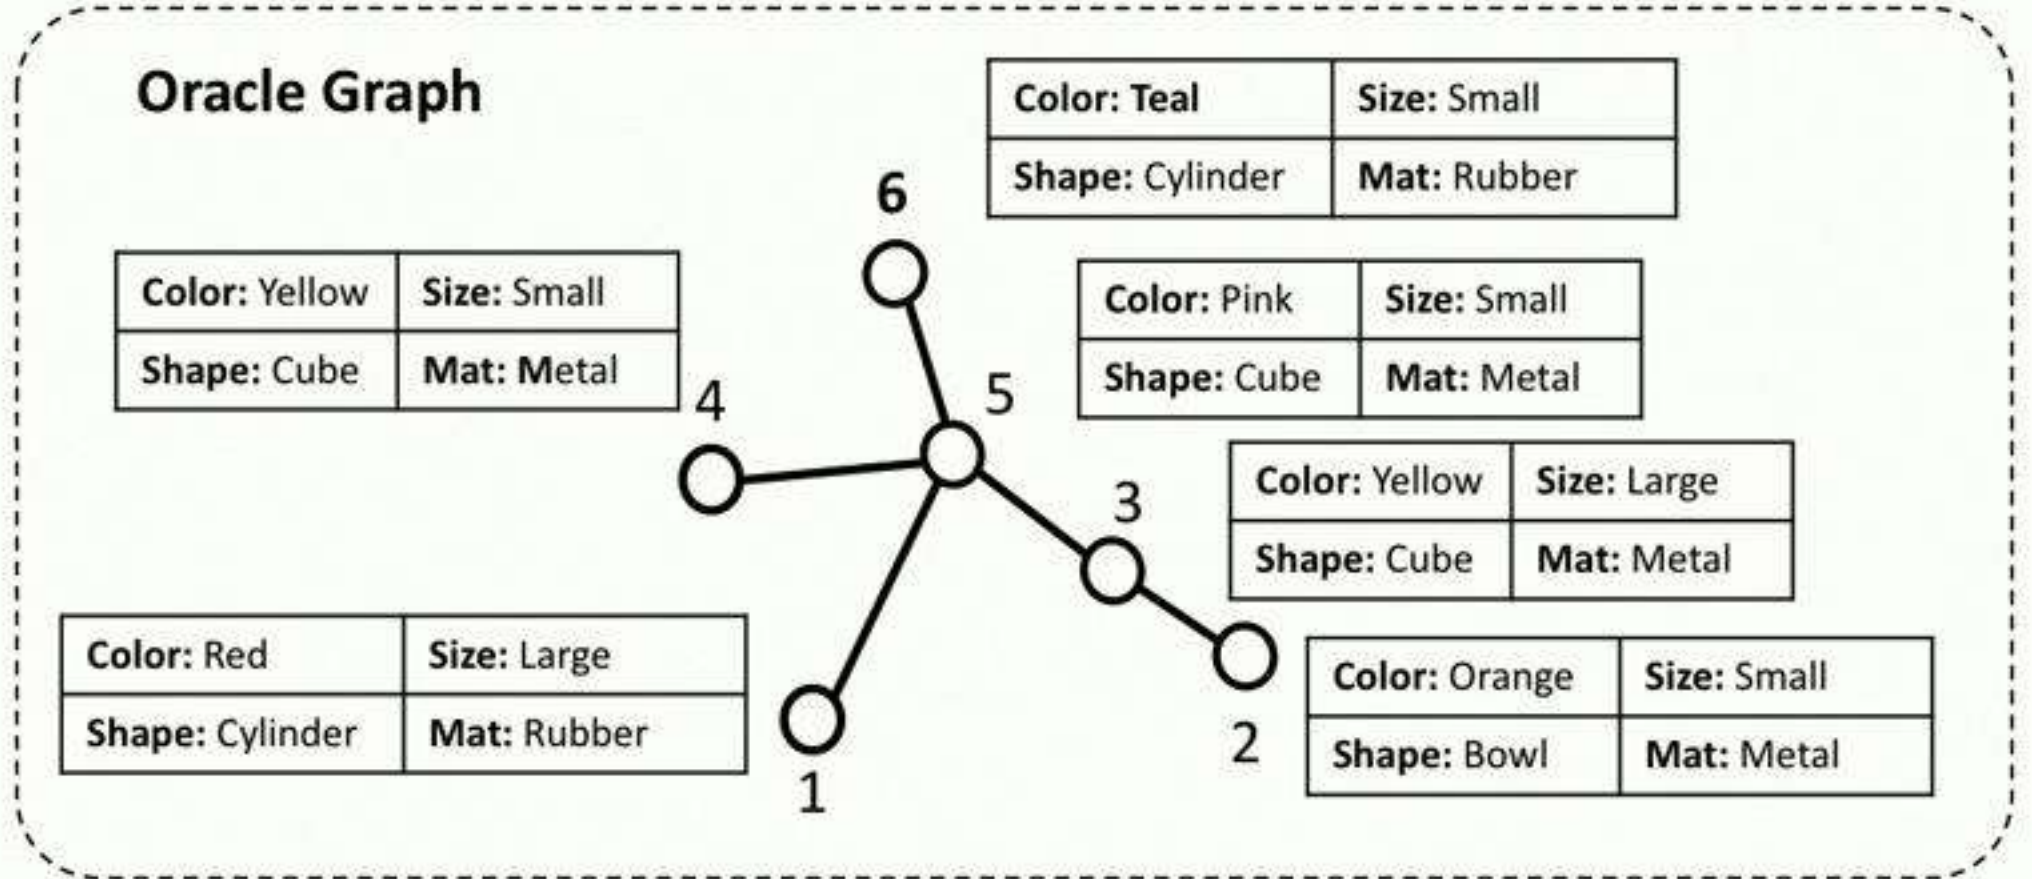
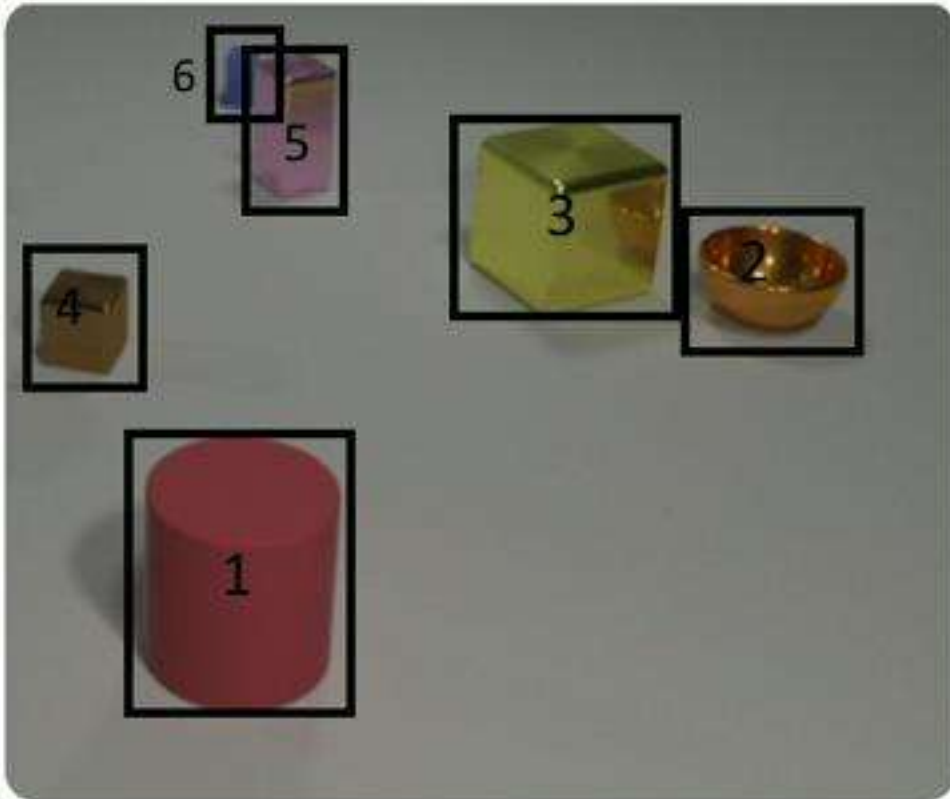
5

Material

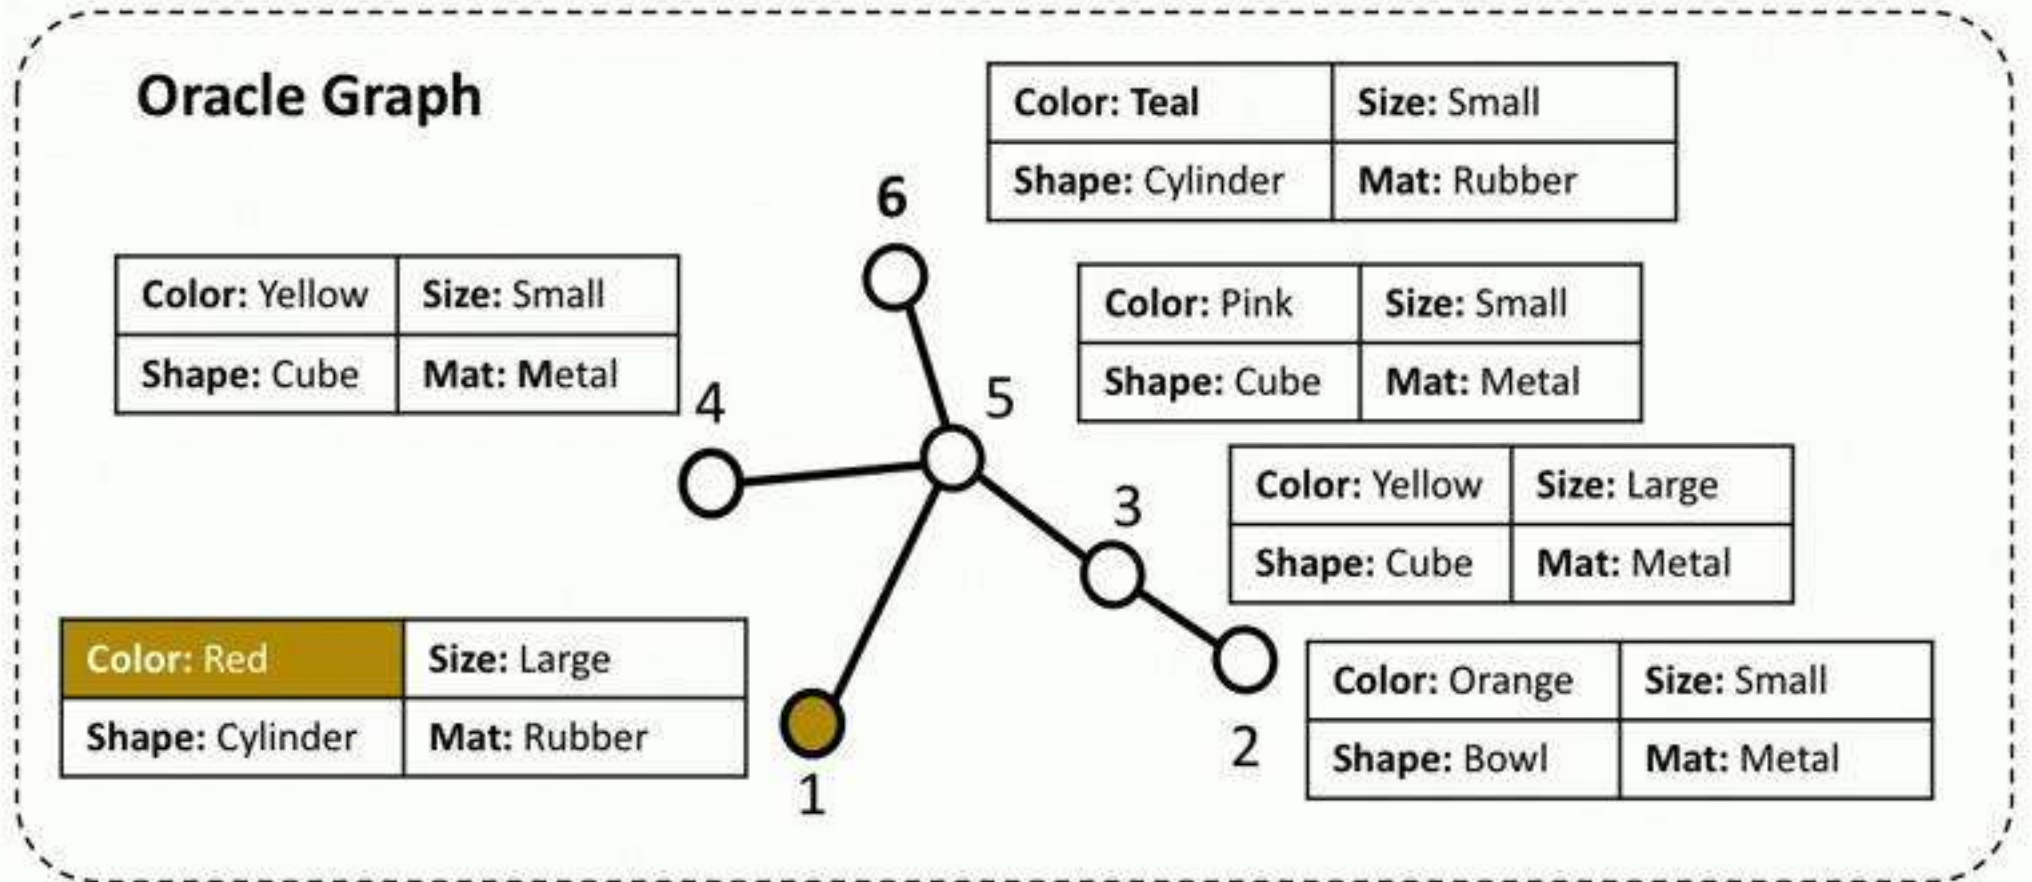
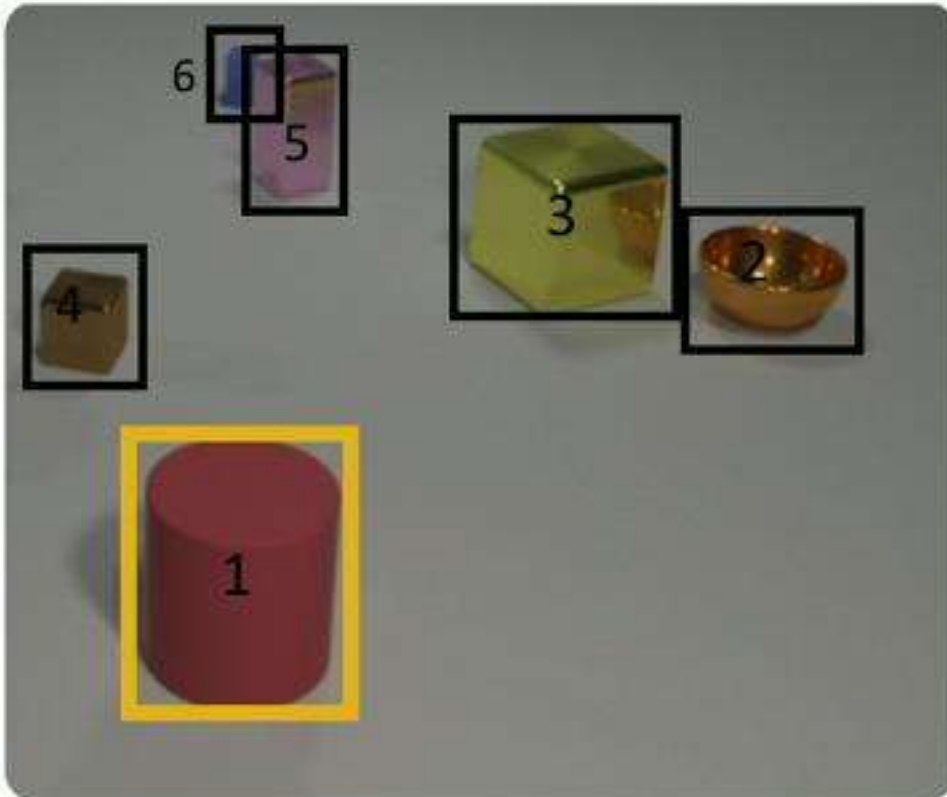
3

What is the material of object at left side of metal cube?

Oracle Question Answering

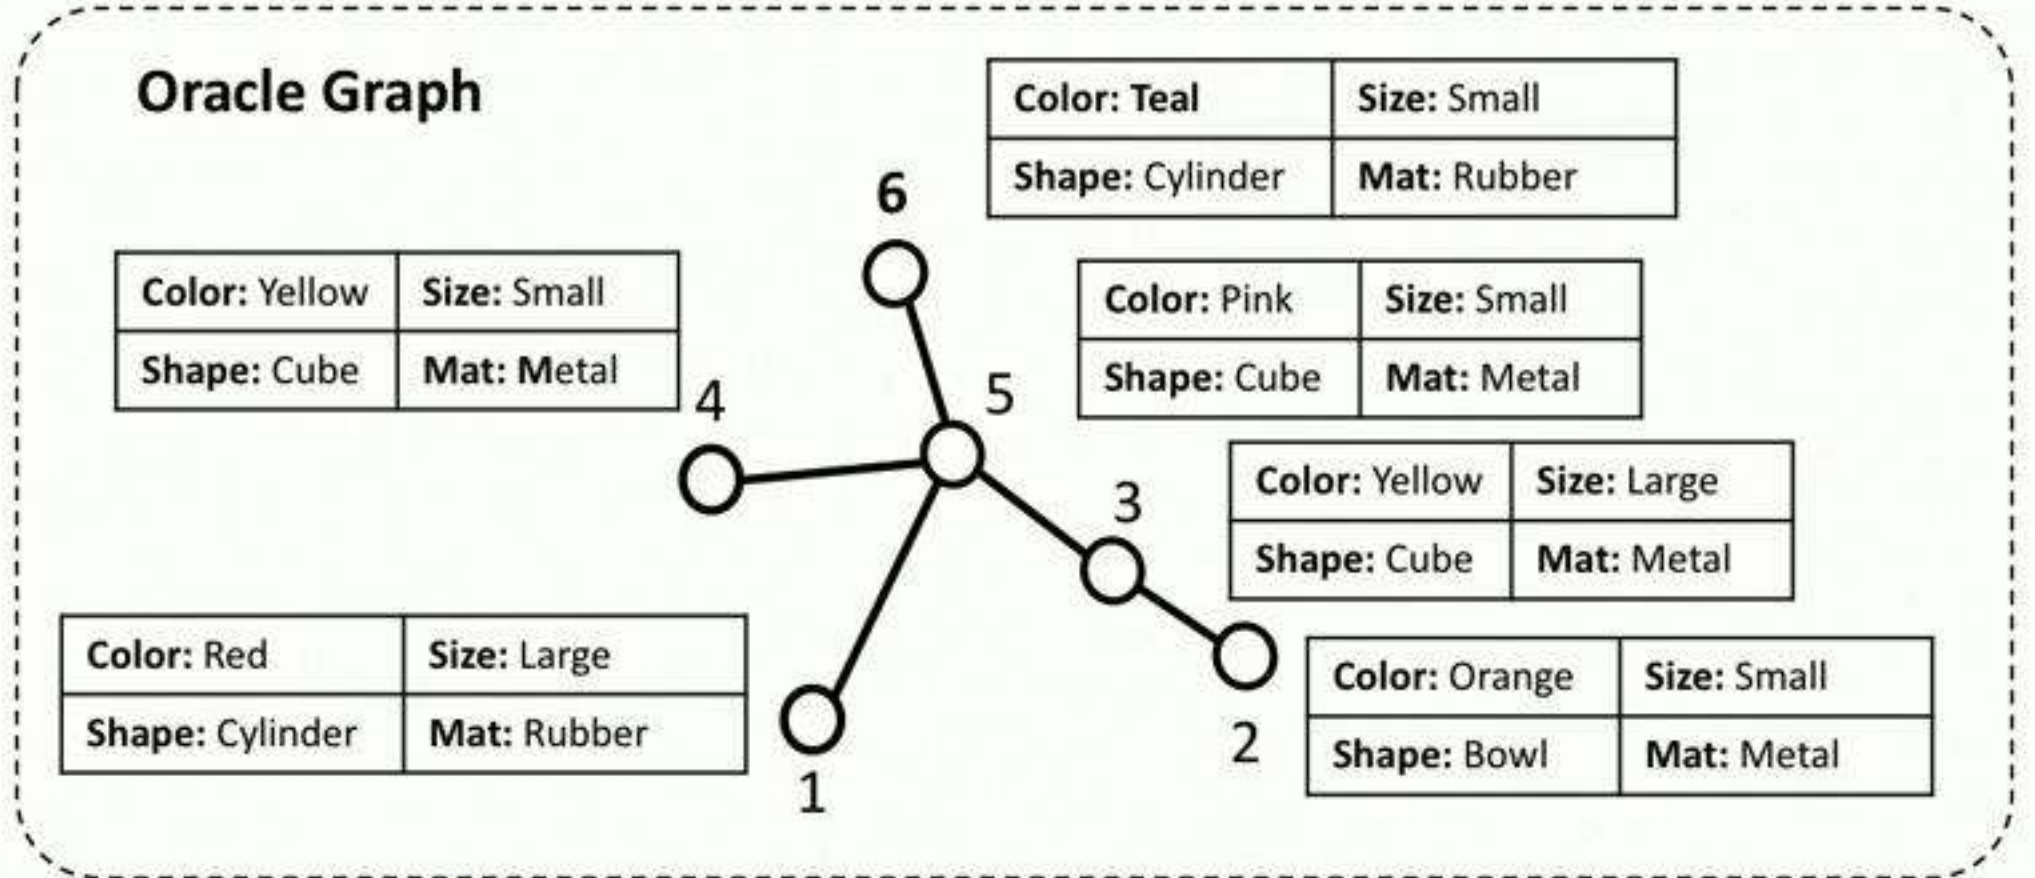
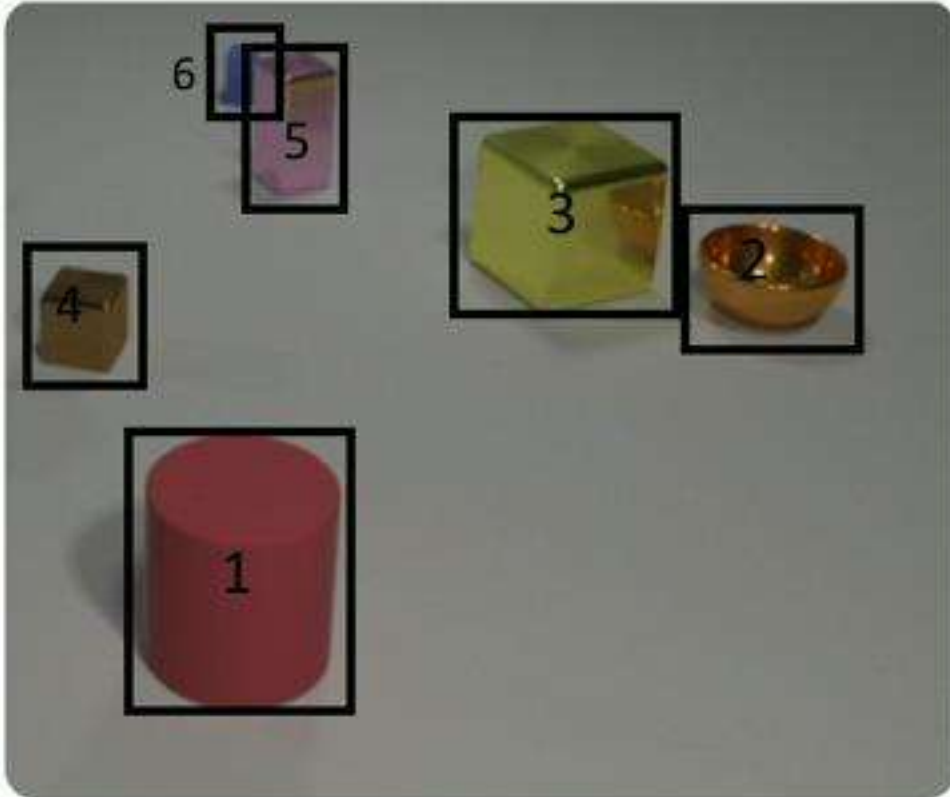


Oracle Question Answering



Q: What is the color of the front most object?

Oracle Question Answering

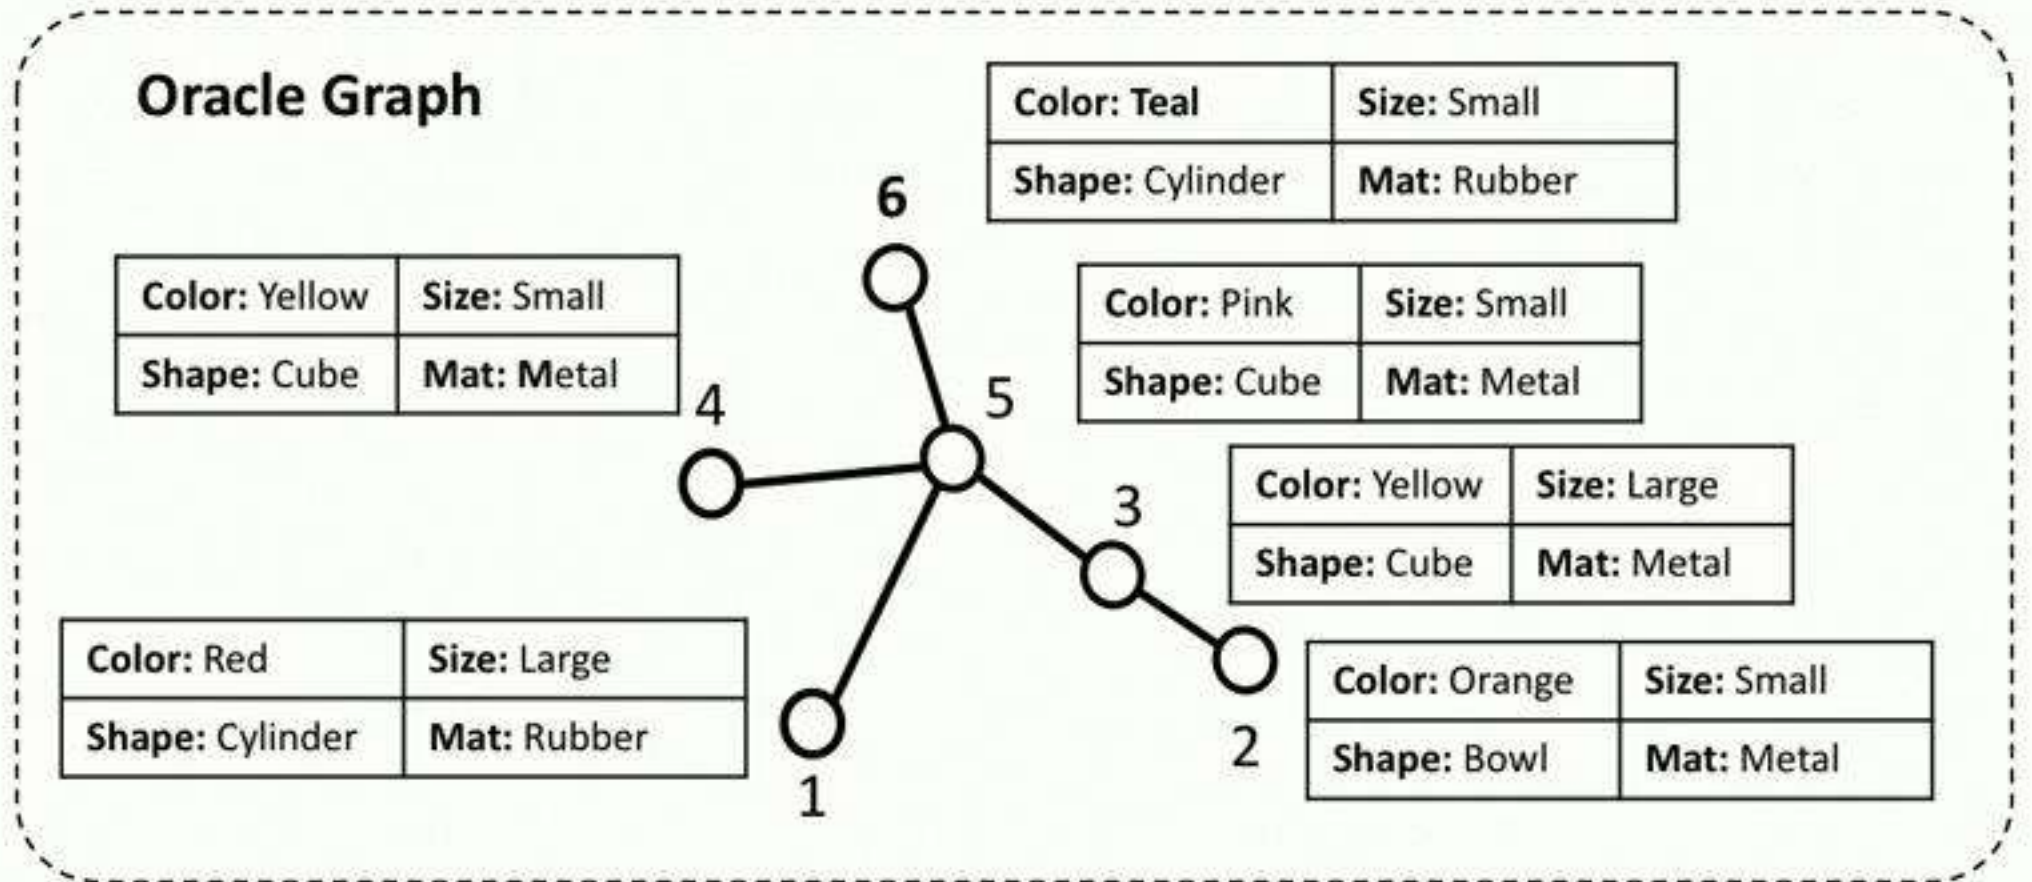
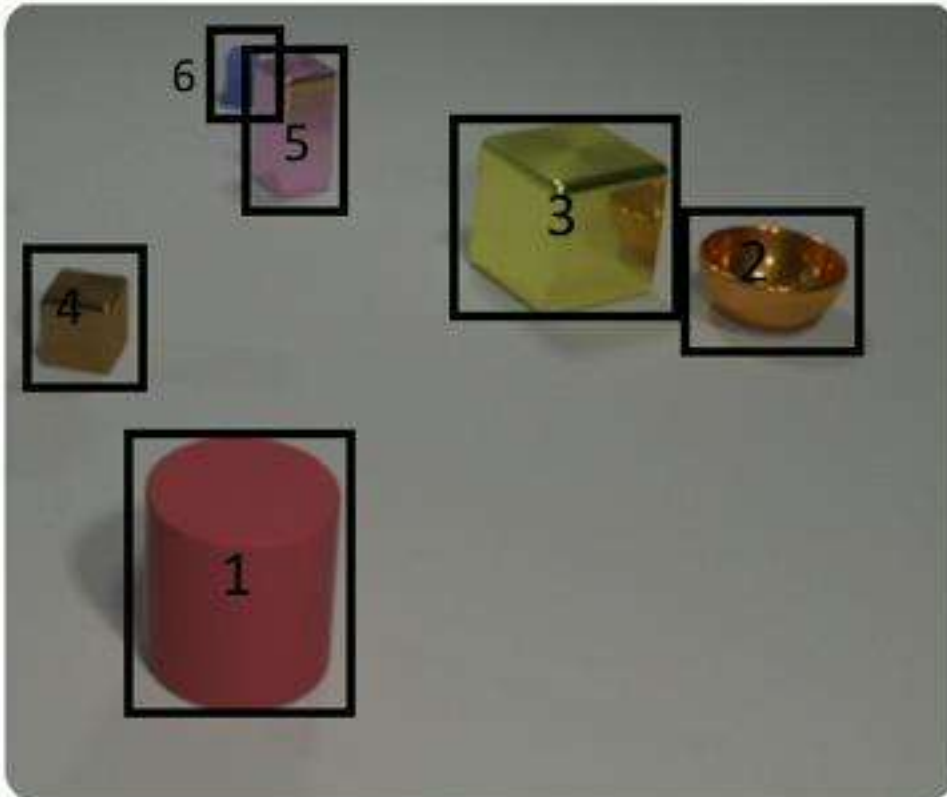


Q: What is the color of the front most object?

A: Red

Q: What is the shape of the object left of the green object?

Oracle Question Answering



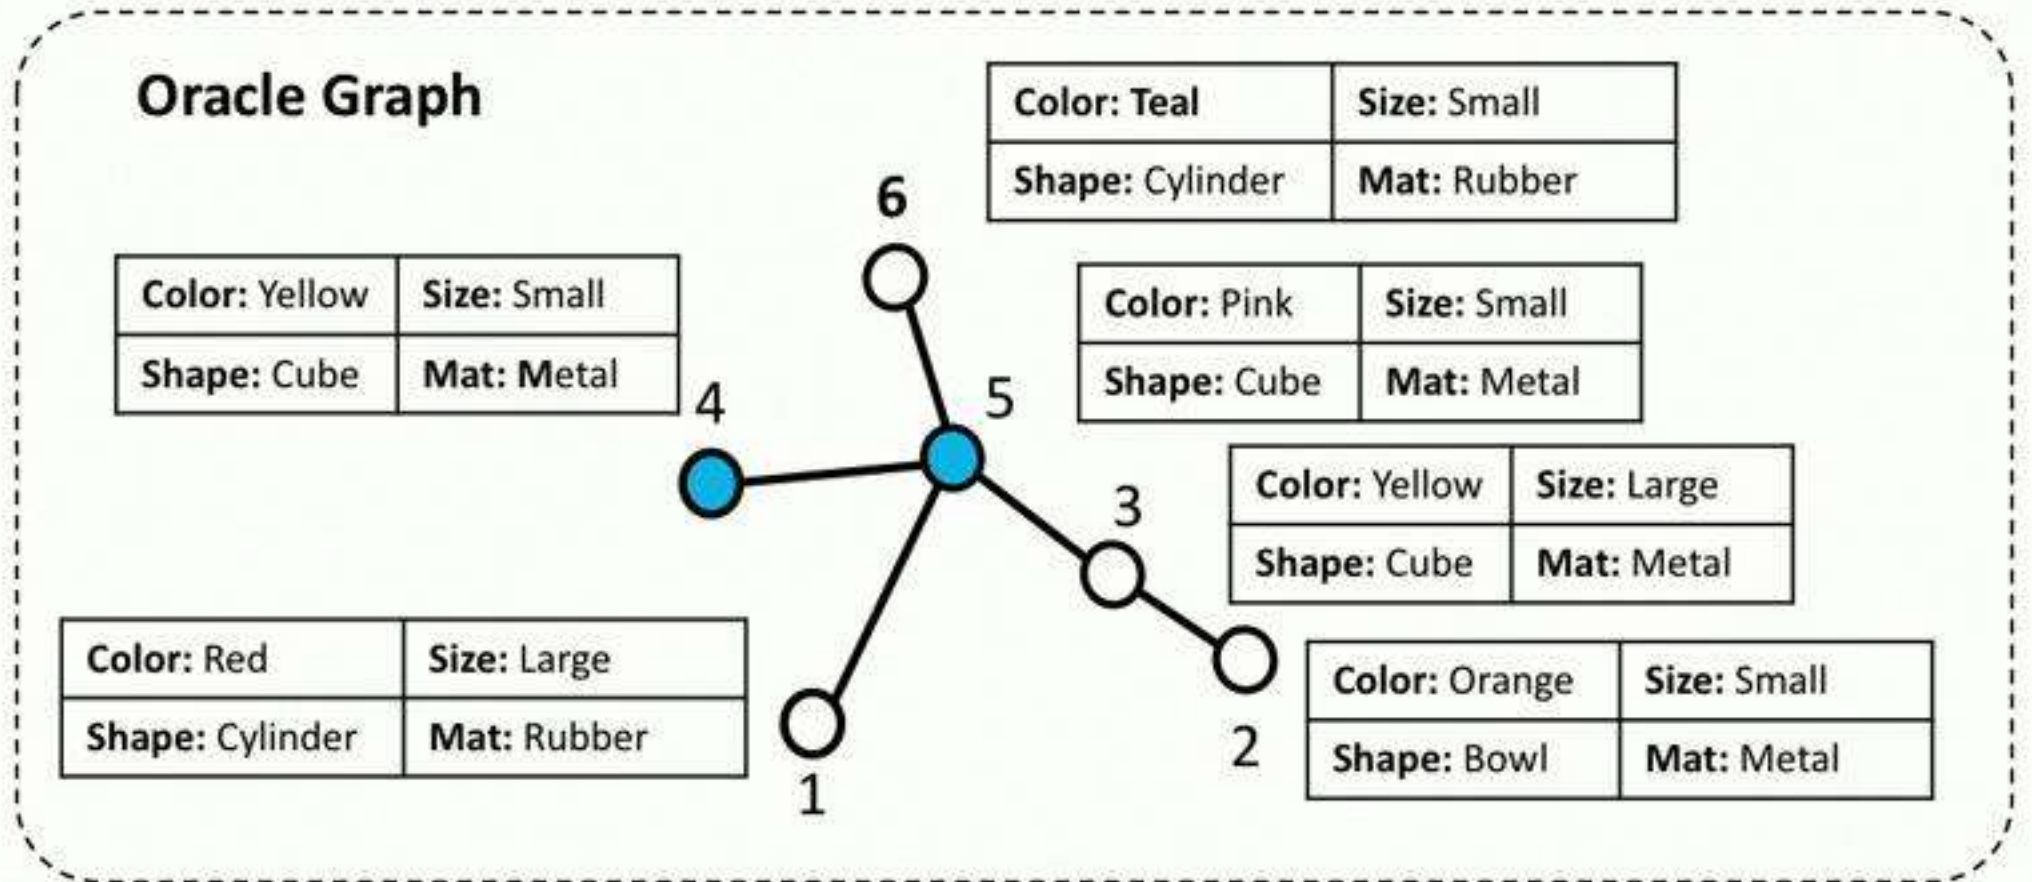
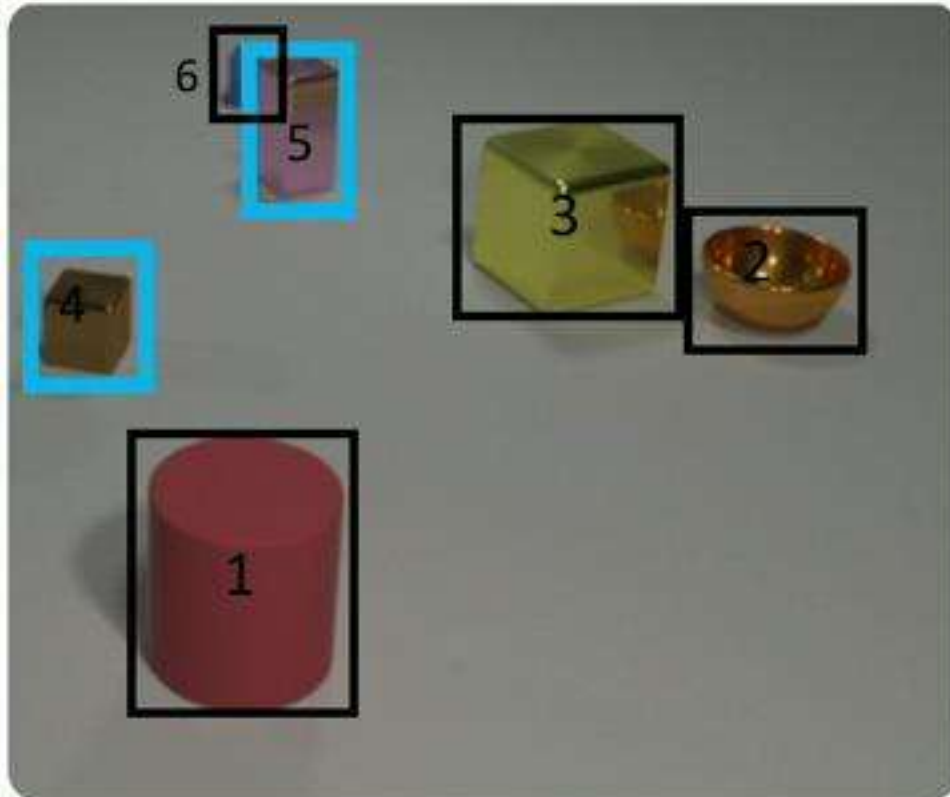
Q: What is the color of the front most object?

A: Red

Q: What is the shape of the object left of the green object?

A: < Invalid >

Oracle Question Answering



Q: What is the color of the front most object?

