Neurosymbolic AI

—

David Cox, Ph.D.
IBM Director, MIT-IBM Watson AI Lab
IBM Research

MIT-IBM
WATSON
AI LAB

“Artificial Intelligence”

The evolution of AI

General AI
Revolutionary

Broad AI
Disruptive and Pervasive

Narrow AI
Emerging

▼ **We are here**

2050 and beyond

The evolution of AI

Narrow AI

Single task, single domain

Superhuman accuracy and speed for certain tasks

Broad AI

Multi-task, multi-domain

Multi-modal

Distributed AI

Explainable

General AI

Cross-domain learning and reasoning

Broad autonomy

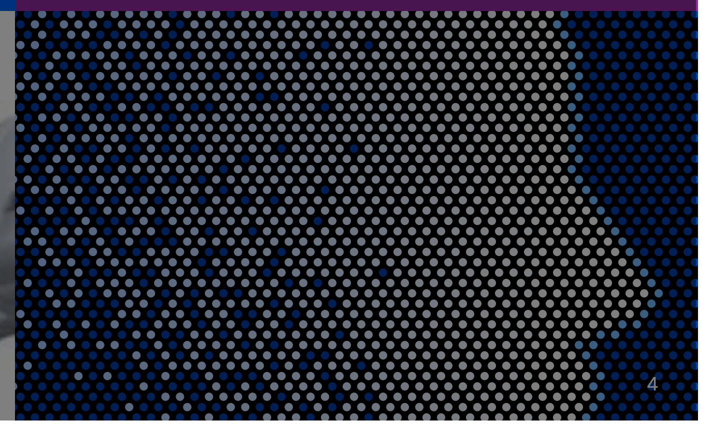

The evolution of AI

Narrow AI

Single task, single domain

Superhuman accuracy and speed for certain tasks

Broad AI

Multi-task, multi-domain

Multi-modal

Distributed AI

Explainable

General AI

Cross-domain learning and reasoning

Broad autonomy

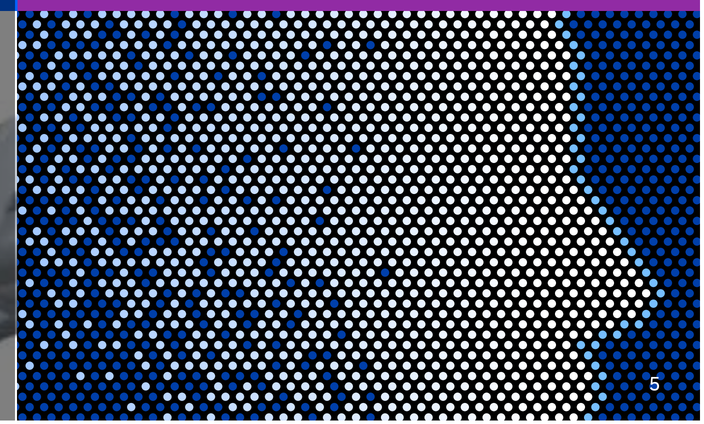

Elon Musk

Elon Musk Compares Building Artificial Intelligence To "Summoning The Demon"

Posted Oct 26, 2014 by Greg Kumparak (@grg)

17.6k SHARES

Next Story

Technology

Stephen Hawking warns artificial intelligence could end mankind

By Rory Cellan-Jones
Technology correspondent

2 December 2014 | Technology

The evolution of AI

Narrow AI

Single task, single domain

Superhuman accuracy and speed for certain tasks

Broad AI

Multi-task, multi-domain

Multi-modal

Distributed AI

Explainable

General AI

Cross-domain learning and reasoning

Broad autonomy

The path to a “Broad AI” toolbox

Explainability

Security

Ethics

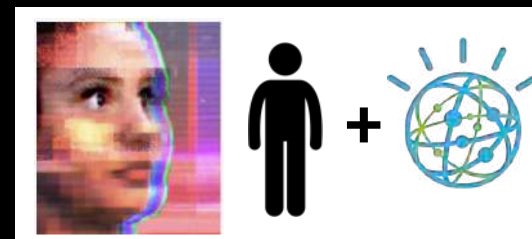

Learn more from small data

Learns to transfer

+

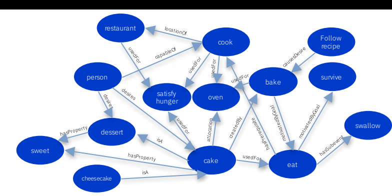

Reasoning

Infrastructure

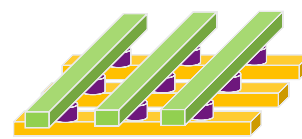


Physics of AI

Platform for AI Lifecycle

Compute | Data & Models | Applications | Workflow

MIT-IBM
WATSON
AI LAB

The evolution of AI

General AI
Revolutionary

Broad AI
Disruptive and Pervasive

Narrow AI
Emerging

▼ **We are here**

2050 and beyond

So what's "narrow" about today's AI toolbox?

Forbes / Tech

DEC 29, 2014 @ 11:37 AM 115,776

Tech 2015: Deep Learning And Machine Intelligence Will Eat The World

Anthony Wing Kosner, CONTRIBUTOR

Quantum of Content and innovations in user experience **FULL BIO**

Opinions expressed by Forbes Contributors are their own.

man in black shirt is playing guitar.

construction worker in orange safety vest is working on road.

Karpathy and Li, 2015


THINK
DENKE
$24,000
$77,147
$21,600
Who is Bram Stoker?
$ 17,973
$1,000

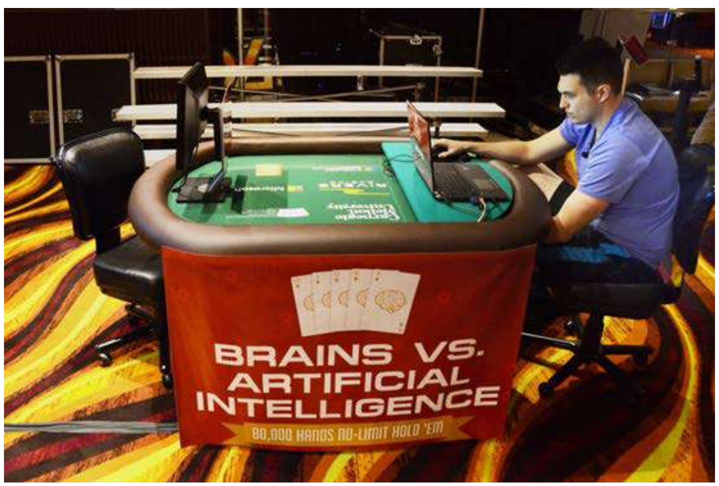
BRAINS VS.
ARTIFICIAL
INTELLIGENCE
80,000 HANDS NO-LIMIT HOLD 'EM

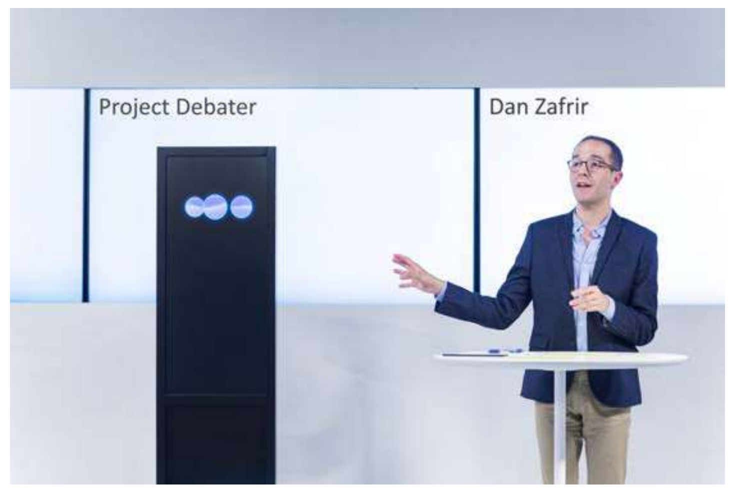
Project Debater
Dan Zafrir

Gatys et al. 2015

Brock et al. 2018

“Teddy Bear”

Meret Oppenheim, *Le Déjeuner en fourrure*

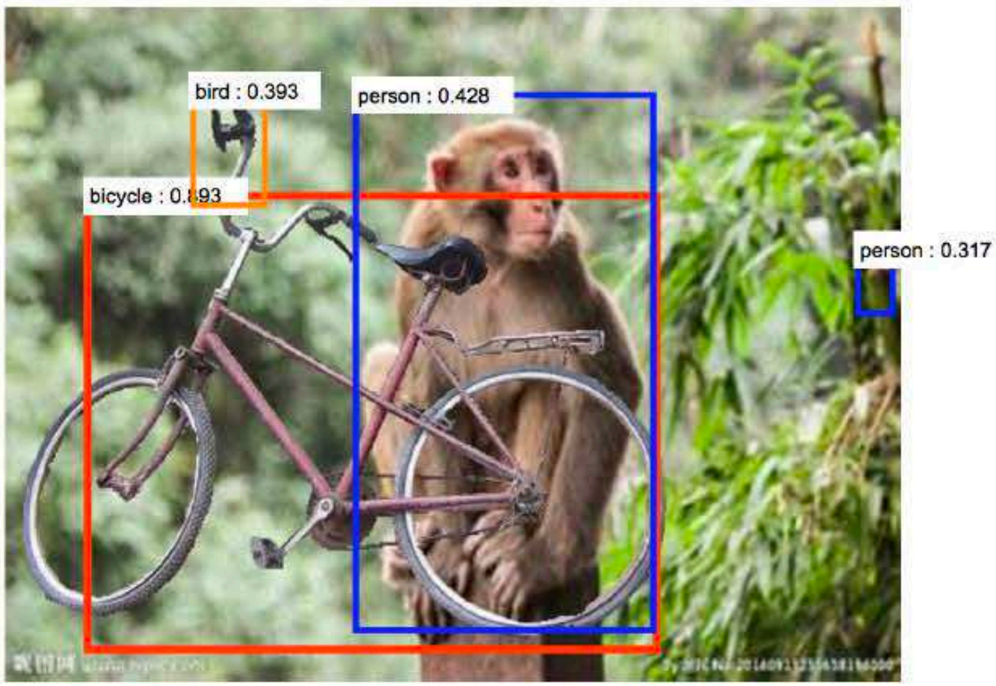

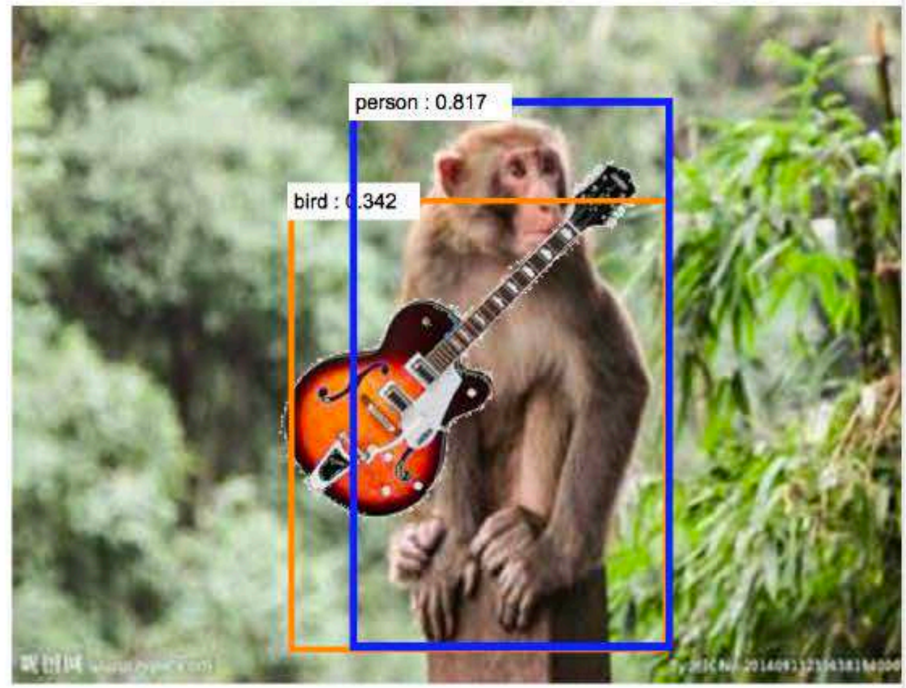

Wang et al. 2018

man in black shirt is playing guitar.

construction worker in orange safety vest is working on road.

Karpathy and Li, 2015

a man riding a
motorcycle on a beach

Lake, Ullman, Tenenbaum & Gershman, 2016

IMAGENET

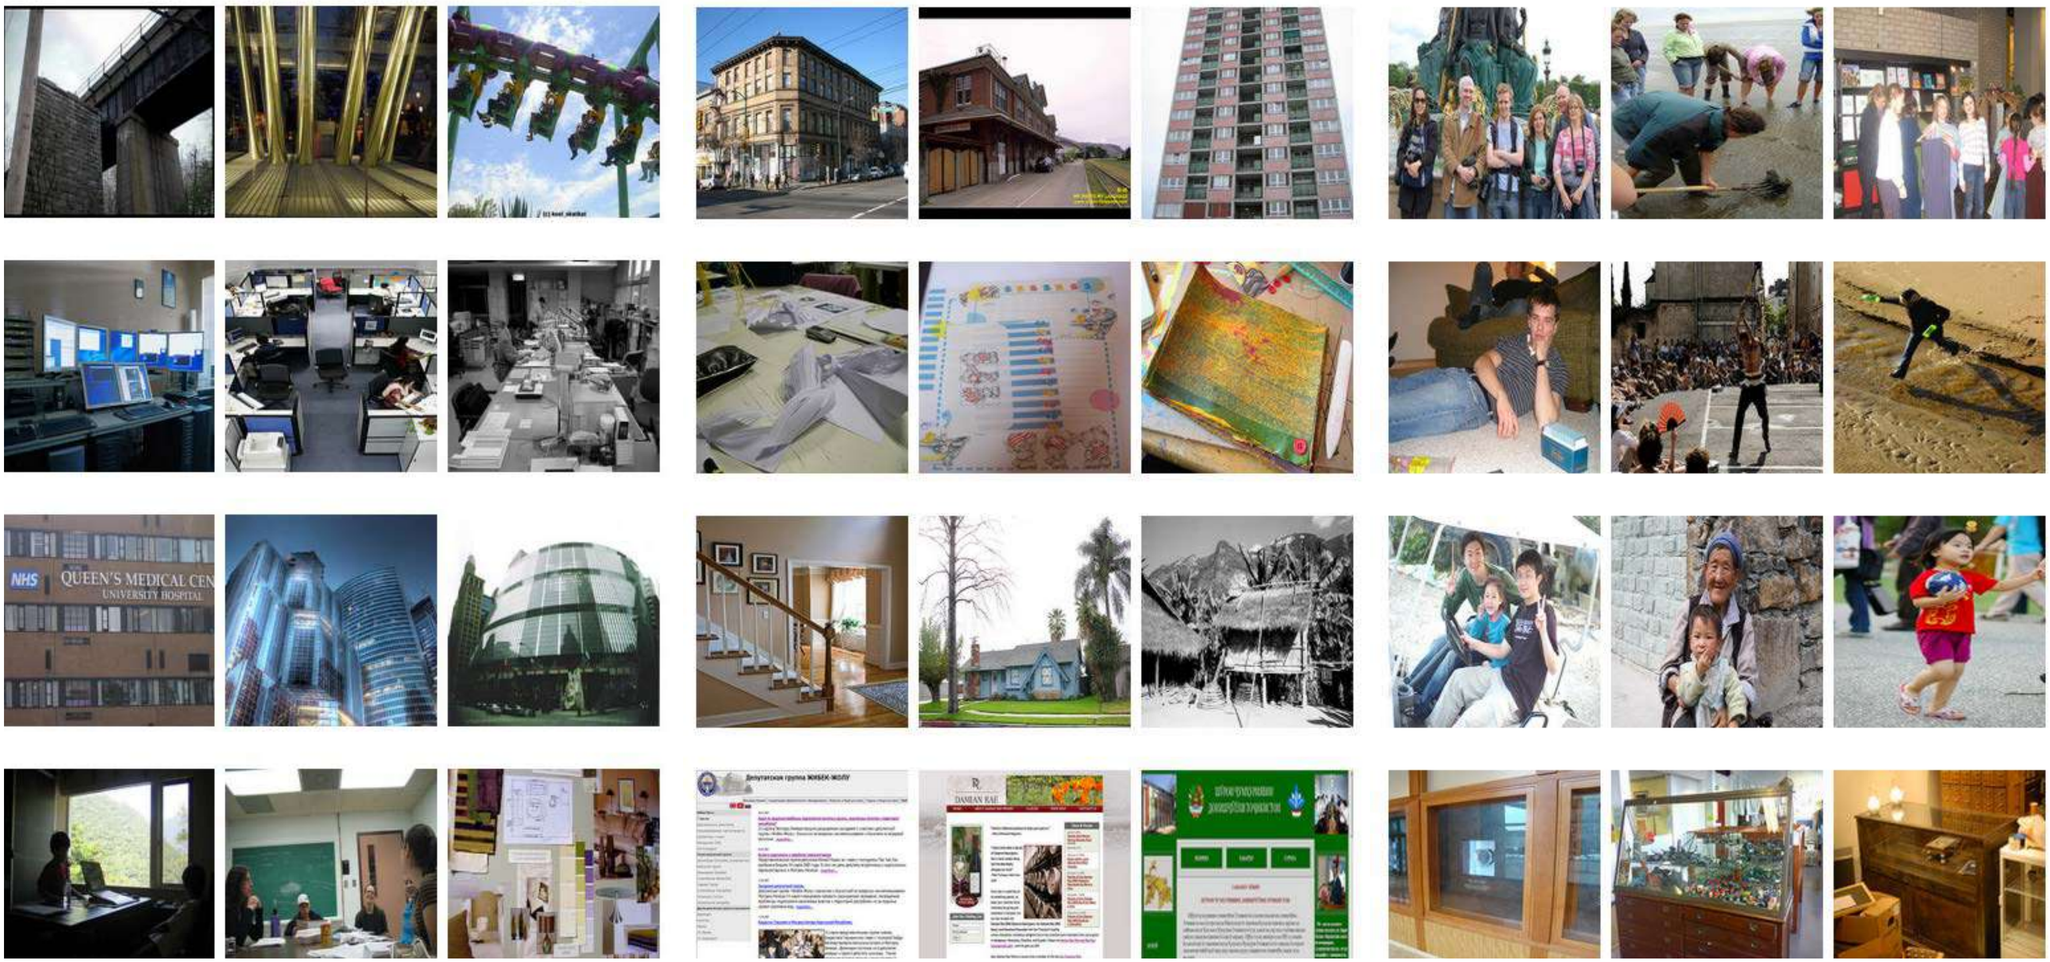

What’s this?

