The Objects of Our Curiosity

Intrinsic Motivation, Intuitive Physics and Self-Supervised Learning

---

NeurIPS Workshop: Modeling the Physical World

2018.12.07

Daniel Yamins

Stanford Neurosciences Institute
Stanford Artificial Intelligence Laboratory
Departments of Psychology and Computer Science
Stanford University

Our work is founded on two mutually reinforcing goals:

Understanding Brains & Cognition

Stanford Neuro AI Lab

Better AI/ML techniques

Computational Models of the Visual System

The primate visual system as a hierarchical, convolutional neural network:

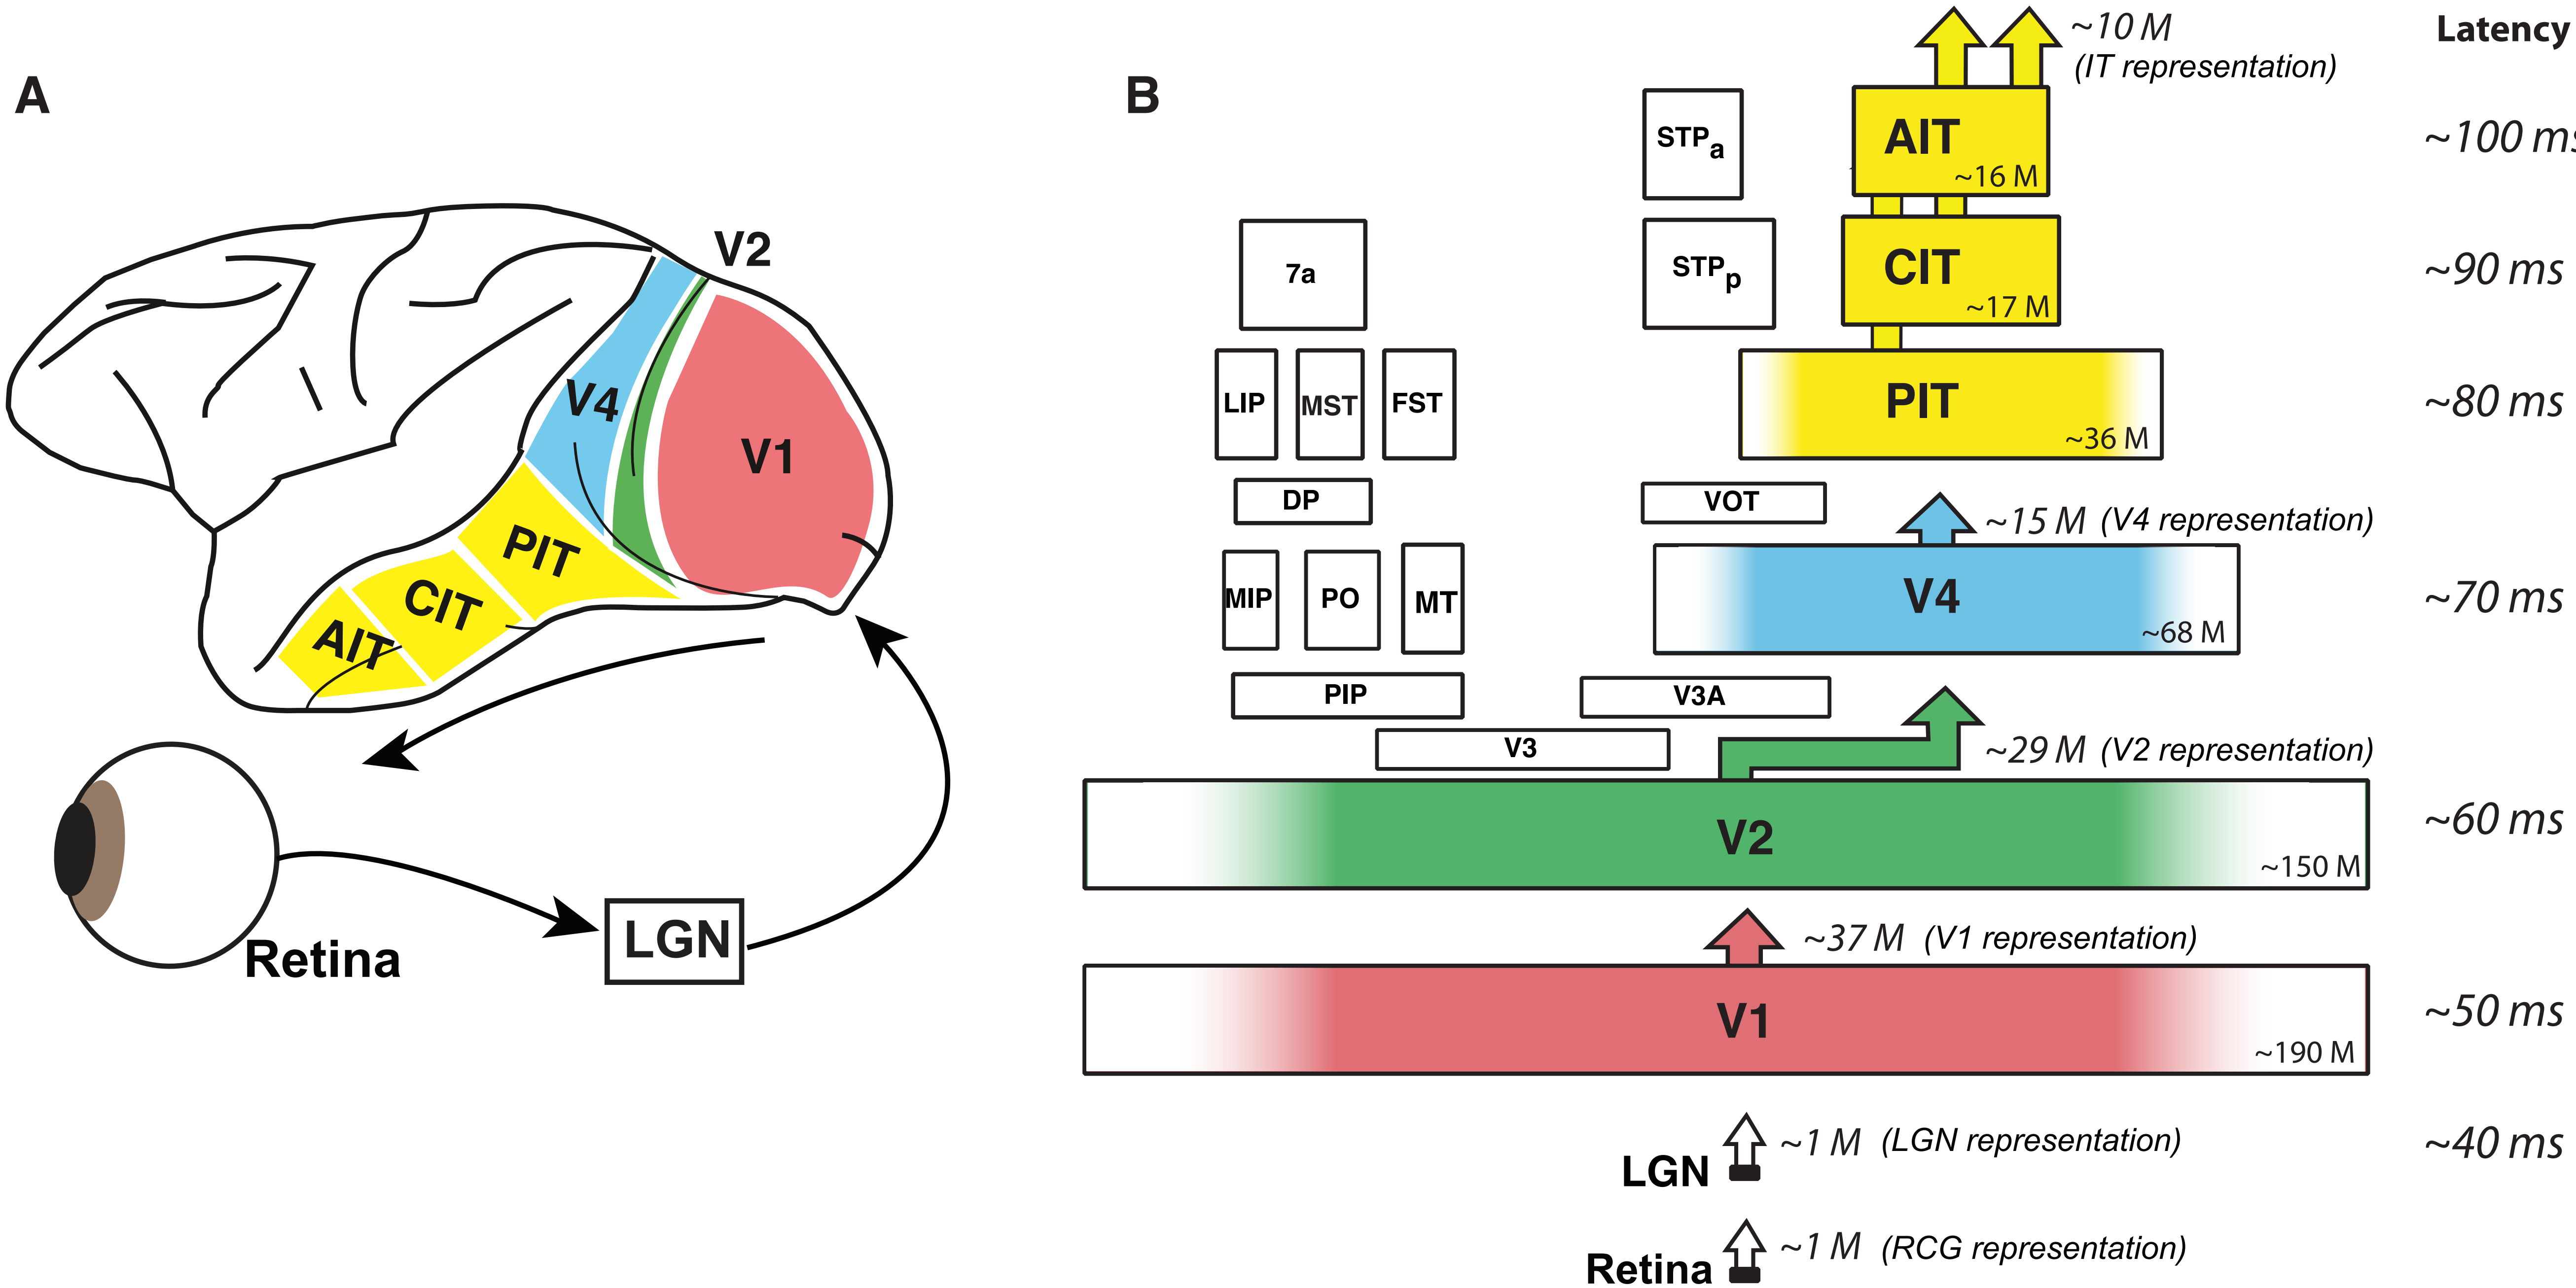

Adapted from DiCarlo et al. 2012

Computational Models of the Visual System

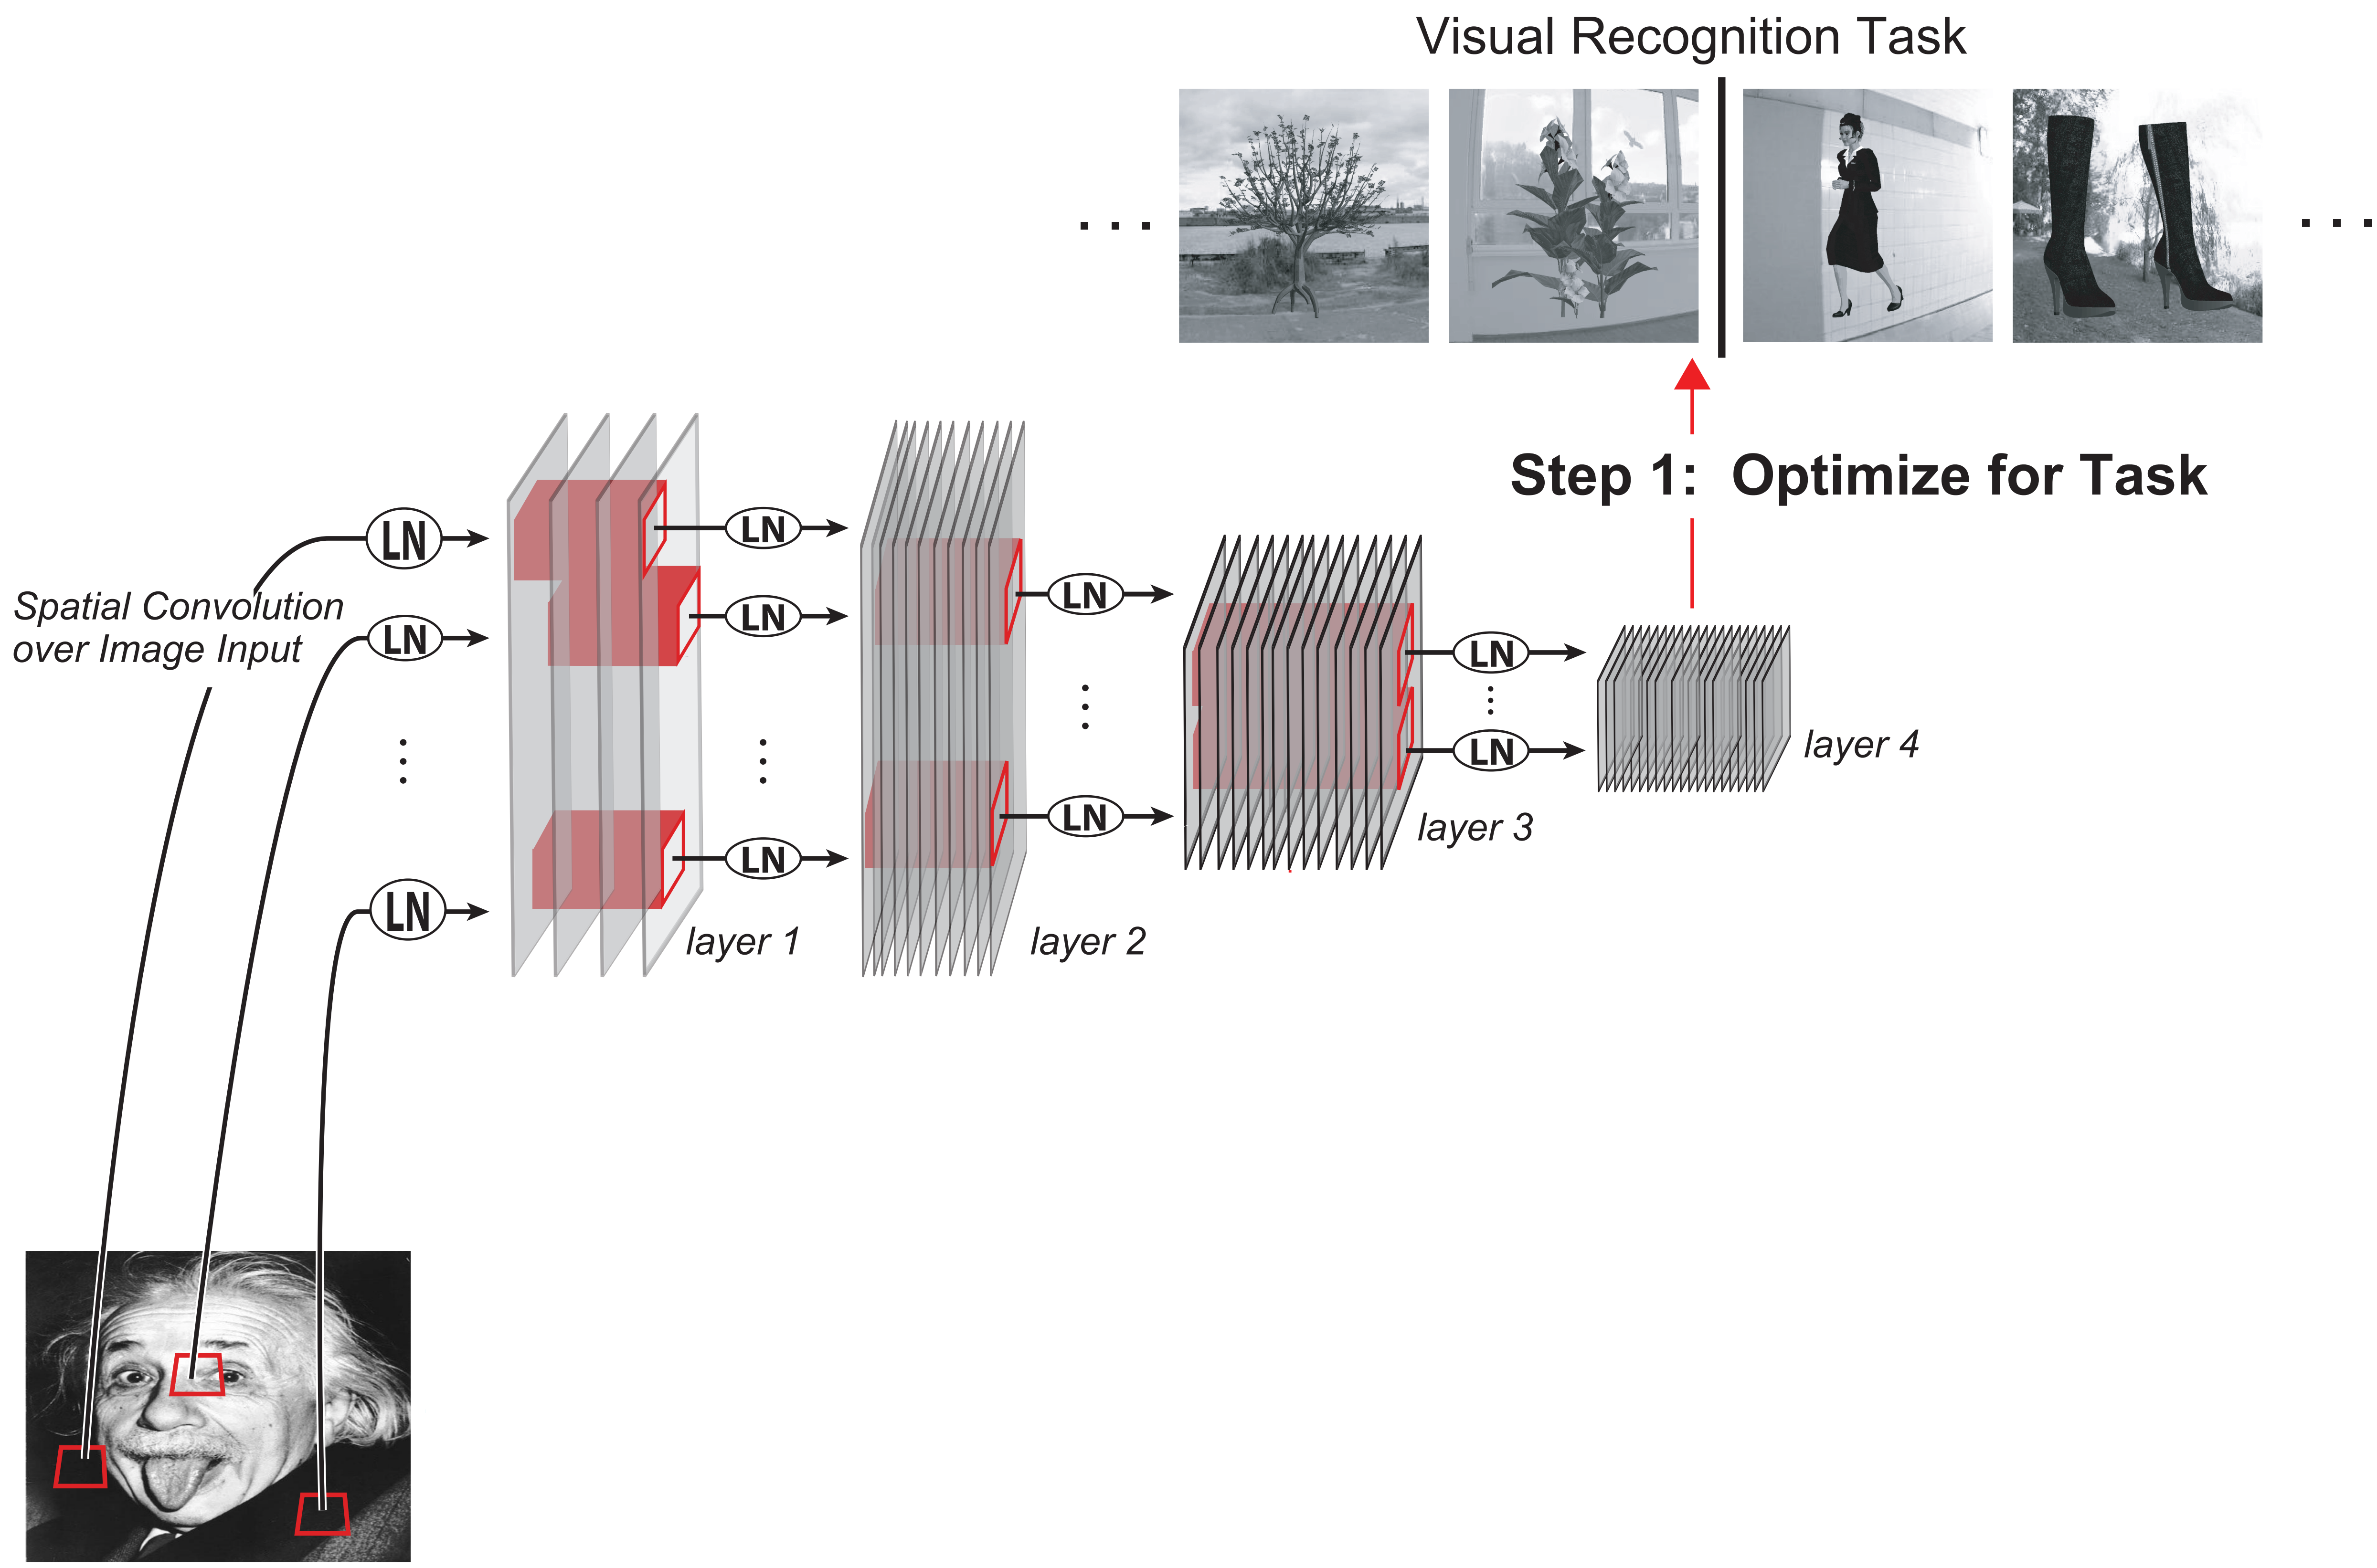

Computational Models of the Visual System

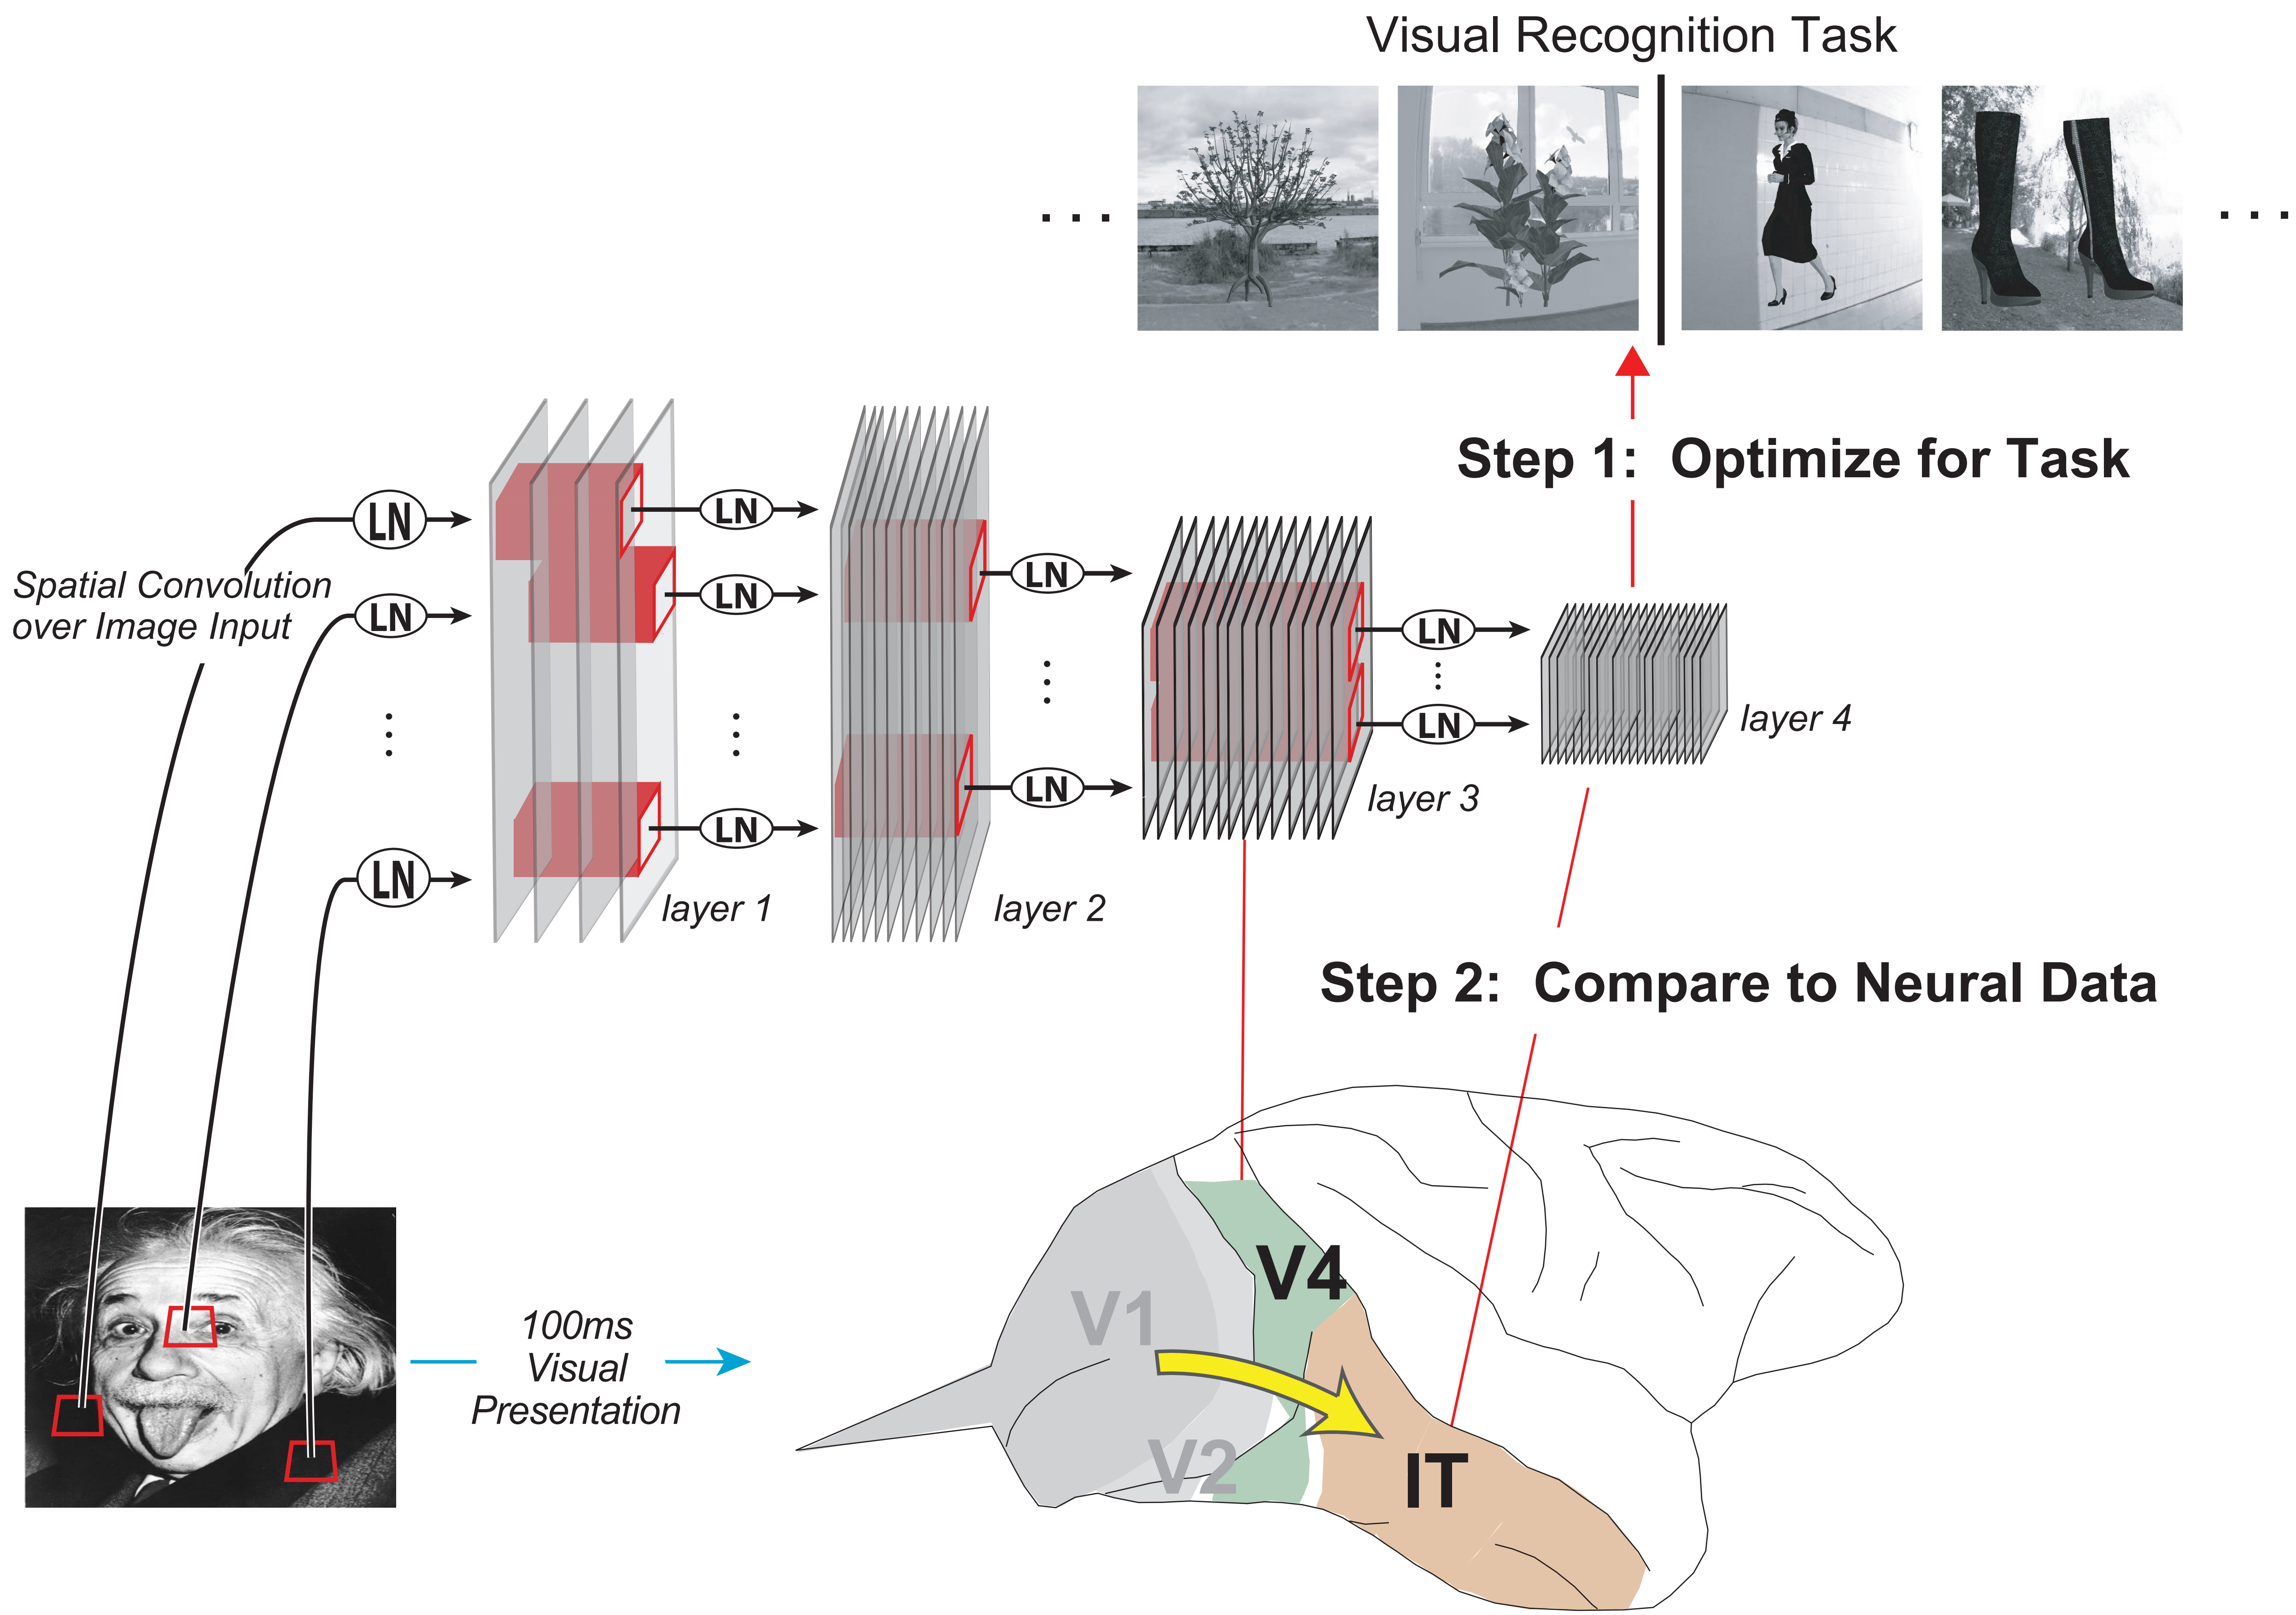

Computational Models of the Visual System

To our knowledge best (in terms of neural prediction) feedforward model is a ~12-layer CNN

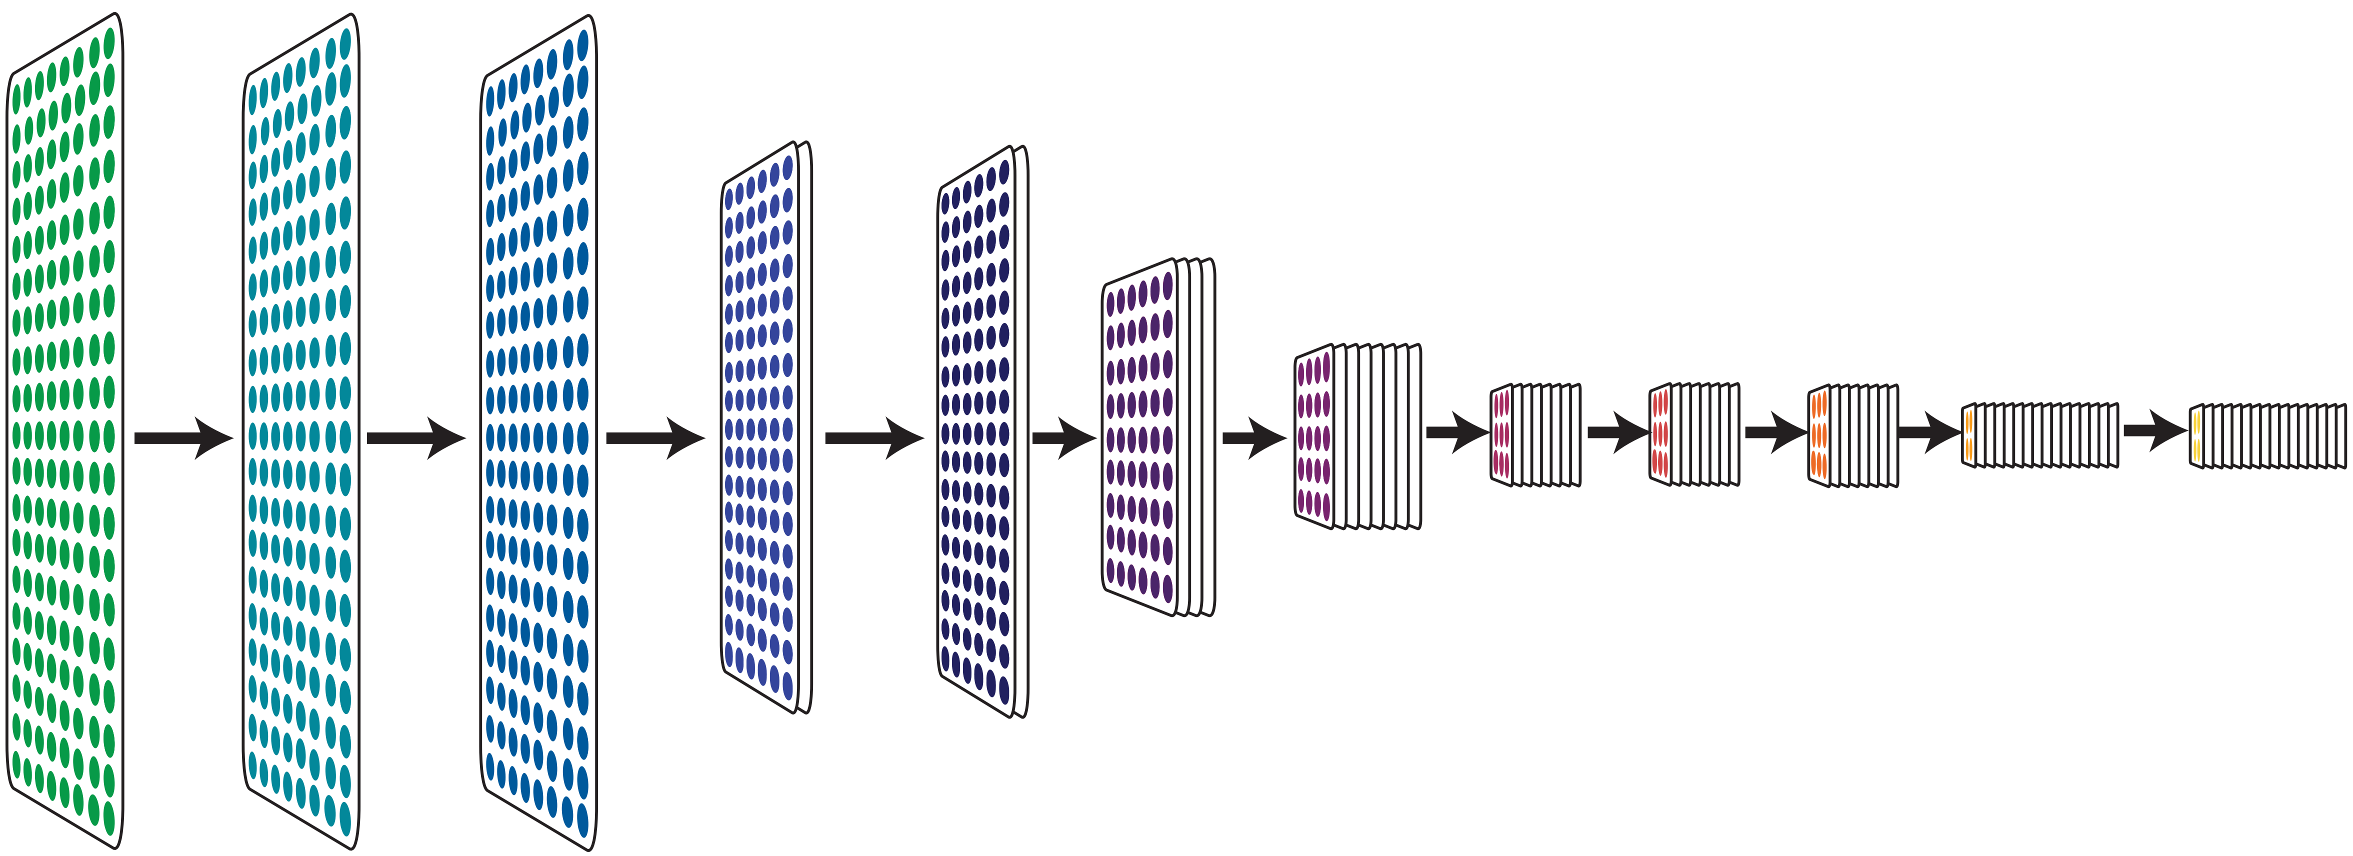

Computational Models of the Visual System

To our knowledge best (in terms of neural prediction) feedforward model is a ~12-layer CNN

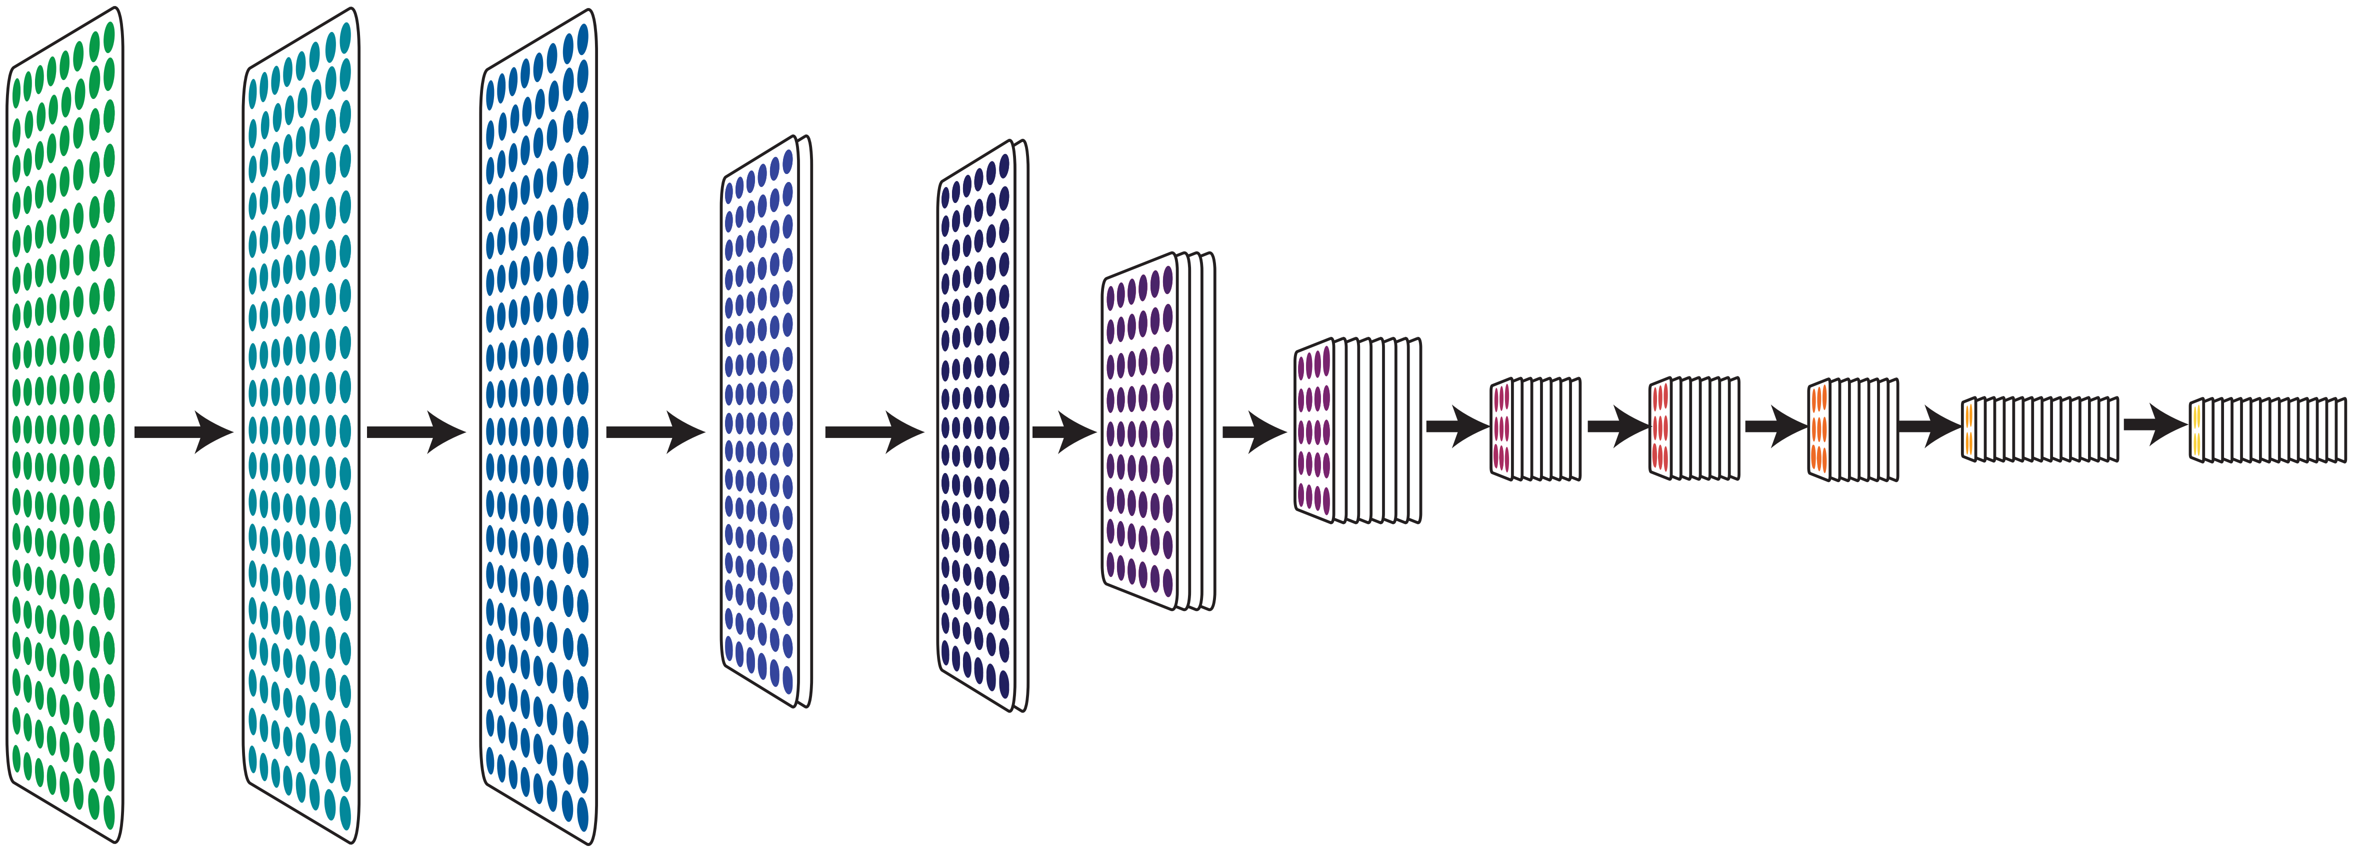

...trained on ImageNet Categorization.