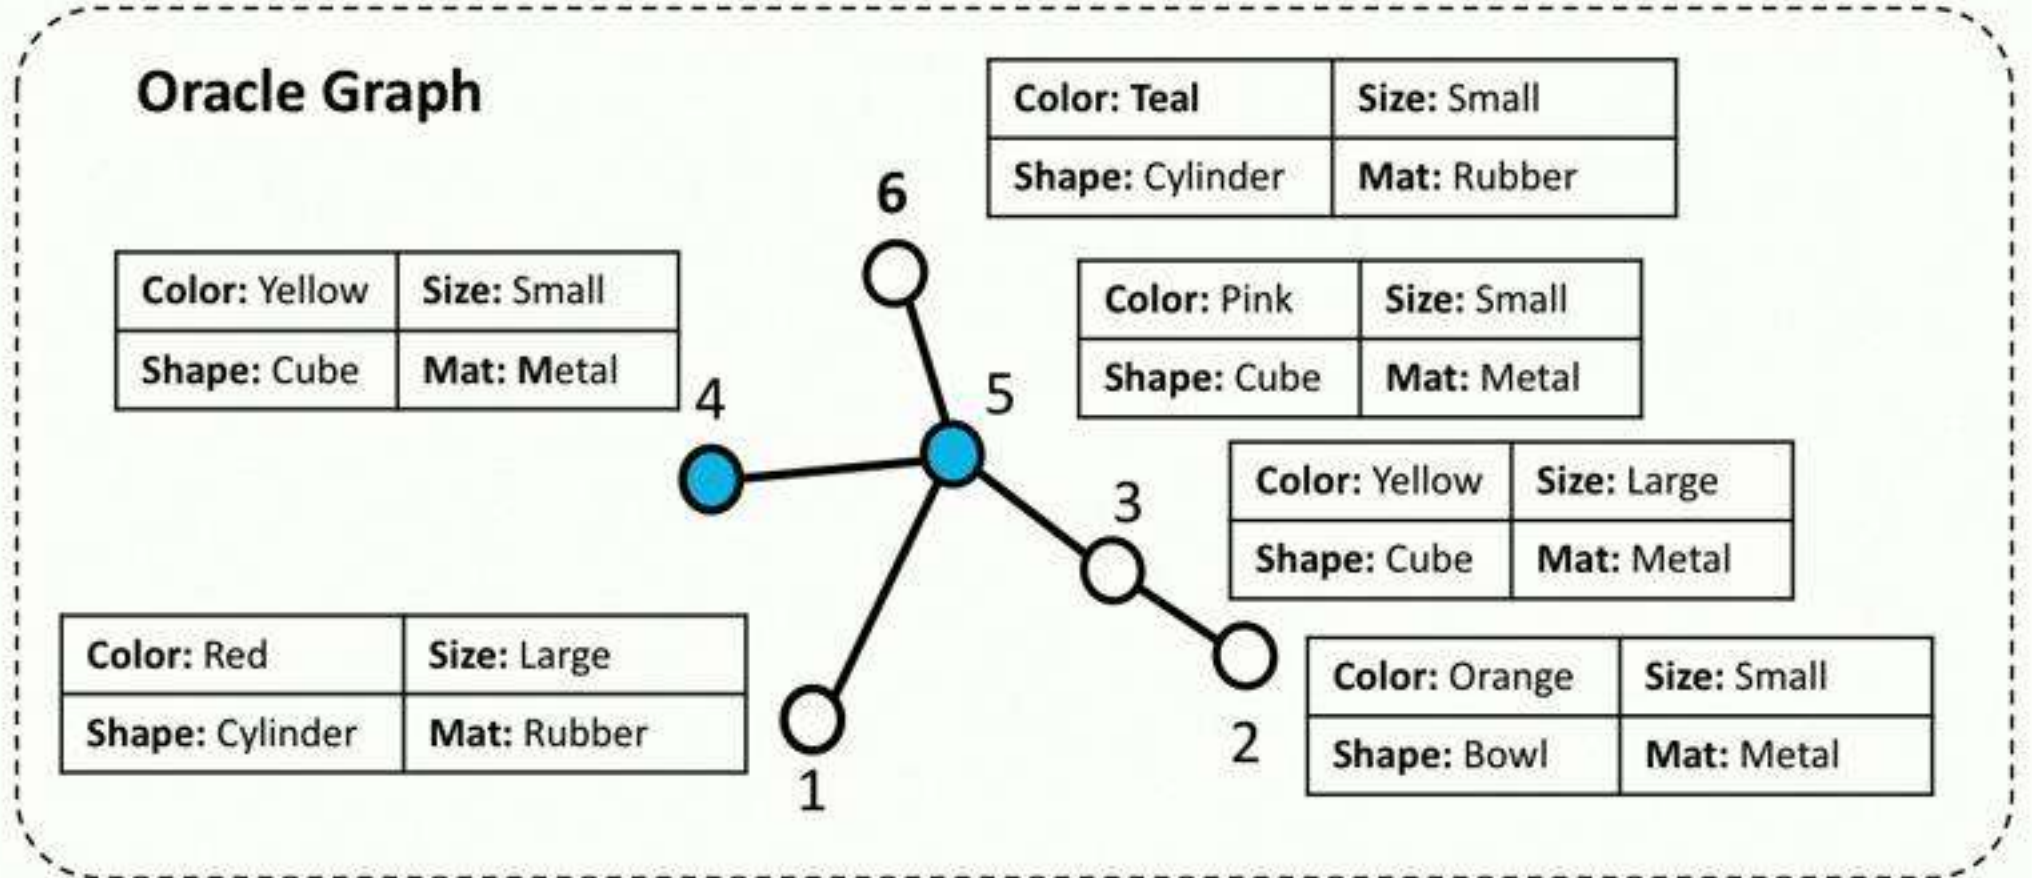
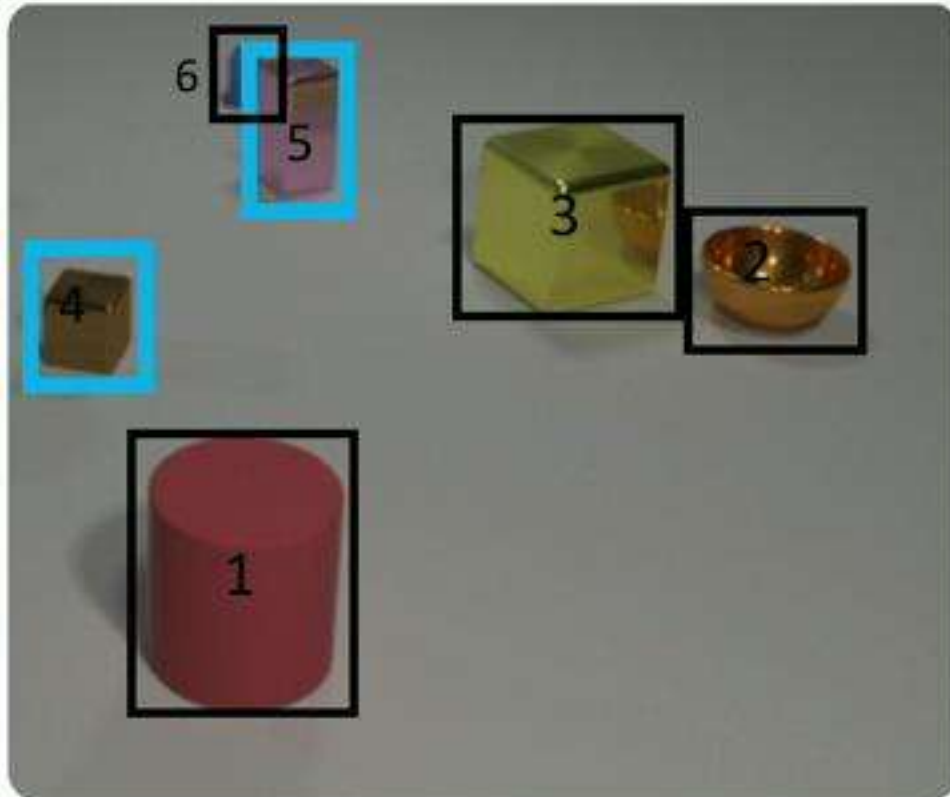
A: Red

Q: What is the shape of the object left of the green object?

A: < Invalid >

Q: What is color is the small cube?

Oracle Question Answering



Q: What is the color of the front most object?

A: Red

Q: What is the shape of the object left of the green object?

A: < Invalid >

Q: What is color is the small cube?

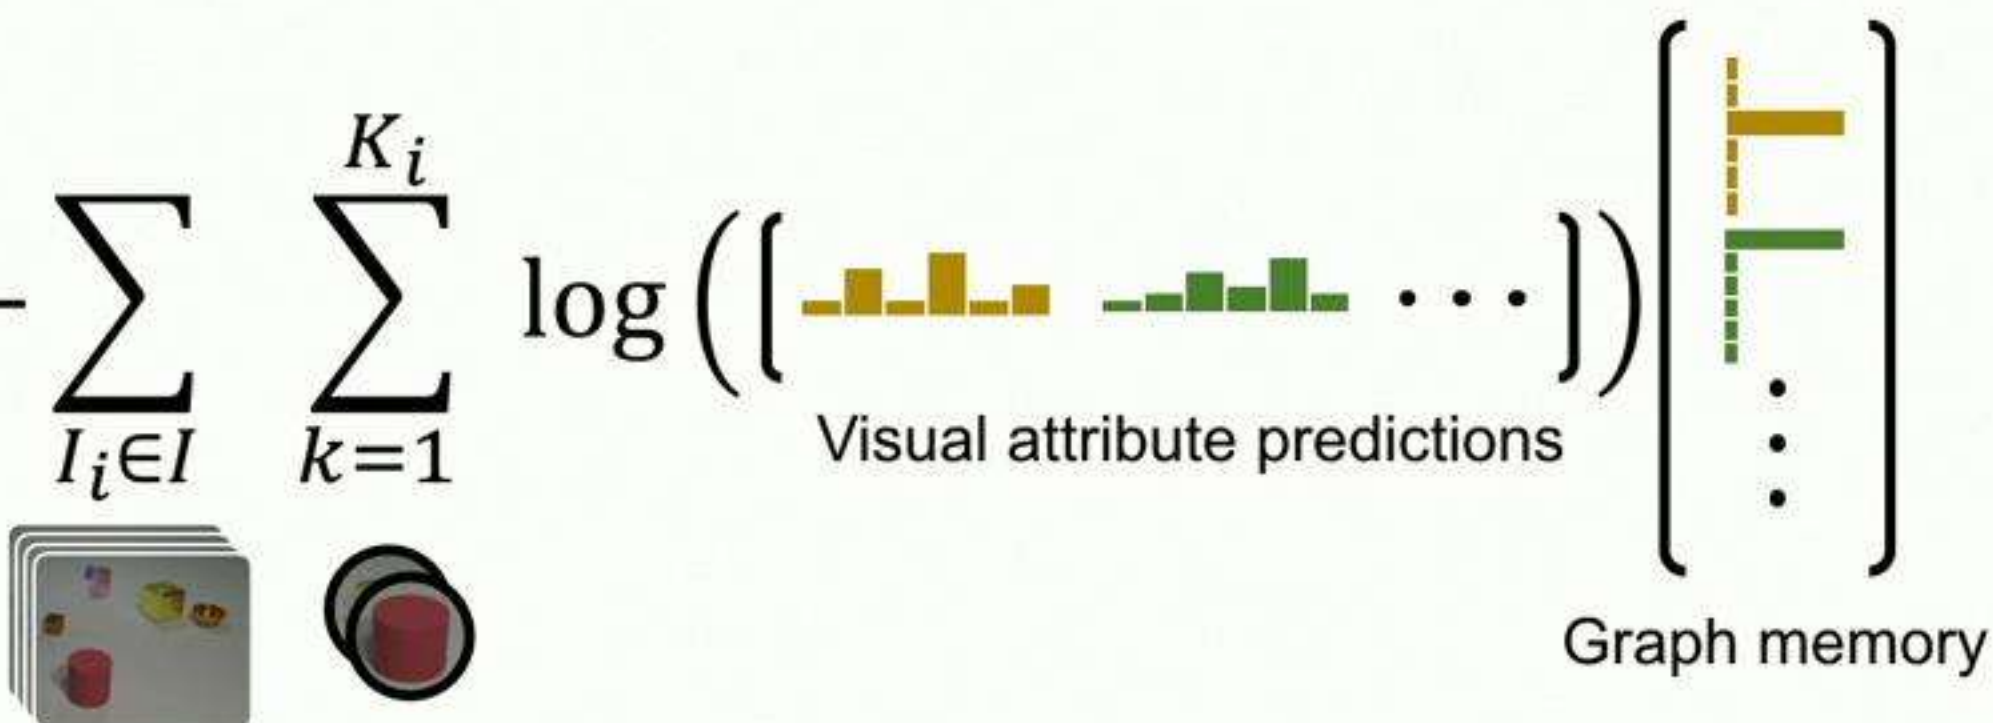
A: < Ambiguous >

Training Objective: Visual System

Cross-Entropy loss between the graph memory and the visual predictions over all images, objects, and attributes

$$\theta_v^* = \arg \min - \sum_{I_i \in I} \sum_{k=1}^{K_i} \log \left(\left(\begin{array}{c} \text{Visual attribute predictions} \end{array} \right) \left(\begin{array}{c} \text{Graph memory} \end{array} \right) \right)$$

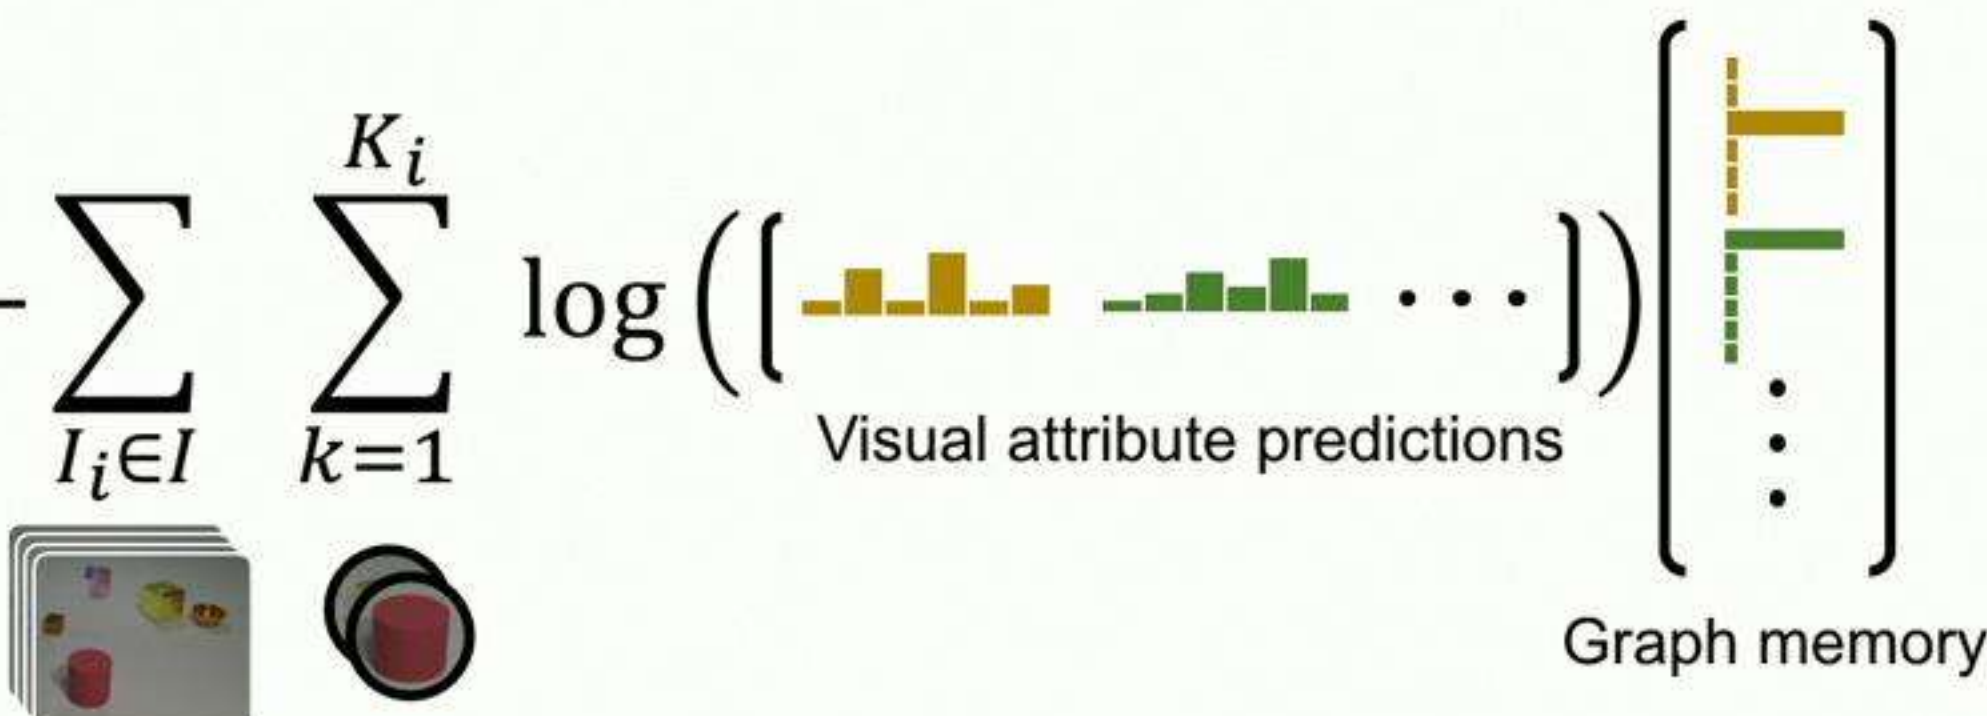
The equation illustrates the training objective for the visual system. It shows the minimization of the negative log-likelihood (Cross-Entropy loss) over all images I_i and their corresponding objects K_i . The visual attribute predictions are represented by a vector of bars (yellow and green), and the graph memory is represented by a vector of bars (yellow and green) with a vertical dashed line indicating the ground truth.



The diagram shows two sets of visual representations. On the left, under the summation, there is a stack of images showing various objects (a red cup, a yellow cup, a green cup, etc.) and a single image of a red cup with a black outline. On the right, under the log function, there are two vectors. The first vector, labeled 'Visual attribute predictions', consists of a series of yellow and green bars of varying heights. The second vector, labeled 'Graph memory', consists of a series of yellow and green bars, with a vertical dashed line indicating the ground truth.

Training Objective: Visual System

Cross-Entropy loss between the graph memory and the visual predictions over all images, objects, and attributes

$$\theta_v^* = \arg \min - \sum_{\substack{I_i \in I \\ \text{stack of images}}} \sum_{k=1}^{K_i} \log \left(\left(\begin{array}{c} \text{Visual attribute predictions} \\ \text{Visual attribute predictions} \end{array} \right) \begin{array}{c} \left(\begin{array}{c} \text{Graph memory} \\ \vdots \end{array} \right) \end{array} \right)$$


Update visual system after each dialog with stochastic gradient descent

Training Objective: Questioner Policy

Train the questioner to maximize expected reward over images and dialog rounds in an episode

$$\theta_{\pi}^* = \arg \max E_V E_{I \sim \mathcal{E}} E_{\pi_q} \left[\sum_{i=1}^n \sum_{t=1}^T r_i^t(q_i^t \sim \pi_q(h_i^t; \theta_{\pi})) \right]$$

Training Objective: Questioner Policy

Train the questioner to maximize expected reward over images and dialog rounds in an episode

Sum of reward over all images and dialog rounds

$$\theta_{\pi}^* = \arg \max \underbrace{E_V E_{I \sim \mathcal{E}} E_{\pi_q}}_{\text{Expectation over visual systems, episodes, and questions}} \left[\sum_{i=1}^n \sum_{t=1}^T r_i^t(q_i^t \sim \pi_q(h_i^t; \theta_{\pi})) \right]$$

Expectation over visual systems, episodes, and questions

Use A2C and update policy after each episode based on all rounds

Training Objective: Questioner Policy Reward

Per round change in graph memory accuracy

$$r_i^t = S(G_i^t, G_i^*) - S(G_i^{t-1}, G_i^*)$$

Training Objective: Questioner Policy

Train the questioner to maximize expected reward over images and dialog rounds in an episode

Sum of reward over all
images and dialog rounds

$$\theta_{\pi}^* = \arg \max E_V E_{I \sim \mathcal{E}} E_{\pi_q} \left[\sum_{i=1}^n \sum_{t=1}^T r_i^t(q_i^t \sim \pi_q(h_i^t; \theta_{\pi})) \right]$$

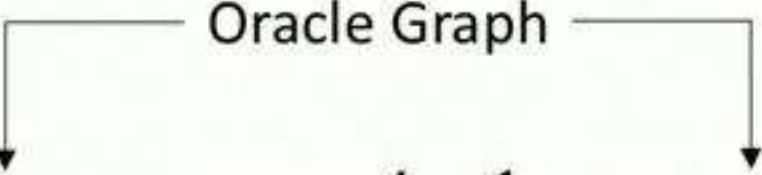
Training Objective: Questioner Policy Reward

Per round change in graph memory accuracy

$$r_i^t = S(G_i^t, G_i^*) - S(G_i^{t-1}, G_i^*)$$

Training Objective: Questioner Policy Reward

Per round change in graph memory accuracy



The diagram shows a horizontal line with the text "Oracle Graph" centered above it. Two vertical arrows point downwards from the ends of this line to the G_i^* term in both the first and second terms of the equation below.

$$r_i^t = S(G_i^t, G_i^*) - S(G_i^{t-1}, G_i^*)$$

Training Objective: Questioner Policy Reward

Per round change in graph memory accuracy

$$r_i^t = S(G_i^t, G_i^*) - S(G_i^{t-1}, G_i^*)$$

Graph Memory at dialog round t

Graph Memory at dialog round $t-1$

Oracle Graph

Training Objective: Questioner Policy Reward

Per round change in graph memory accuracy

$$r_i^t = S(G_i^t, G_i^*) - S(G_i^{t-1}, G_i^*)$$

Graph Memory at dialog round t

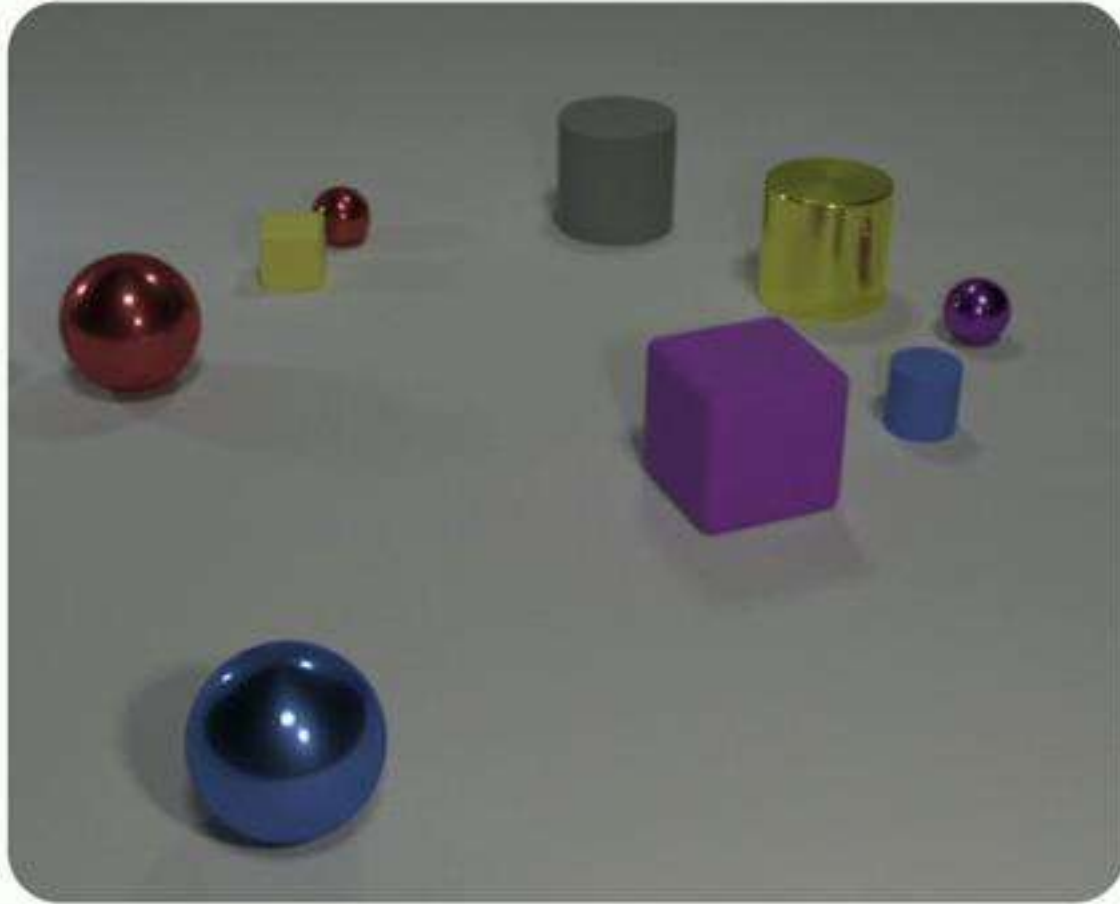
Graph Memory at dialog round $t-1$

Oracle Graph

Can be improved by:

- Asking unambiguous, informative questions (top-down)
- Improving the visual system quickly (bottom-up)

Experiments: Environments



Synthesized Dataset

Different shapes, colors, materials and sizes. Extended from CLEVR dataset [1]

[1] CLEVR. Johnson et al.



Realistic Dataset

Various real indoor scenes.
Annotated based on the ARID dataset [2]

[2] Recognizing Objects In-the-Wild. Loghmani et al.

Experiments: Baselines

Baseline Heuristic Questioners:

Target object

Target attribute

Reference object

Experiments: Baselines

Baseline Heuristic Questioners:

	Random
Target object	Uniform
Target attribute	Uniform
Reference object	Uniform

Experiments: Baselines

Baseline Heuristic Questioners:

	<u>Random</u>	<u>Entropy</u>
Target object	Uniform	Highest entropy
Target attribute	Uniform	Highest entropy
Reference object	Uniform	Lowest entropy

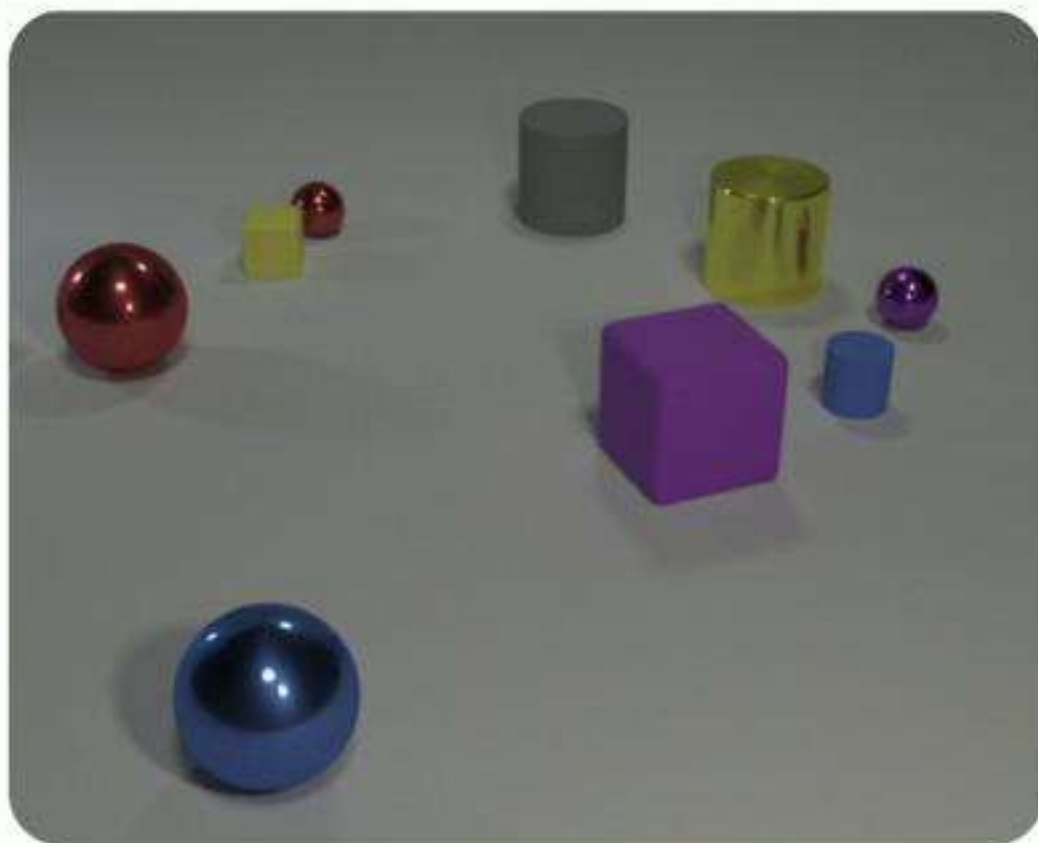
Experiments: Baselines

Baseline Heuristic Questioners:

	<u>Random</u>	<u>Entropy</u>	<u>Entropy + Context</u>
Target object	Uniform	Highest entropy	Highest Entropy + Spatial
Target attribute	Uniform	Highest entropy	Highest Entropy
Reference object	Uniform	Lowest entropy	Lowest Entropy + Spatial