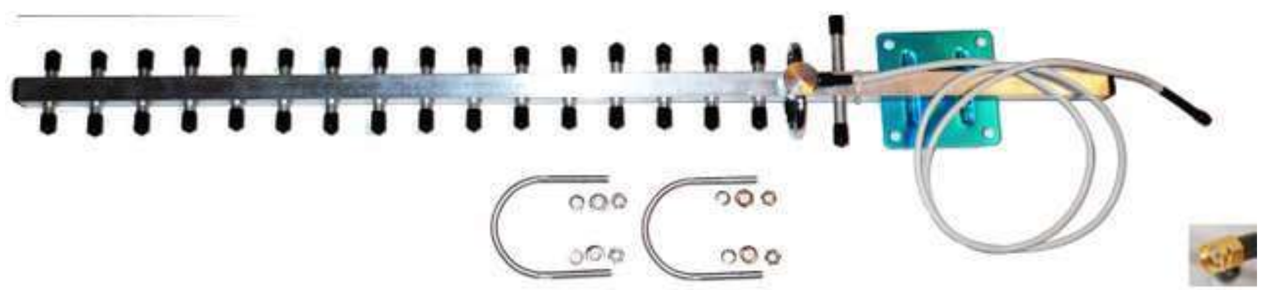

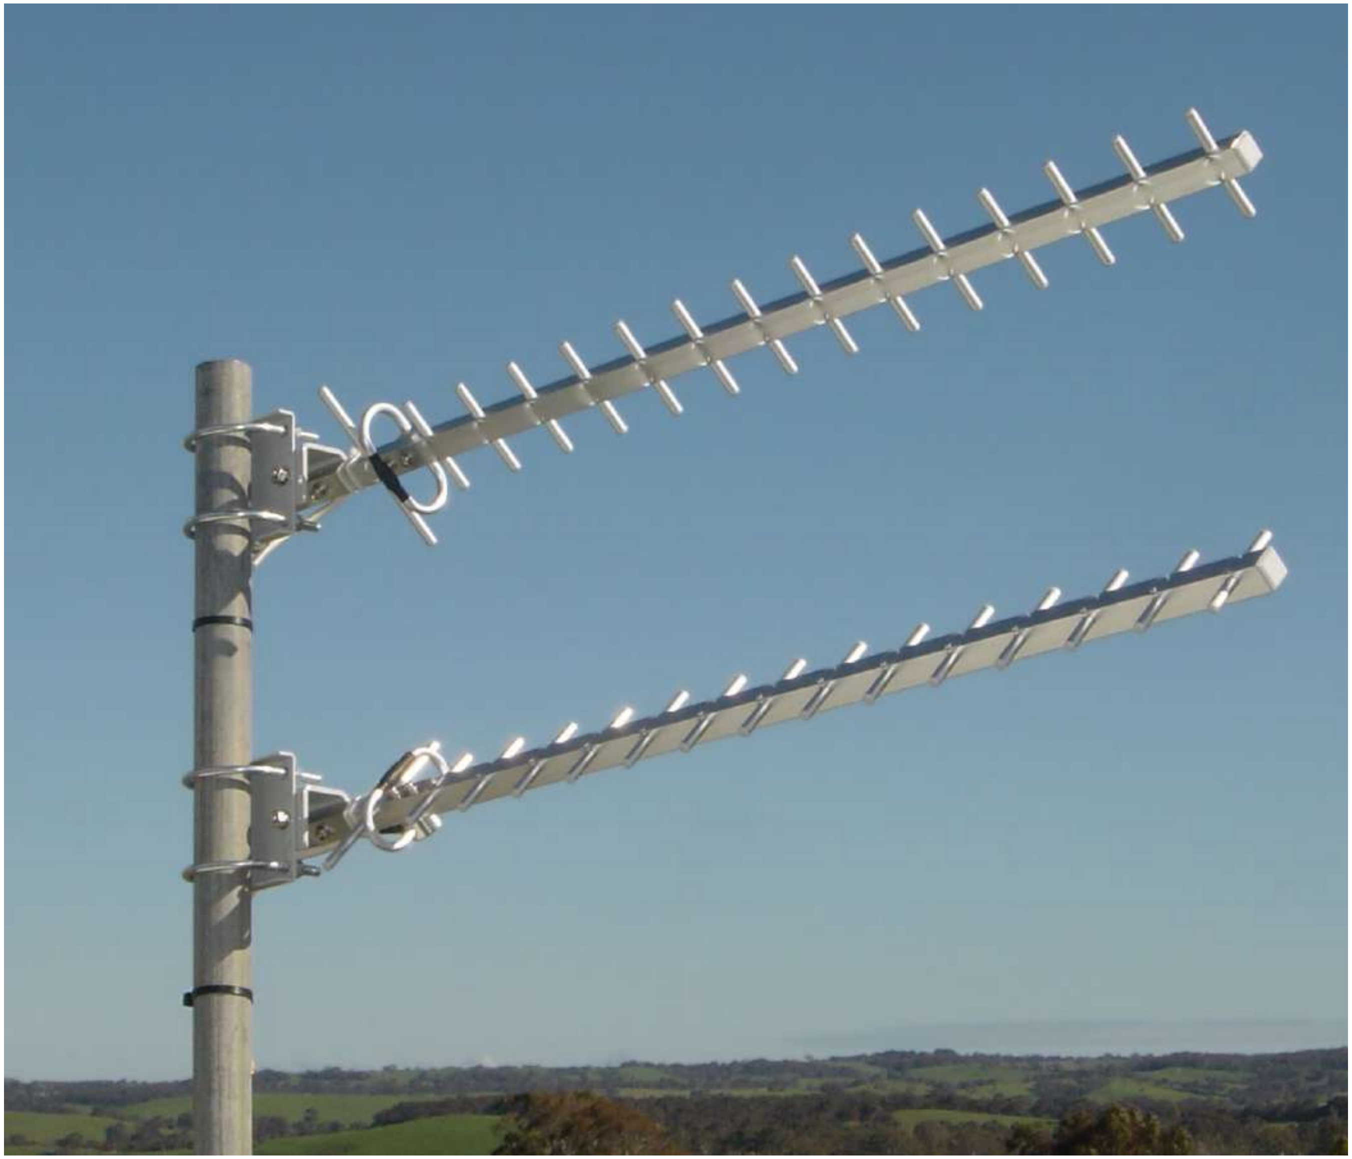

IMAGENET

   

ObjectNet

Boris Katz
MIT

Andrei Barbu
MIT

Dan Gutfreund
IBM

ObjectNet

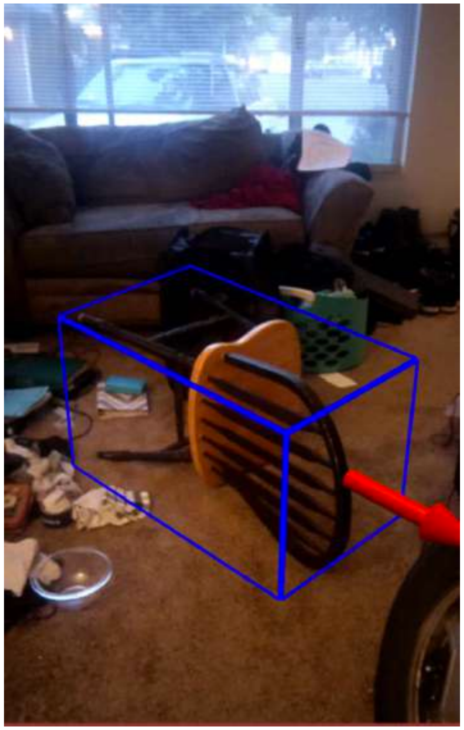

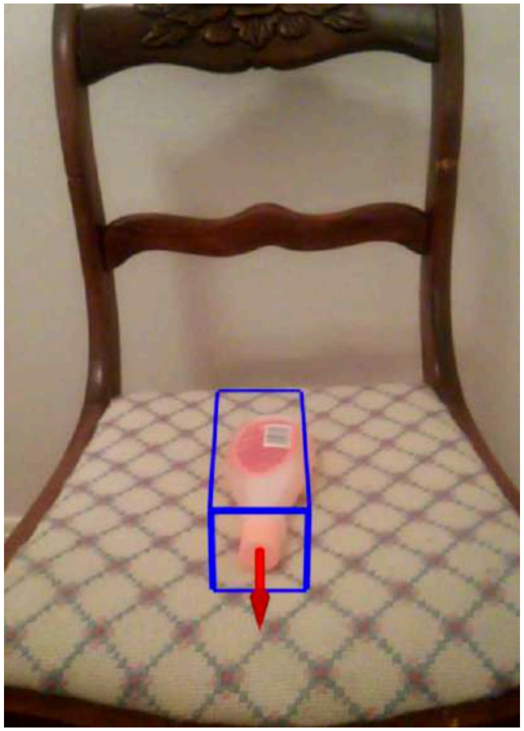

- ~50K images
- ~300 object classes
- 4 different room types

Testing ImageNet-trained models on ObjectNet

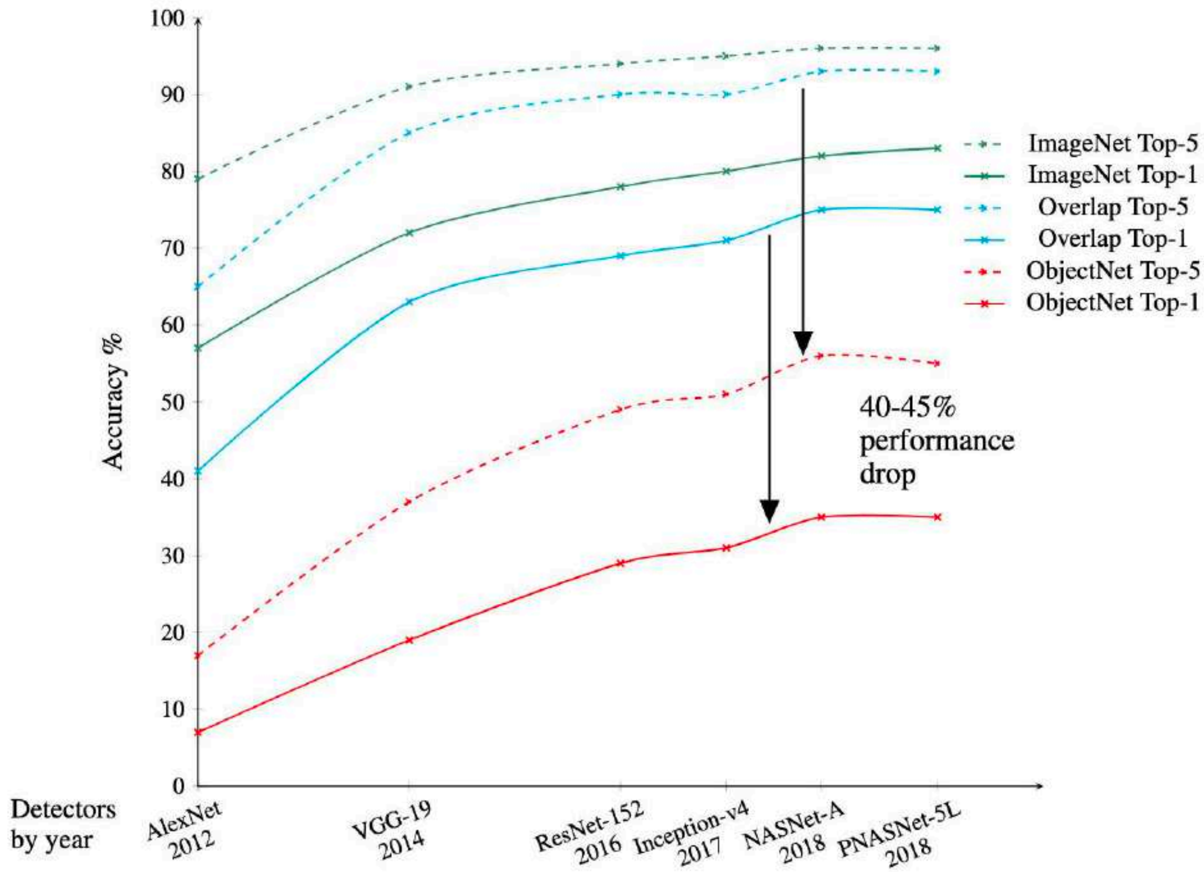

Original Top-3 inferred captions:

1. A red stop sign sitting on the side of a road.
2. A stop sign on the corner of a street.
3. A red stop sign sitting on the side of a street.

Adversarial Top-3 captions:

1. A brown teddy bear laying on top of a bed.
2. A brown teddy bear sitting on top of a bed.
3. A large brown teddy bear laying on top of a bed.

Chen et al. 2018

Pin-yu Chen
IBM

Xu et al. 2019

How many blocks are on the right of the three-level tower?

Will the block tower fall if the top block is removed?

Are there more trees than animals?

What is the shape of the object closest to the large cylinder?

Neural Networks / Deep Learning

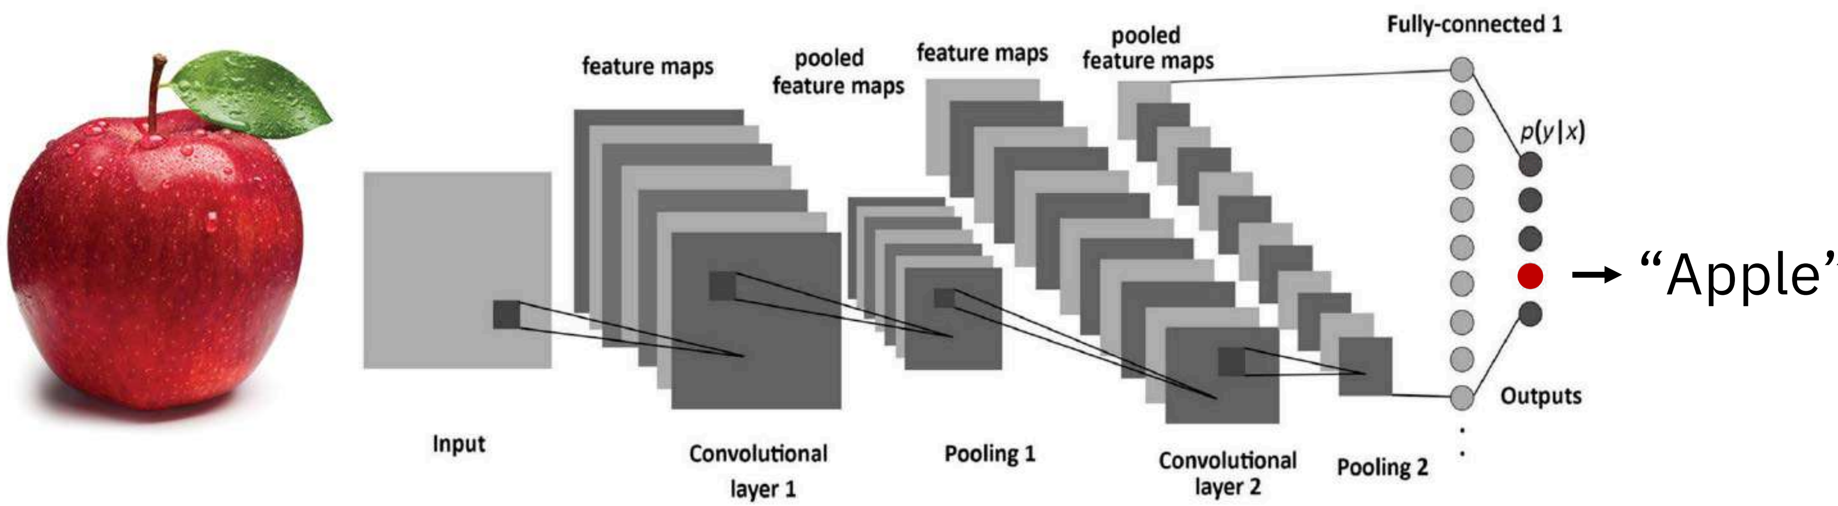

Symbolic AI

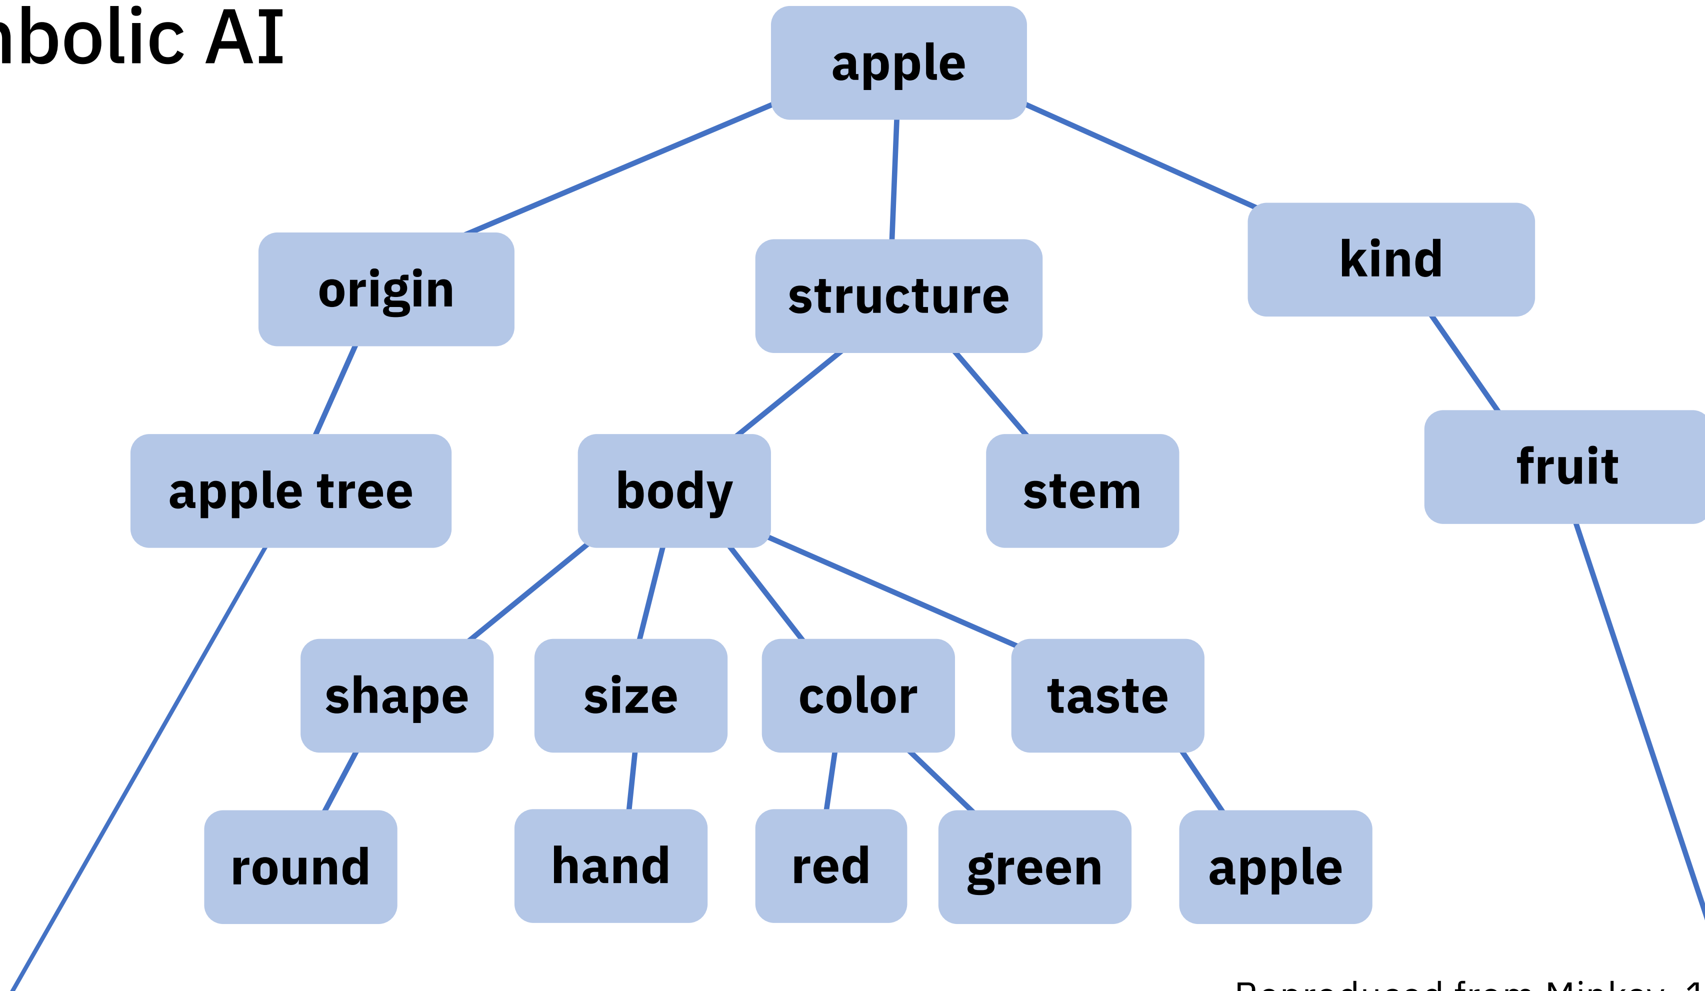

Reproduced from Minksy, 1991

Neural-symbolic AI

Disentangling reasoning from vision and language understanding

Jiajun Wu

Chuang Gan

Joshua Tenenbaum

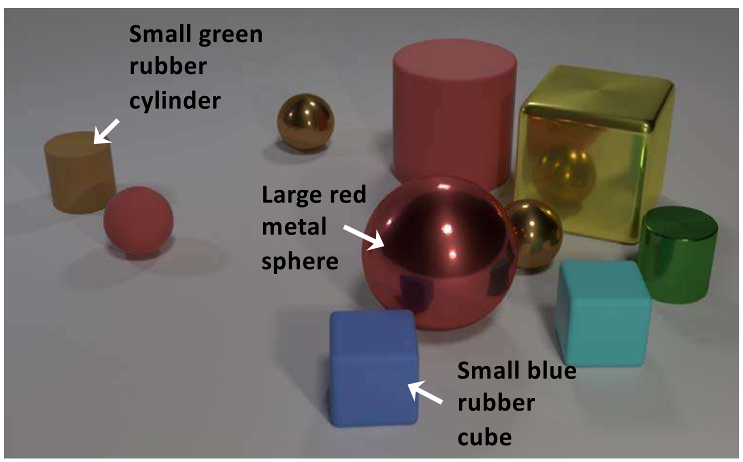

Question: *Are there an equal number of large things and metal spheres?*

Program: equal_number(count(filter_size(Scene, Large)), count(filter_material(filter_shape(Scene, Sphere), Metal)))