The Problem

There's just no way that these creatures receive millions of high-level semantic labels during learning.

ImageNet is a pretty effective proxy, but just obviously deeply wrong.

The Problem

There's just no way that these creatures receive millions of high-level semantic labels during learning.

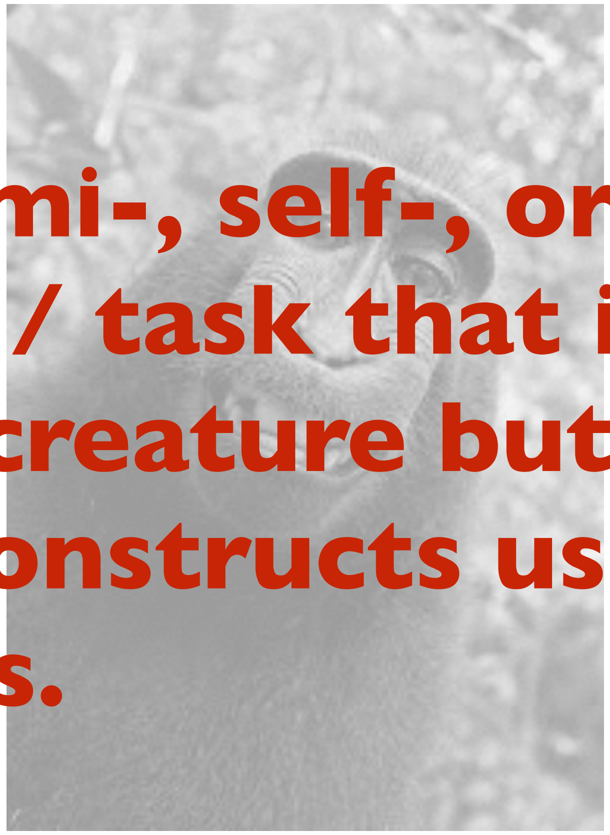

Must find some sort of semi-, self-, or unsupervised loss function / task that is "realistically costly" to the creature but is sufficiently powerful that it constructs useful representations.

ImageNet is a pretty effective proxy, but just obviously deeply wrong.

Self-supervised learning

Olshausen & Field (1996)

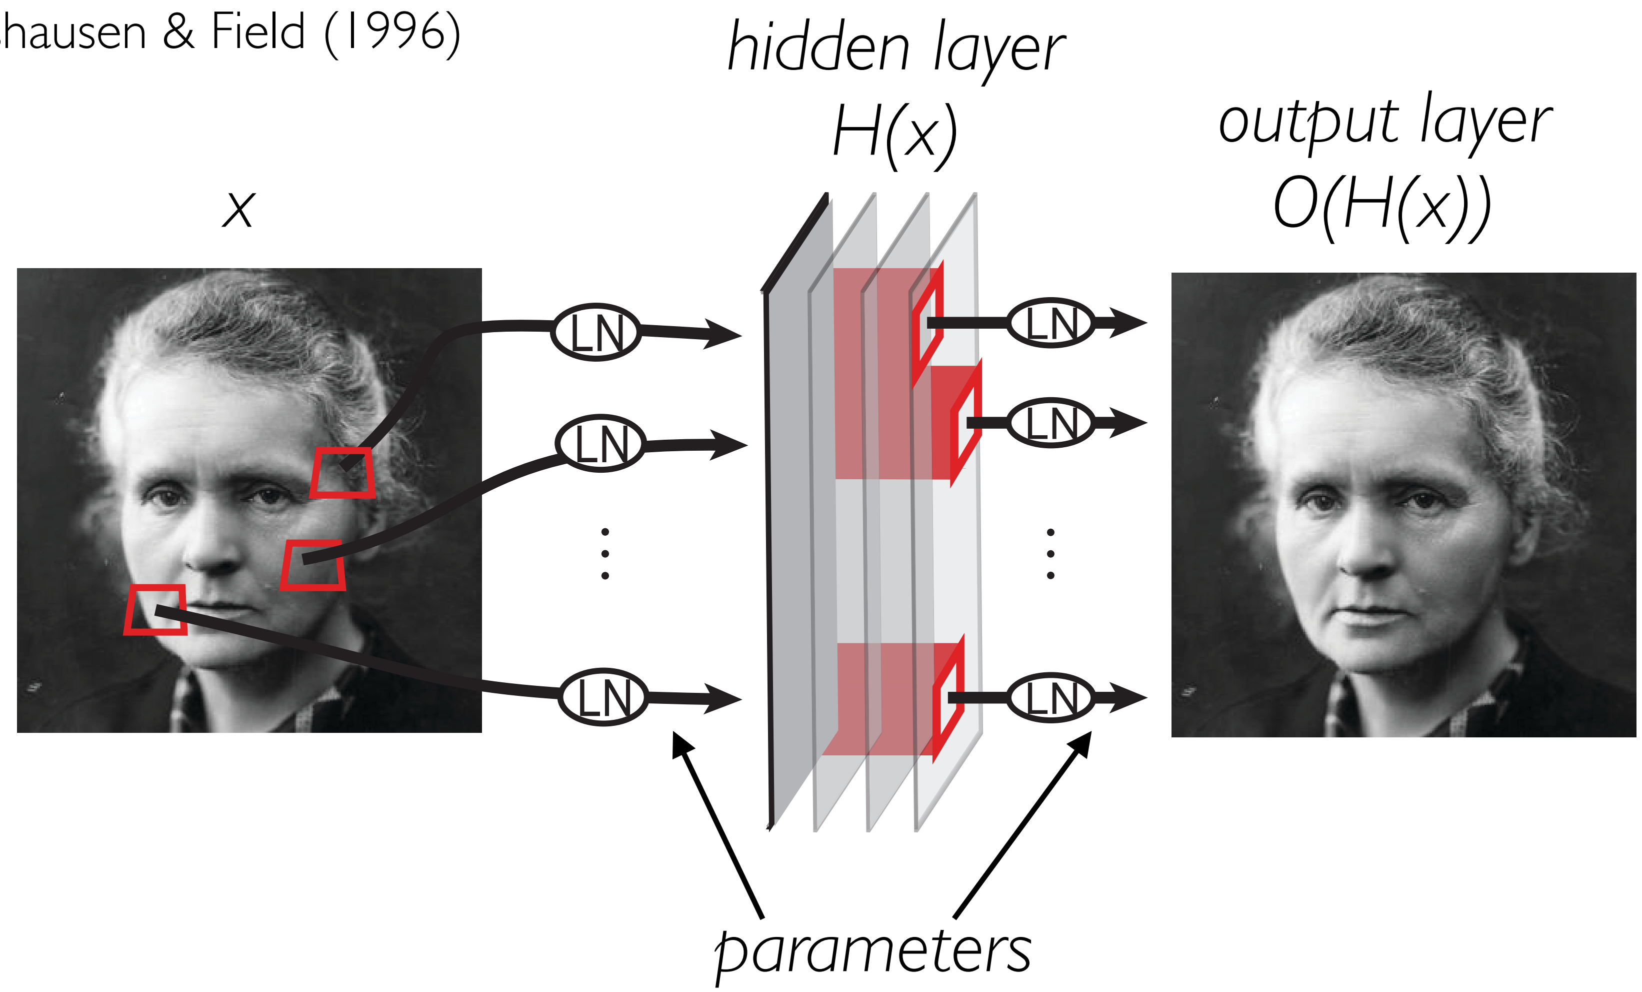

Self-supervised learning

Olshausen & Field (1996)

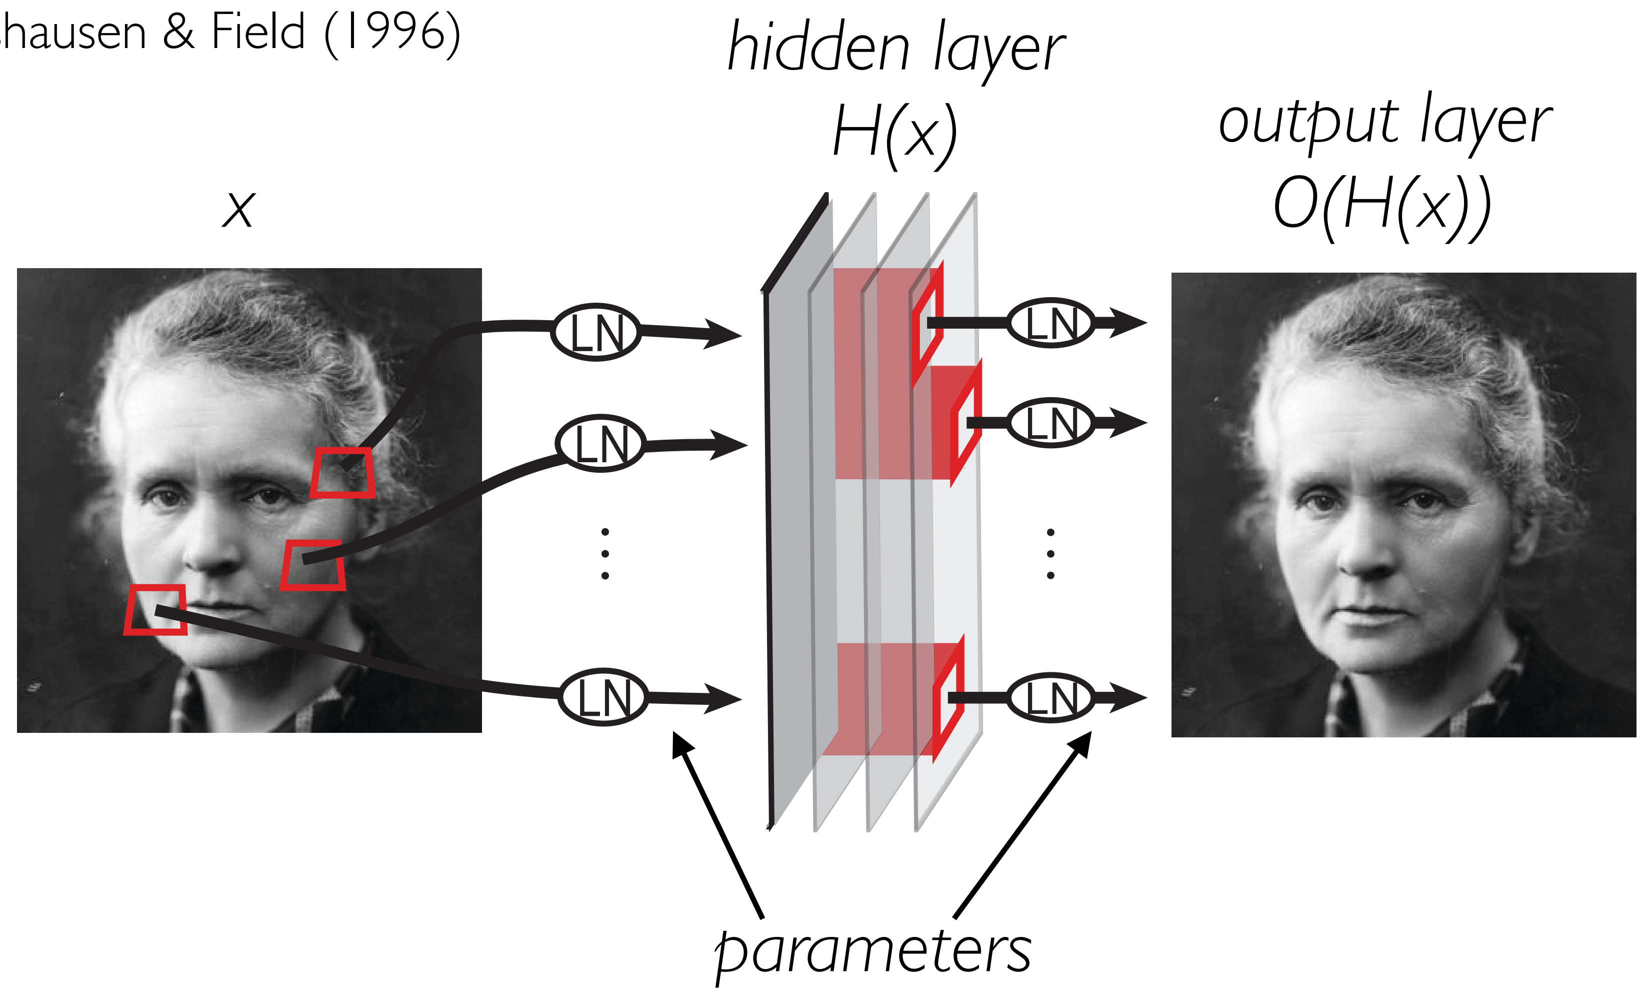

$$L(x) = |x - O(H(x))|^2 + \lambda \cdot |H(x)|$$

reconstruction loss

complexity penalty

Self-supervised learning

Olshausen & Field (1996)

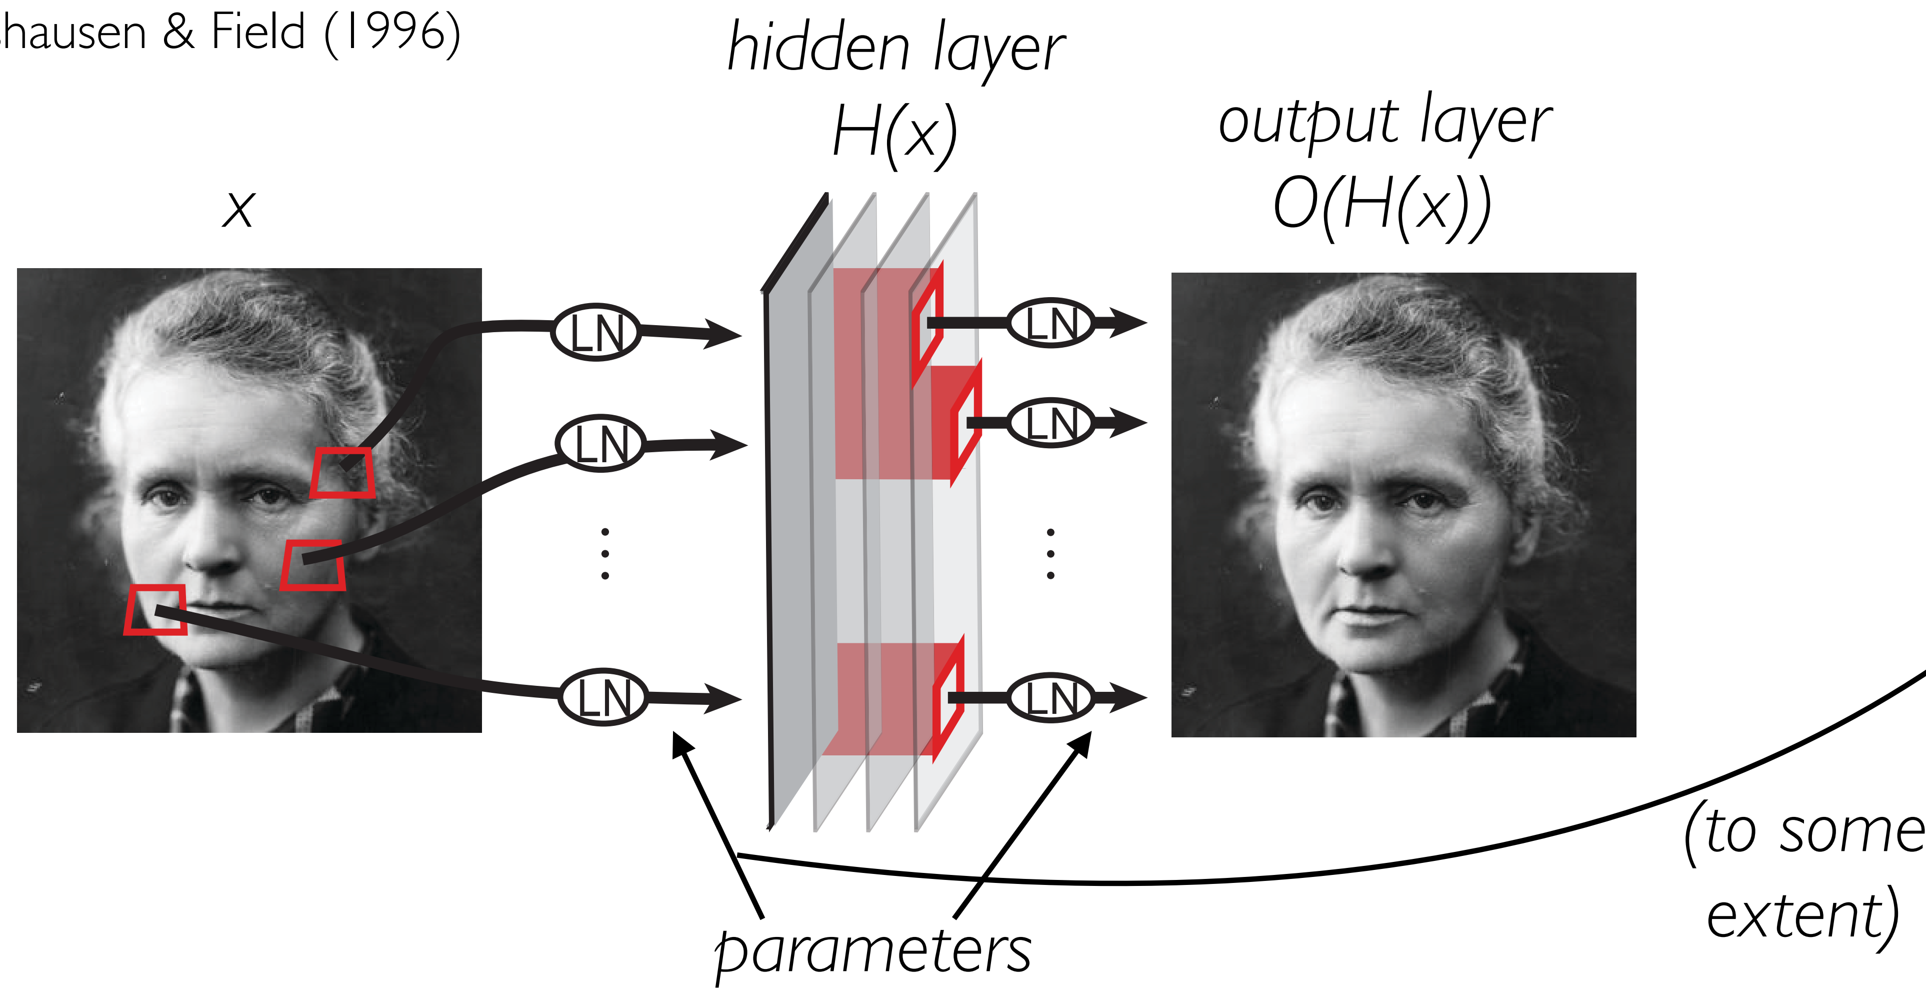

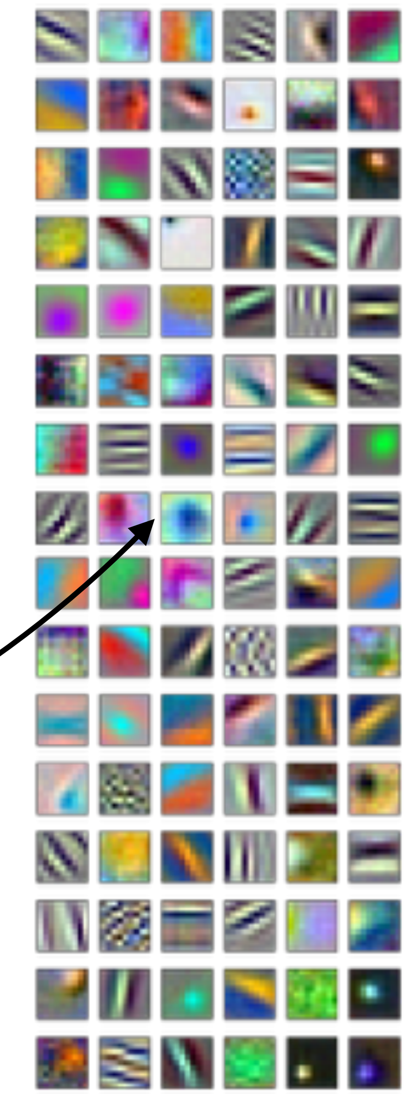

$$L(x) = |x - O(H(x))|^2 + \lambda \cdot |H(x)|$$

reconstruction loss

complexity penalty

Self-supervised learning

Dynamics might give richer signal

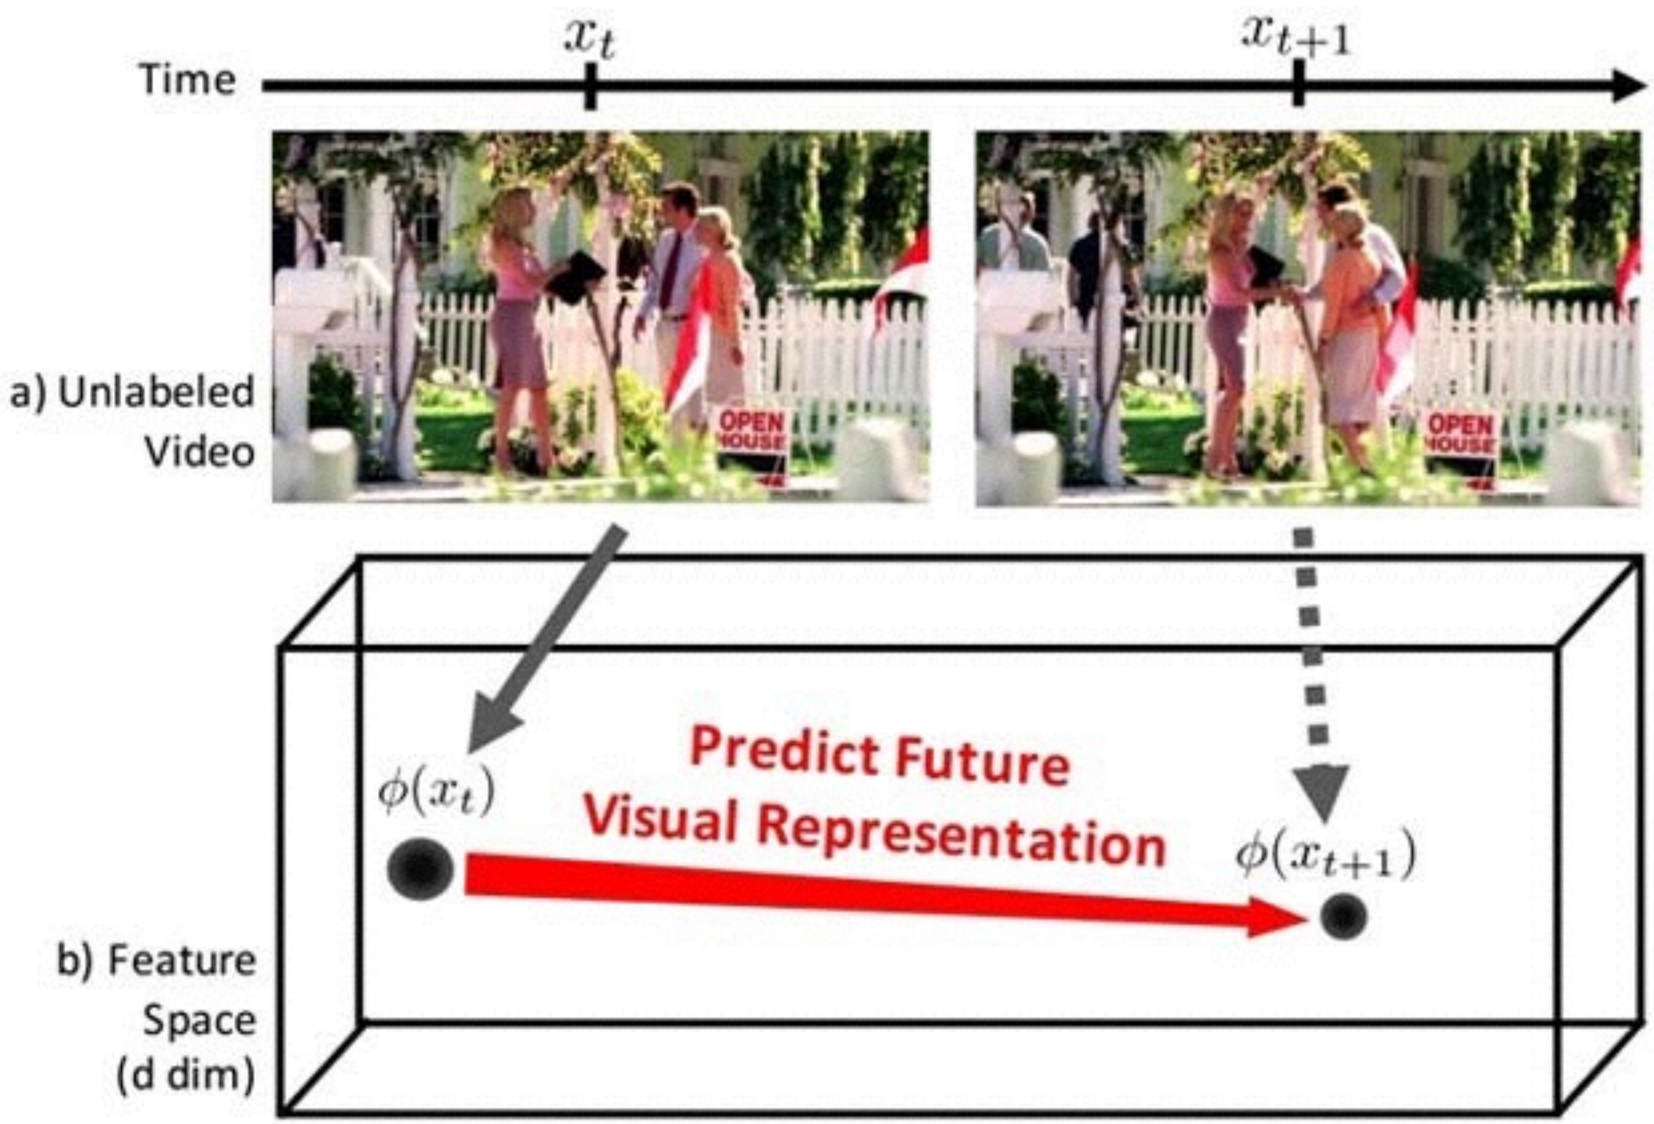

Self-supervised learning

Dynamics might give richer signal

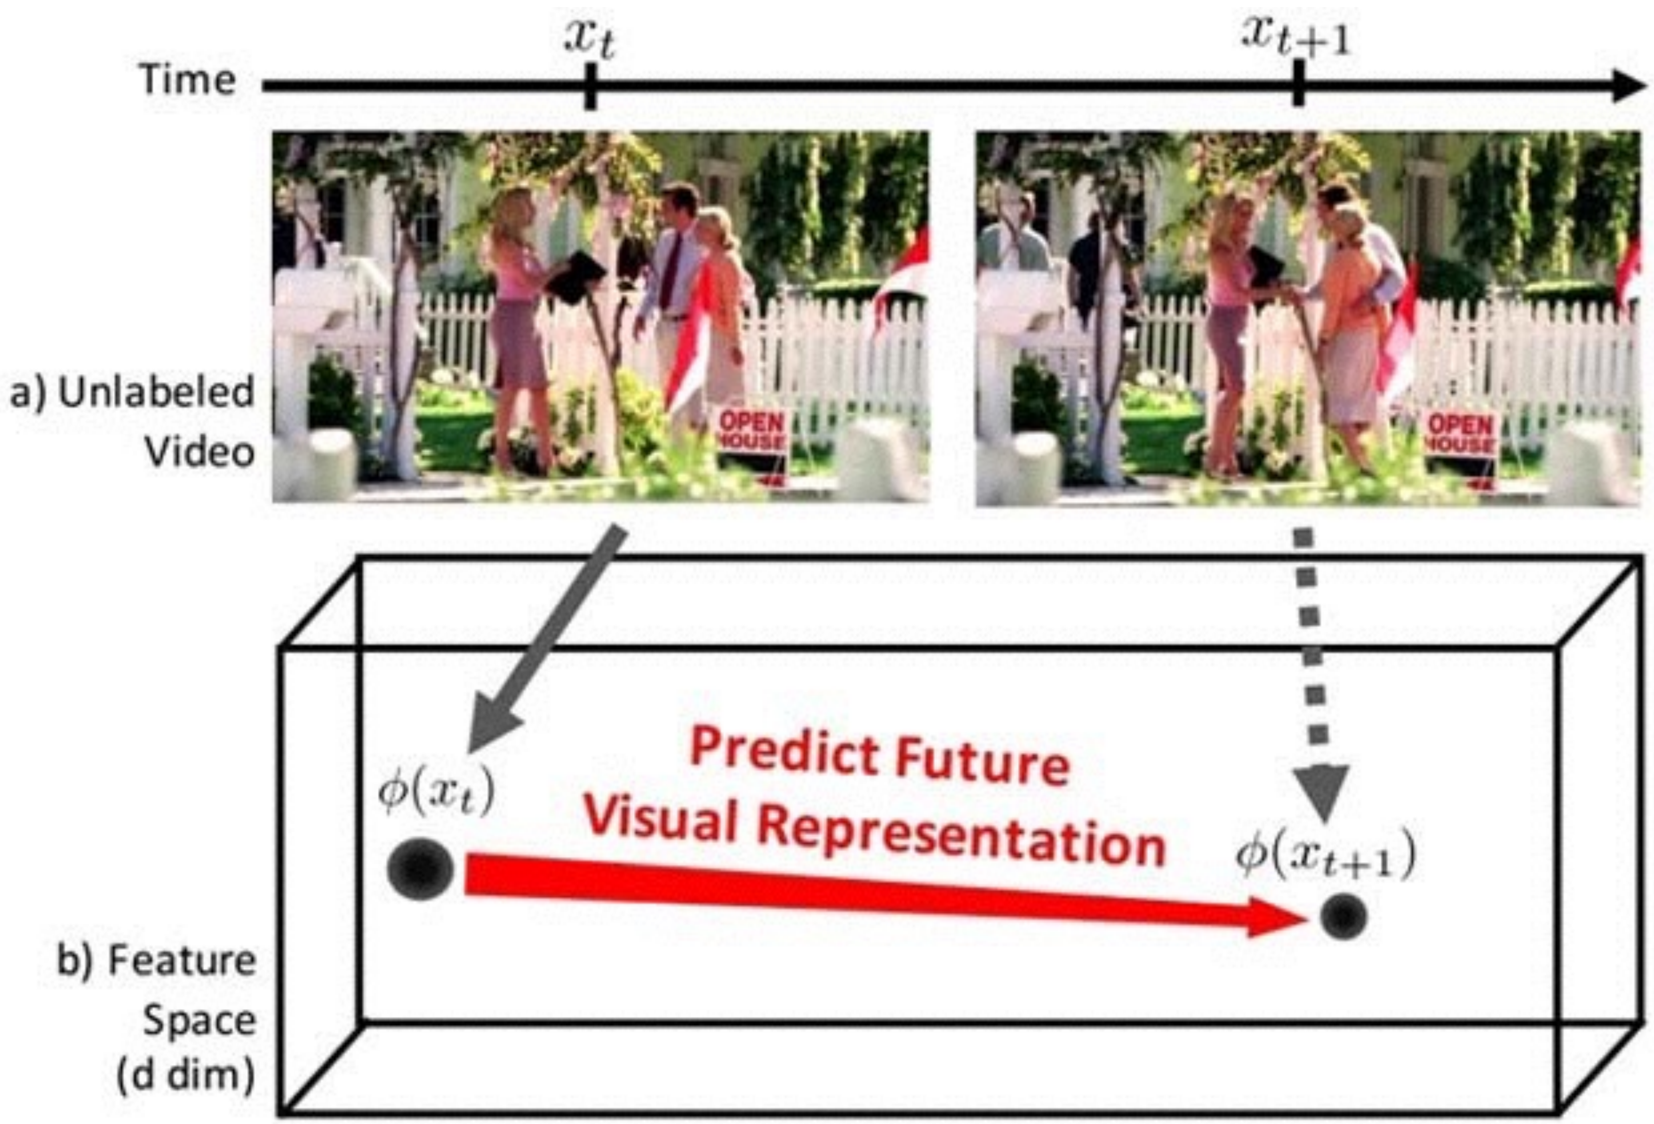

$$L(x) = |x_{t+1} - \mathrm{Decode}(\mathrm{Encode}(x_t))|^2 + \lambda \cdot \mathrm{Penalty}(\mathrm{Encode}(x_t))$$

Self-supervised learning

Dynamics might give richer signal *. . . but most passive video sequences are **quite boring***

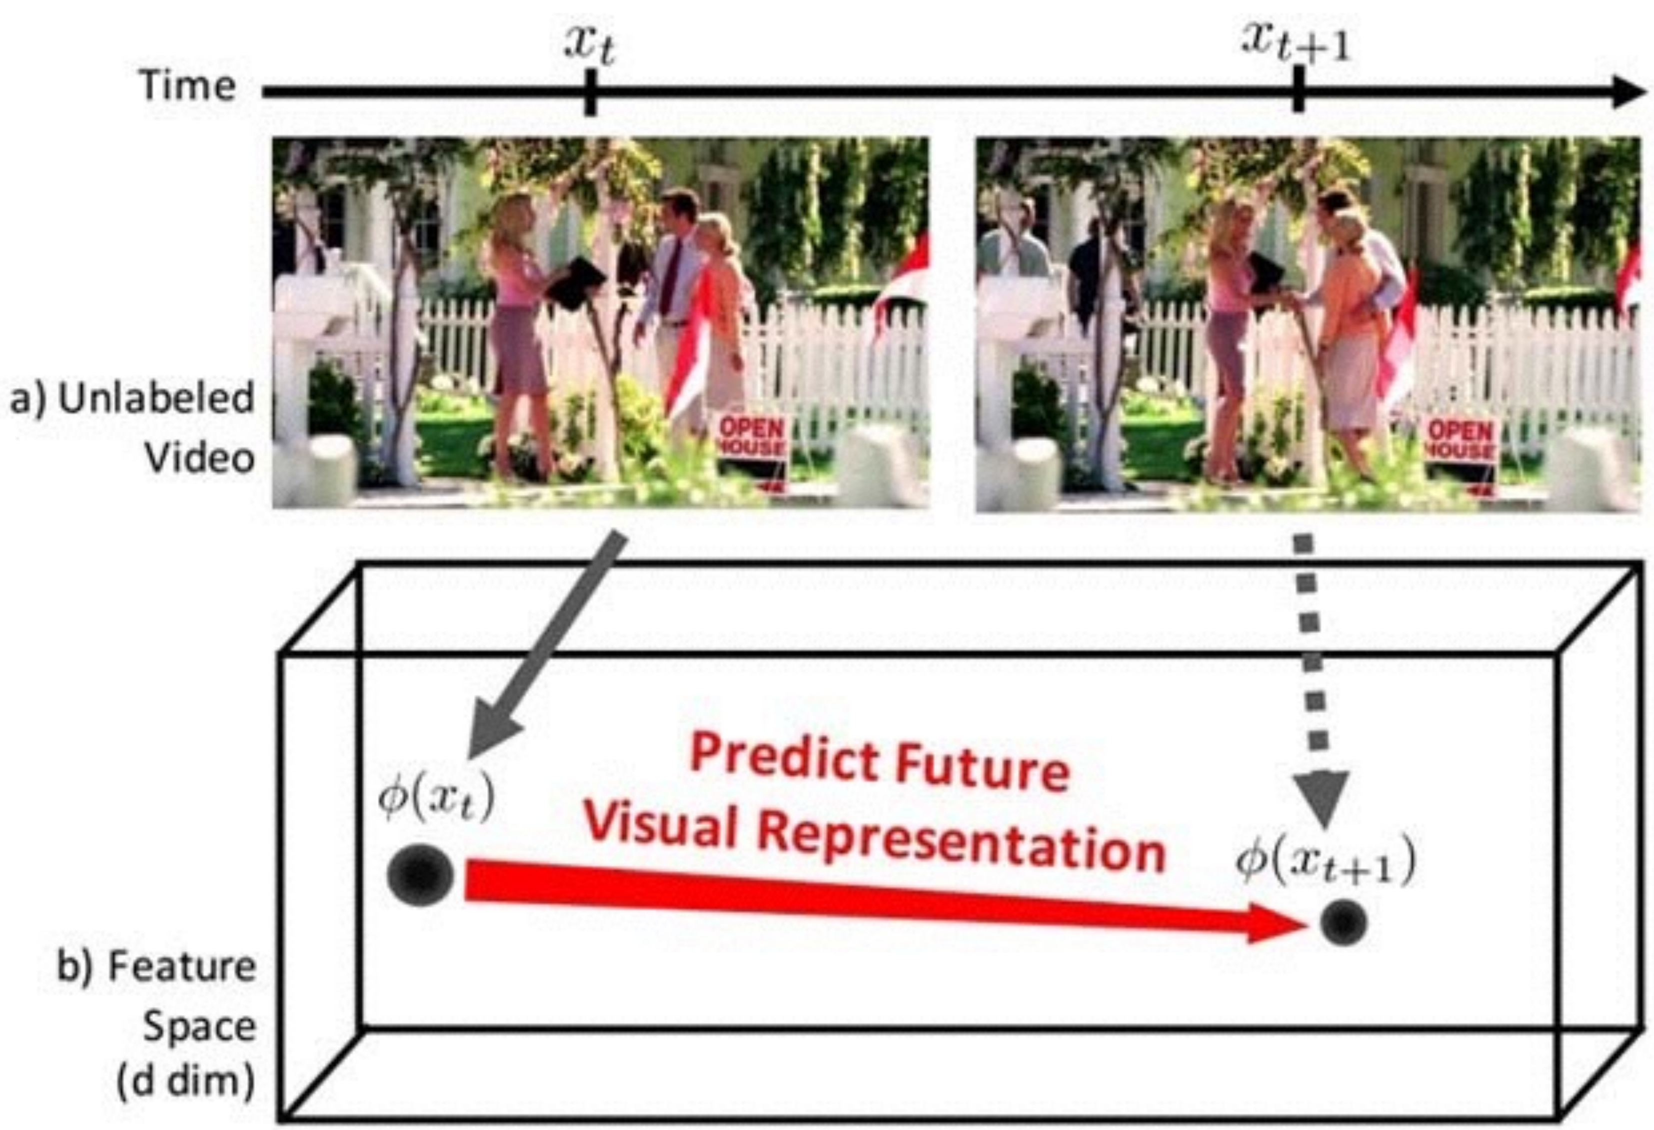

fairly trivial features

$$L(x) = |x_{t+1} - \mathrm{Decode}(\mathrm{Encode}(x_t))|^2 + \lambda \cdot \mathrm{Penalty}(\mathrm{Encode}(x_t))$$

Children learn through **play.**

How does this work?

Self-supervised learning

Give agent some kind of volition to take actions

$$L(x) = |x_{t+1}^{\text{action}} - \text{Decode}(\text{Encode}(x_t))|^2 + \lambda \cdot \text{Penalty}(\text{Encode}(x_t))$$

Self-supervised learning

Give agent some kind of volition to take actions . . . but now the agent will be ***lazy***

$$L(x) = |x_{t+1}^{\text{action}} - \text{Decode}(\text{Encode}(x_t))|^2 + \lambda \cdot \text{Penalty}(\text{Encode}(x_t))$$

Self-supervised learning

*Give agent some kind of volition to take actions . . . but now the agent will be **lazy***

$$
\begin{aligned}
L(x) = |x^{\text{action}}{}_{t+1} - \text{Decode}(\text{Encode}(x_t))|^2 + \lambda \cdot \text{Penalty}(\text{Encode}(x_t)) \\
+ \text{Intrinsic Motivation}
\end{aligned}
$$