Experiment: Standard Training

Standard Dataset



Shapes: cube, sphere, cylinder

Colors: gray, red, blue, green, yellow, purple

Materials: rubber, metal

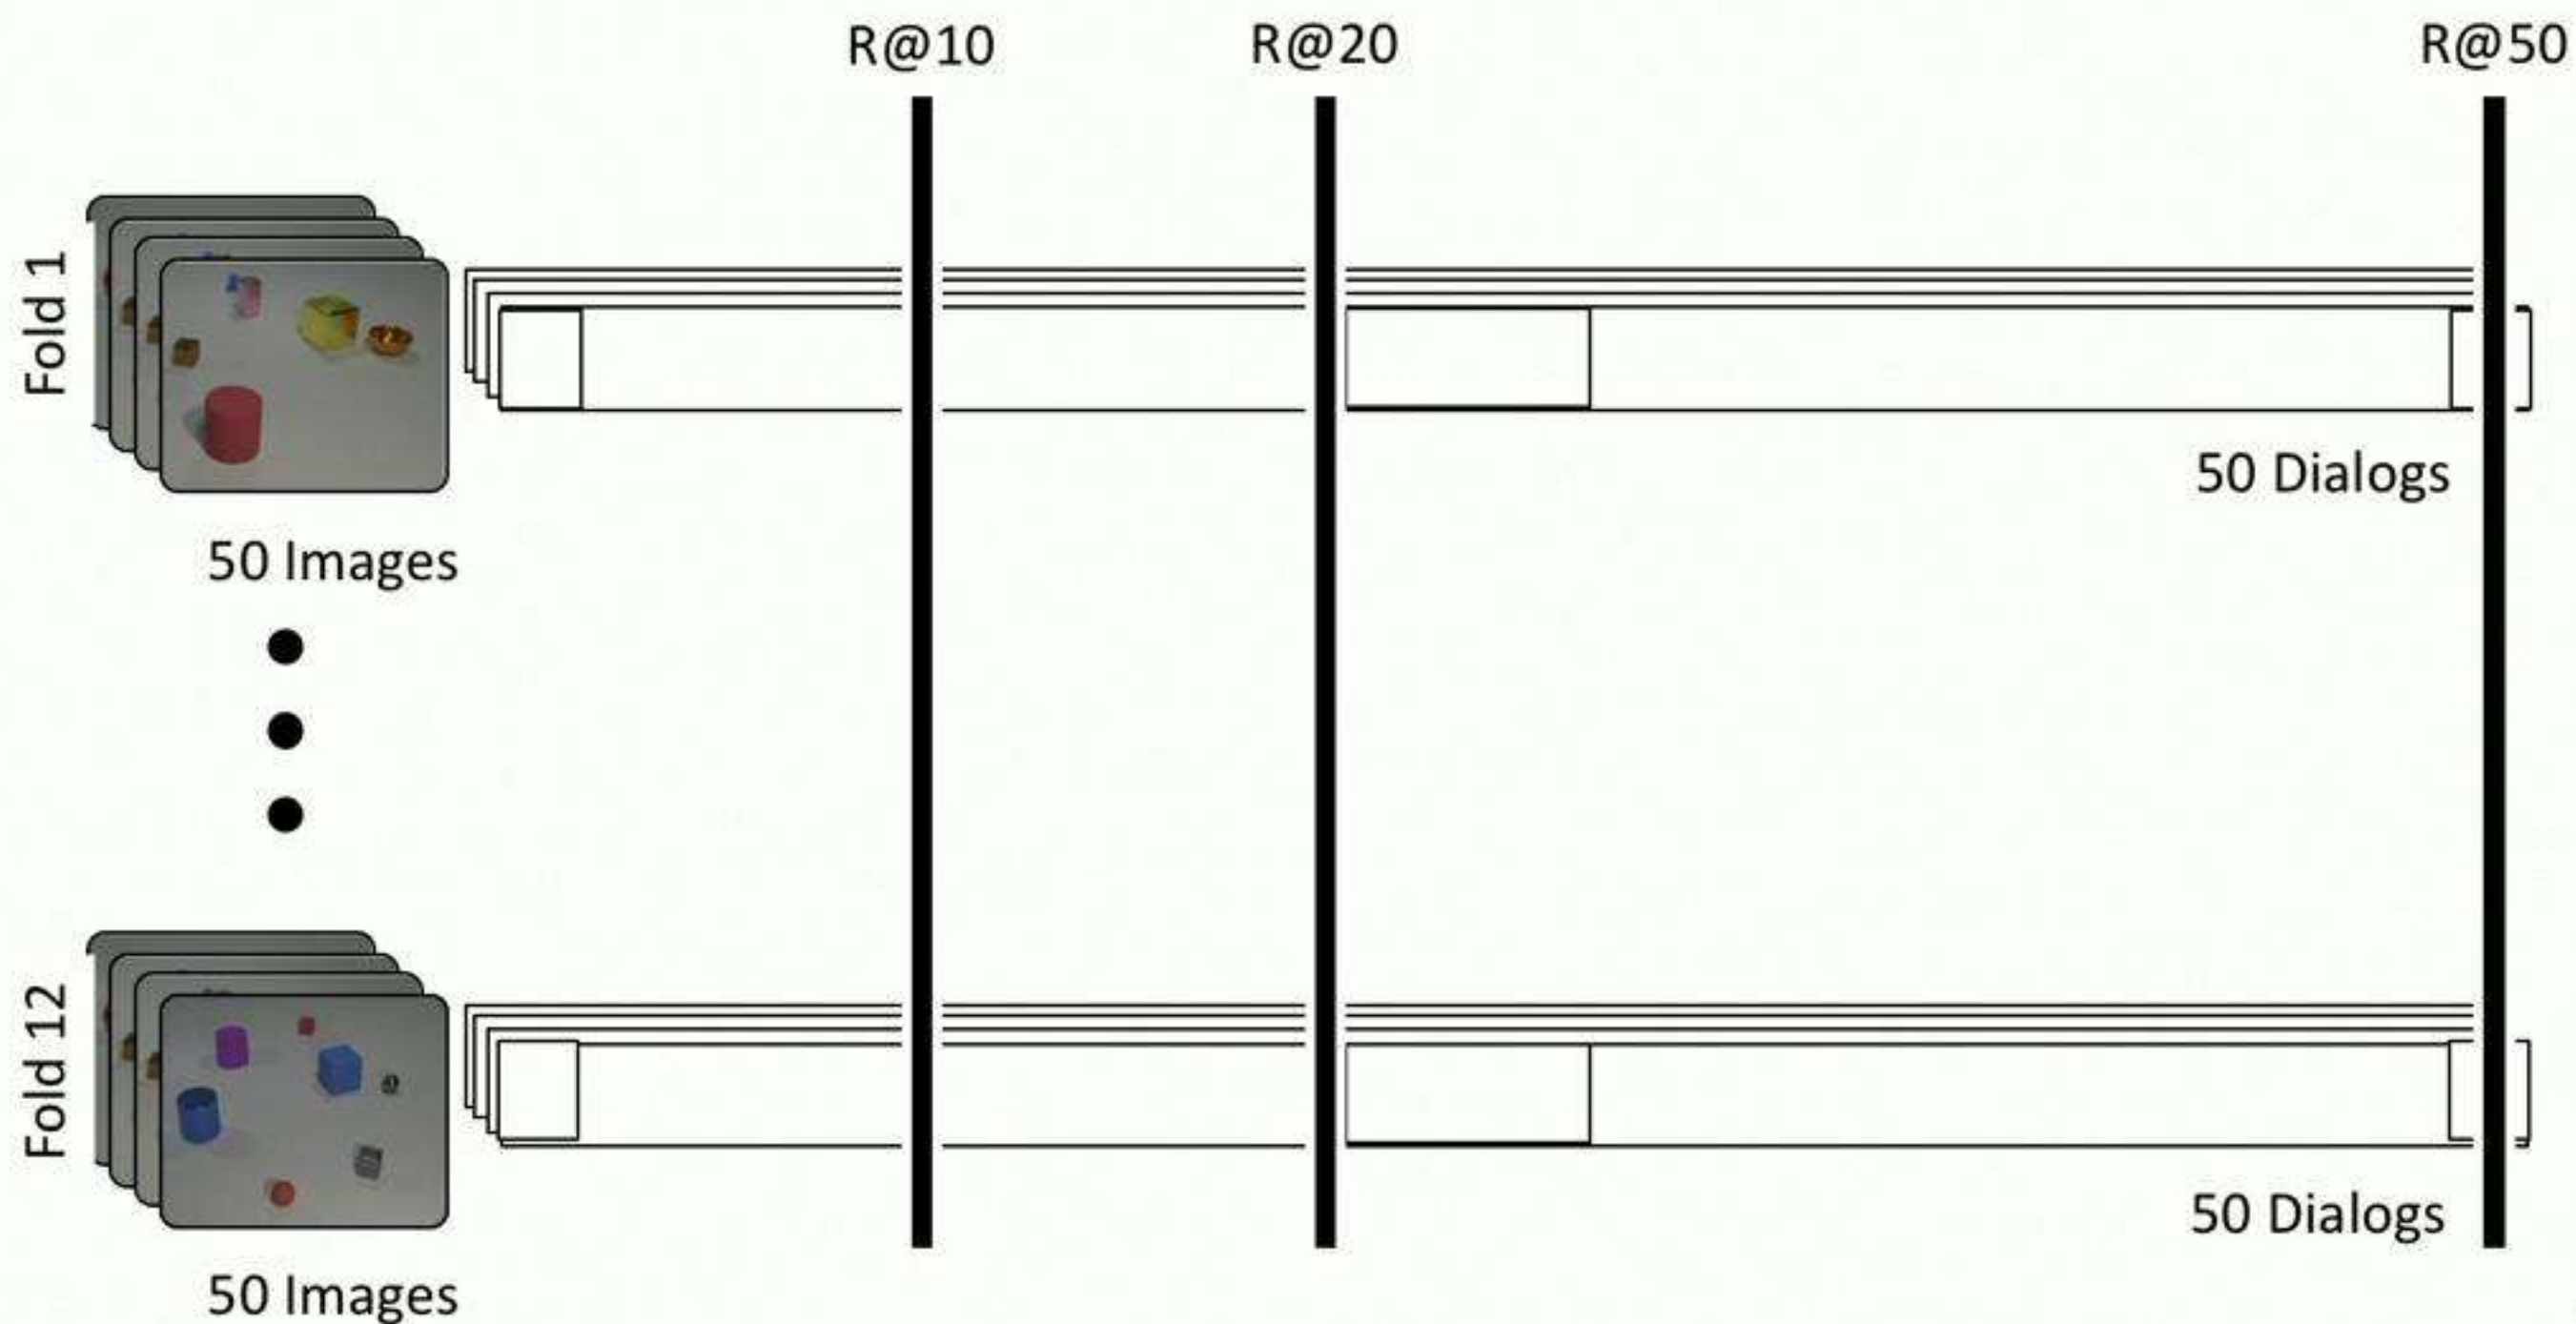
Sizes: small, large

#Objects: 5-10

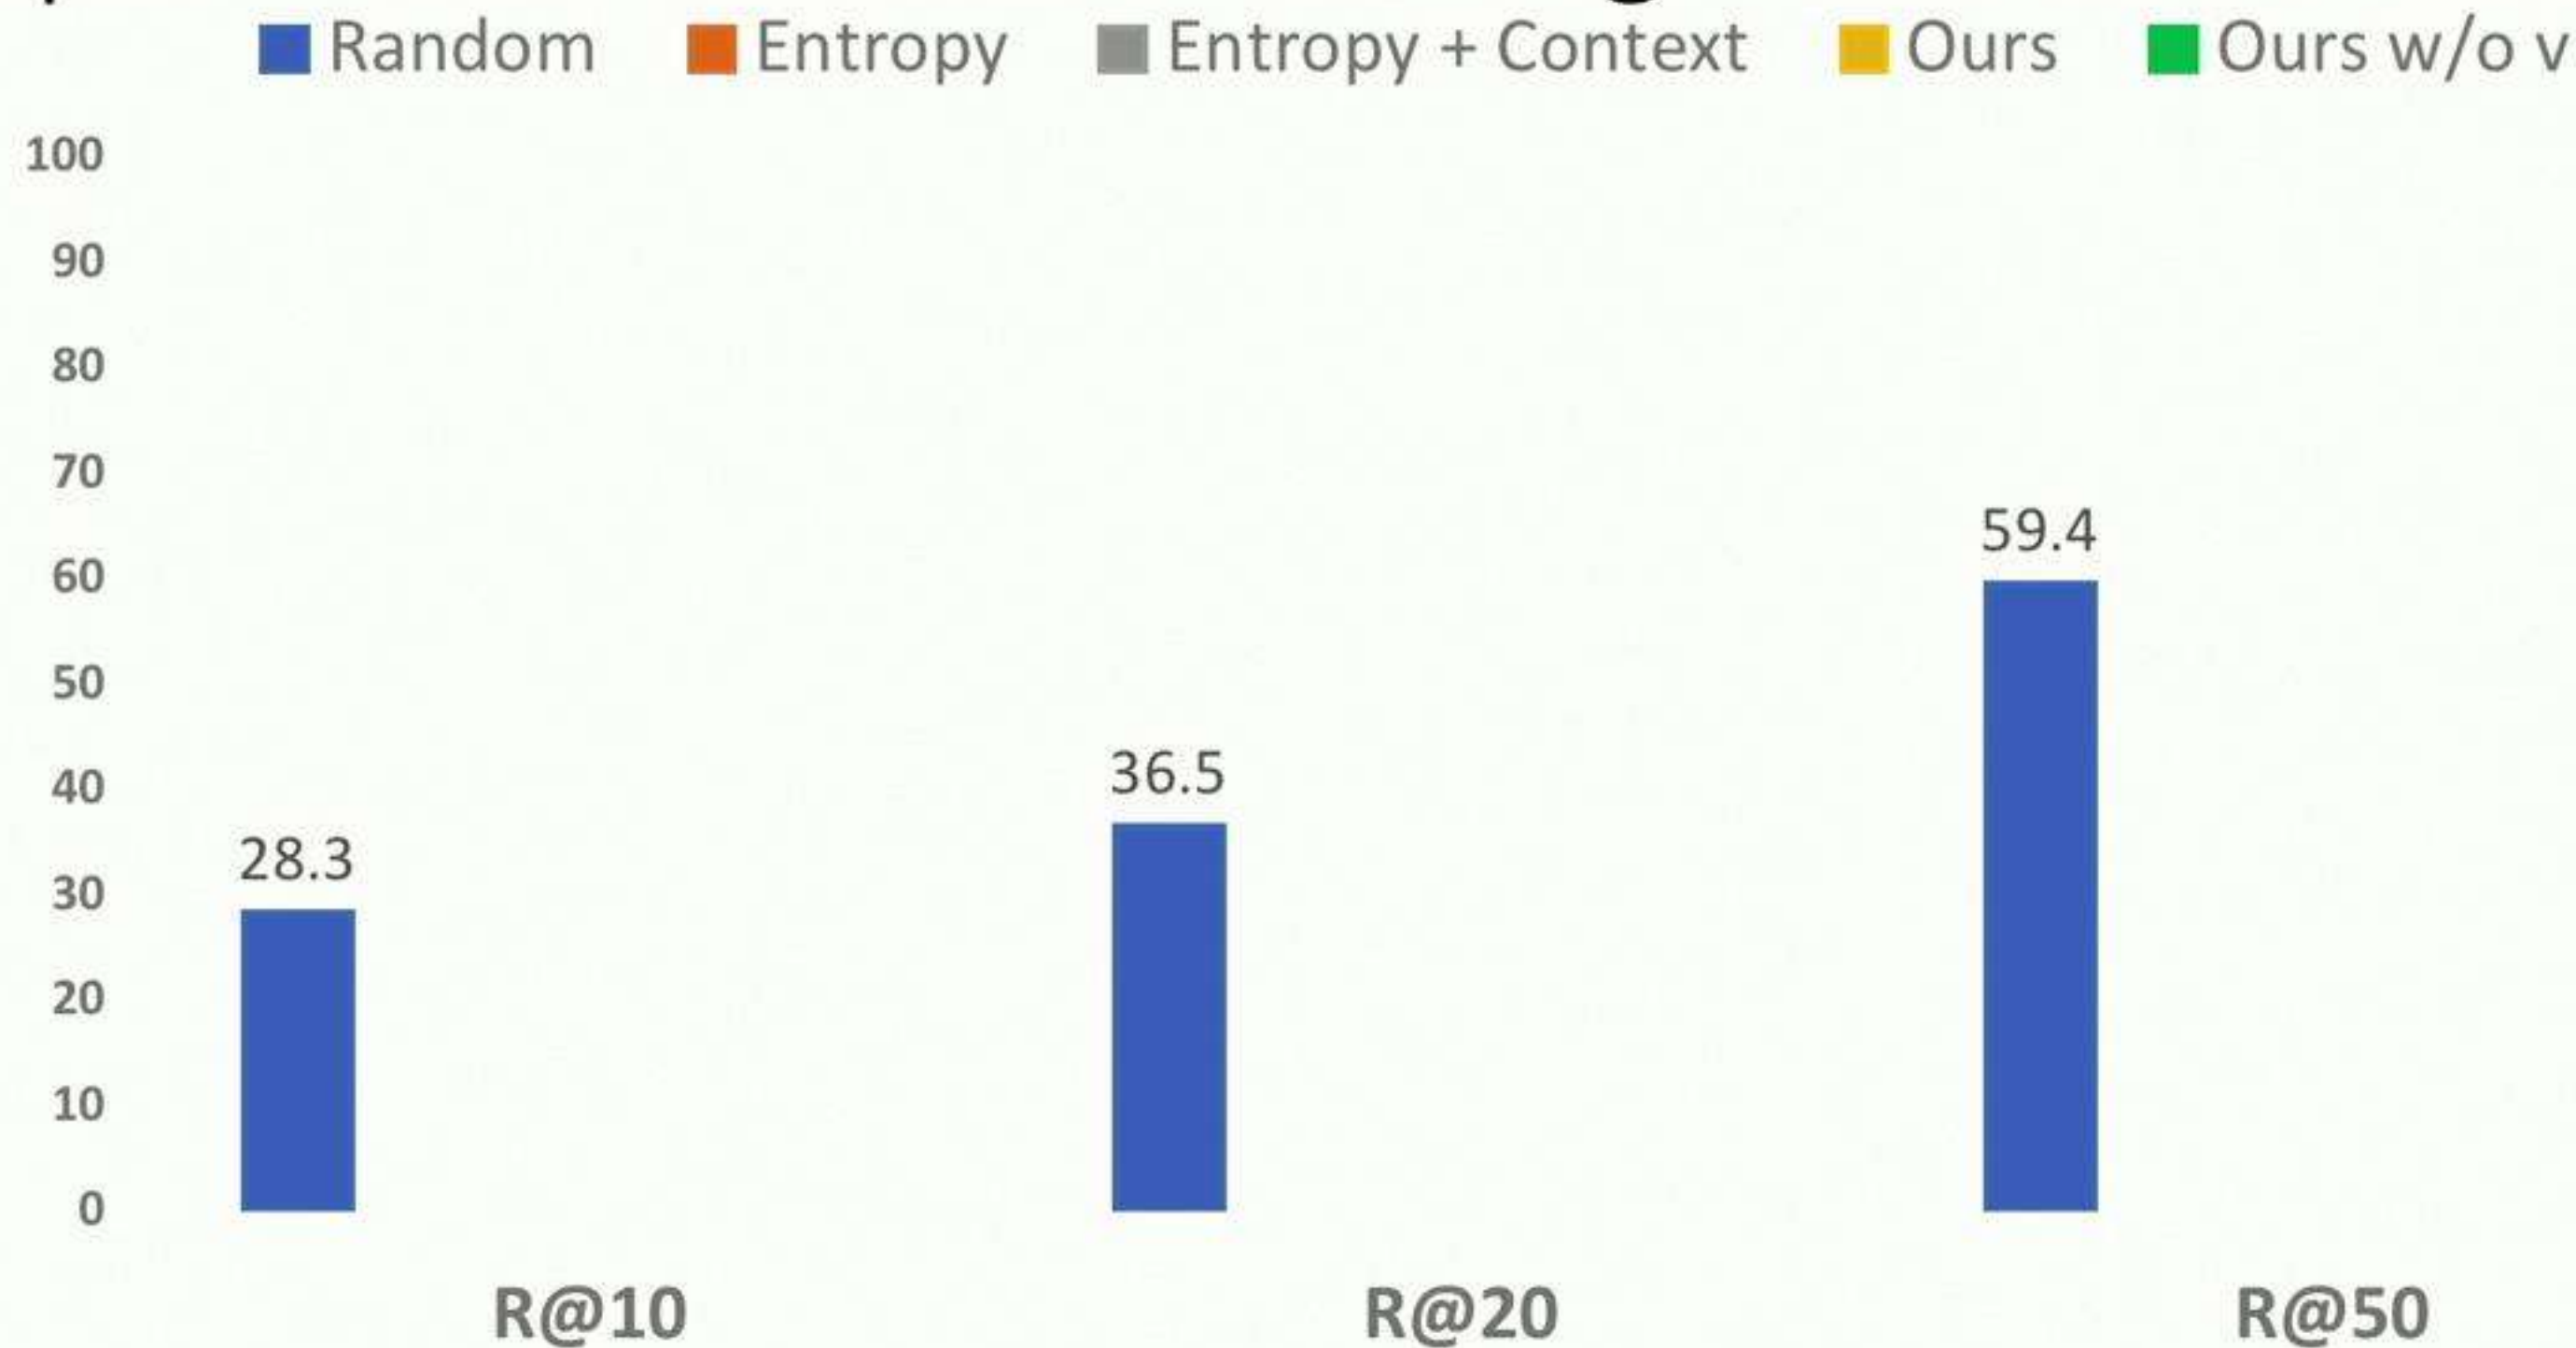
	Train	Val	Test
# Images	900	300	600

Experiments: Metrics

Graph Recovery $R@k$: Average accuracy of graph memory at dialog round k

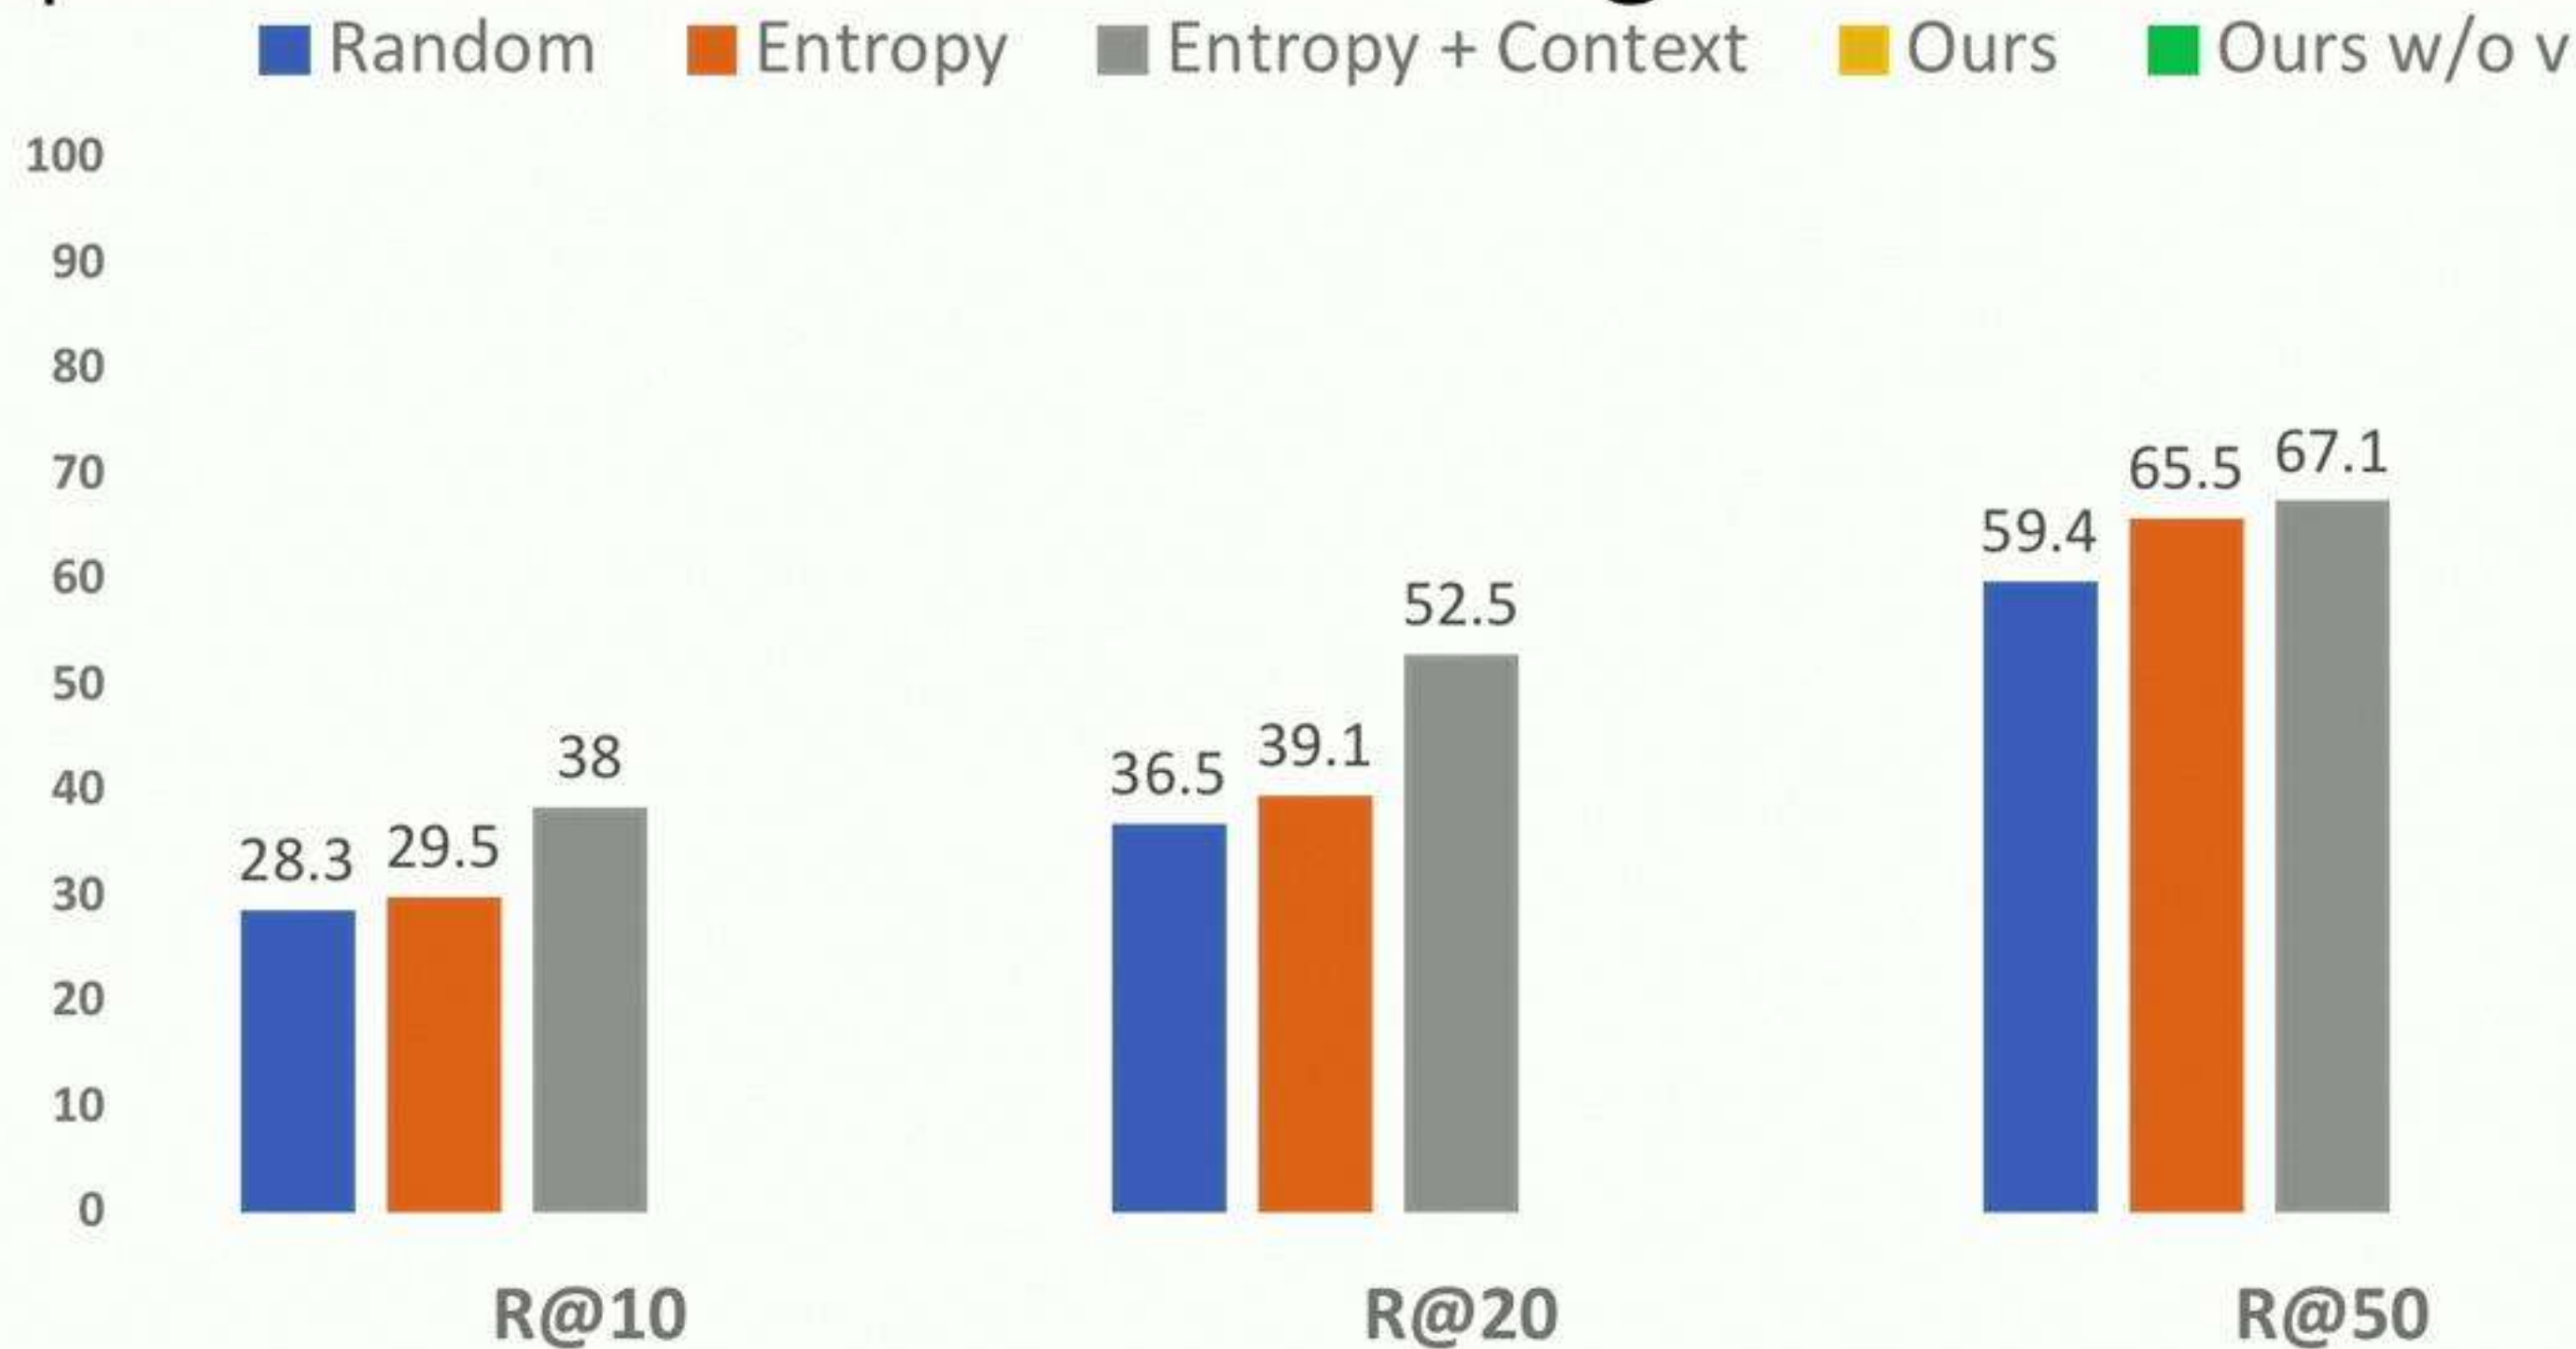


Experiment: Standard Training + Standard Testing



Graph Recovery

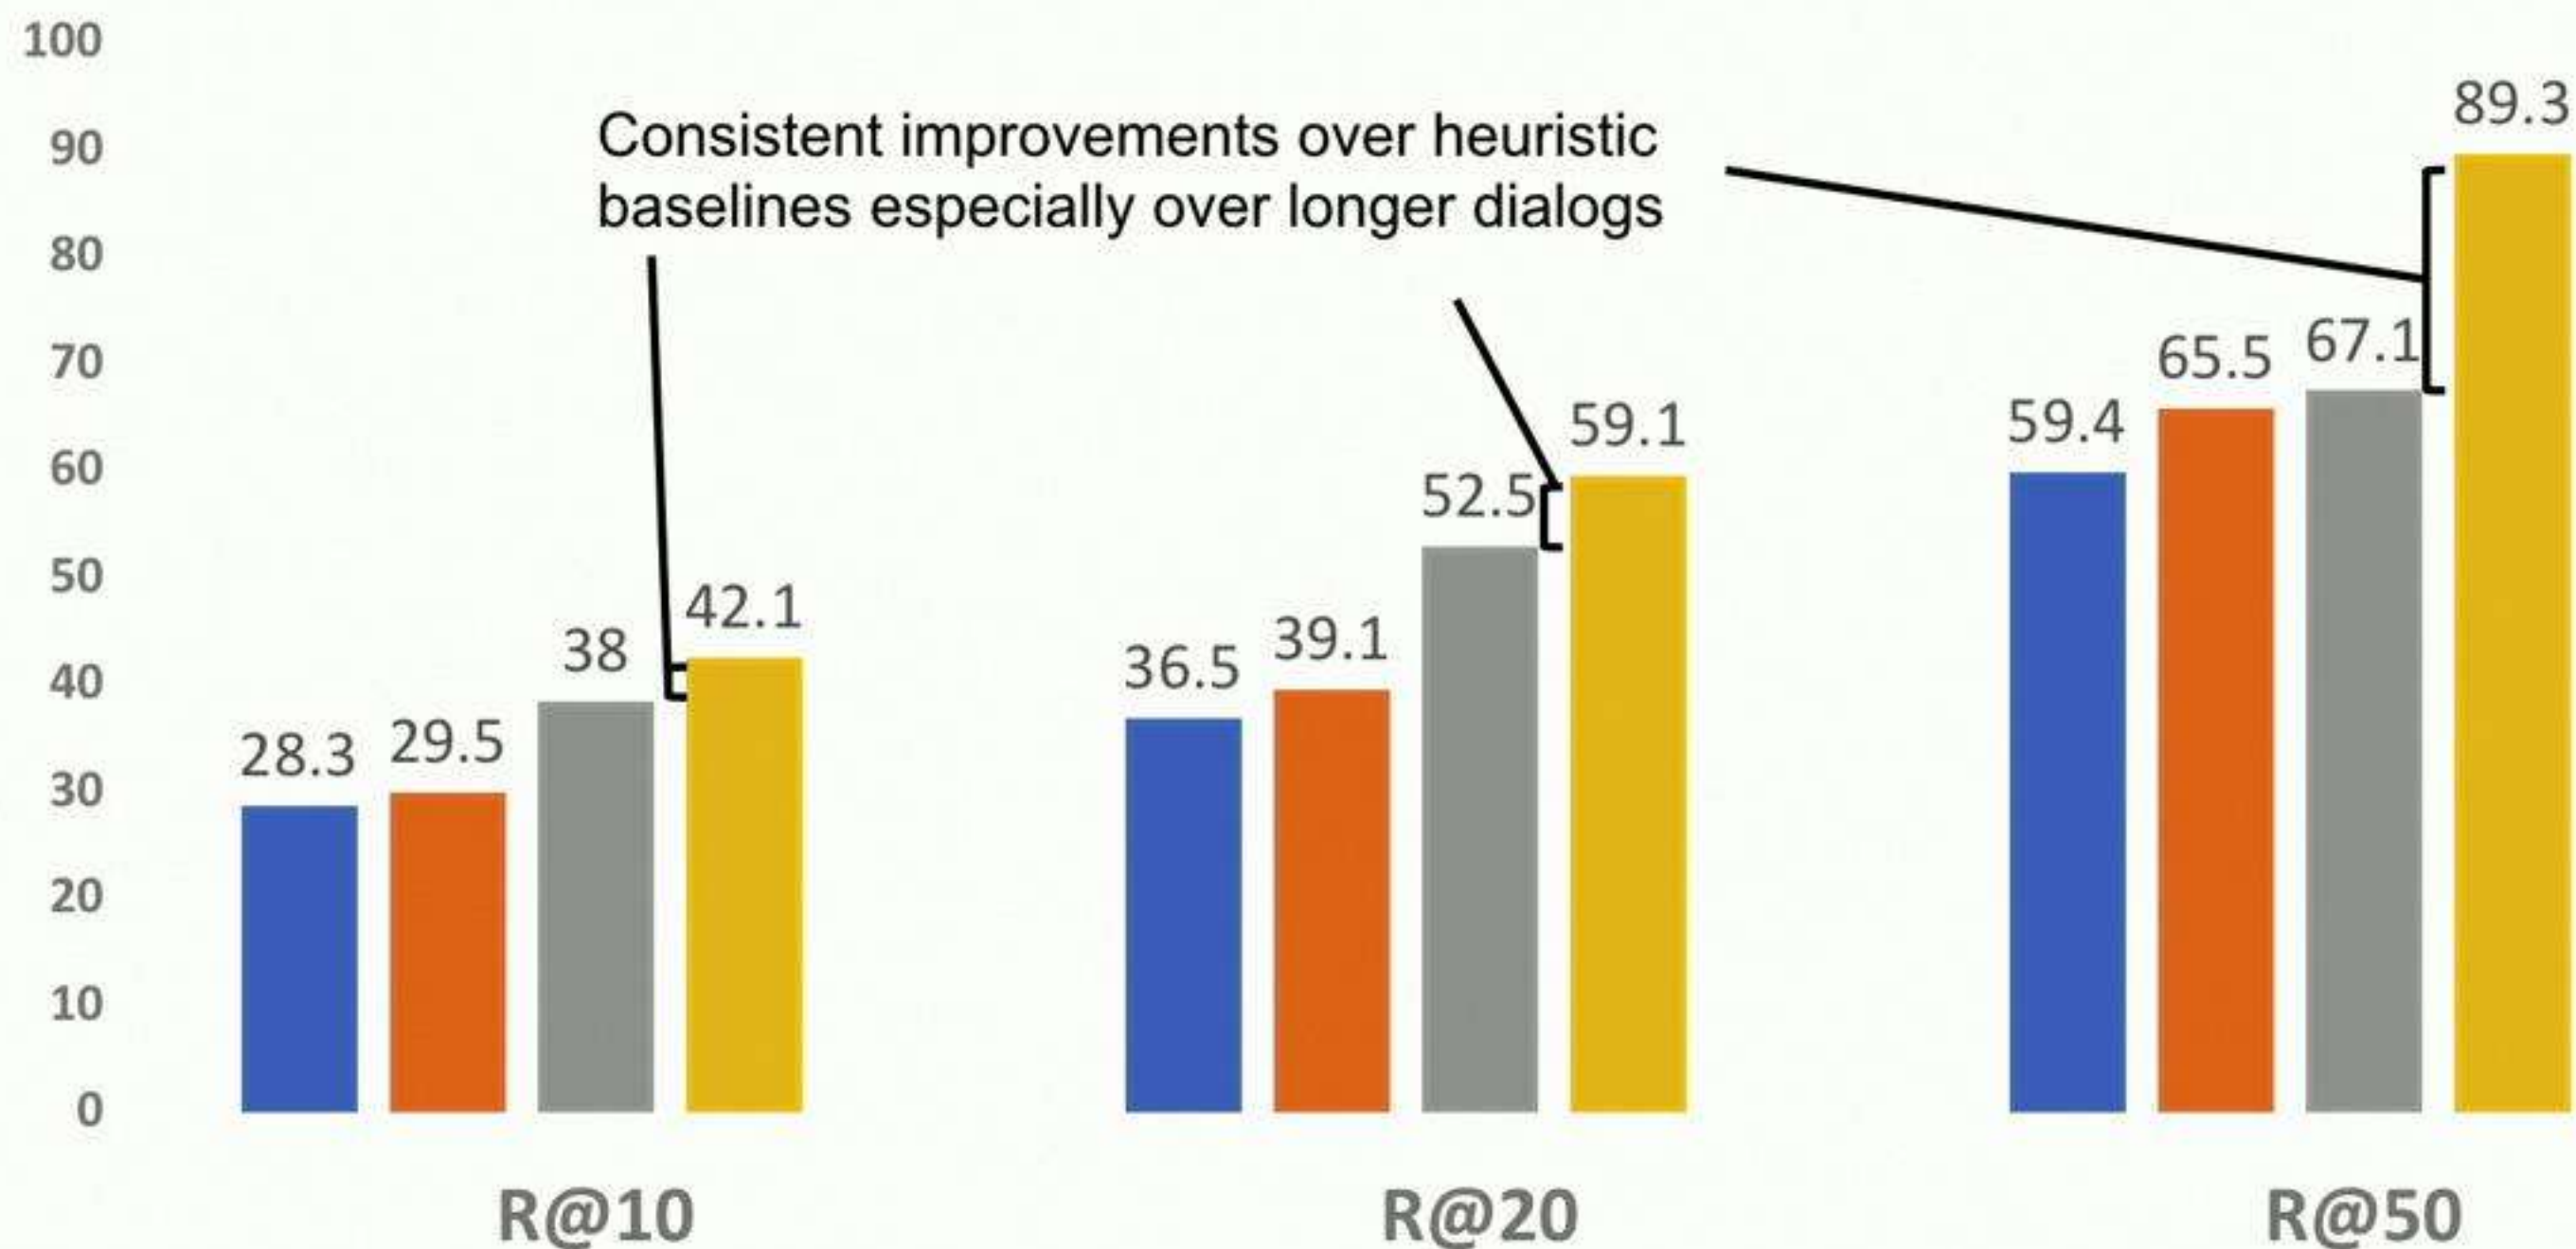
Experiment: Standard Training + Standard Testing



Graph Recovery

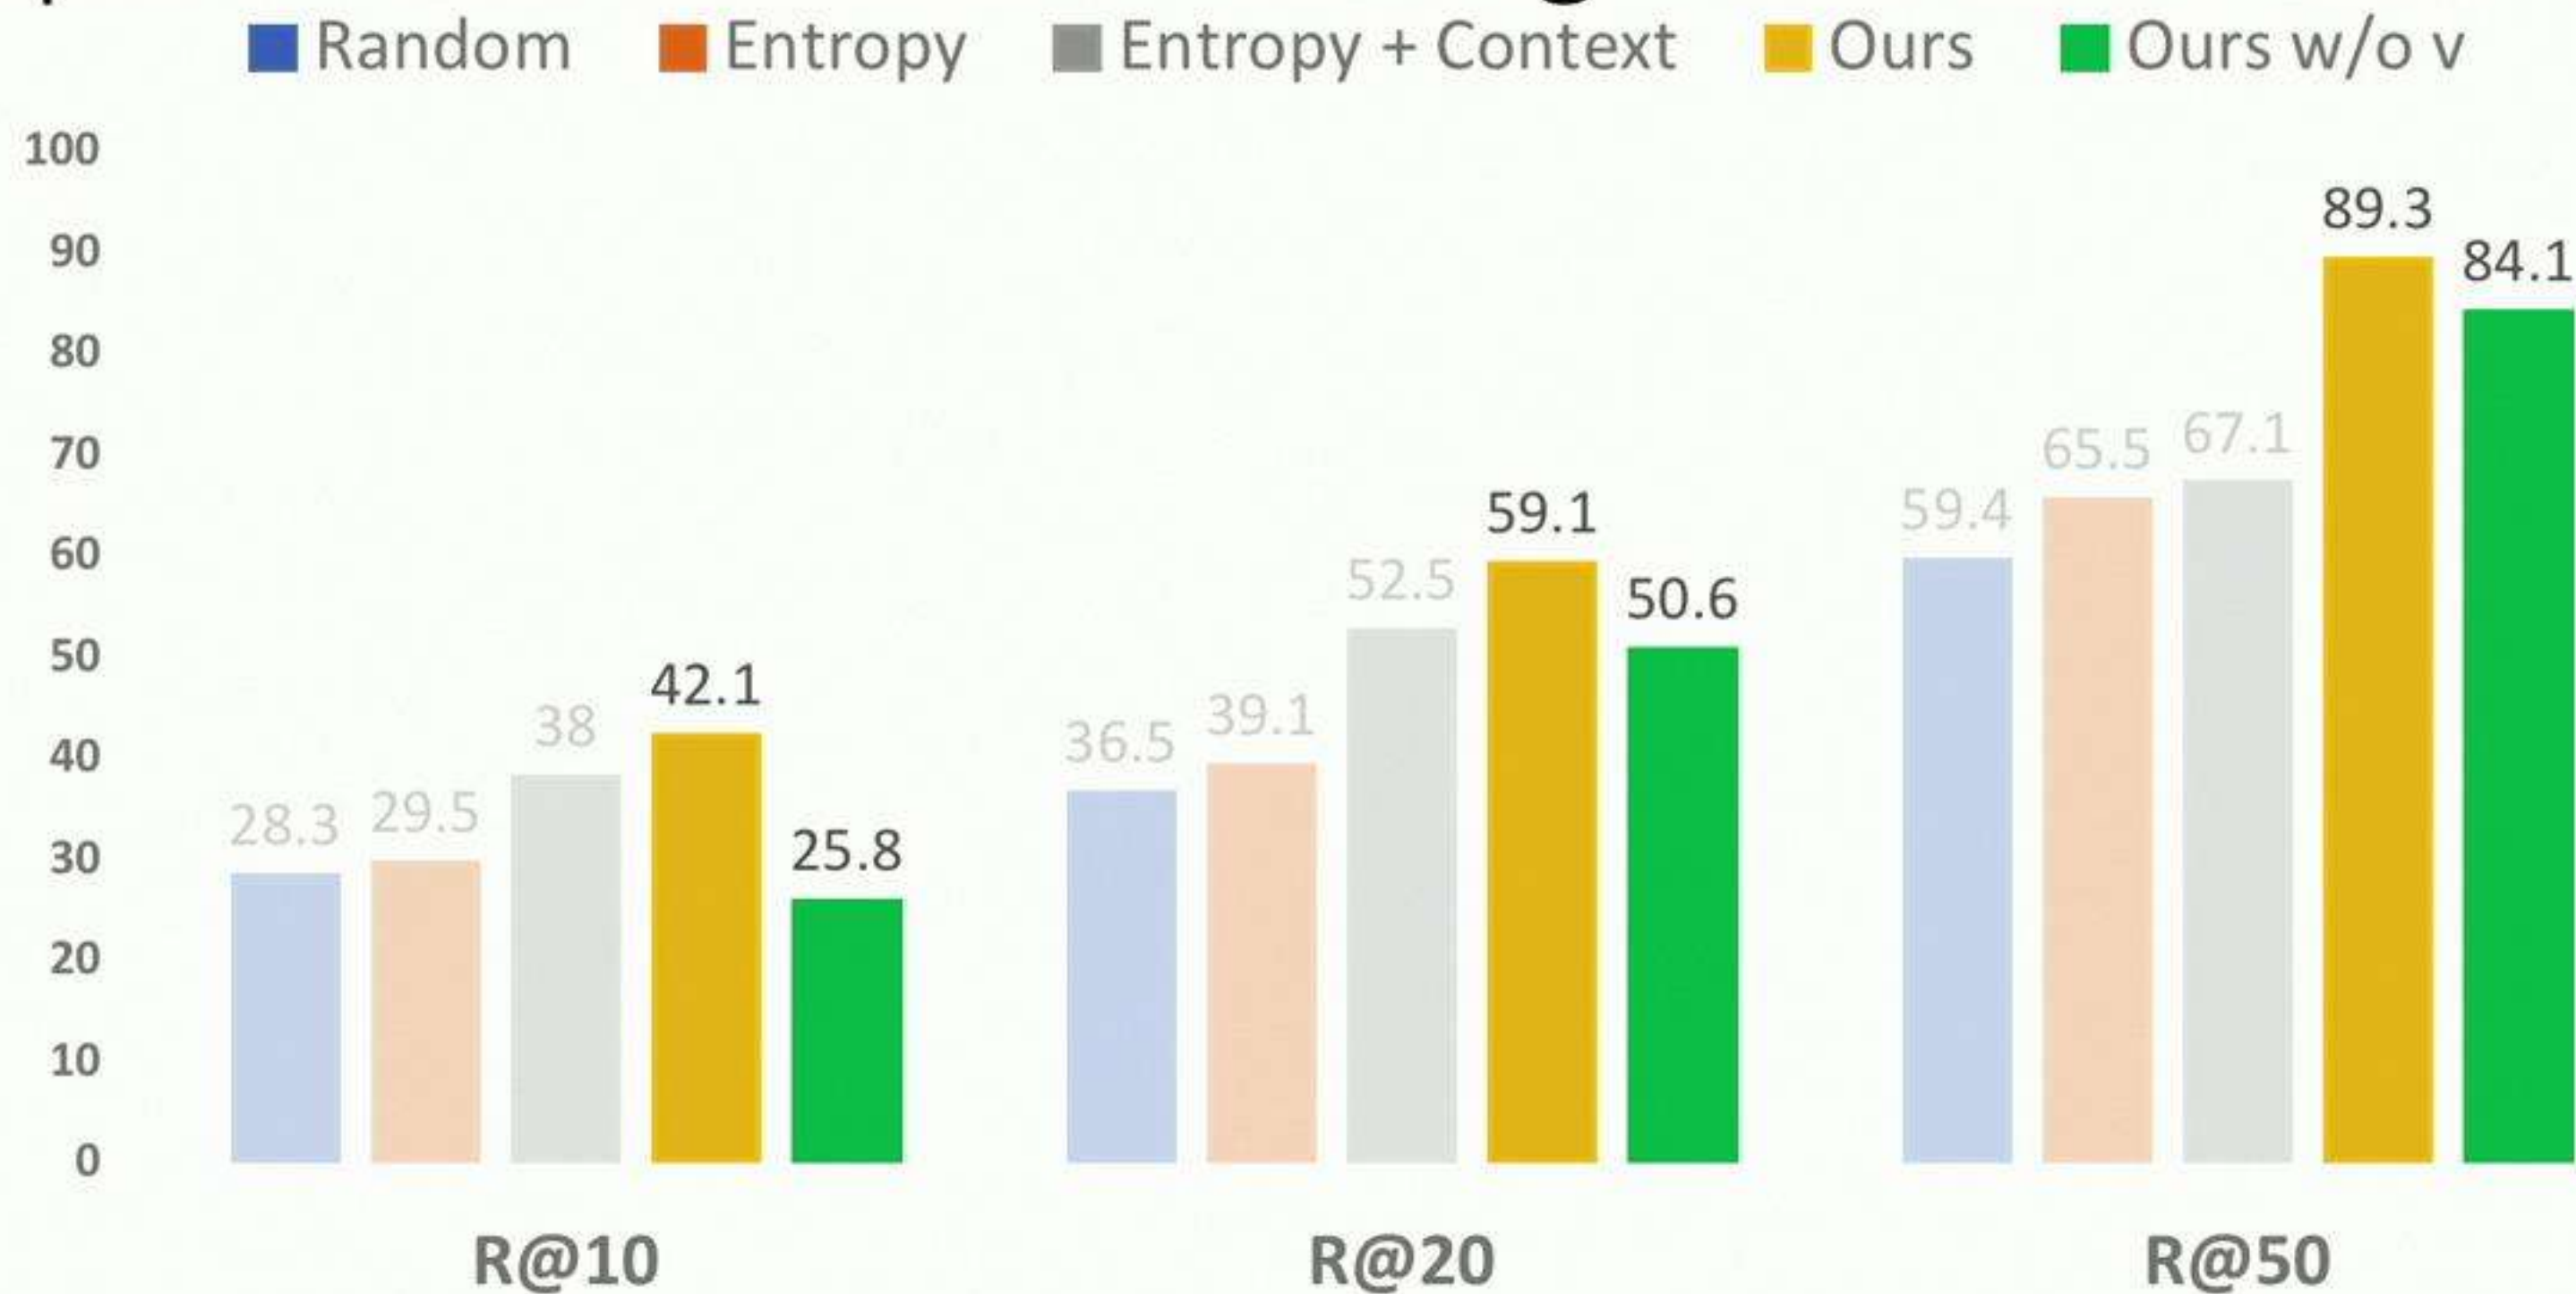
Experiment: Standard Training + Standard Testing

■ Random ■ Entropy ■ Entropy + Context ■ Ours ■ Ours w/o v



Graph Recovery

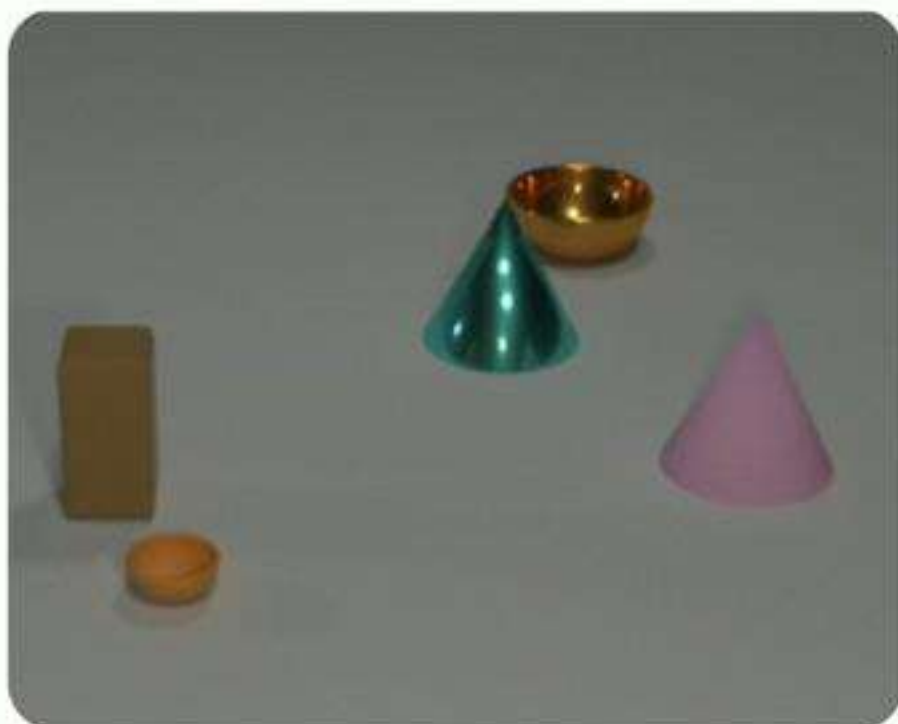
Experiment: Standard Training + Standard Testing



Graph Recovery

Experiments: Novel Object Environments

Novel



New colors and shapes

600 images for test
(12 episodes)

Mixed



Mix of novel and standard
colors and shapes

600 images for test
(12 episodes)

Realistic



51 categories, 11 colors,
6 materials

1200 images for test
(24 episodes)

Experiments: Novel Object Environments

Novel

Mixed

Realistic

Note that questioners are trained on Standard and then evaluated in these new settings with randomly initialized visual systems