Answer: *Yes*

End-to-End Visual Reasoning

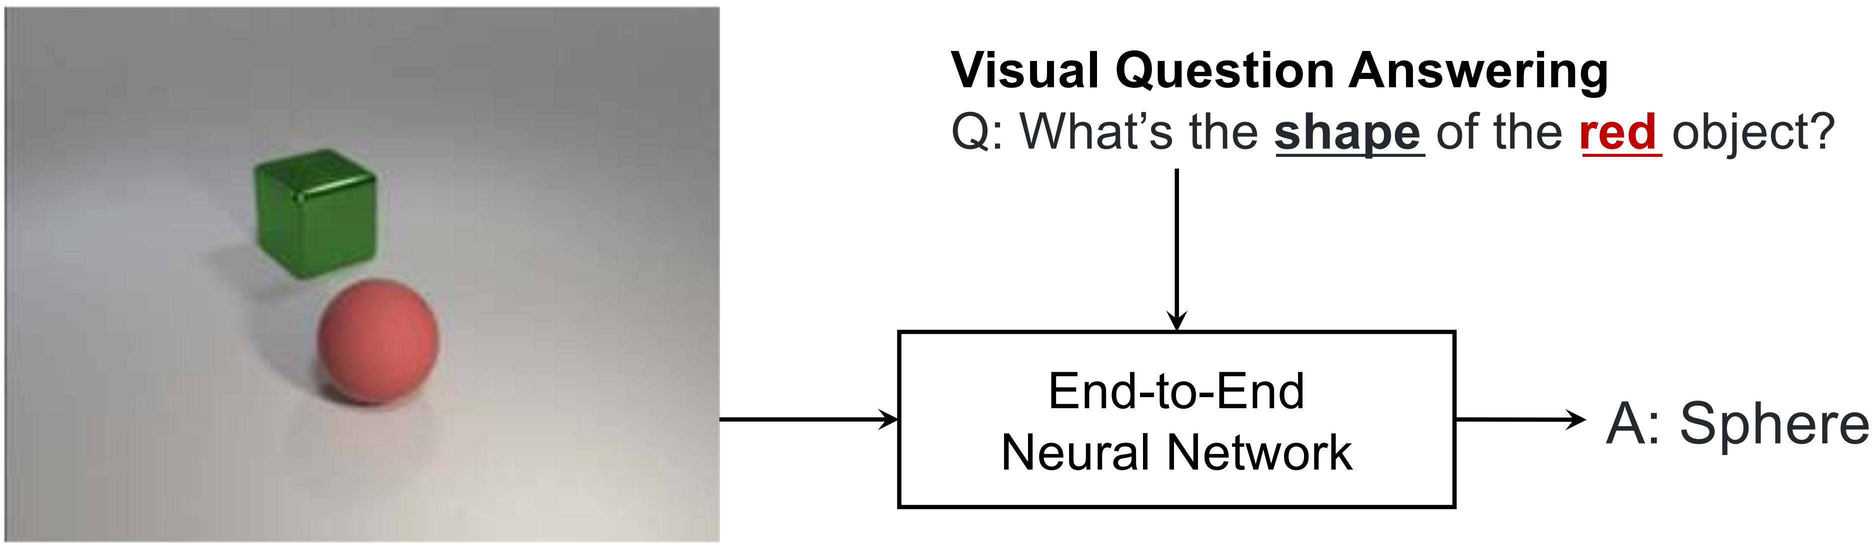

NMN [Andreas et al., 2016]
IEP [Johnson et al., 2017]
FiLM [Perez et al., 2018],
MAC [Hudson & Manning, 2018]
Stack-NMN [Hu et al., 2018]
TbD [Mascharka et al. 2018]

End-to-End Visual Reasoning

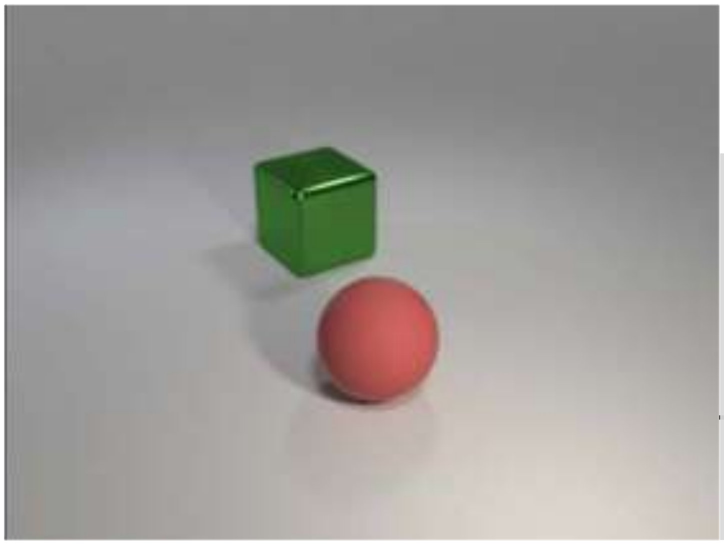

Visual Question Answering

Q: What's the **shape** of the **red** object?

Concept
(e.g., colors, shapes)

Reasoning
(e.g., count)

NMN [Andreas et al., 2016]
IEP [Johnson et al., 2017]
FiLM [Perez et al., 2018],
MAC [Hudson & Manning, 2018]
Stack-NMN [Hu et al., 2018]
TbD [Mascharka et al. 2018]

End-to-End Visual Reasoning

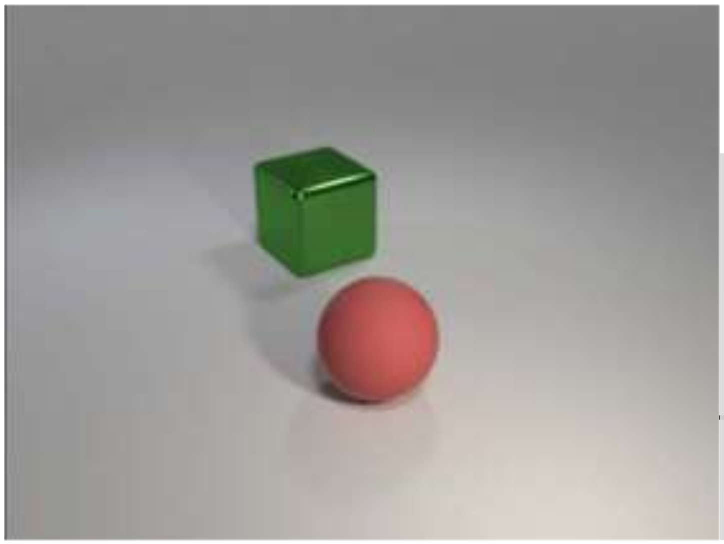

Visual Question Answering
Q: What's the **shape** of the **red** object?

Entangled

Concept ⟷ Reasoning

(e.g., colors, shapes) (e.g., count)

NMN [Andreas et al., 2016]
IEP [Johnson et al., 2017]
FiLM [Perez et al., 2018],
MAC [Hudson & Manning, 2018]
Stack-NMN [Hu et al., 2018]
TbD [Mascharka et al. 2018]

End-to-End Visual Reasoning

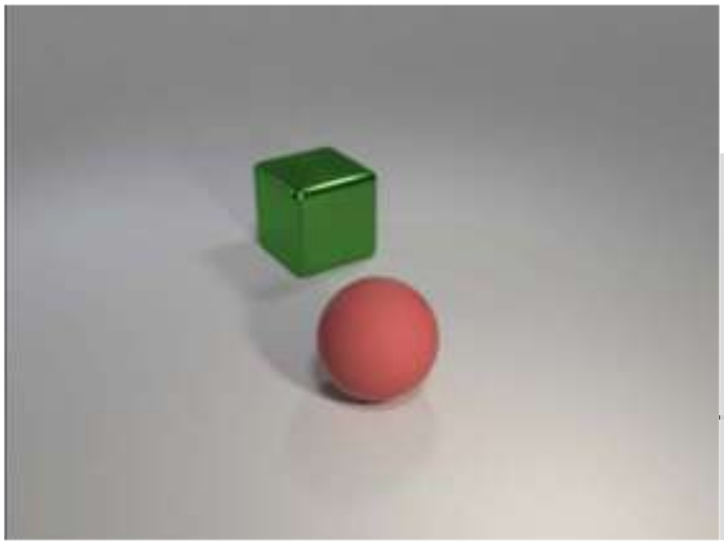

Visual Question Answering

Q: What's the **shape** of the **red** object?

Entangled

Concept ⟷ Reasoning

(e.g., colors, shapes) (e.g., count)

Hard to transfer

Image Captioning

Instance Retrieval

NMN [Andreas et al., 2016]
IEP [Johnson et al., 2017]
FiLM [Perez et al., 2018],
MAC [Hudson & Manning, 2018]
Stack-NMN [Hu et al., 2018]
TbD [Mascharka et al. 2018]