Environment

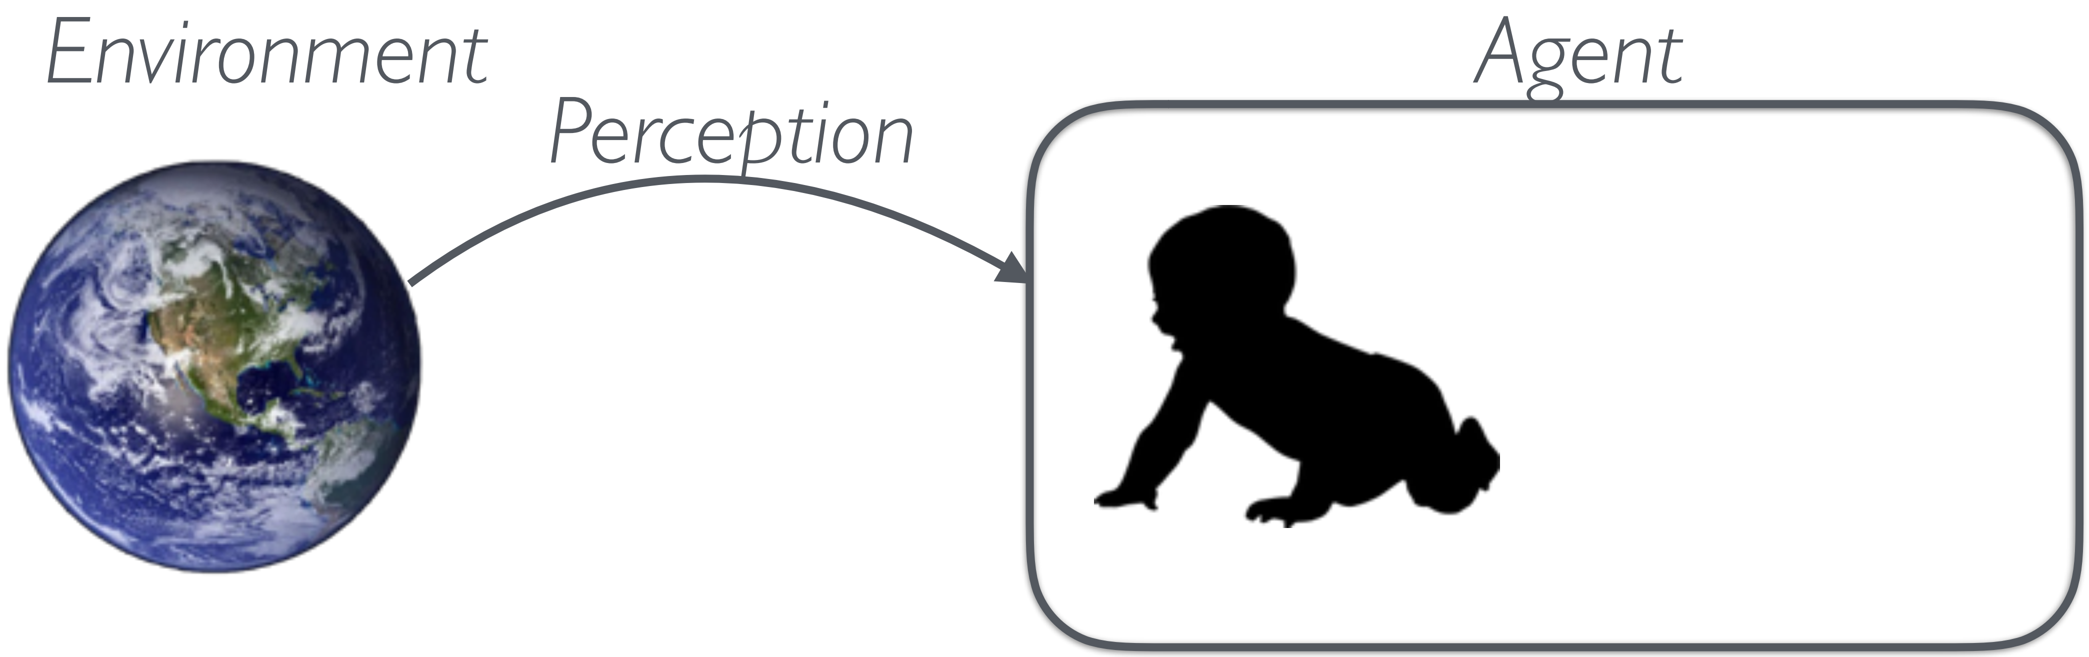
Environment
Perception
Agent

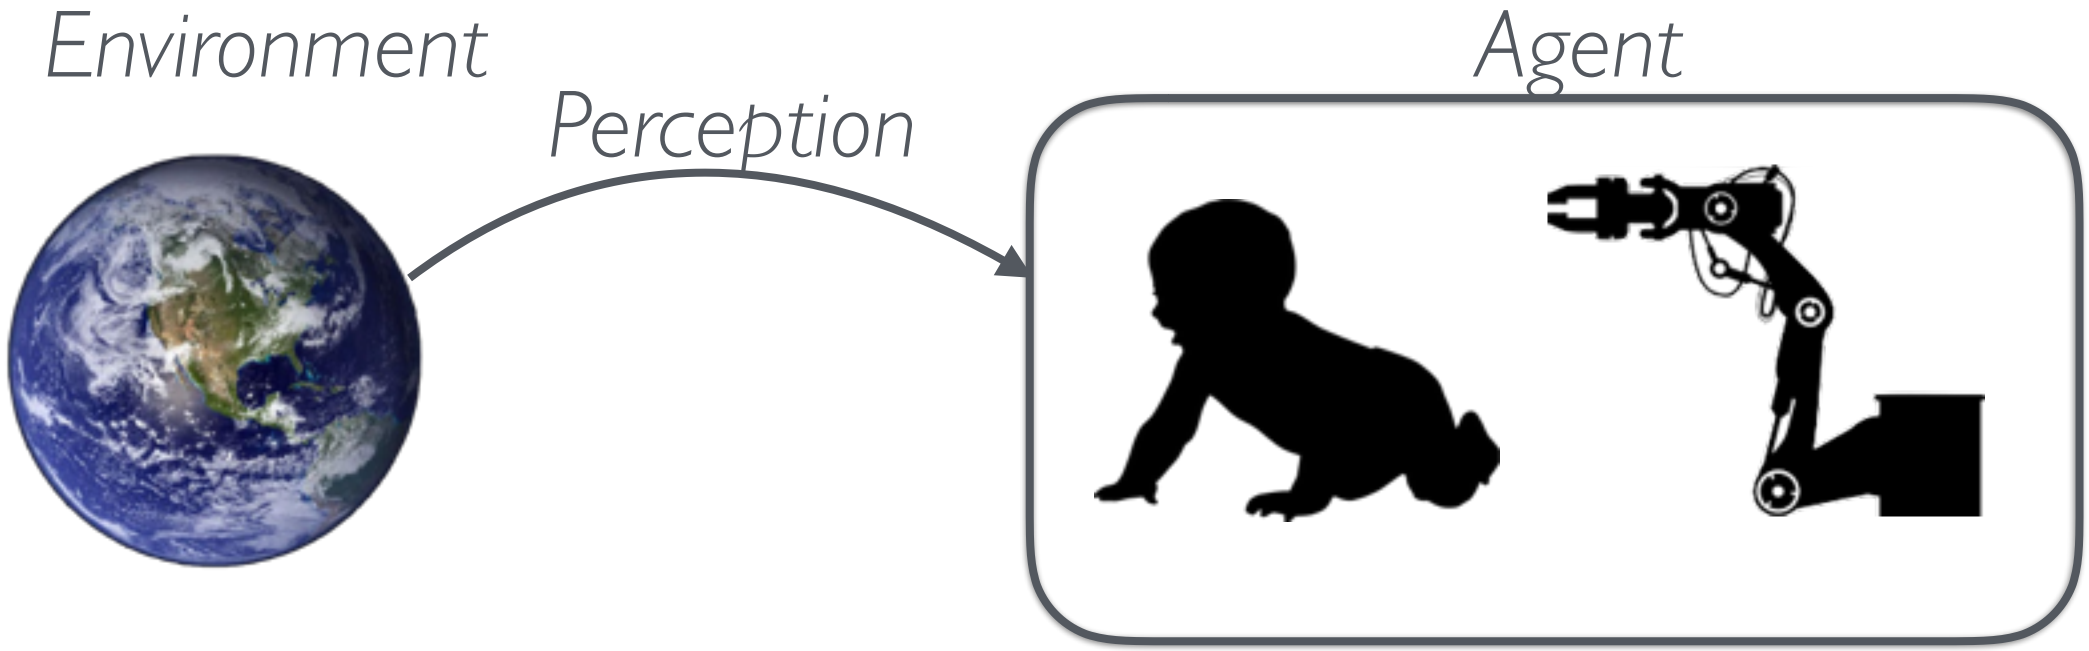
Environment
Perception
Agent

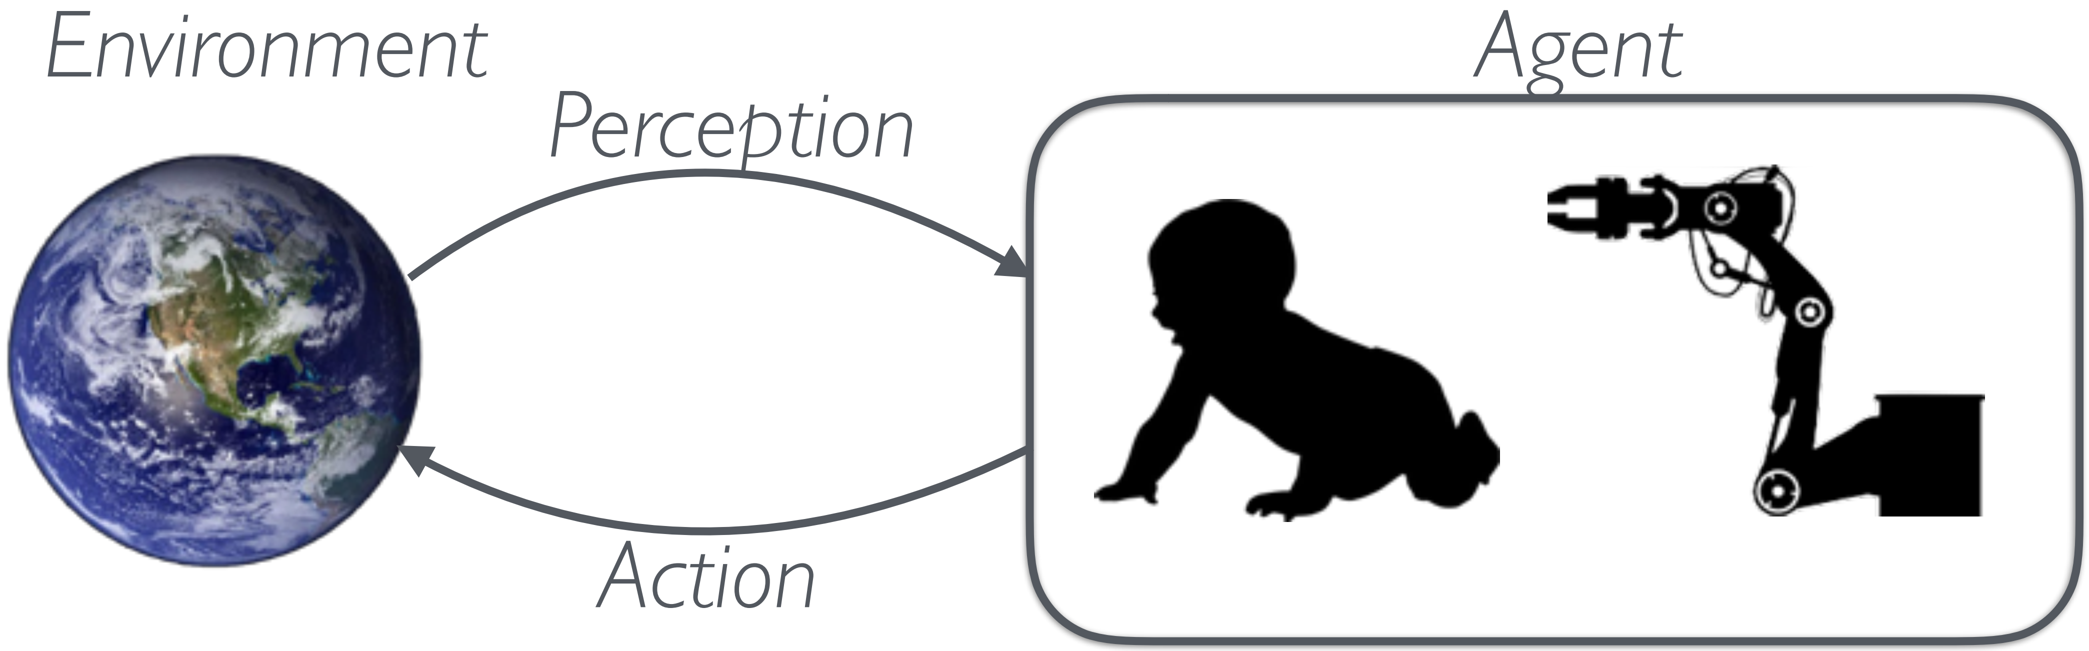
Environment
Perception
Agent
Action

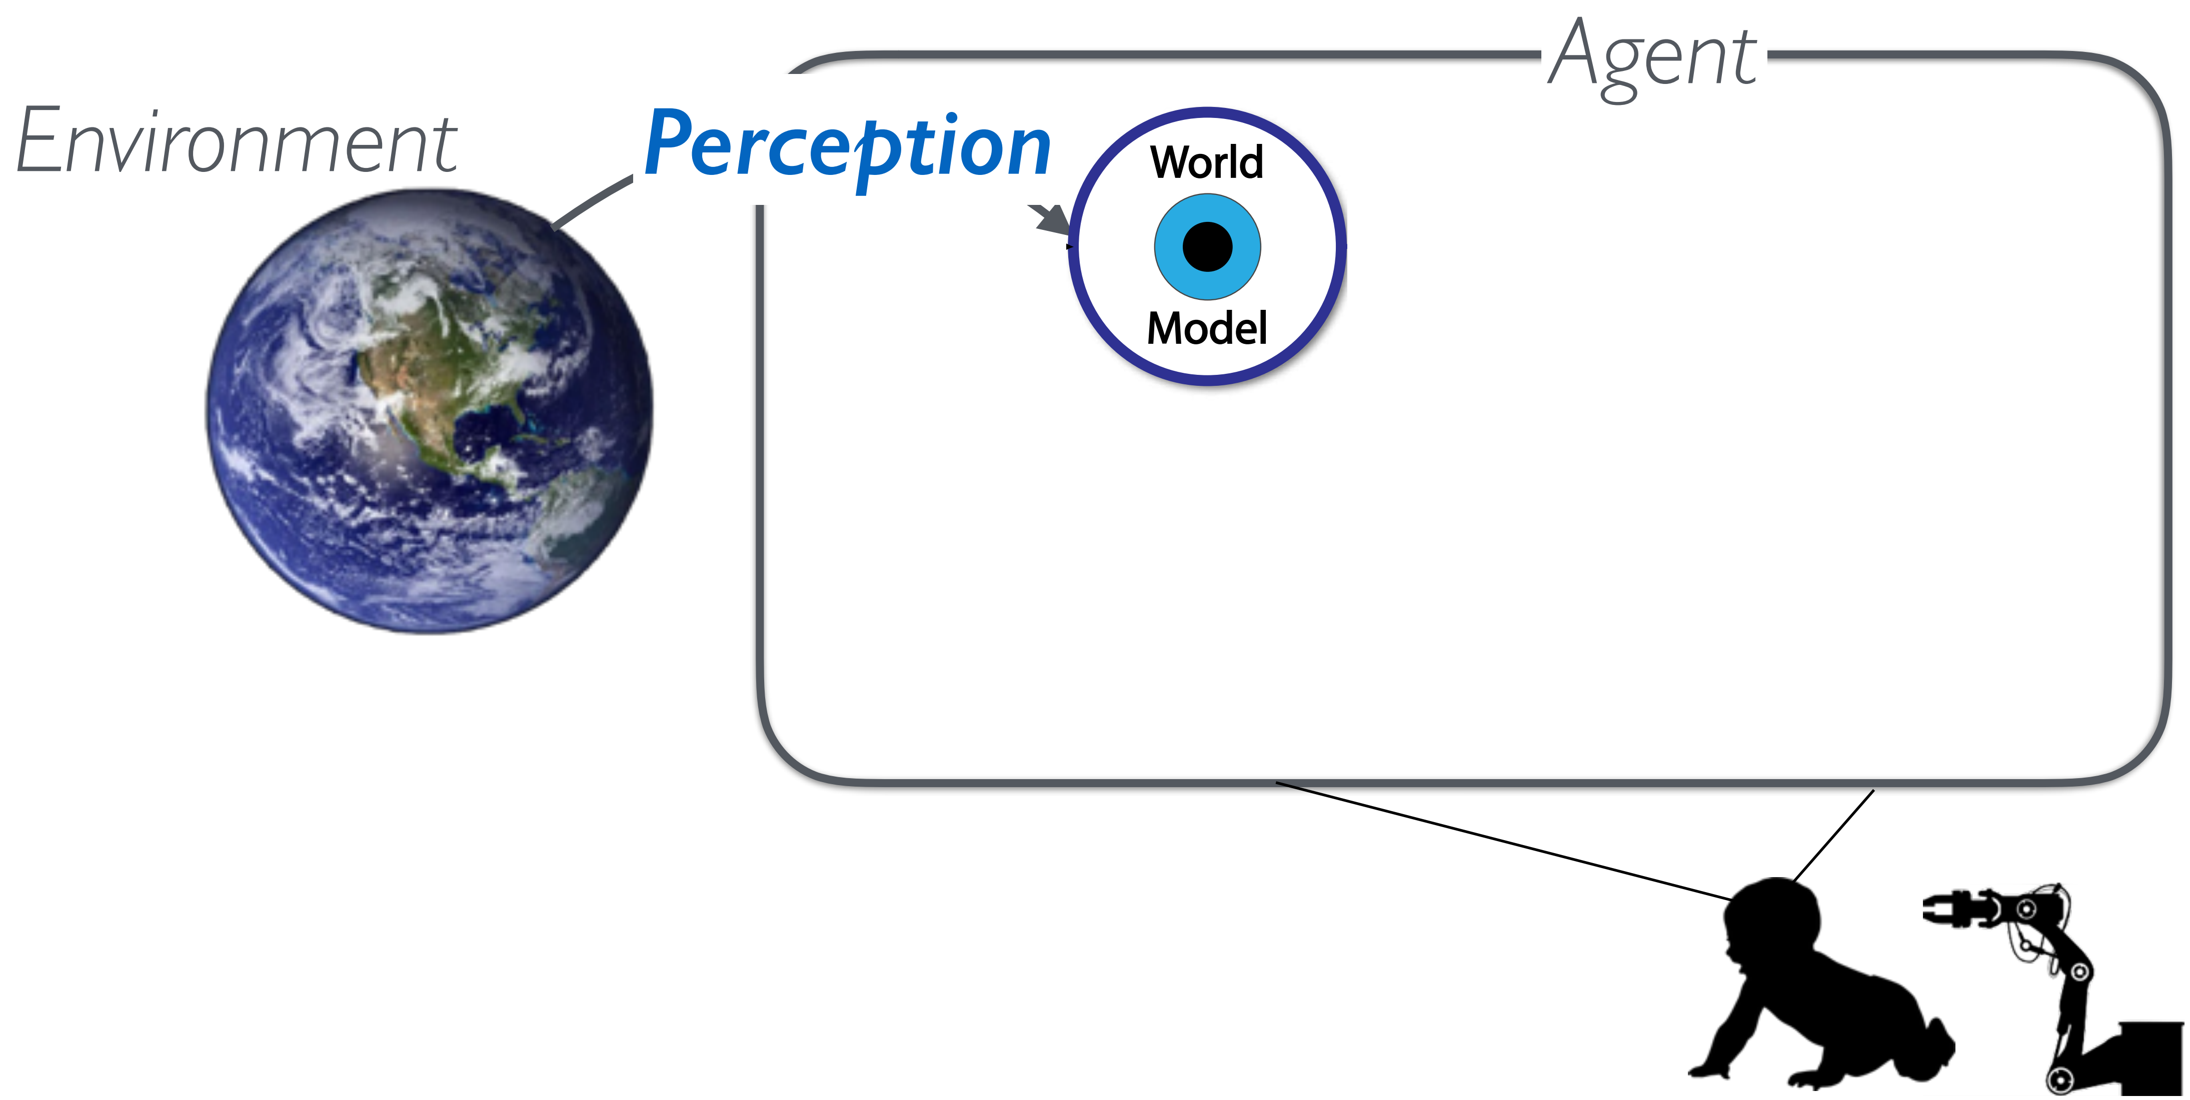
Environment
Perception
Agent
World
Model

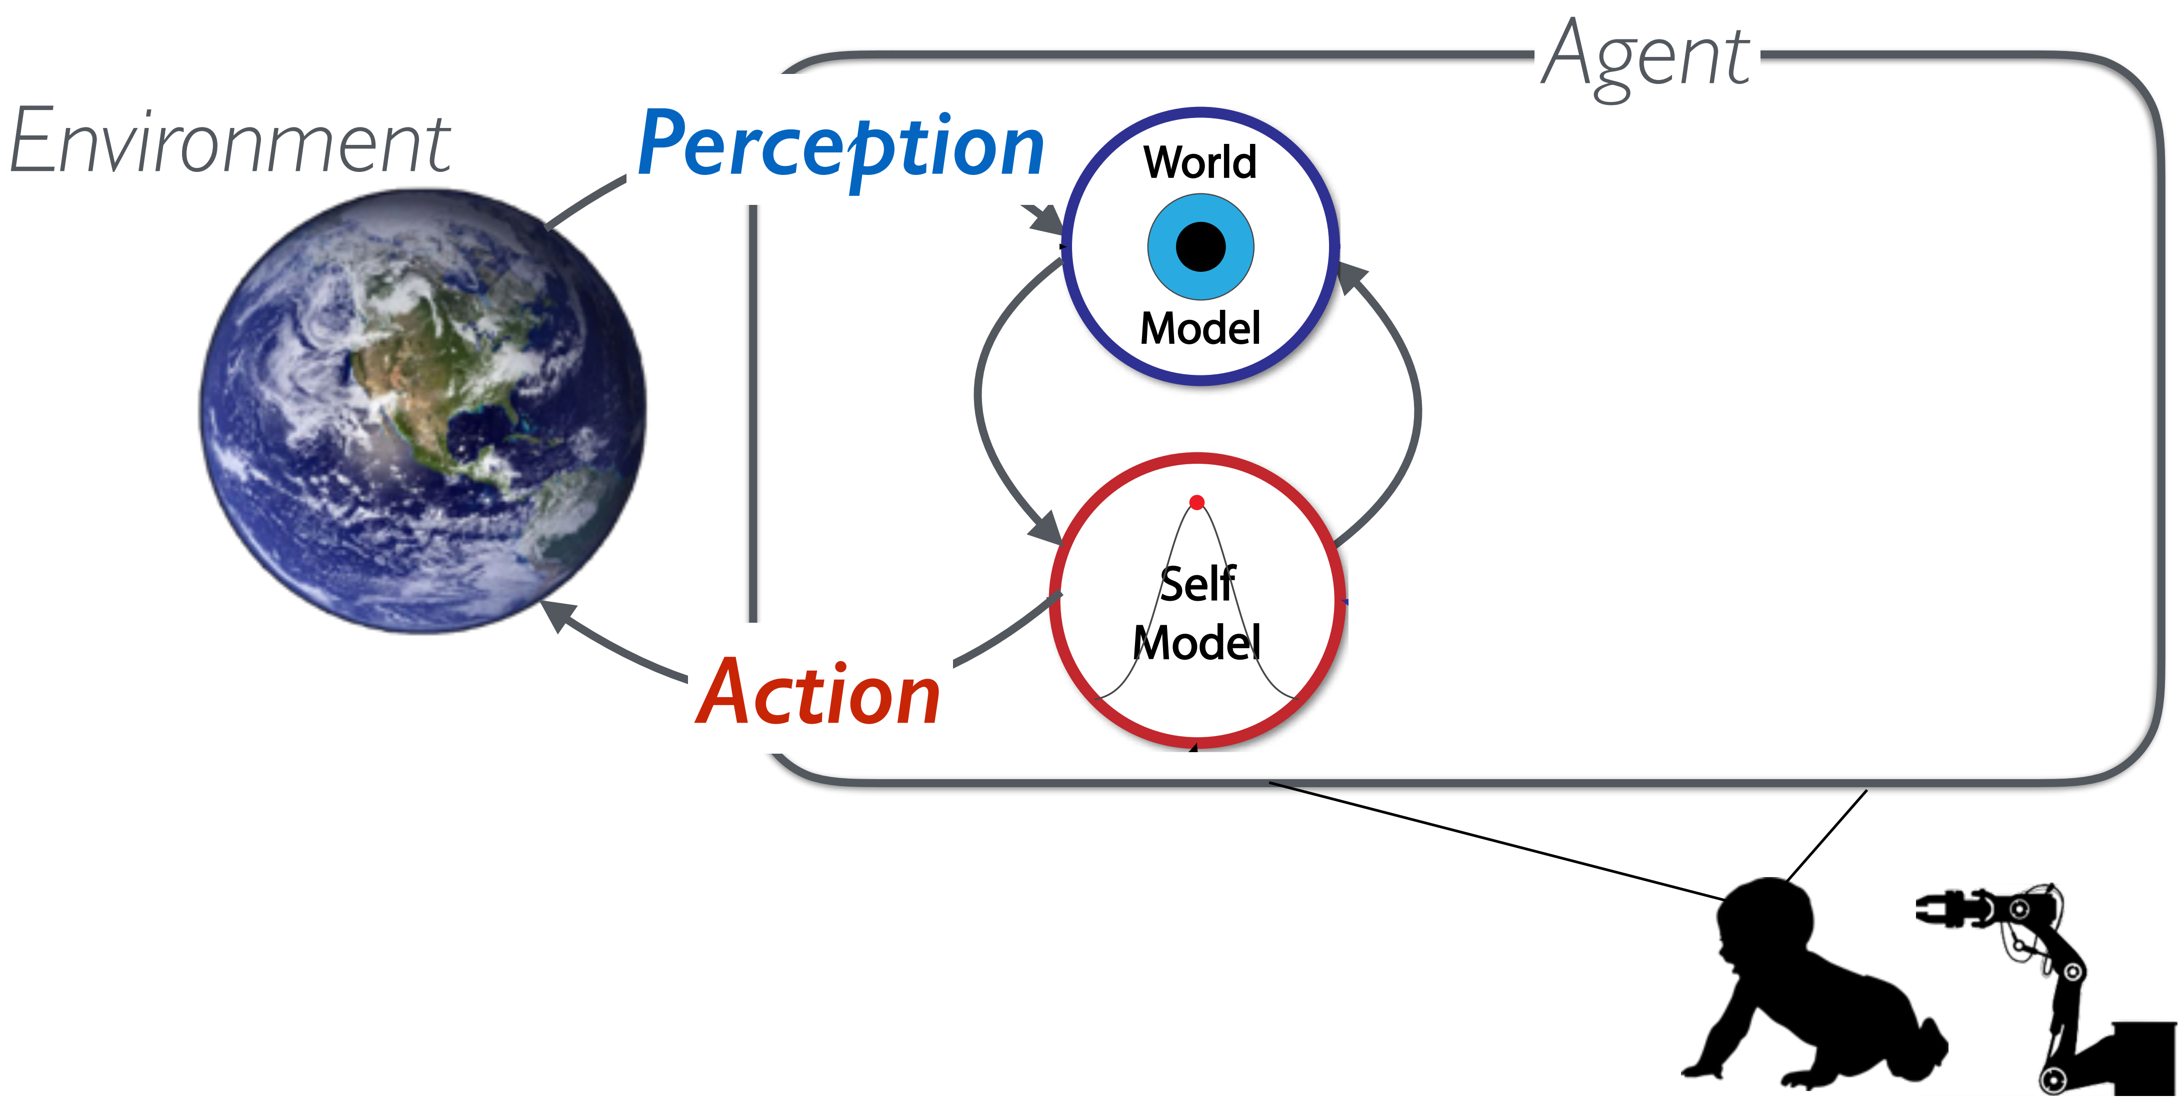

Environment
Perception
Agent
World
Model
Self
Model
Action

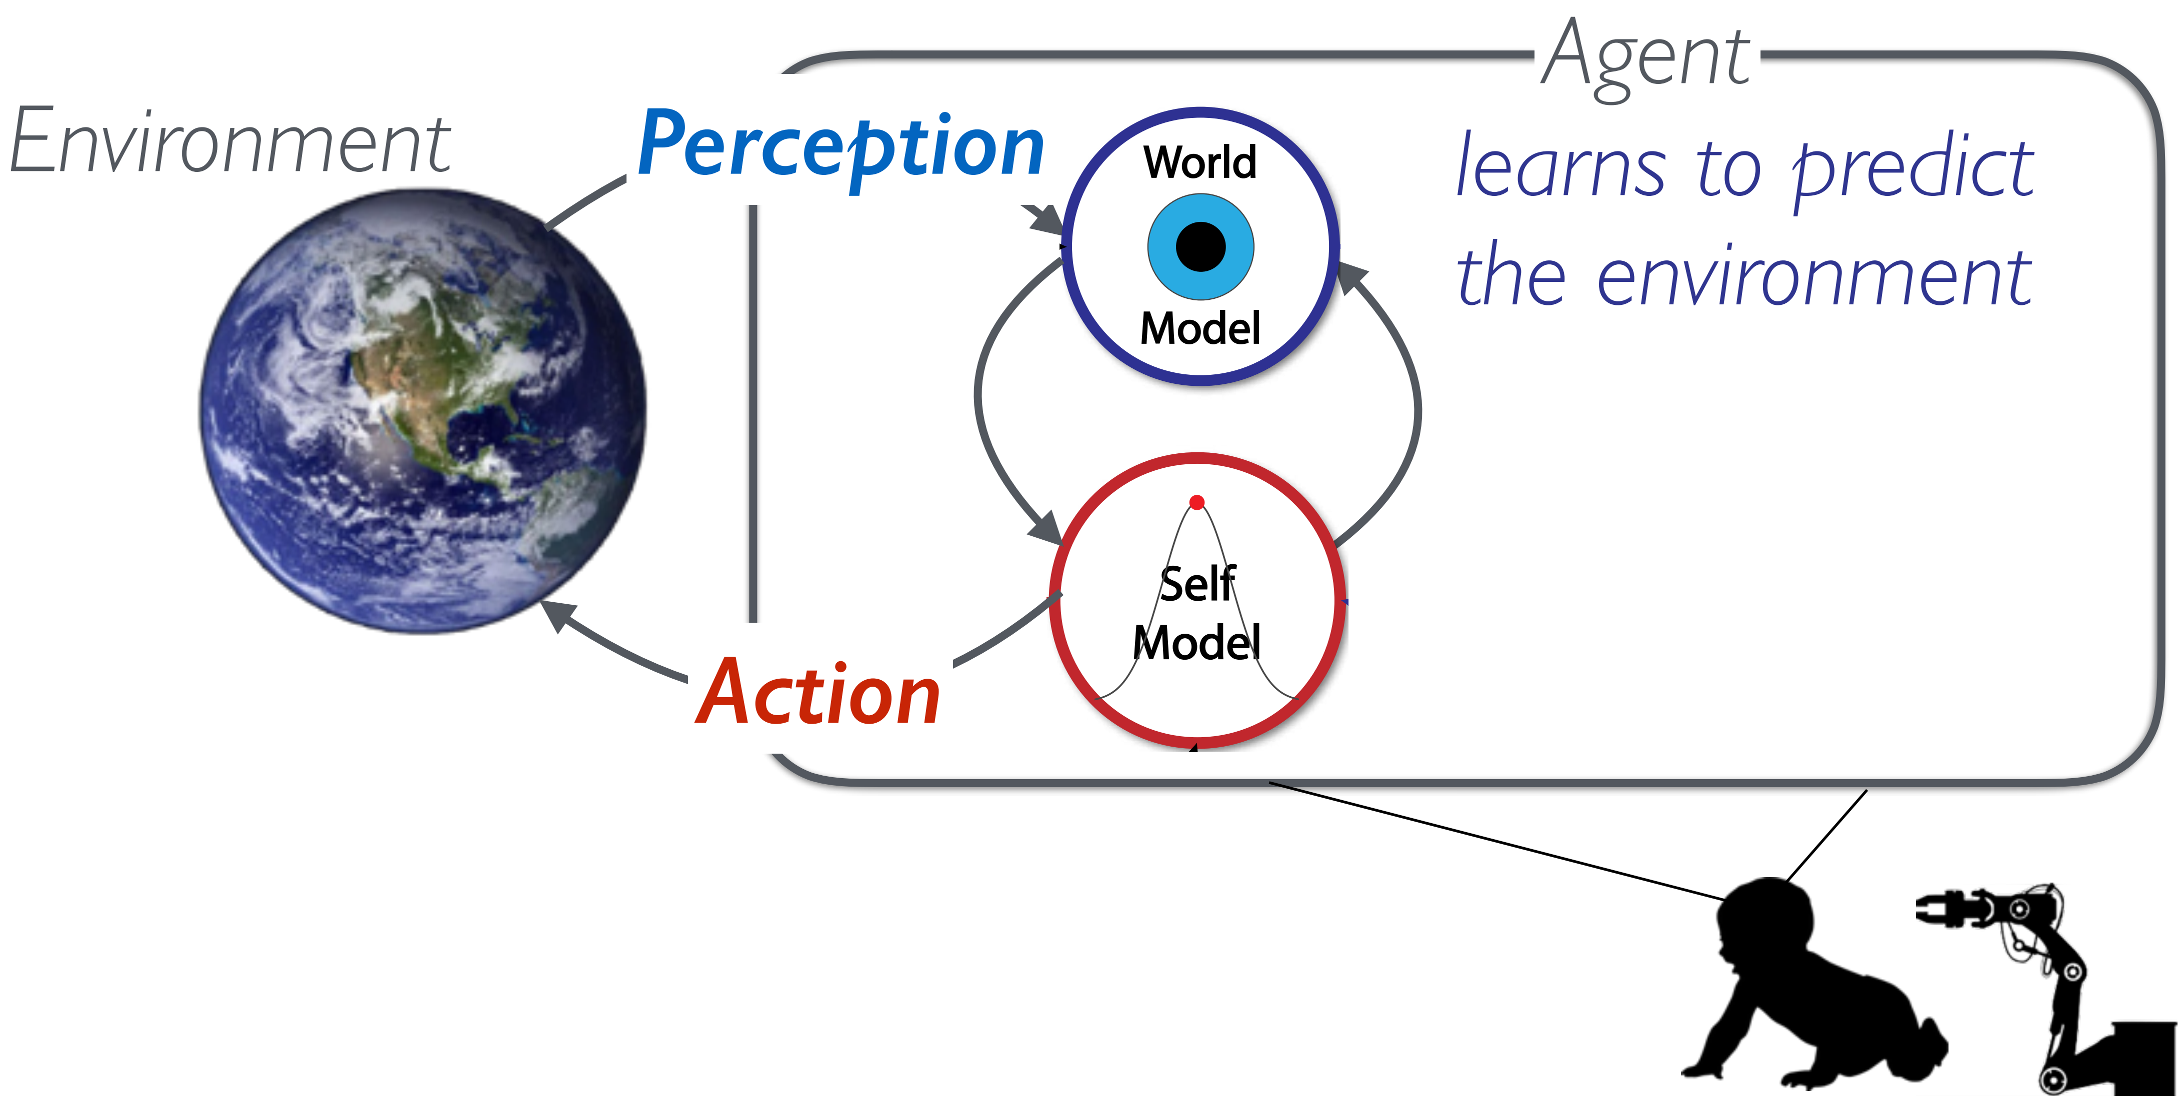
Environment
Perception
World
Model
Agent
learns to predict
the environment
Self
Model
Action

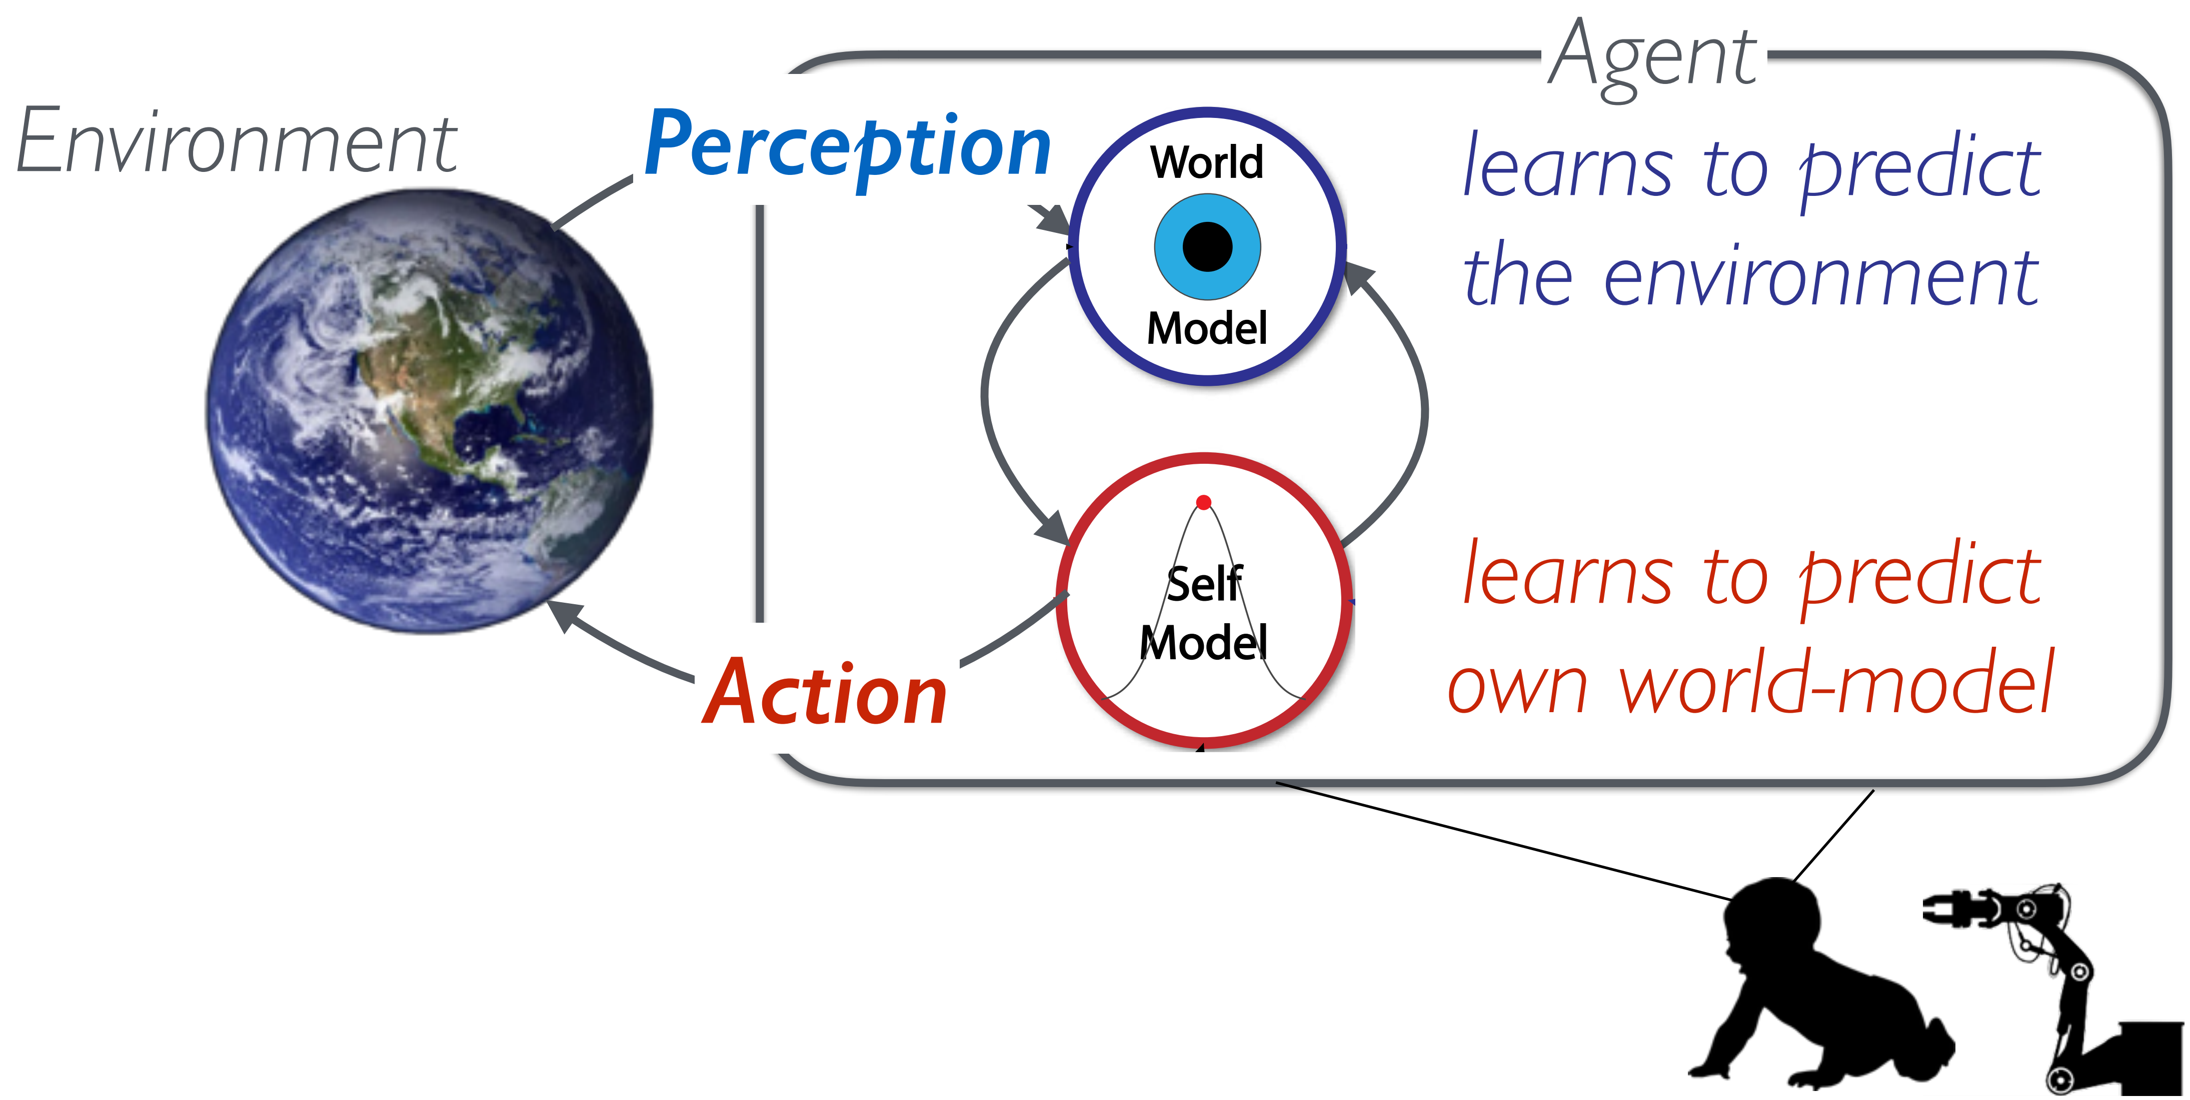
Agent
Environment
Perception
World
Model
learns to predict
the environment
Self
Model
Action
learns to predict
own world-model

A curiosity principle:

The ***self-model*** directs the agent toward ***interesting*** actions — the ones that the ***world-model*** doesn't yet fully understand

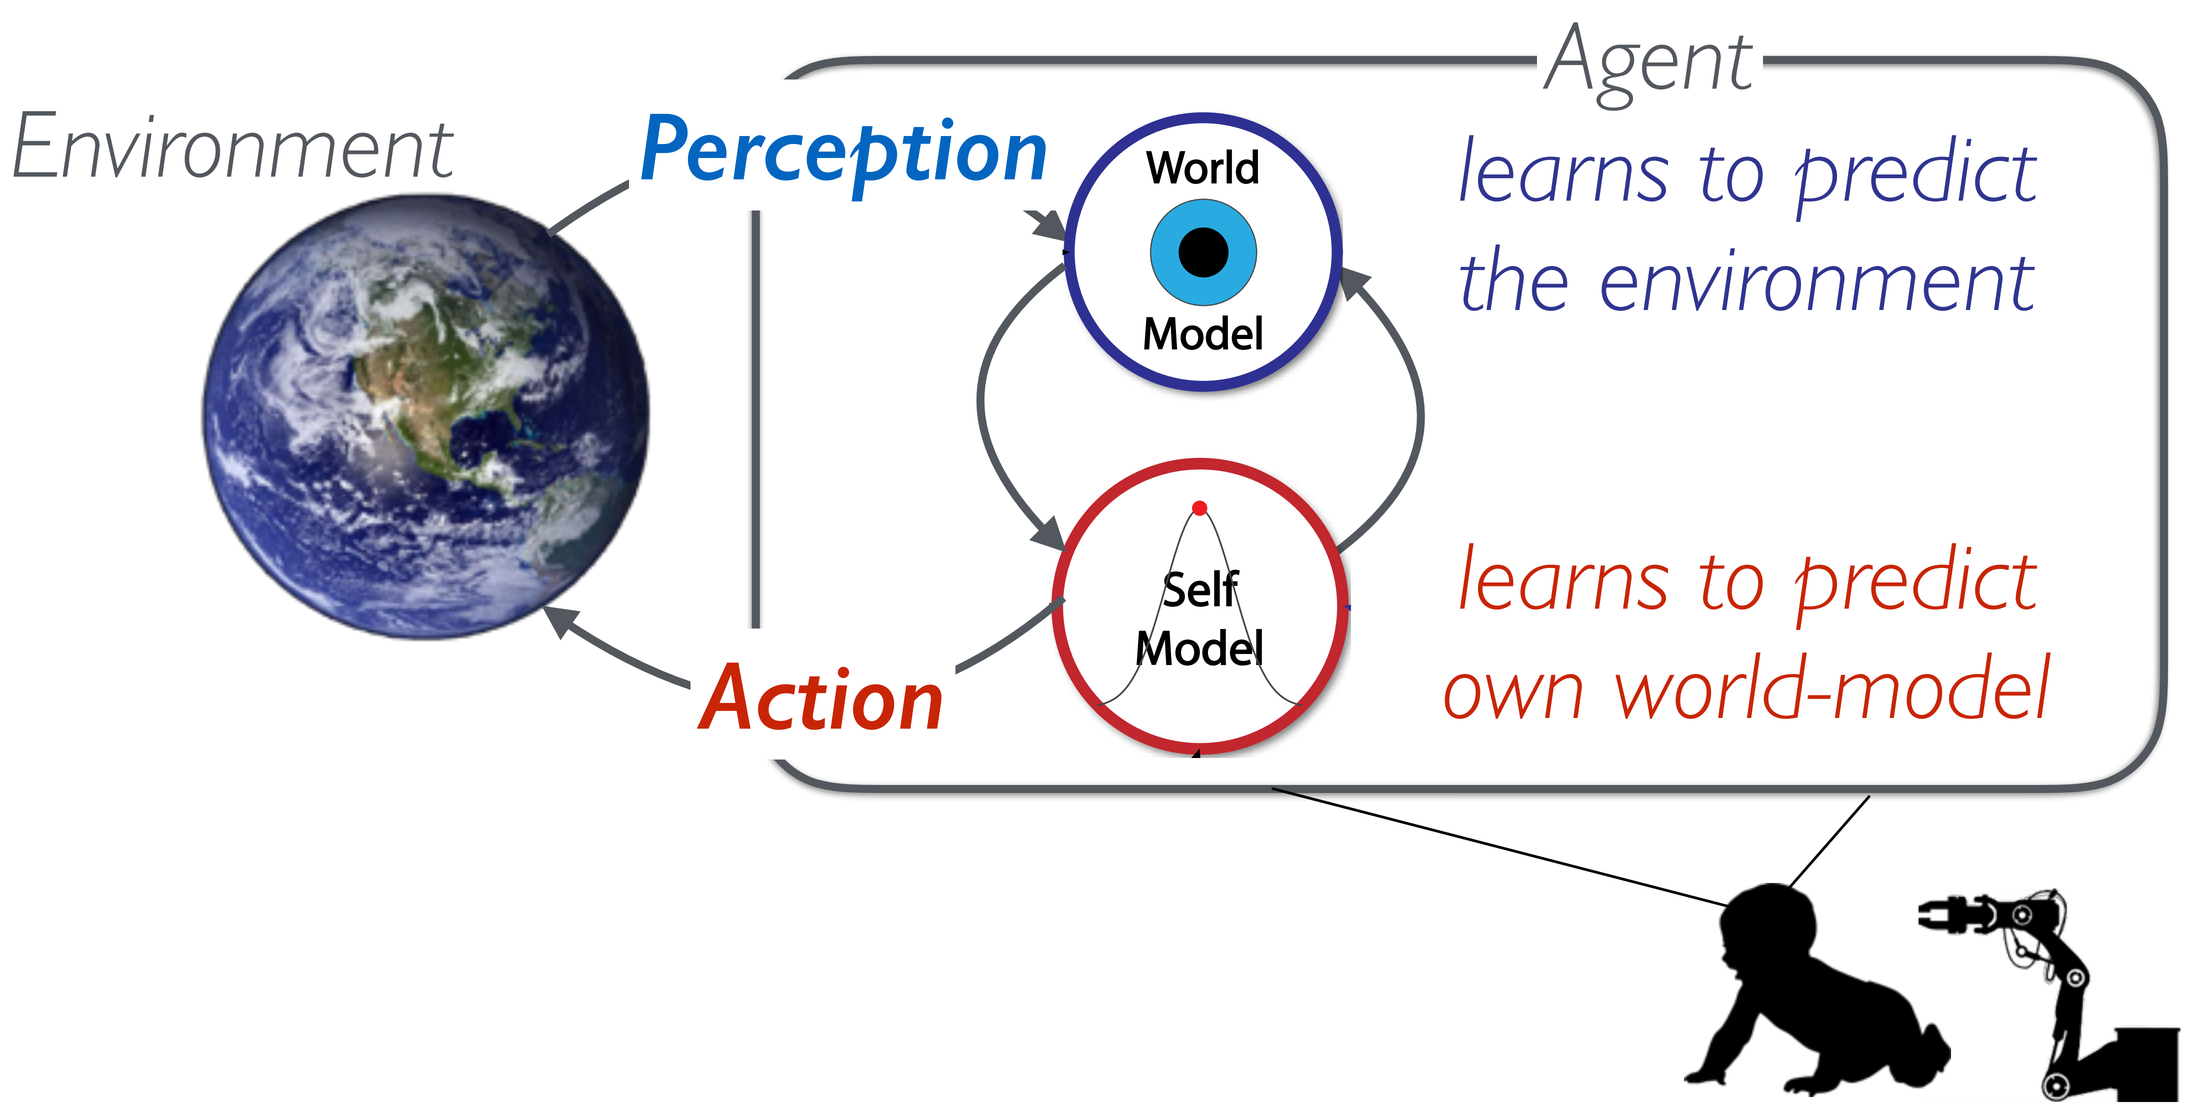

A curiosity principle:

The ***self-model*** directs the agent toward ***interesting*** actions — the ones that the ***world-model*** doesn't yet fully understand