New colors and shapes

Mix of novel and standard
colors and shapes

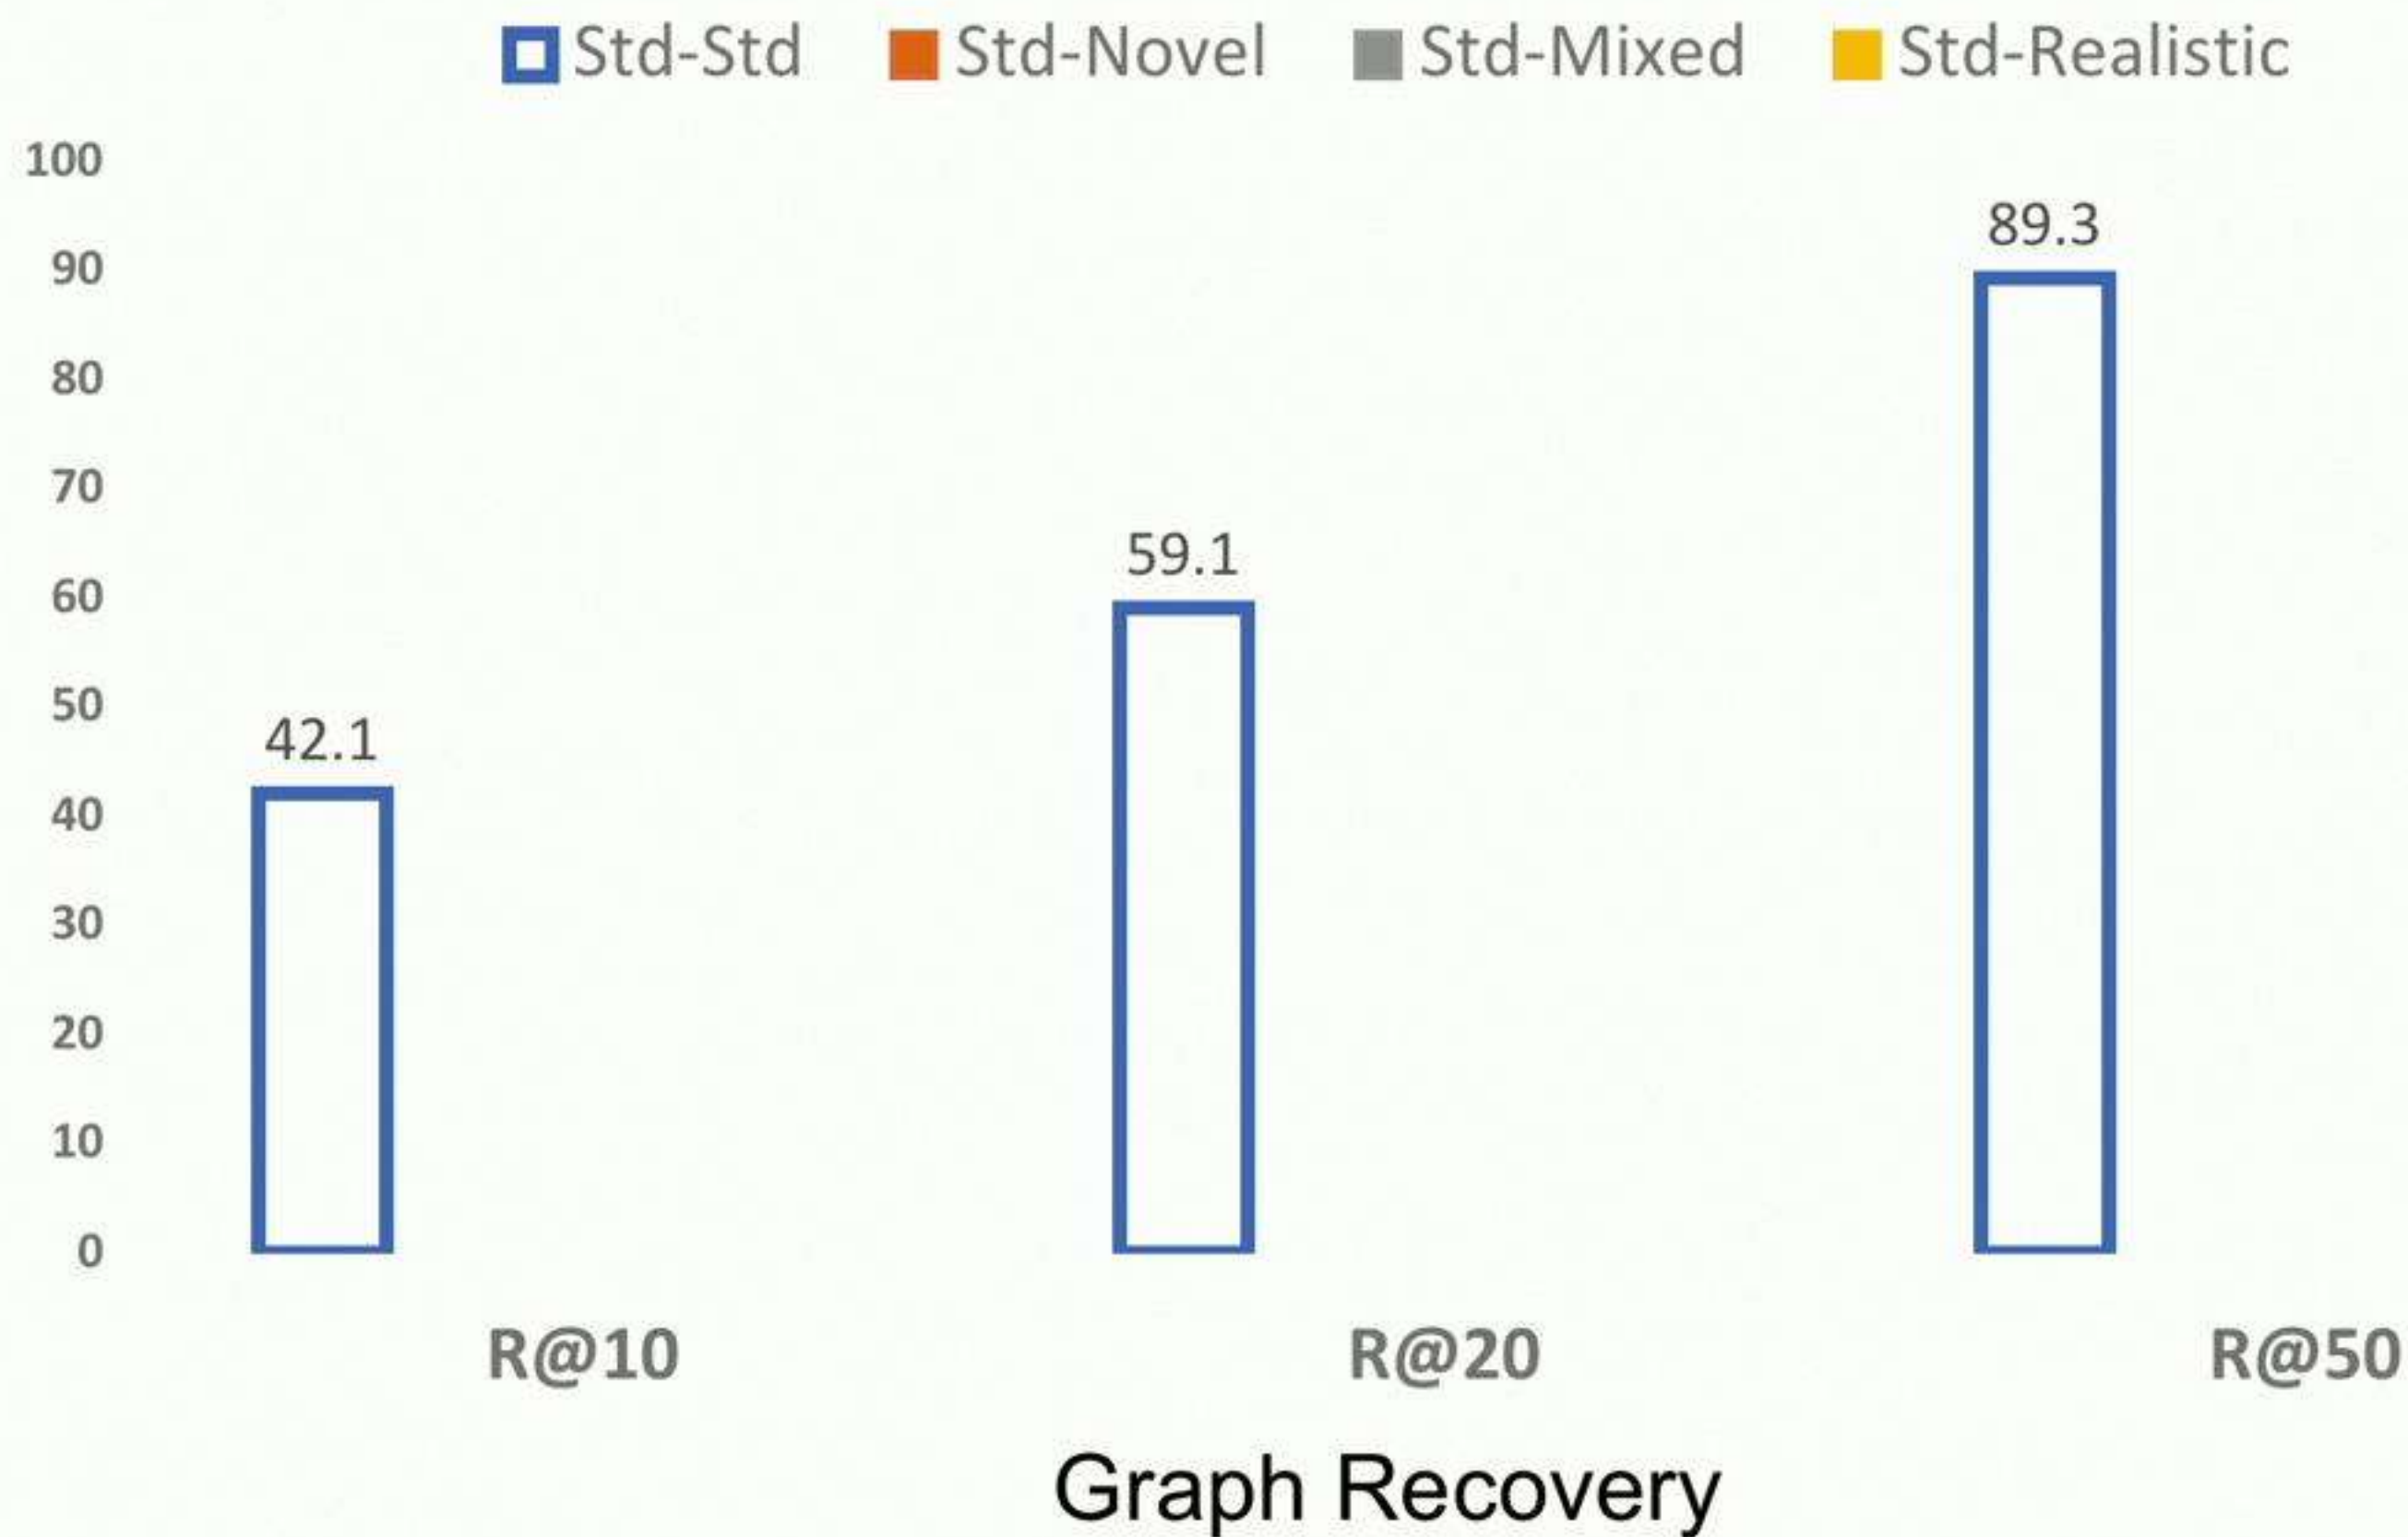
51 categories, 11 colors,
6 materials

600 images for test
(12 episodes)

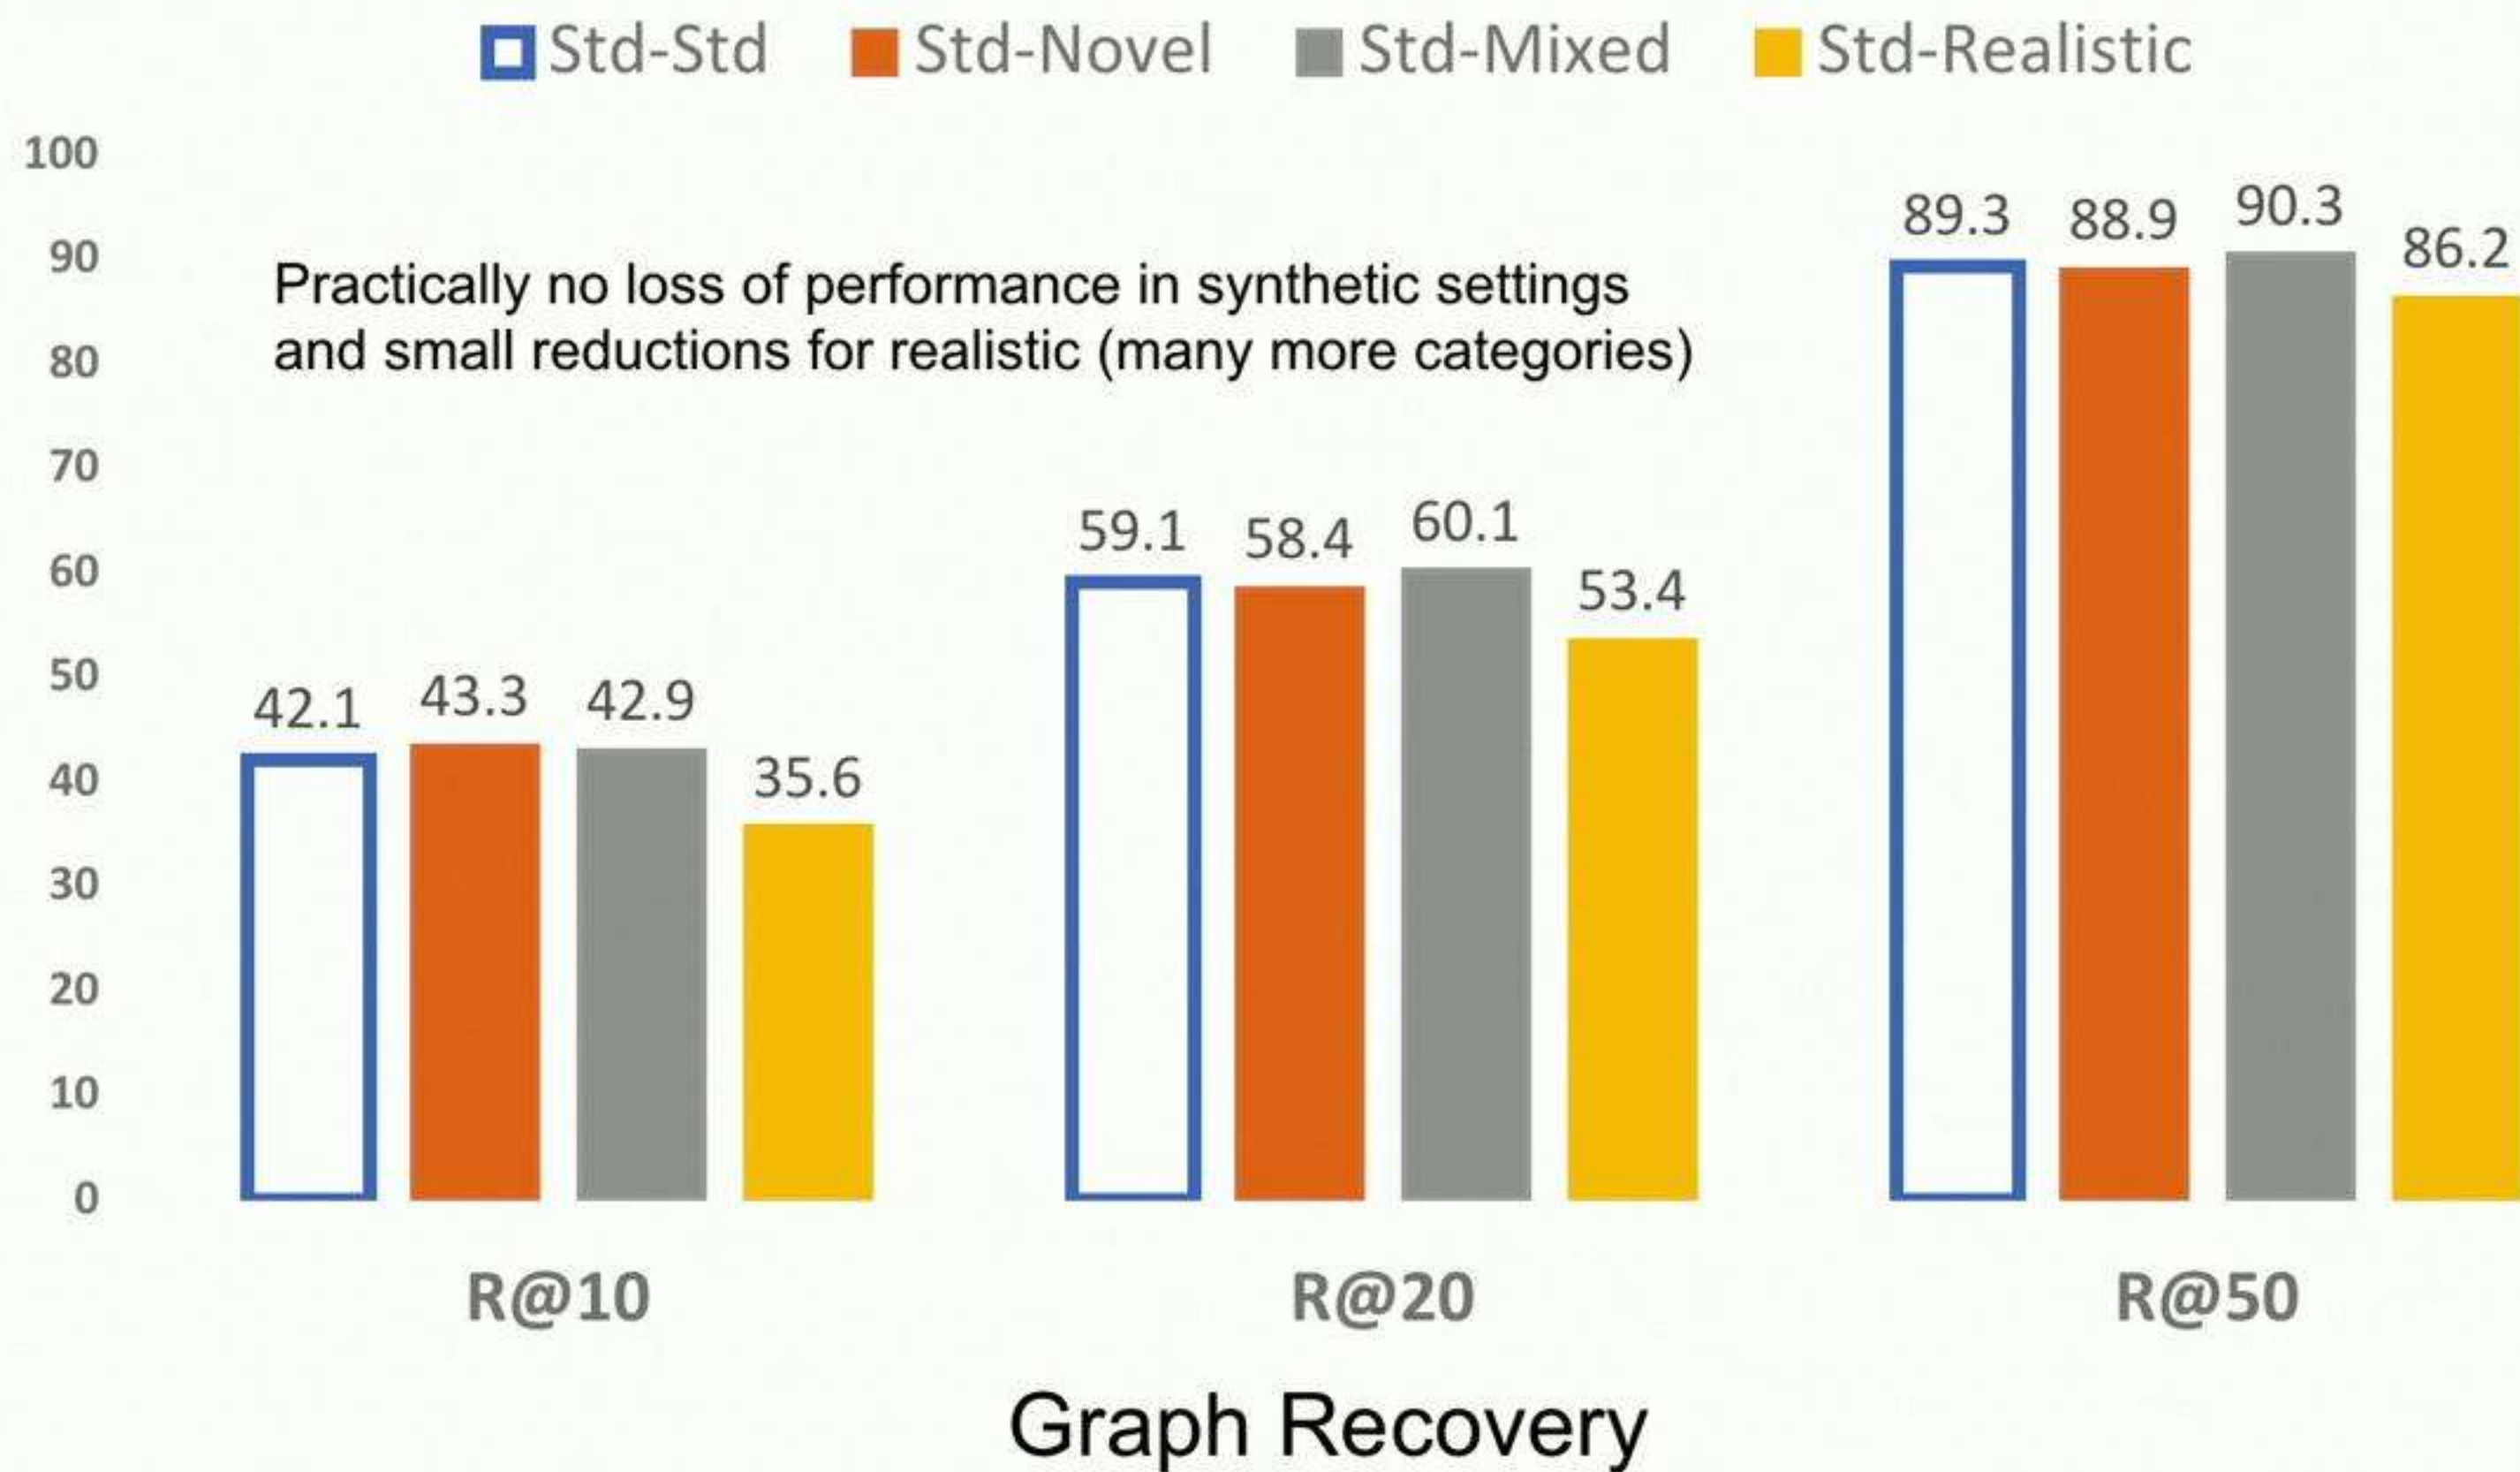
600 images for test
(12 episodes)

1200 images for test
(24 episodes)

Experiments: Standard Train – New Test Environments



Experiments: Standard Train – New Test Environments



Experiments: Qualitative Example

What material is the leftmost thing?



Experiments: Qualitative Example



What material is the leftmost thing?

food

There is a leftmost object; what is it?

Experiments: Qualitative Example



What material is the leftmost thing?

food

There is a leftmost object; what is it?

potato

The leftmost object is what color?

Experiments: Qualitative Example



What material is the leftmost thing?

food

There is a leftmost object; what is it?

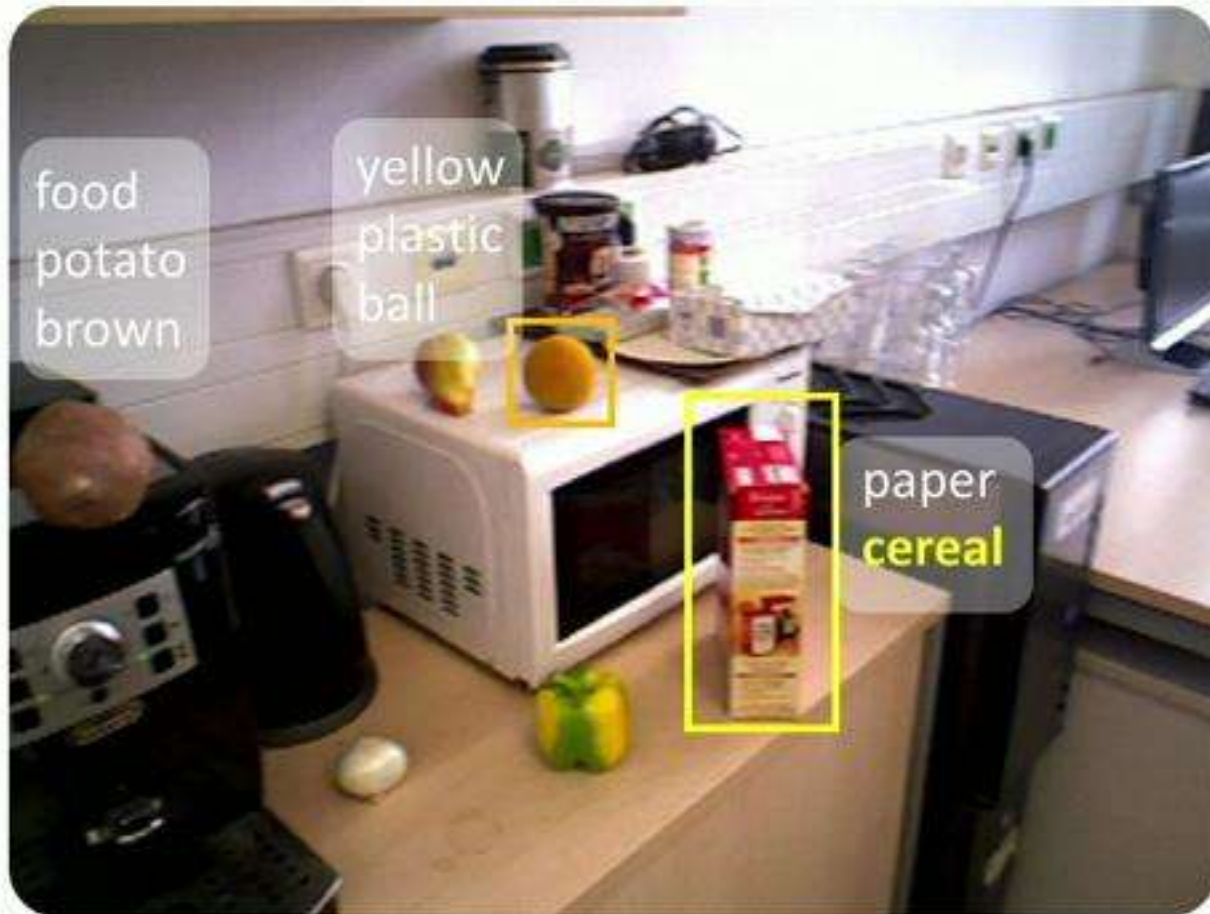
potato

The leftmost object is what color?

brown

What is the closest thing that is in front of the yellow plastic ball made of?

Experiments: Qualitative Example



What material is the leftmost thing?

food

There is a leftmost object; what is it?

potato

The leftmost object is what color?

brown

What is the closest thing that is in front of the yellow plastic ball made of?

paper

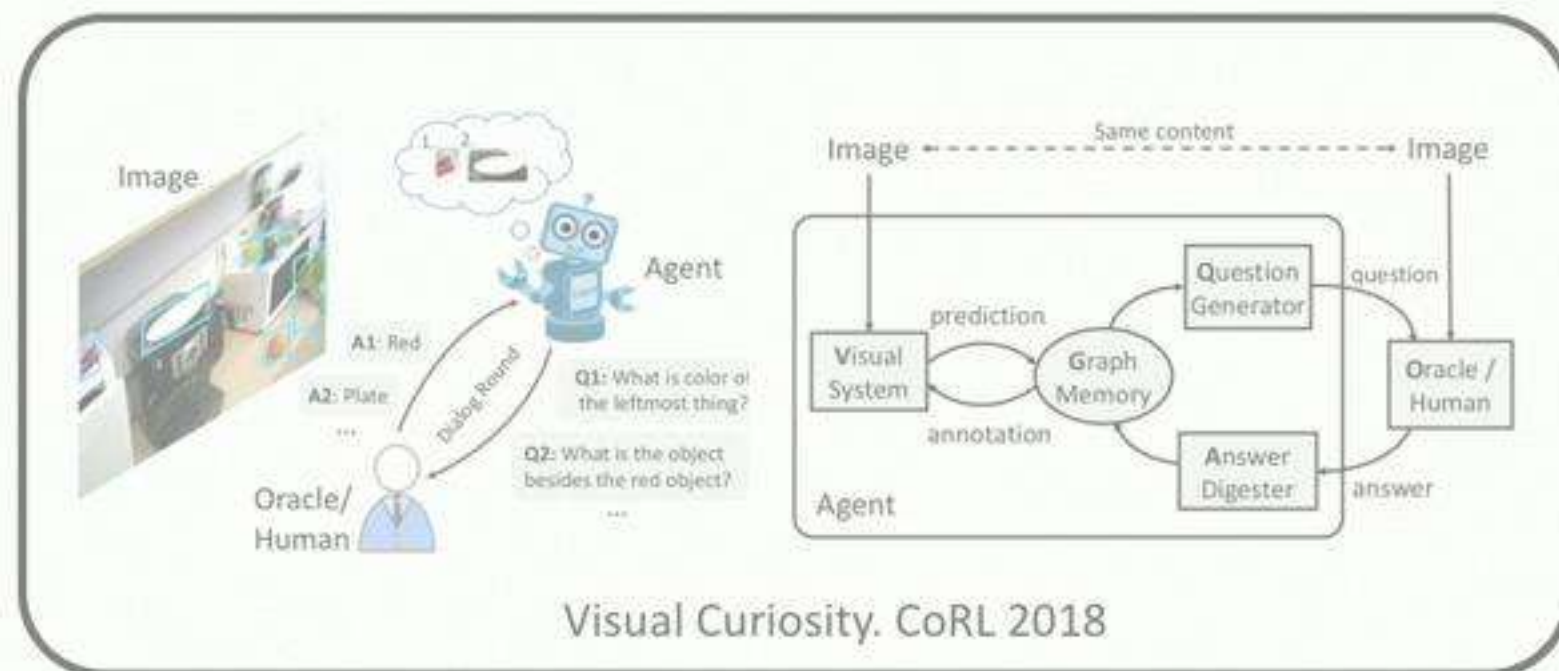
What is the closest thing that is in front of the yellow plastic ball?

cereal

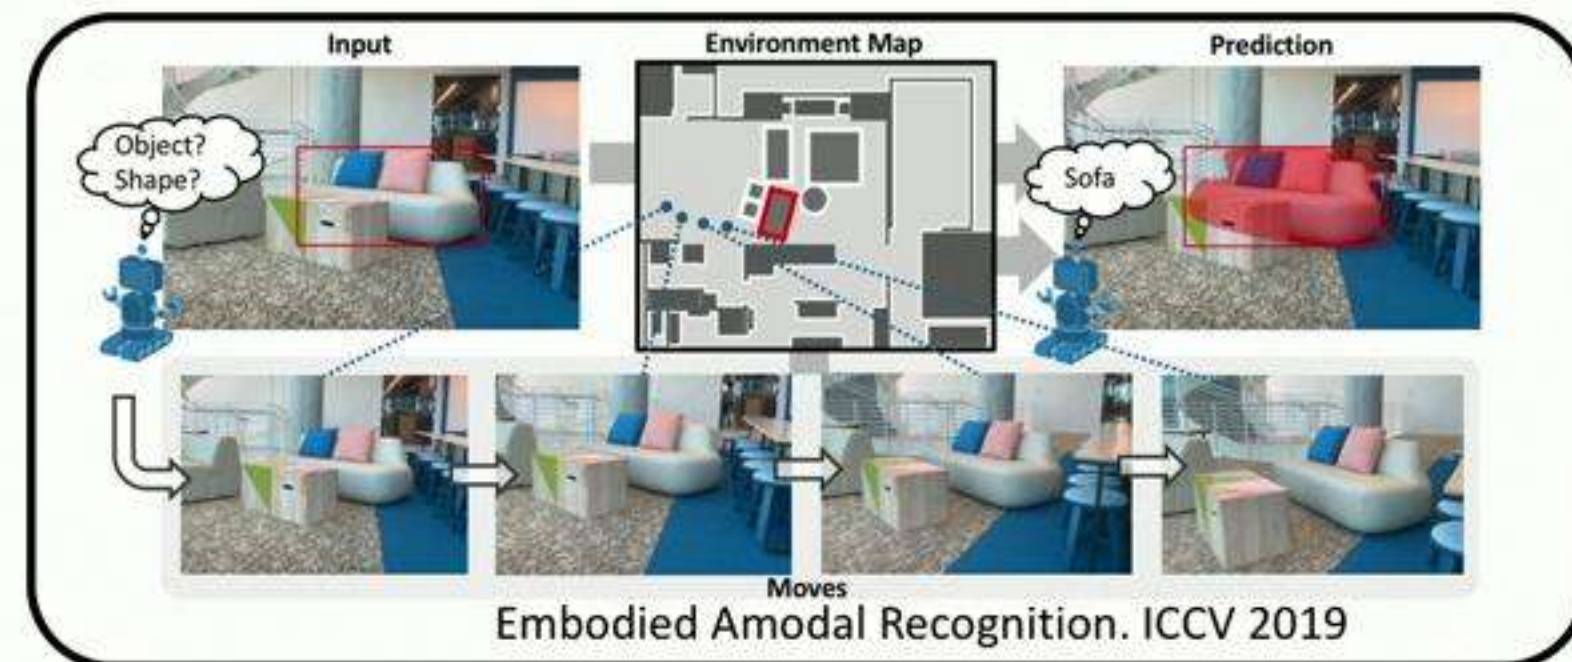
Takeaways

- A new neural-symbolic pipeline is proposed to learn visual curiosity for an agent
- Through interaction with humans (Oracle), the agent improves its visual understanding capacity gradually
- The learned questions generation policy can directly adapt from synthetic dataset to realistic dataset

In this talk

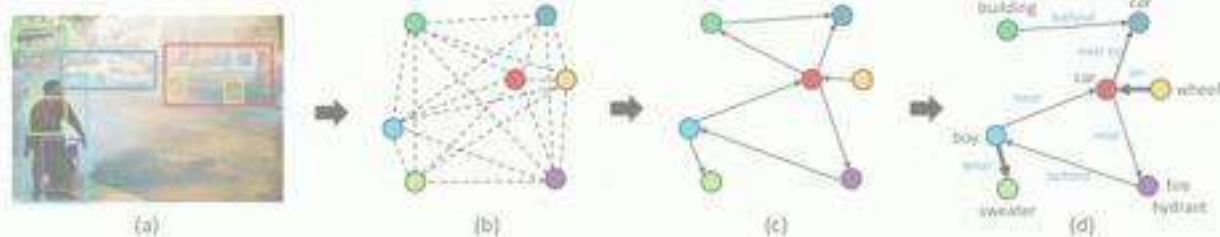


Interact with Human



Interact with Environment

Per-Image Structure



Graph R-CNN for Scene Graph Generation. ECCV 2018

Structured Visual Understanding

Visual Understanding by Moving in 3D Environment

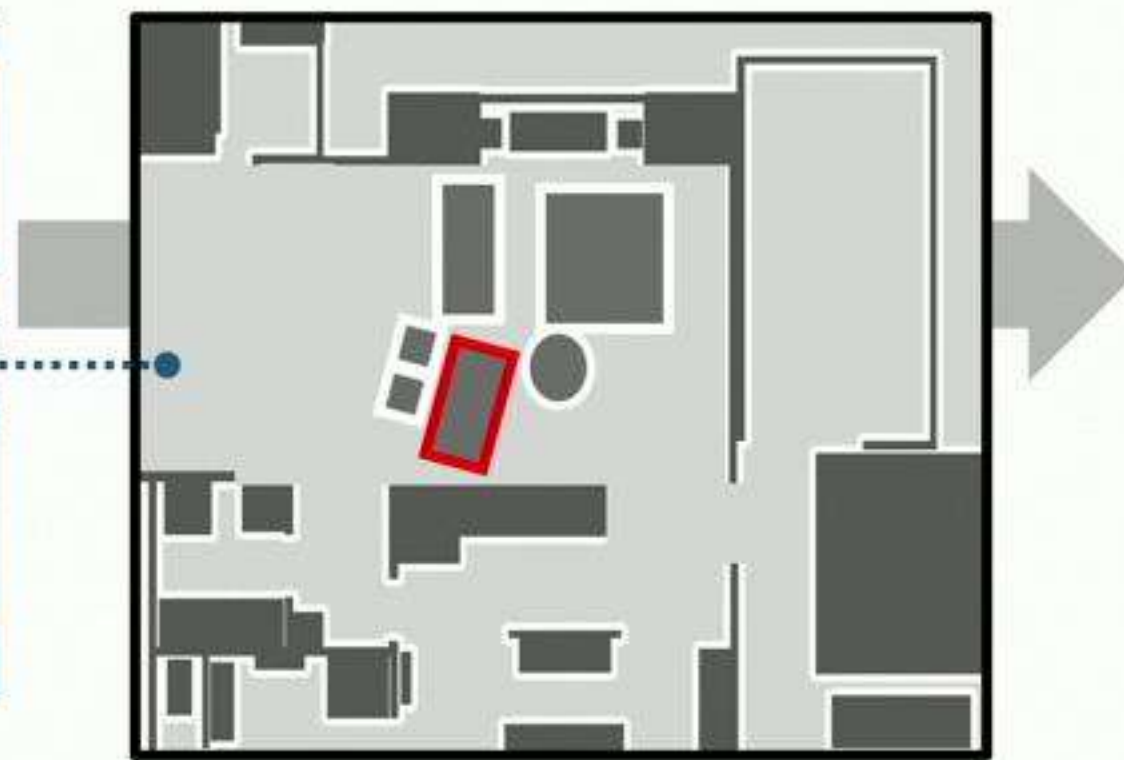
Embodied Amodal Recognition: Learning to Move to Perceive Object. ICCV 2019