Task: **Visual Reasoning**

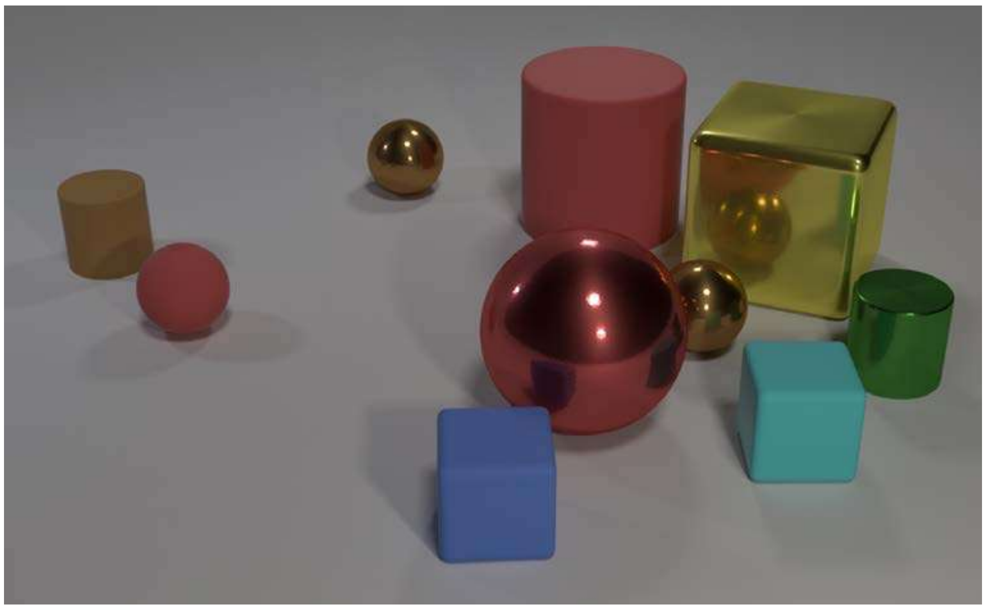

Question: *Are there an equal number of large things and metal spheres?*

Task: **Visual Reasoning**

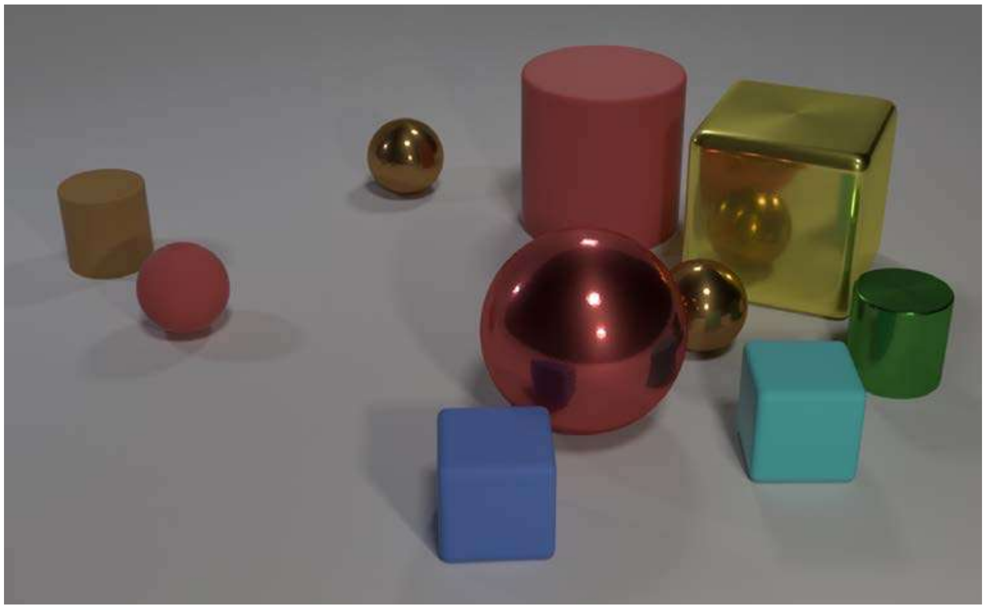

Question: *Are there an equal number of large things and metal spheres?*

Task: **Visual Reasoning**

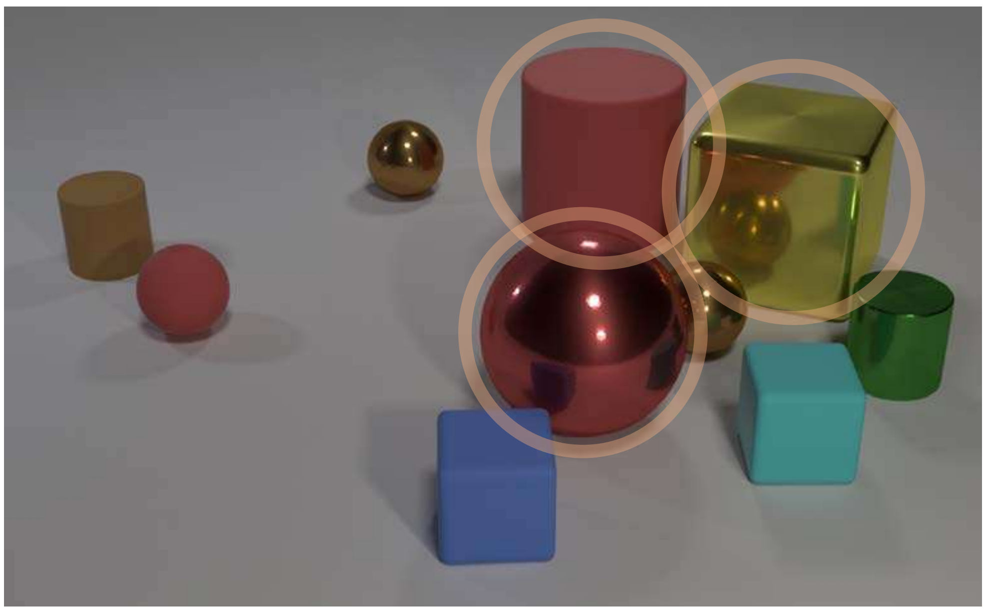

Question: *Are there an equal number of* ***large things*** *and metal spheres?*

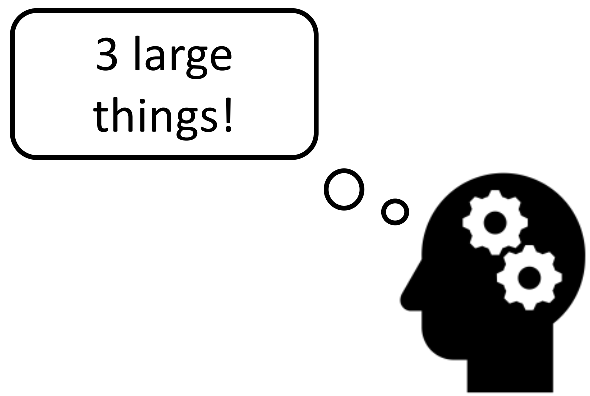

Task: **Visual Reasoning**

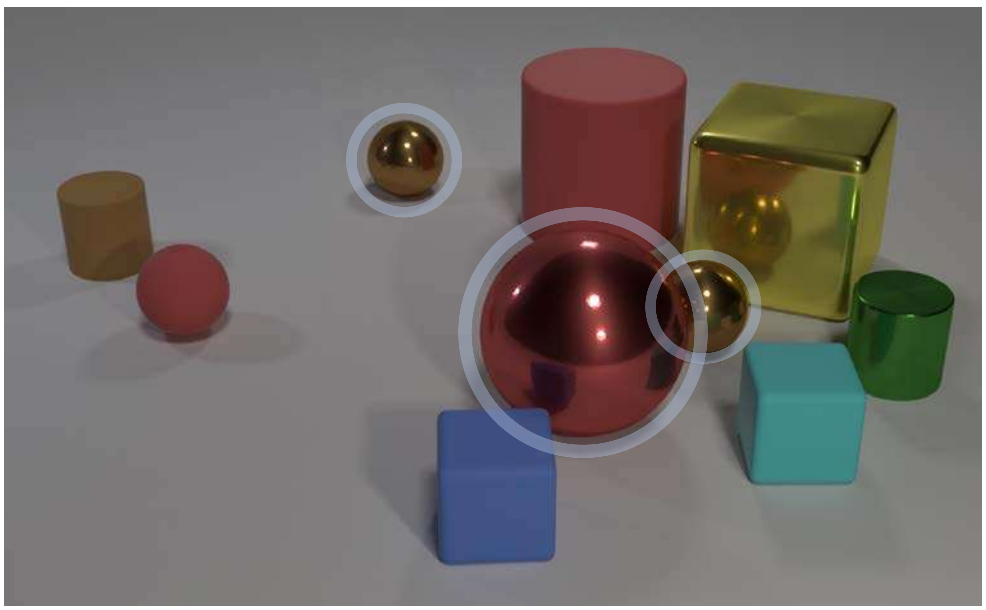

Question: *Are there an equal number of large things and* ***metal spheres****?*

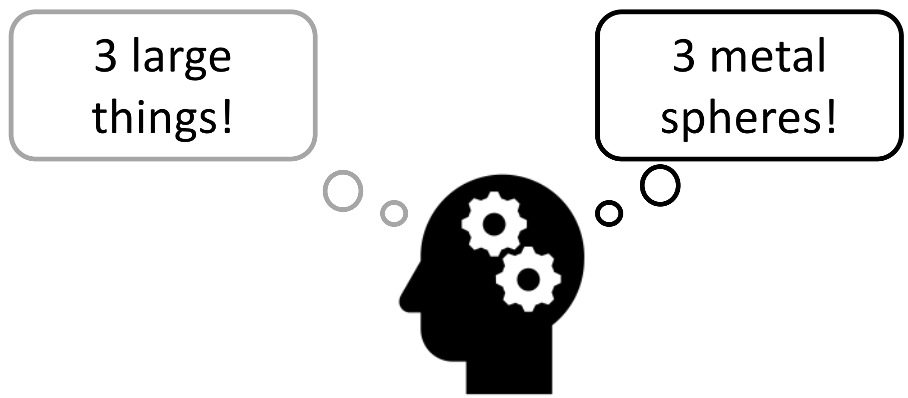

Task: **Visual Reasoning**

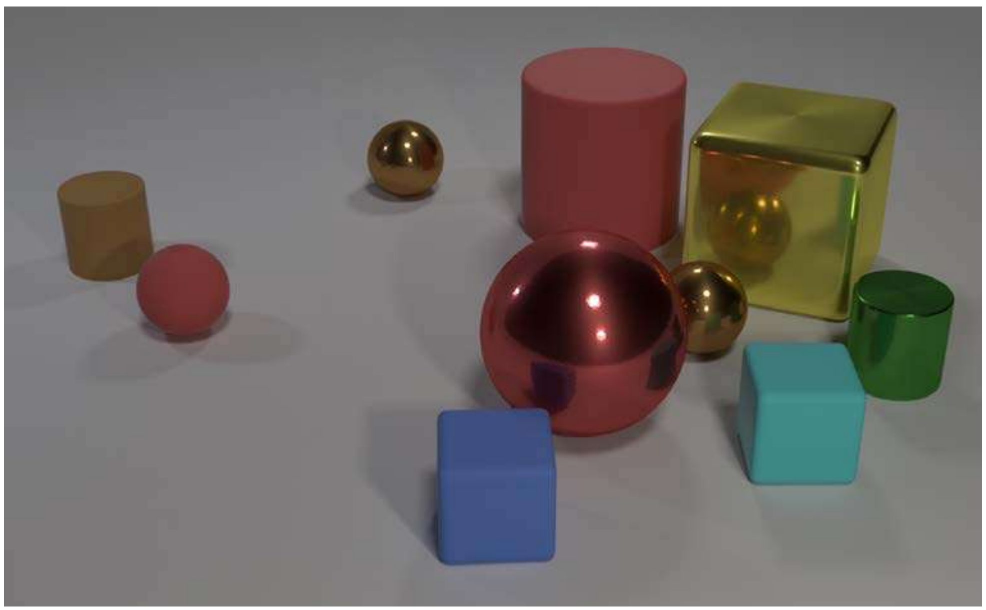

Question: *Are there an* ***equal number*** *of large things and metal spheres?*

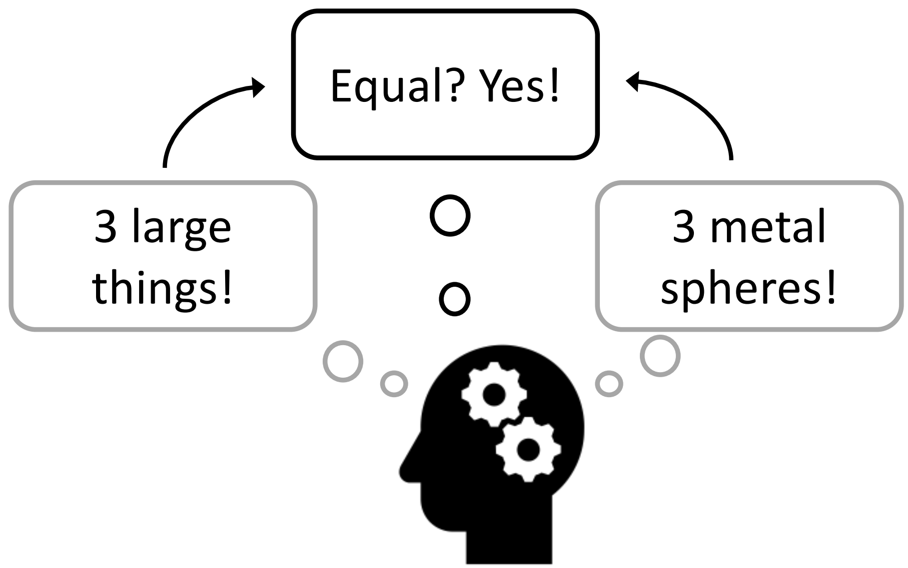

Task: **Visual Reasoning**

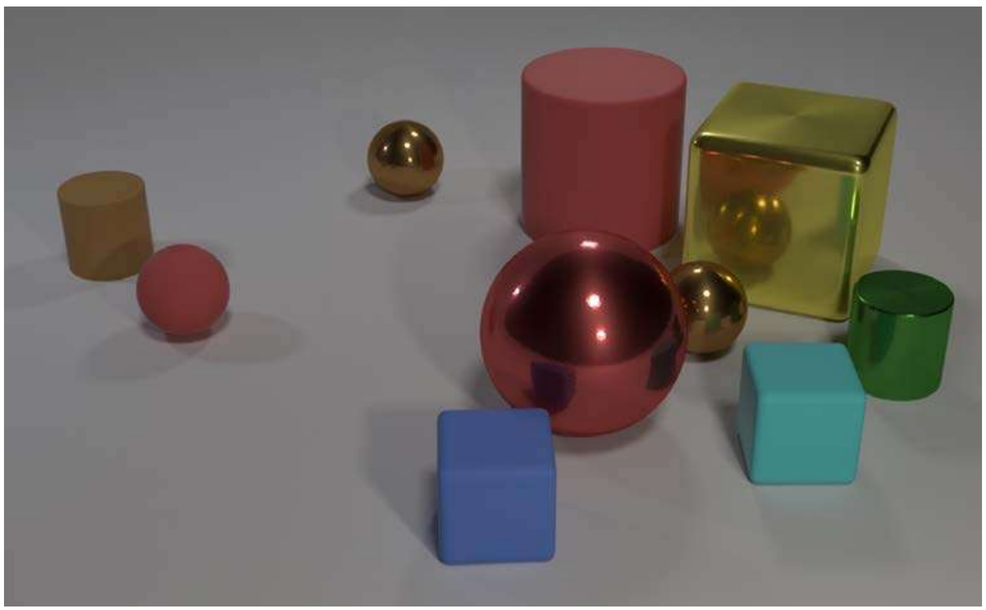

Visual Perception

Question Understanding

Question: *Are there an equal number of large things and metal spheres?*

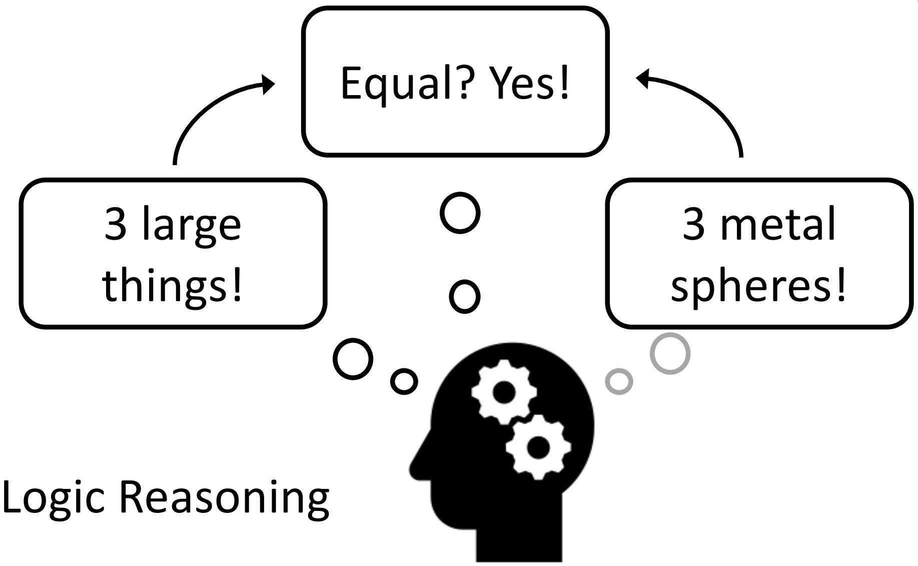

Logic Reasoning

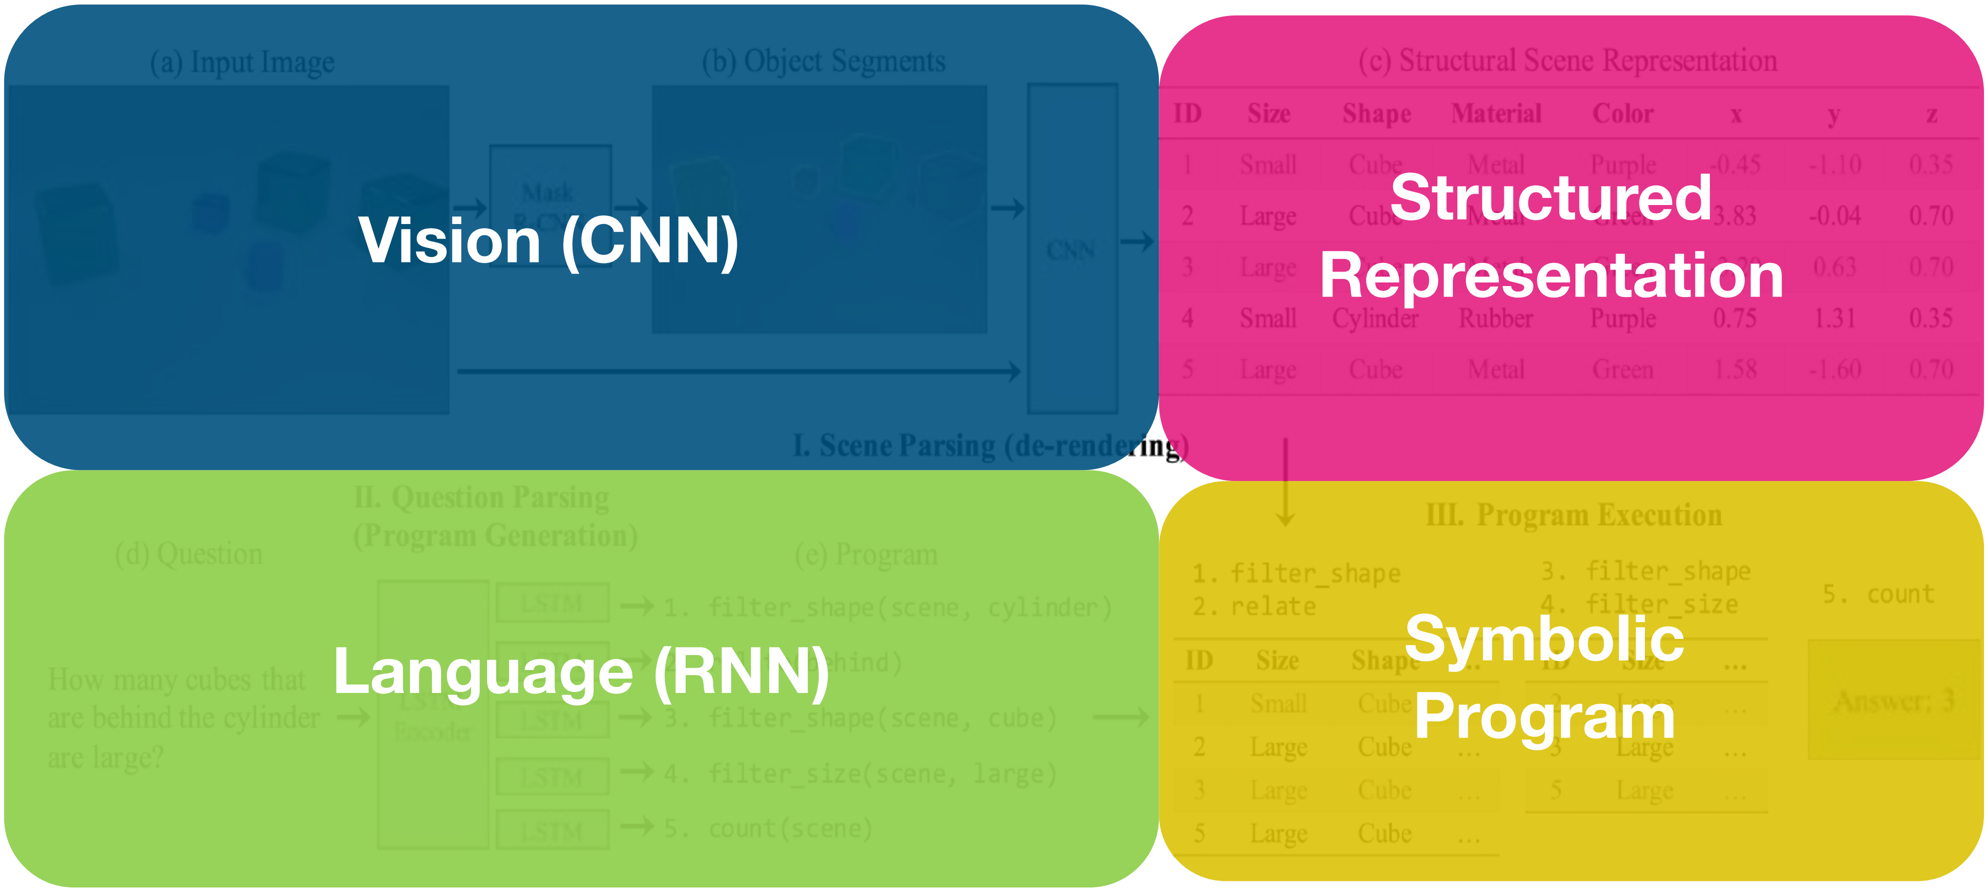
Vision (CNN)
Structured Representation
Language (RNN)
Symbolic Program

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

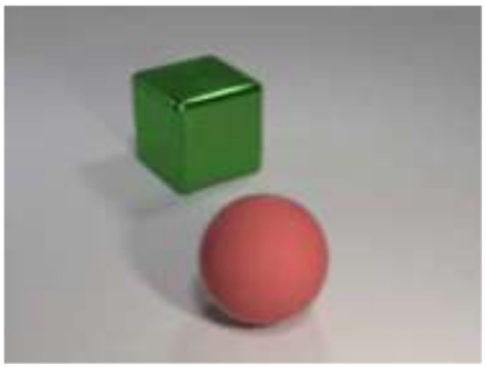

Scene Parsing →

Language

Q: What's the **shape** of the **red** object?

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

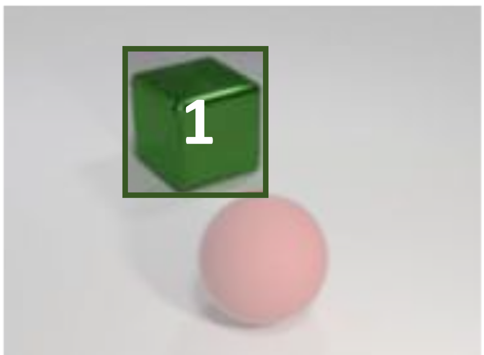

Scene Parsing

ID	Color	Shape	Material
1	Green	Cube	Metal

Language

Q: What's the **shape** of the **red** object?

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

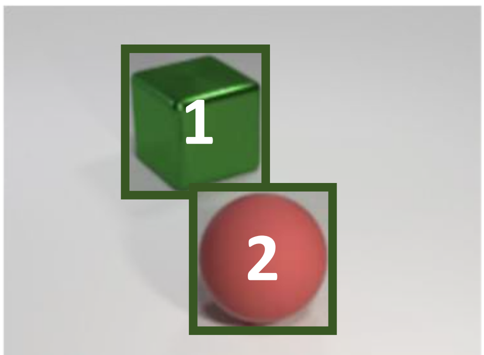

Scene Parsing

ID	Color	Shape	Material
1	Green	Cube	Metal
2	Red	Sphere	Rubber

Language

Q: What's the **shape** of the **red** object?

Semantic Parsing

```
Filter(Red)
↓
Query(Shape)
```

Program

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

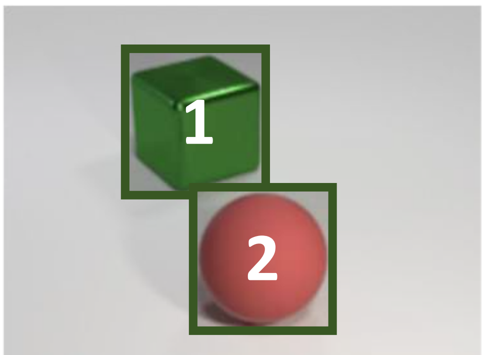

Scene Parsing →

ID	Color	Shape	Material
1	Green	Cube	Metal
2	Red	Sphere	Rubber

Language

Q: What's the **shape** of the **red** object?

Semantic Parsing →

```
Filter(Red)
↓
Query(Shape)
```

Program

Symbolic Reasoning

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

Scene Parsing

ID	Color	Shape	Material
1	Green	Cube	Metal
2	Red	Sphere	Rubber

Symbolic Reasoning

Language

Q: What's the shape of the red object?

Semantic Parsing

```
Filter(Red)
Query(Shape)
```

Program

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

Scene Parsing

ID	Color	Shape	Material
1	Green	Cube	Metal
2	*Red*	Sphere	Rubber

Symbolic Reasoning

Language

Q: What's the shape of the red object?

Semantic Parsing

```
Filter(Red)
Query(Shape)
```

Program

NS-VQA [Yi et al. 2018]

Incorporate Concepts in Visual Reasoning

Vision

Scene Parsing

ID	Color	Shape	Material
1	Green	Cube	Metal
2	Red	*Sphere*	Rubber

Language

Q: What's the shape of the red object?