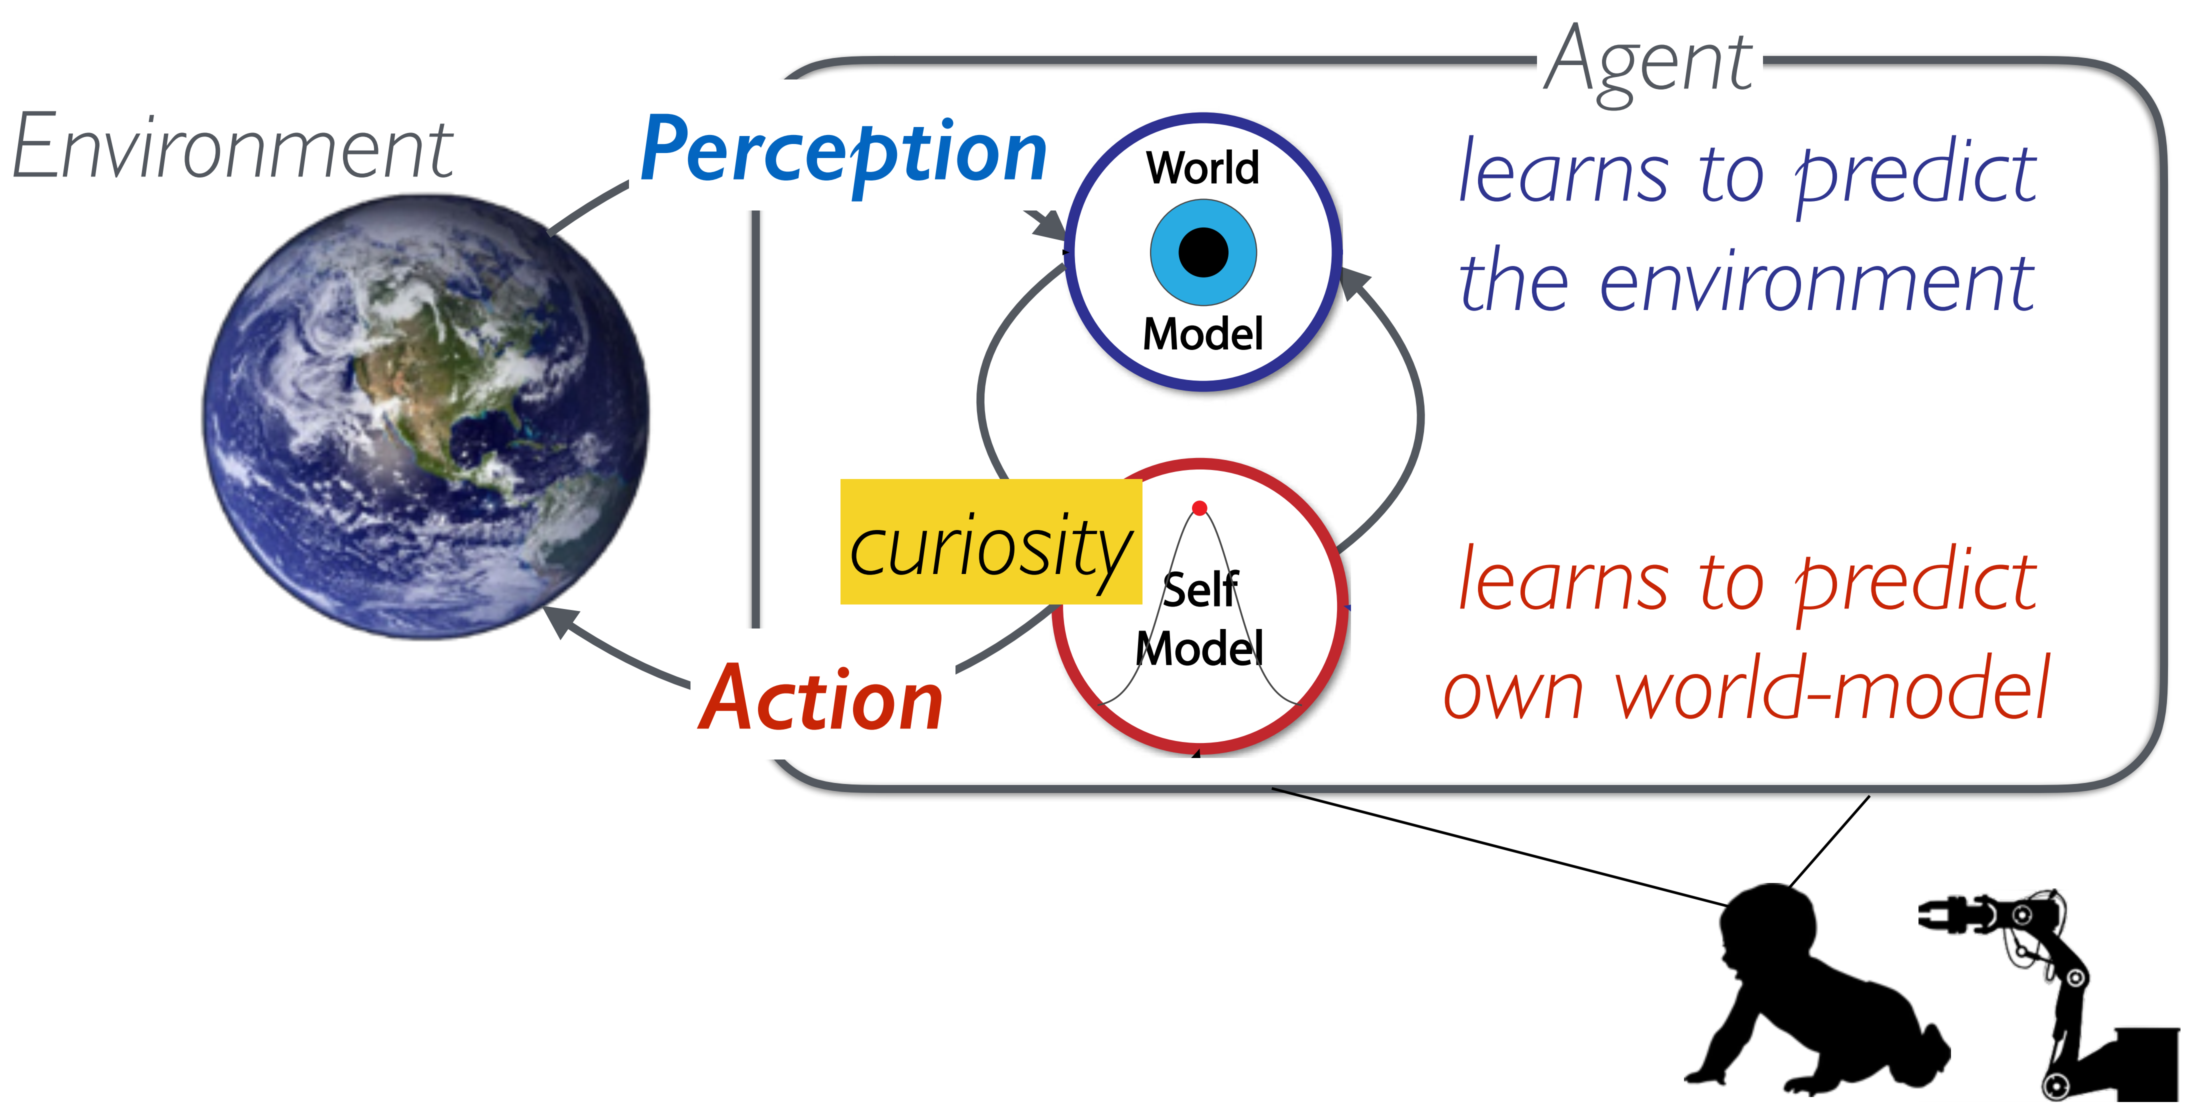

A curiosity principle:

The ***self-model*** directs the agent toward ***interesting*** actions — the ones that the ***world-model*** doesn't yet fully understand

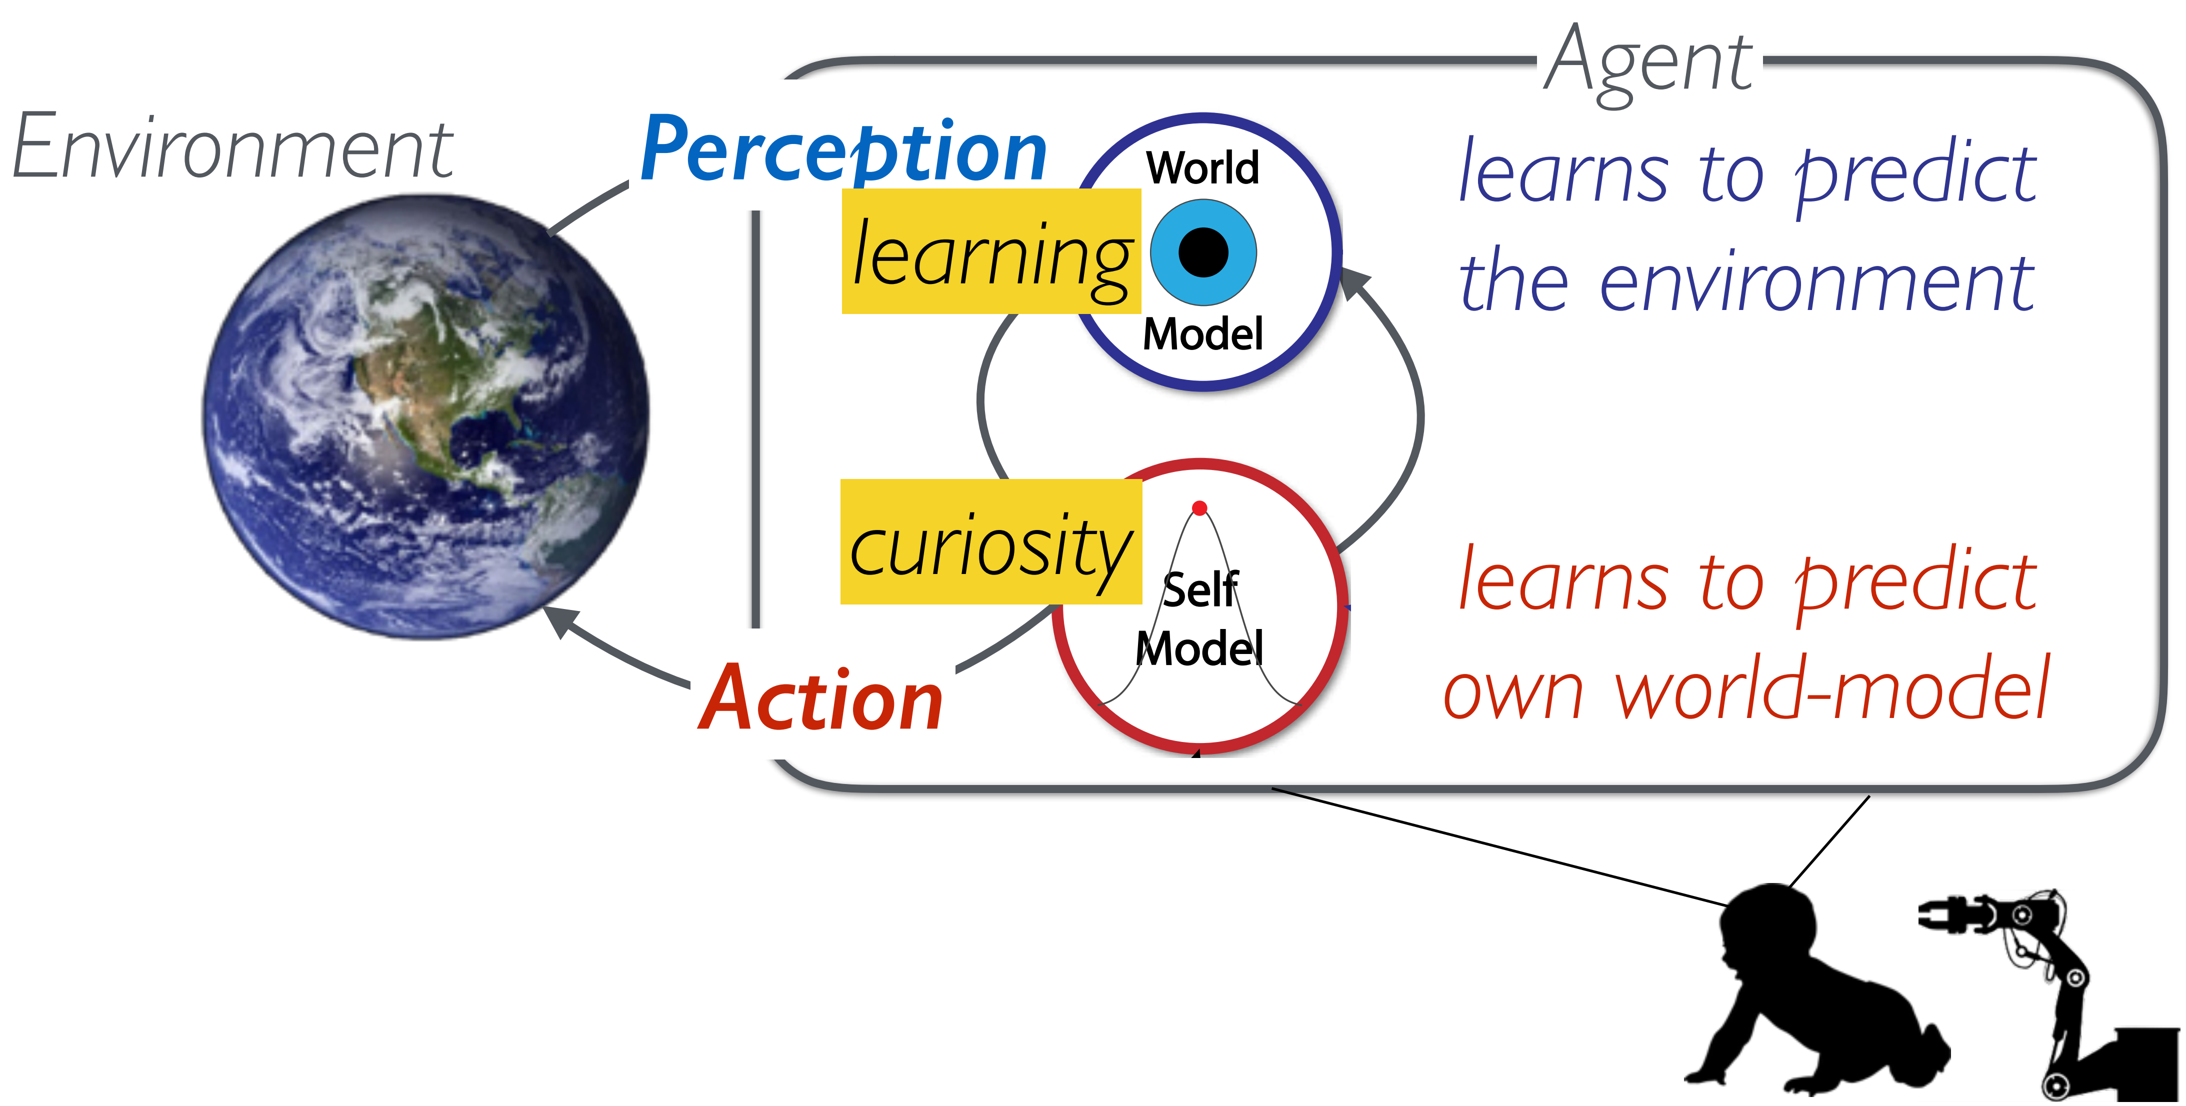

A curiosity principle:

The ***self-model*** directs the agent toward ***interesting*** actions — the ones that the ***world-model*** doesn't yet fully understand

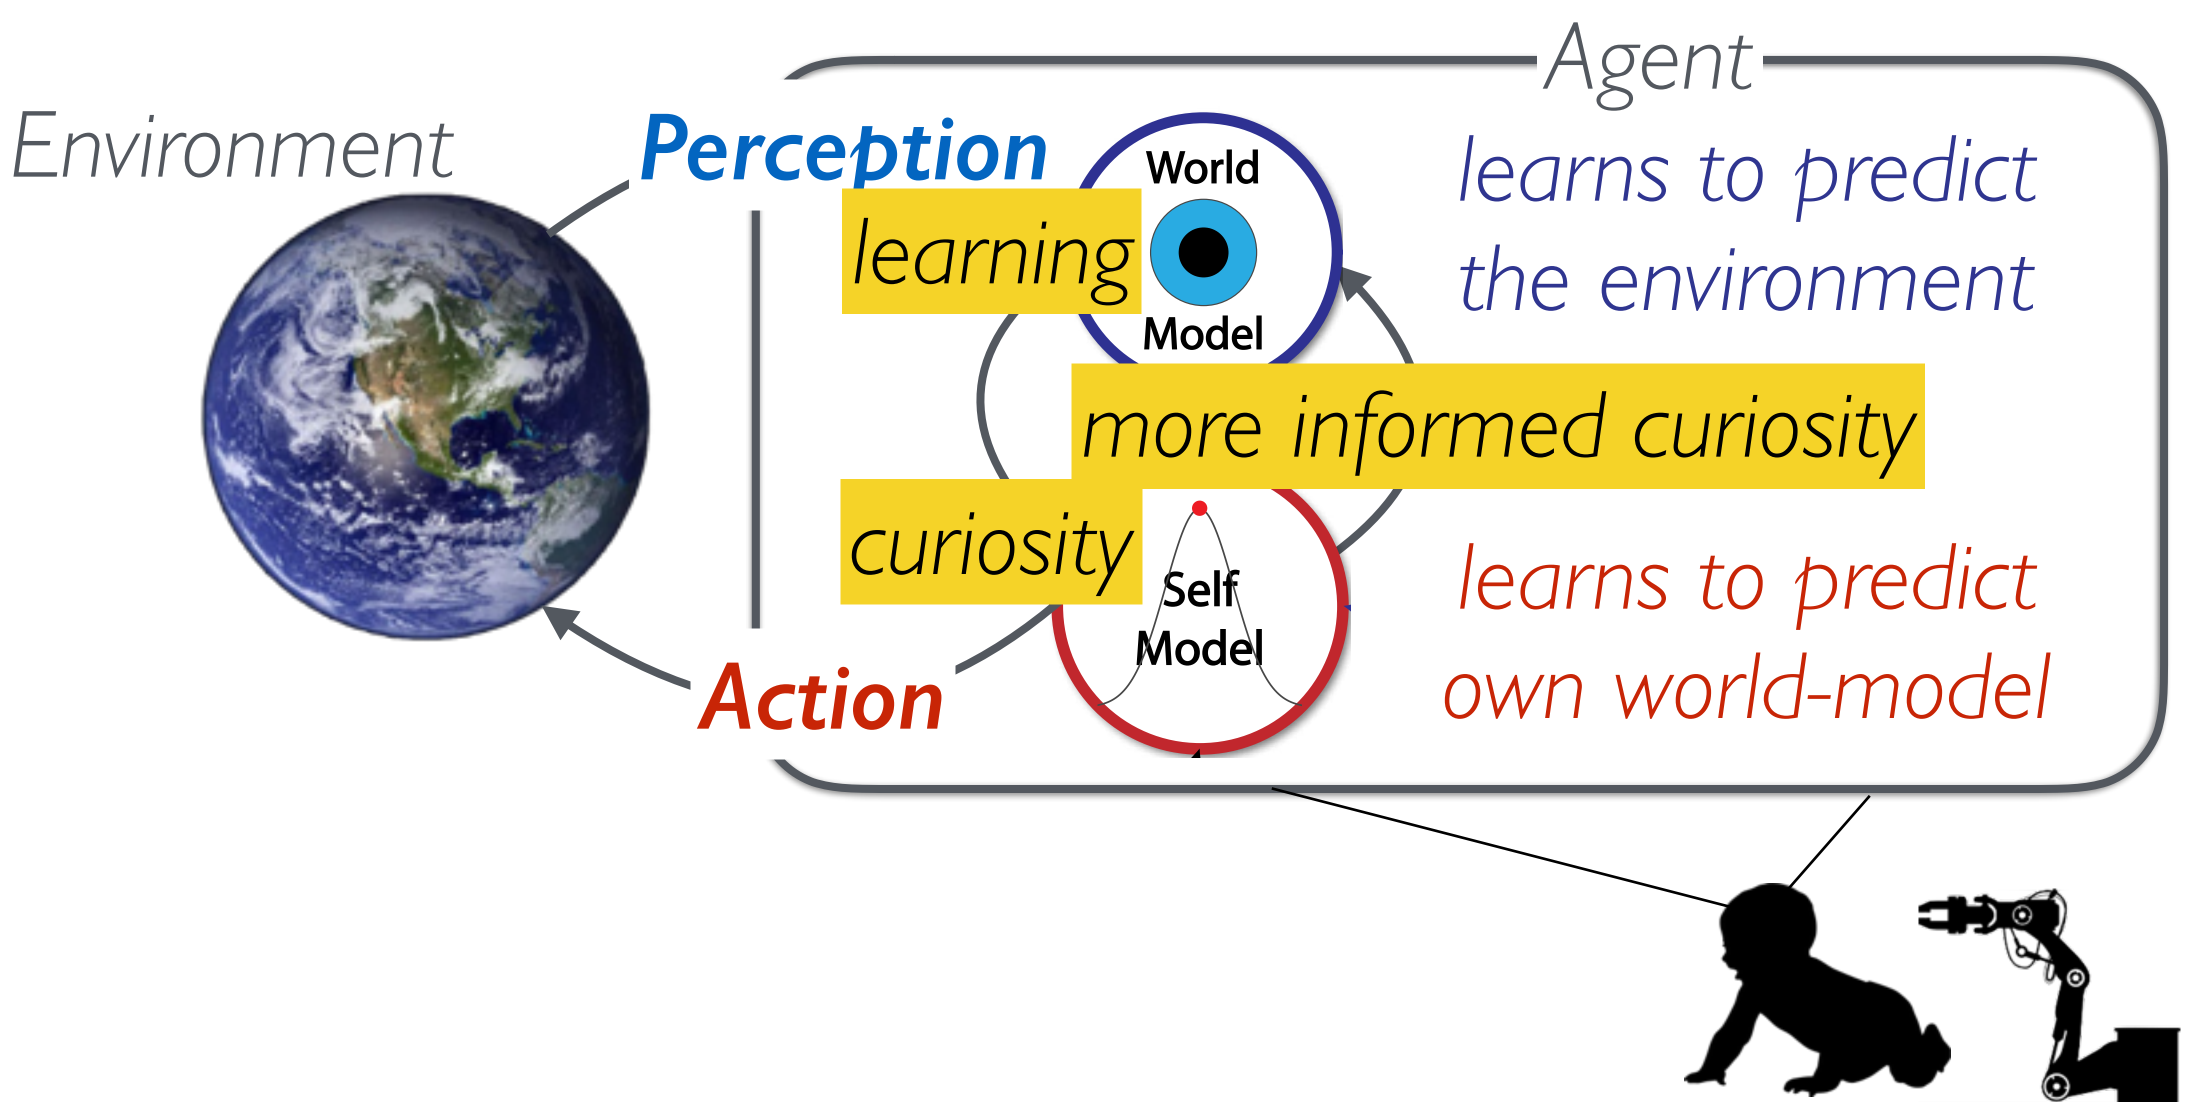

Learning to Play

Nick Haber

Damian Mrowca

Stephanie Wang

Fei-Fei Li

NIPS 2018

Learning to Play With Intrinsically-Motivated, Self-Aware Agents

Nick Haber[1,2,3,*], **Damian Mrowca**[4,*], **Stephanie Wang**[4], **Li Fei-Fei**[4], **and Daniel L. K. Yamins**[1,4,5]

Departments of Psychology[1], Pediatrics[2], Biomedical Data Science[3], Computer Science[4], and Wu Tsai Neurosciences Institute[5], Stanford, CA 94305

{nhaber, mrowca}@stanford.edu

Abstract

Infants are experts at playing, with an amazing ability to generate novel structured behaviors in unstructured environments that lack clear extrinsic reward signals. We seek to mathematically formalize these abilities using a neural network that implements curiosity-driven intrinsic motivation. Using a simple but ecologically naturalistic simulated environment in which an agent can move and interact with objects it sees, we propose a "world-model" network that learns to predict the dynamic consequences of the agent's actions. Simultaneously, we train a separate explicit "self-model" that allows the agent to track the error map of its world-model. It then uses the self-model to adversarially challenge the developing world-model. We demonstrate that this policy causes the agent to explore novel and informative interactions with its environment, leading to the generation of a spectrum of complex behaviors, including ego-motion prediction, object attention, and object gathering. Moreover, the world-model that the agent learns supports improved performance on object dynamics prediction, detection, localization and recognition tasks. Taken together, our results are initial steps toward creating flexible autonomous agents that self-supervise in realistic physical environments.

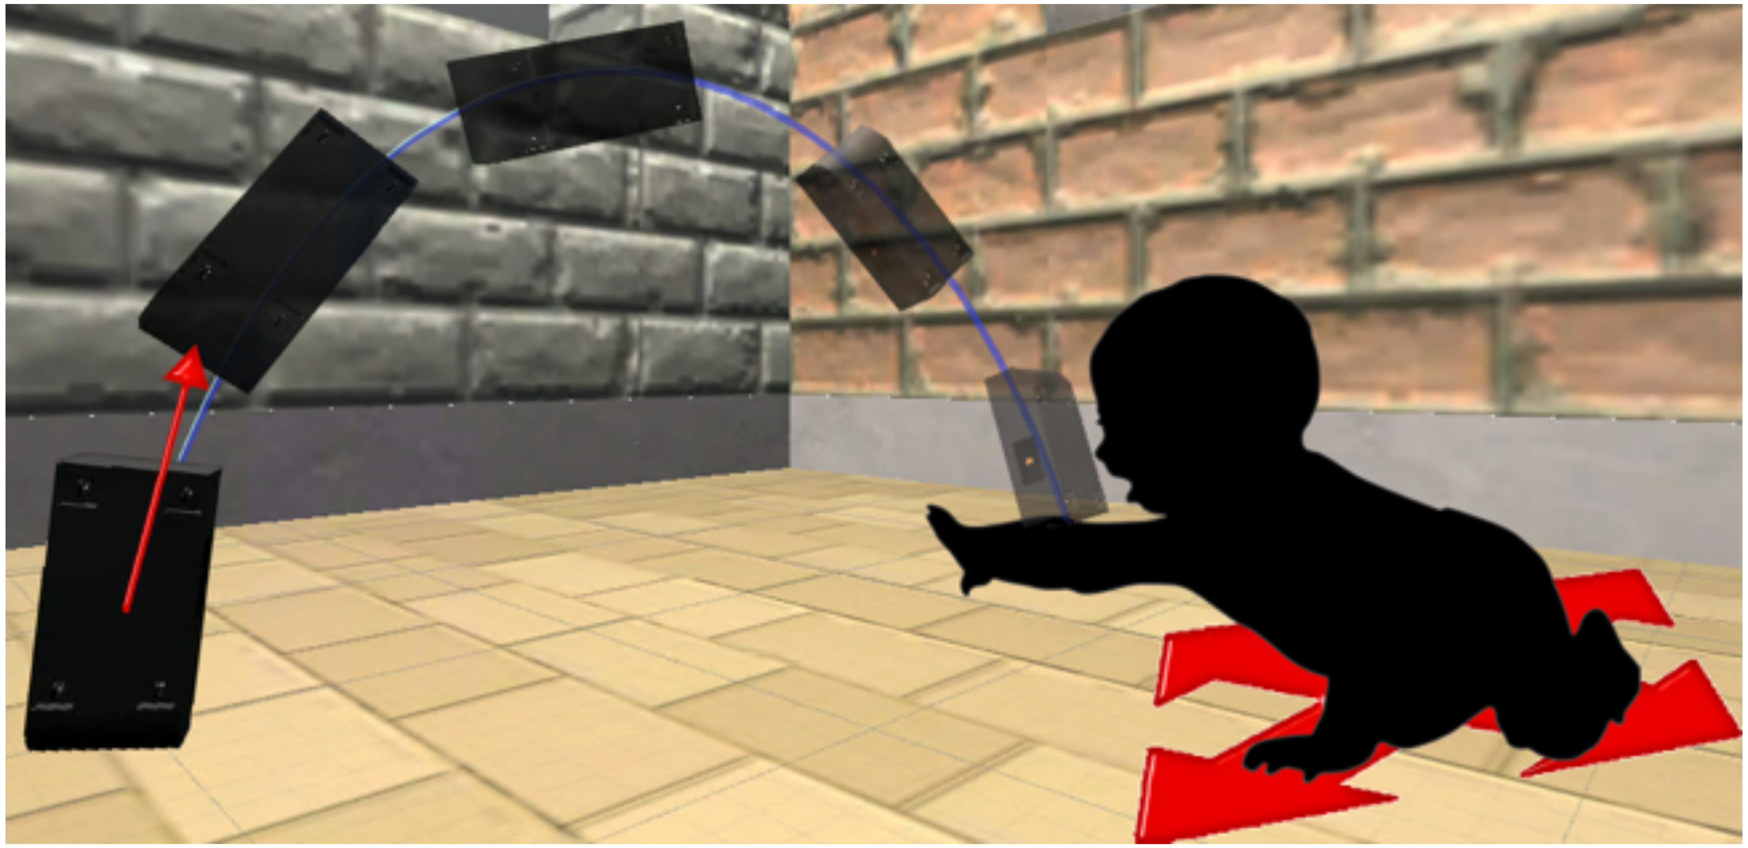

Agent ("baby") can (a) swivel its head

(b) move around the room

(c) apply forces to objects

. . . and receives back images of what happened, given action

Learning to Play — Overview

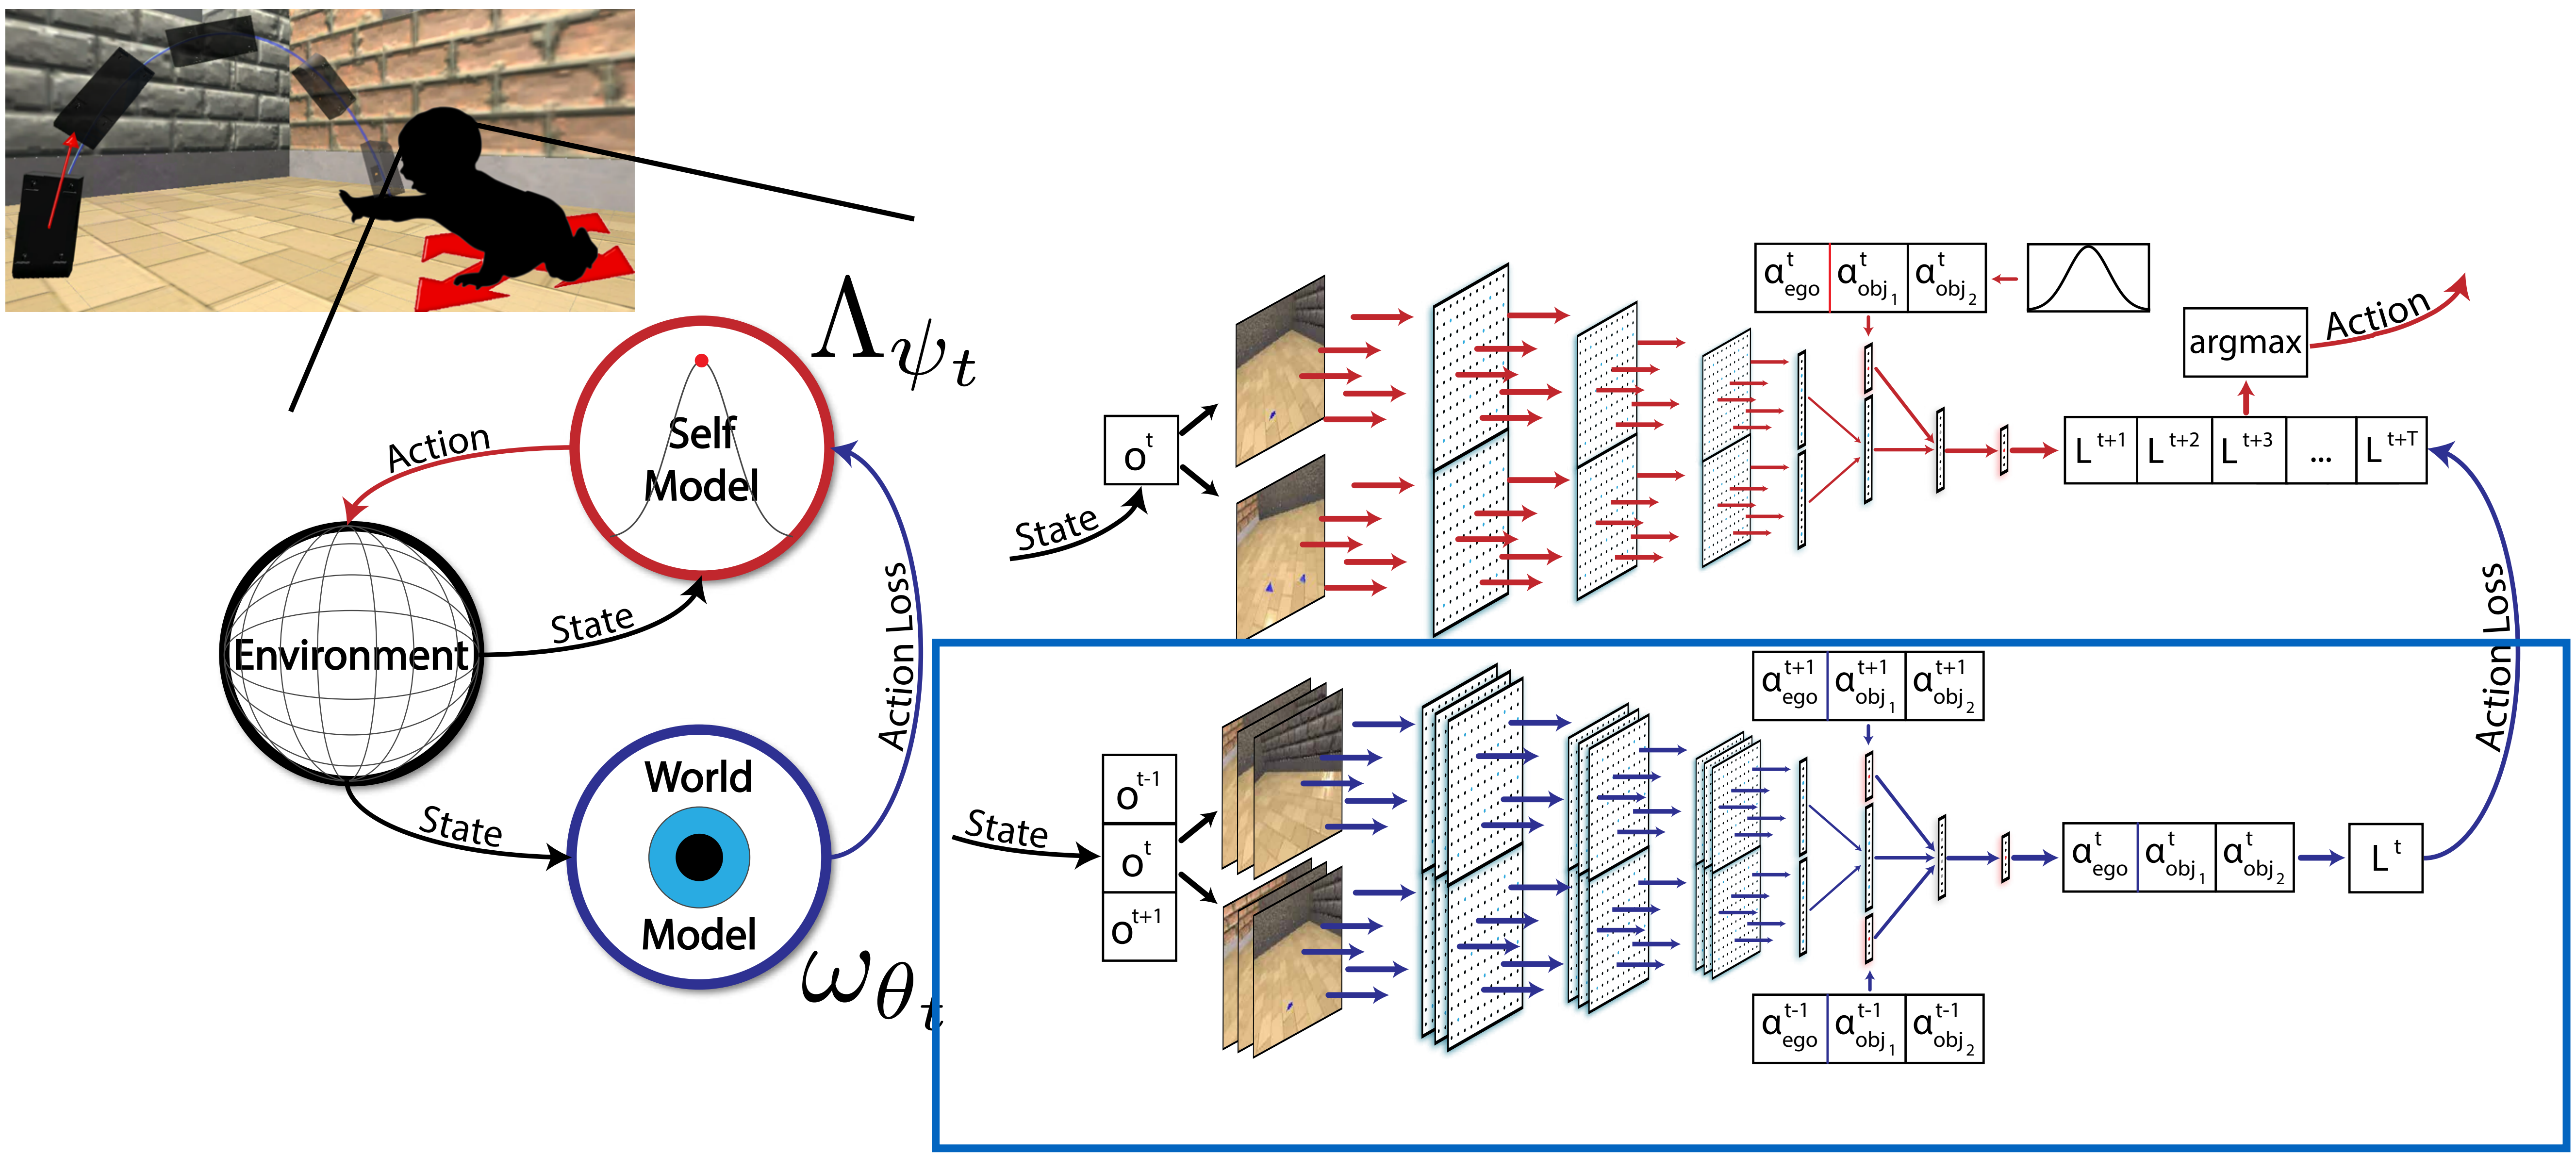

Model has two pieces: **(1) World-Model**

Learning to Play — Overview

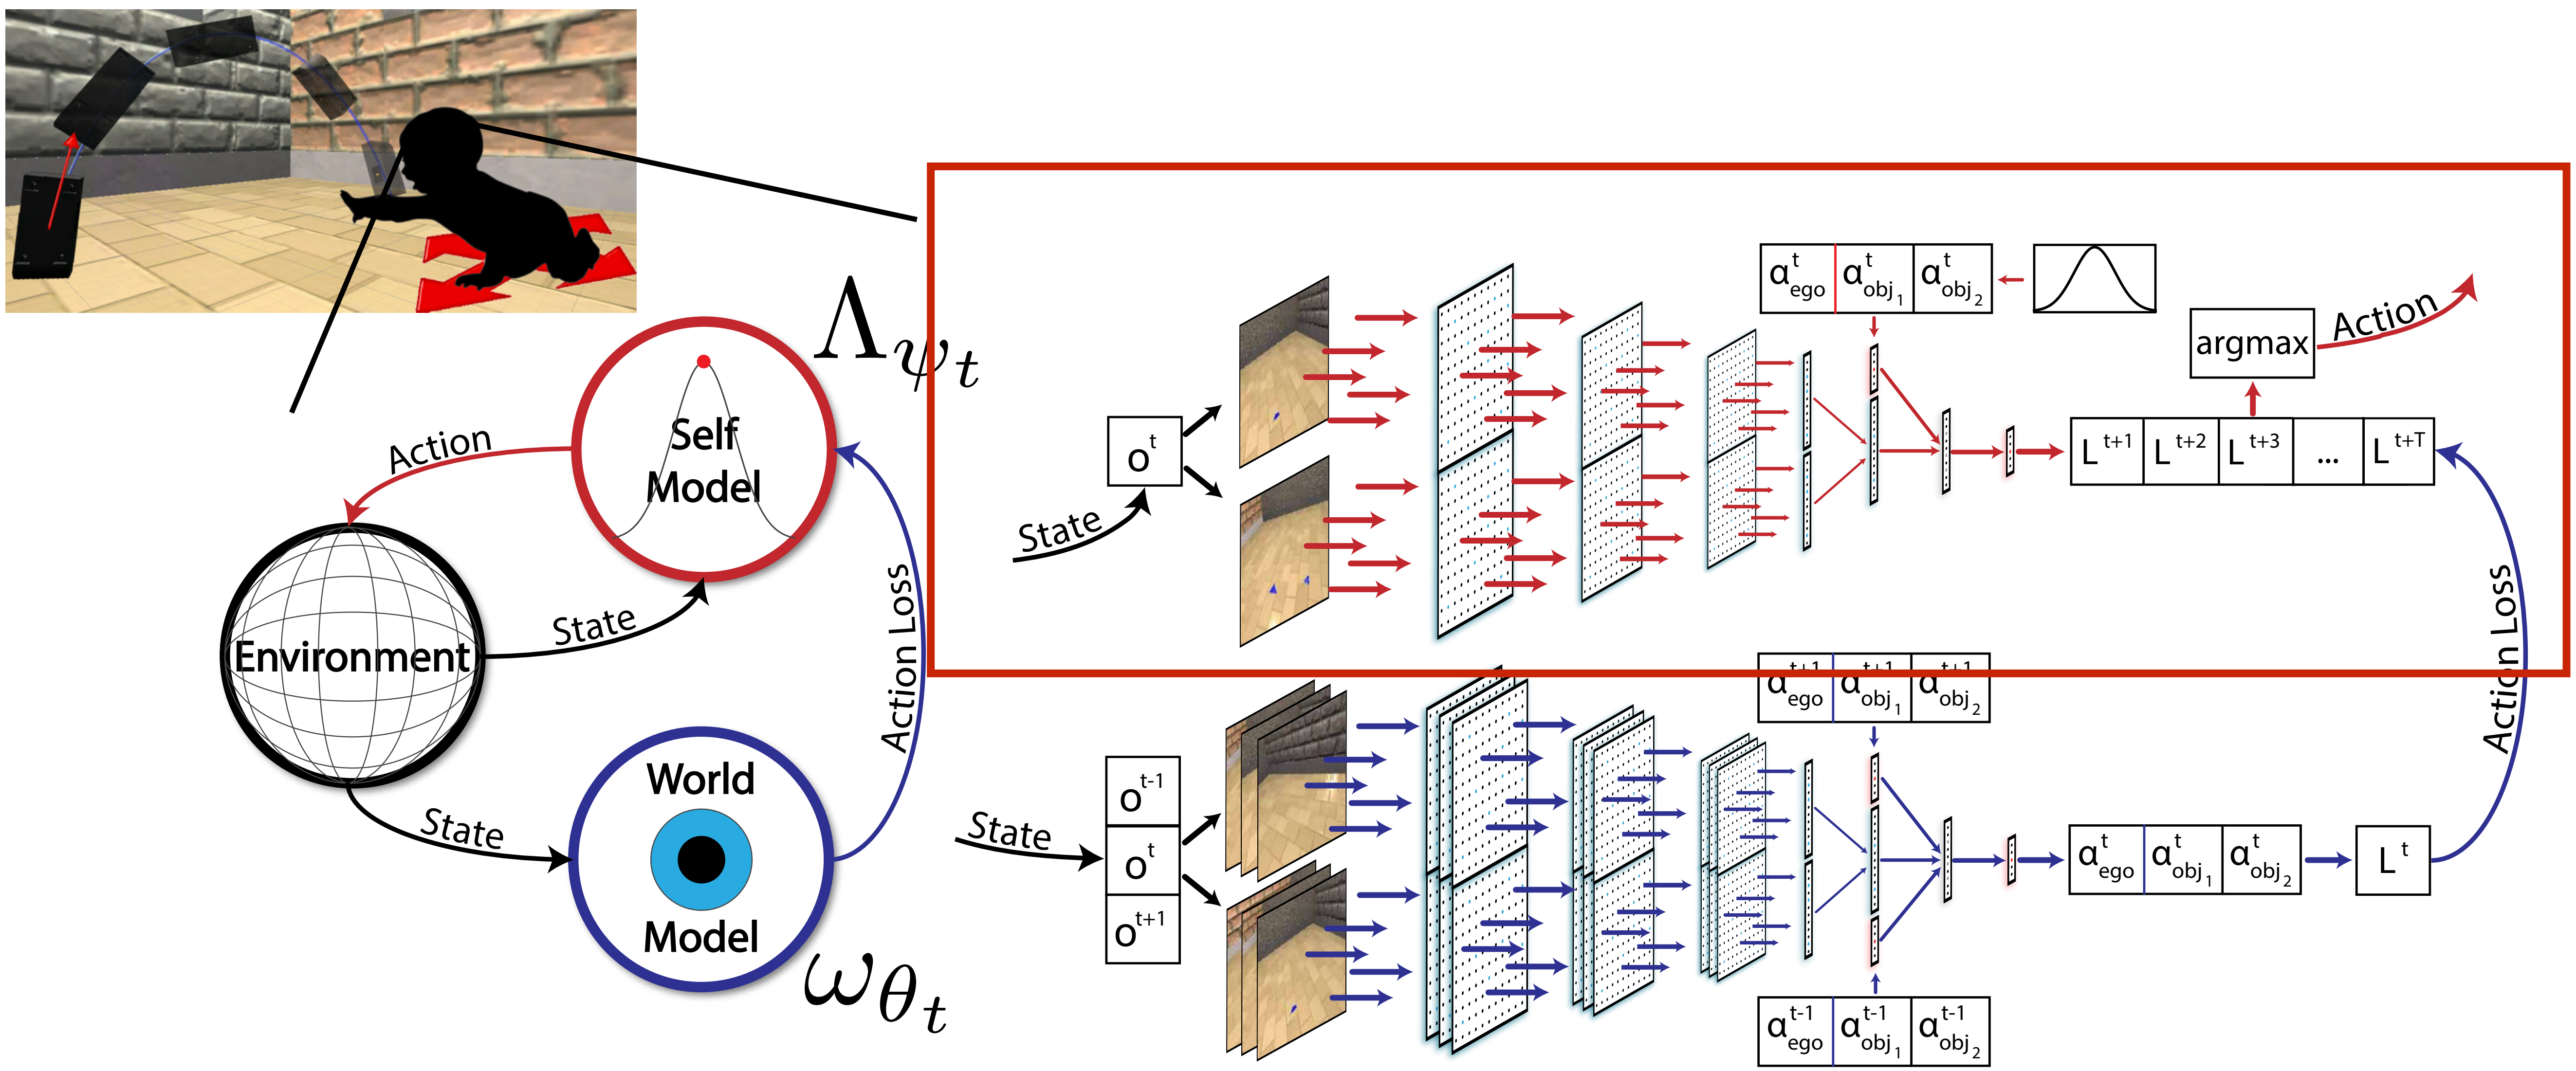

Model has two pieces: **(1) World-Model**

(2) Self-Model

Learning to Play — World-Model

Goal of ***world-model****:*
"Post-dict" the action taken given past and future states and actions