Motivation

Input



Environment Map



Prediction

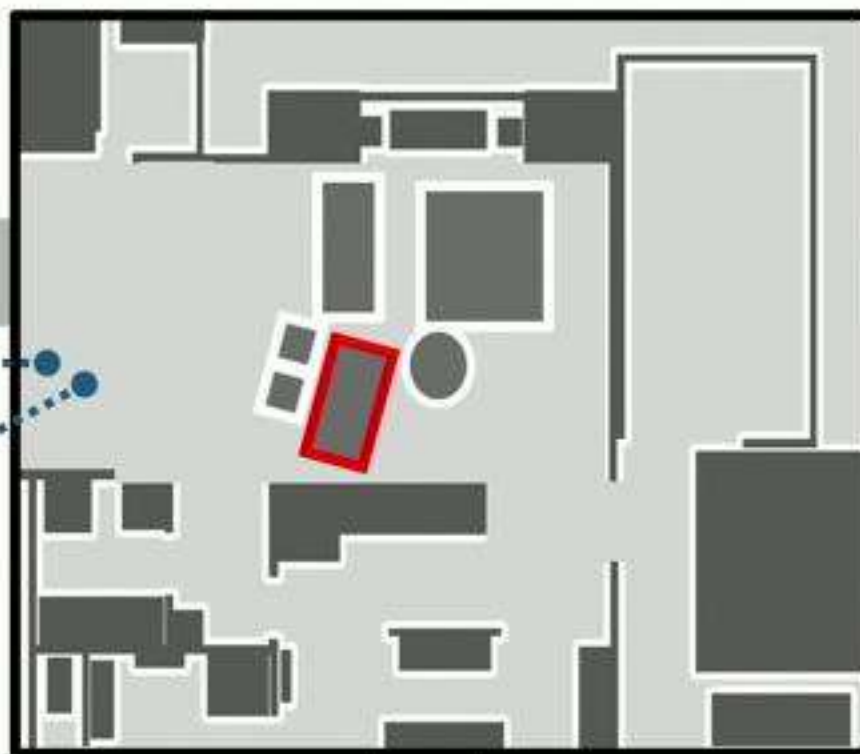


Motivation

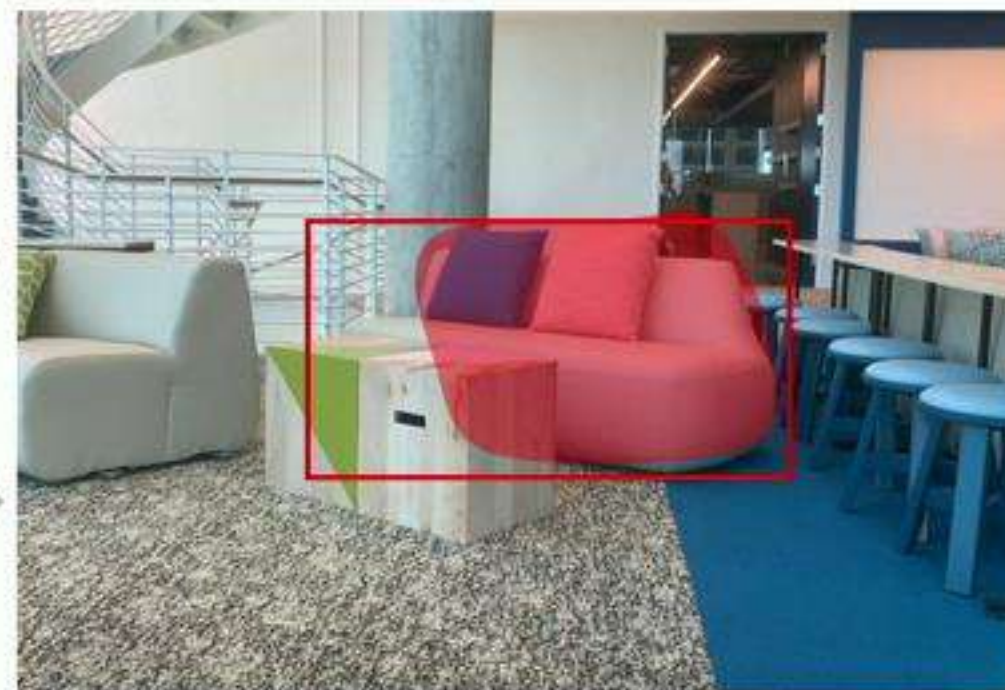
Input



Environment Map



Prediction

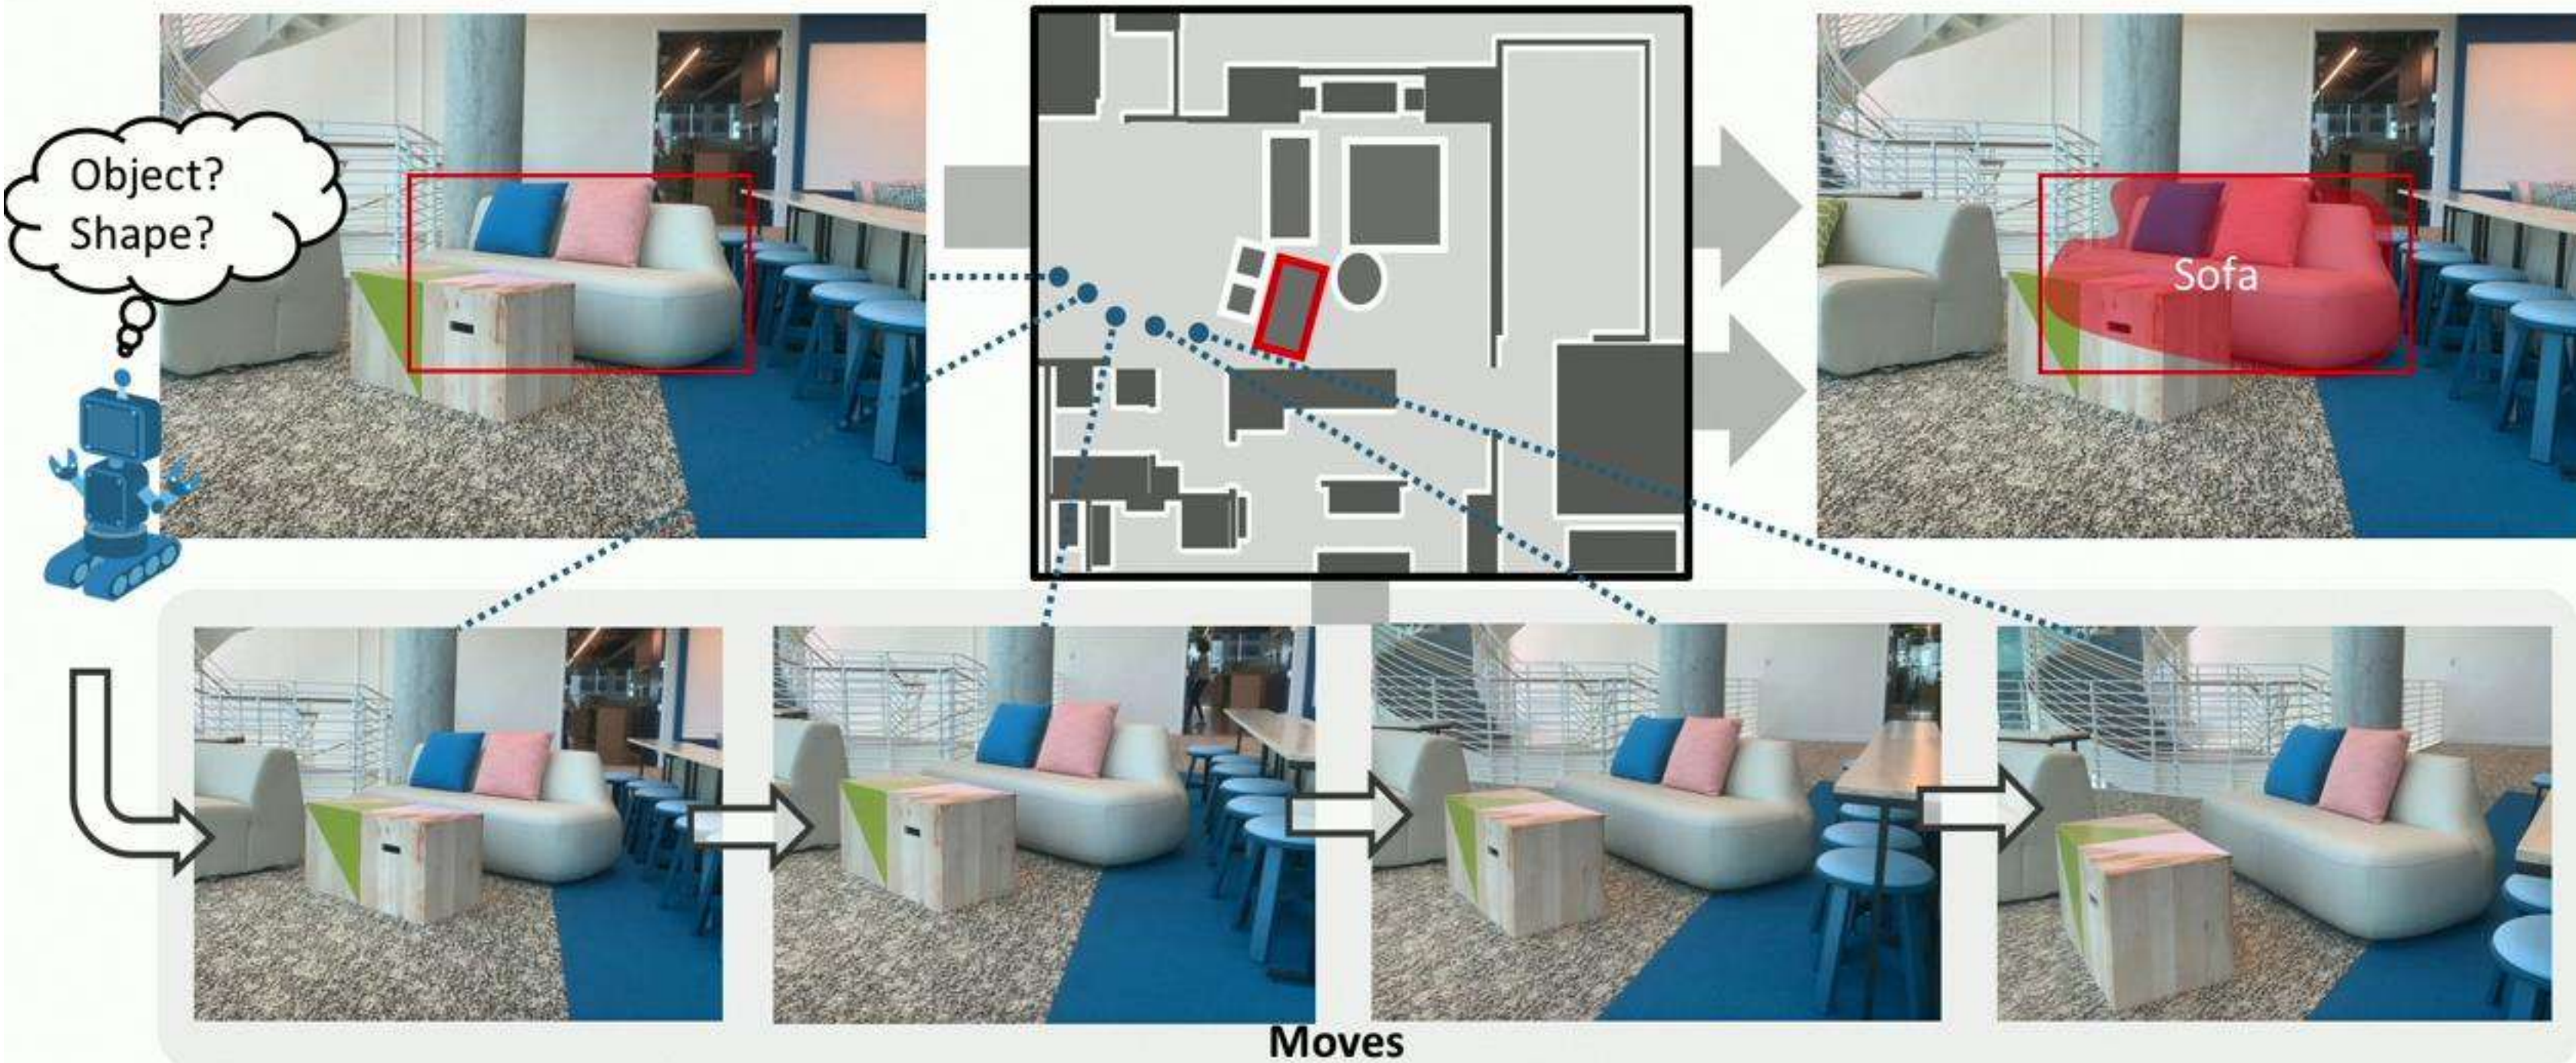


Motivation

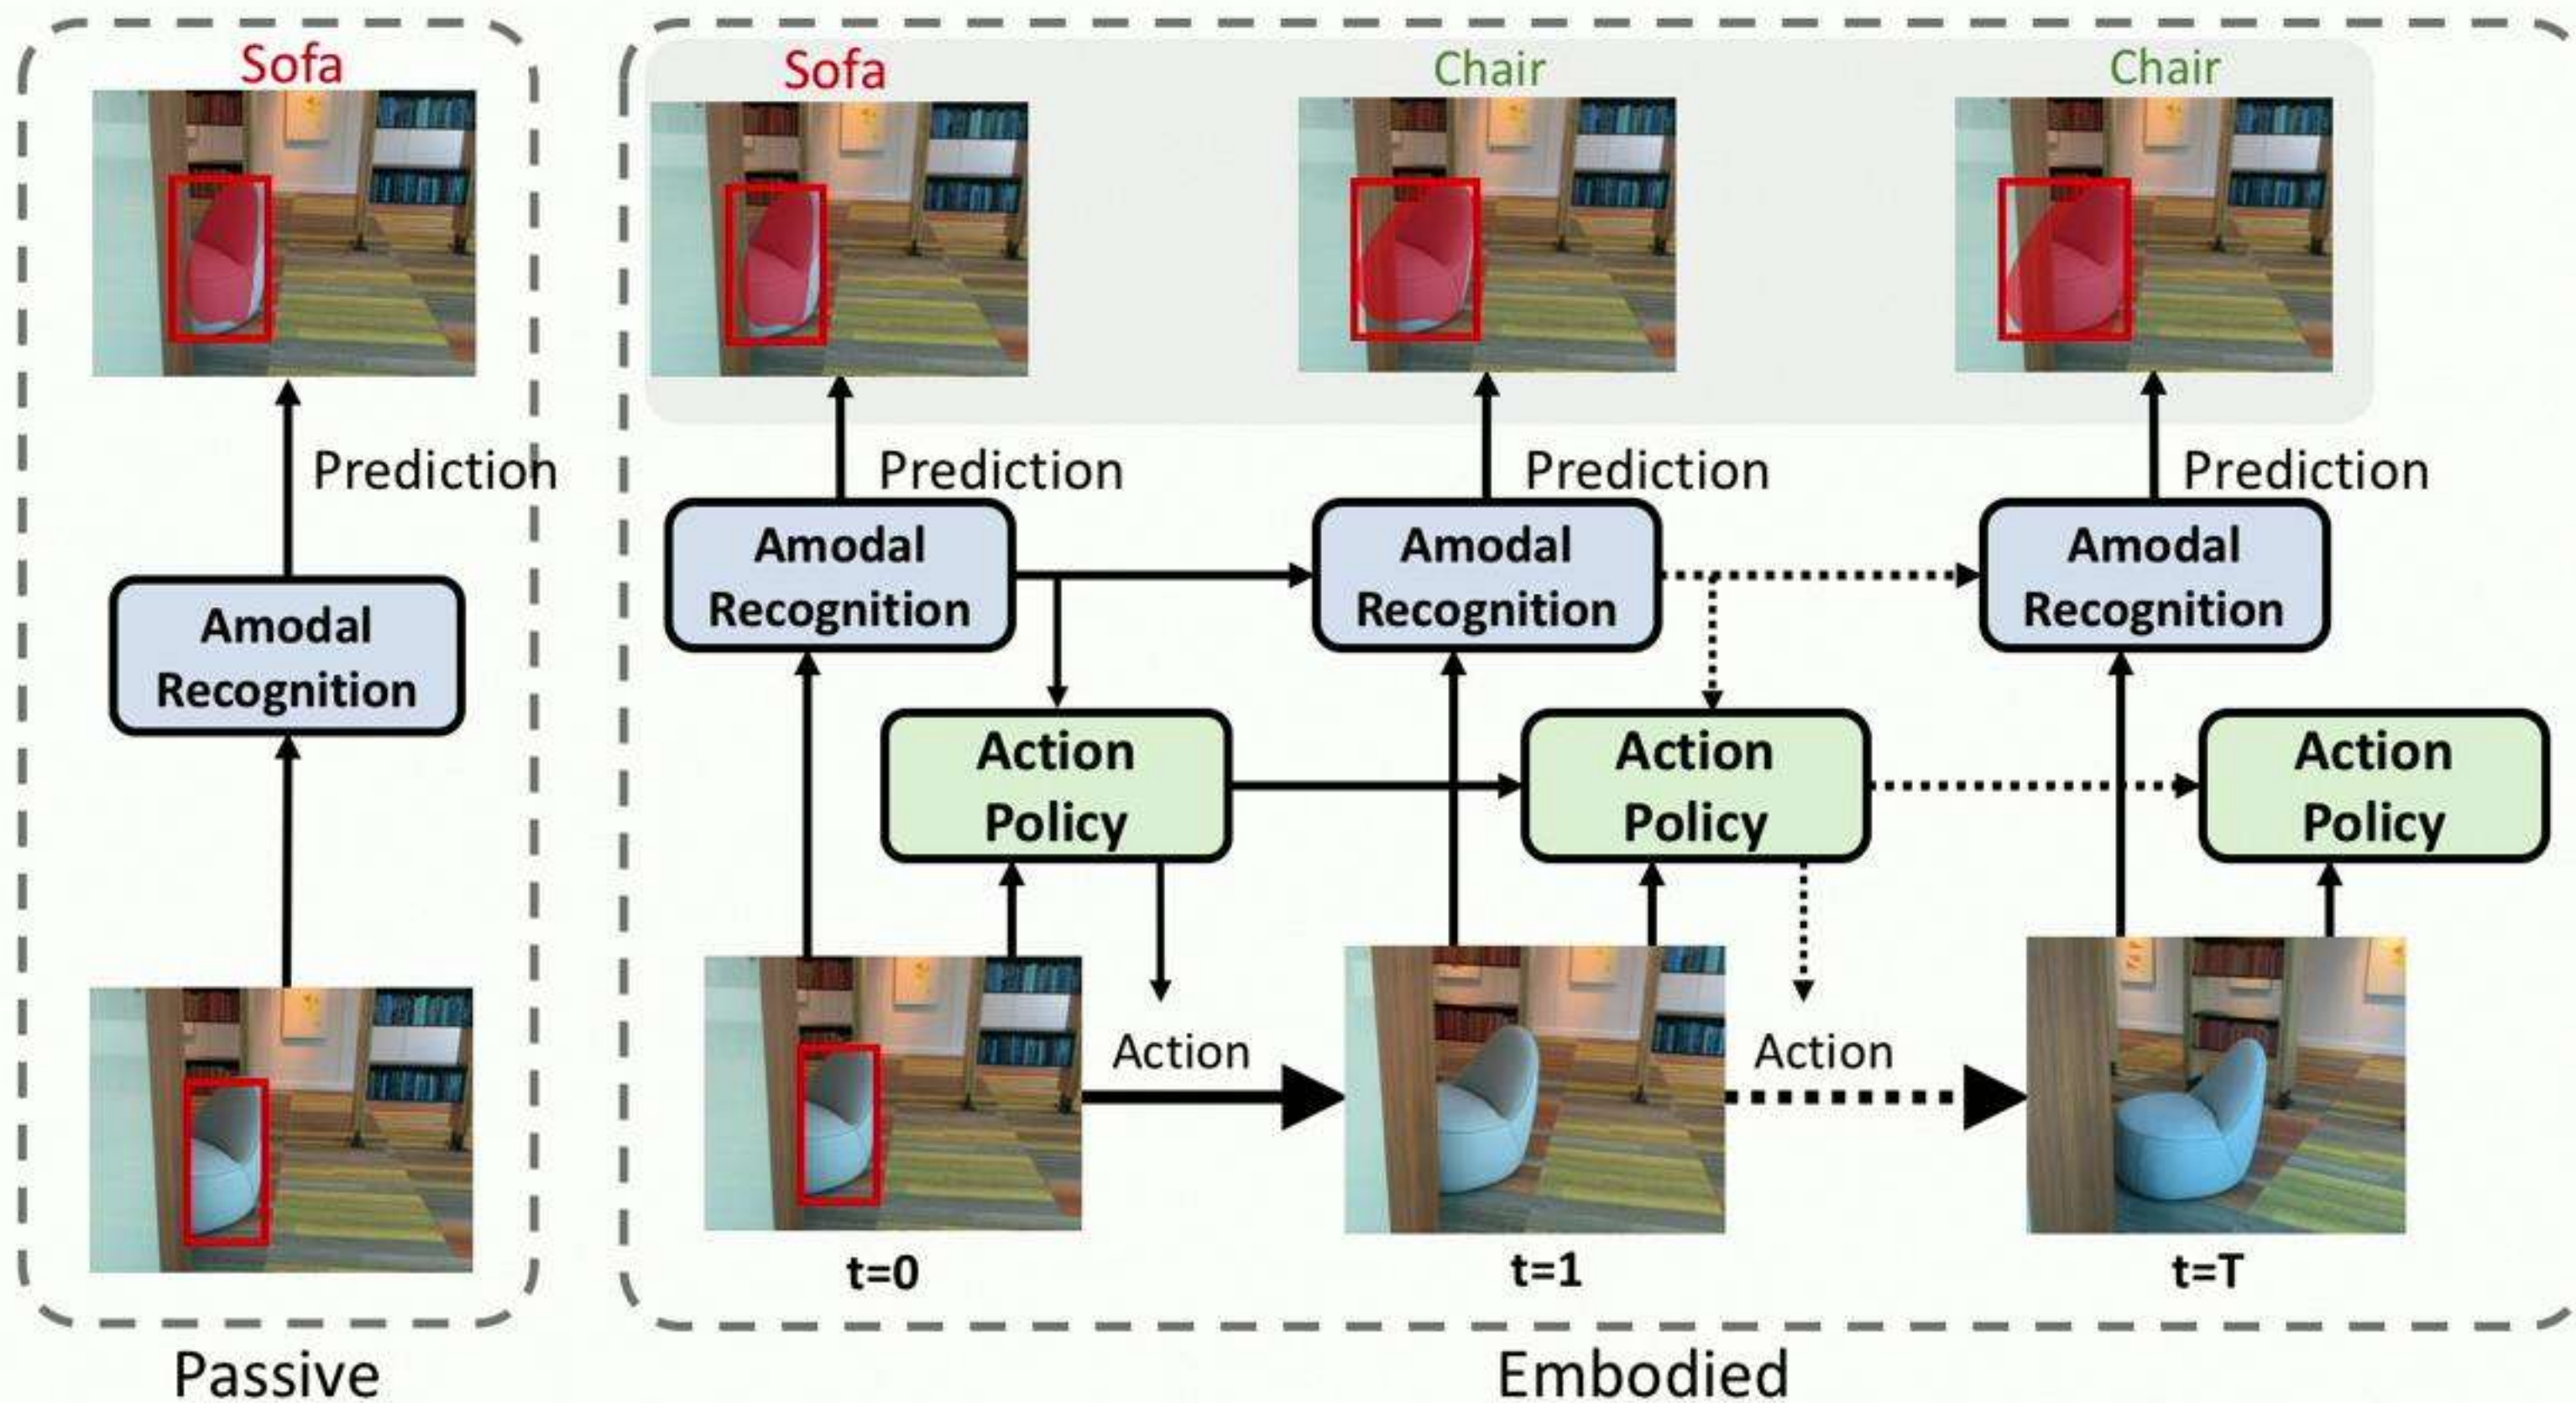
Input

Environment Map

Prediction

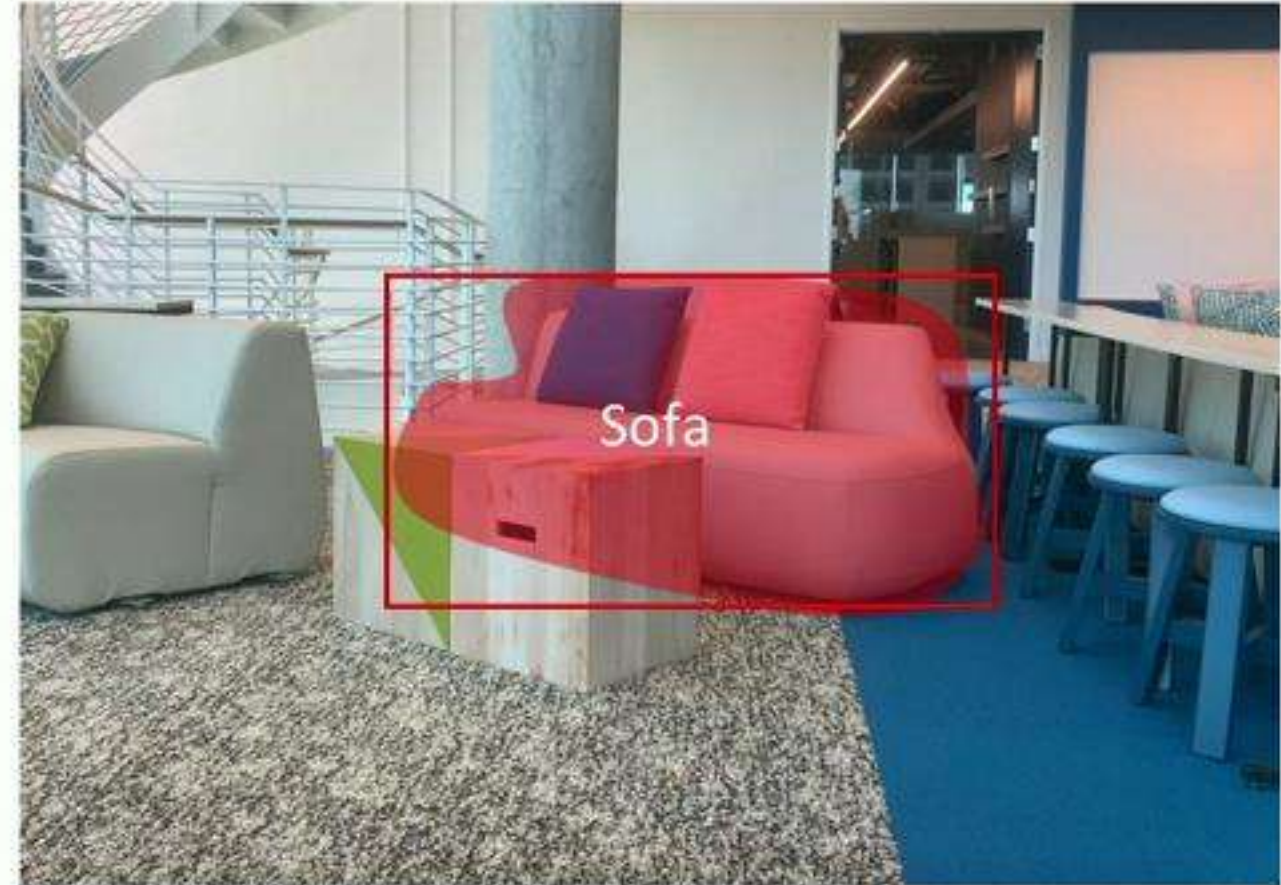


Embodied Amodal Recognition (EAR)

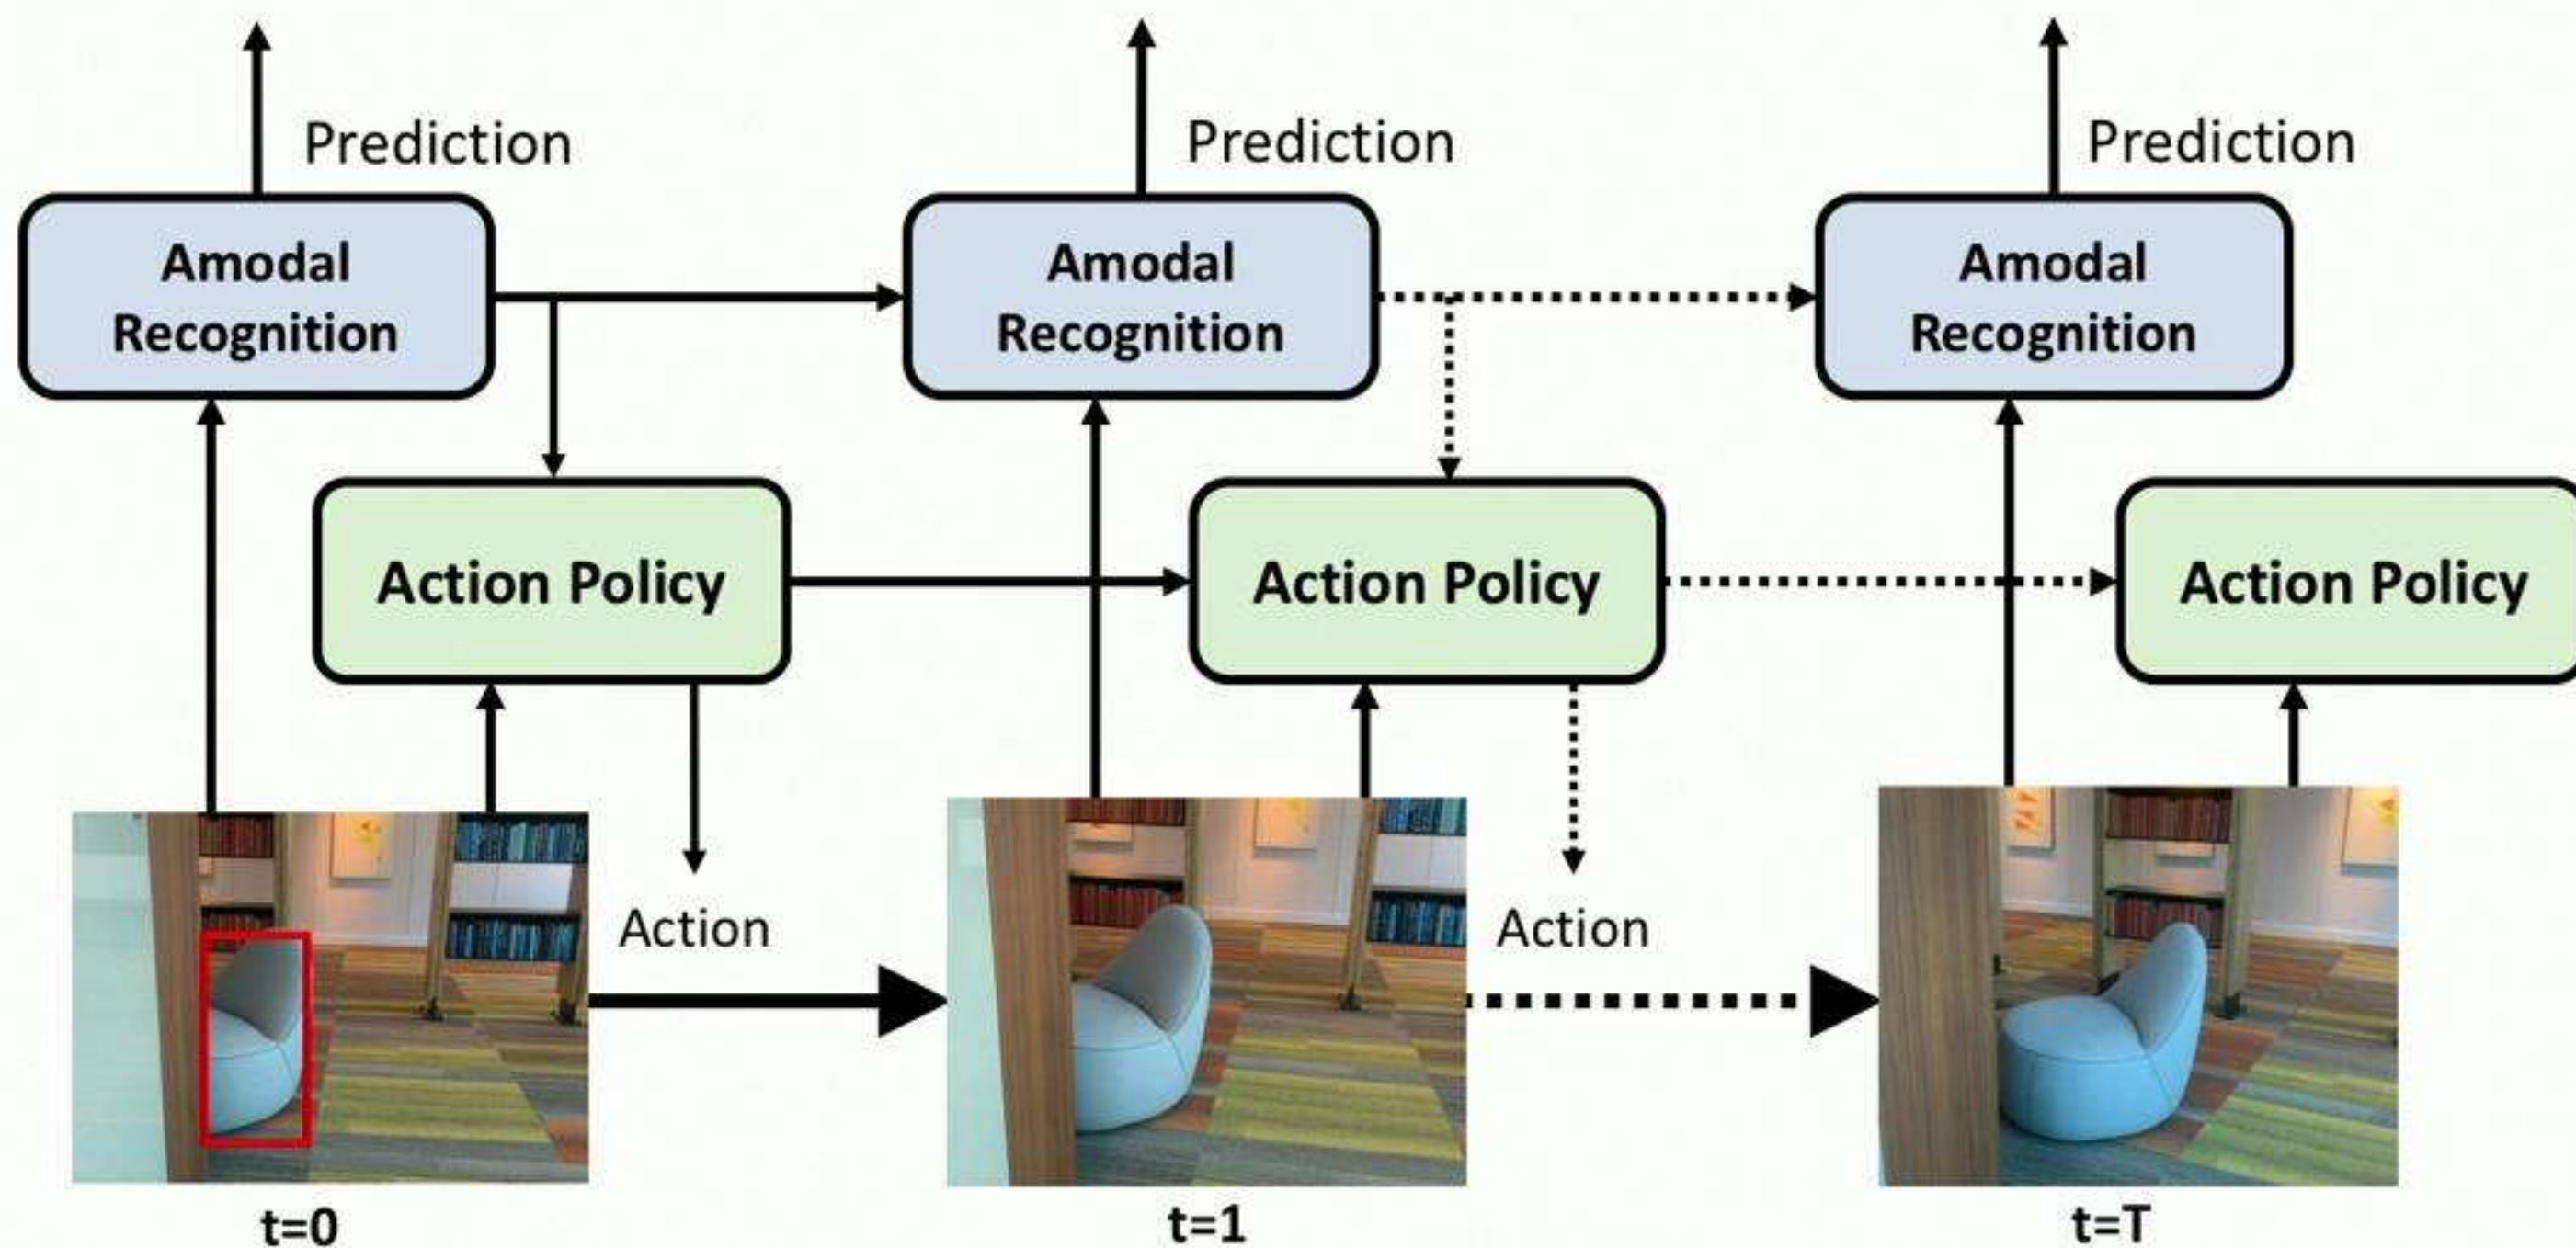


EAR Task

- **Three sub-tasks**
 - Object recognition
 - 2D amodal localization
 - 2D amodal segmentation
- **Single target object**
 - Specify one object as the target
- **Predict for the first frame**



EVR Model



Amodal Recognition

Objective:

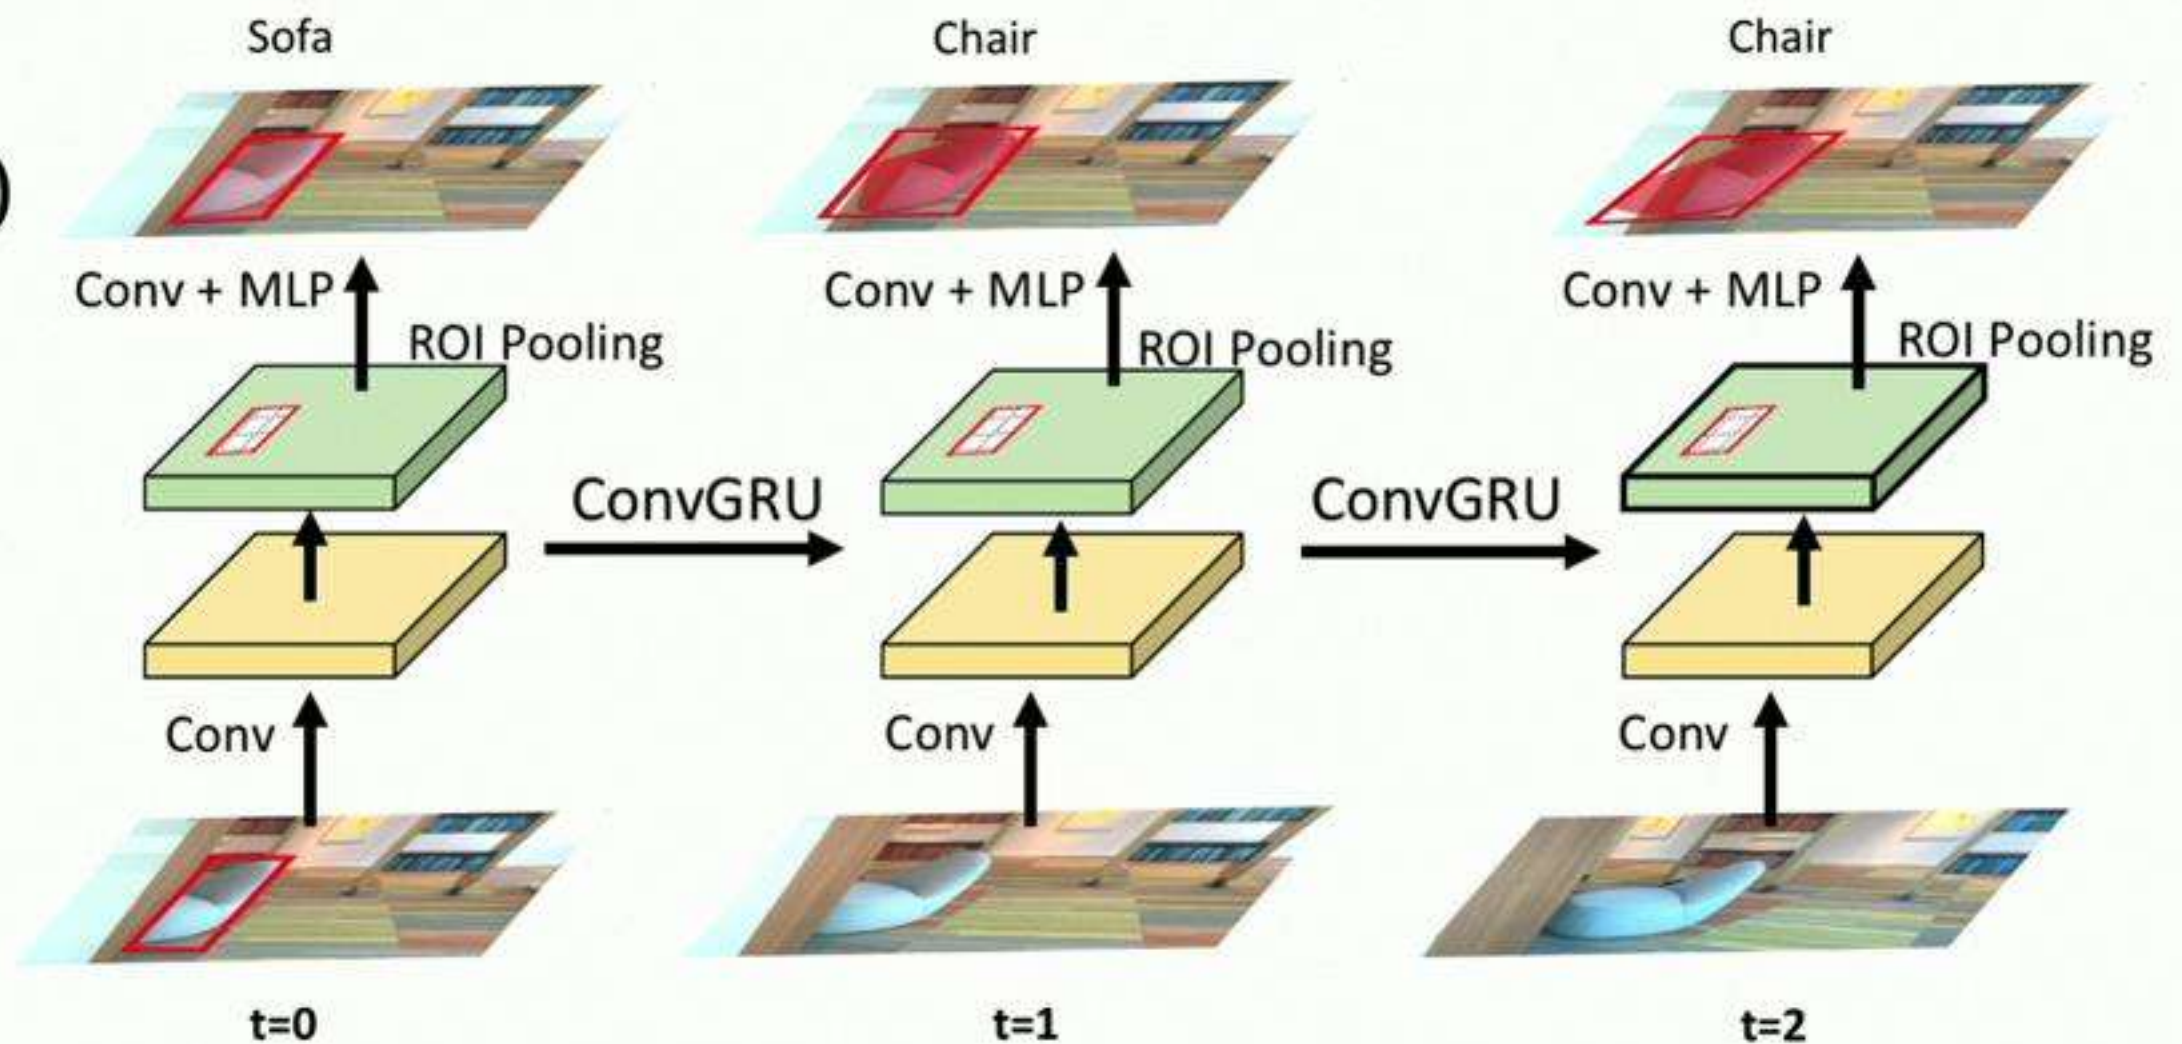
$$y_t = f(\overset{\text{Visible Box}}{b_0}, \underbrace{I_0, I_1, \dots, I_t}_{\text{Observations}})$$

Temporal Aggregation

$$h_t = GRU(x_t, h_{t-1})$$

Three losses:

$$L = L_c + L_b + L_m$$



Learn to Move (Policy)