Semantic Parsing

`Filter(Red)`

`Query(Shape)`

Program

Symbolic Reasoning

Sphere

Advantage 1: **High Accuracy**

Method	Accuracy (%)
Human	92.6
RN	95.5
IEP	96.9
FiLM	97.6
MAC	98.9
TbD	99.1
NS-VQA (Ours)	**99.8**

Effectively perfect!

[Yi et al. NeurIPS 2018, Johnson et al. ICCV 2017, Santoro et al. NIPS 2017, Perez et al. AAAI 2018, Hudson et al. ICLR 2018, Mascharka et al. CVPR 2018]

Advantage 2: **Data Efficiency**

High accuracy when trained with just 1% the of the data that other methods require

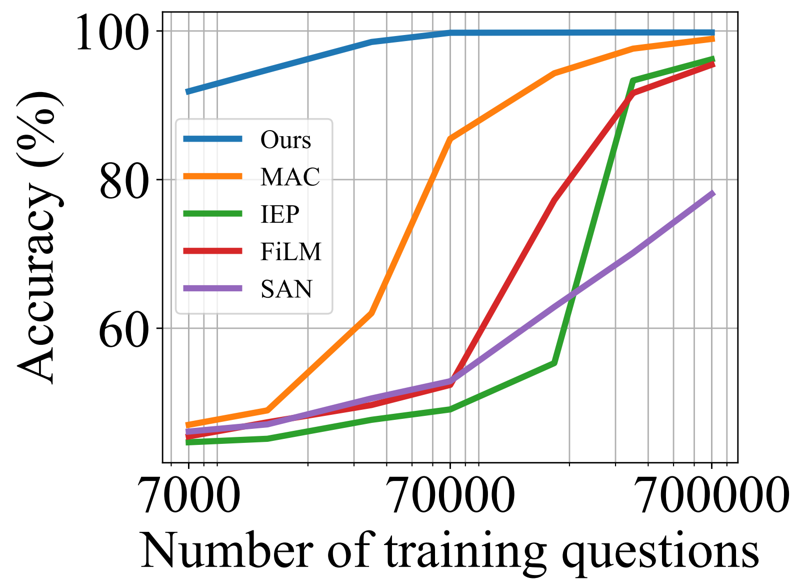

[Yi et al. NeurIPS 2018]

Advantage 3: **Transparency and Interpretability**

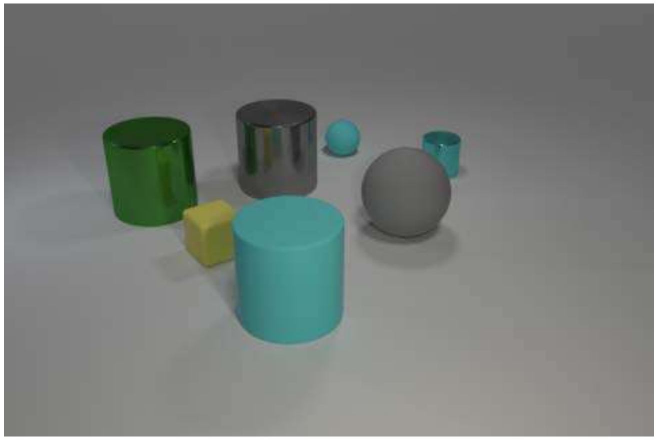

Question: Are there more yellow matte things that are right of the gray ball than cyan metallic objects?

```
scene
filter_cyan
filter_metal
count
...(4 modules)
scene
filter_yellow
filter_rubber
count
greater_than
```

Answer: no

[Yi et al. NeurIPS 2018, Johnson et al. ICCV 2017]

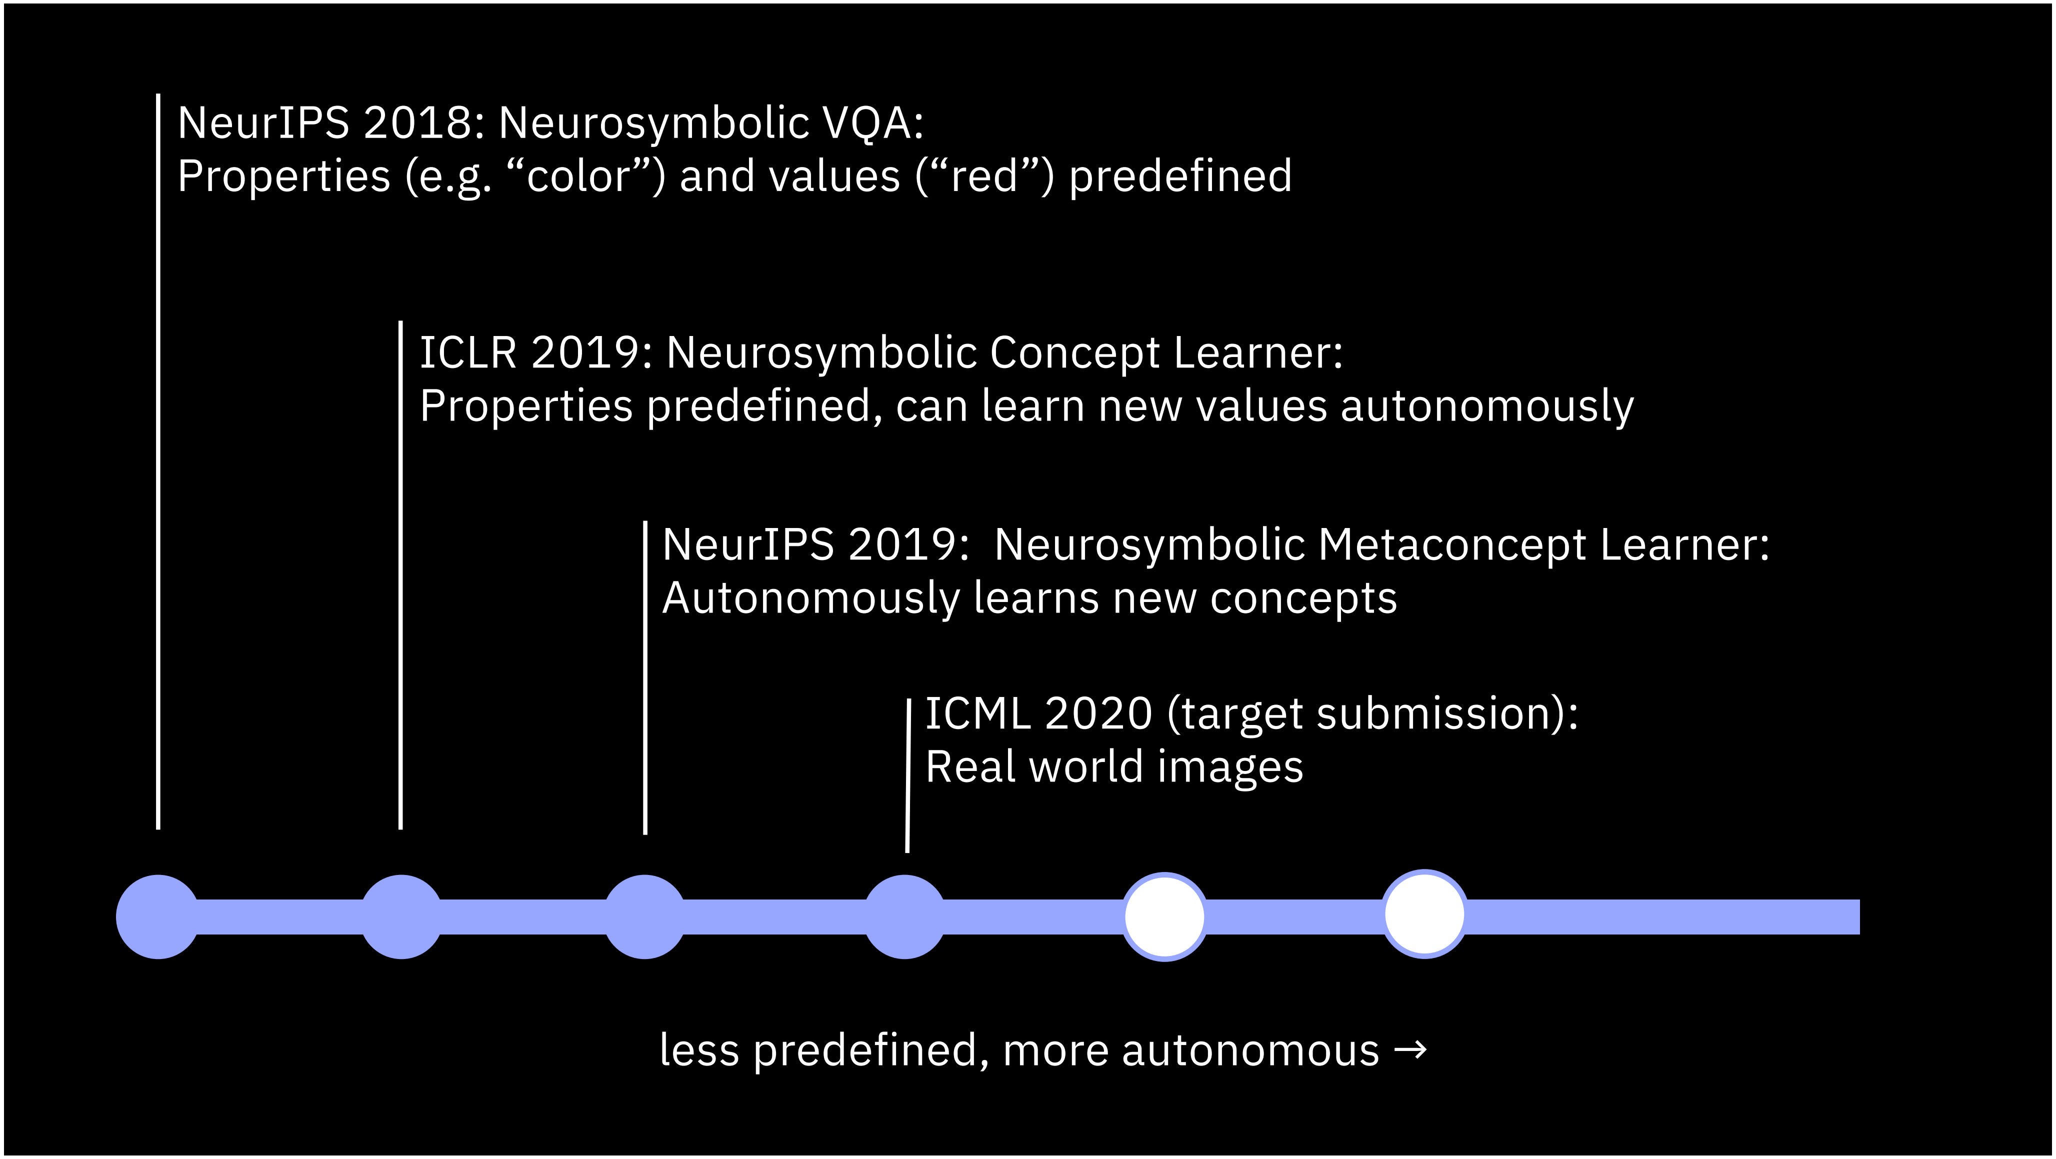
NeurIPS 2018: Neurosymbolic VQA:
Properties (e.g. “color”) and values (“red”) predefined
ICLR 2019: Neurosymbolic Concept Learner:
Properties predefined, can learn new values autonomously
NeurIPS 2019: Neurosymbolic Metaconcept Learner:
Autonomously learns new concepts
ICML 2020 (target submission):
Real world images
less predefined, more autonomous →

Neuro-Symbolic Concept Learning

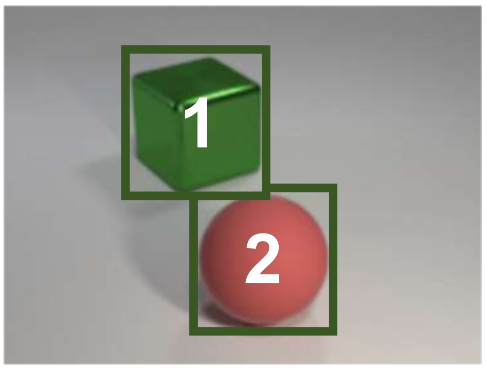

Q: What's the **shape** of the **red** object?