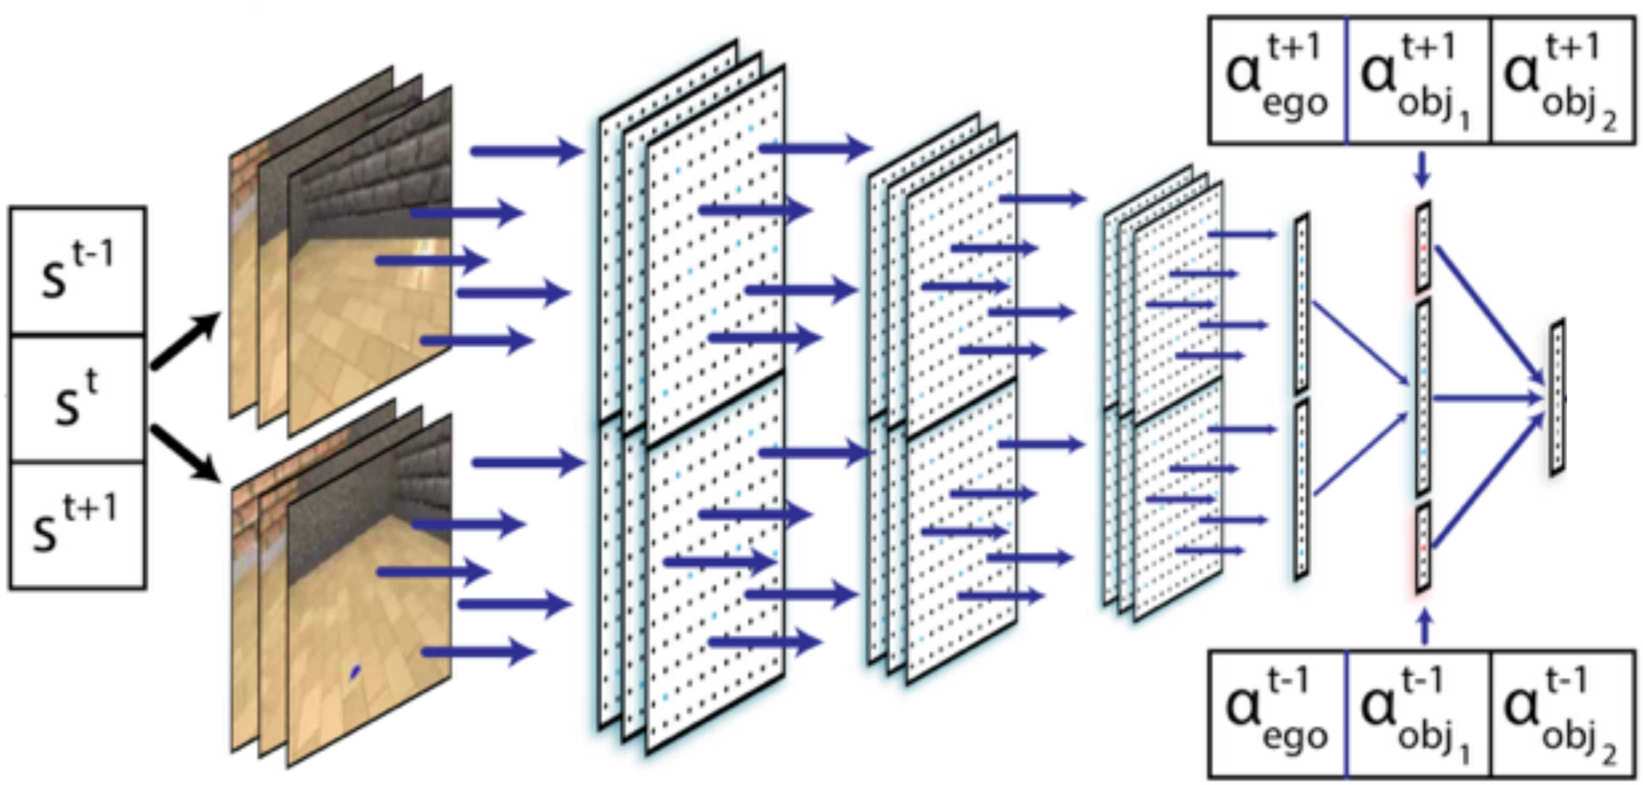

Learning to Play — World-Model

*Goal of **world-model**:*
"Post-dict" the action taken given past and future states and actions

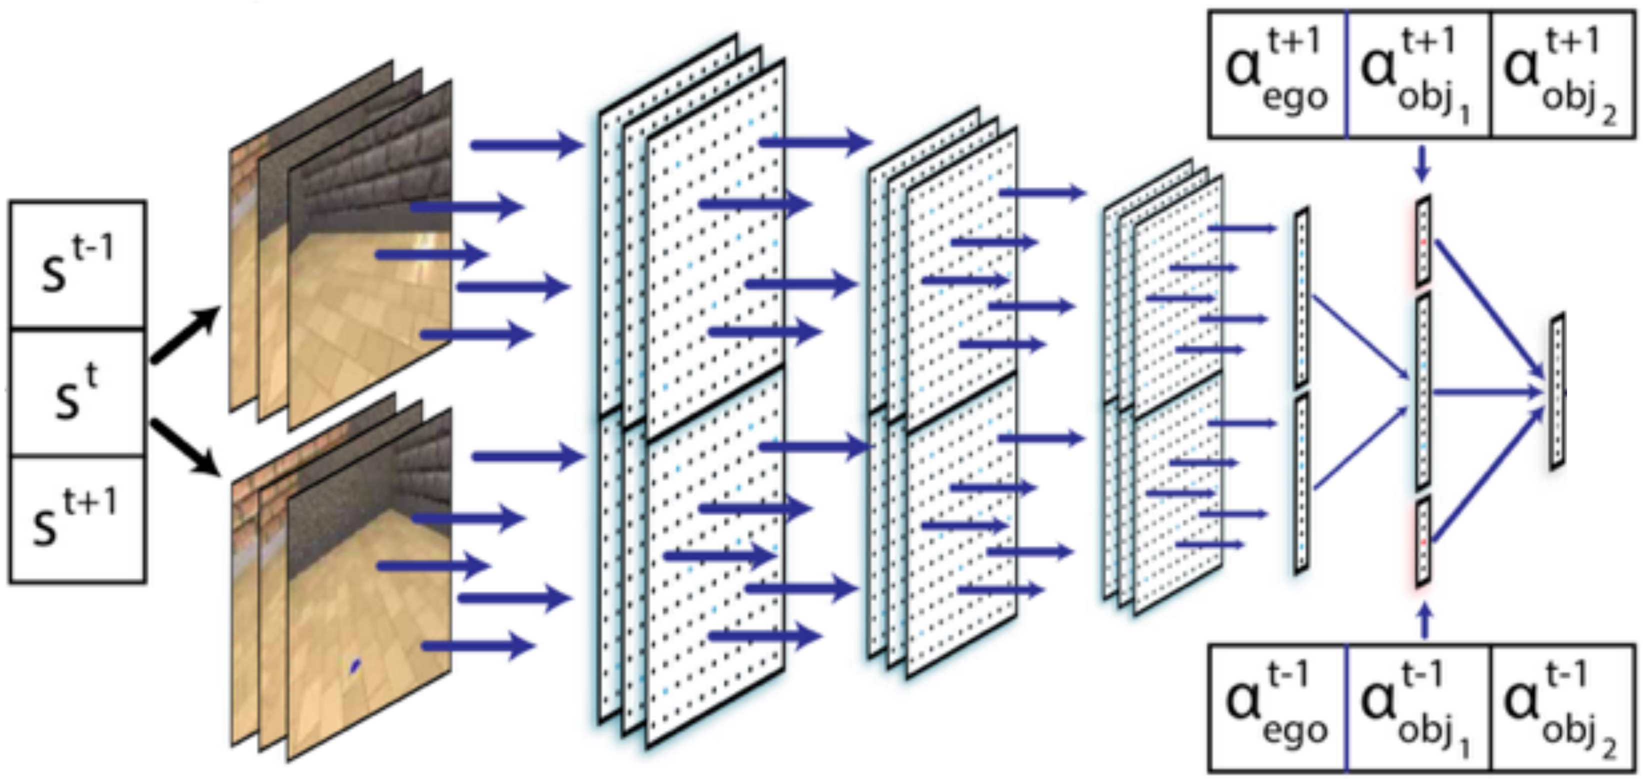

ID

Inverse Dynamics

Learning to Play — Self-Model

*Goal of **self-model**: Predict errors ("loss") of World-Model*

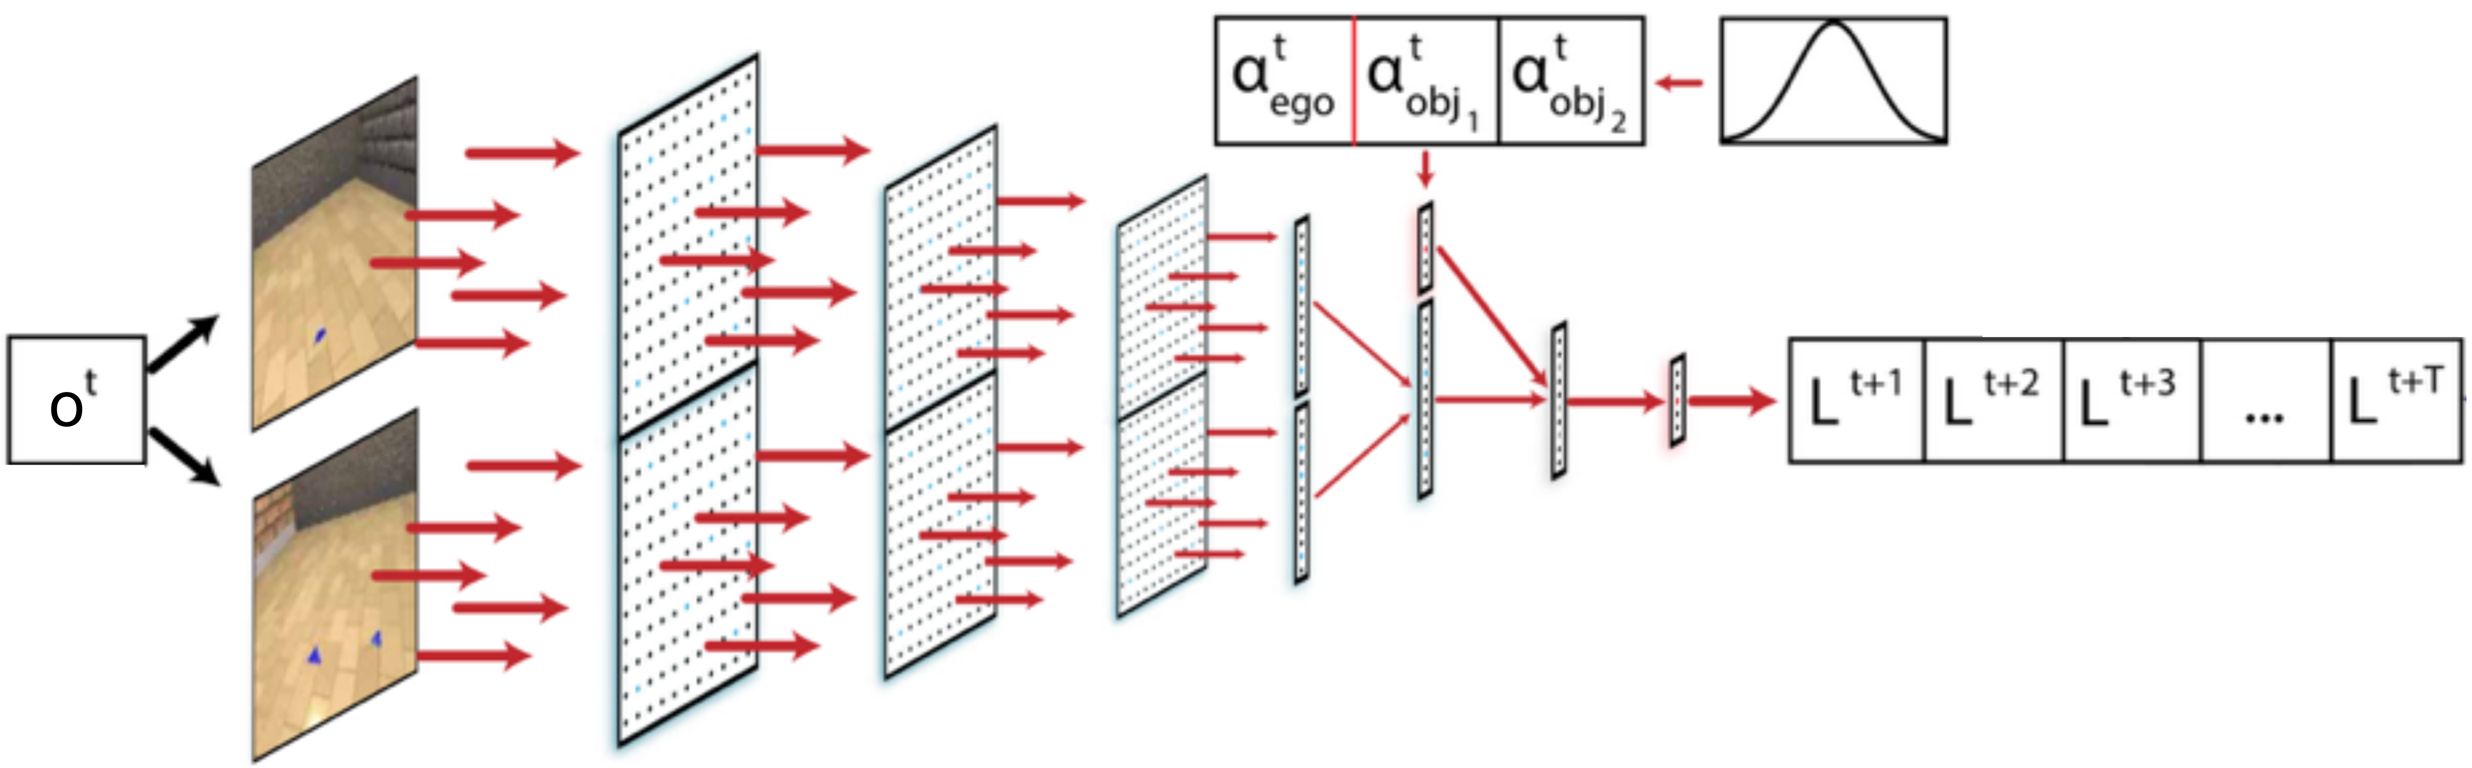

Learning to Play — Self-Model

*Goal of **self-model**: Predict errors ("loss") of World-Model*

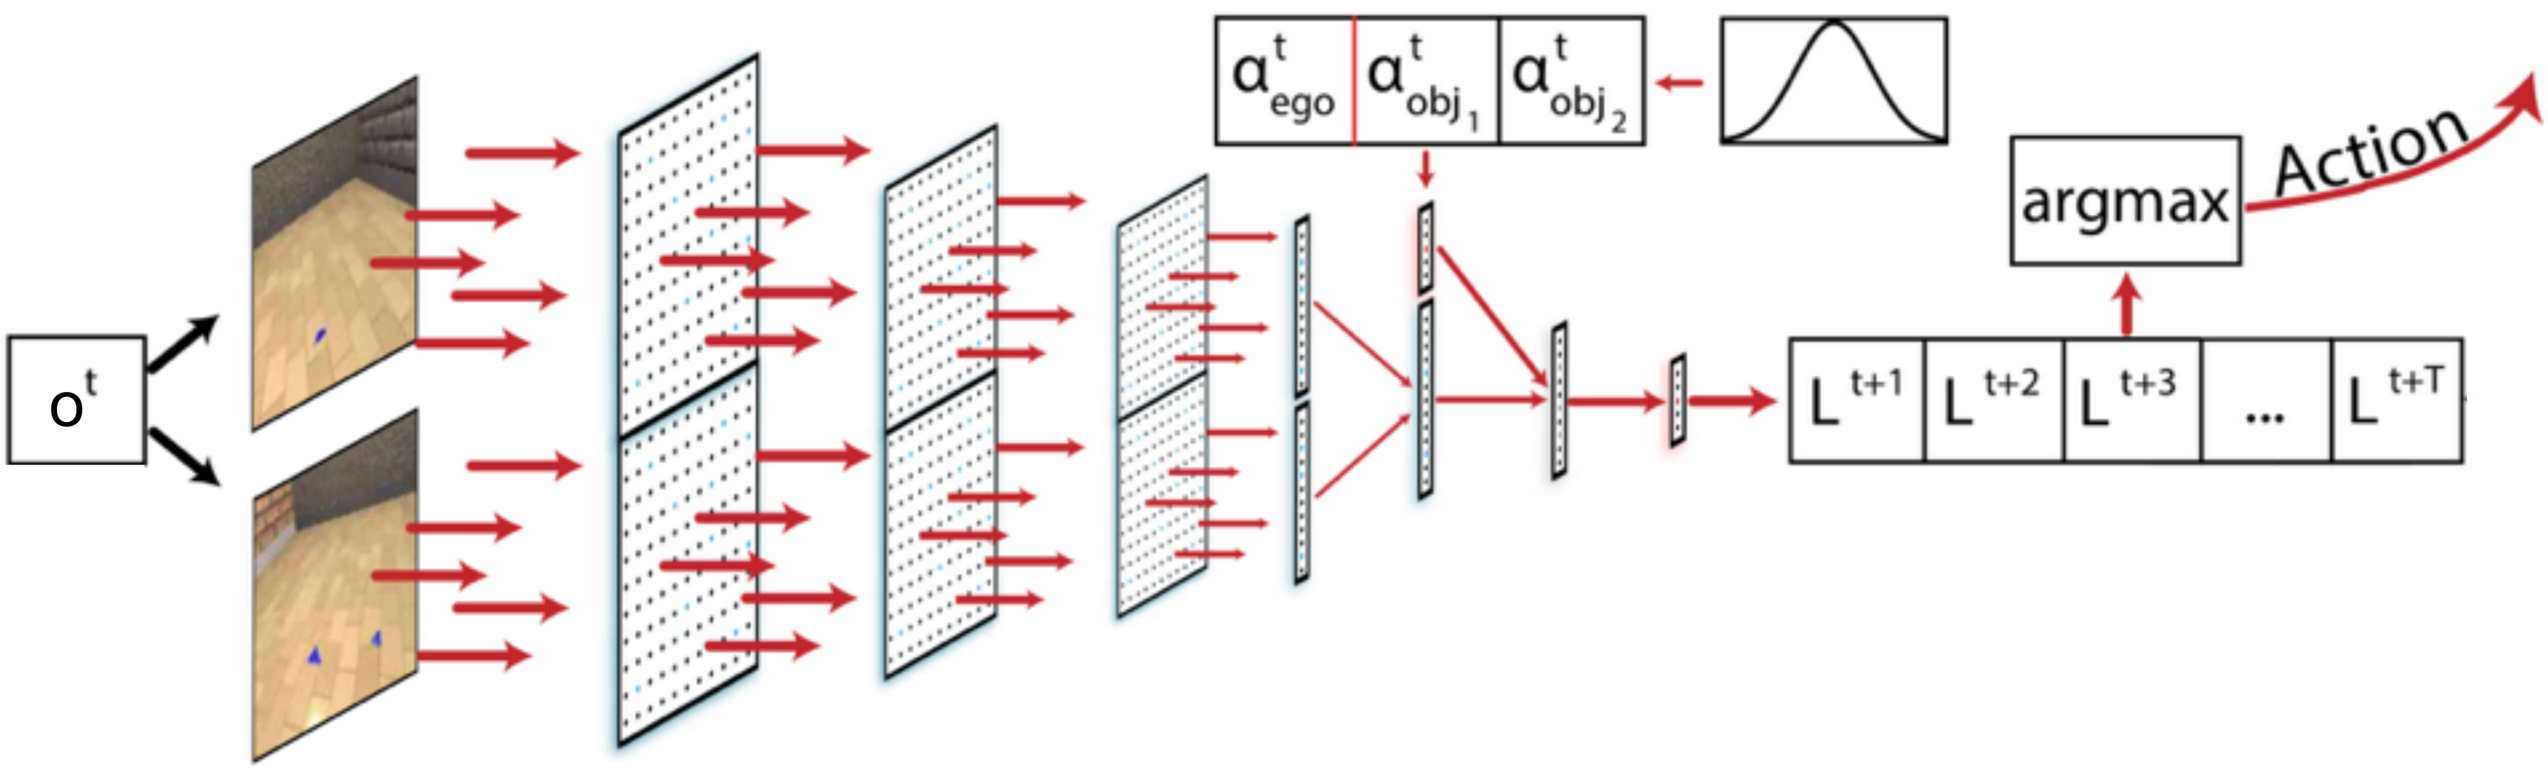

Learning to Play - Model overview

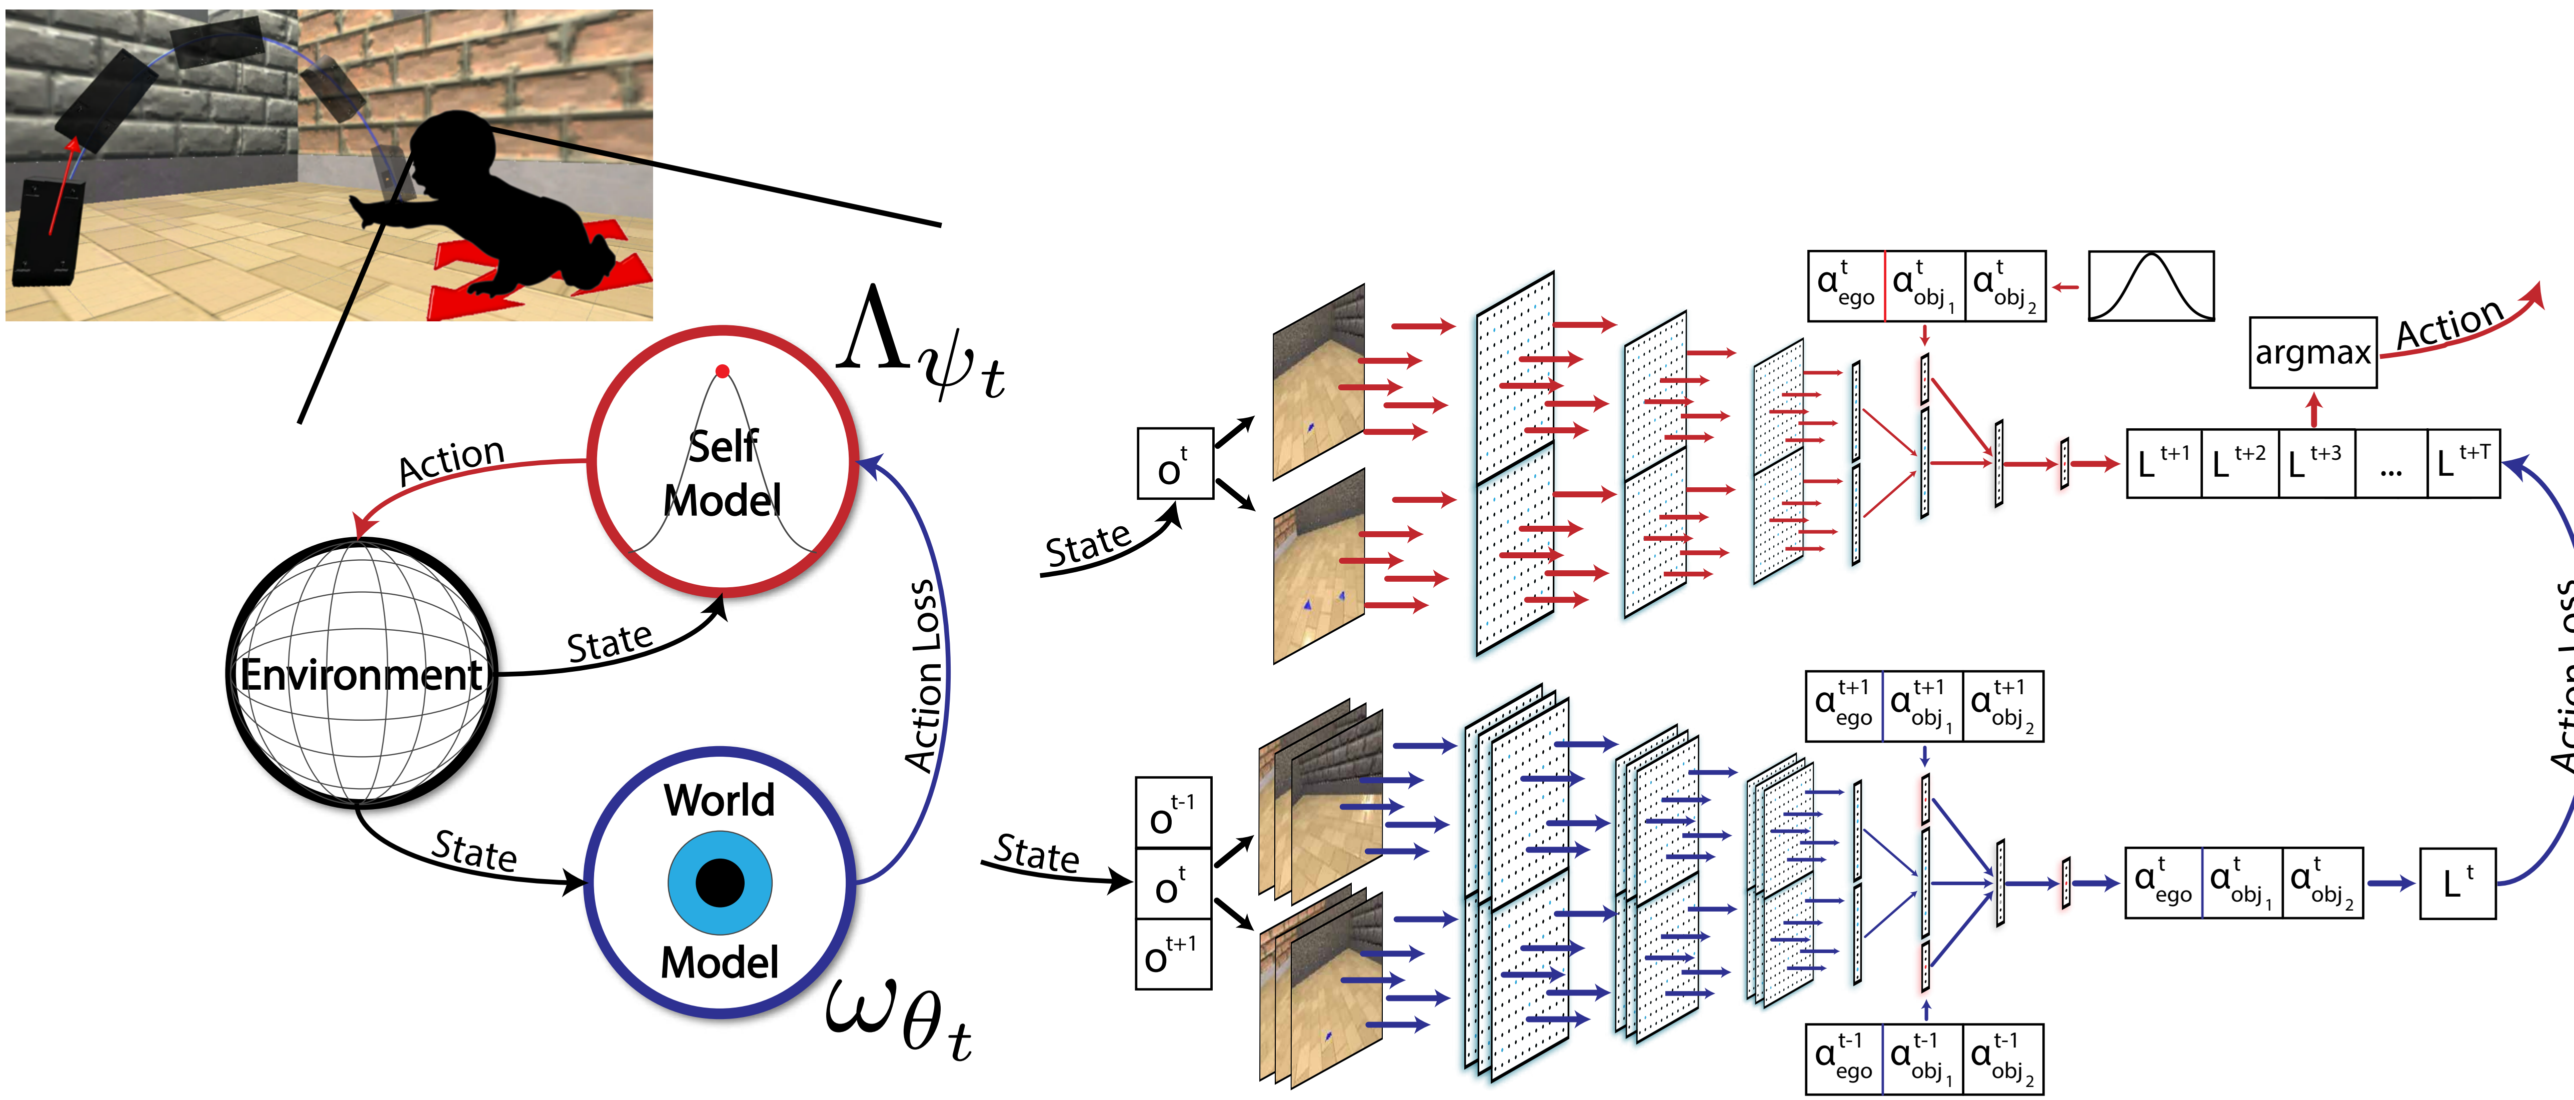

Goal of **world-model** network is to predict consequences of actions

Goal of **self-model** network is to predict errors of world-model ("self-aware")

Learning to Play — Self-Model

*Goal of **self-model**: Predict errors ("loss") of World-Model*

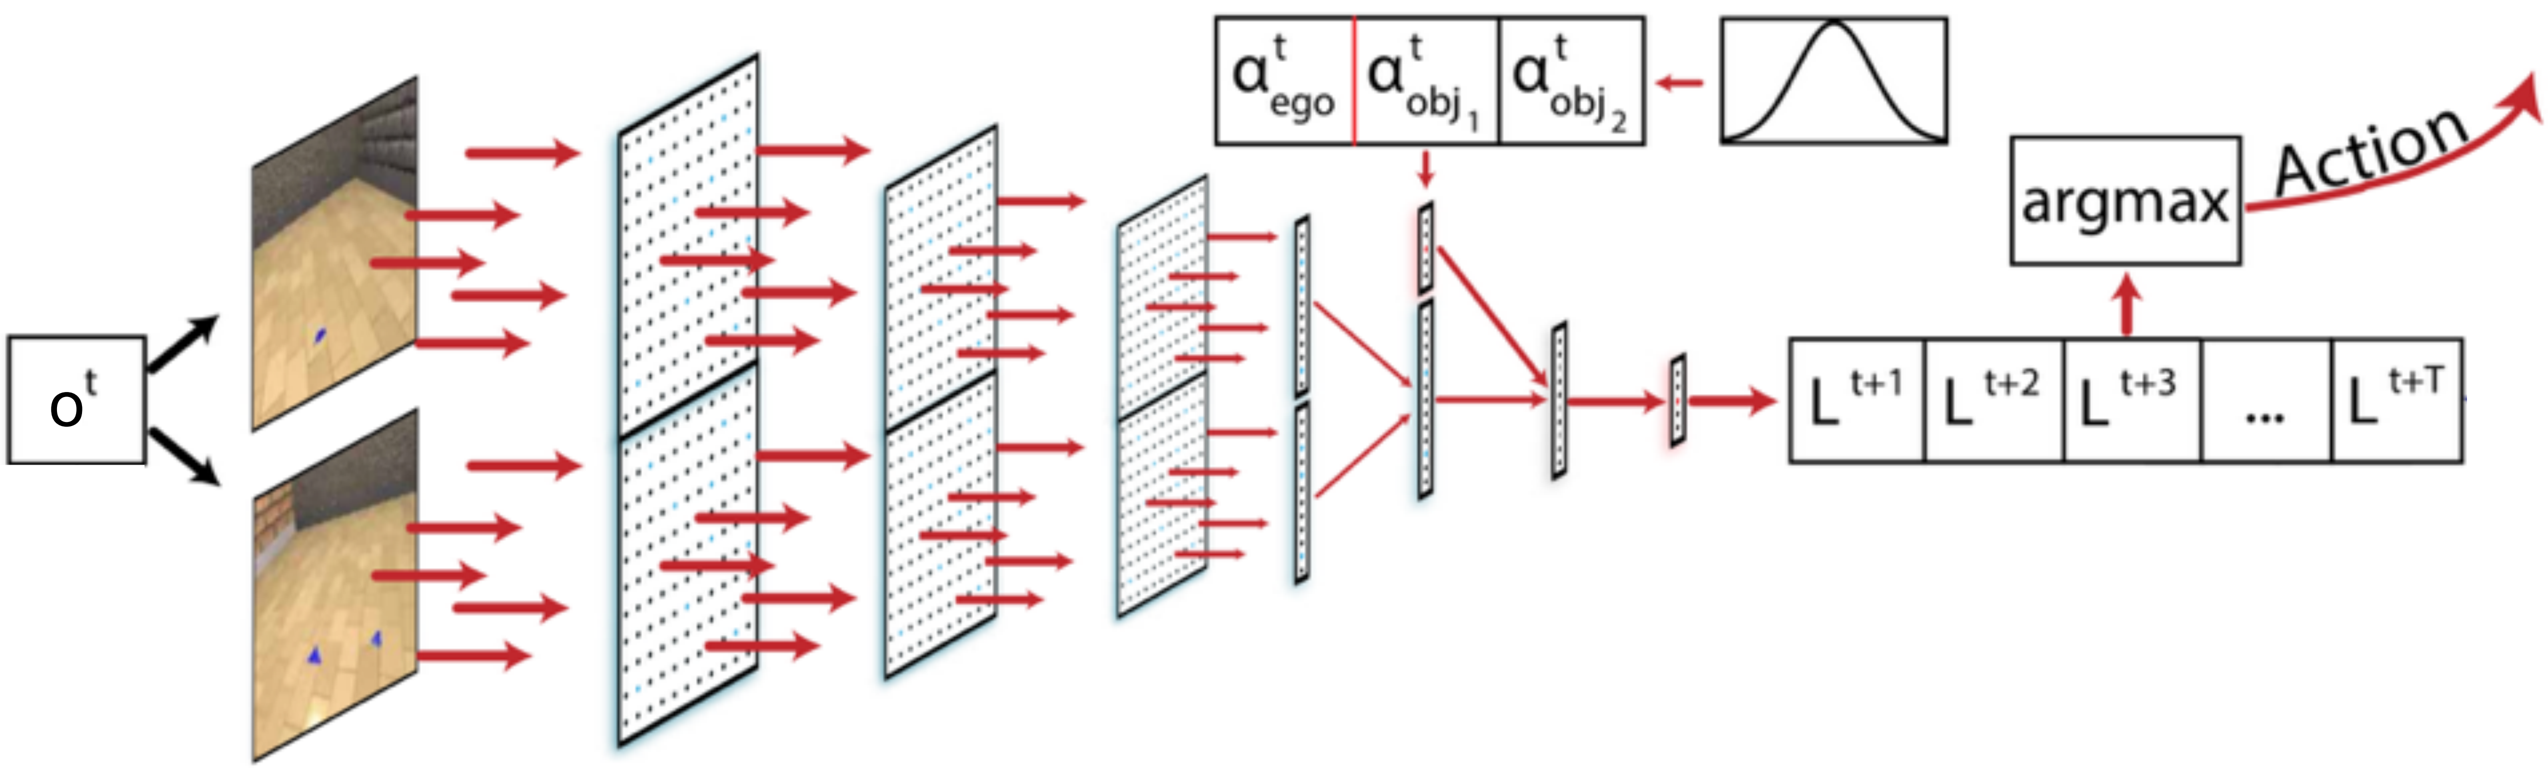

Sample 1000x actions and choose the one that maximizes the World-Model loss

$$\pi(a) \sim \exp(\beta\sigma_\Lambda(a))$$

Policy mechanism

Learning to Play — Adversarial Policy

Action choice: self-model is **adversarial** to world-model ("curious intrinsic motivation")

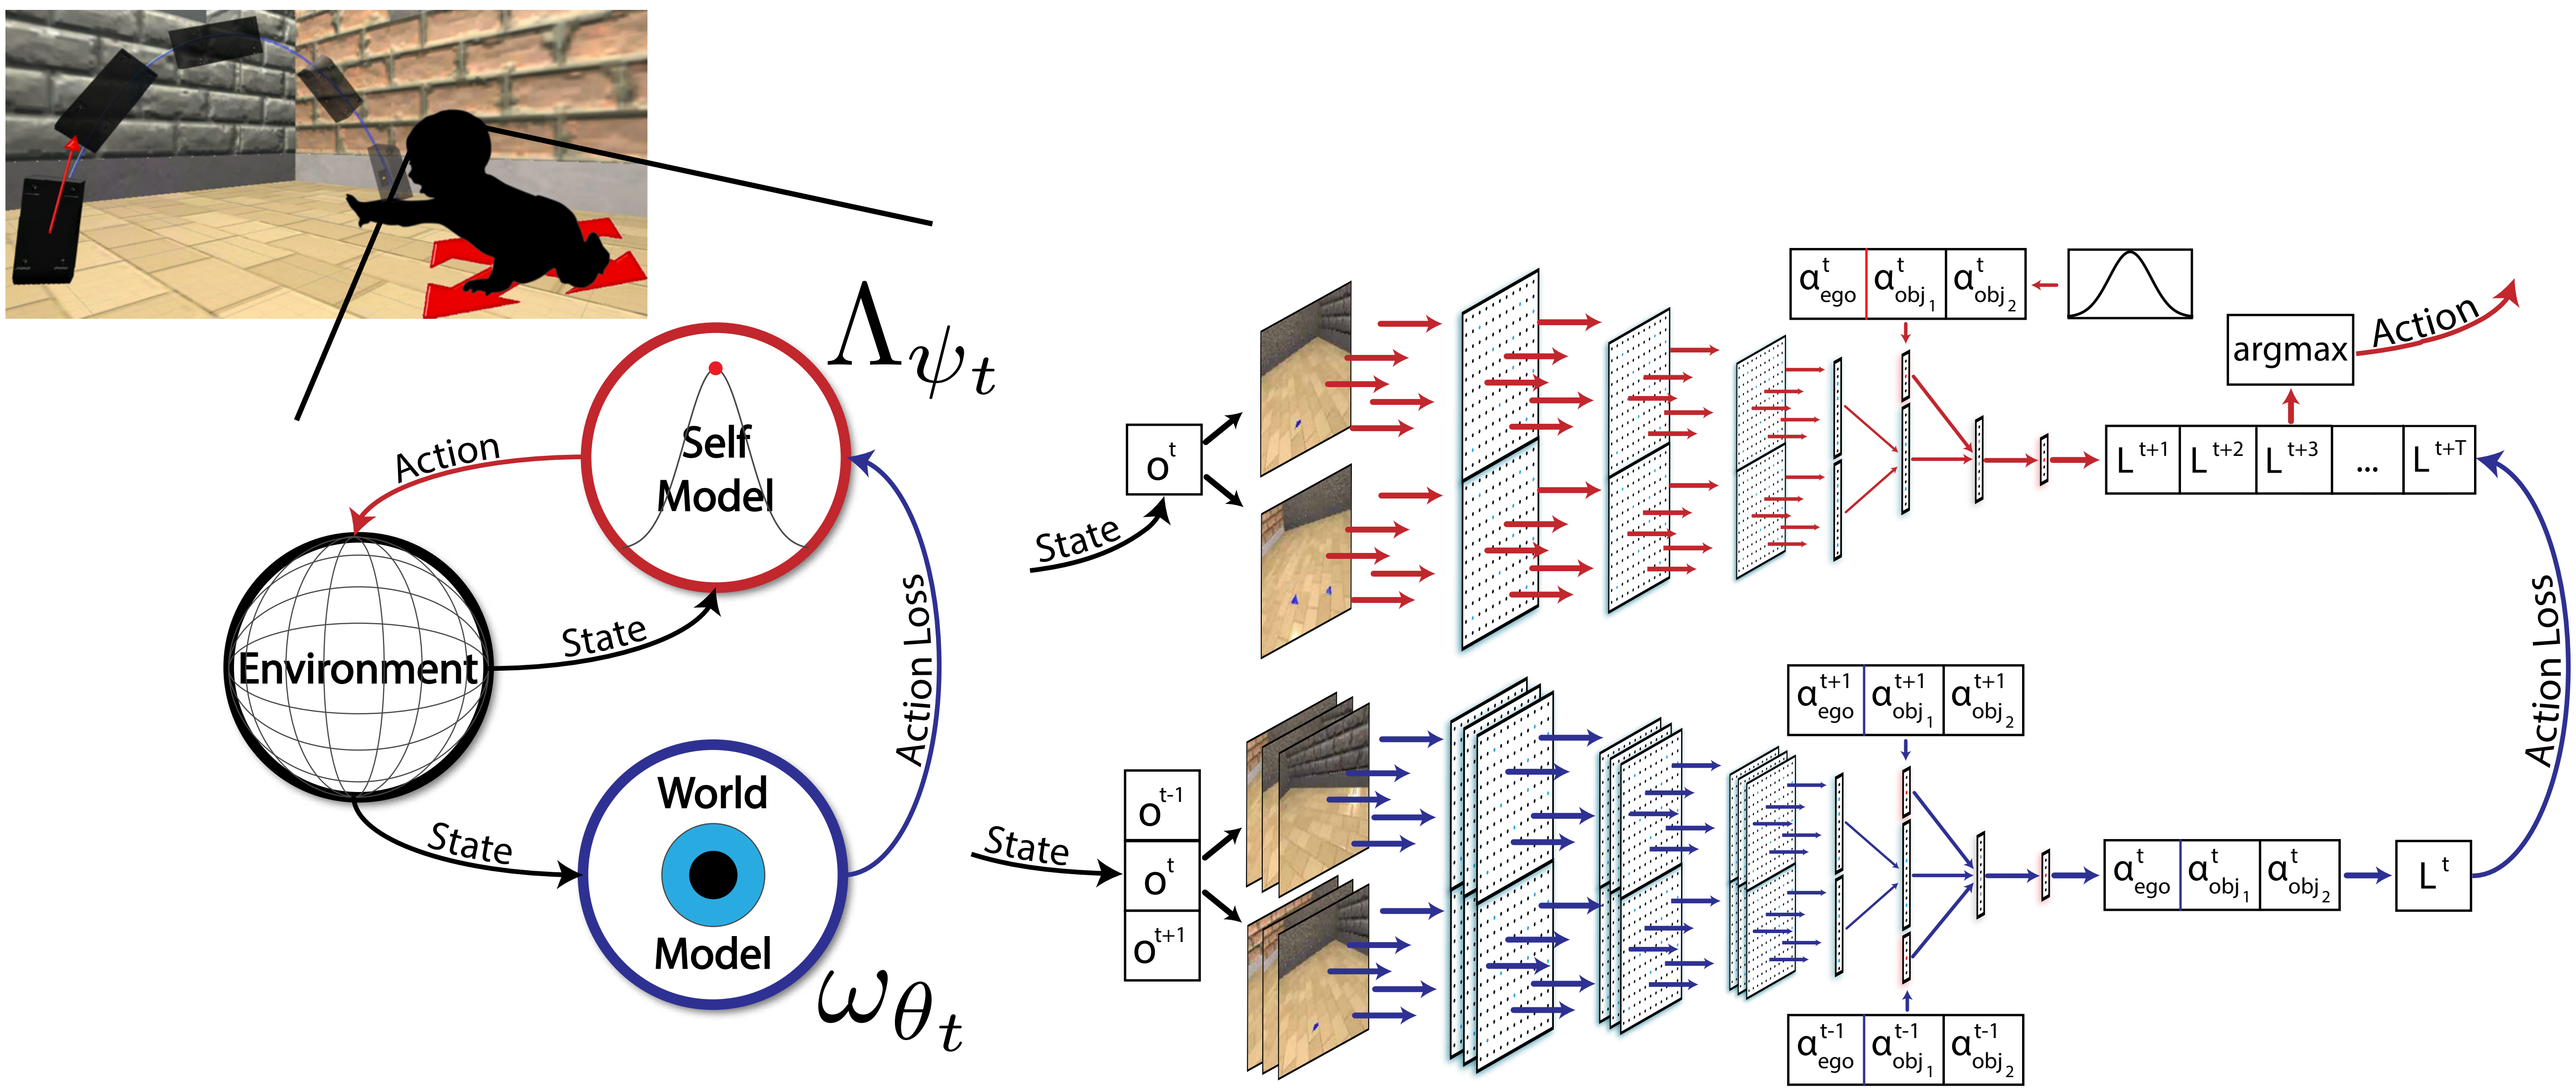

Goal of **world-model** network is to predict consequences of actions

Goal of **self-model** network is to predict errors of world-model ("self-aware")

Place agent in room with a single object.