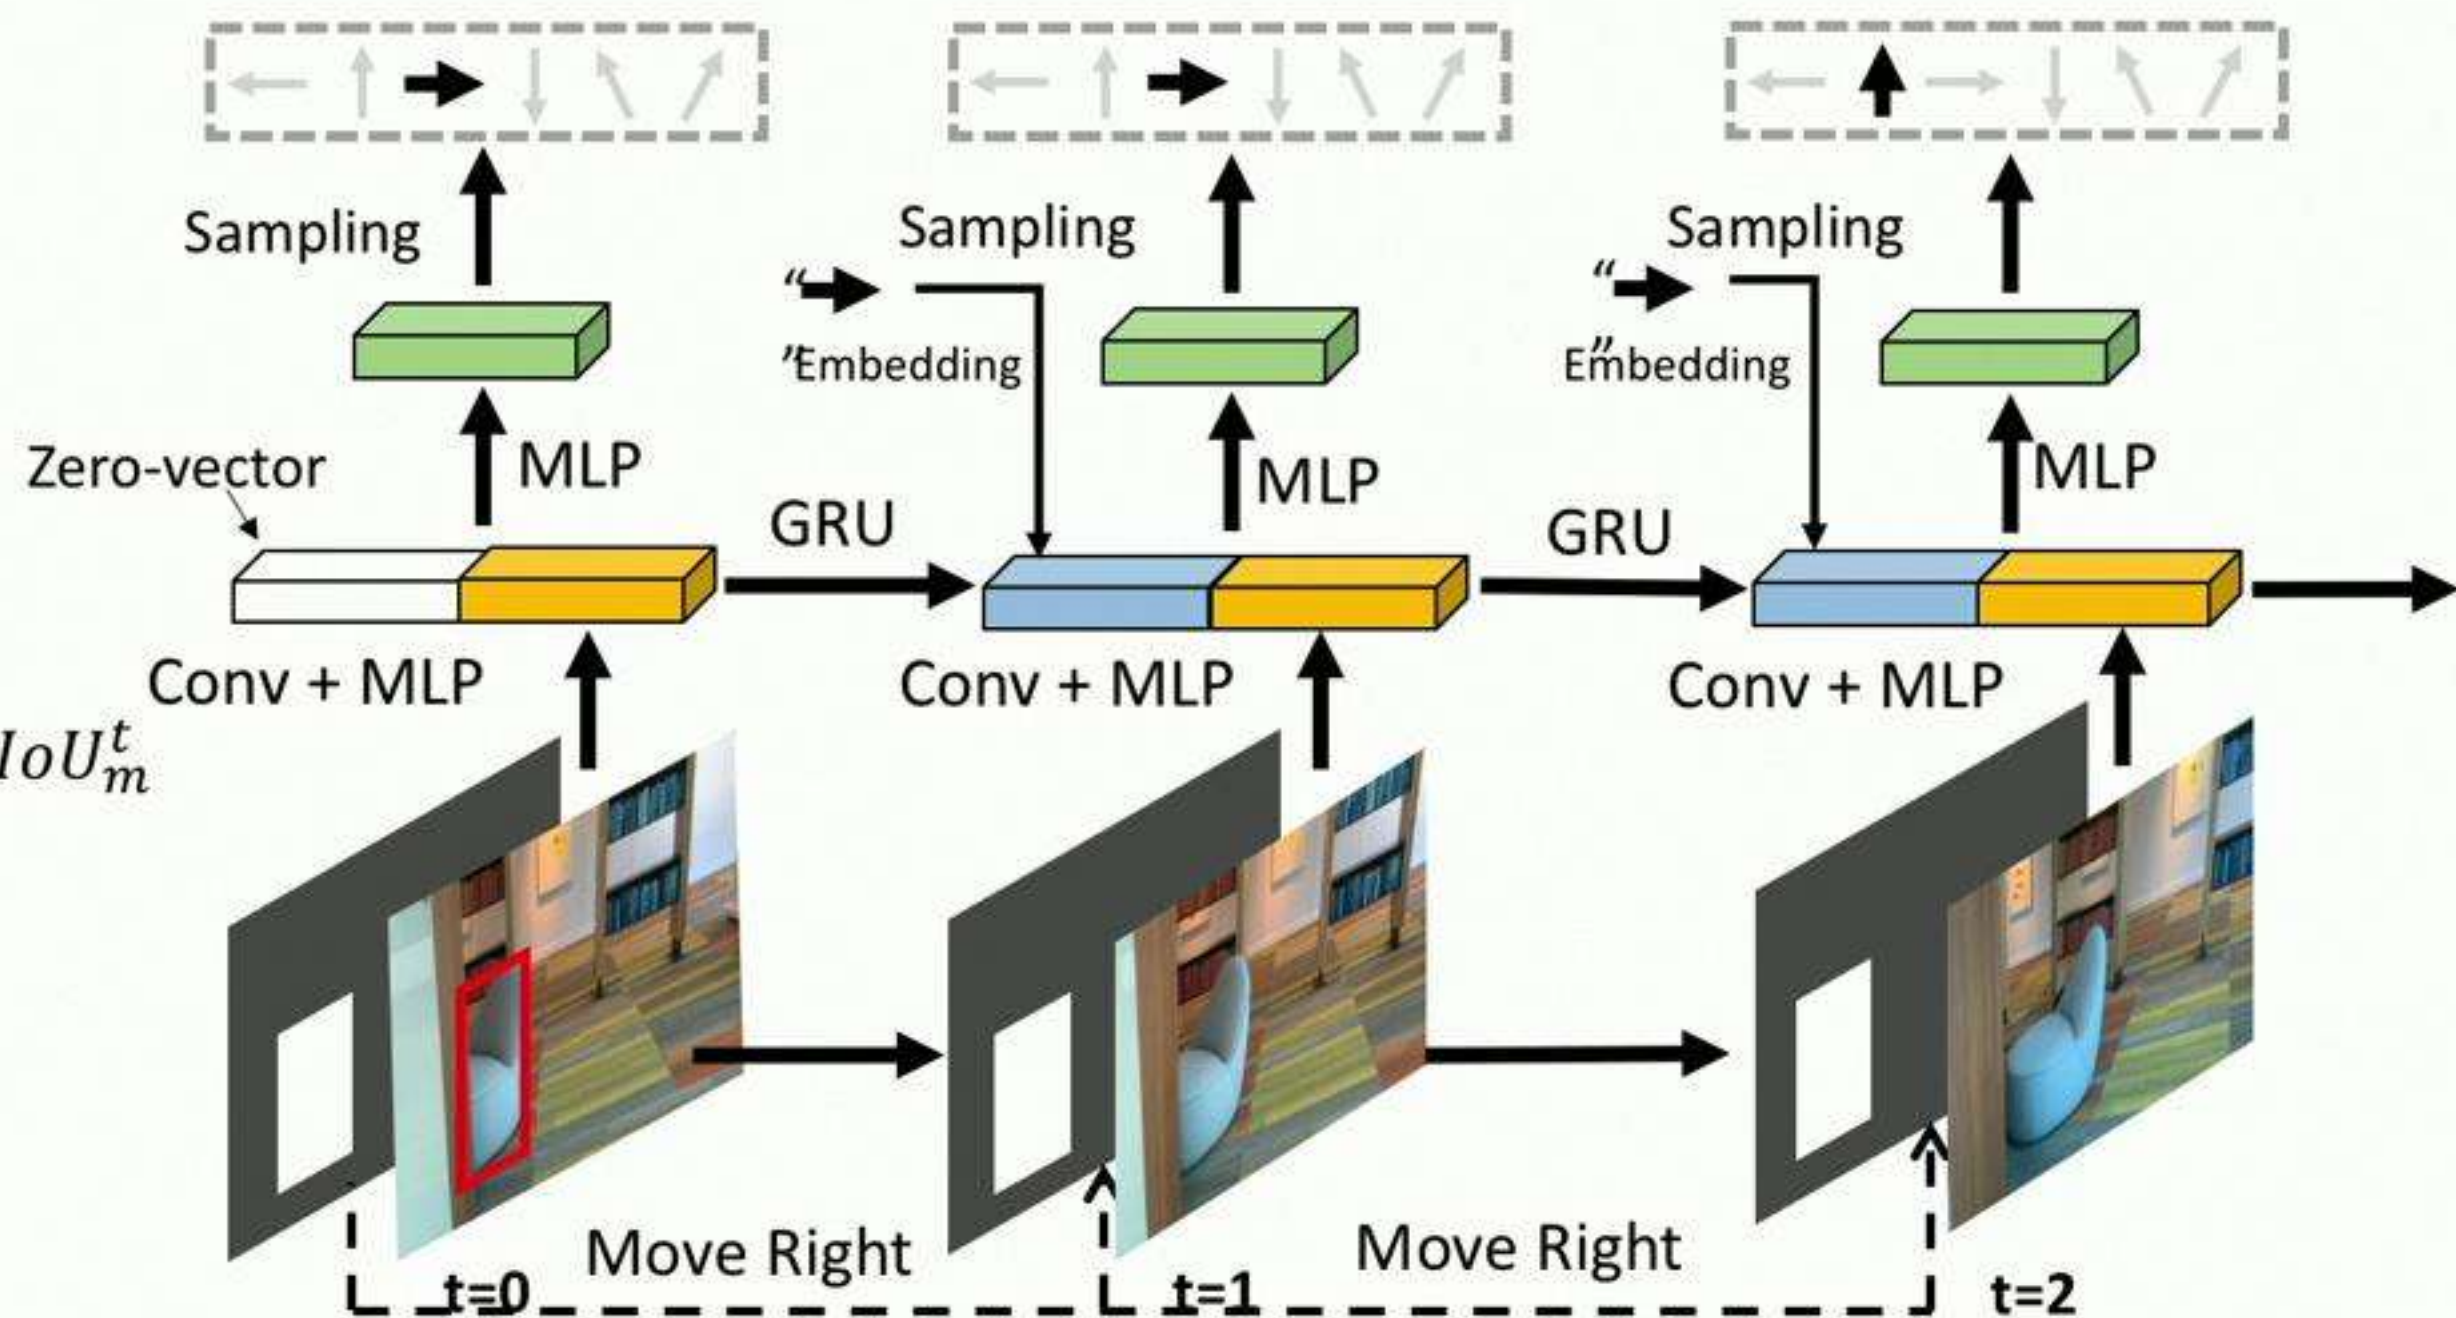
Objective:

$$a_t = \pi(b_0, I_0, I_1, \dots, I_t)$$

Reward:

$$r_t = \lambda_c Acc_c^t + \lambda_b IoU_b^t + \lambda_m IoU_m^t$$

Reward reshaping
+ REINFORCEMENT



Amodal Recognition

Objective:

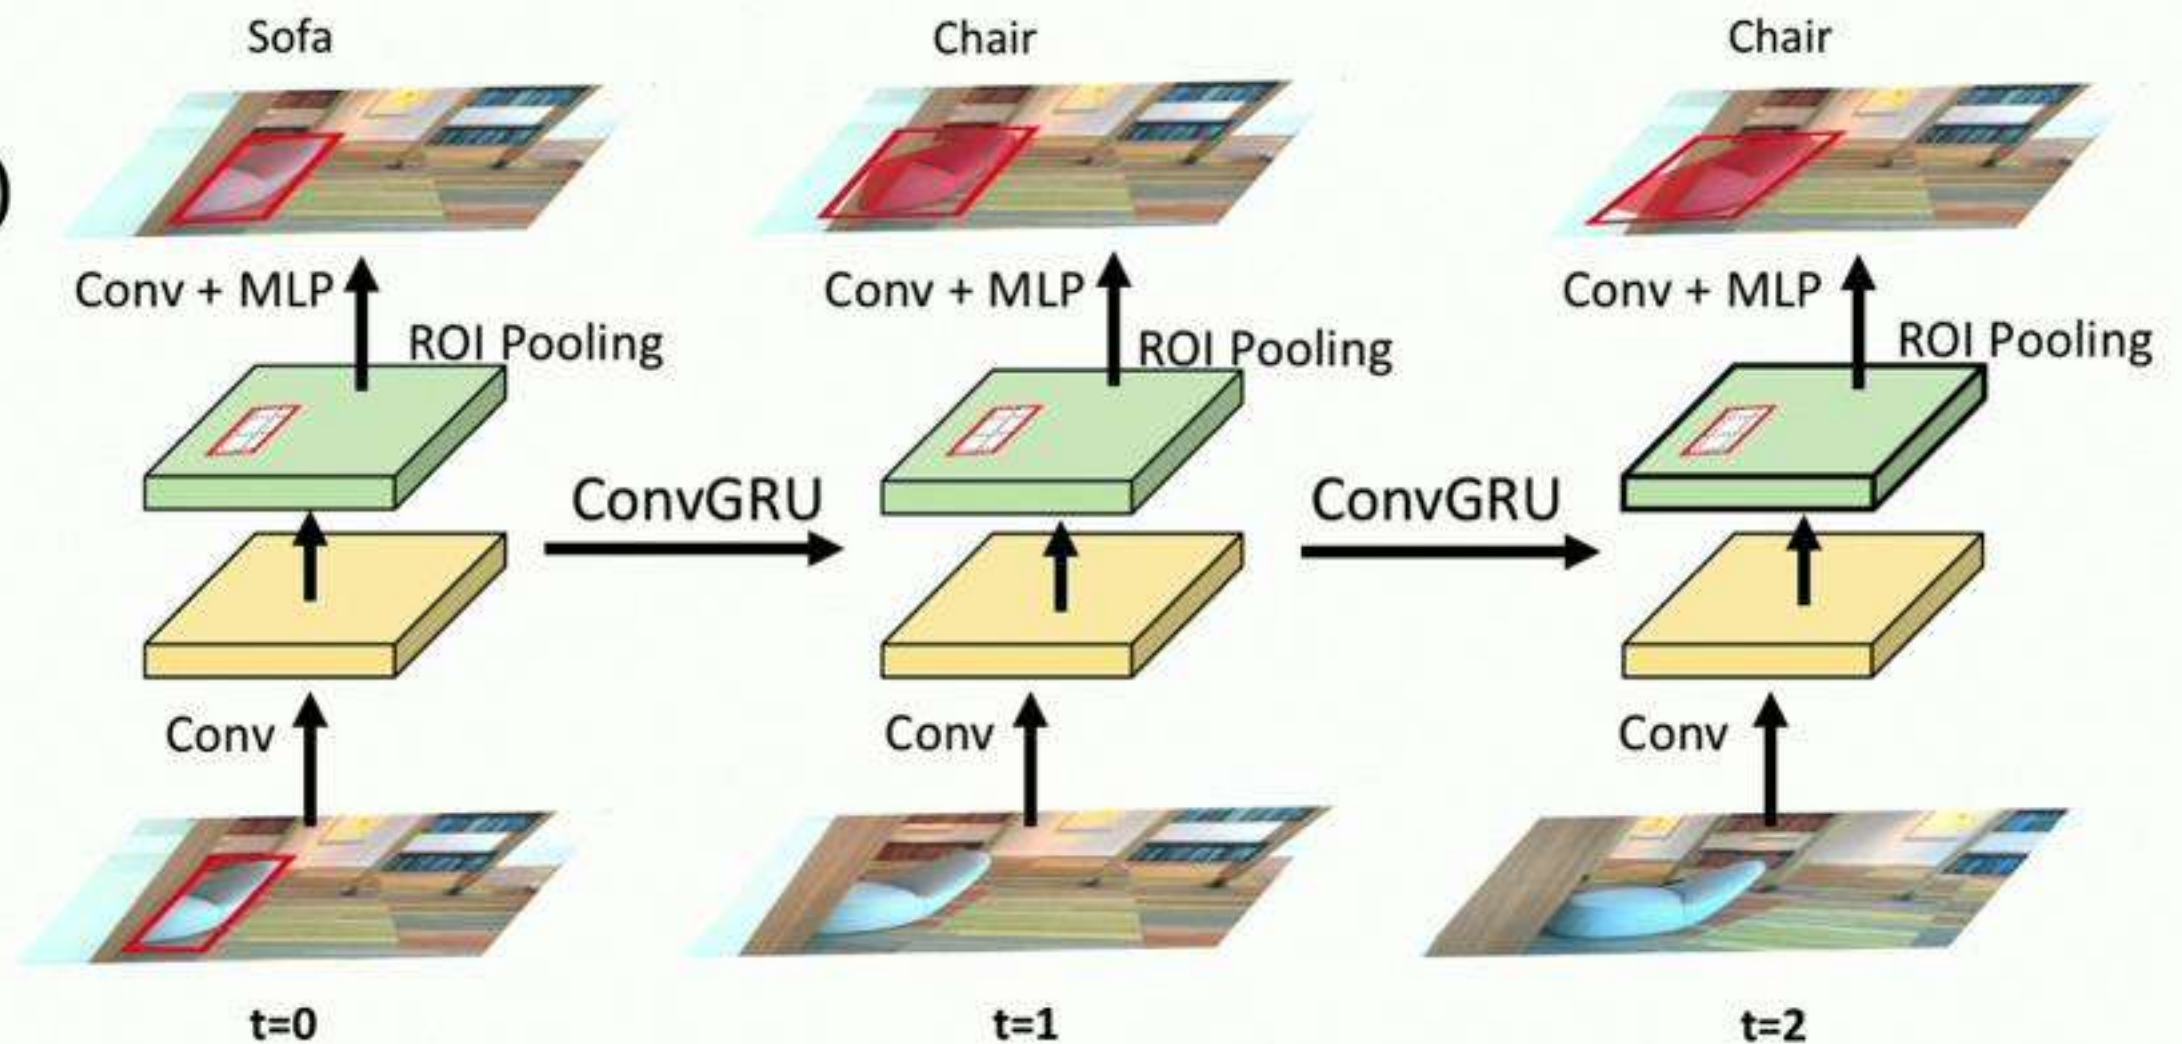
$$y_t = f(\overset{\text{Visible Box}}{b_0}, \underbrace{I_0, I_1, \dots, I_t}_{\text{Observations}})$$

Temporal Aggregation

$$h_t = GRU(x_t, h_{t-1})$$

Three losses:

$$L = L_c + L_b + L_m$$



Learn to Move (Policy)

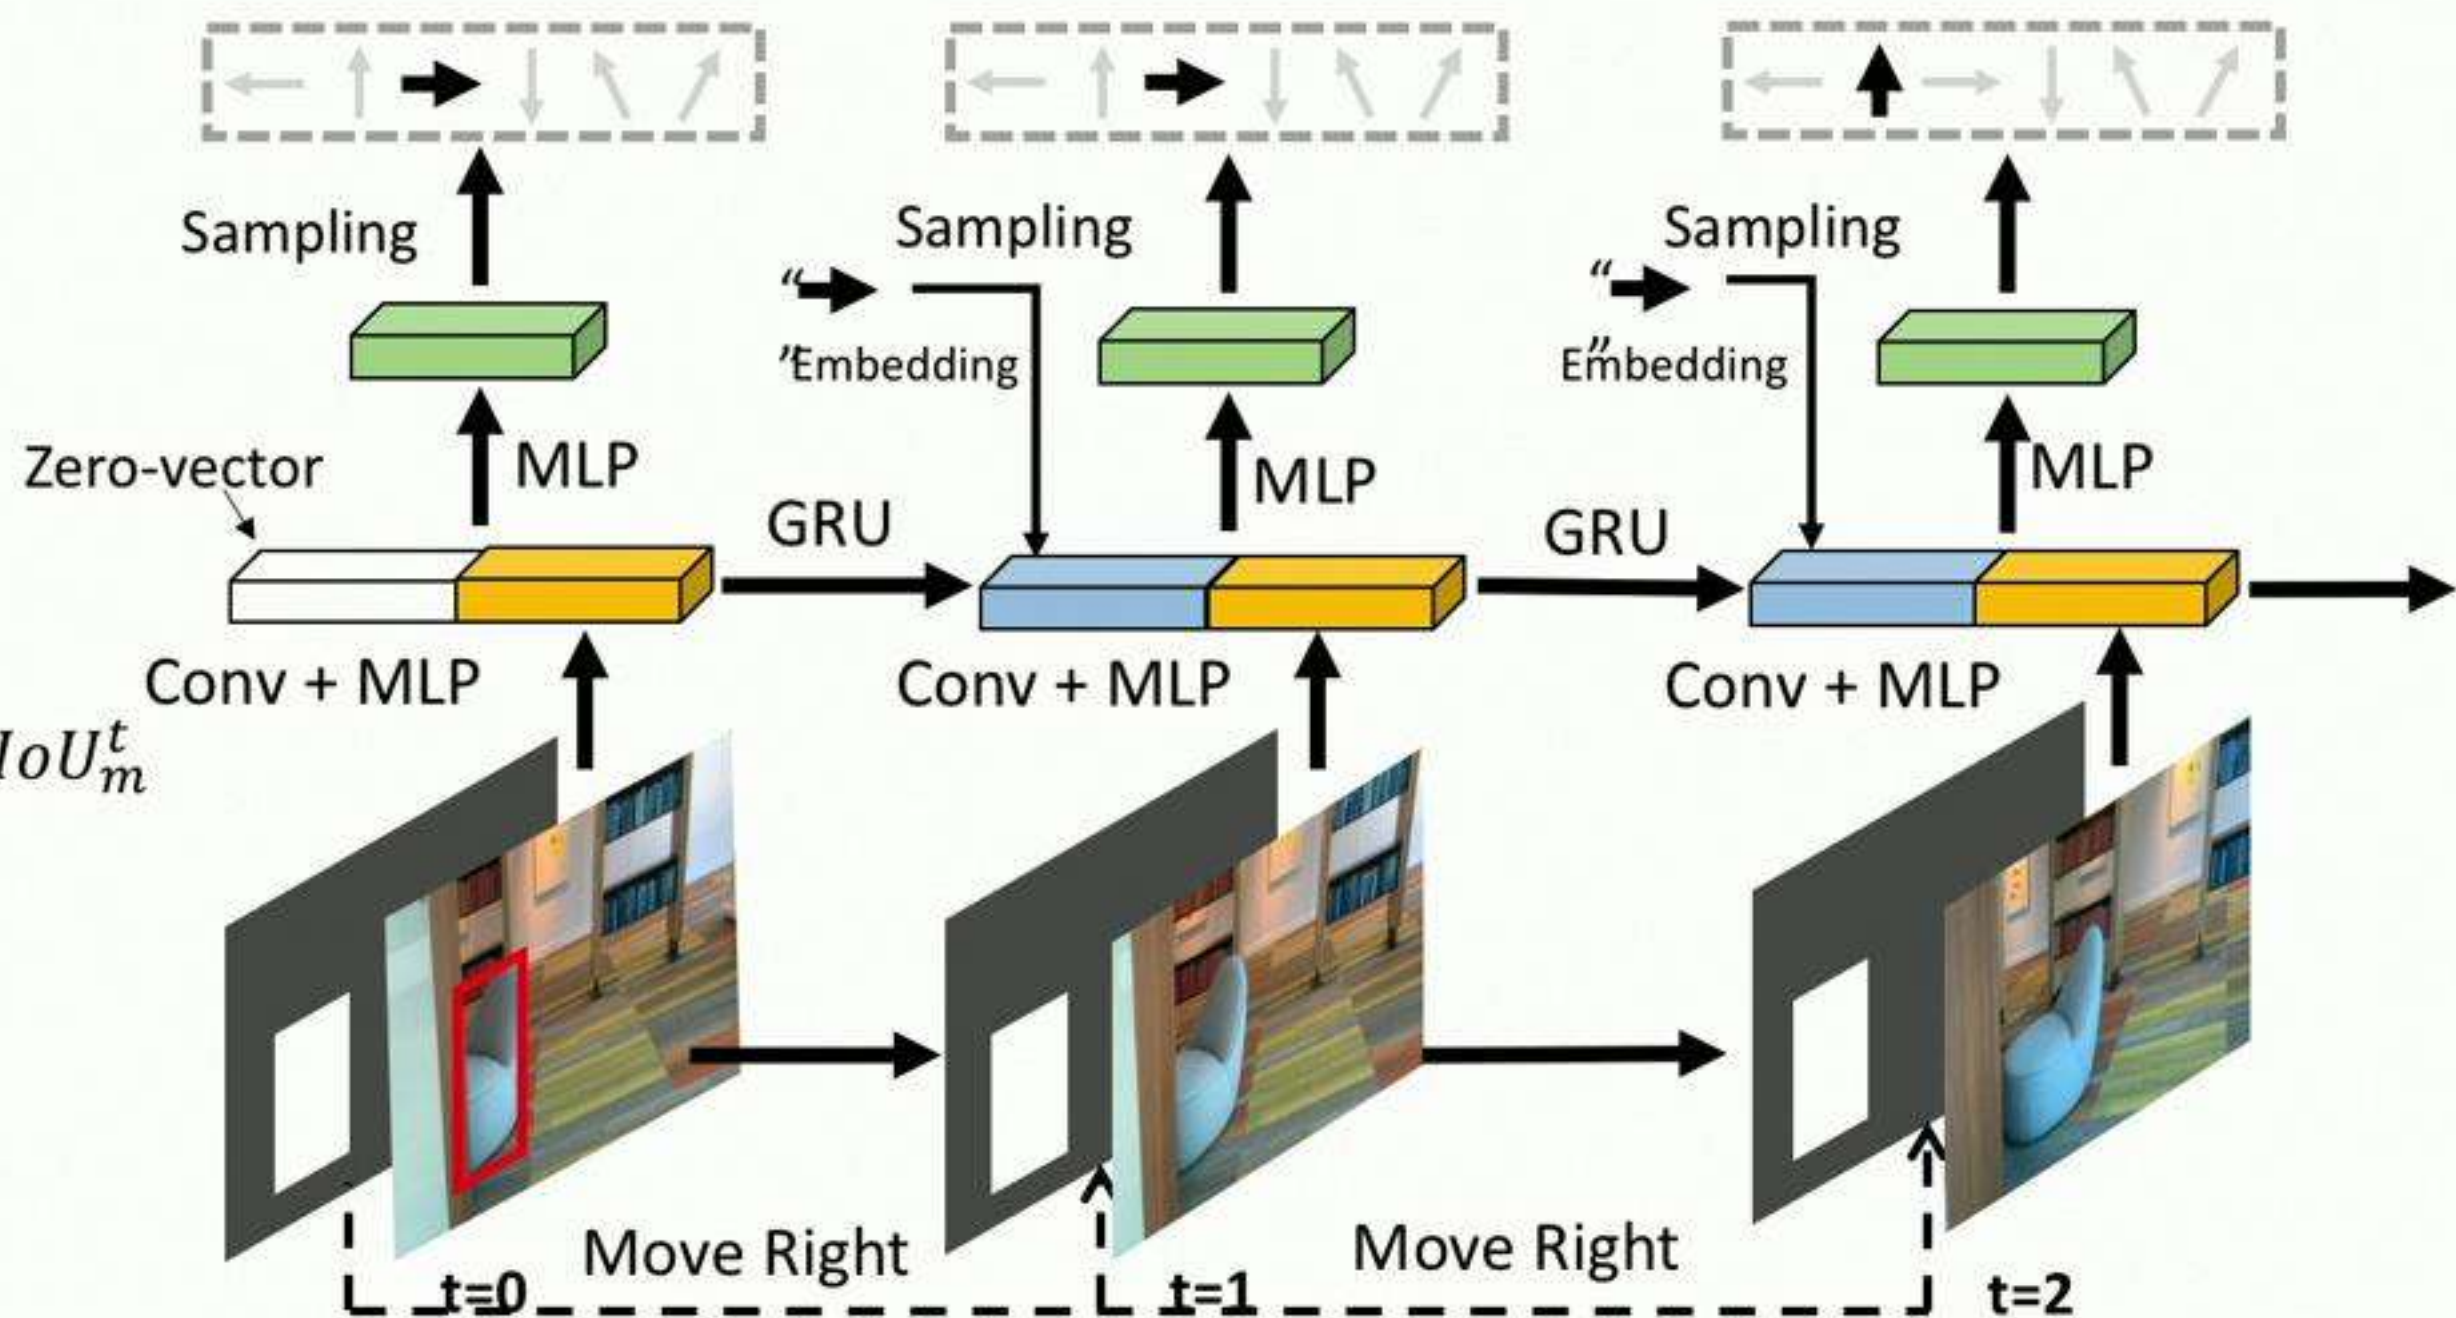
Objective:

$$a_t = \pi(b_0, I_0, I_1, \dots, I_t)$$

Reward:

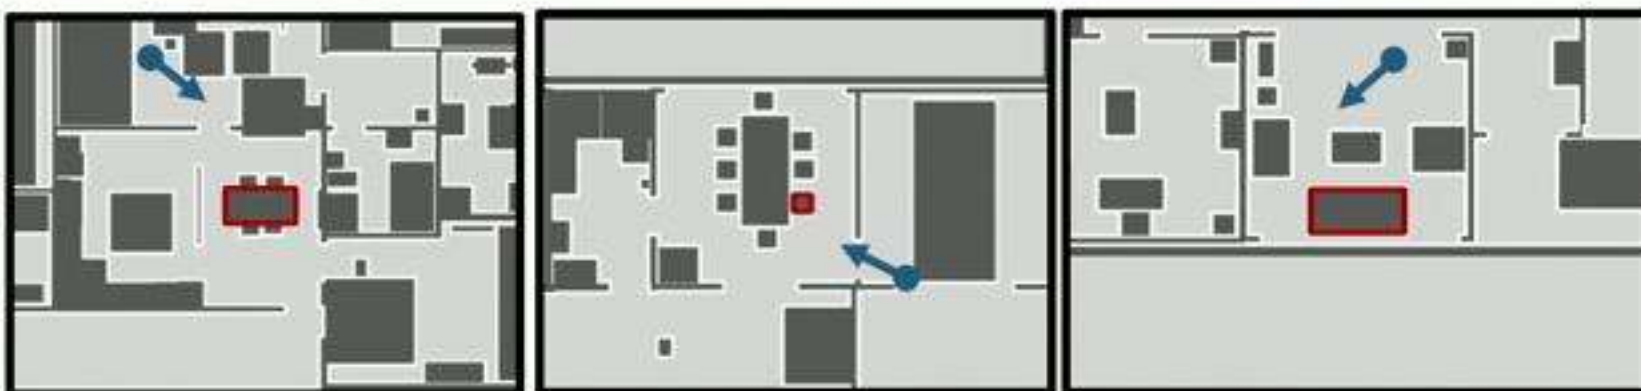
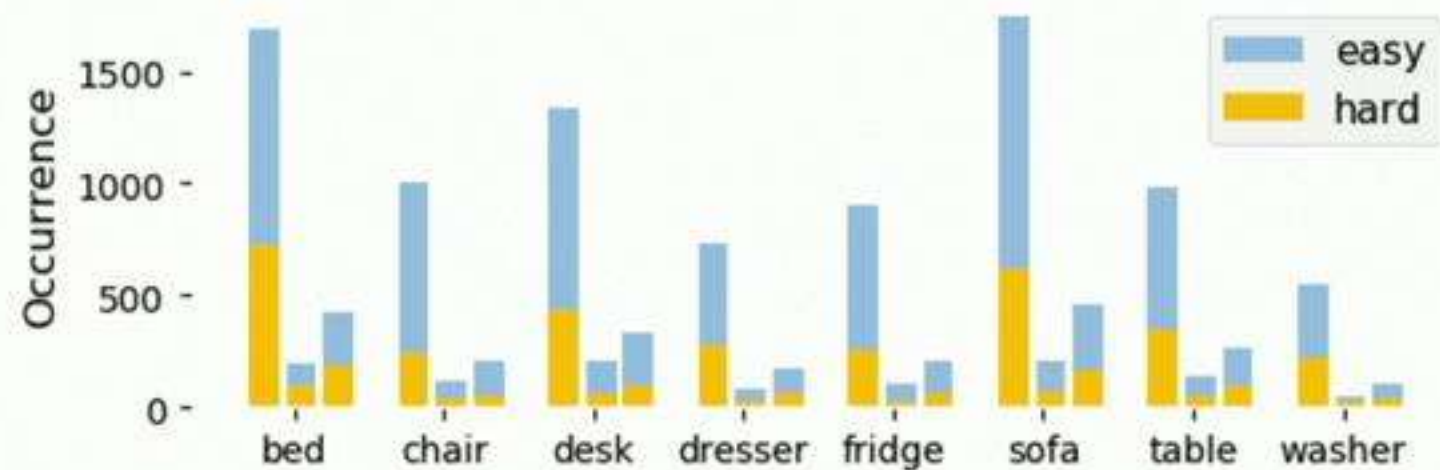
$$r_t = \lambda_c Acc_c^t + \lambda_b IoU_b^t + \lambda_m IoU_m^t$$

Reward reshaping
+ REINFORCEMENT

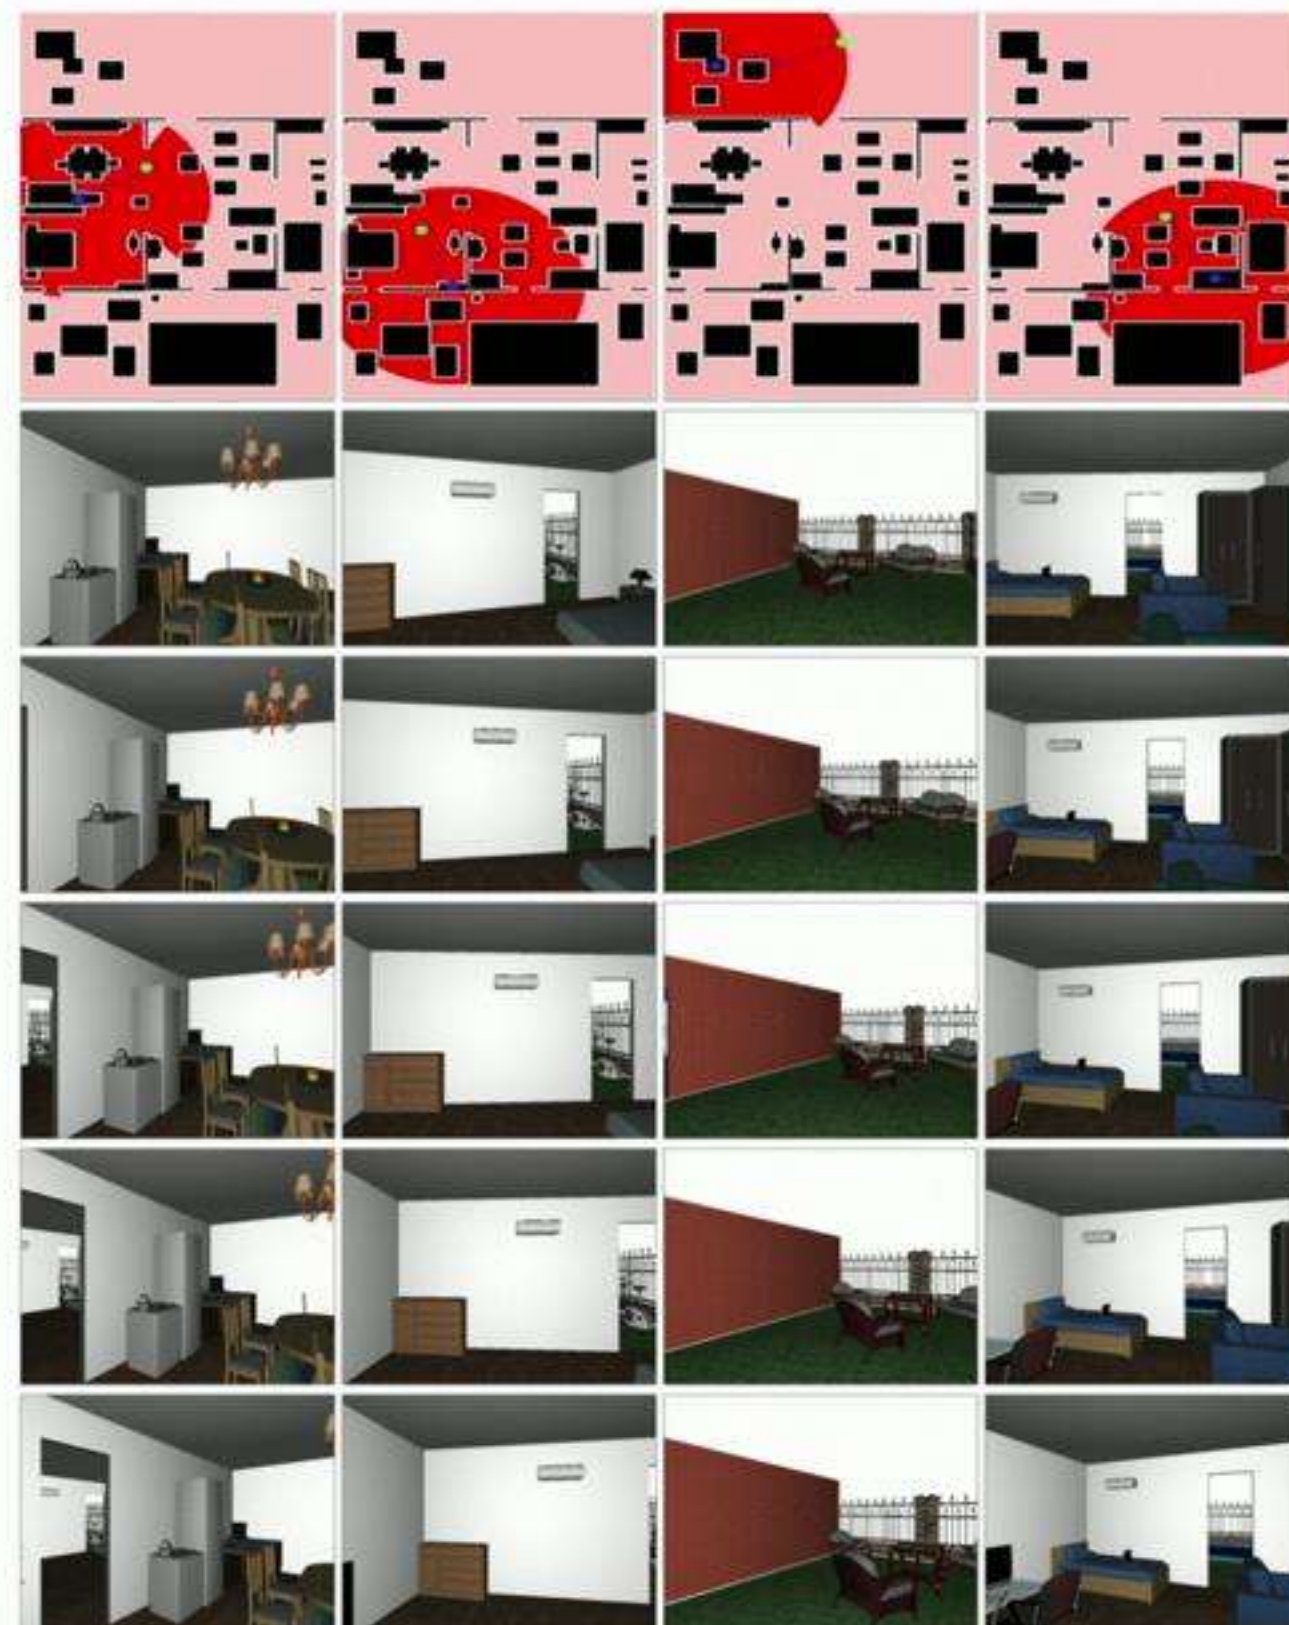


Dataset

Object Categories



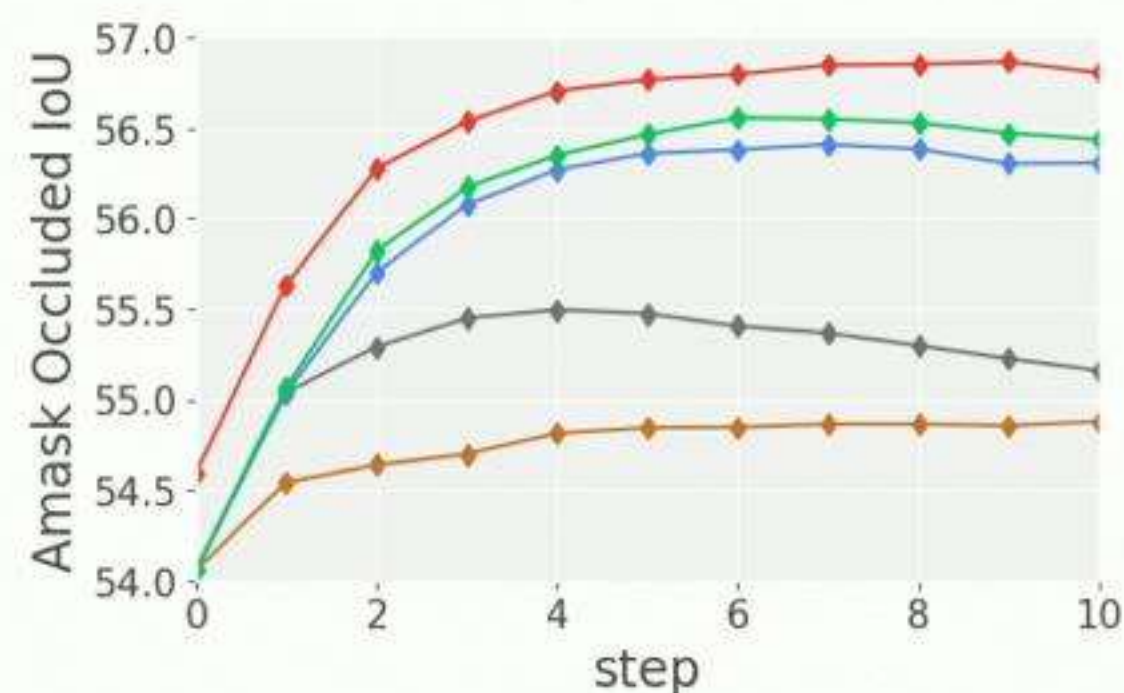
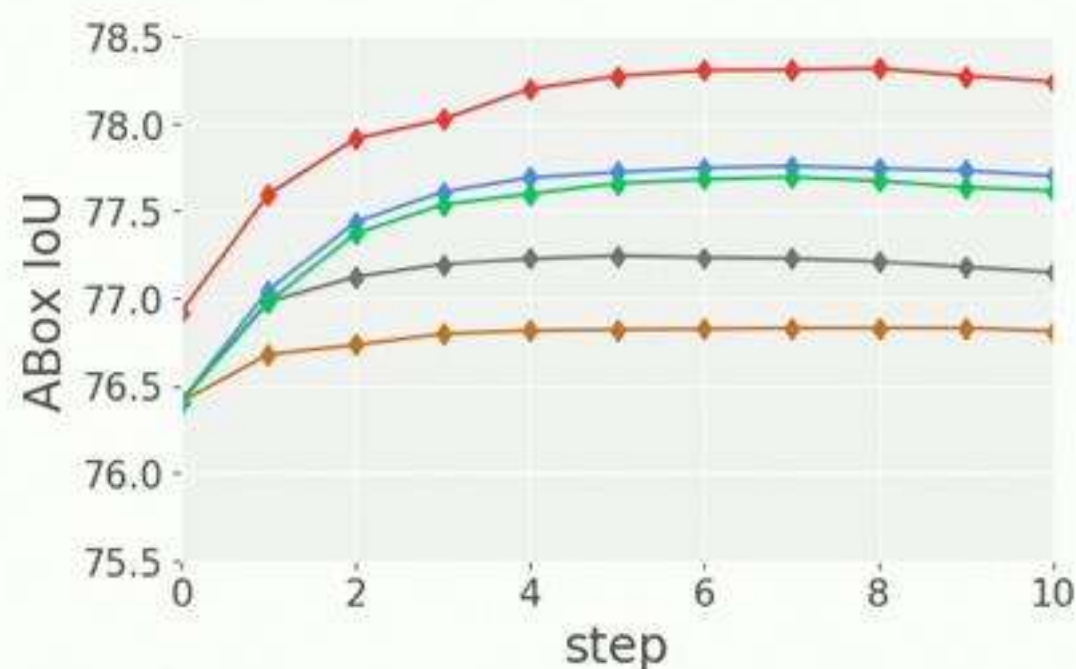
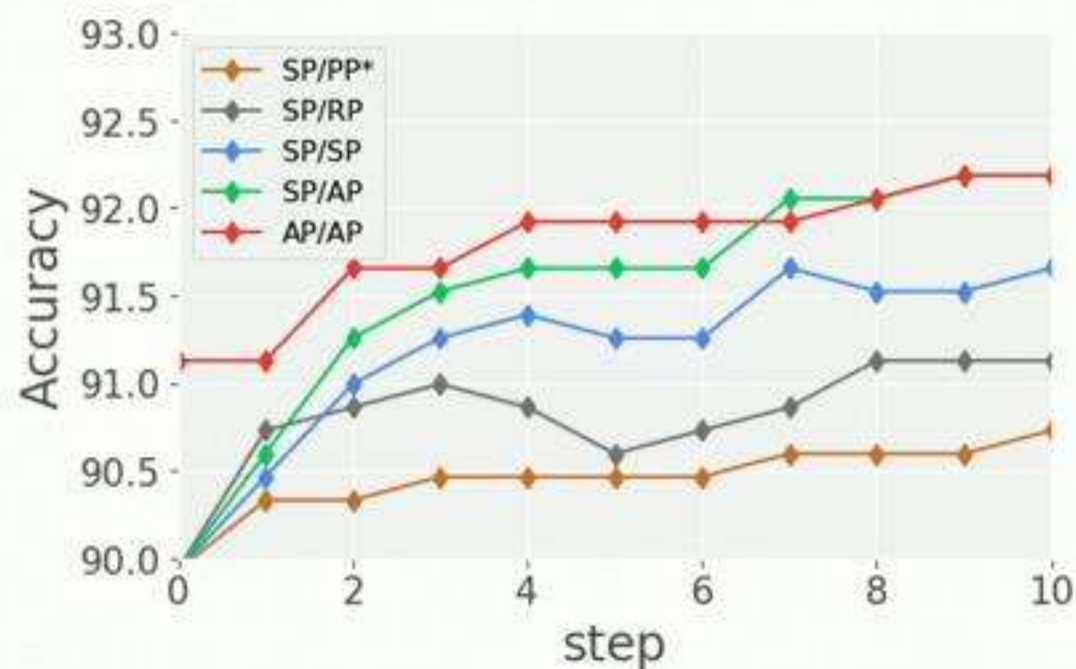
Shortest path toward target