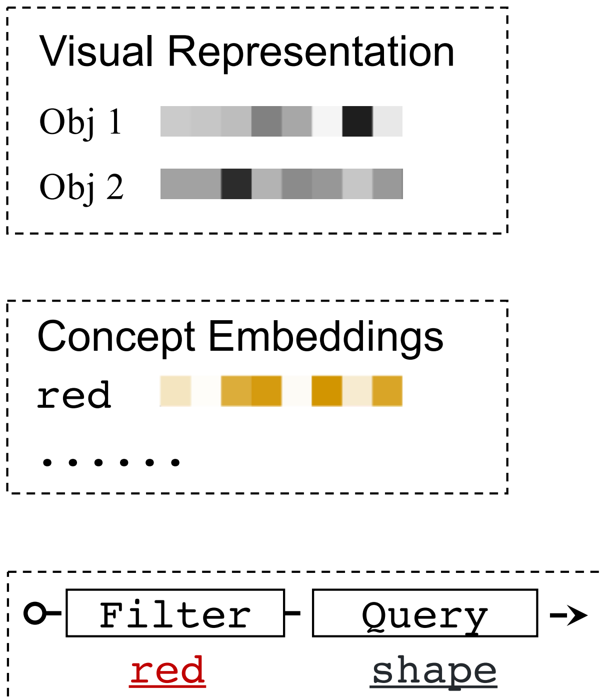

Neuro-Symbolic Concept Learning

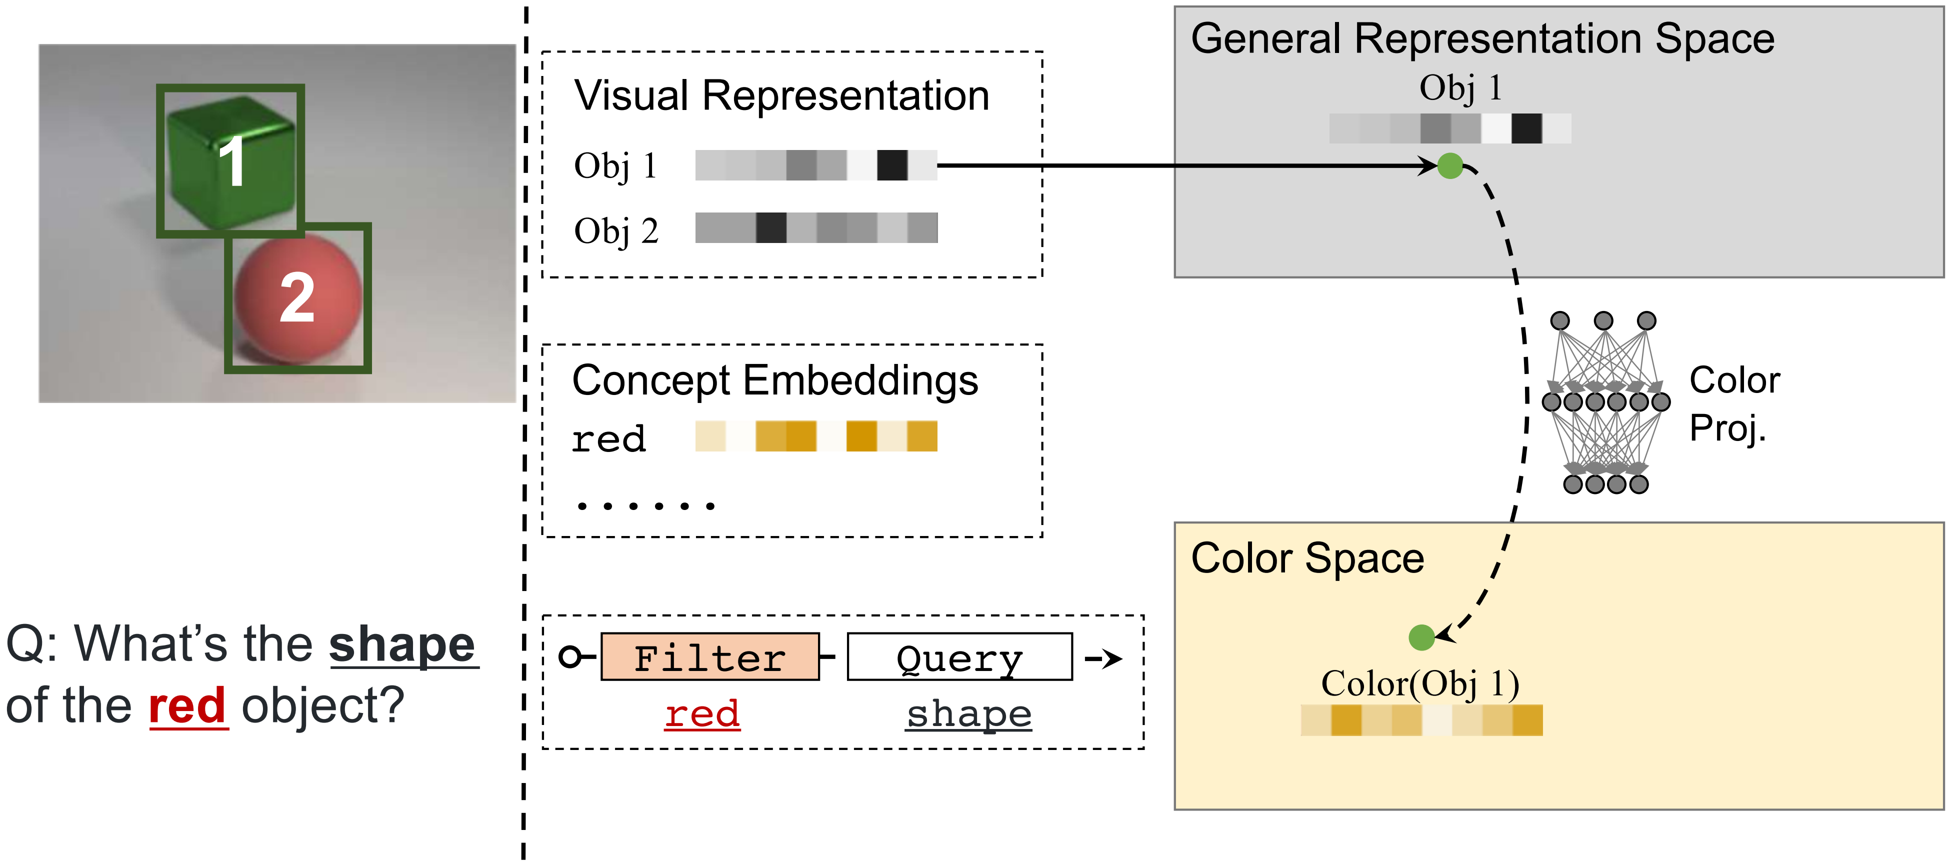

Neuro-Symbolic Concept Learning

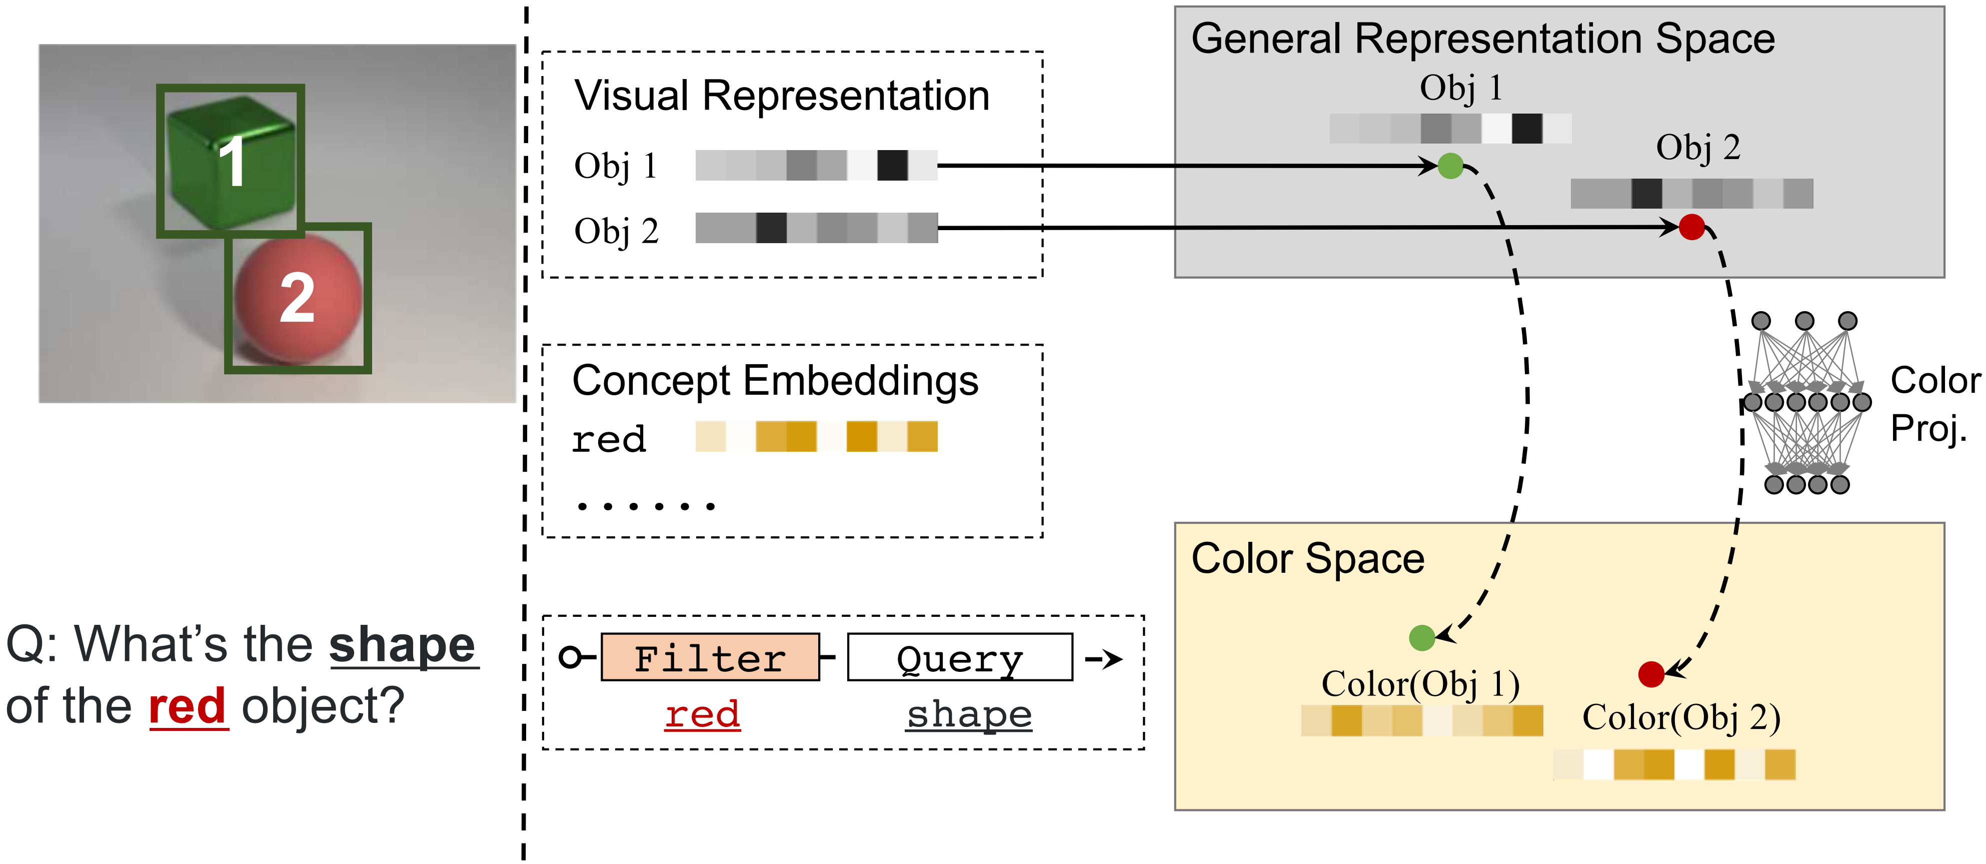

Neuro-Symbolic Concept Learning

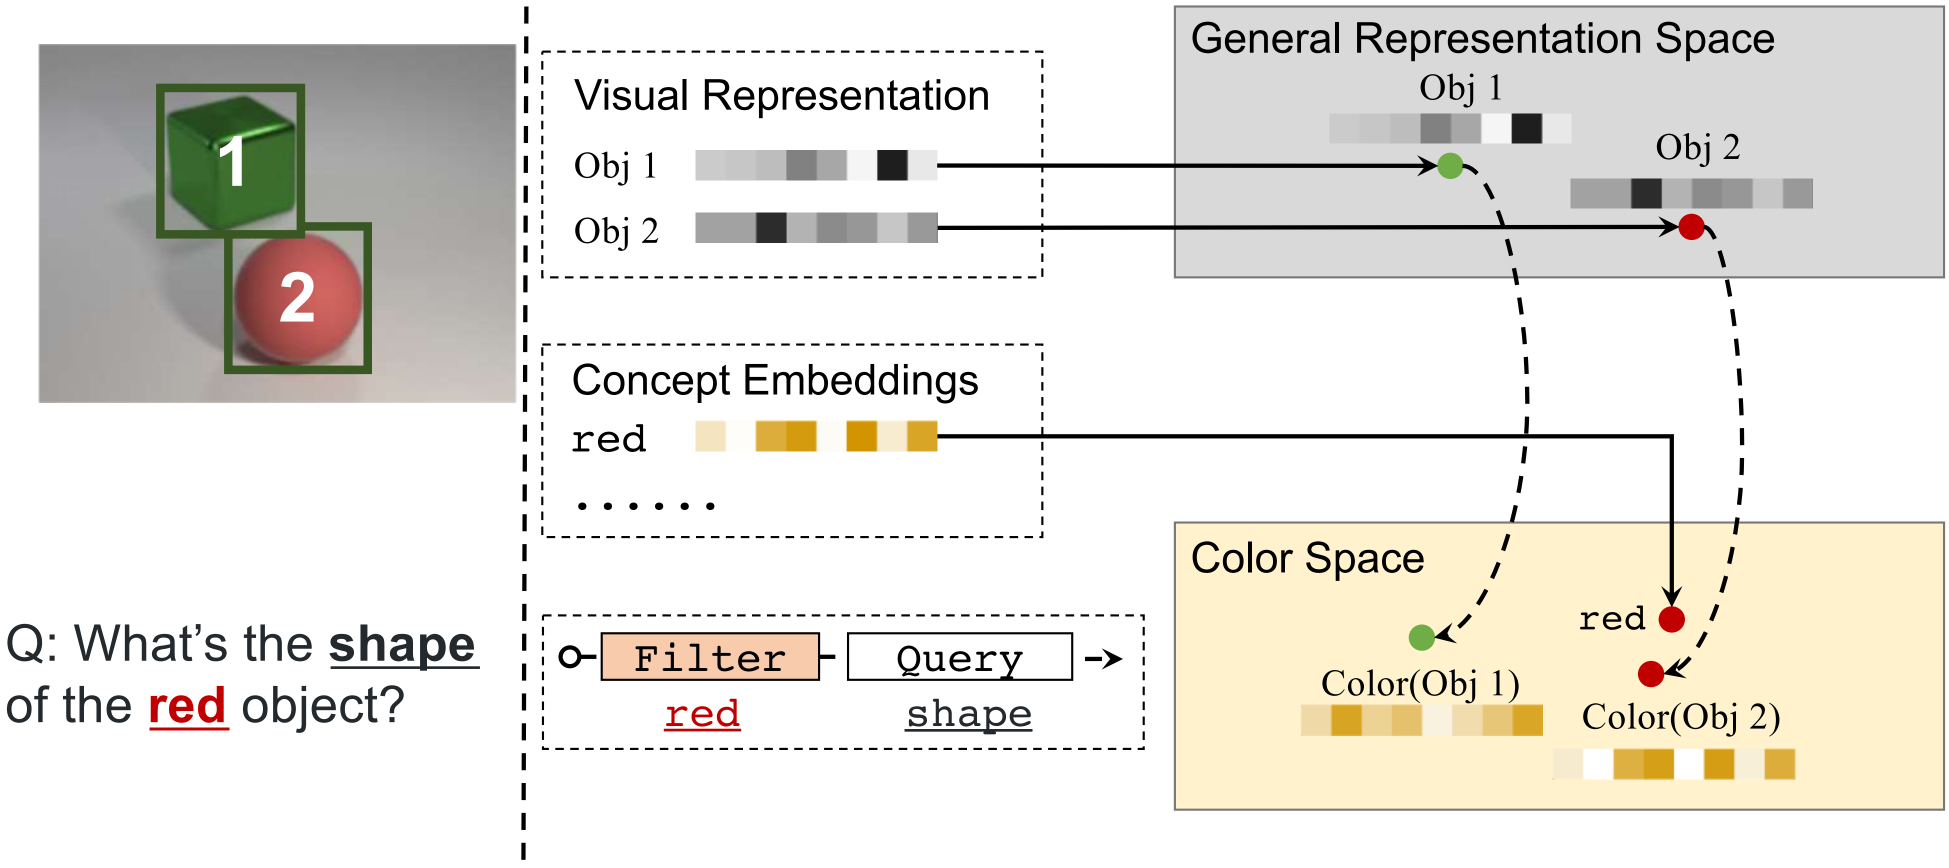

Symbolic AI

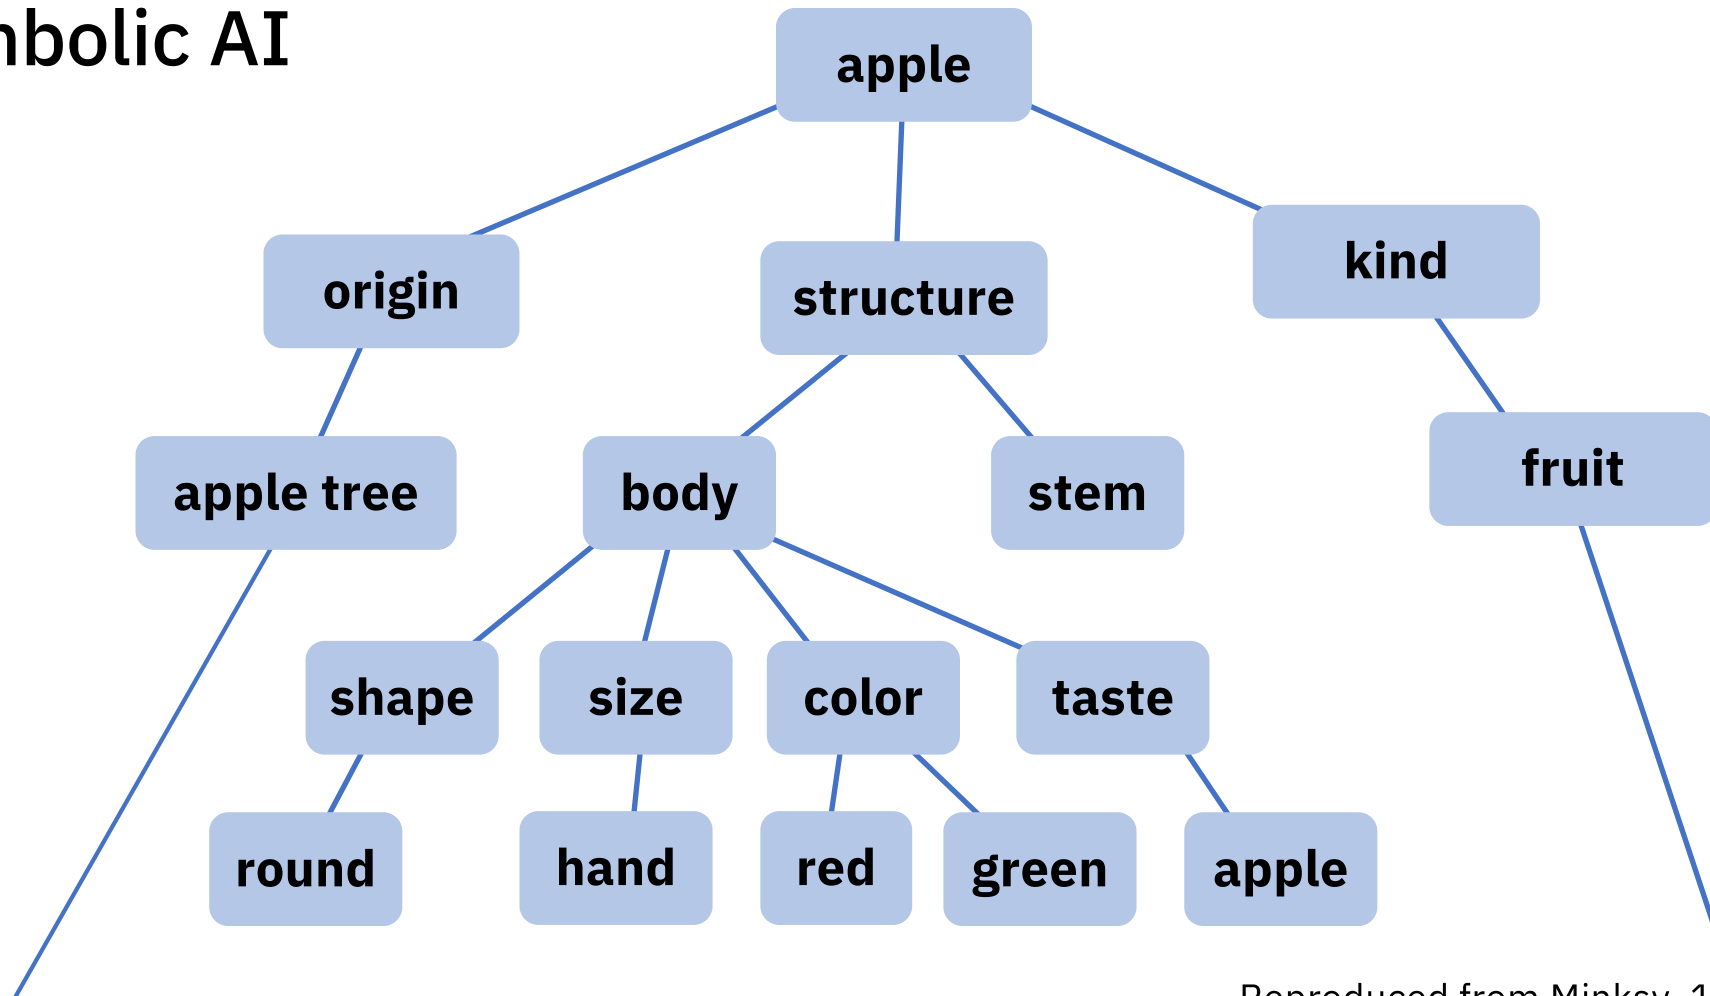

Reproduced from Minksy, 1991

Meta-concept Learning

Han et al. NeurIPS 2019

Visual reasoning questions

color:
red

Q: Is there any **red cube**?
A: Yes.

color:
green

Q: Is there any **green block**?
A: Yes

CLEVR
(Johnson et al. 2017)

Laridae

Ivory Gull

Q: Is there any **Ivory Gull**?
A: Yes.
Q: Is there any **Laridae**?
A: Yes.

Black Tern

Q: Is there any **Black Tern**?
A: Yes.
Q: Is there any **Laridae**?
A: Yes.

CUB
(Wah et al. 2011)

+ Metaconcept questions

Q: Is red a **same kind** of concept as green?
A: Yes.

Q: Is cube a **synonym** of block?
A: Yes.

Q: Is Laridae a **hypernym** of Ivory gull?
A: Yes.

Augmenting VQA with Metaconcepts

Visual reasoning questions

color:
red

Q: Is there any **red cube**?
A: Yes.

color:
green

Q: Is there any **green block**?
A: Yes

CLEVR
(Johnson et al. 2017)

Laridae

Ivory Gull

Q: Is there any **Ivory Gull**?
A: Yes.
Q: Is there any **Laridae**?
A: Yes.

Black Tern

Q: Is there any **Black Tern**?
A: Yes.
Q: Is there any **Laridae**?
A: Yes.

CUB
(Wah et al. 2011)

+ Metaconcept questions

Q: Is red a **same kind** of concept as green?
A: Yes.

Q: Is cube a **synonym** of block?
A: Yes.

Q: Is Laridae a **hypernym** of Ivory gull?
A: Yes.