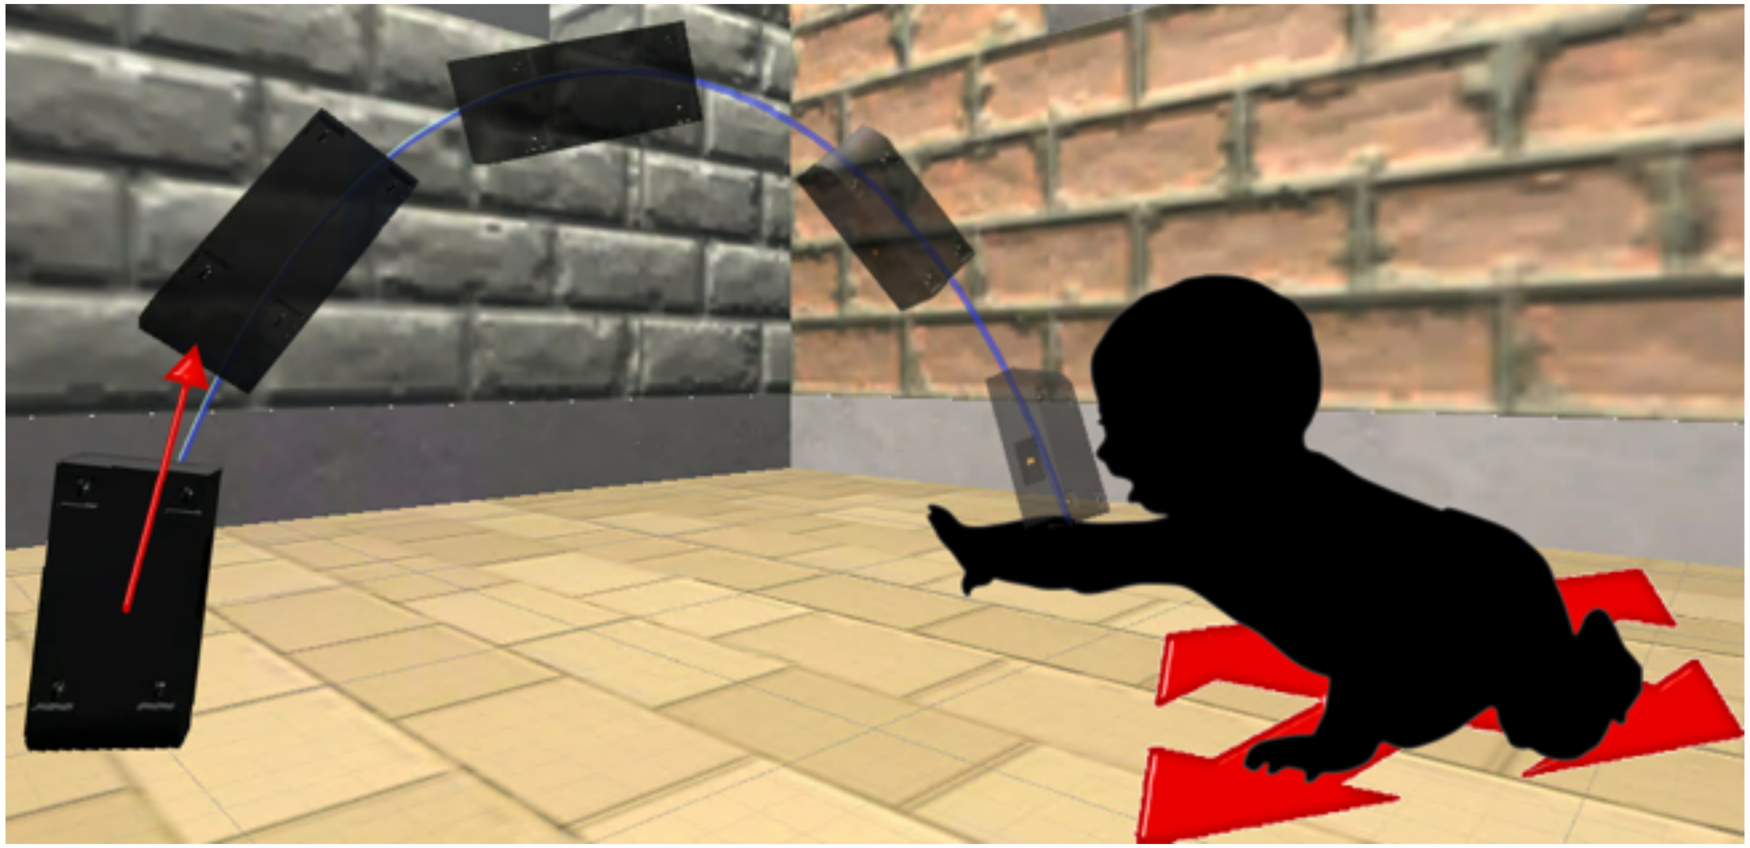

Self-supervised learning

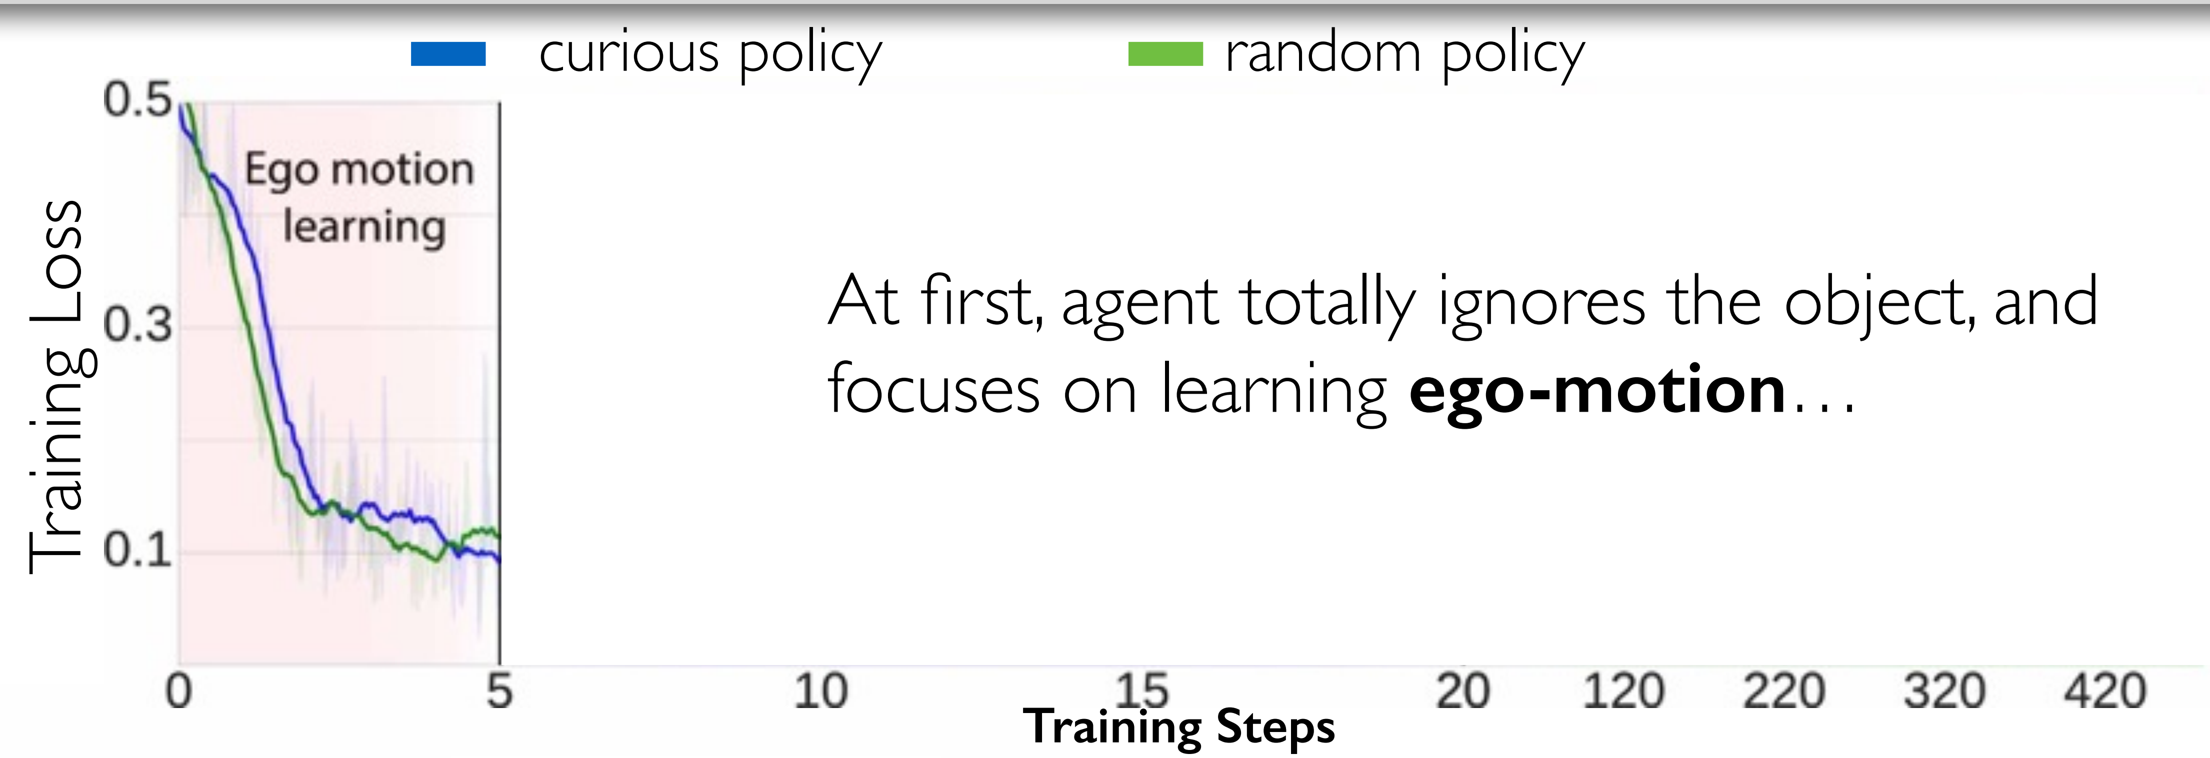

At first, agent totally ignores the object, and focuses on learning **ego-motion**…

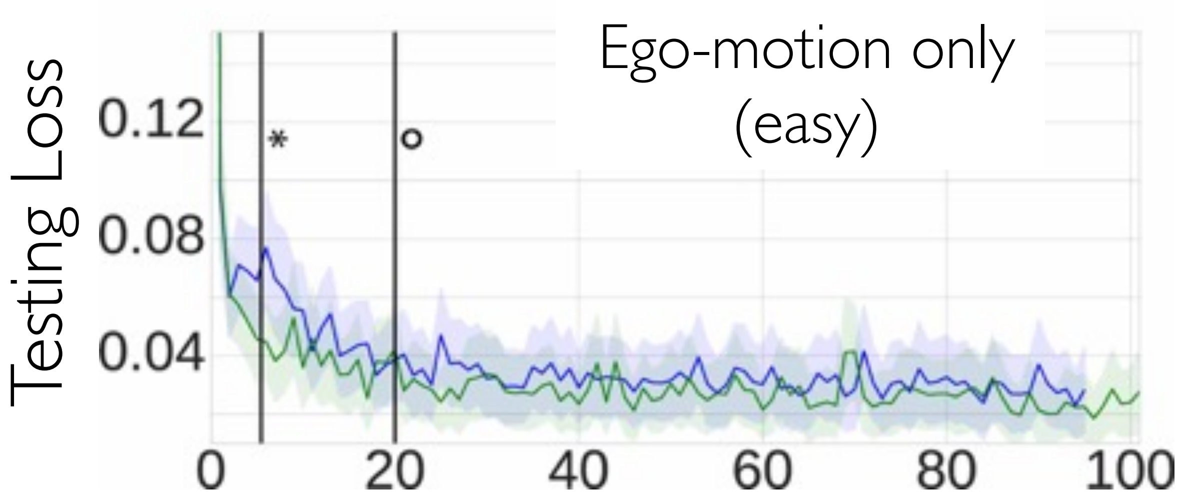

Self-supervised learning

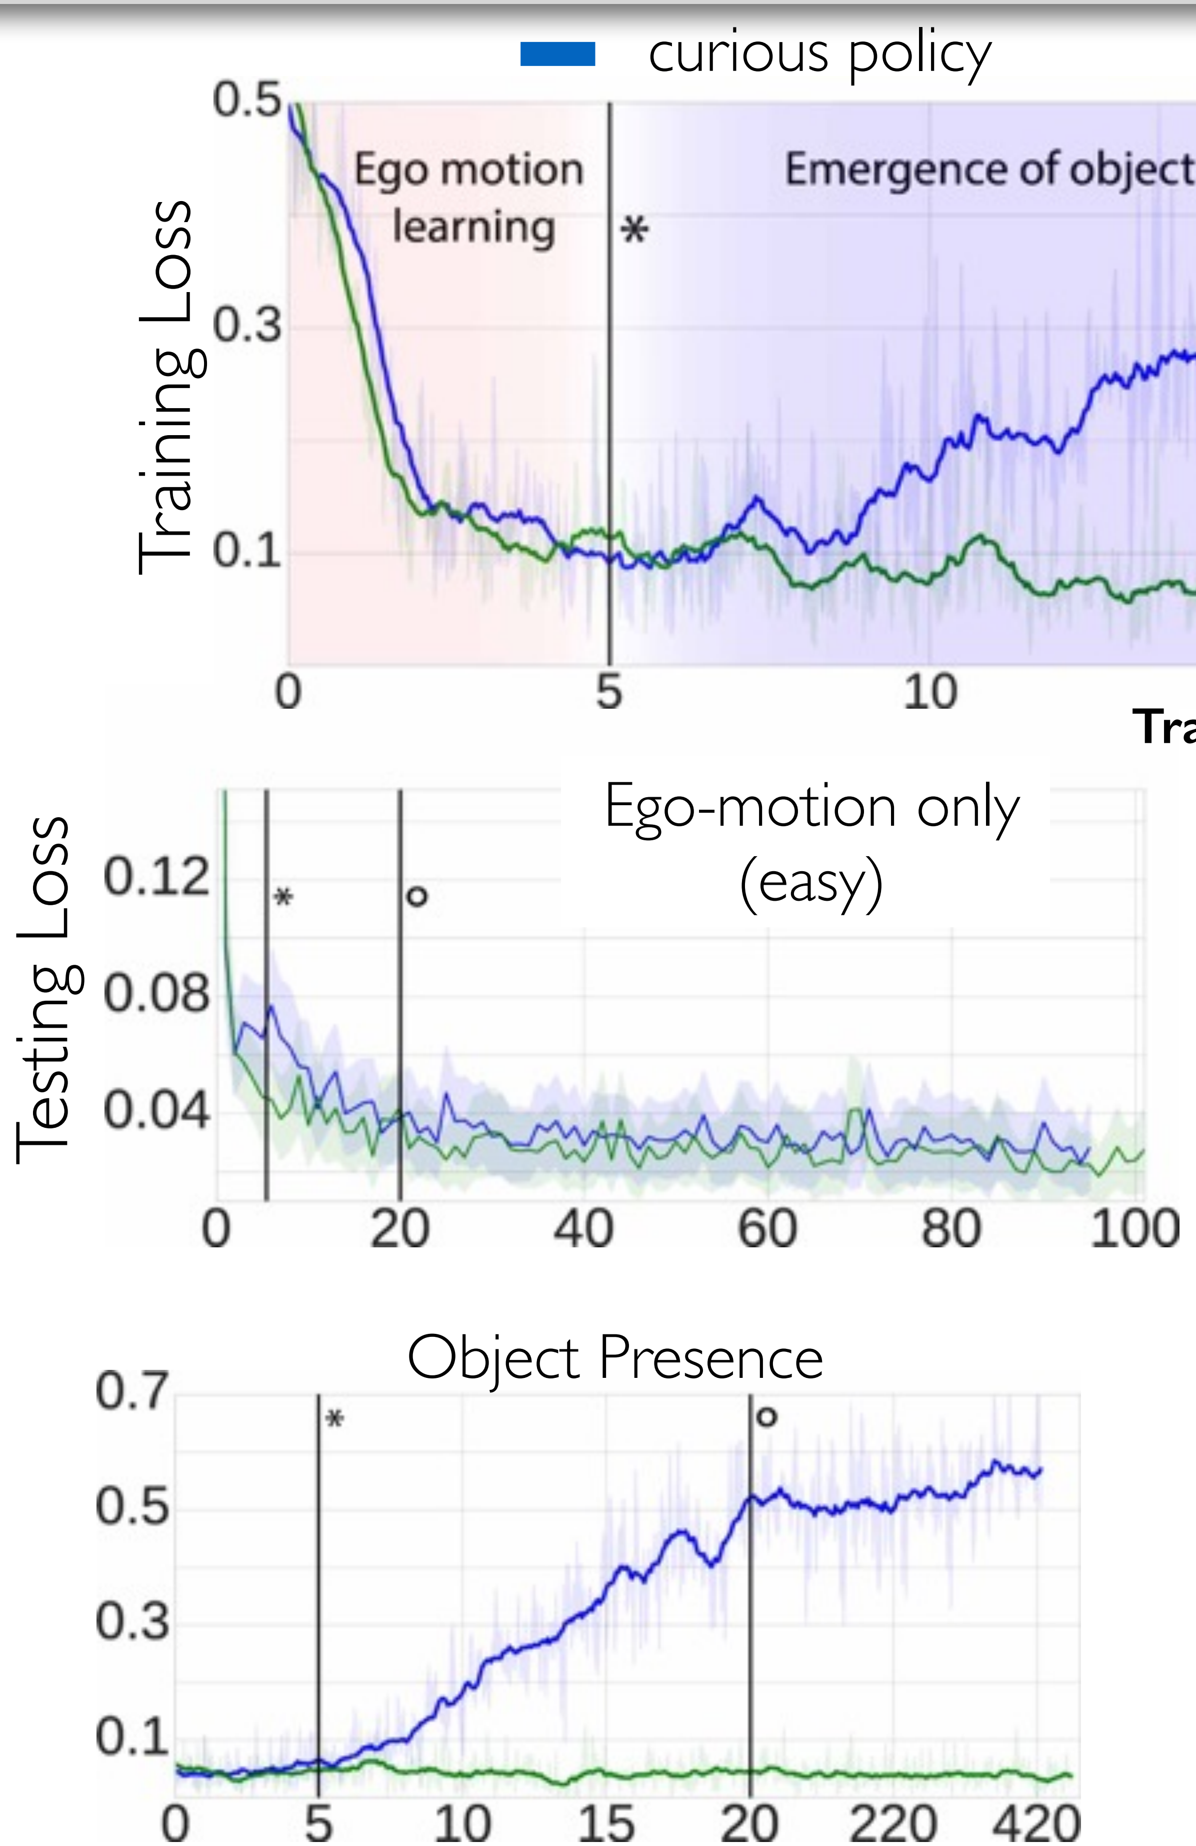

But the curious agent eventually "gets bored" of ego-motion prediction and starts to focus on the object!

Self-supervised learning

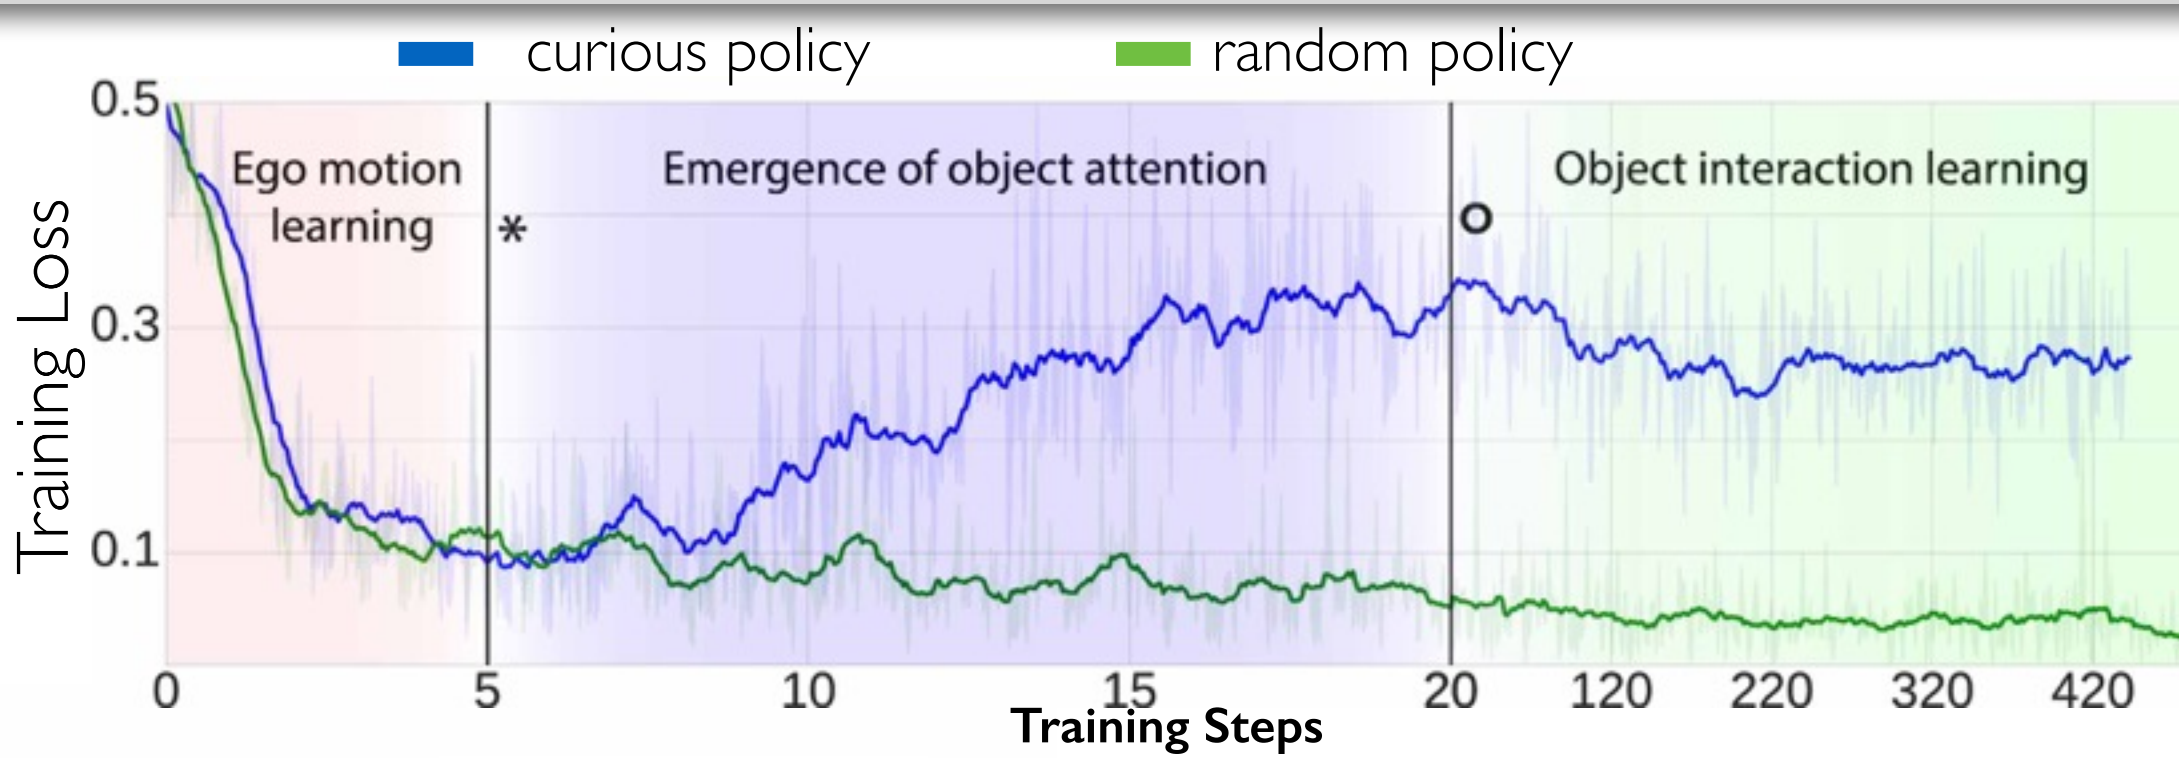

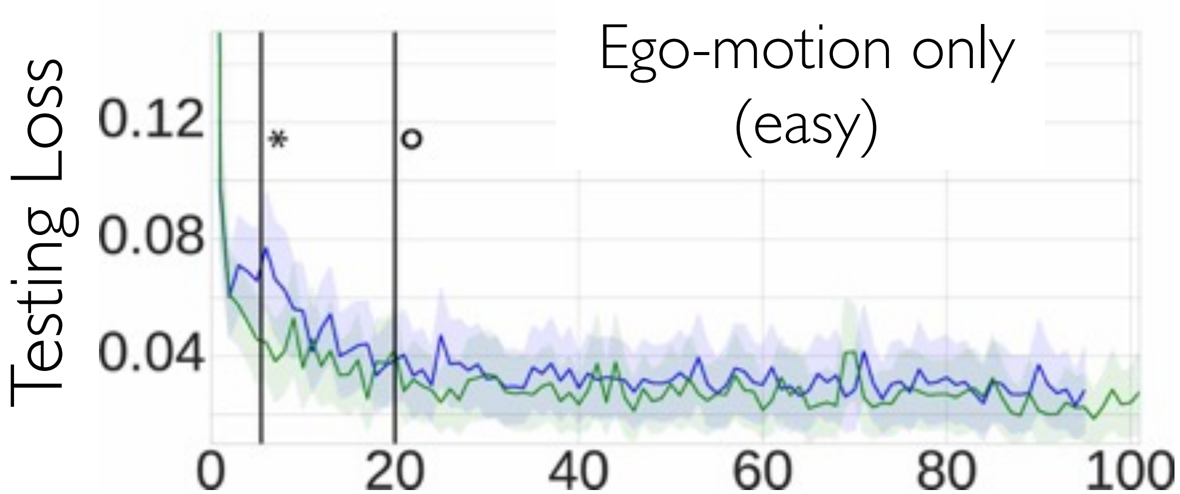

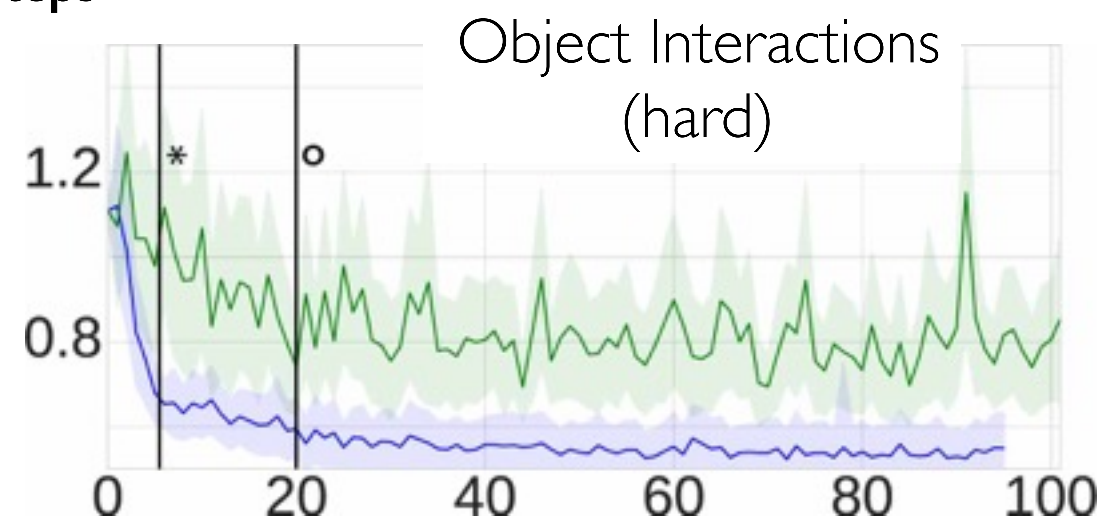

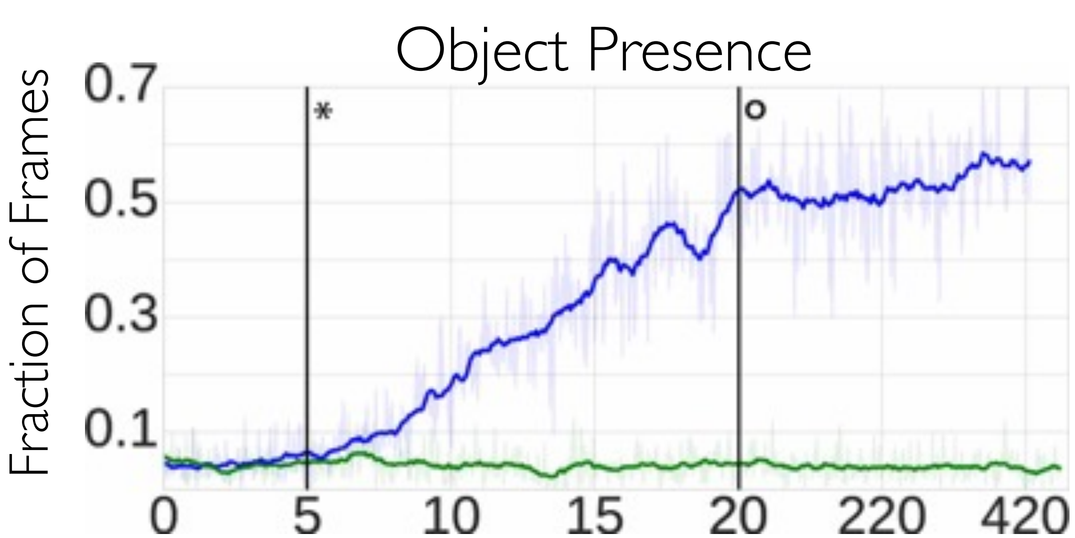

Emergence of:

(a) ego-motion understanding

(b) object attention

(c) improved world-model

Self-supervised learning

Simple navigation and planning behavior emerges …

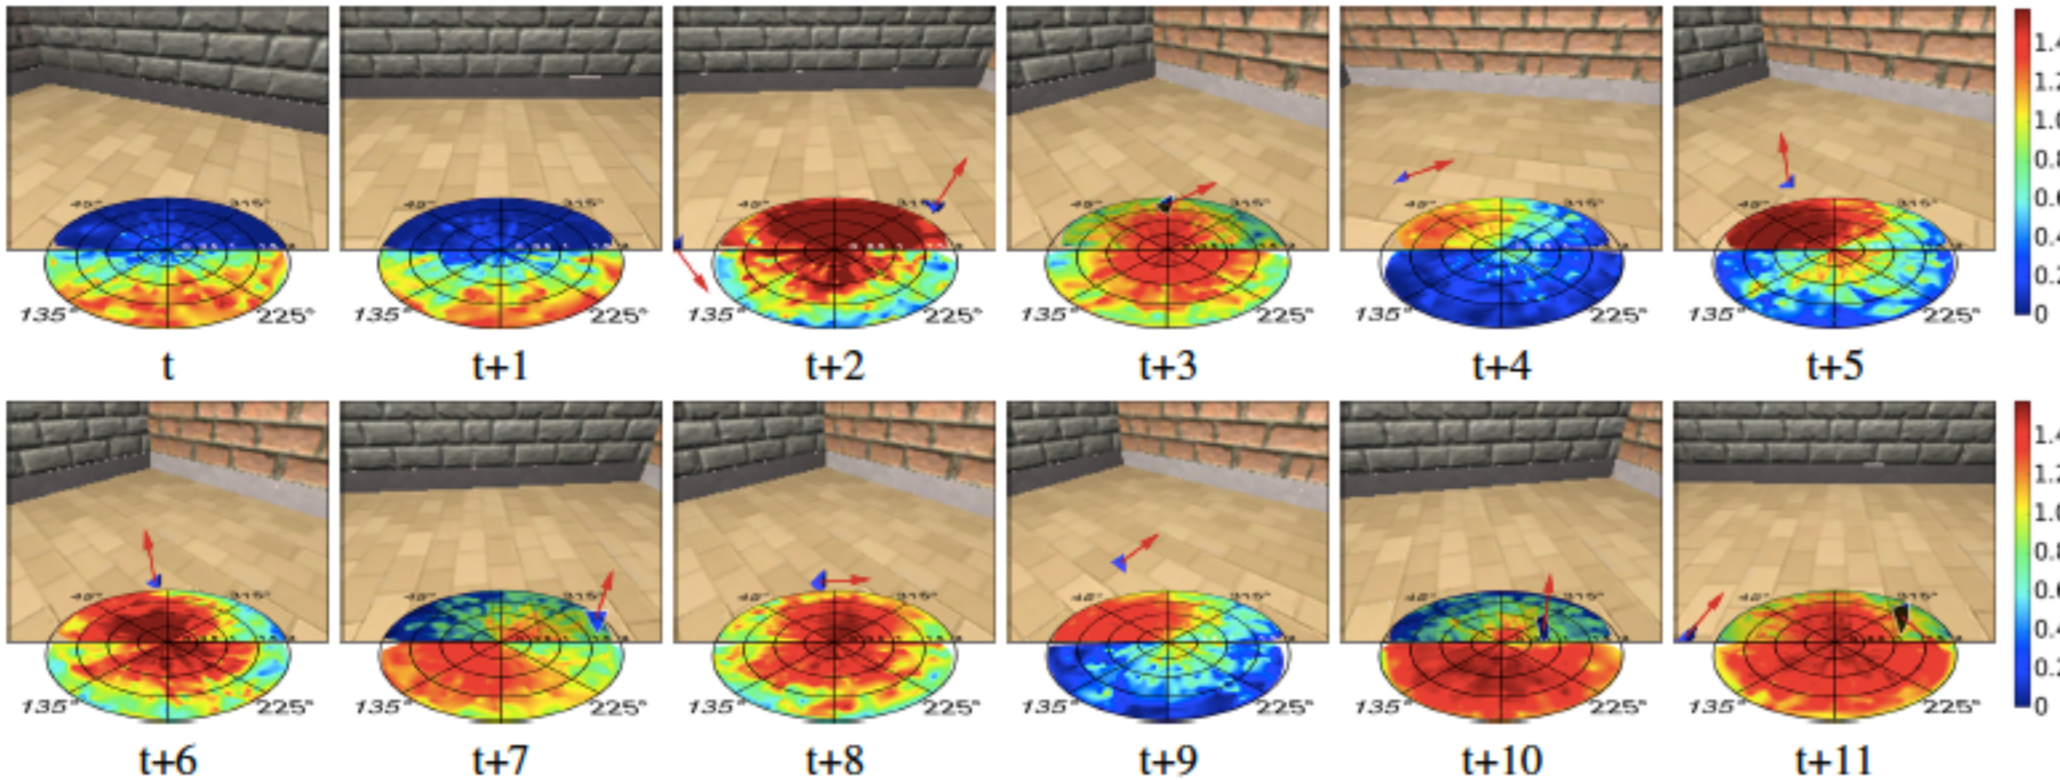

If an object is not in view, the agent turns to find one…

Self-supervised learning

Simple navigation and planning behavior emerges …

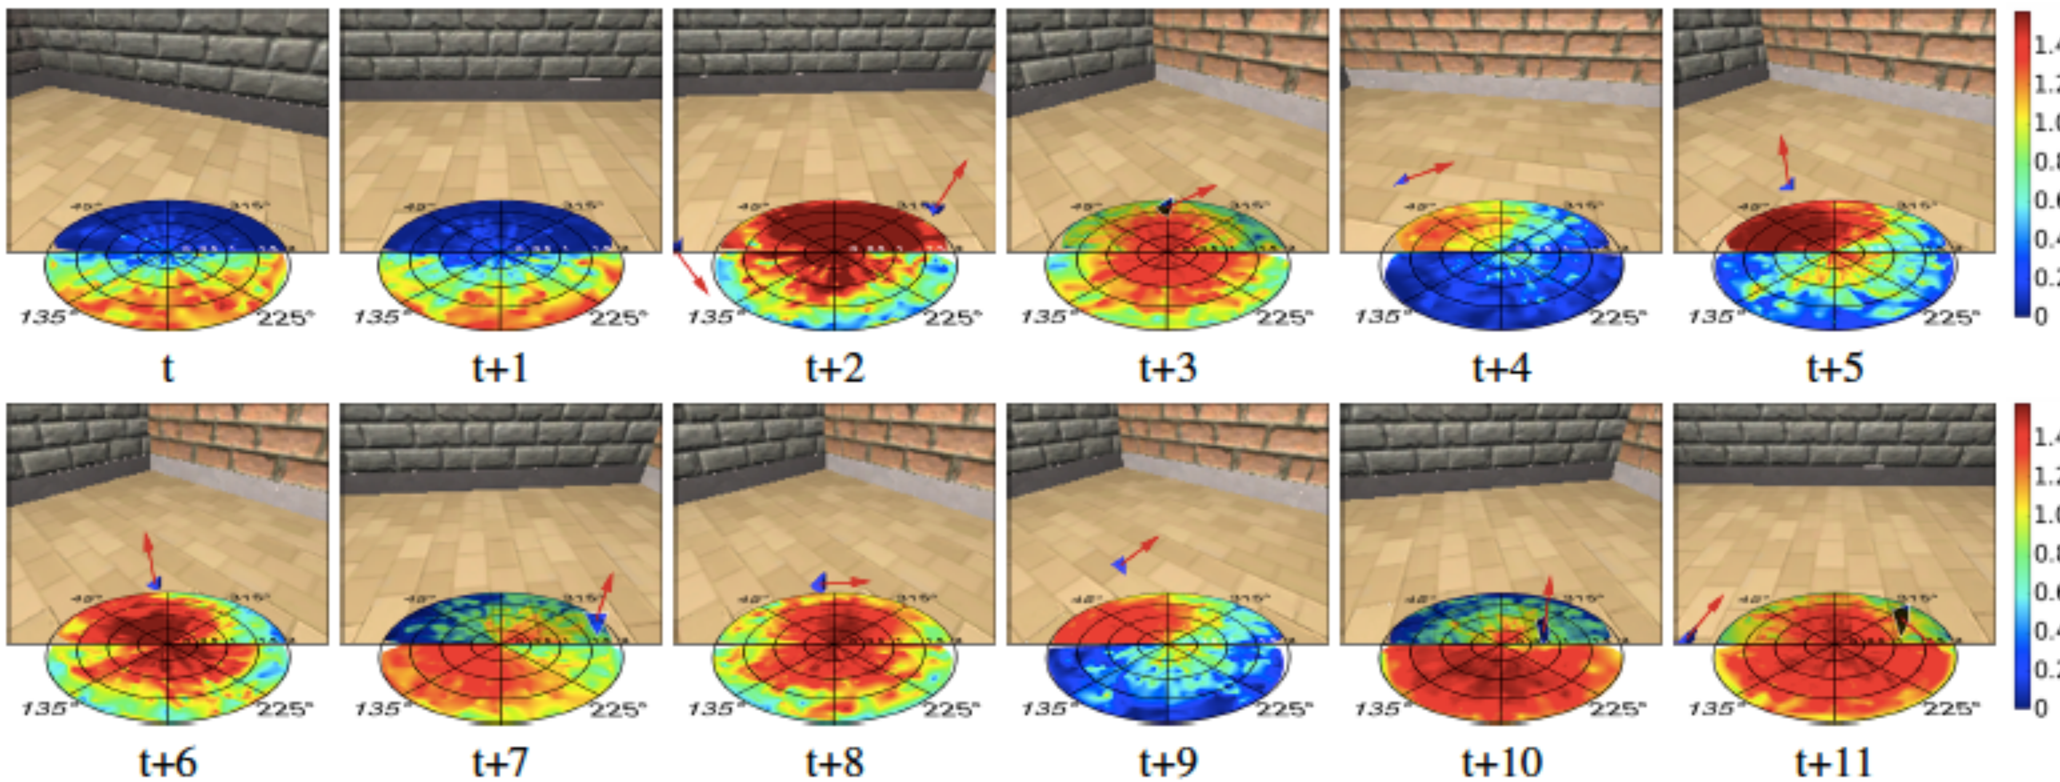

If an object is not in view, the agent turns to find one…

… if an object is too far to touch, the agent moves toward one.

Self-supervised learning

Simple navigation and planning behavior emerges …

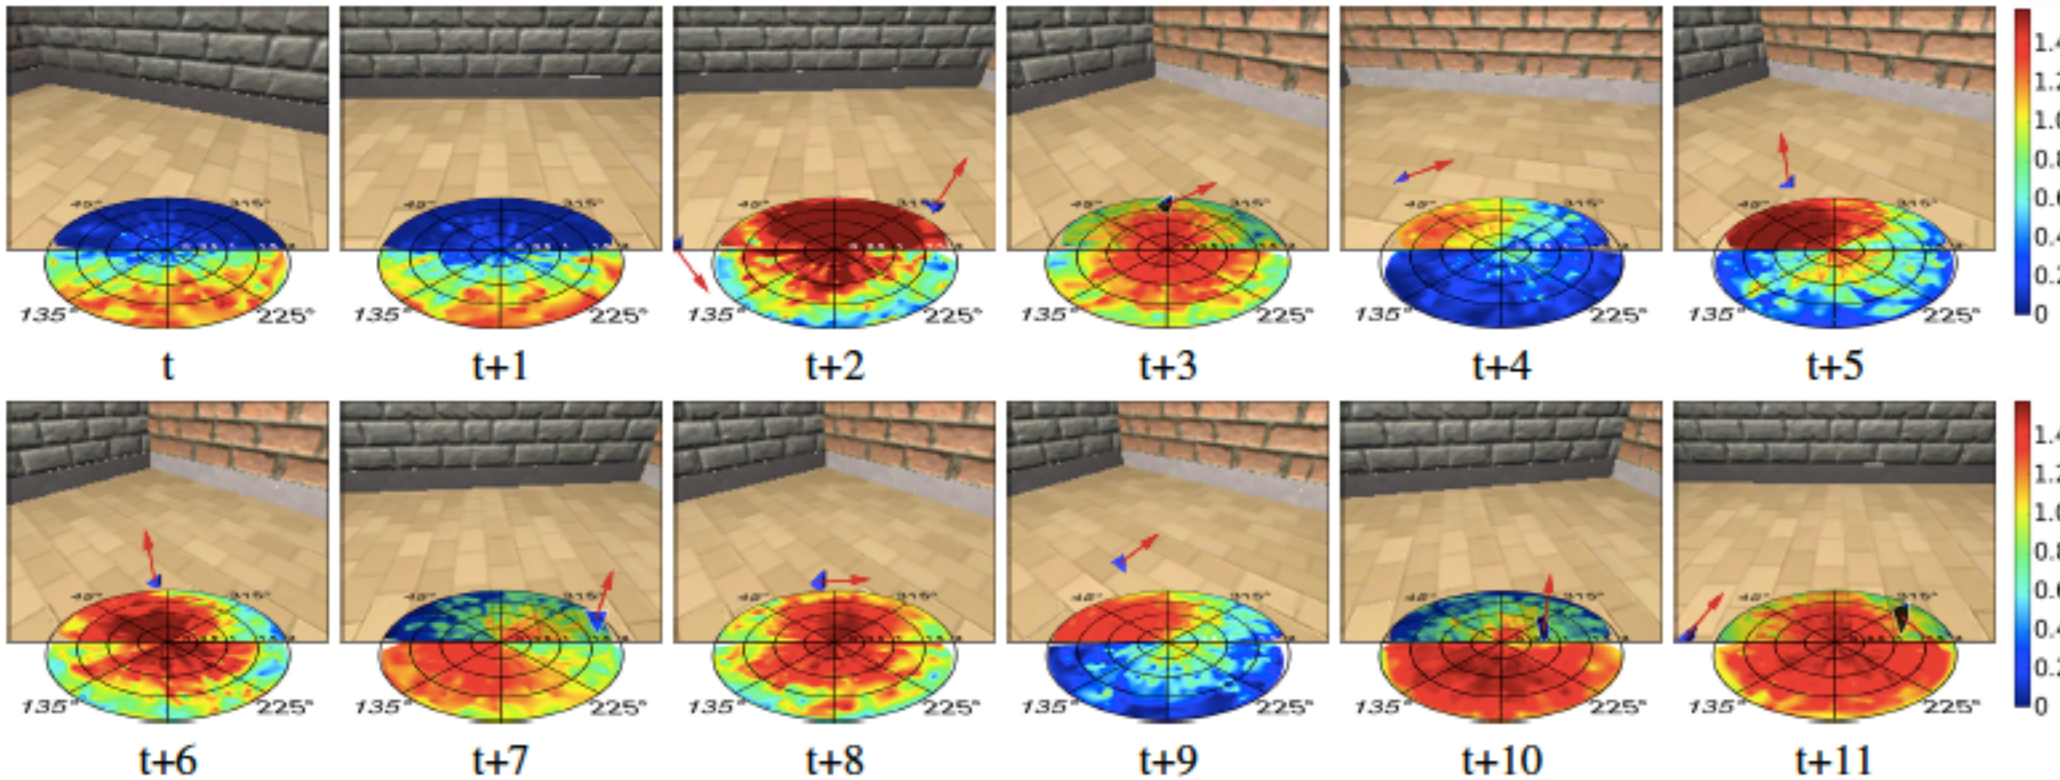

If an object is not in view, the agent turns to find one…

… if an object is too far to touch, the agent moves toward one.

… and once the agent is close to an object, it stays close and interacts with it.