Training

- Stage-wised Training:
 - First train amodal visual recognition with shortest path
 - Then fix amodal visual recognition module, train policy network
 - Afterwards, train amodal visual recognition with learned path

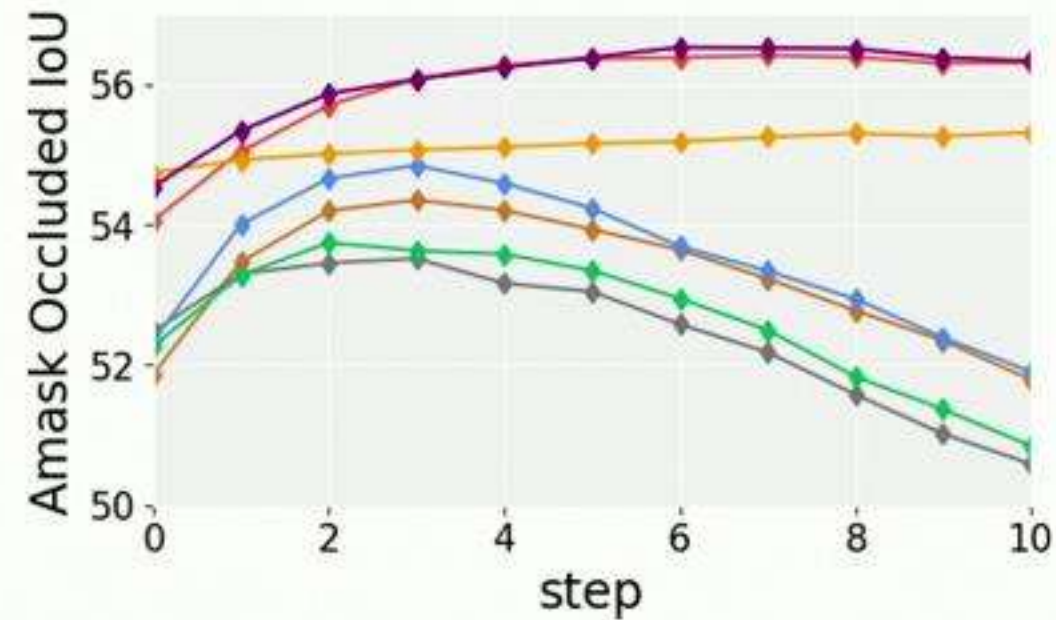
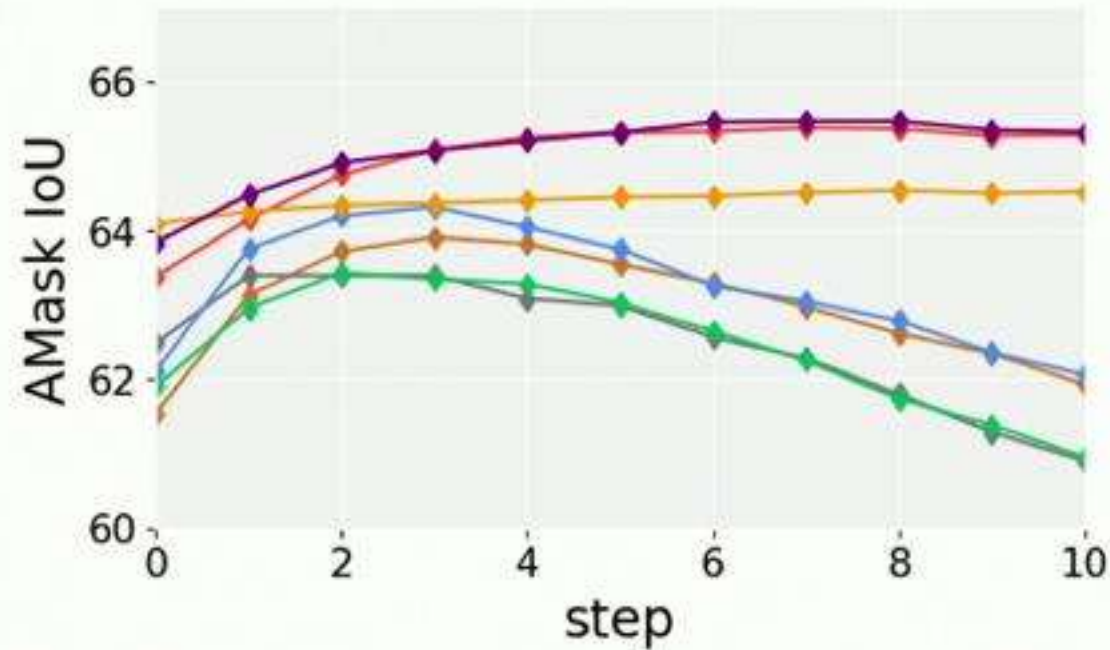
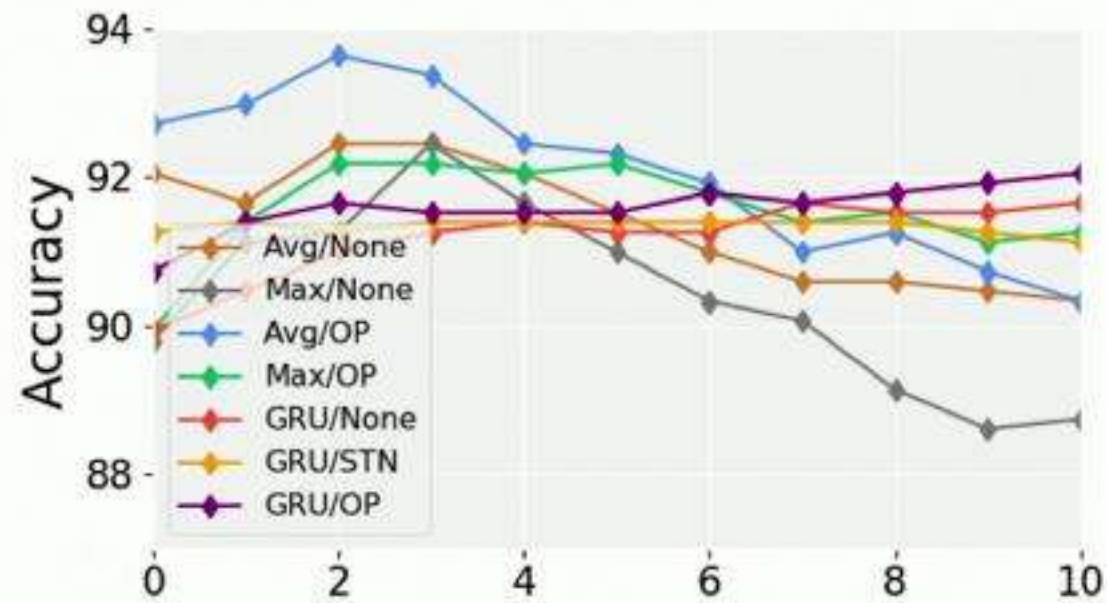
Results



1. Embodiment helps to improve amodal visual recognition performance
2. Our learned moving strategy for agent outperforms other moving strategy and also static visual system.
3. Amodal recognition performance tends to saturate at the end of moving

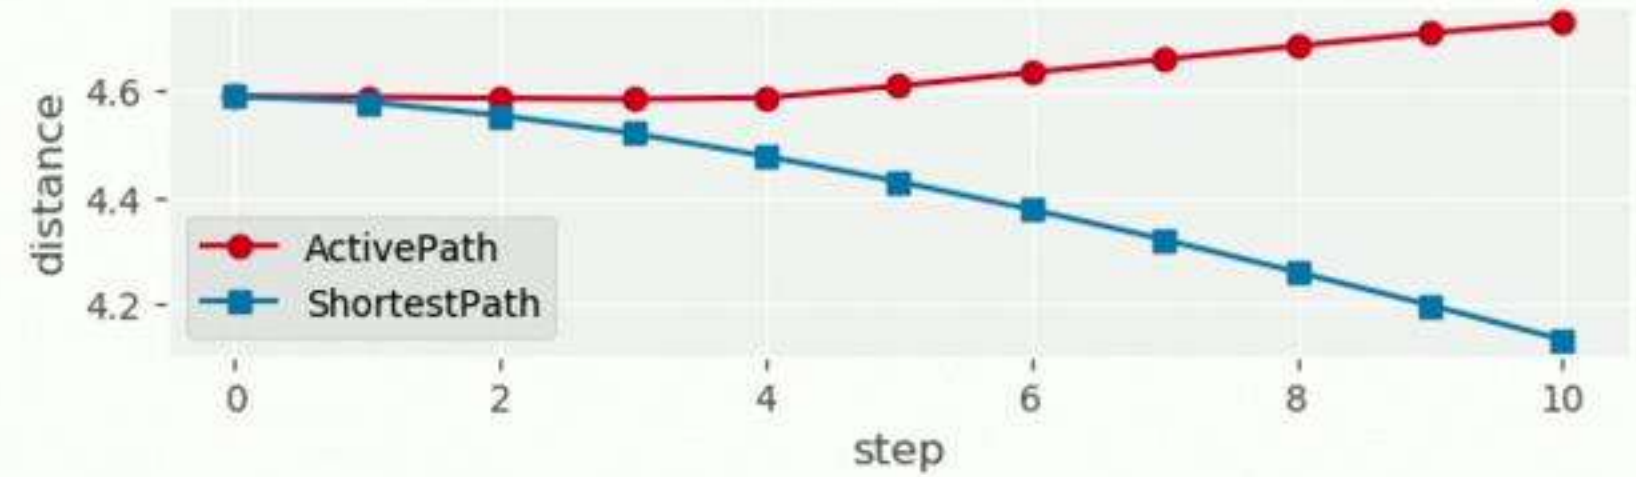
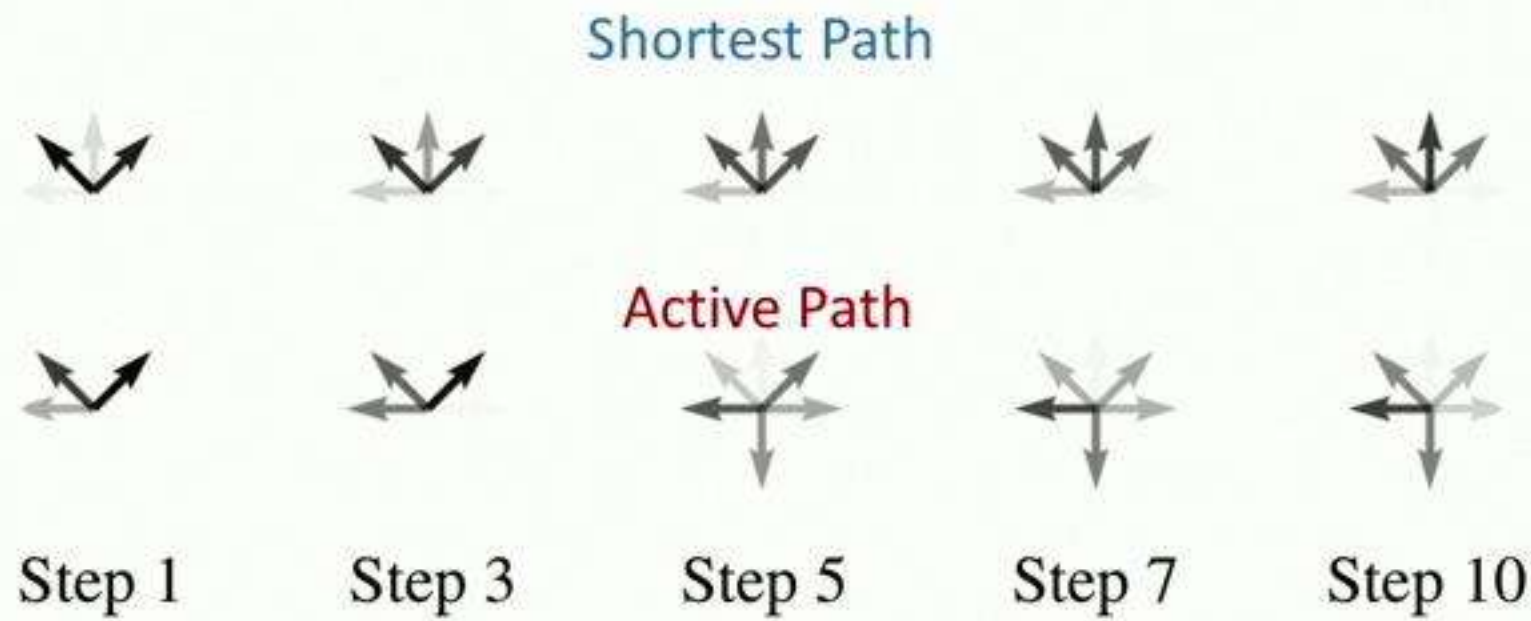
Ablated Study

Different feature aggregation and feature warping methods:



1. GRU works much better than average or max pooling for aggregation
2. Warping feature using optical flow helps to improve the performance
3. Combine GRU and optical flow works slightly better than either

Learned Actions

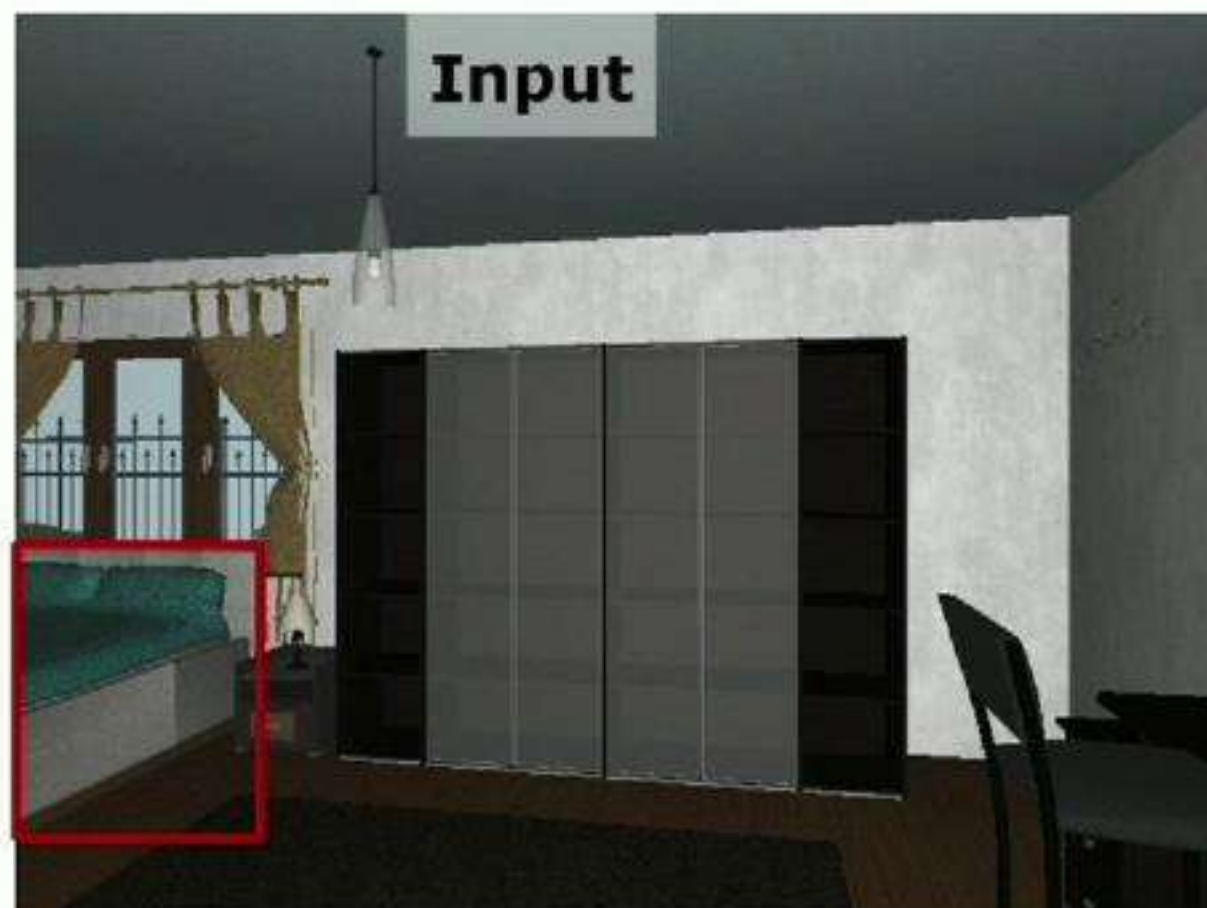


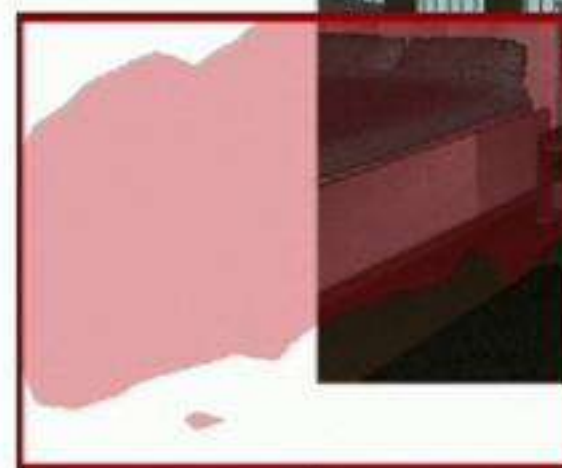
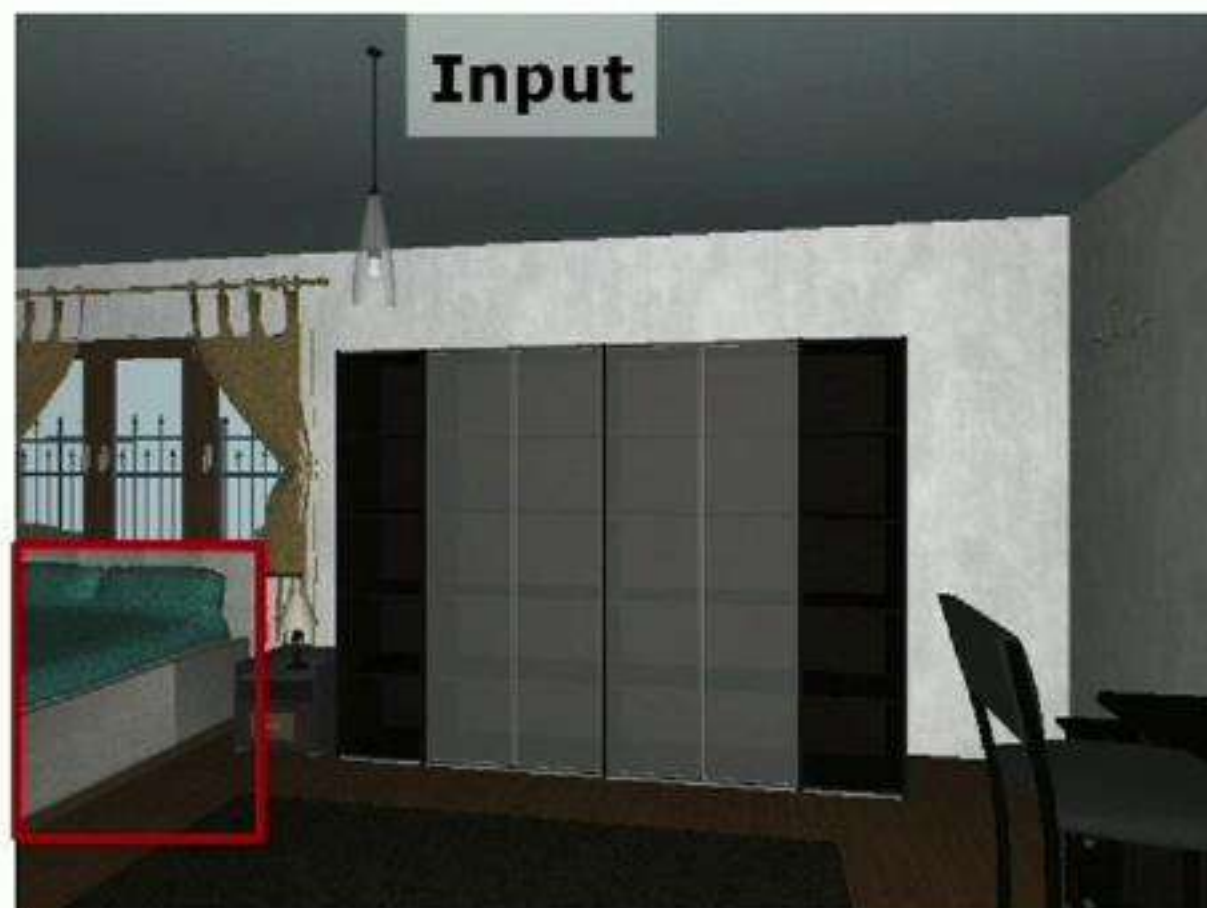
1. Our policy network has learned different moving strategies from shortest path
2. In general, the learned policy keeps the agent in a distance to the target object

Passive Perception vs. Active Perception

Shortest Path vs. Learned Active Path

Shortest-Path





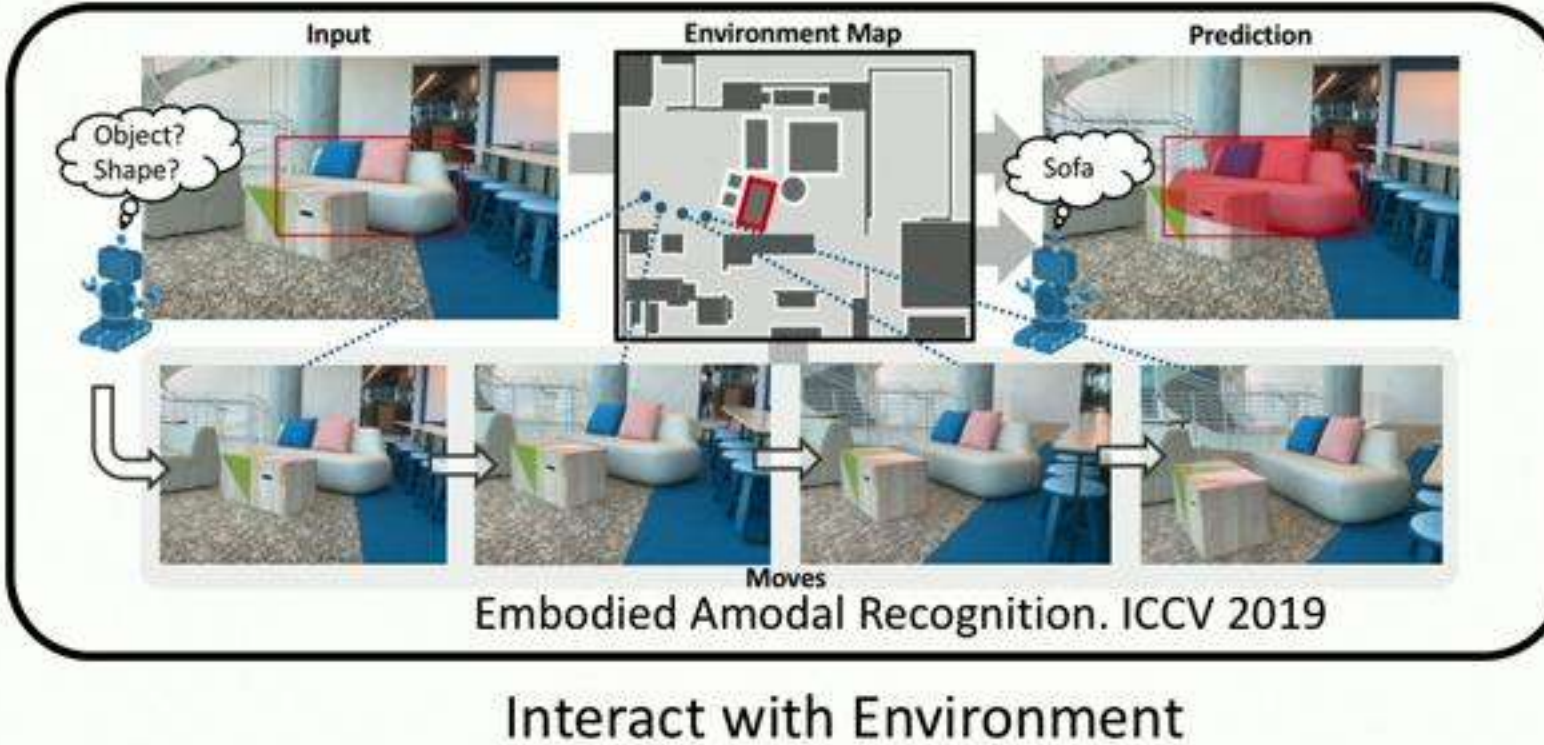
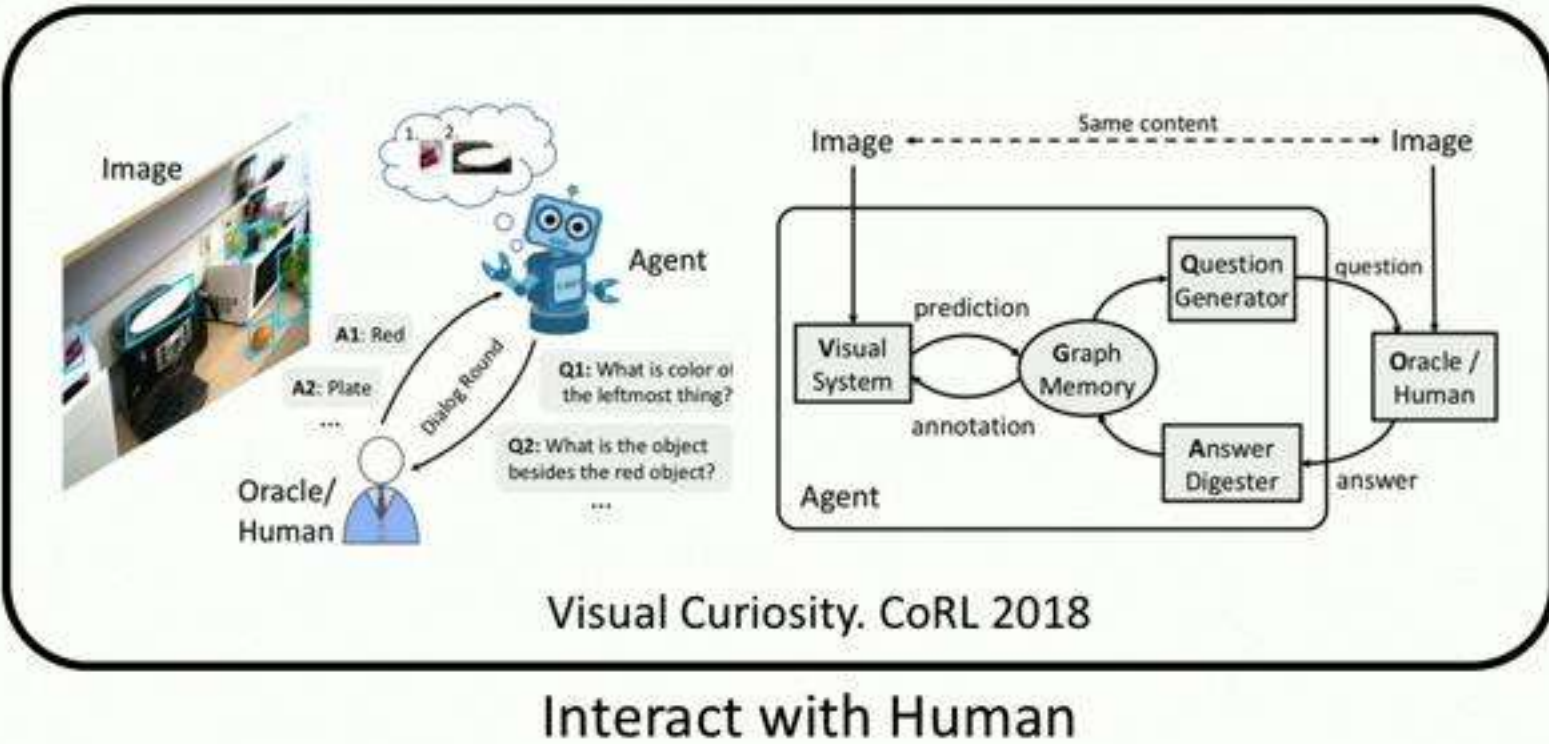
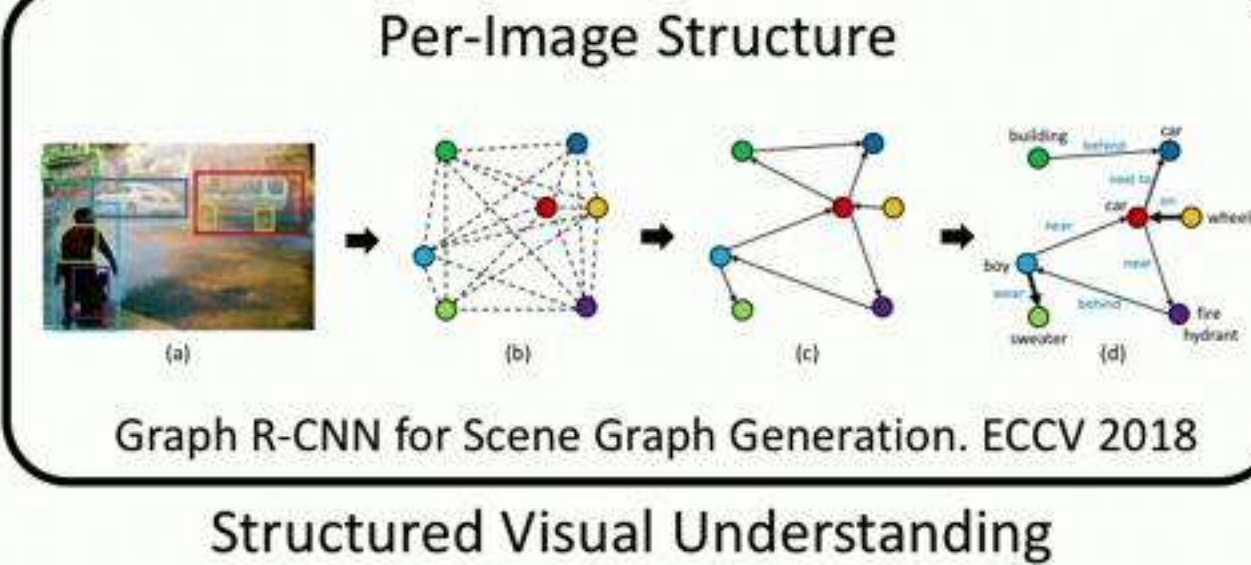
Active Path

Shortest-Path

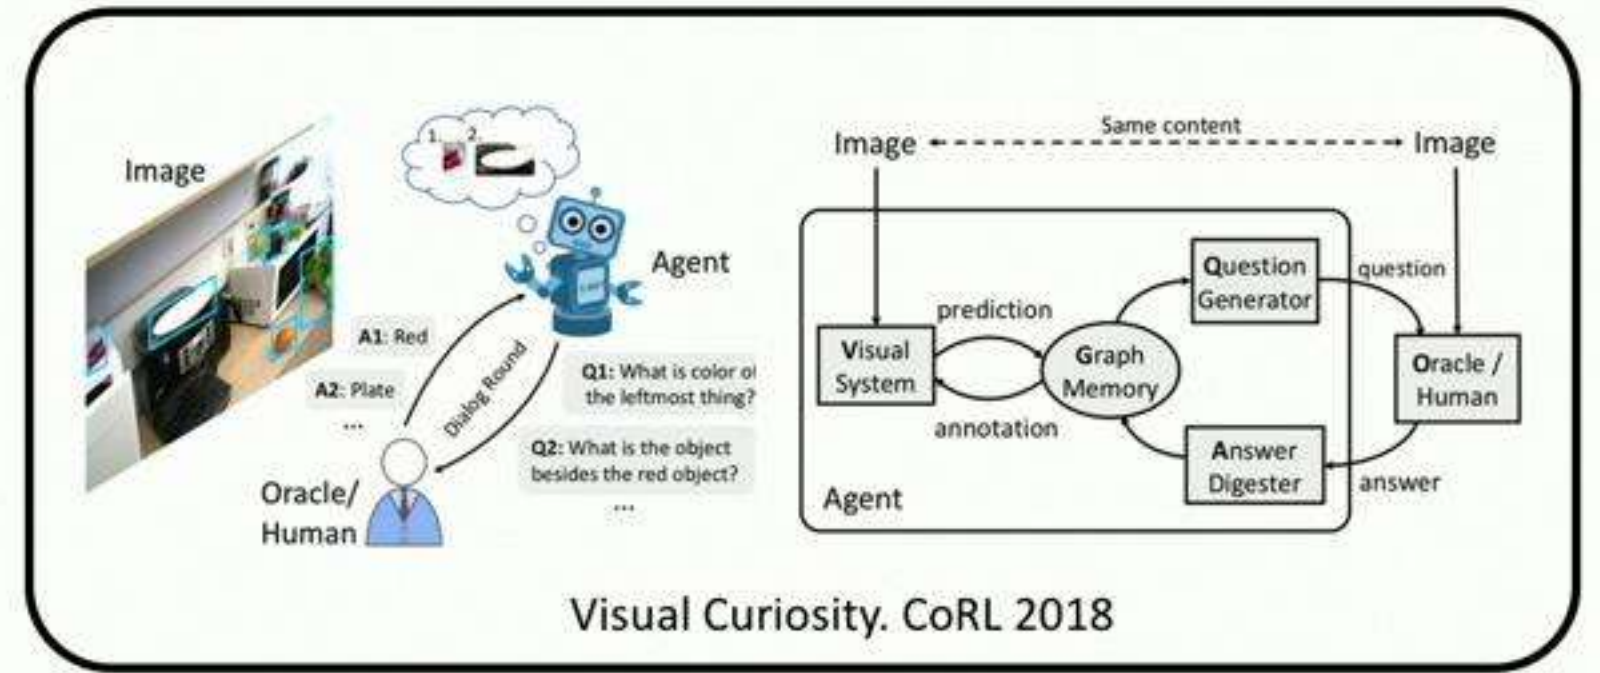
Takeaways

- A embodied visual recognition system is proposed as an initial step toward an intelligent agent system
- Embodiment helps to get better visual recognition of objects
- Learning a better moving strategy is challenging but helpful to improve visual recognition
- It would be interesting to enable full understanding of the whole environment after moves

In this talk

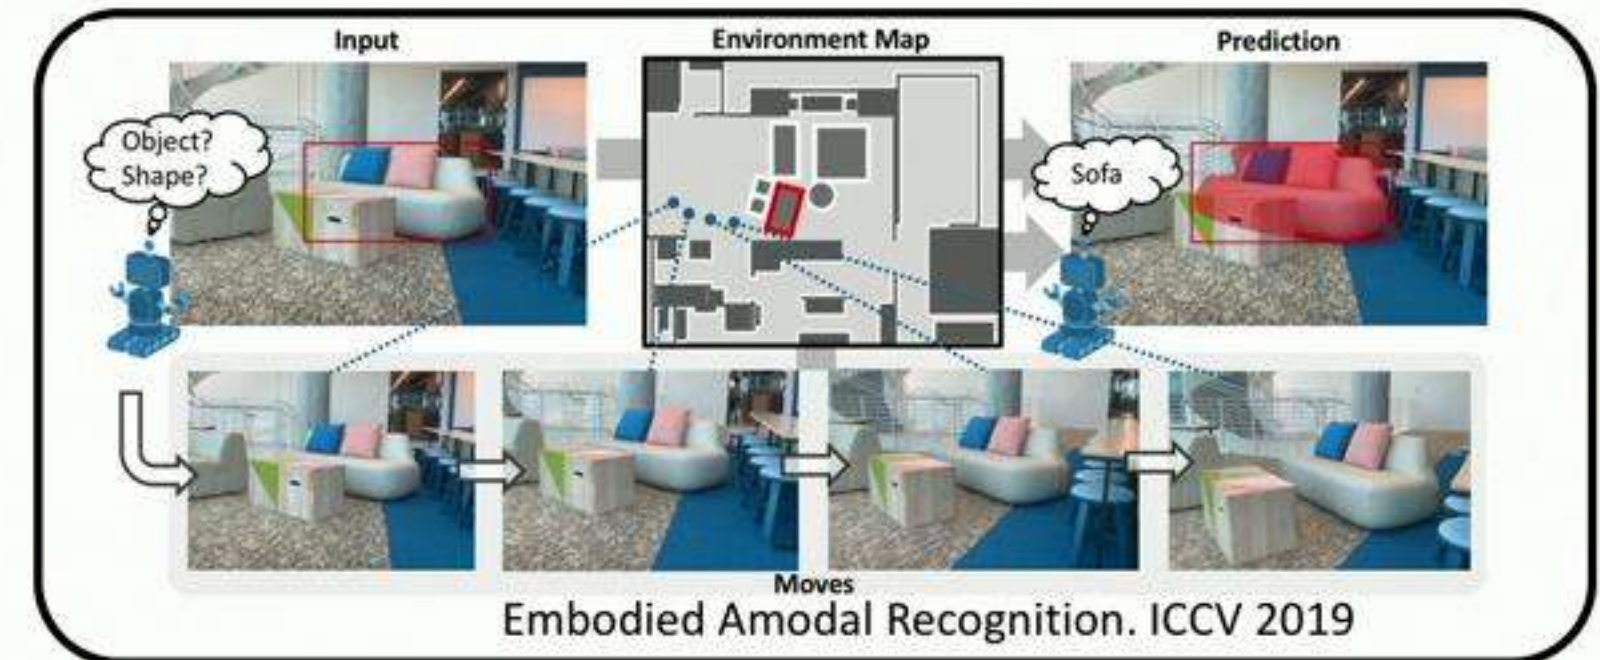
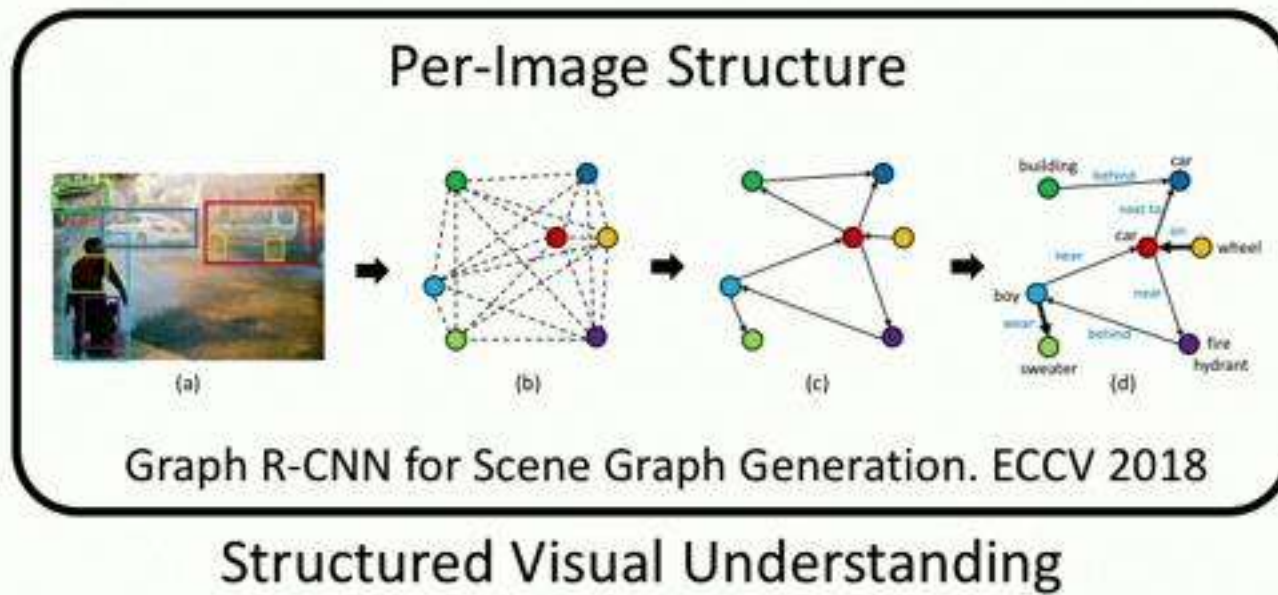