Program Execution Animated

Visual reasoning questions

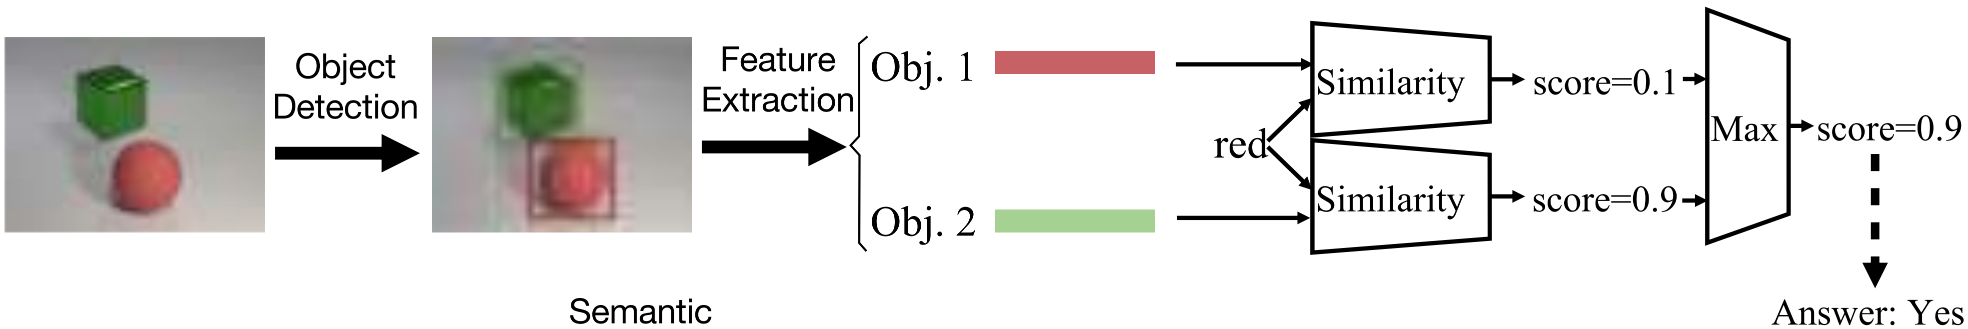

Q: Is there any red object? — Semantic Parsing → P: **Exist(Filter(red))**

Metaconcept questions

Q: Is red a same kind of concept as yellow? — Semantic Parsing → P: **MetaVerify(**
red, yellow, same-kind
)

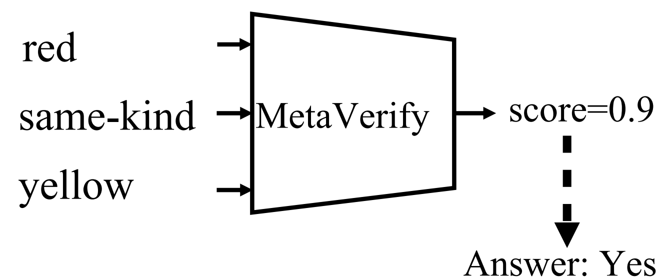

Generalization

Metaconcept Generalization

Q: Is there any *airplane*?
A: Yes

Q: Is there any *kid*?
A: Yes

Q: Is airplane a *synonym* of plane?
A: Yes

Q: Is there any *plane*?
A: Yes

Q: Is there any *child*?
A: Yes

Q: Is kid a *synonym* of child?
A: Yes

Training

Testing: metaconcepts on unseen pairs of concepts

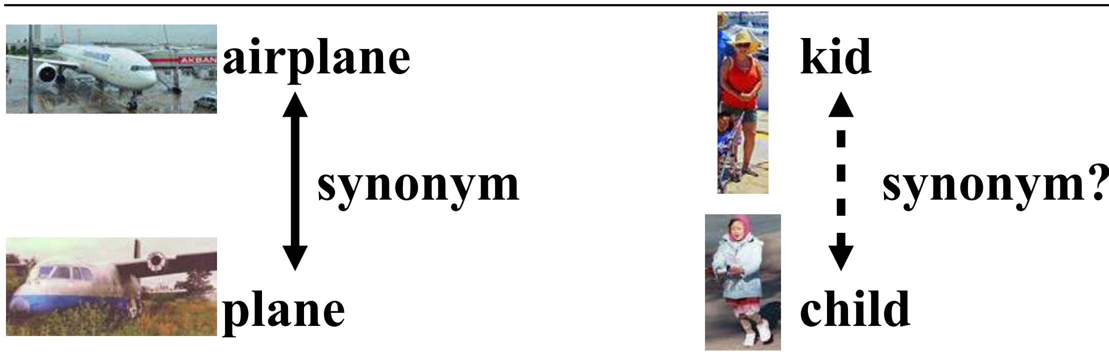

Generalization

Metaconcept Generalization: Results

Q: Is there any *airplane*?
A: Yes

Q: Is there any *plane*?
A: Yes

Training

Q: Is there any *kid*?
A: Yes

Q: Is there any *child*?
A: Yes

Q: Is airplane a *synonym* of plane?
A: Yes

Q: Is kid a *synonym* of child?
A: Yes

Testing

		Q.Type	GRU (Lang. Only) [Cho et al., 2014]	GRU-CNN [Zhou et al., 2015]	BERT (question ; concept) [Jacob Devlin, 2018]	NS-CL [Mao et al. 2019]	VCML
CLEVR	Synonym	50.0	$66.3_{\pm 1.4}$	$60.9_{\pm 10.6}$	$76.2_{\pm 10.2}$; $80.2_{\pm 16.1}$	$\mathbf{100.0}_{\pm 0.0}$	$\mathbf{100.0}_{\pm 0.0}$
	Same-kind	50.0	$64.7_{\pm 5.1}$	$61.5_{\pm 6.6}$	$75.4_{\pm 5.4}$; $80.1_{\pm 10.0}$	$92.3_{\pm 4.9}$	$\mathbf{99.3}_{\pm 1.0}$
GQA	Synonym	50.0	$80.8_{\pm 1.0}$	$76.2_{\pm 0.8}$	$76.2_{\pm 2.4}$; $83.1_{\pm 1.5}$	$81.2_{\pm 2.8}$	$\mathbf{91.1}_{\pm 1.7}$
	Same-kind	50.0	$56.3_{\pm 2.3}$	$57.3_{\pm 5.3}$	$59.5_{\pm 2.7}$; $68.2_{\pm 4.0}$	$66.8_{\pm 4.1}$	$\mathbf{69.1}_{\pm 1.7}$
CUB	Hypernym	50.0	$74.3_{\pm 5.2}$	$76.7_{\pm 8.8}$	$75.6_{\pm 1.2}$; $61.7_{\pm 10.3}$	$80.1_{\pm 7.3}$	$\mathbf{94.8}_{\pm 1.3}$
	Meronym	50.0	$80.1_{\pm 5.9}$	$78.1_{\pm 4.8}$	$63.1_{\pm 3.2}$; $72.9_{\pm 9.9}$	$\mathbf{97.7}_{\pm 1.1}$	$92.5_{\pm 1.0}$

CLEVERER: CoLlision Events for Video REpresentation and Reasoning

- Descriptive

Q: *What is the material of the last object to collide with the cyan cylinder?*

A: *Metal*

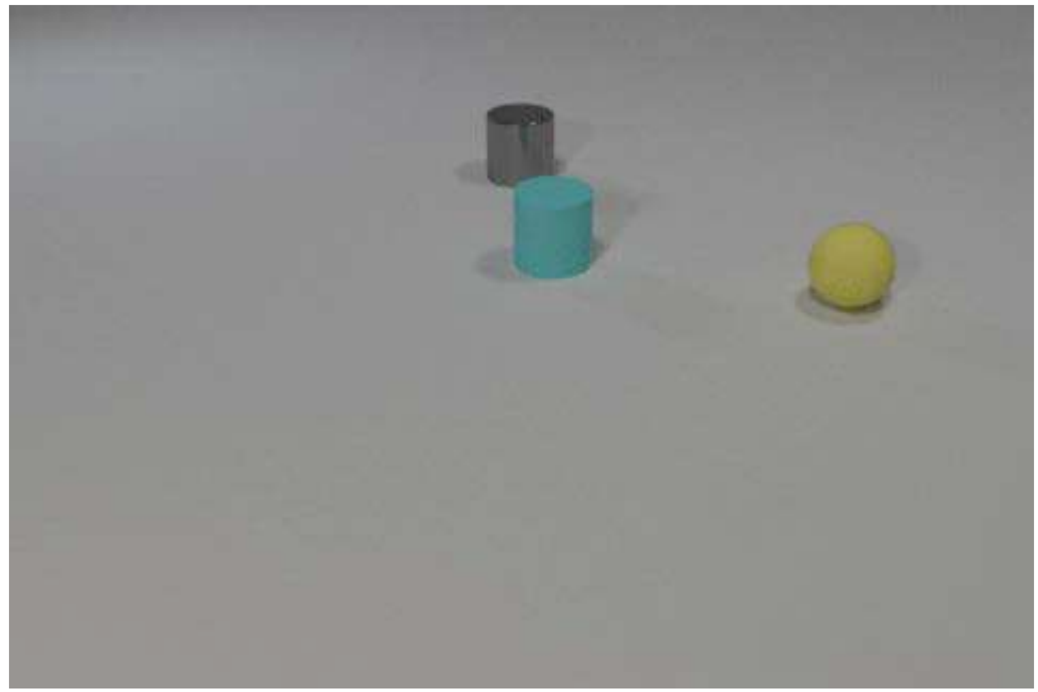

Chuang Gan w/ Kevin Xi, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba & Josh Tenenbaum

- Explanatory

Q: *What is responsible for the collision between the rubber and metal cylinder?*

A. The presence of the yellow sphere
B. The collision between the rubber cylinder and the red rubber sphere

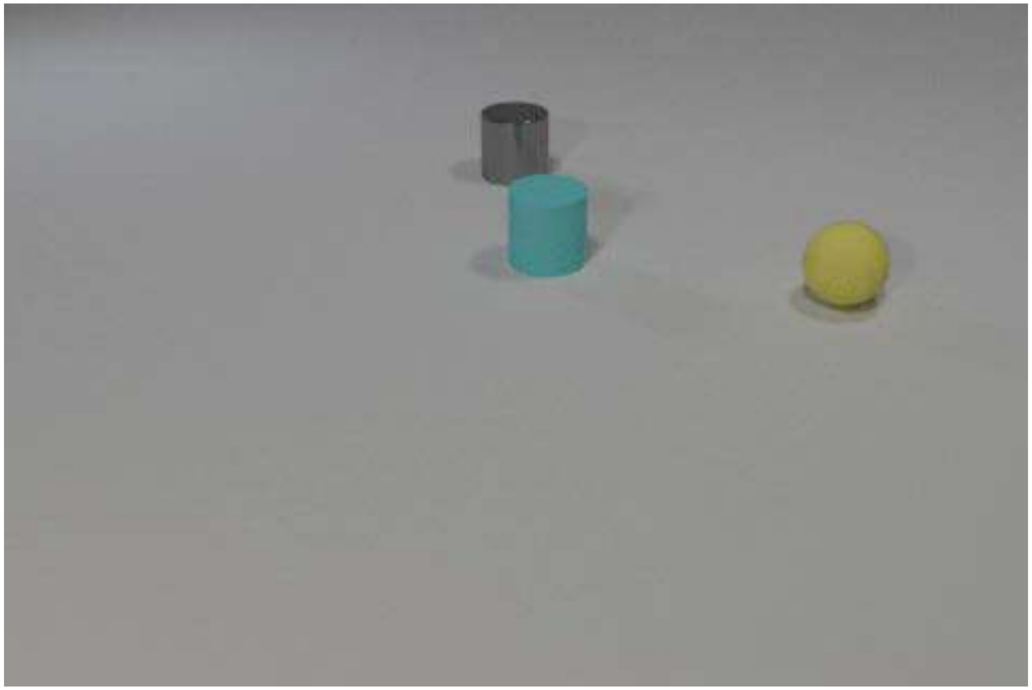

- Counterfactual

Q: *What will happen without the cyan cylinder?*

A. The red rubber sphere and the metal sphere collide
B. The red rubber sphere and the gray object collide

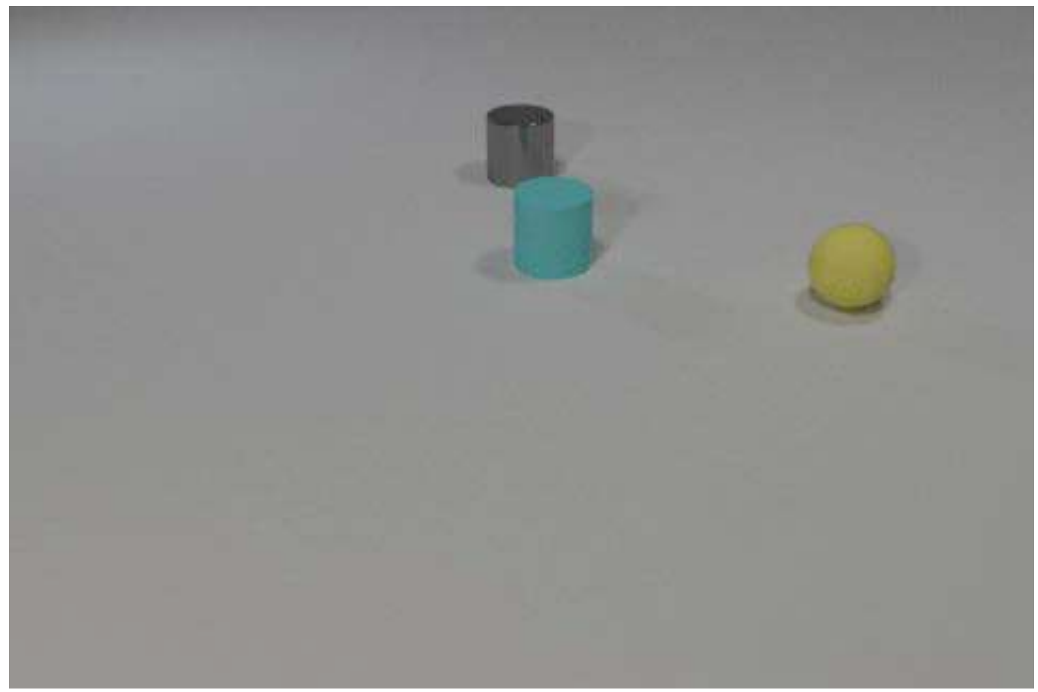

Looking Ahead

How many employees have over 10 years experience but have moved location in the last year?

What factors might contribute to better output from Factory A vs. Factory B?

Why is our database down?

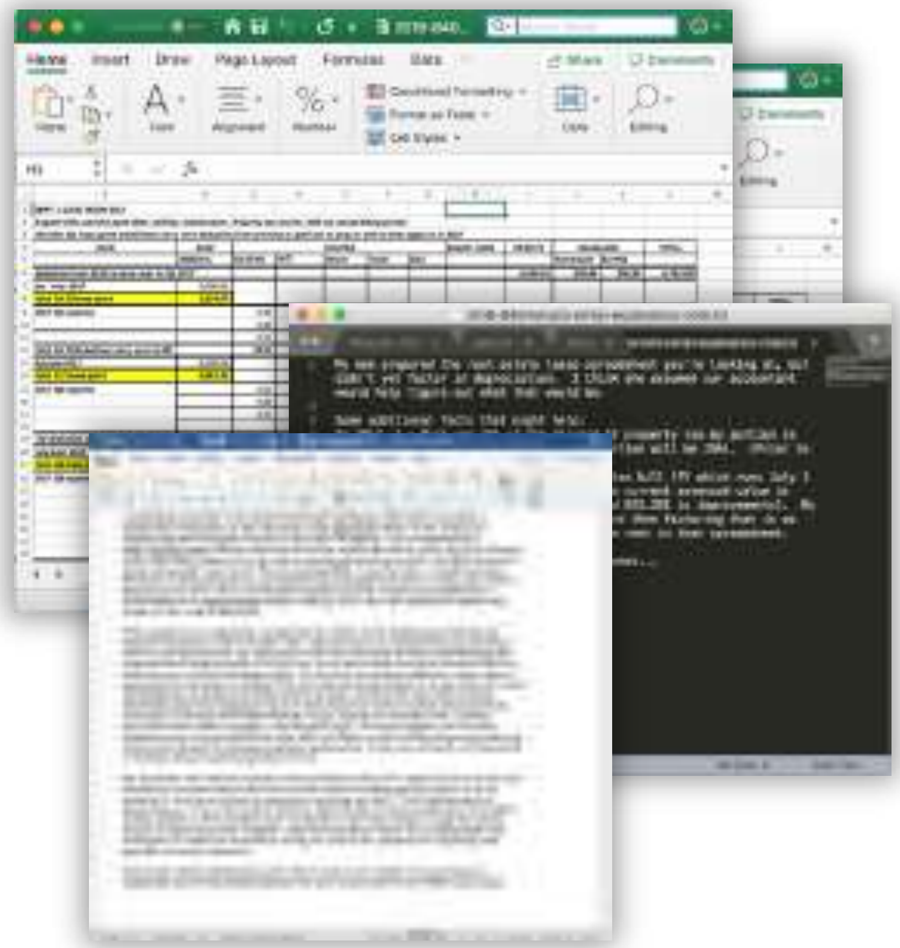

Neurosymbolic **Generative Models**

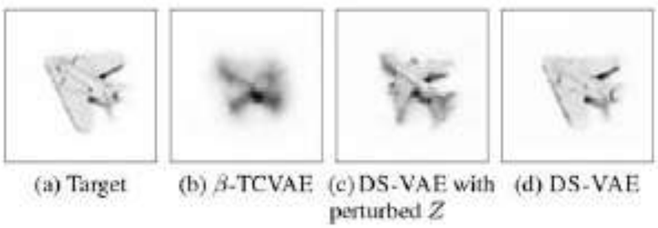

Srivastava et al. 2020 (submitted)

Neurosymbolic **Safe ML/RL**

Fulton et al AAAI 2018

Neurosymbolic **NLU**

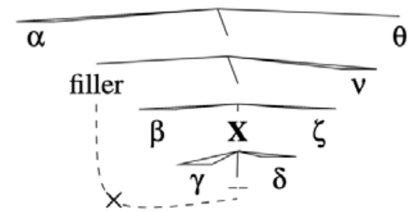

Wilcox et al. NAACL 2019

Neurosymbolic **Planning**

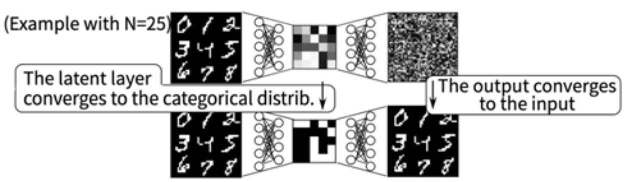

Asai et al. AAAI 2018

Neurosymbolic **Code Optimization**

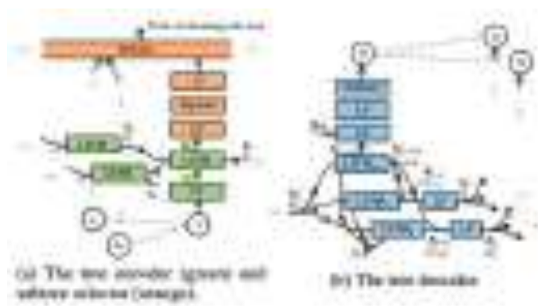

Shi et al. ICLR 2019

Neurosymbolic **Machine Common Sense**

Smith et al. NeurIPS 2019

Inducing Behavioral Insight

Inferring flexible behavioral plans/policies from temporal observation data

Julie Shah
MIT

Christian Muise
IBM

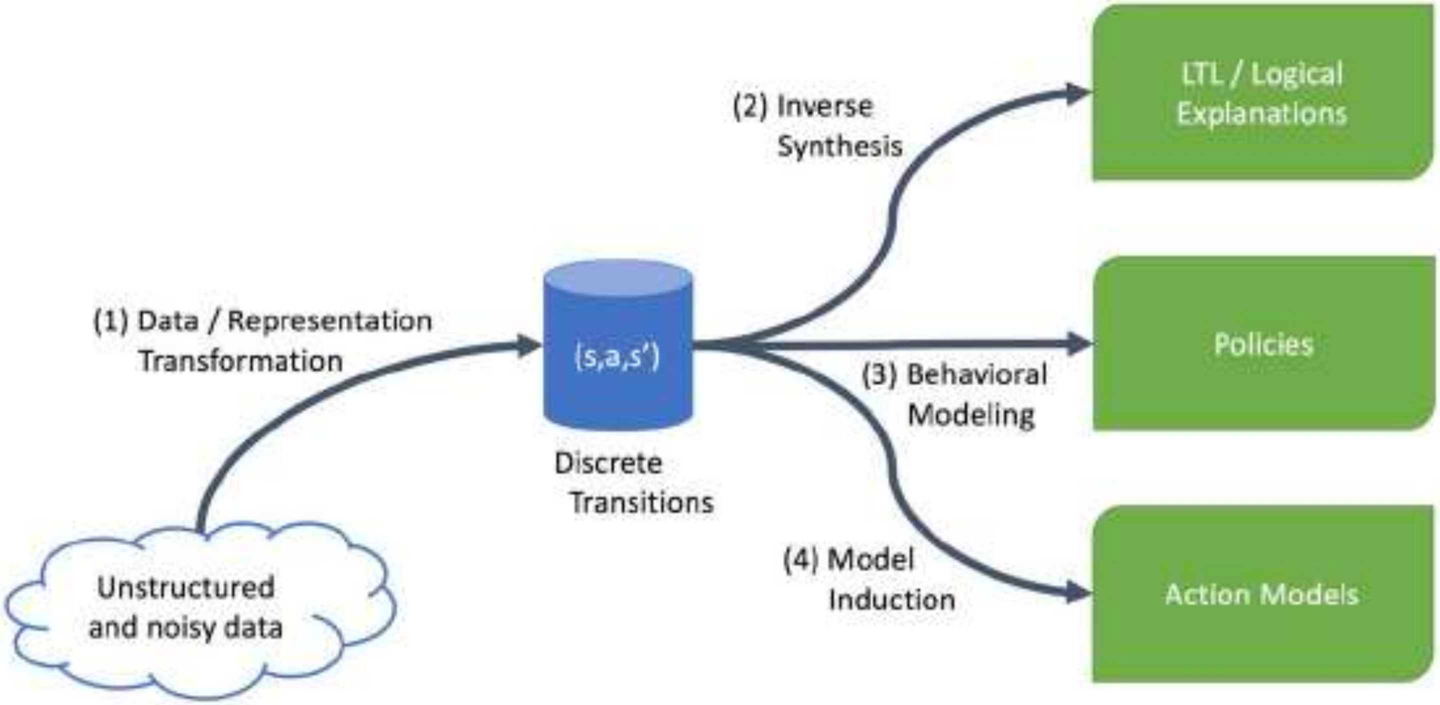

```
(:action pickup

 :parameters (?b1 ?b2 - block)

 :precondition (and (on ?b1 ?b2)
                    (hand-clear))

 :effect (and (not (hand-clear))
              (not (on ?b1 ?b2))
              (holding ?b1))
)
```