Self-supervised learning

Simple navigation and planning behavior emerges …

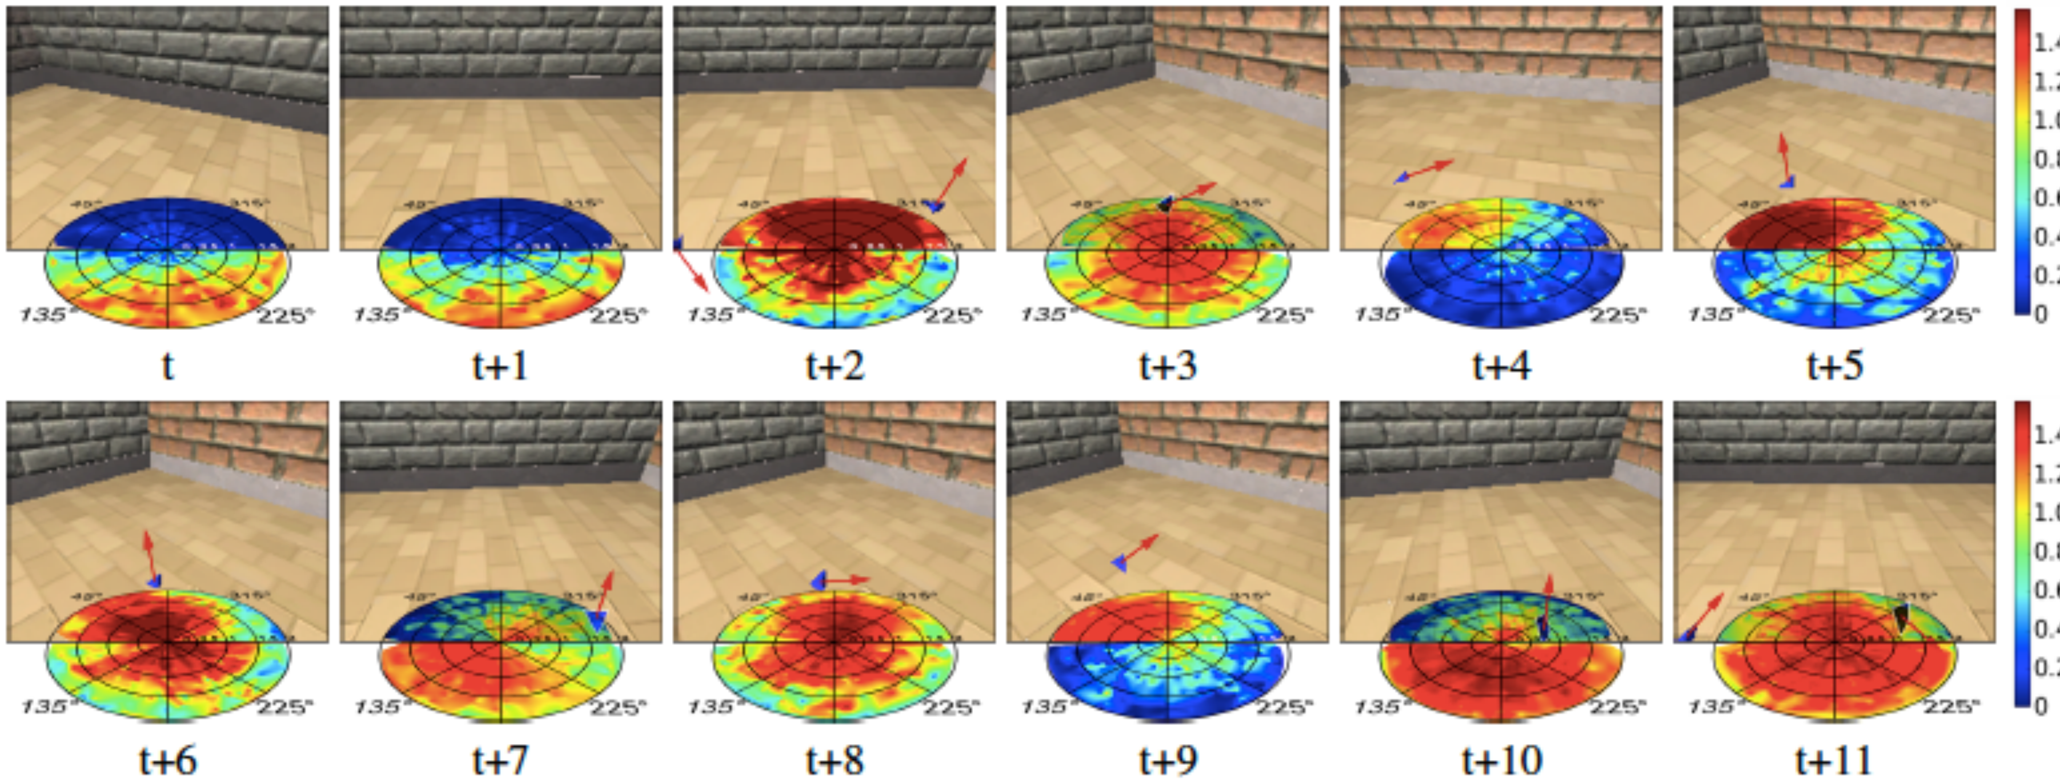

Moreover, substantially improved transfer learning accuracy:

(a) object detection (present or not): ~8% vs ~40% accuracy

(b) object position: ~6px vs ~4px error

(c) object recognition (among 16 geometries): ~12% vs ~30% accuracy

Self-supervised learning

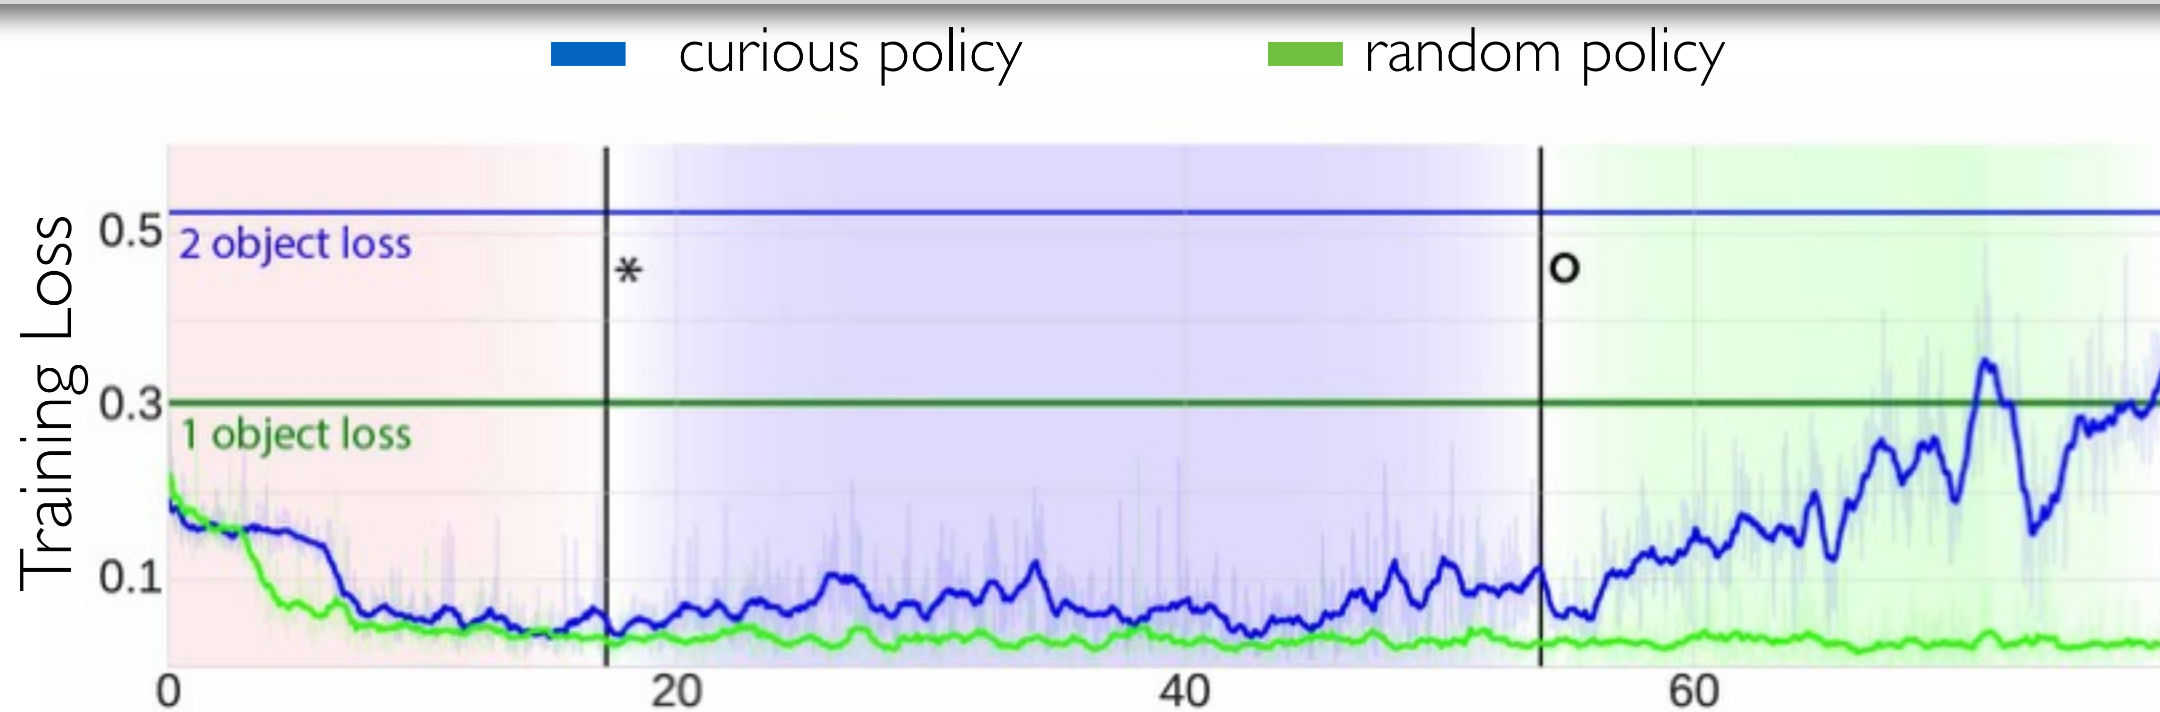

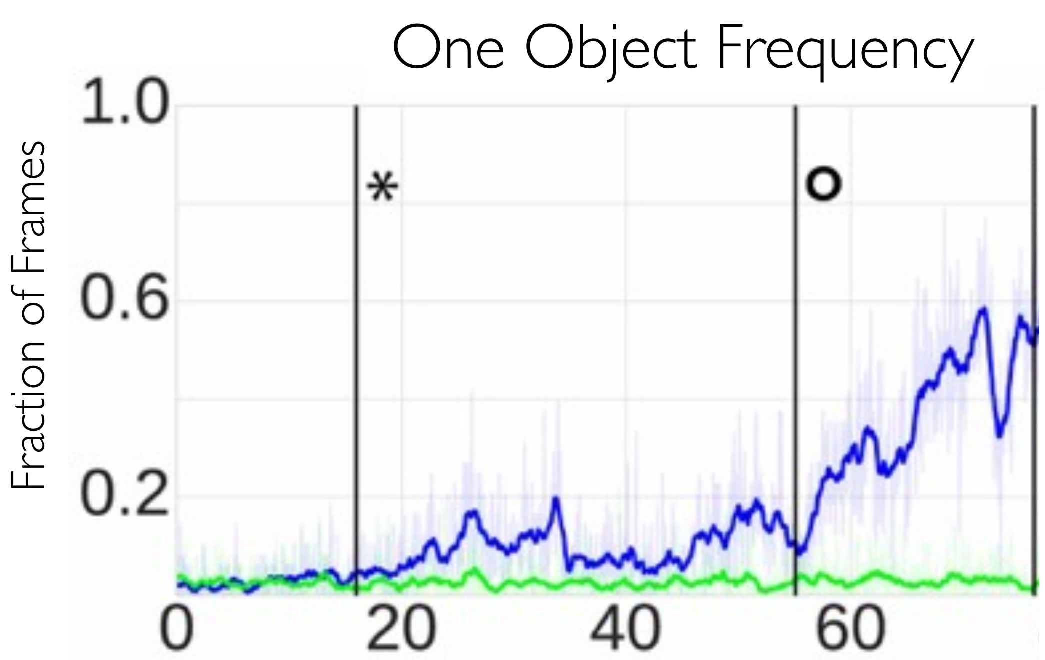

When multiple objects are present, the agent at first recapitulates its behavior with a single object …

Self-supervised learning

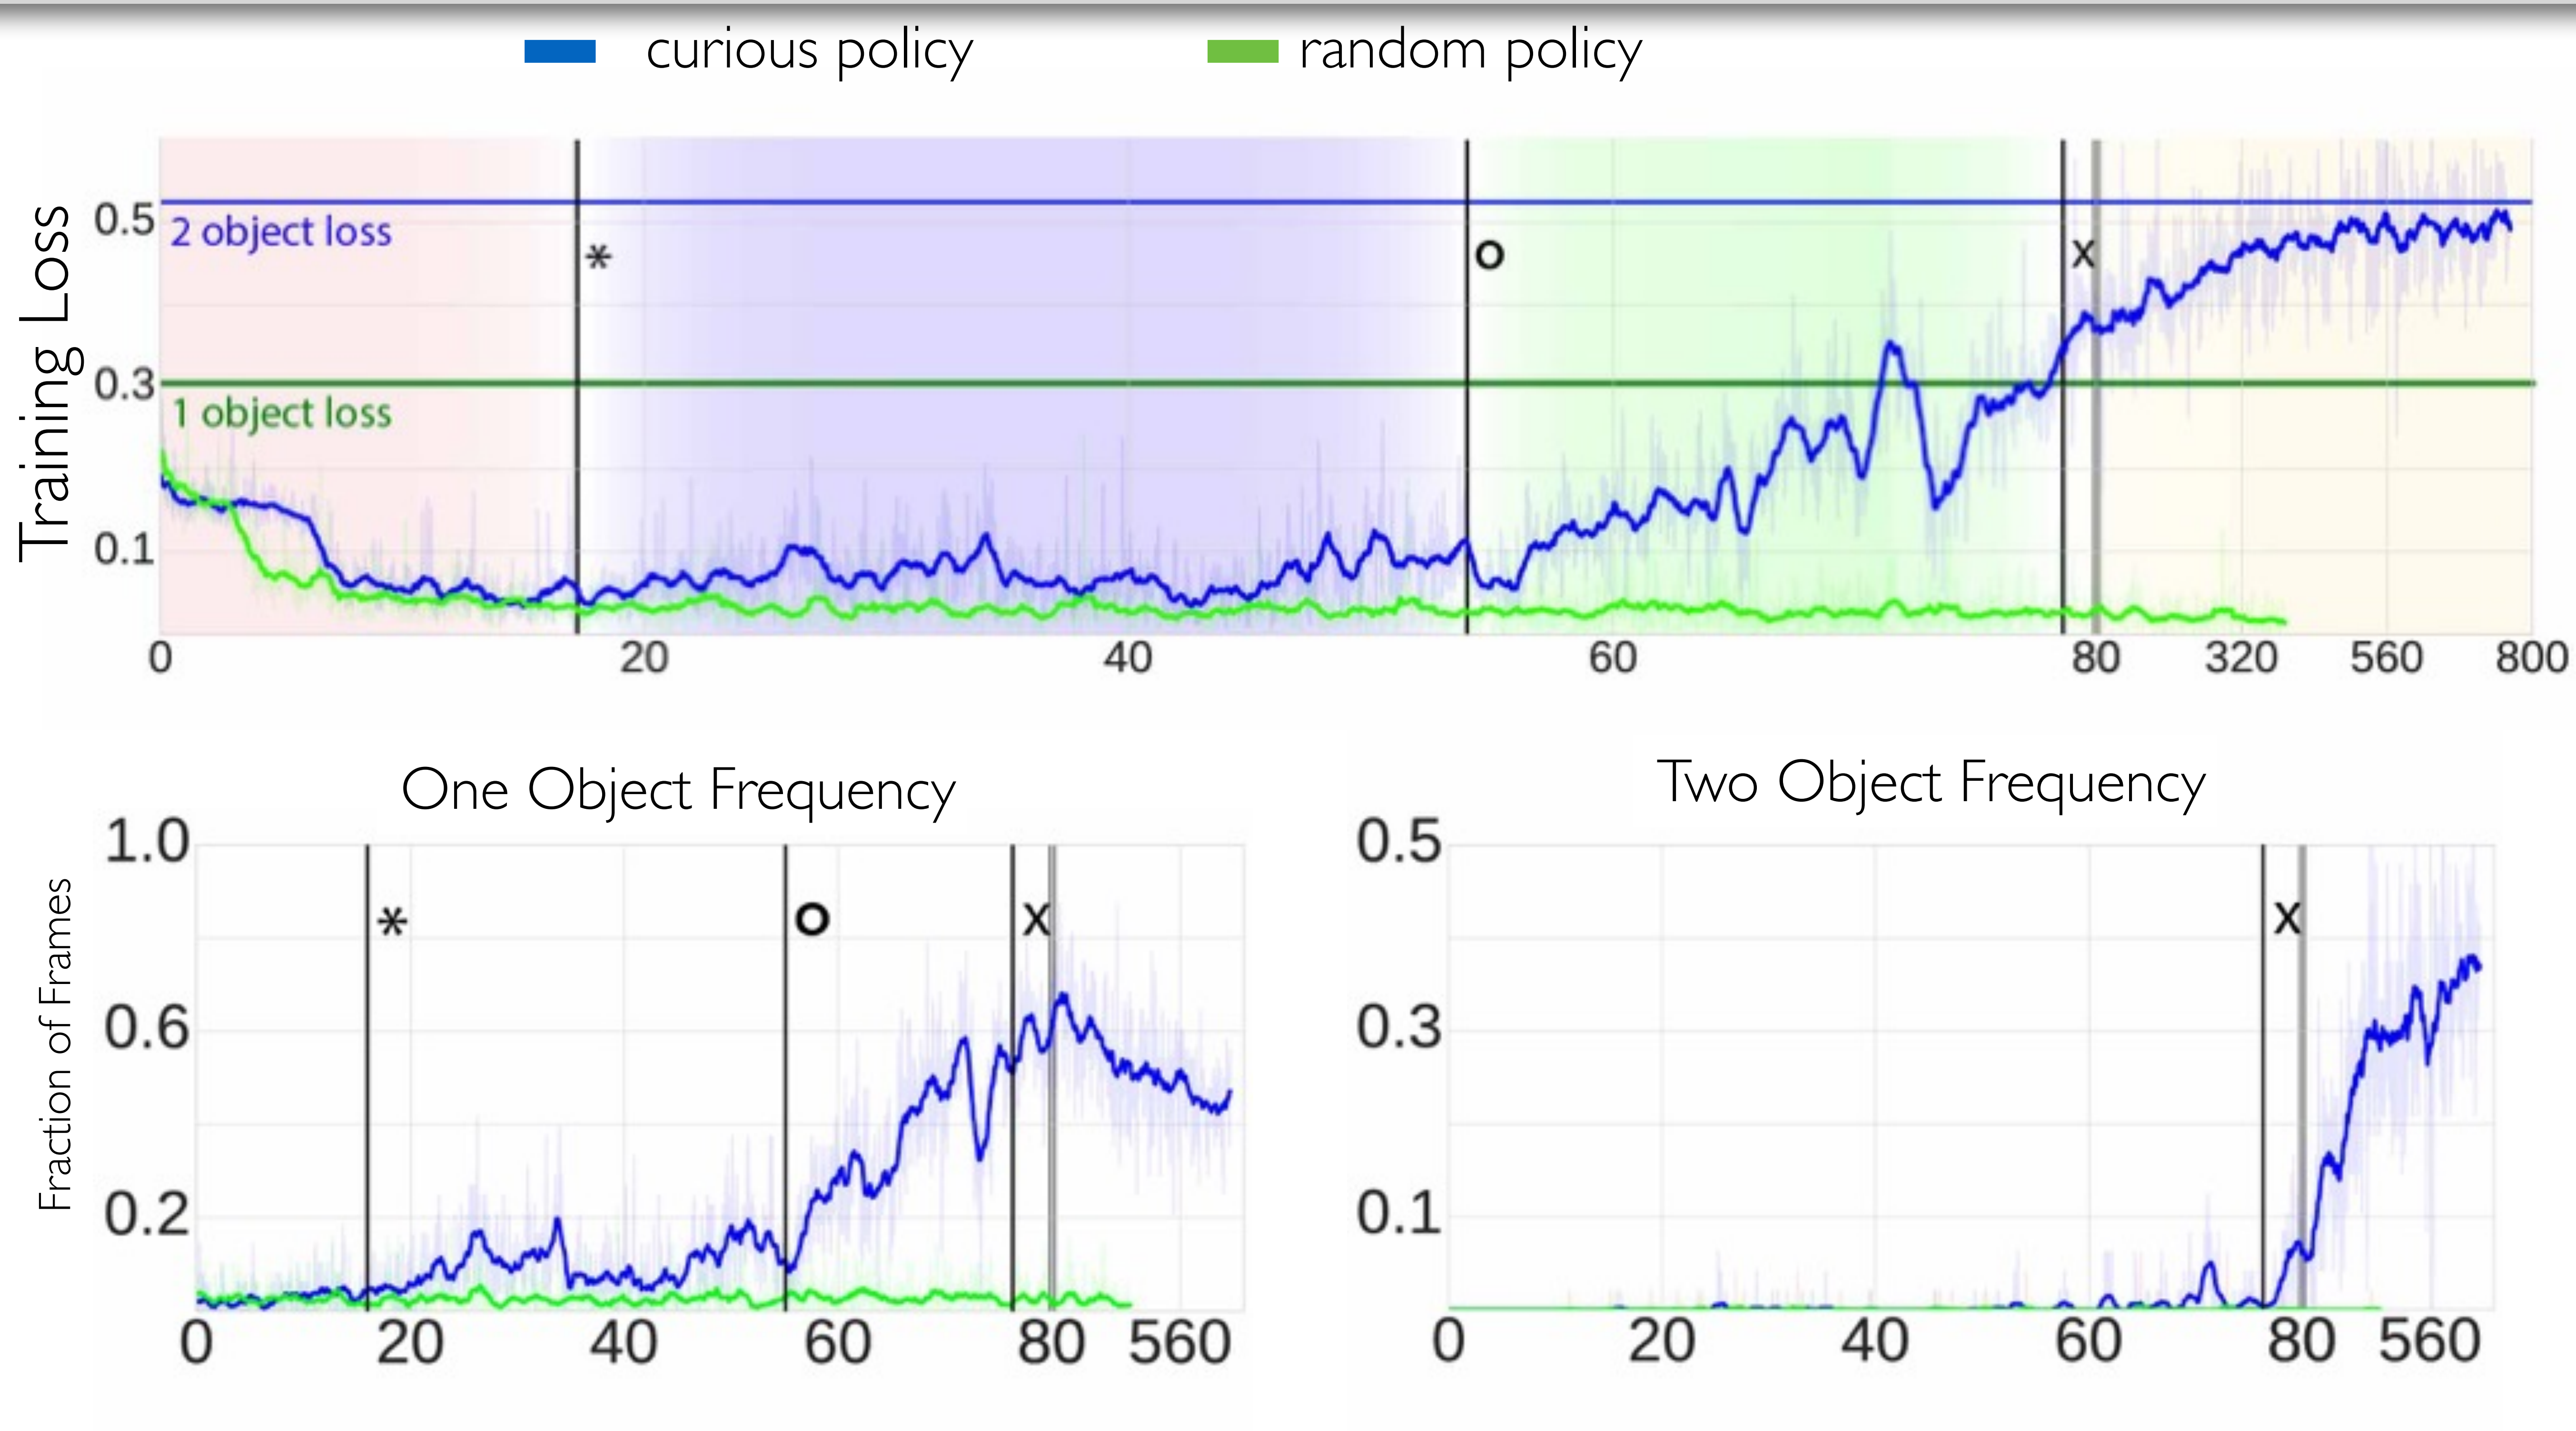

… but then discovers the interest of bringing objects together.

Self-supervised learning

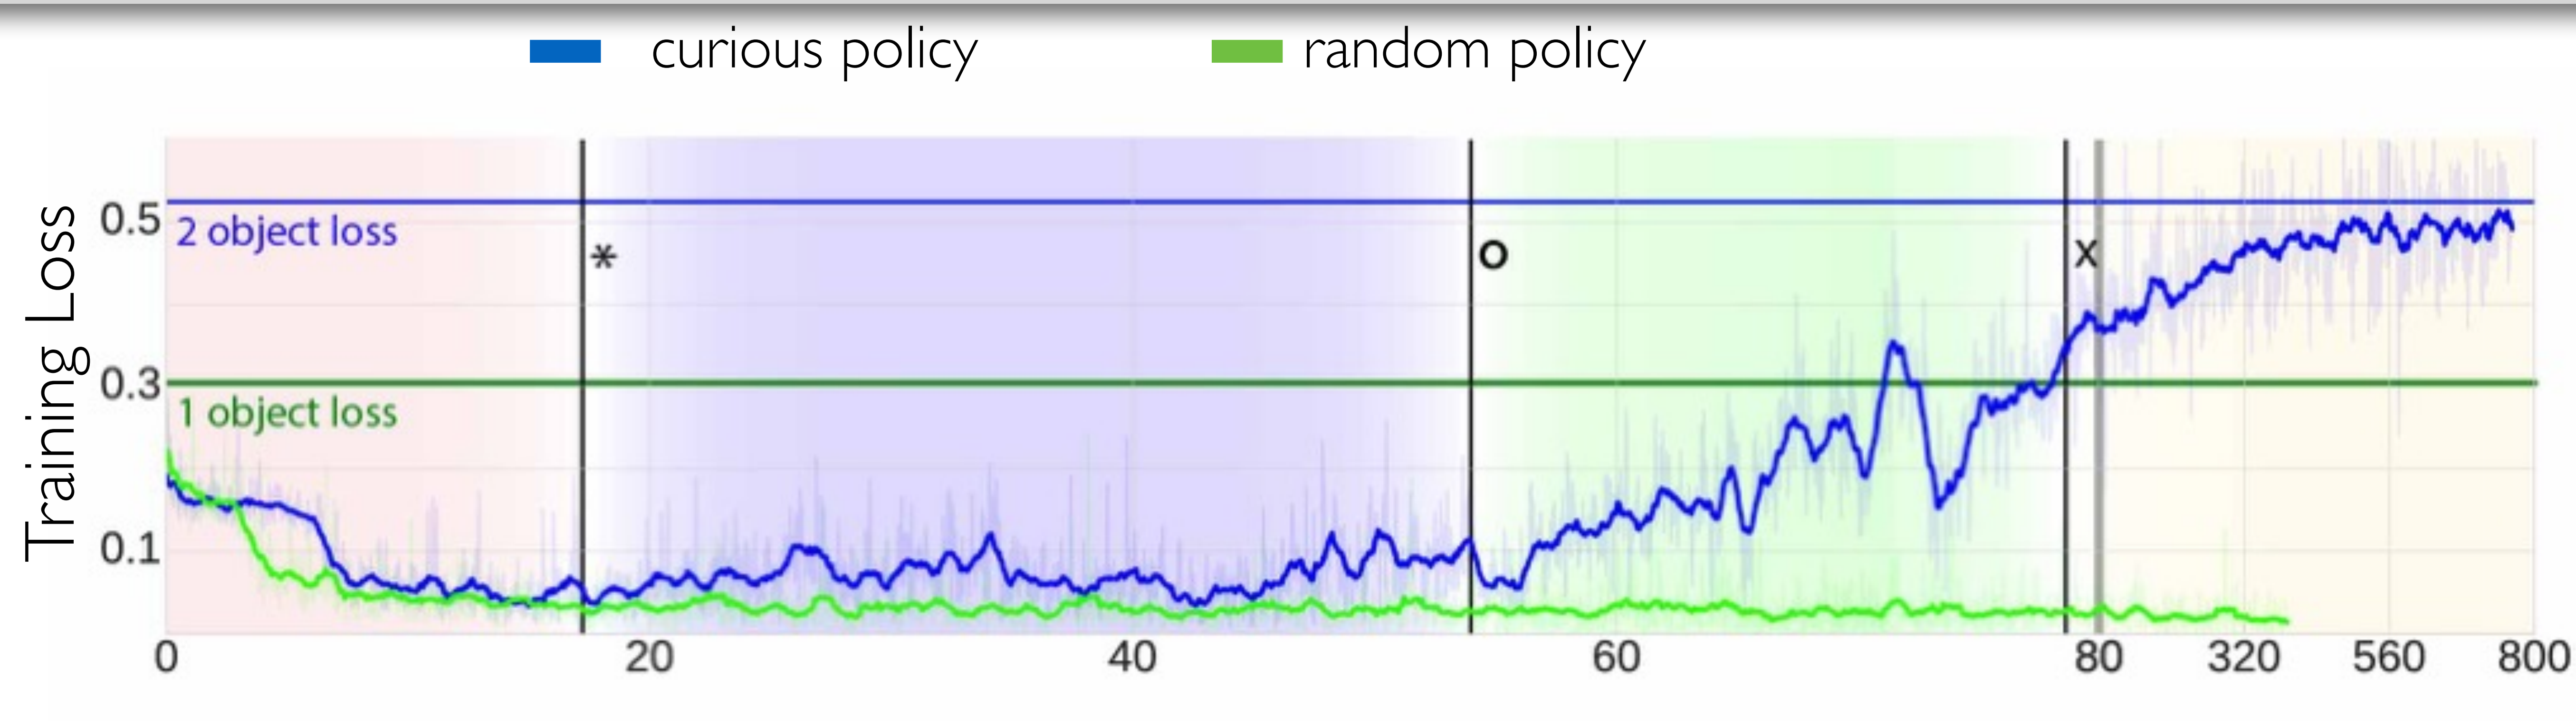

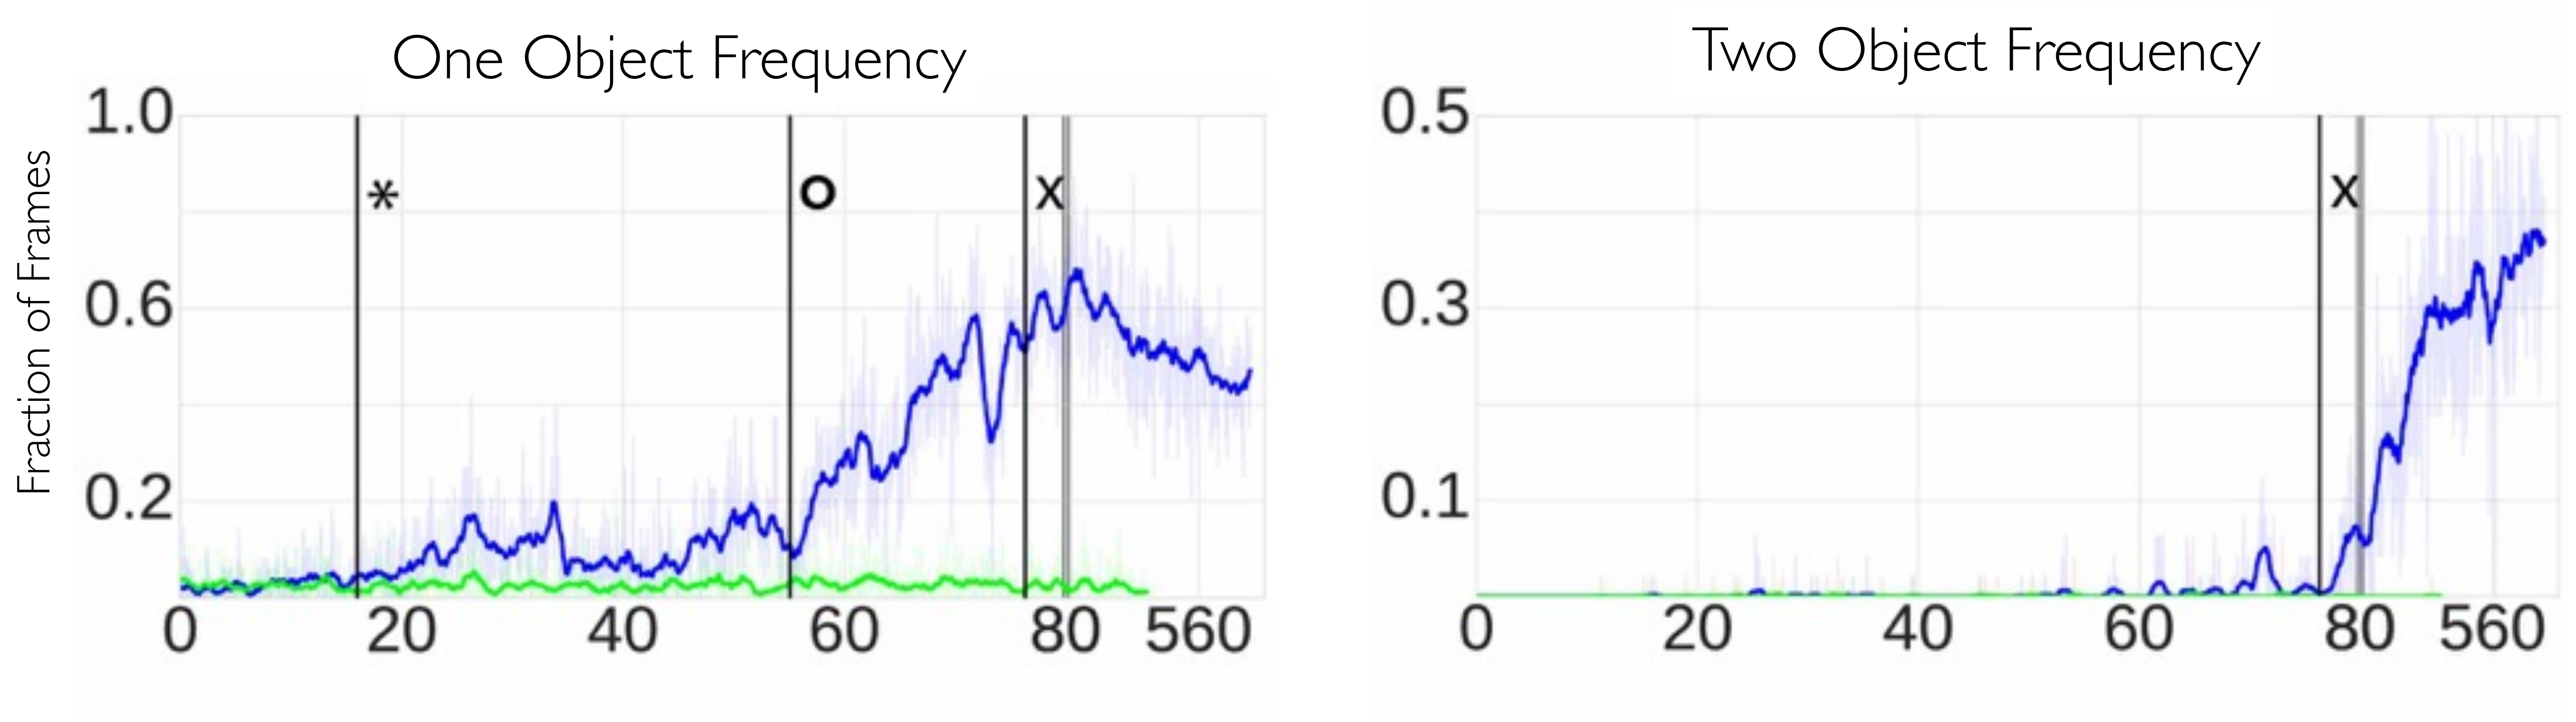

Object recognition in testing (one object per image): ~16% vs ~40% accuracy

… especially large gain compared to training in single-obj case

Glossing over a key problem:

The above ideas rely on having the agent solve a dynamics prediction problem about the world.

Self-supervised learning: General Formulation

Start with some data H…

$$H$$

Self-supervised learning: General Formulation

…*Create input data X from H*…

$$
\begin{tikzcd}
& H \arrow[dl, "\xi"'] \\
X &
\end{tikzcd}
$$

Self-supervised learning: General Formulation

…Create output data Y from H…

$$\begin{tikzcd} & H \arrow[dl, "\xi"'] \arrow[dr, "\eta"] & \\ X & & Y \end{tikzcd}$$

Self-supervised learning: General Formulation

…Predict Y from X.

$$\begin{tikzcd} & H \arrow[dl, "\xi"'] \arrow[dr, "\eta"] & \\ X \arrow[rr, dotted, "\omega"'] & & Y \end{tikzcd}$$

Self-supervised learning: General Formulation

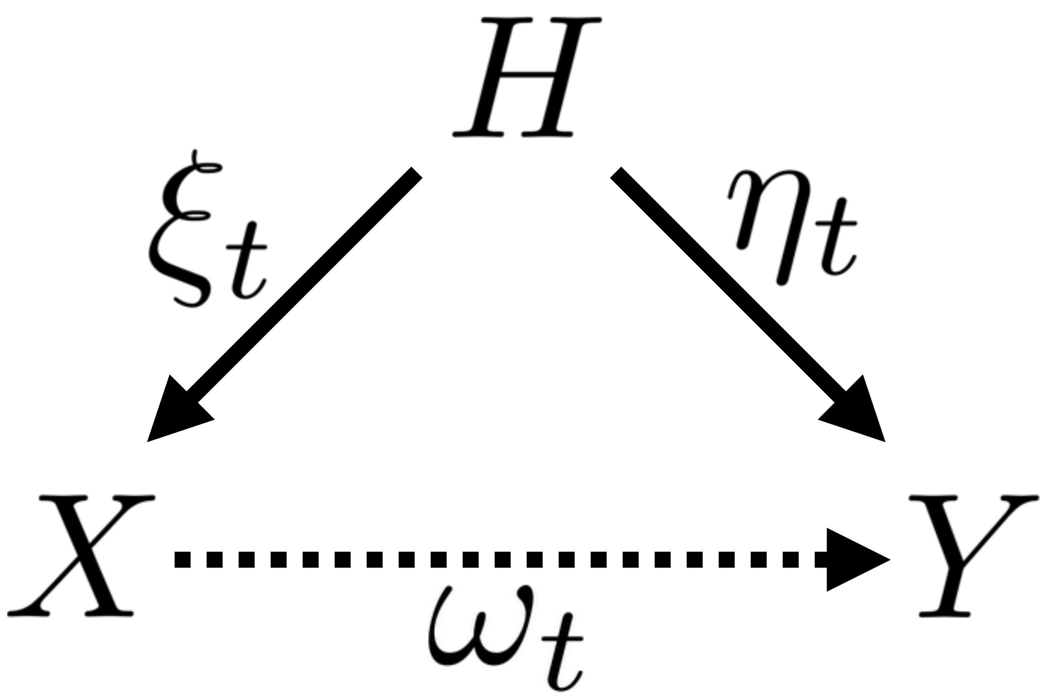

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

Self-supervised learning: General Formulation

$$\begin{tikzcd} & H \arrow[dl, "\xi_t"'] \arrow[dr, "\eta_t"] & \\ X \arrow[rr, dashed, "\omega_t"'] & & Y \end{tikzcd}$$

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

Self-supervised learning: General Formulation

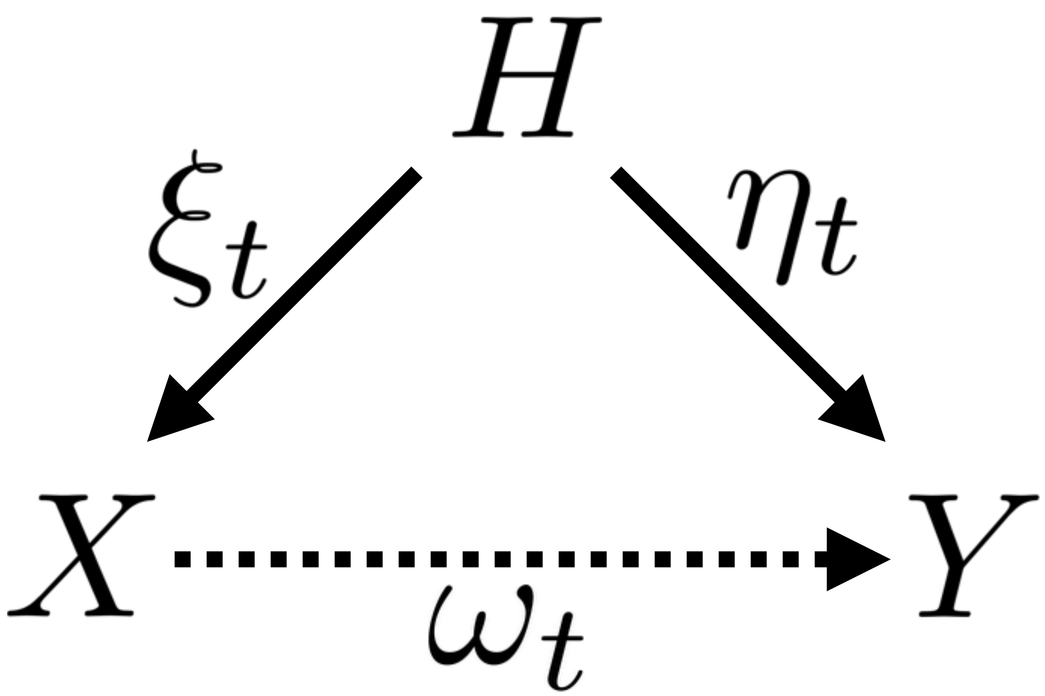

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

1) is hard, because … pixel prediction is hard!

Intuitive Physics as Underlying Goal

Obvious idea: just predict future pixels

Finn et. al (2016)

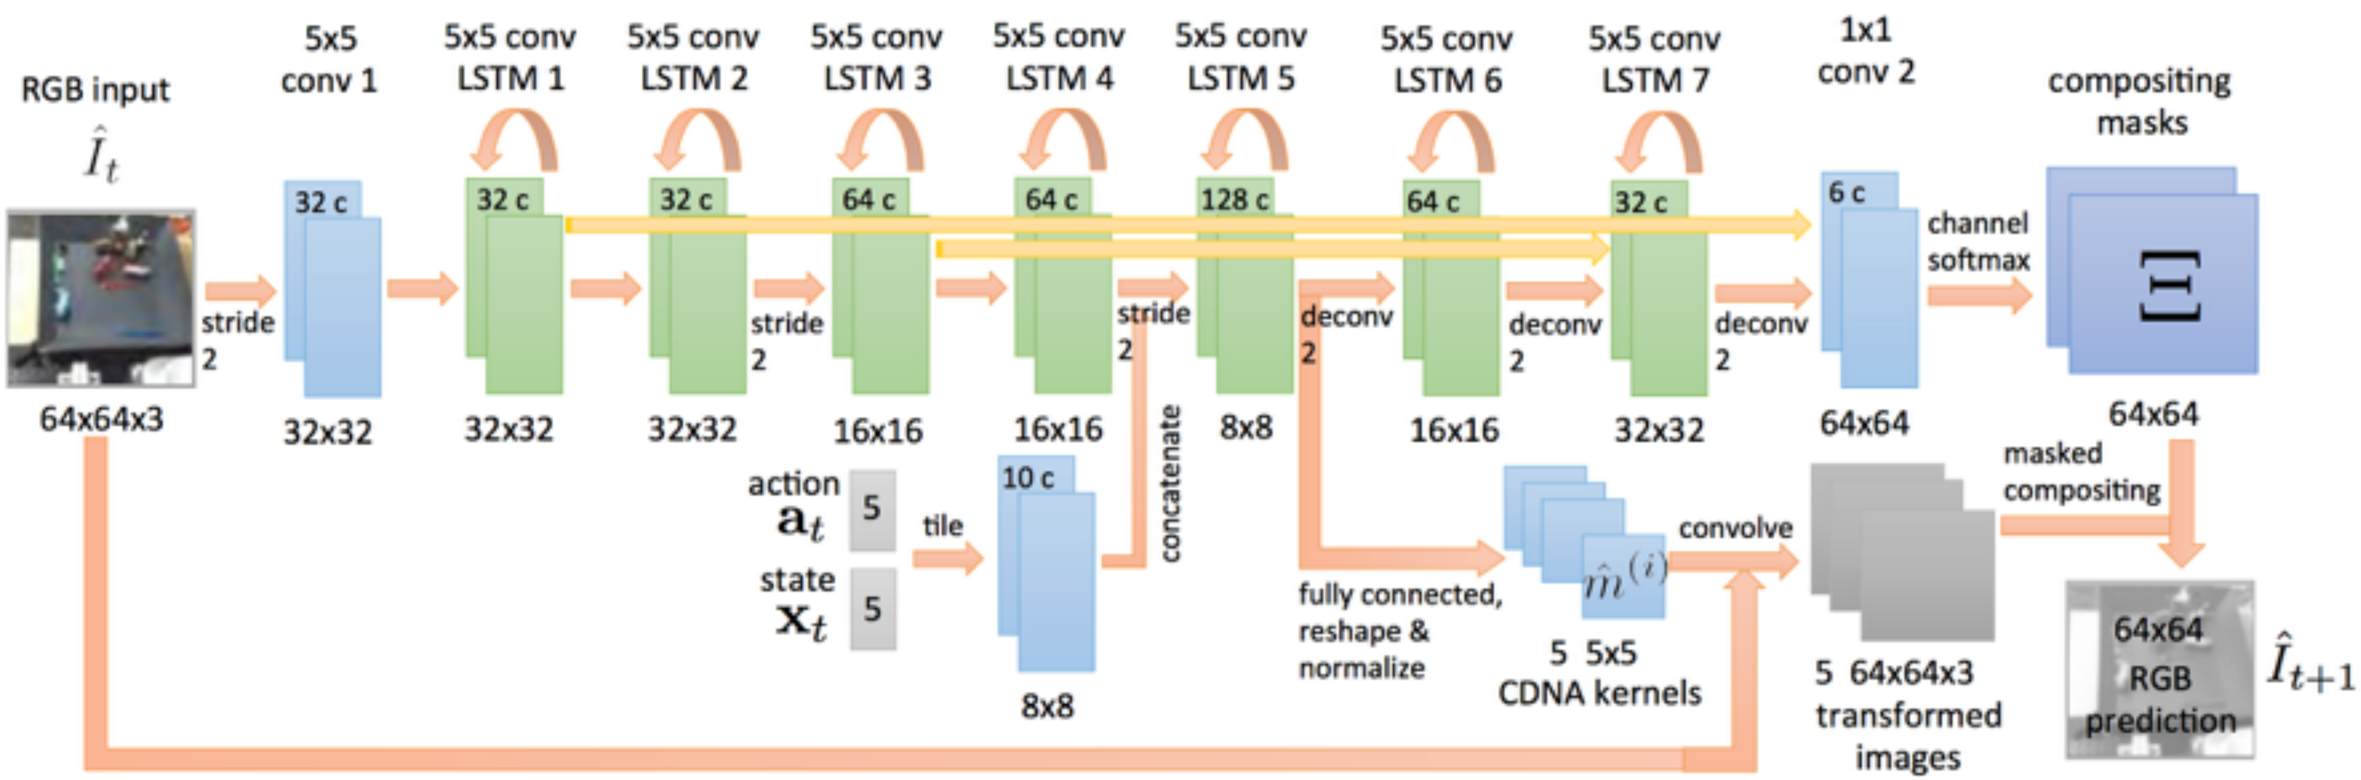

PredRNN(2017) ; Wang (2018) ; among many others

Intuitive Physics as Underlying Goal

Pixel prediction is hard.

t=1

t = 2

Two blue objects in a room

Intuitive Physics as Underlying Goal

Pixel prediction is hard.

t=1

t = 2

Two blue objects in a room

Objects acted on and camera moves

Intuitive Physics as Underlying Goal

Pixel prediction is hard.

t=1

t = 2

Two blue objects in a room

Objects acted on and camera moves

Ground truth

t=3

t = 4

t=5

t=6

Intuitive Physics as Underlying Goal

Pixel prediction is hard.