In this talk



Interact with Human

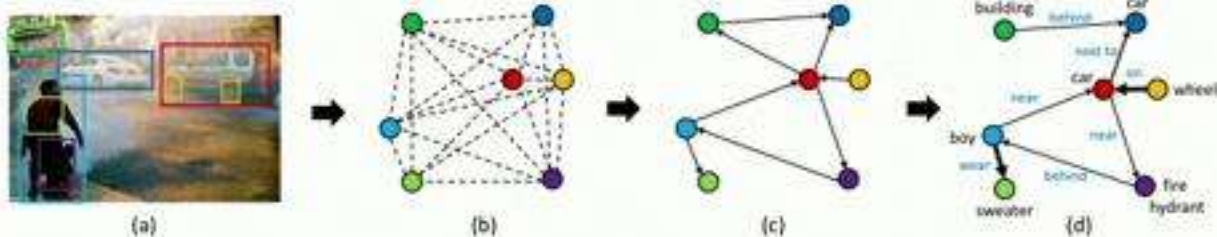
Learn from interactions with human and environment



Interact with Environment

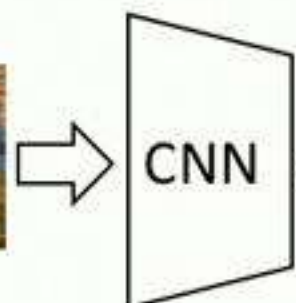
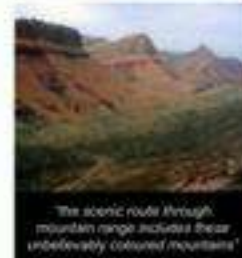
As Future Works

Per-Image Structure



Graph R-CNN for Scene Graph Generation. ECCV 2018

Structured Visual Understanding

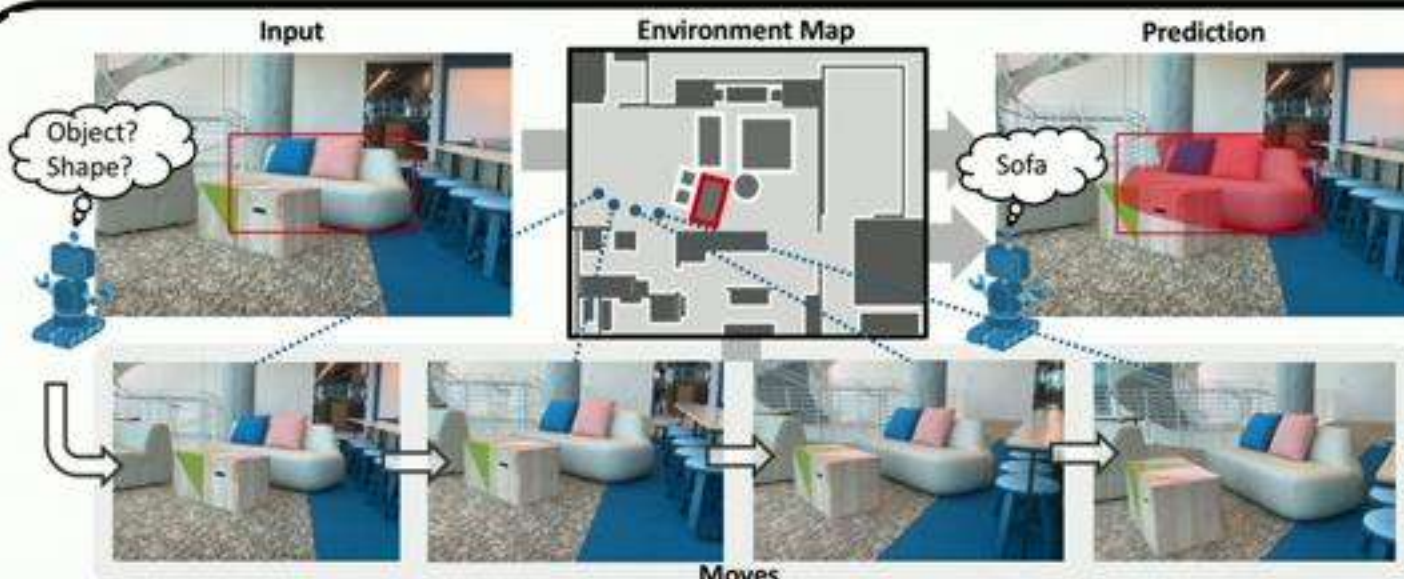


Two cars are parked beside the road

Learning from image corpus

Interact with Human

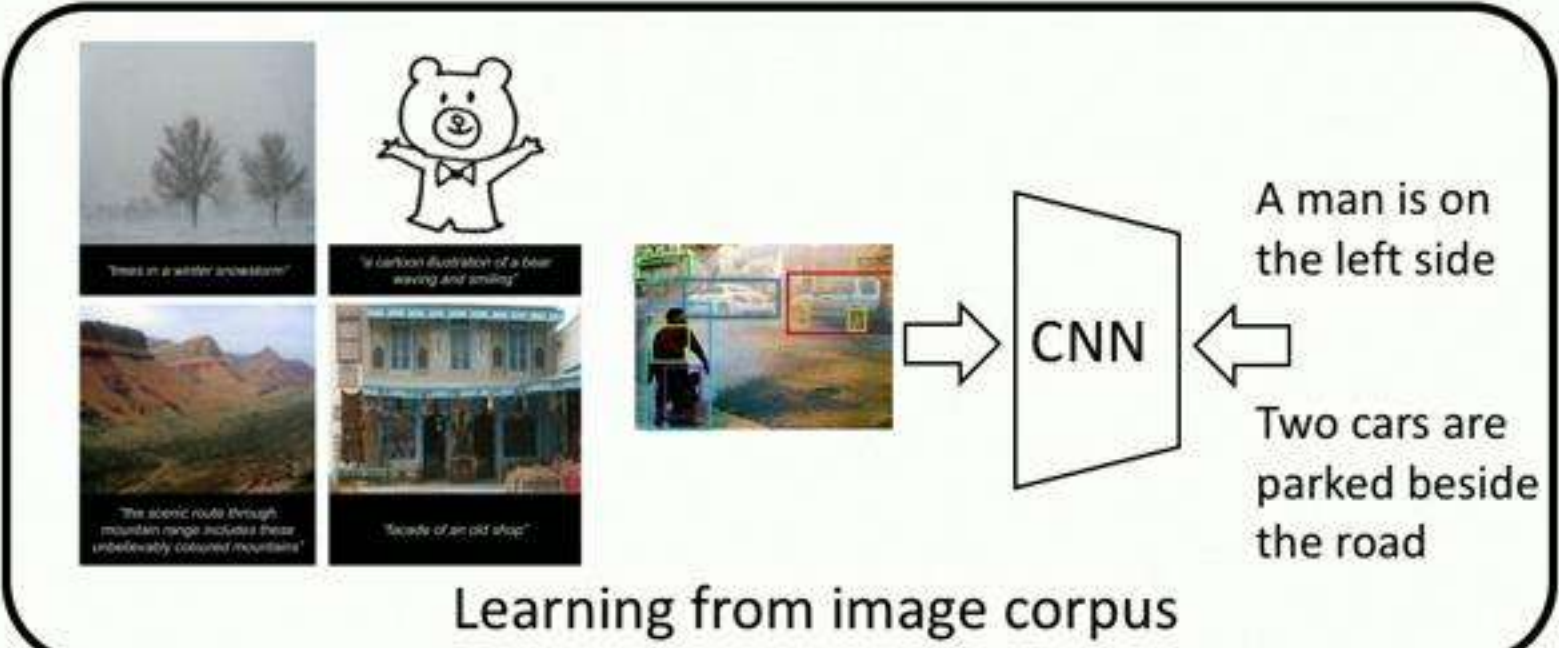
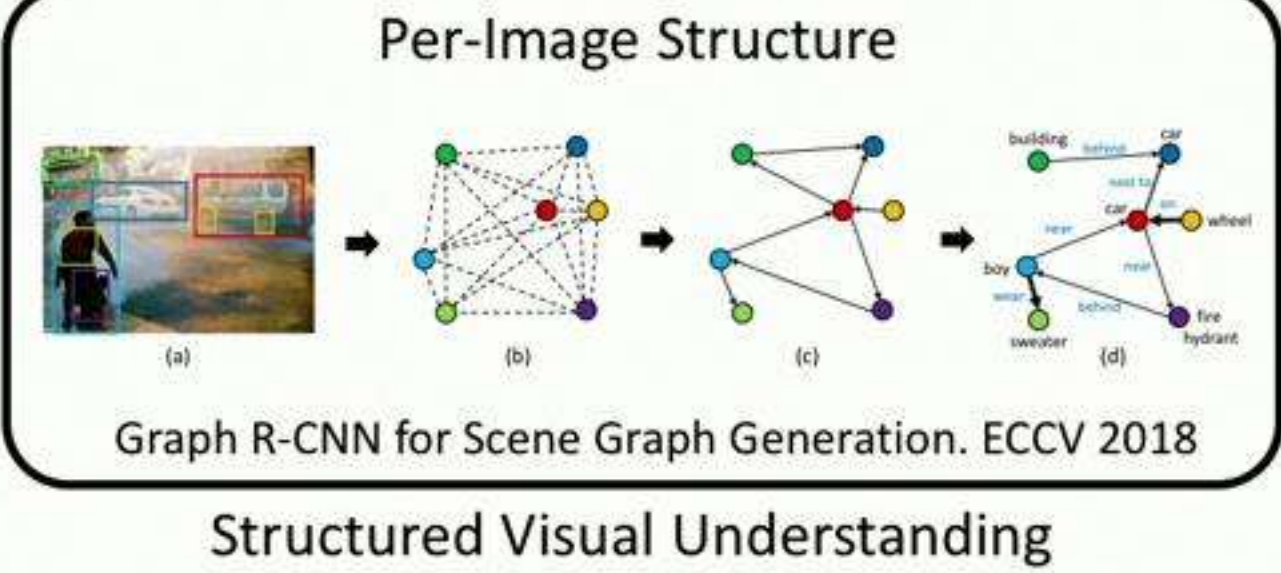
Learn from interactions with human and environment



Embodied Amodal Recognition. ICCV 2019

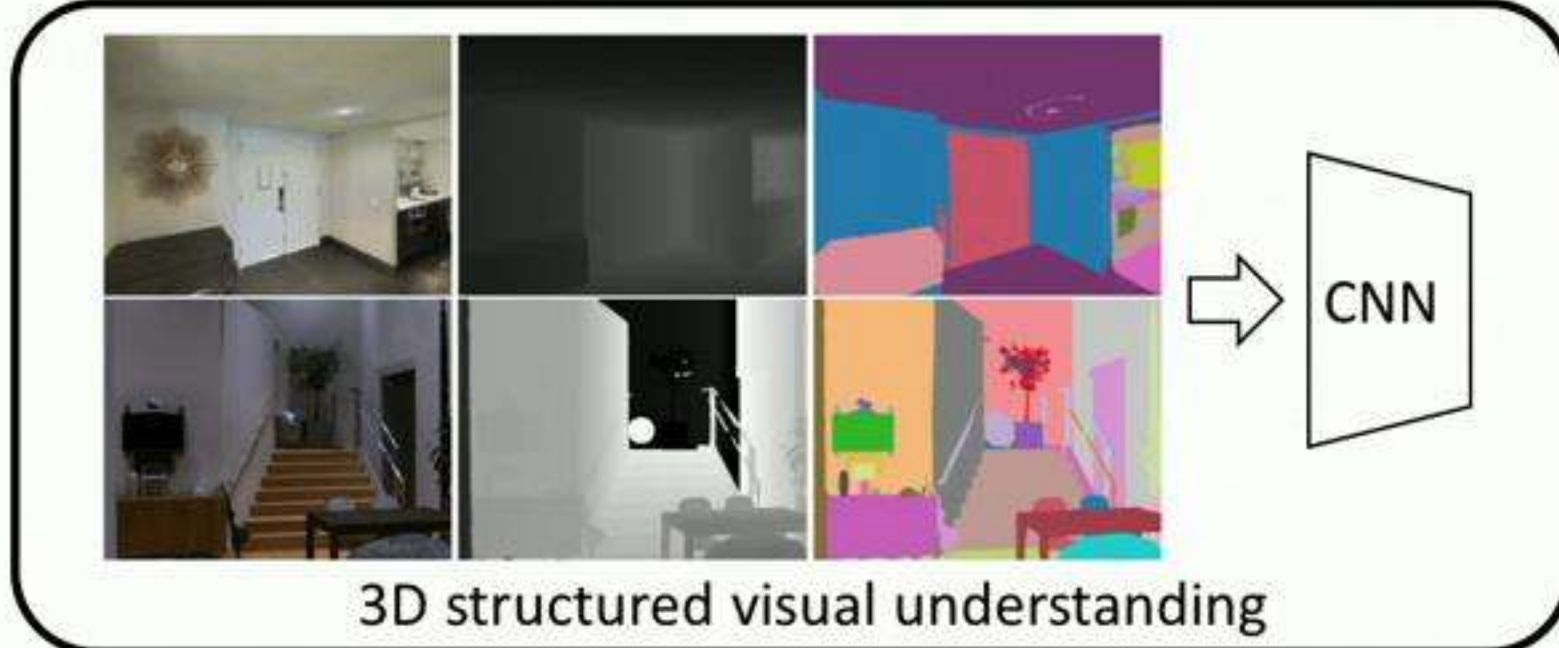
Interact with Environment

As Future Works



Interact with Human

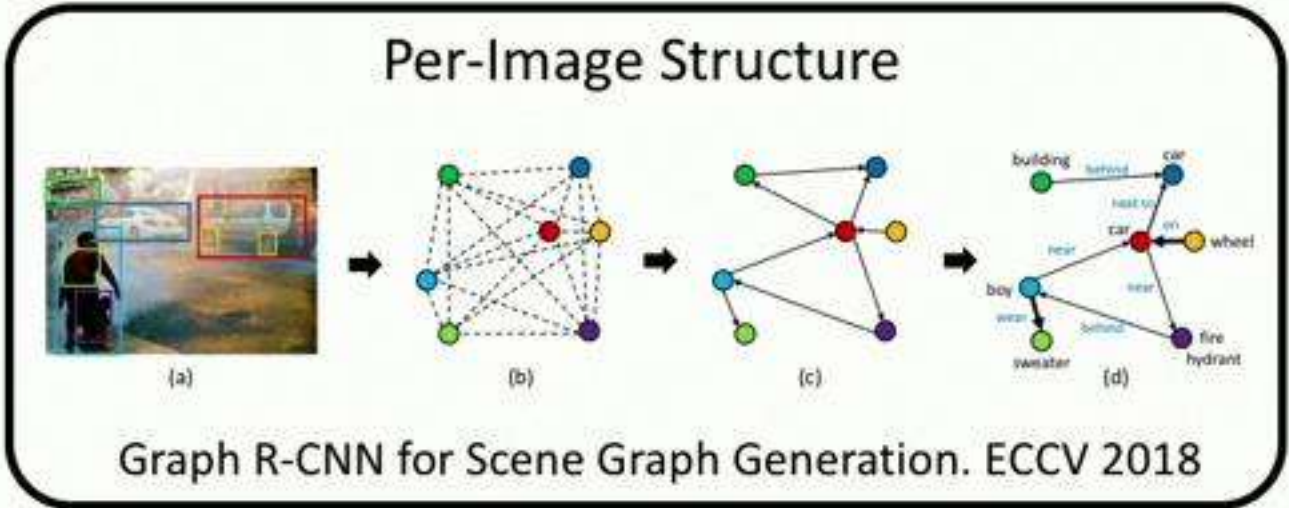
Learn from interactions with human and environment



Interact with Environment

As Future Works

Learn for interactions with human and environment



Structured Visual Understanding



Learn



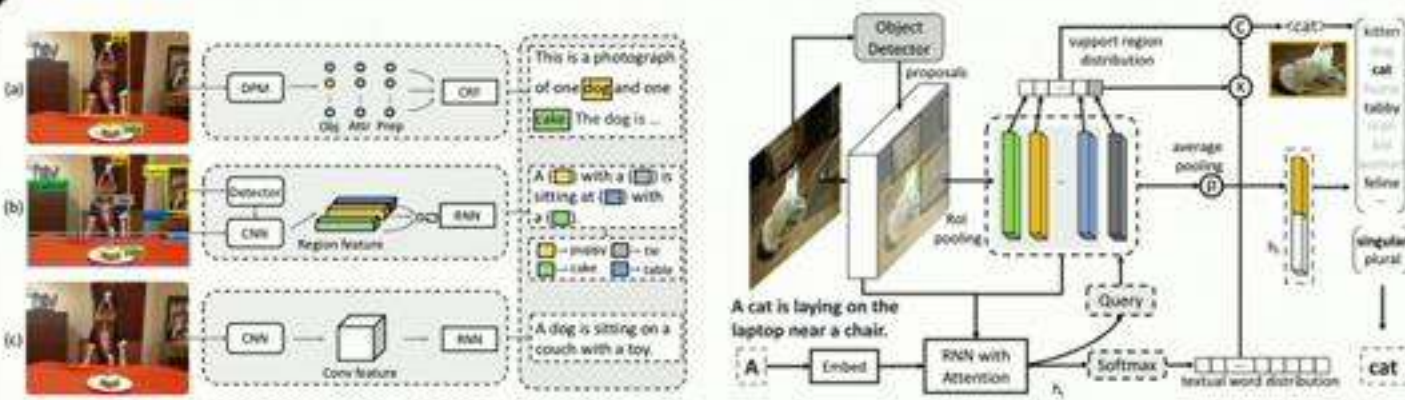
The diagram illustrates the Neural Baby Talk architecture for video captioning. It shows three input video frames (a, b, c) being processed by different neural networks (DPM, CNN, RNN) to generate captions. The main architecture involves an Object Detector, Region Pooling, and an RNN with Attention mechanism to produce a textual word distribution for the final caption.

Input Frames and Captions:

- (a) This is a photograph of one **box** and one **lake**. The dog is ...
- (b) A **box** with a **lake** is sitting at **lake** with a **lake**.
- (c) A dog is sitting on a couch with a toy.

Architecture Components:

- Object Detector:** Takes an input image and outputs proposals.
- Region Pooling:** Processes proposals to generate support region distribution and average pooling.
- RNN with Attention:** Takes an input image (A) and a query (h) to produce a textual word distribution.
- Softmax:** Outputs the final textual word distribution.
- Embed:** Takes an input image (A) and outputs an embedding.
- Query:** Takes an input image (h) and outputs a query.
- Support Region Distribution:** Takes an input image (h) and outputs a support region distribution.
- Average Pooling:** Takes an input image (h) and outputs an average pooling result.
- Textual Word Distribution:** The final output of the RNN with Attention.

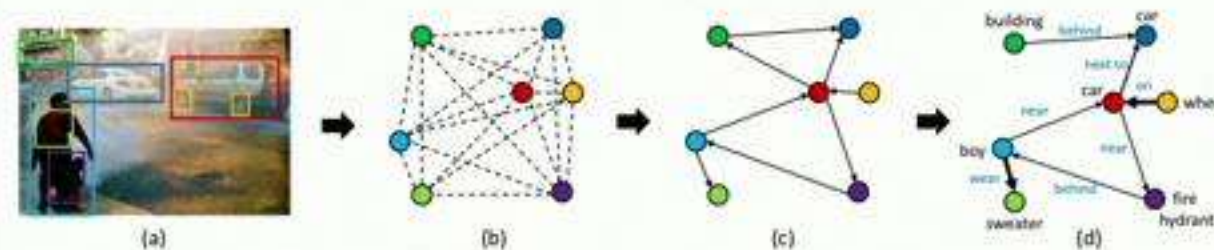


Neural Baby Talk. CVPR 2018

Per-Image Structure

(a) Input image with bounding boxes. (b) Dense graph representation. (c) Sparse graph representation. (d) Final scene graph with labeled edges (e.g., near, behind, on, wearing, has part).

Graph R-CNN for Scene Graph Generation. ECCV 2018

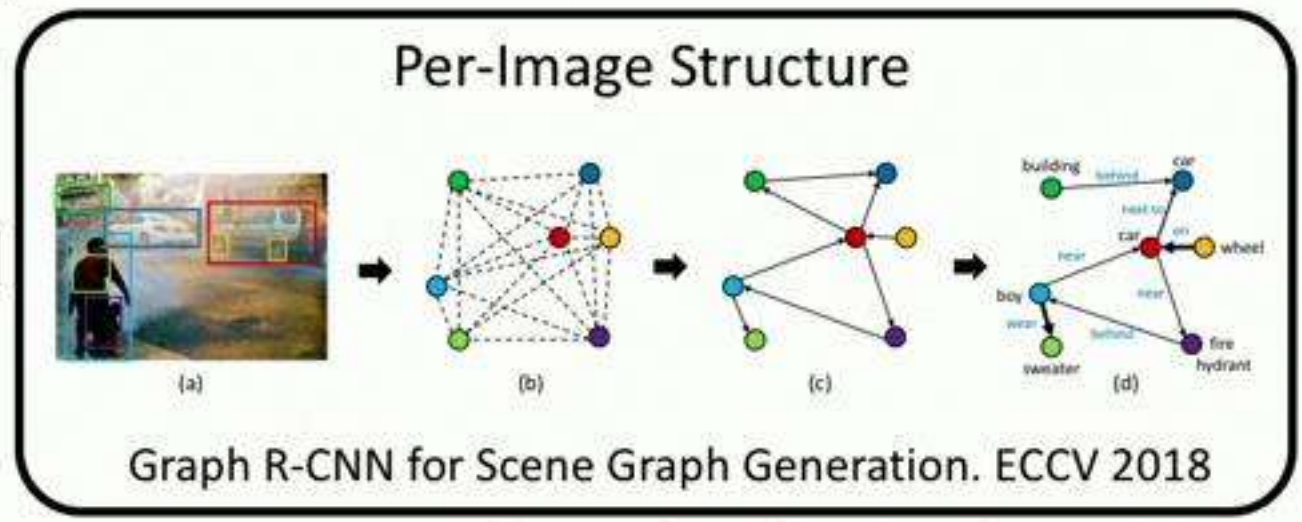
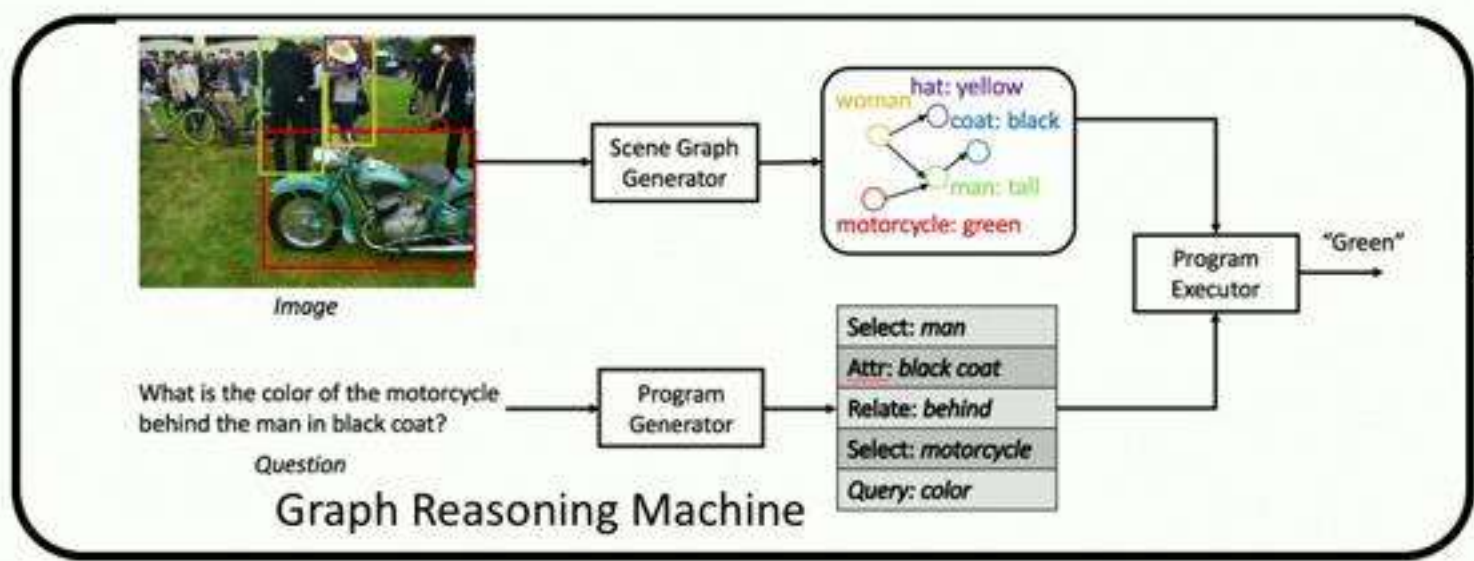
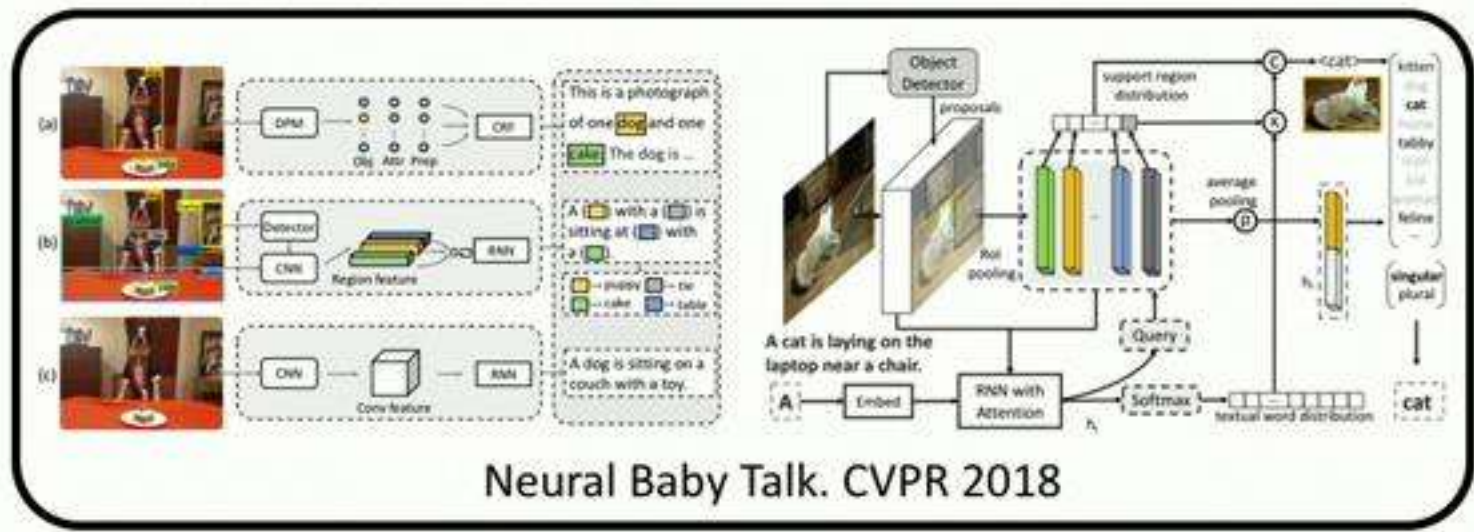


Graph R-CNN for Scene Graph Generation. ECCV 2018

Structured Visual Understanding

Learr

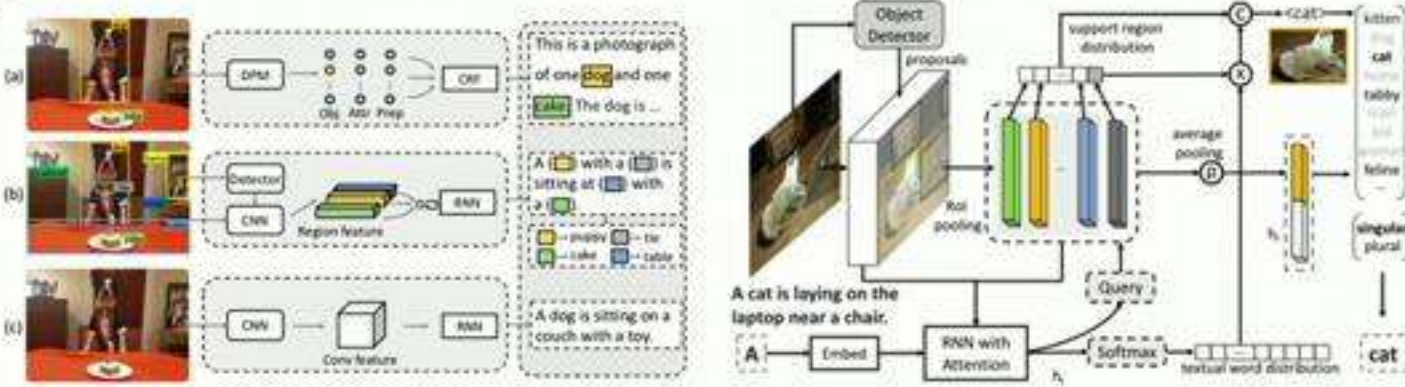
As Future Works



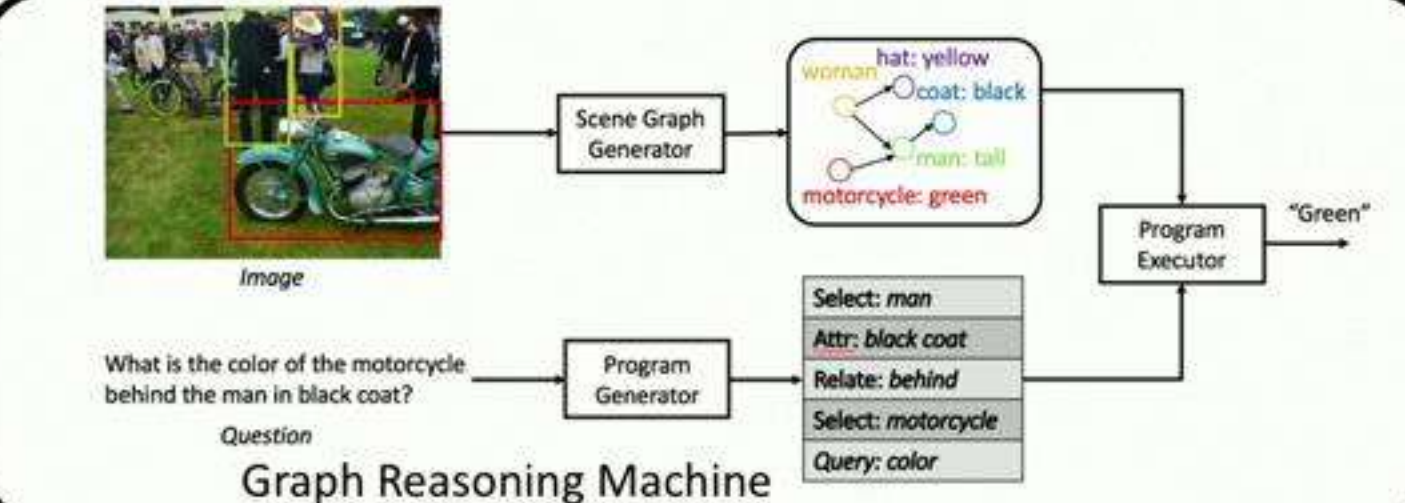
Structured Visual Understanding




As Future Works



Neural Baby Talk. CVPR 2018

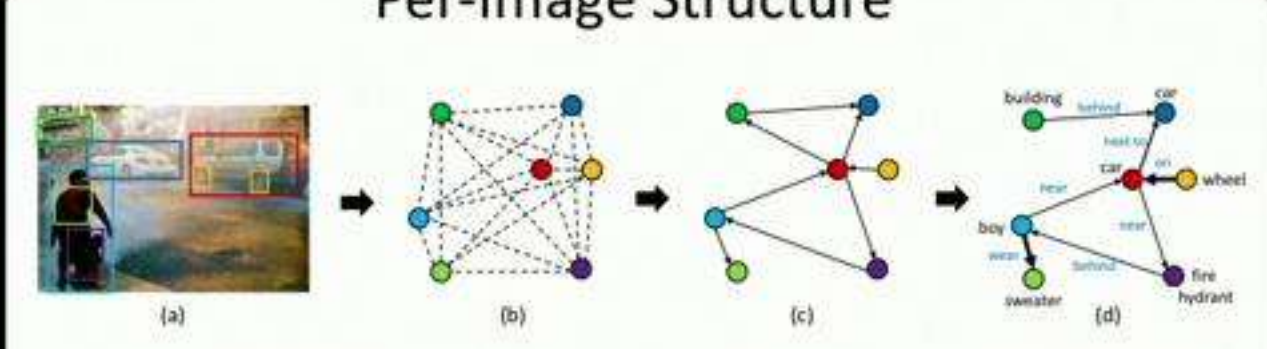


Graph Reasoning Machine



Visual Language Navigation. Anderson et al.

Per-Image Structure

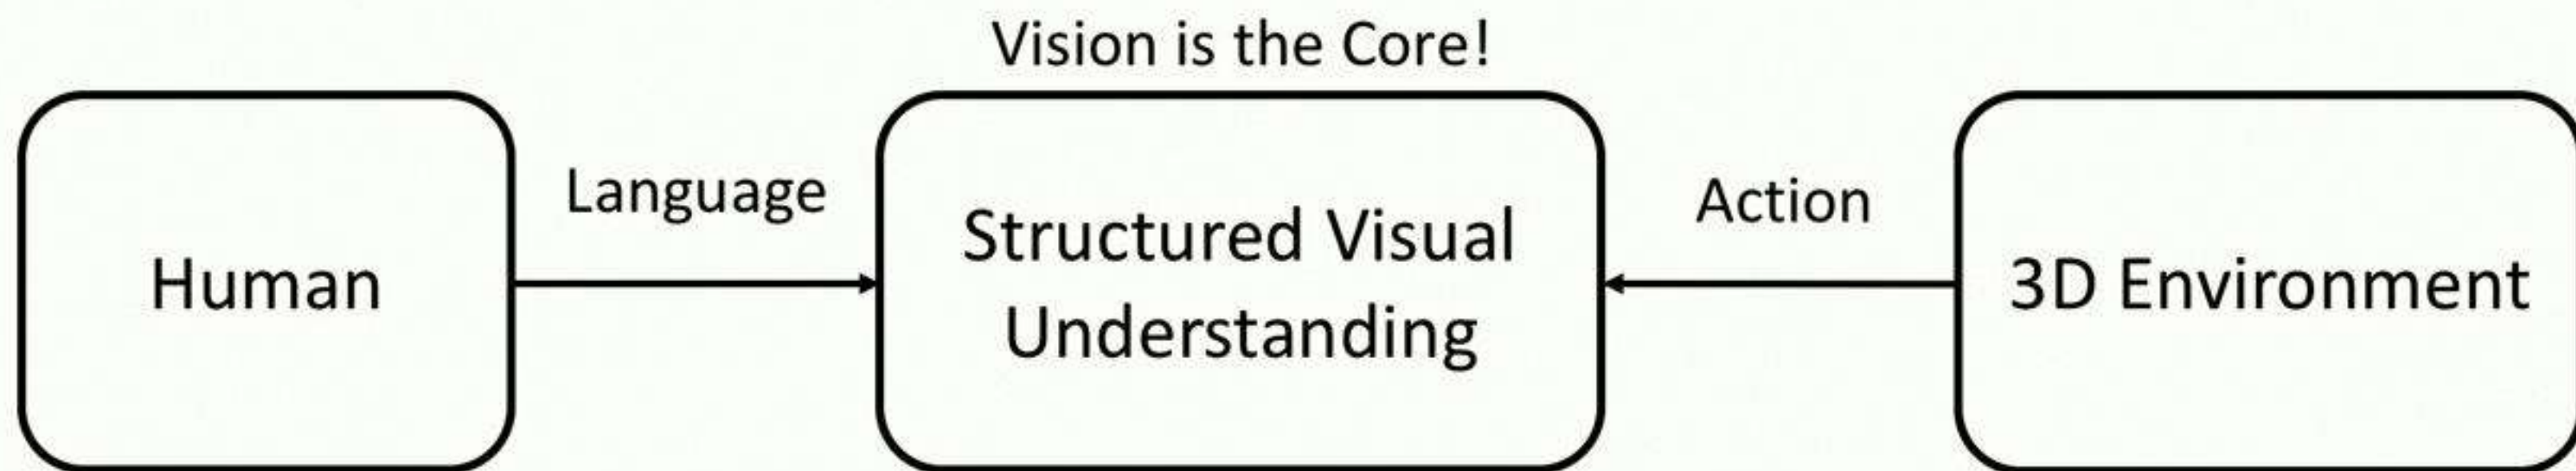


Graph R-CNN for Scene Graph Generation. ECCV 2018

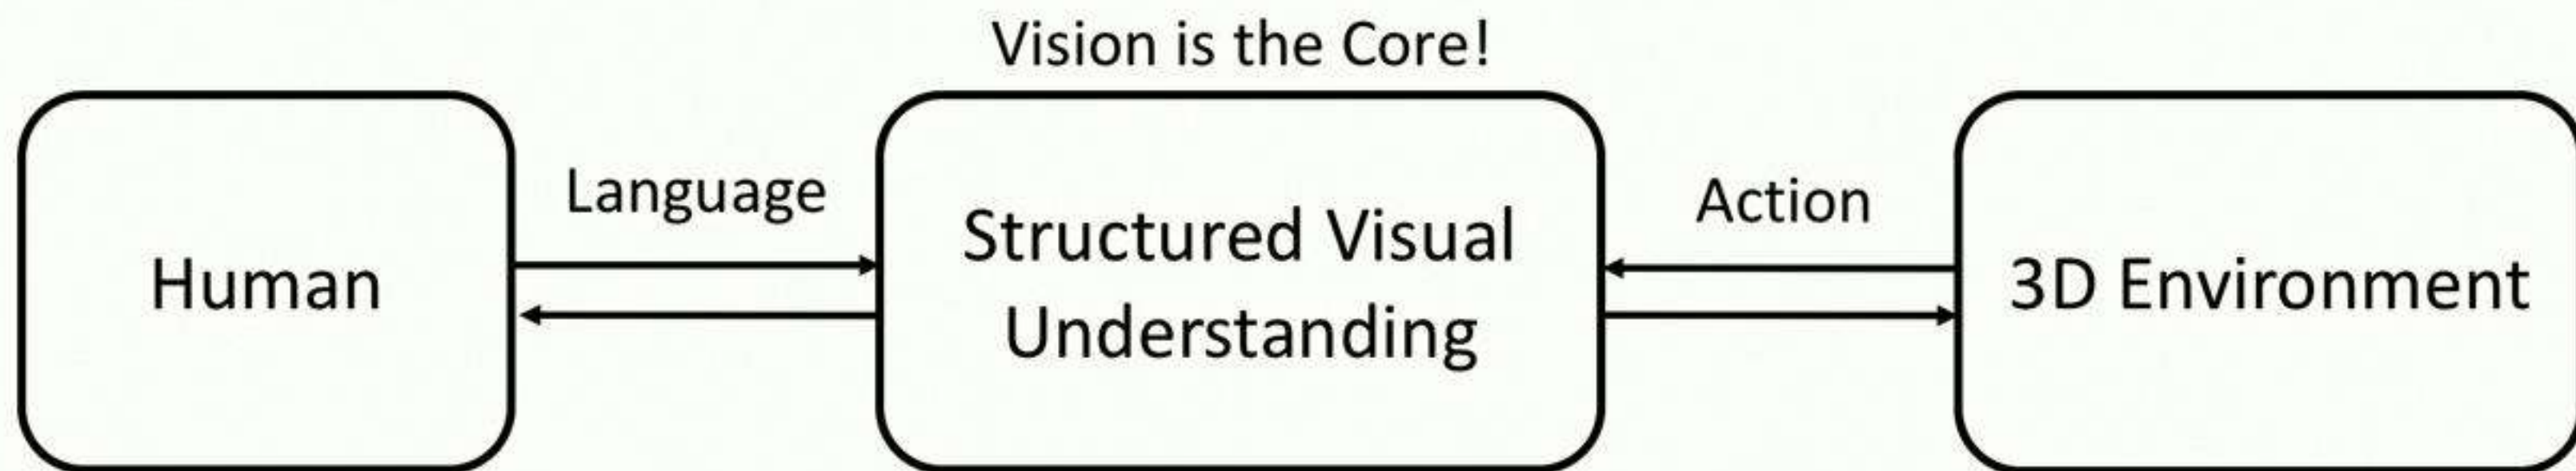
Structured Visual Understanding



To summarize



To summarize



Collaborators



Devi Parikh



Dhruv Batra



Jiasen Lu



Stefan Lee



Zhile Ren



Minze Xu



Anitha Kannan



Xinlei Chen



Ji Lin



Chuang Gan



Hongyuan Zhu



David Crandall

Thanks!