Task: Induce the action theory of an environment through observations

LatPlan

Mixing symbolic planning with neural networks

Masataro Asai
IBM

LatPlan

Mixing symbolic planning with neural networks

Masataro Asai
IBM

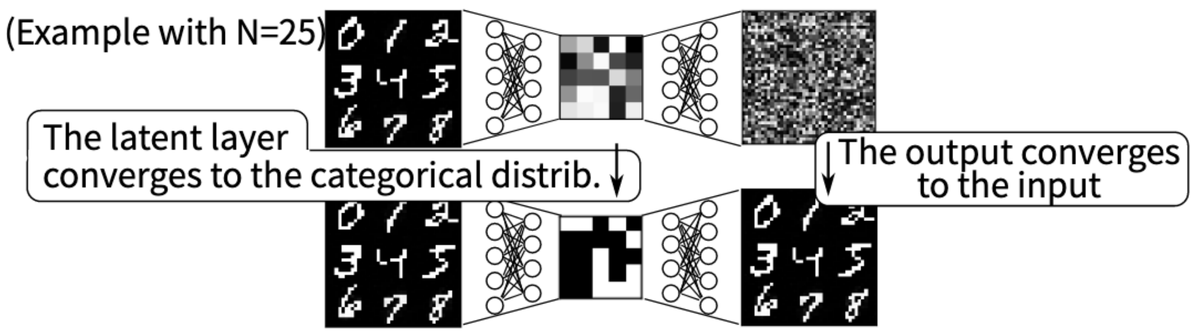

When
$Empty(x, y_{old}) \wedge$
$at(x, y_{new}, p) \wedge$
$up(y_{new}, y_{old})$;
then
$\neg Empty(x, y_{old}) \wedge$
$Empty(x, y_{new}) \wedge$
$\neg at(x, y_{new}, p) \wedge$
$at(x, y_{old}, p)$

```
;; Translates to a PDDL model below:
(:action slide-up ...
 :precondition
 (and (empty ?x ?y-old)
      (at ?x ?y-new ?p) ...)
 :effects
 (and (not (empty ?x ?y-old))
      (empty ?x ?y-new)
      (not (at ?x ?y-new ?p))
      (at ?x ?y-old ?p)))
```

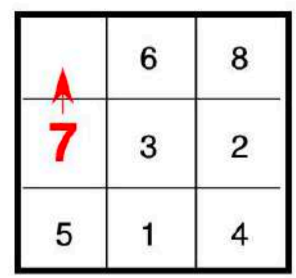

Verifiably Safe Reinforcement Learning

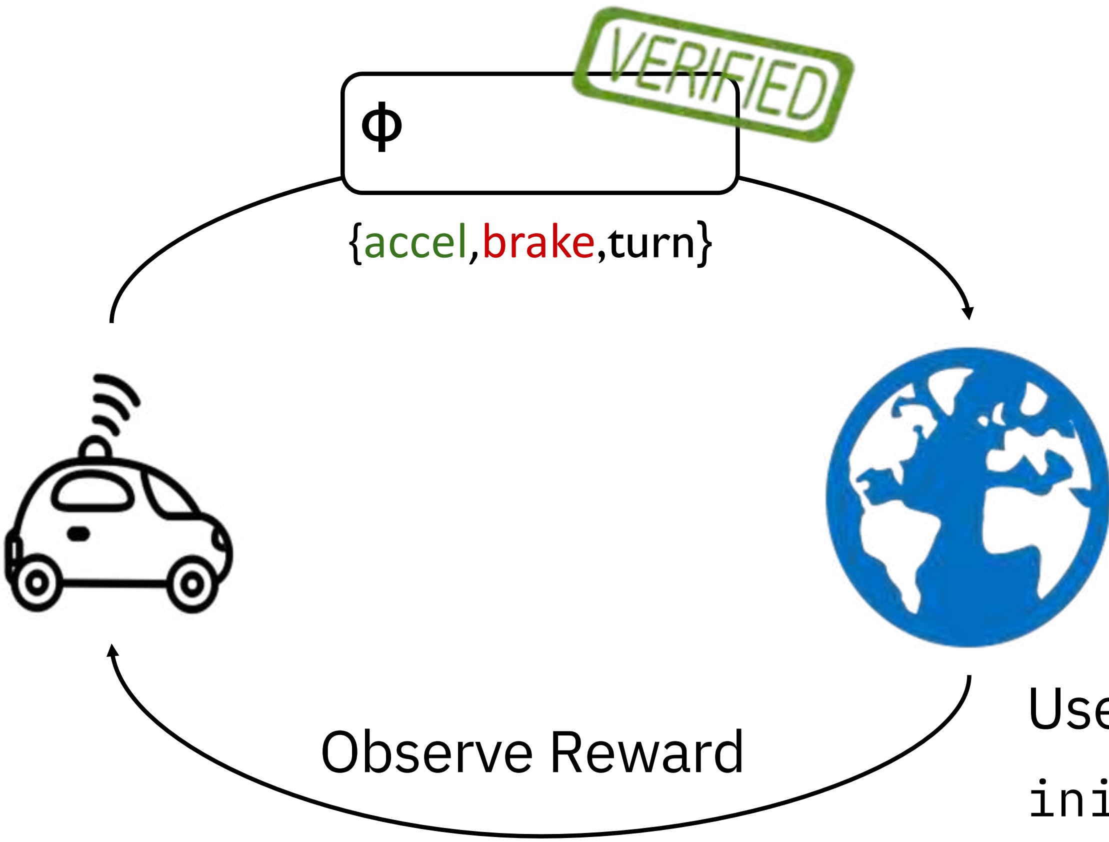

Nathan Fulton
IBM

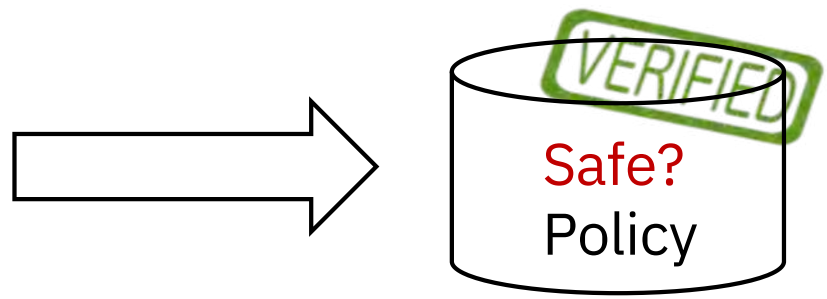

Use a theorem prover to prove:

$$\mathtt{init} \rightarrow [\{\{\mathtt{accel} \cup \mathtt{brake}\};\mathtt{ODEs}\}^*]\mathtt{safe}$$

is correctly monitored by **ϕ**.

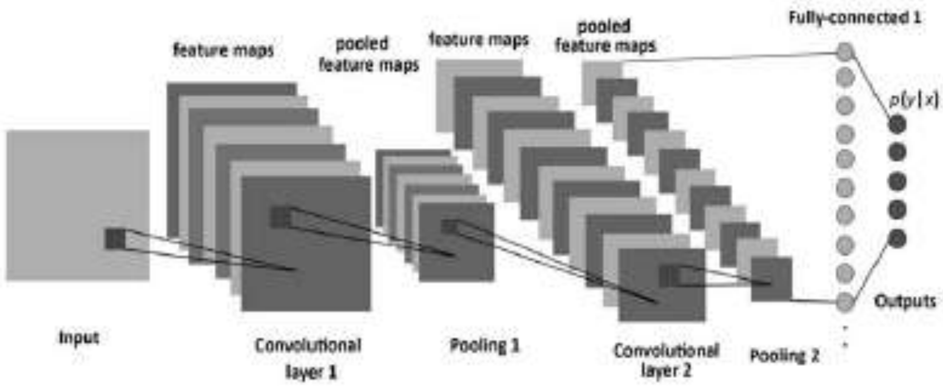

+

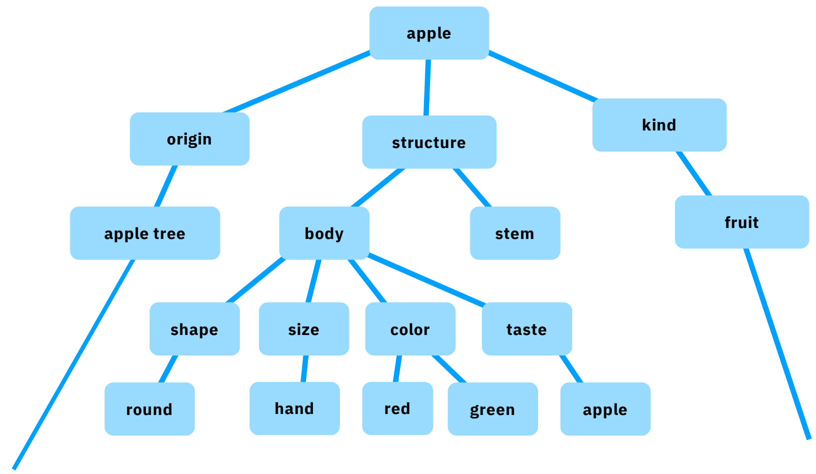


```
(:action pickup
 :parameters (?b1 ?b2 - block)
 :precondition (and (on ?b1 ?b2)
                    (hand-clear))
 :effect (and (not (hand-clear))
              (not (on ?b1 ?b2))
              (holding ?b1))
)
```

NEURAL NETWORKS

SYMBOLIC AI

MIT-IBM Watson AI Lab

Causal Inference

Beyond Correlation—inferring and testing for causal relationships in complex systems

Caroline Uhler
MIT

Guy Bresler
MIT

Karthikeyan Shanmugam
IBM

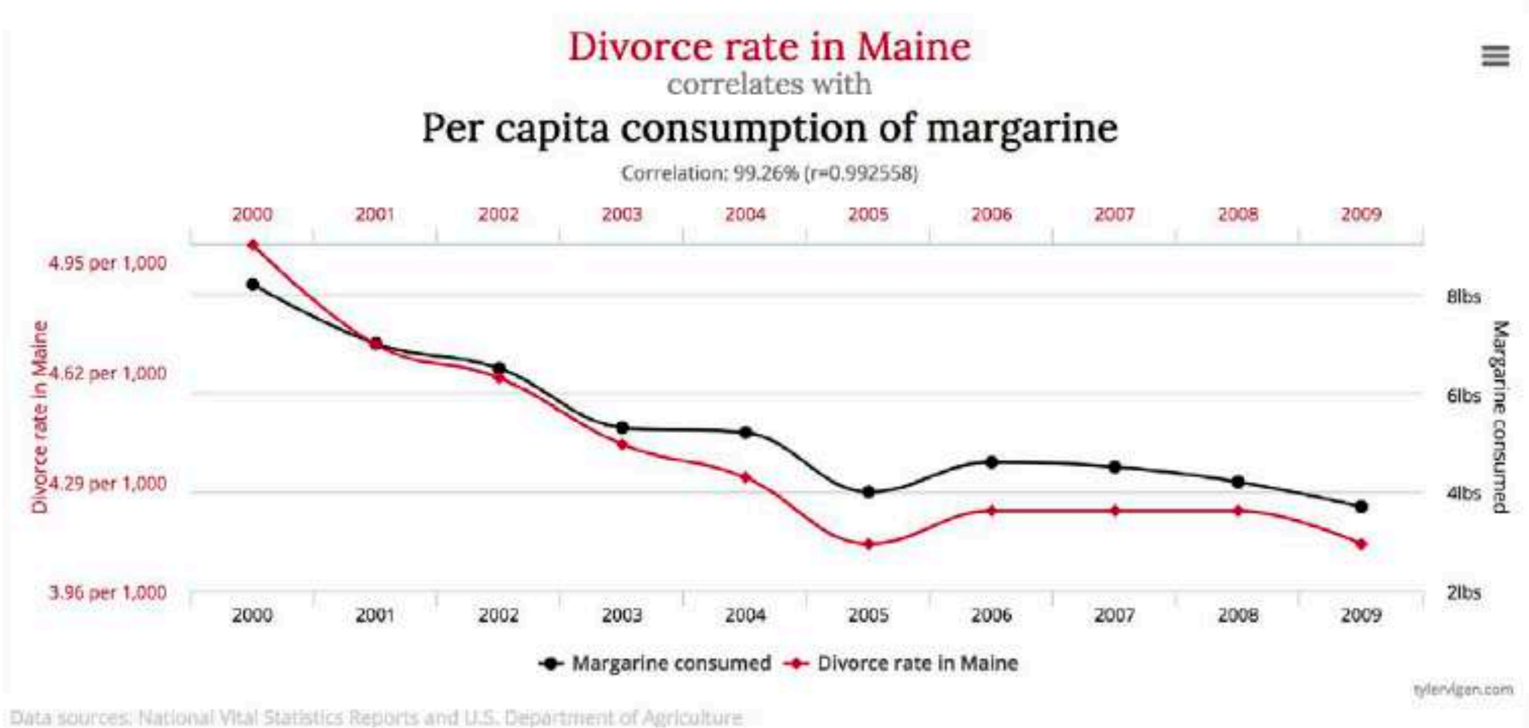

http://tylervigen.com/spurious-correlations

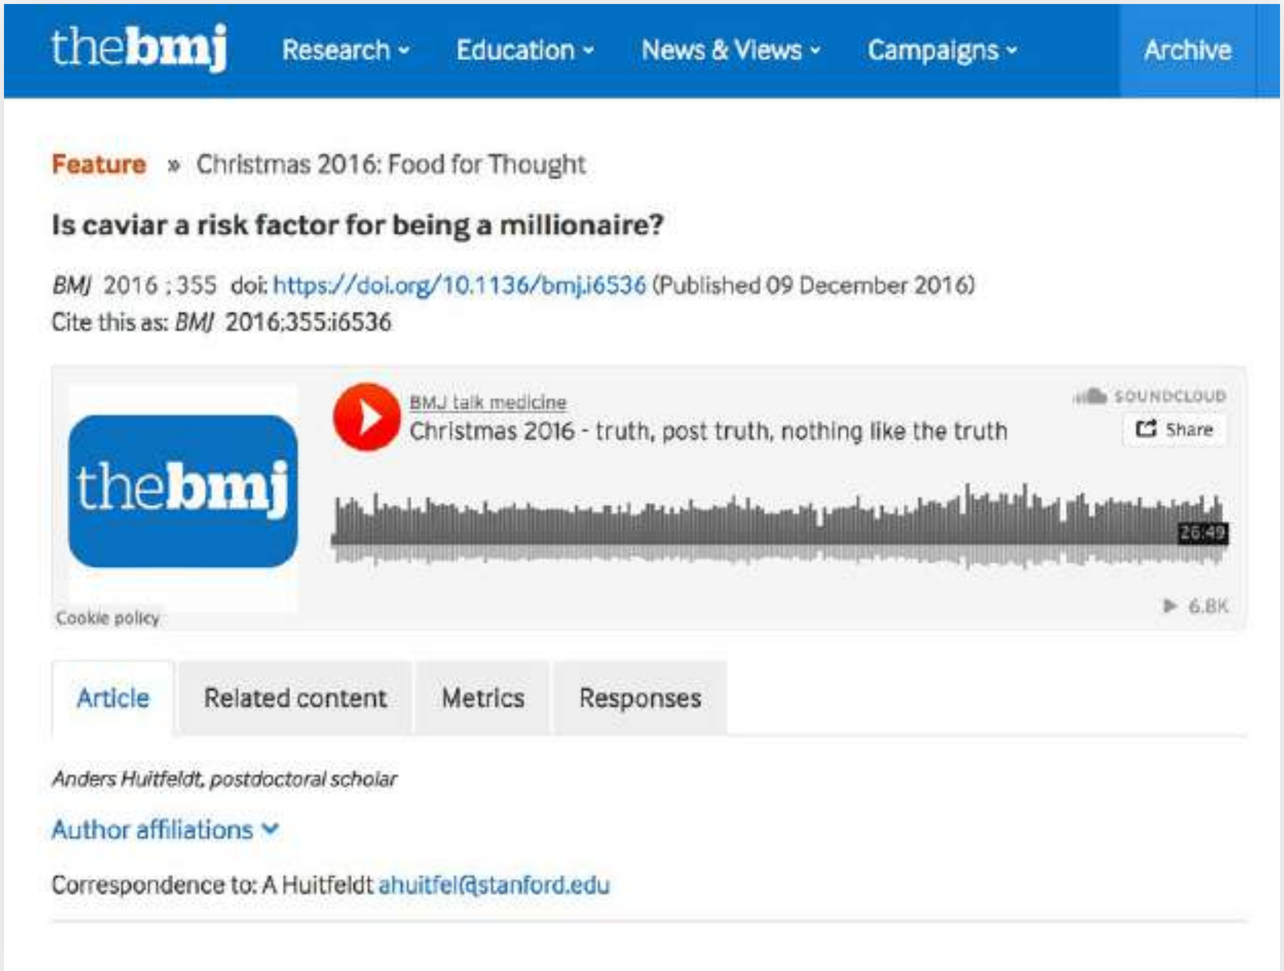

thebmj Research Education News & Views Campaigns Archive

Feature » Christmas 2016: Food for Thought

Is caviar a risk factor for being a millionaire?

BMJ 2016 ; 355 doi: https://doi.org/10.1136/bmj.i6536 (Published 09 December 2016)
Cite this as: *BMJ* 2016;355:i6536

BMJ talk medicine
Christmas 2016 - truth, post truth, nothing like the truth
SOUNDCLOUD
Share
26:49
6.8K
Cookie policy

Article | Related content | Metrics | Responses

Anders Huitfeldt, postdoctoral scholar

Author affiliations

Correspondence to: A Huitfeldt ahuitfel@stanford.edu