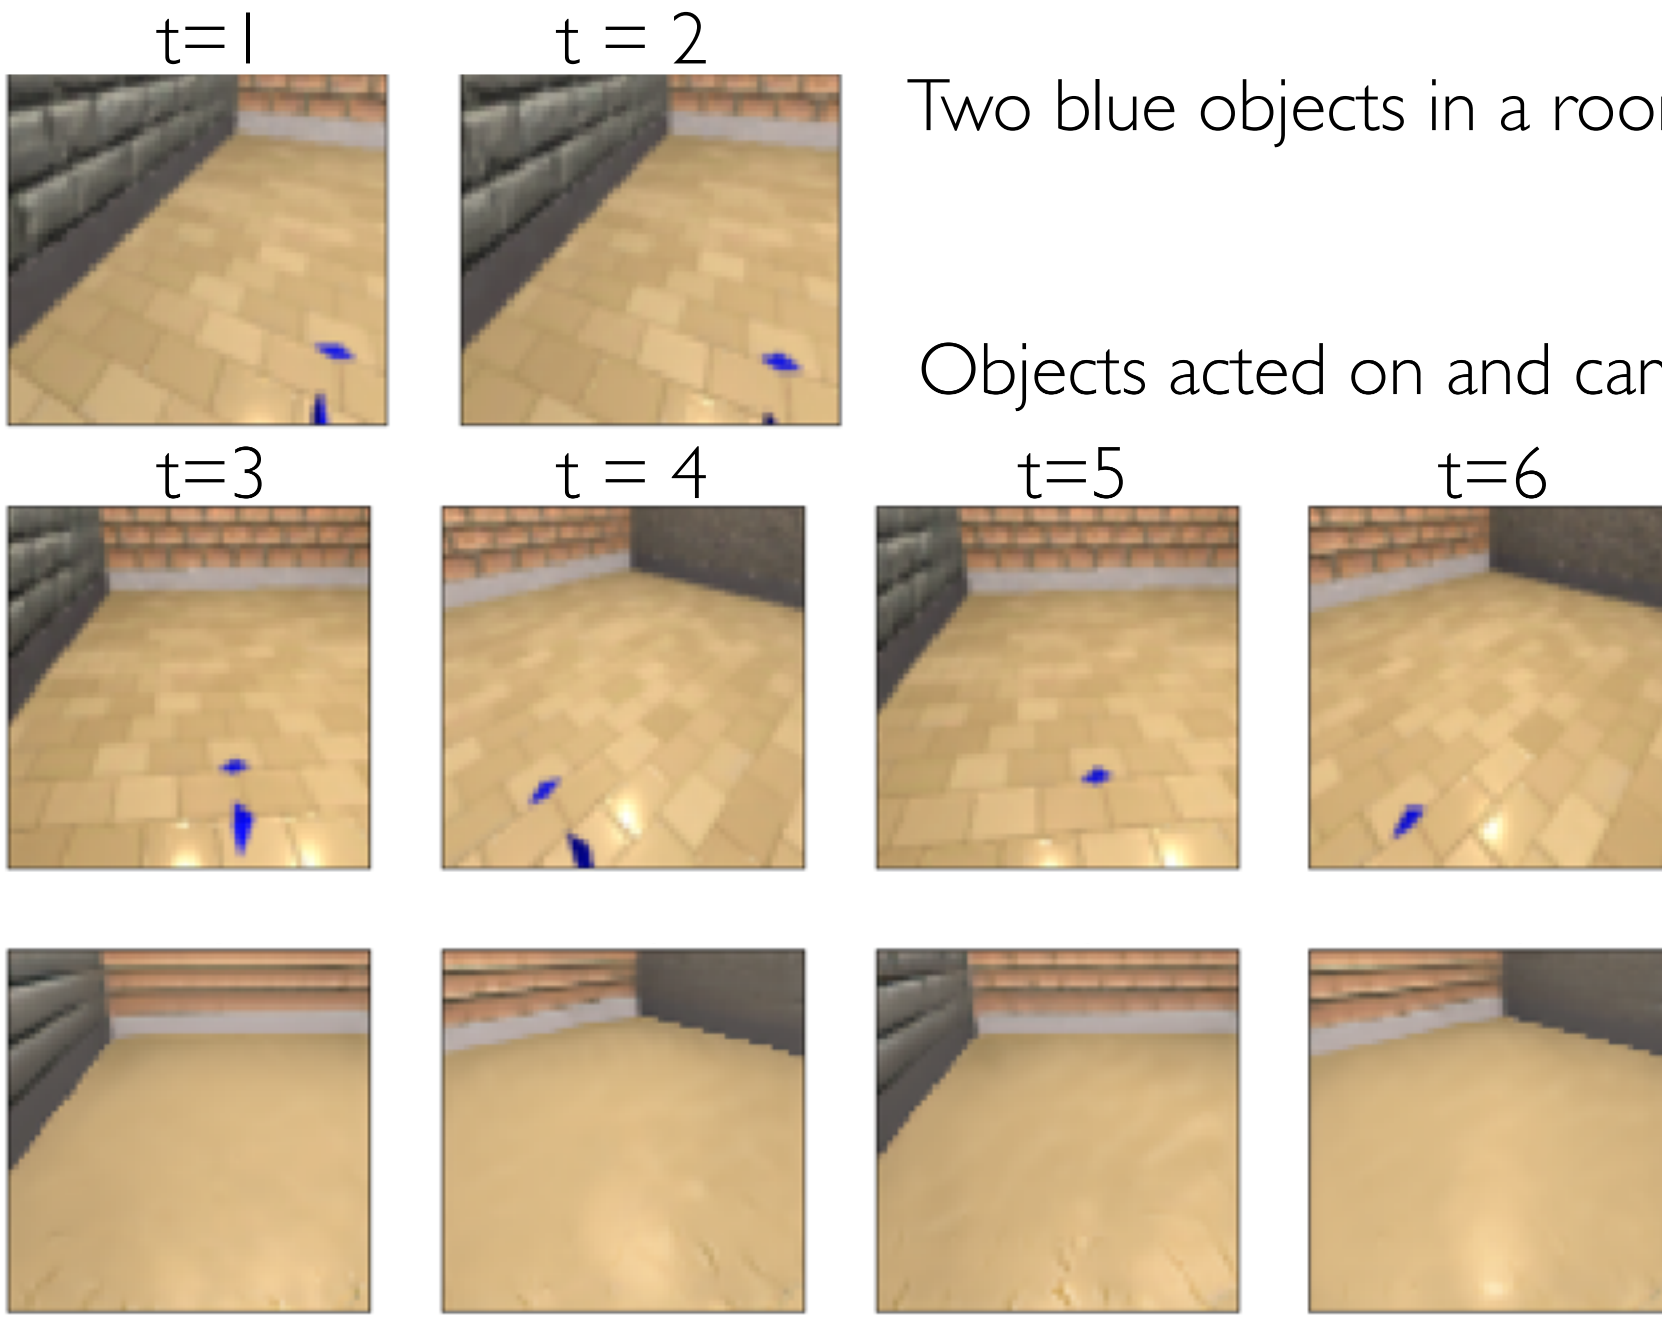

Two blue objects in a room

Objects acted on and camera moves

Ground truth

Prediction

Self-supervised learning: General Formulation

$$\begin{tikzcd} & H \arrow[dl, "\xi_t"'] \arrow[dr, "\eta_t"] & \\ X \arrow[rr, dashed, "\omega_t"'] & & Y \end{tikzcd}$$

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

1) is hard, because … pixel prediction is hard!

2) is mostly what we did in the work described above because it's easier …

Self-supervised learning: General Formulation

$$\begin{tikzcd} & H \arrow[dl, "\xi_t"'] \arrow[dr, "\eta_t"] & \\ X \arrow[rr, dashed, "\omega_t"'] & & Y \end{tikzcd}$$

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

1) is hard, because … pixel prediction is hard!

2) is mostly what we did in the work described above because it's easier …

BUT DEGENERATE!

Self-supervised learning: General Formulation

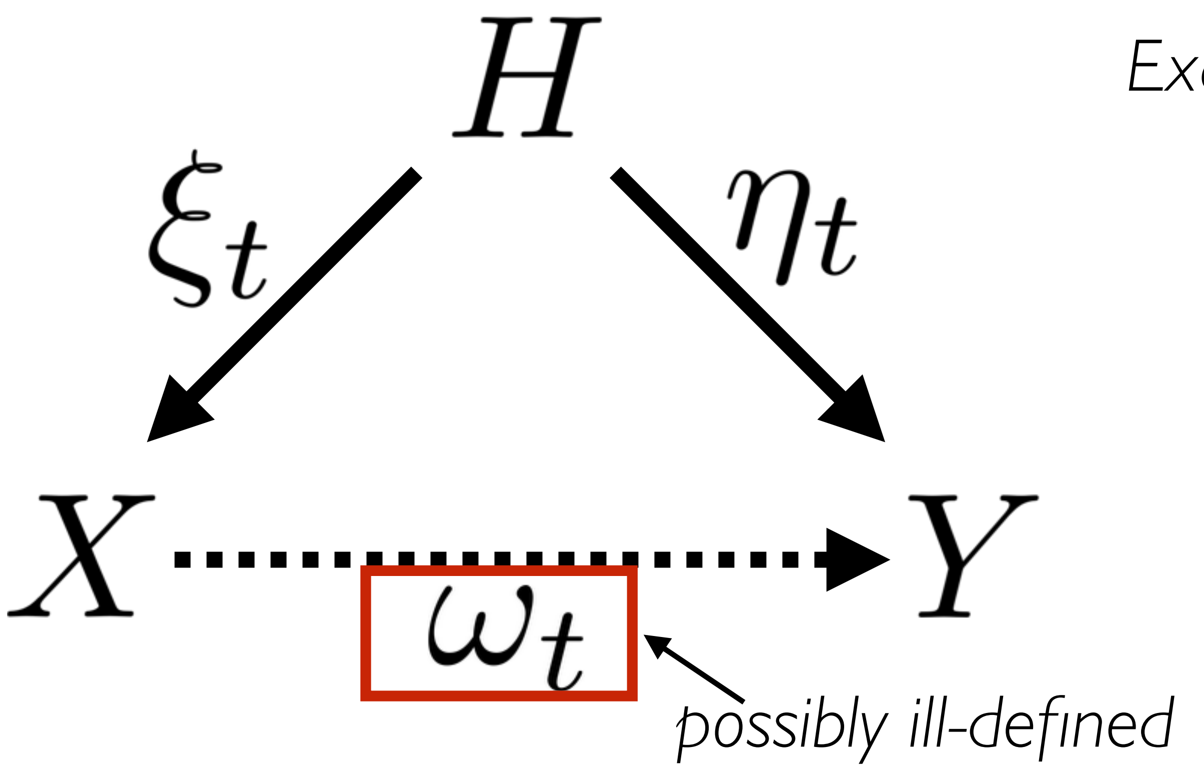

THE DREADED

WHITE-NOISE PROBLEM

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

1) is hard, because … pixel prediction is hard!

2) is mostly what we did in the work described above because it's easier …

BUT DEGENERATE!

Self-supervised learning: General Formulation

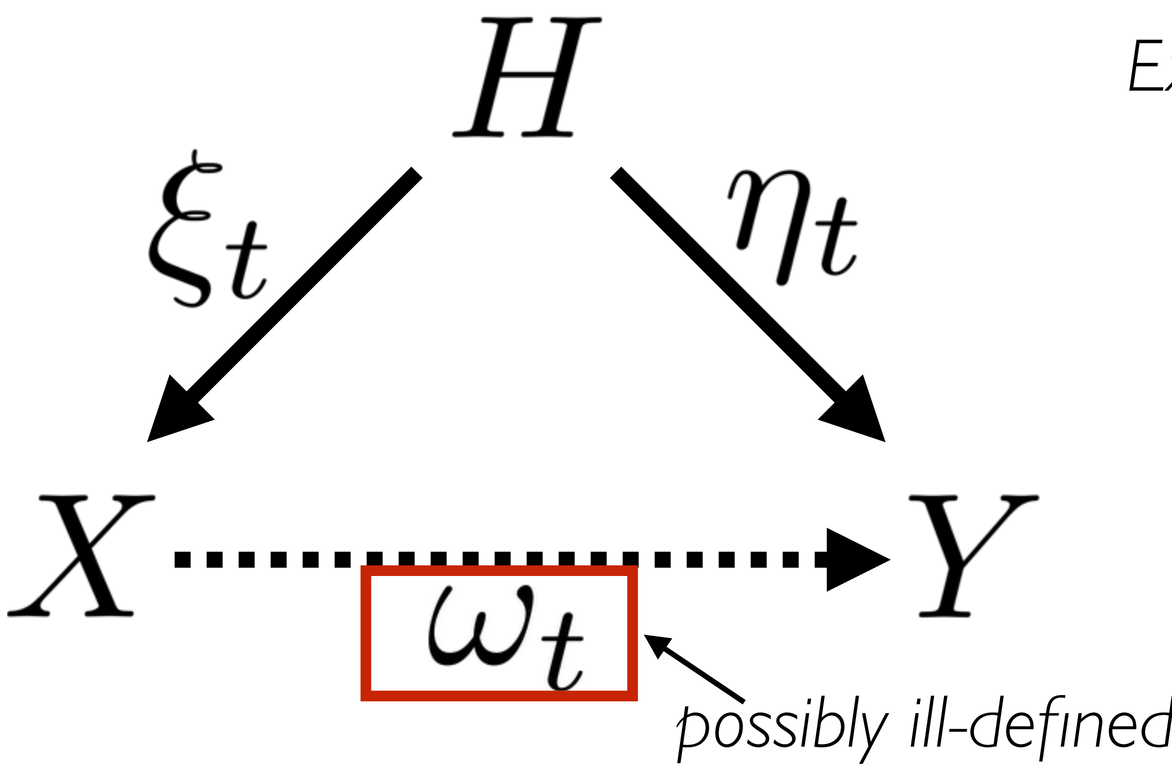

THE DREADED
WHITE-NOISE PROBLEM

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

2) inverse dynamics prediction

state(t), state(t+1) $\Longrightarrow$ *action(t)*

1) is hard, because … pixel prediction is hard!

2) is mostly what we did in the previous work because it's easier …

BUT DEGENERATE!

Ex: pushing down on an object

Self-supervised learning: General Formulation

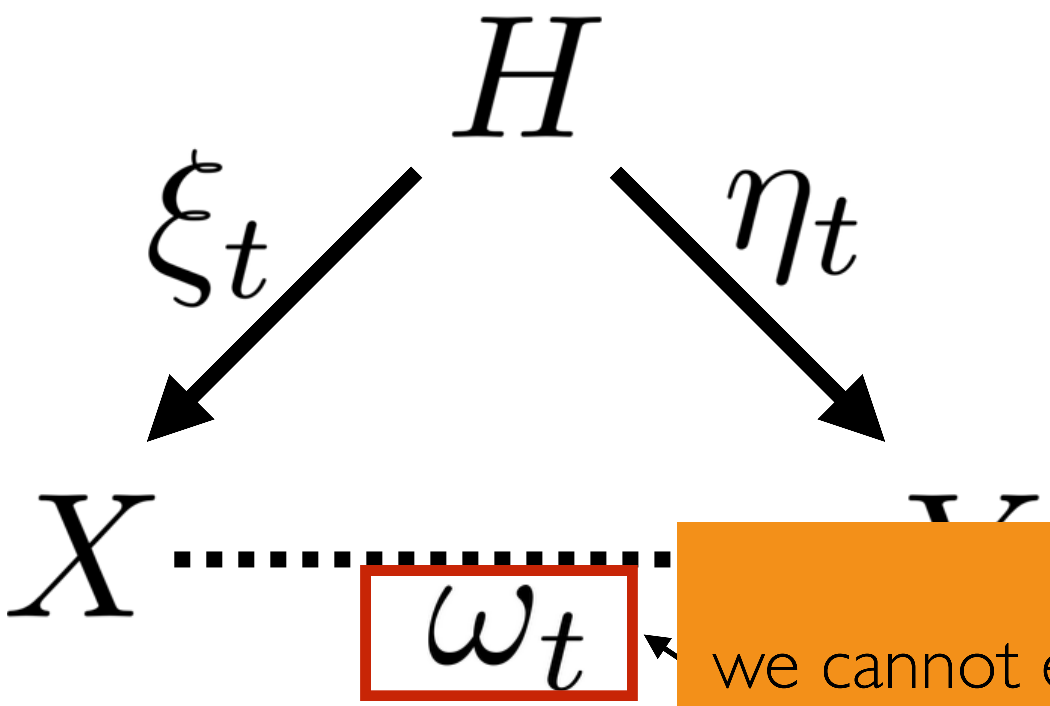

Examples:

1) Forward future prediction

state(t), action(t) $\Longrightarrow$ *state(t+1)*

rediction

$\Longrightarrow$ *action(t)*

Conclusion:
we cannot escape having to do better future prediction
— so let's attack the problem directly.

THE DREADE

WHITE-NOISE PROBLEM

1) is hard, because … pixel prediction is hard!

2) is mostly what we did in the previous work because it's easier …

BUT DEGENERATE!

Ex: pushing down on an object

Damian Mrowca*

Chengxu Zhuang*

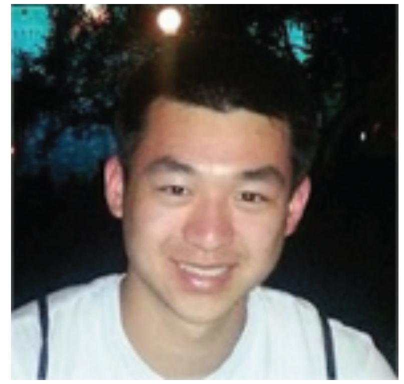

Eli Wang

Nick Haber

Fei-Fei Li

Josh Tenenbaum

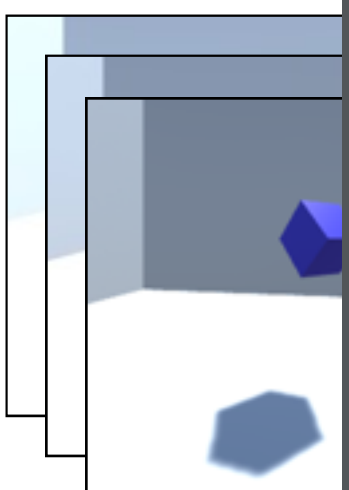

t-n …

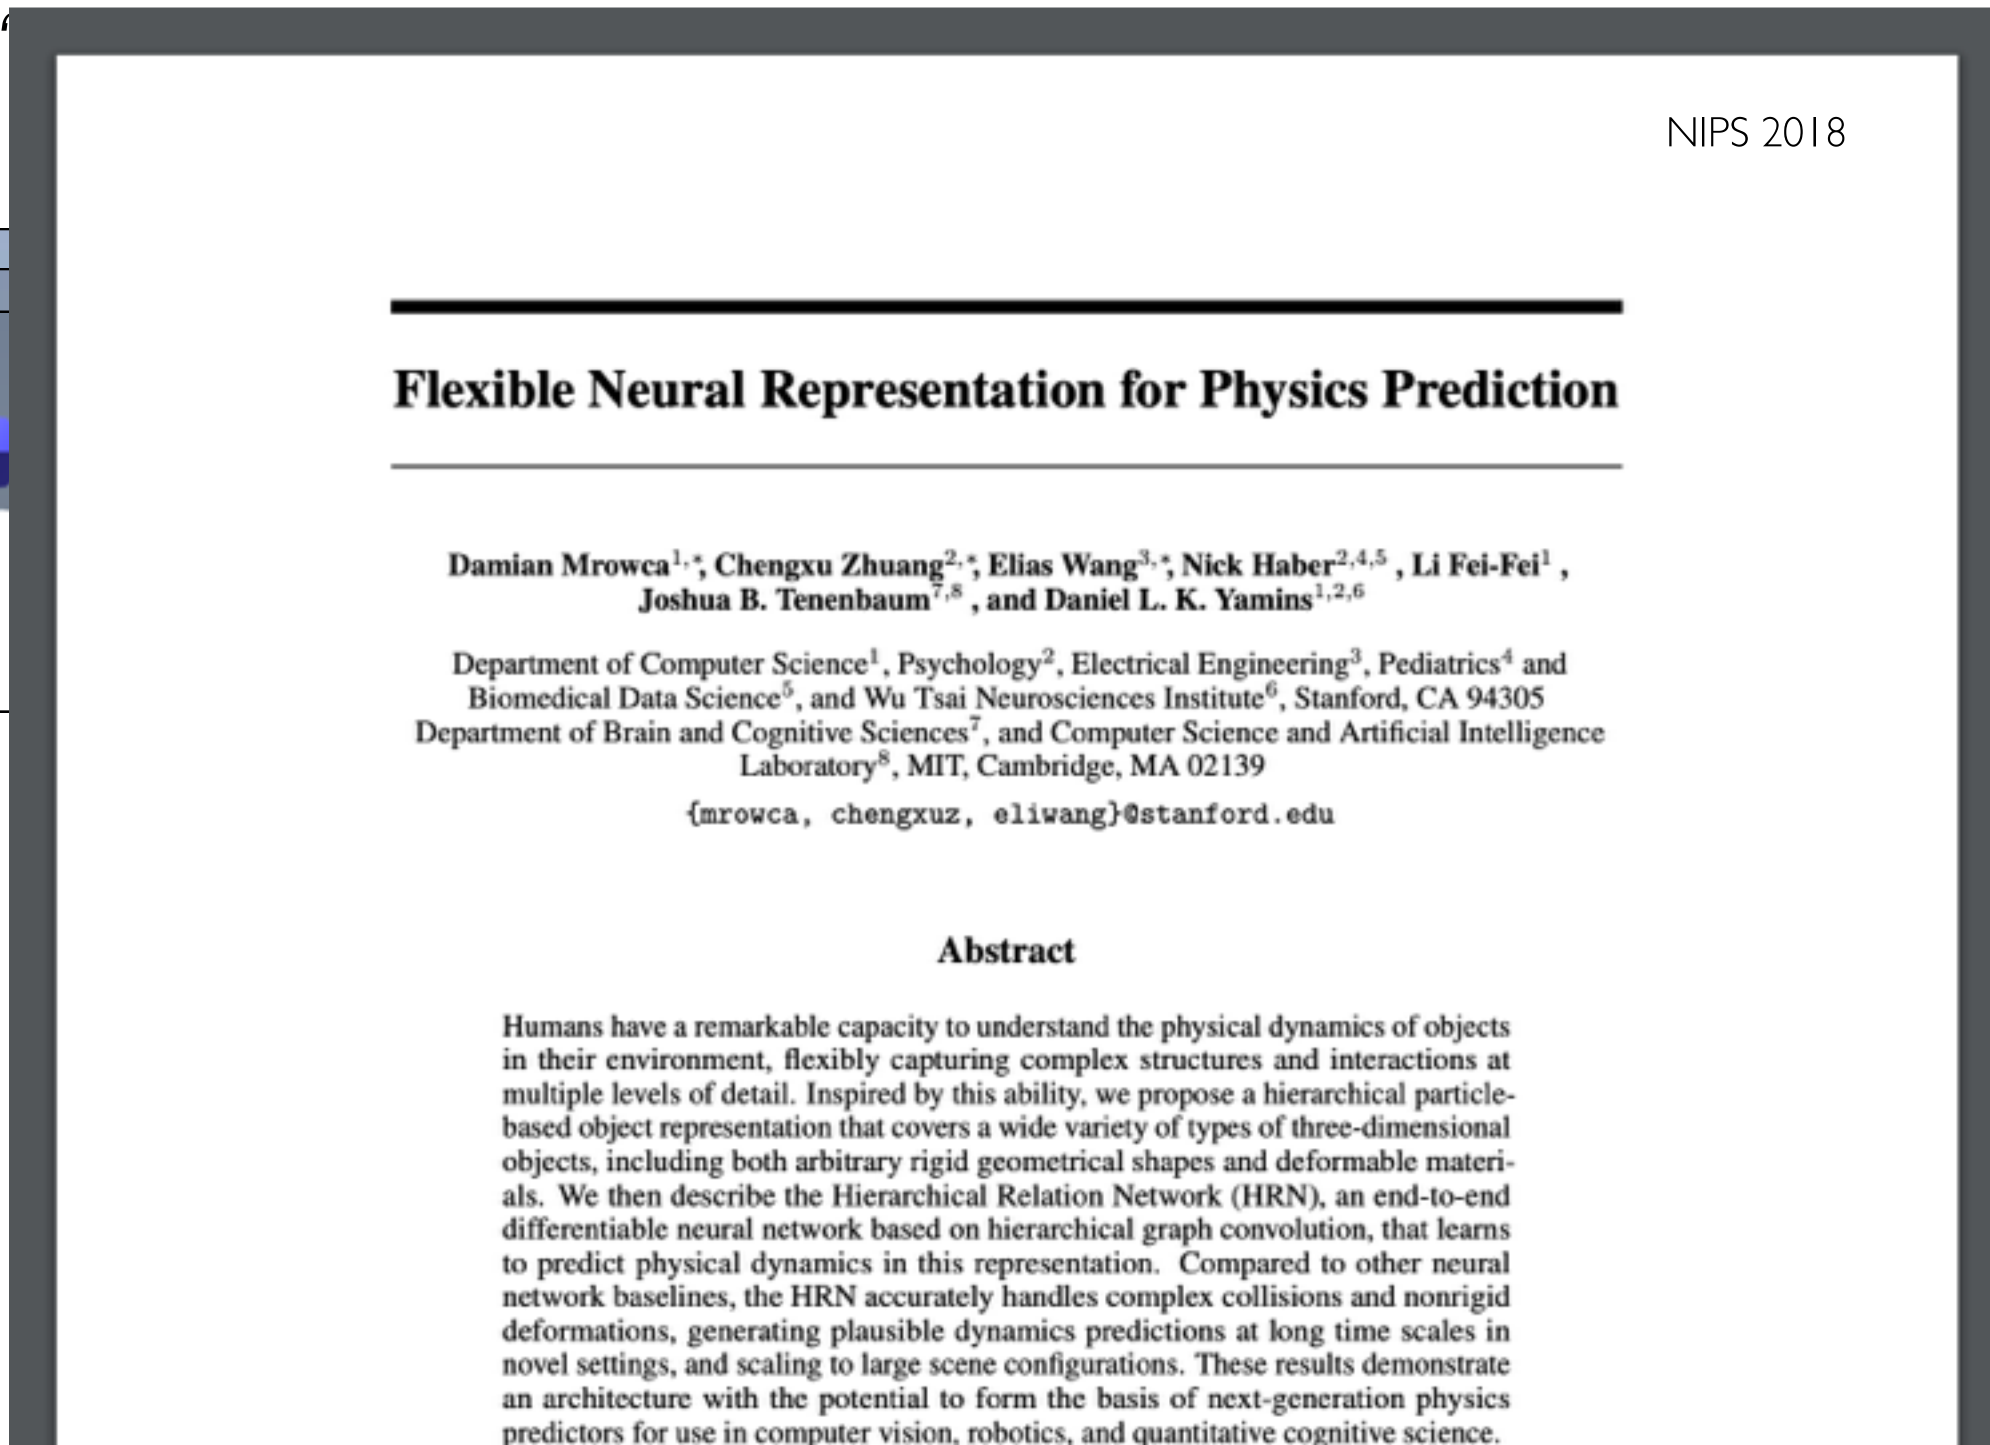

NIPS 2018

Flexible Neural Representation for Physics Prediction

Damian Mrowca[1,*], **Chengxu Zhuang**[2,*], **Elias Wang**[3,*], **Nick Haber**[2,4,5], **Li Fei-Fei**[1], **Joshua B. Tenenbaum**[7,8], **and Daniel L. K. Yamins**[1,2,6]

Department of Computer Science[1], Psychology[2], Electrical Engineering[3], Pediatrics[4] and Biomedical Data Science[5], and Wu Tsai Neurosciences Institute[6], Stanford, CA 94305
Department of Brain and Cognitive Sciences[7], and Computer Science and Artificial Intelligence Laboratory[8], MIT, Cambridge, MA 02139

{mrowca, chengxuz, eliwang}@stanford.edu

Abstract

Humans have a remarkable capacity to understand the physical dynamics of objects in their environment, flexibly capturing complex structures and interactions at multiple levels of detail. Inspired by this ability, we propose a hierarchical particle-based object representation that covers a wide variety of types of three-dimensional objects, including both arbitrary rigid geometrical shapes and deformable materials. We then describe the Hierarchical Relation Network (HRN), an end-to-end differentiable neural network based on hierarchical graph convolution, that learns to predict physical dynamics in this representation. Compared to other neural network baselines, the HRN accurately handles complex collisions and nonrigid deformations, generating plausible dynamics predictions at long time scales in novel settings, and scaling to large scene configurations. These results demonstrate an architecture with the potential to form the basis of next-generation physics predictors for use in computer vision, robotics, and quantitative cognitive science.

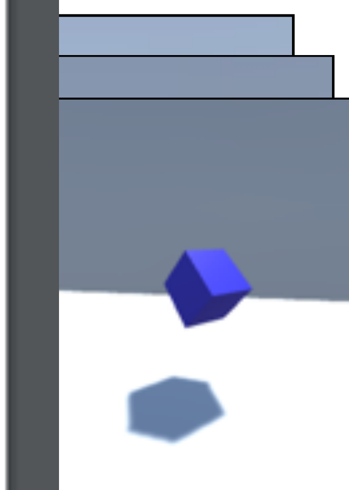

-1 … t+k

Discovering the proper latent space for physical prediction…

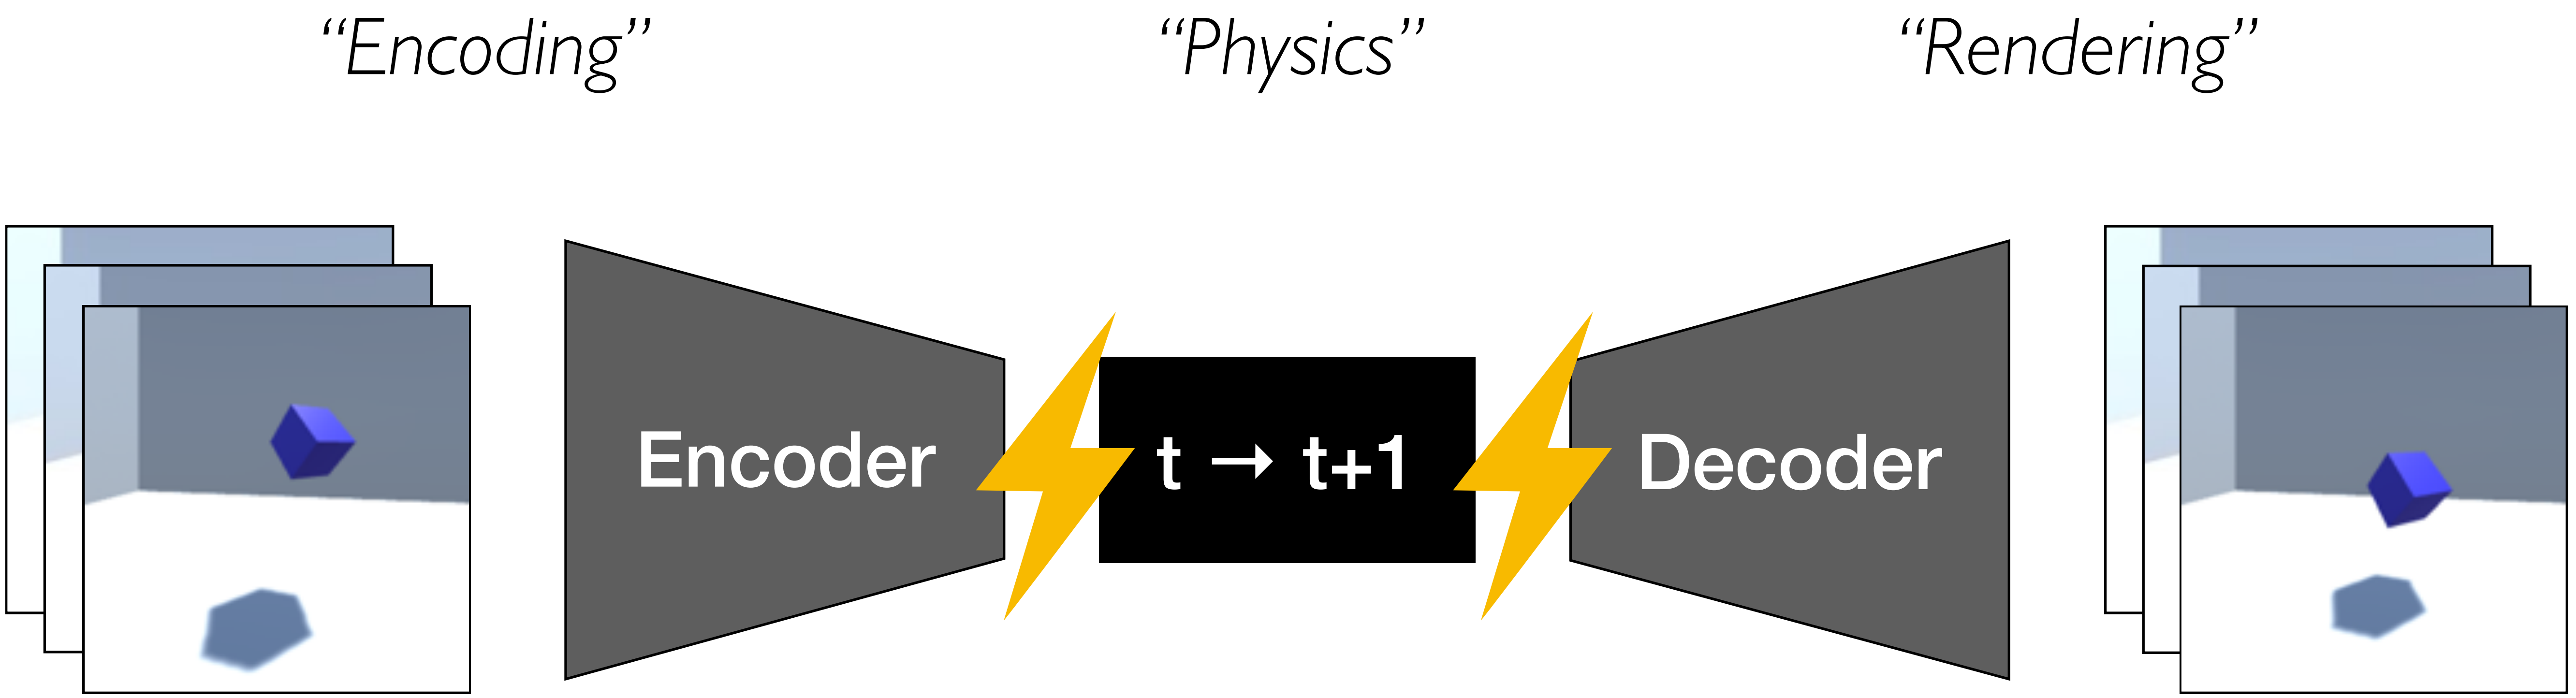

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early, perhaps by 3 months.

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early . . .

ELSEVIER

Cognition

Volume 20, Issue 3, 1985, Pages 191-208

COGNITION

Object permanence in five-month-old infants ☆

Renée Baillargeon *, Elizabeth S. Spelke *, Stanley Wasserman *

⊞ Show more

https://doi.org/10.1016/0010-0277(85)90008-3

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early …

ELSEVIER

Cognition

Volume 20, Issue 3, 1985, Pages 191-208

COGNITION

Object permanence in five-month-old infants ☆

Renée Baillargeon *, Elizabeth S. Spelke *, Stanley Wasserman *

⊞ Show more

https://doi.org/10.1016/0010-0277(85)90008-3

Get rights and content

Conv2d structures, even with RNNs, have trouble with object permanence.

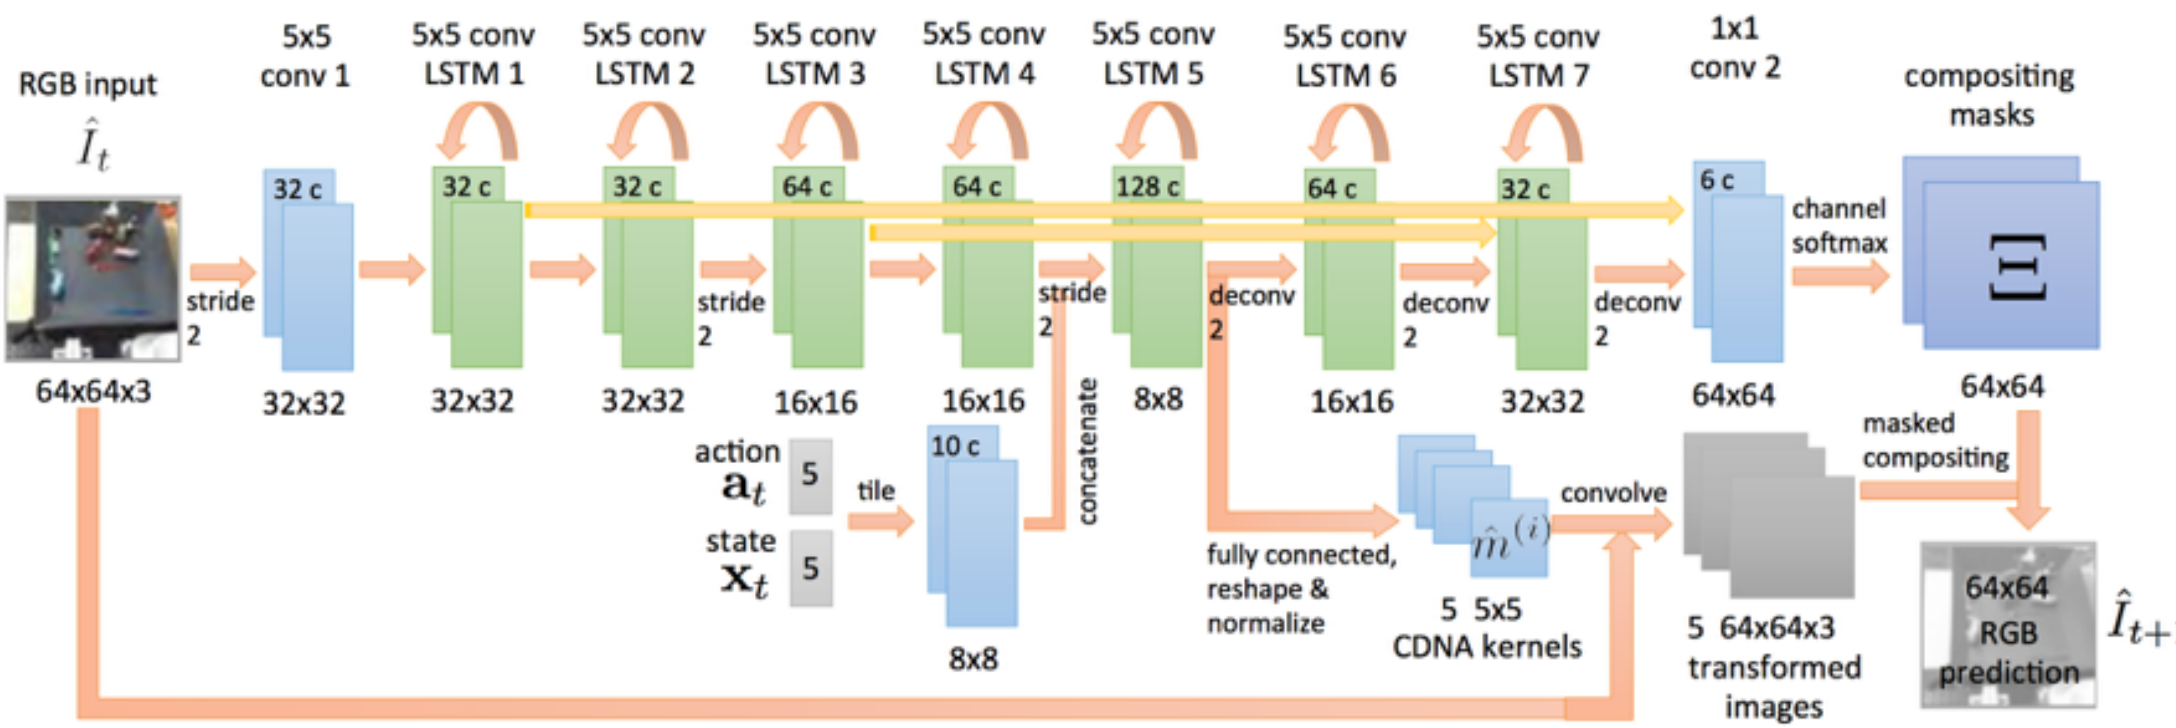

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early, perhaps by 3 months.

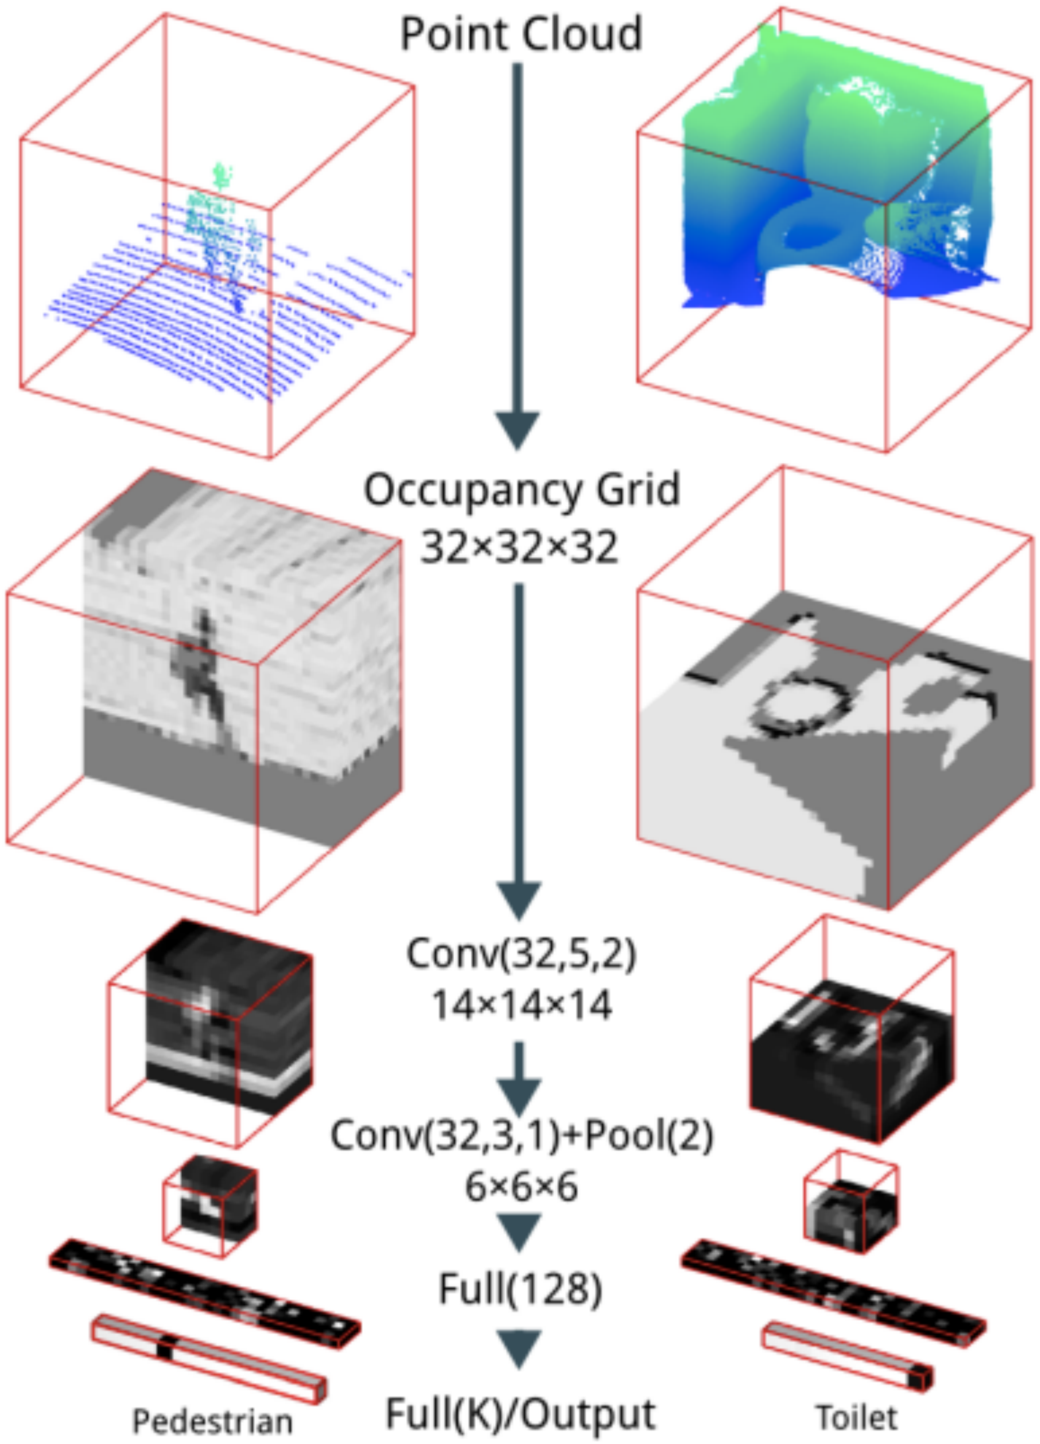

Conv3d structures are better for object permanence, but very inefficient: hard to achieve high resolution.

Spatial convolutions are not ideal for physics propagation

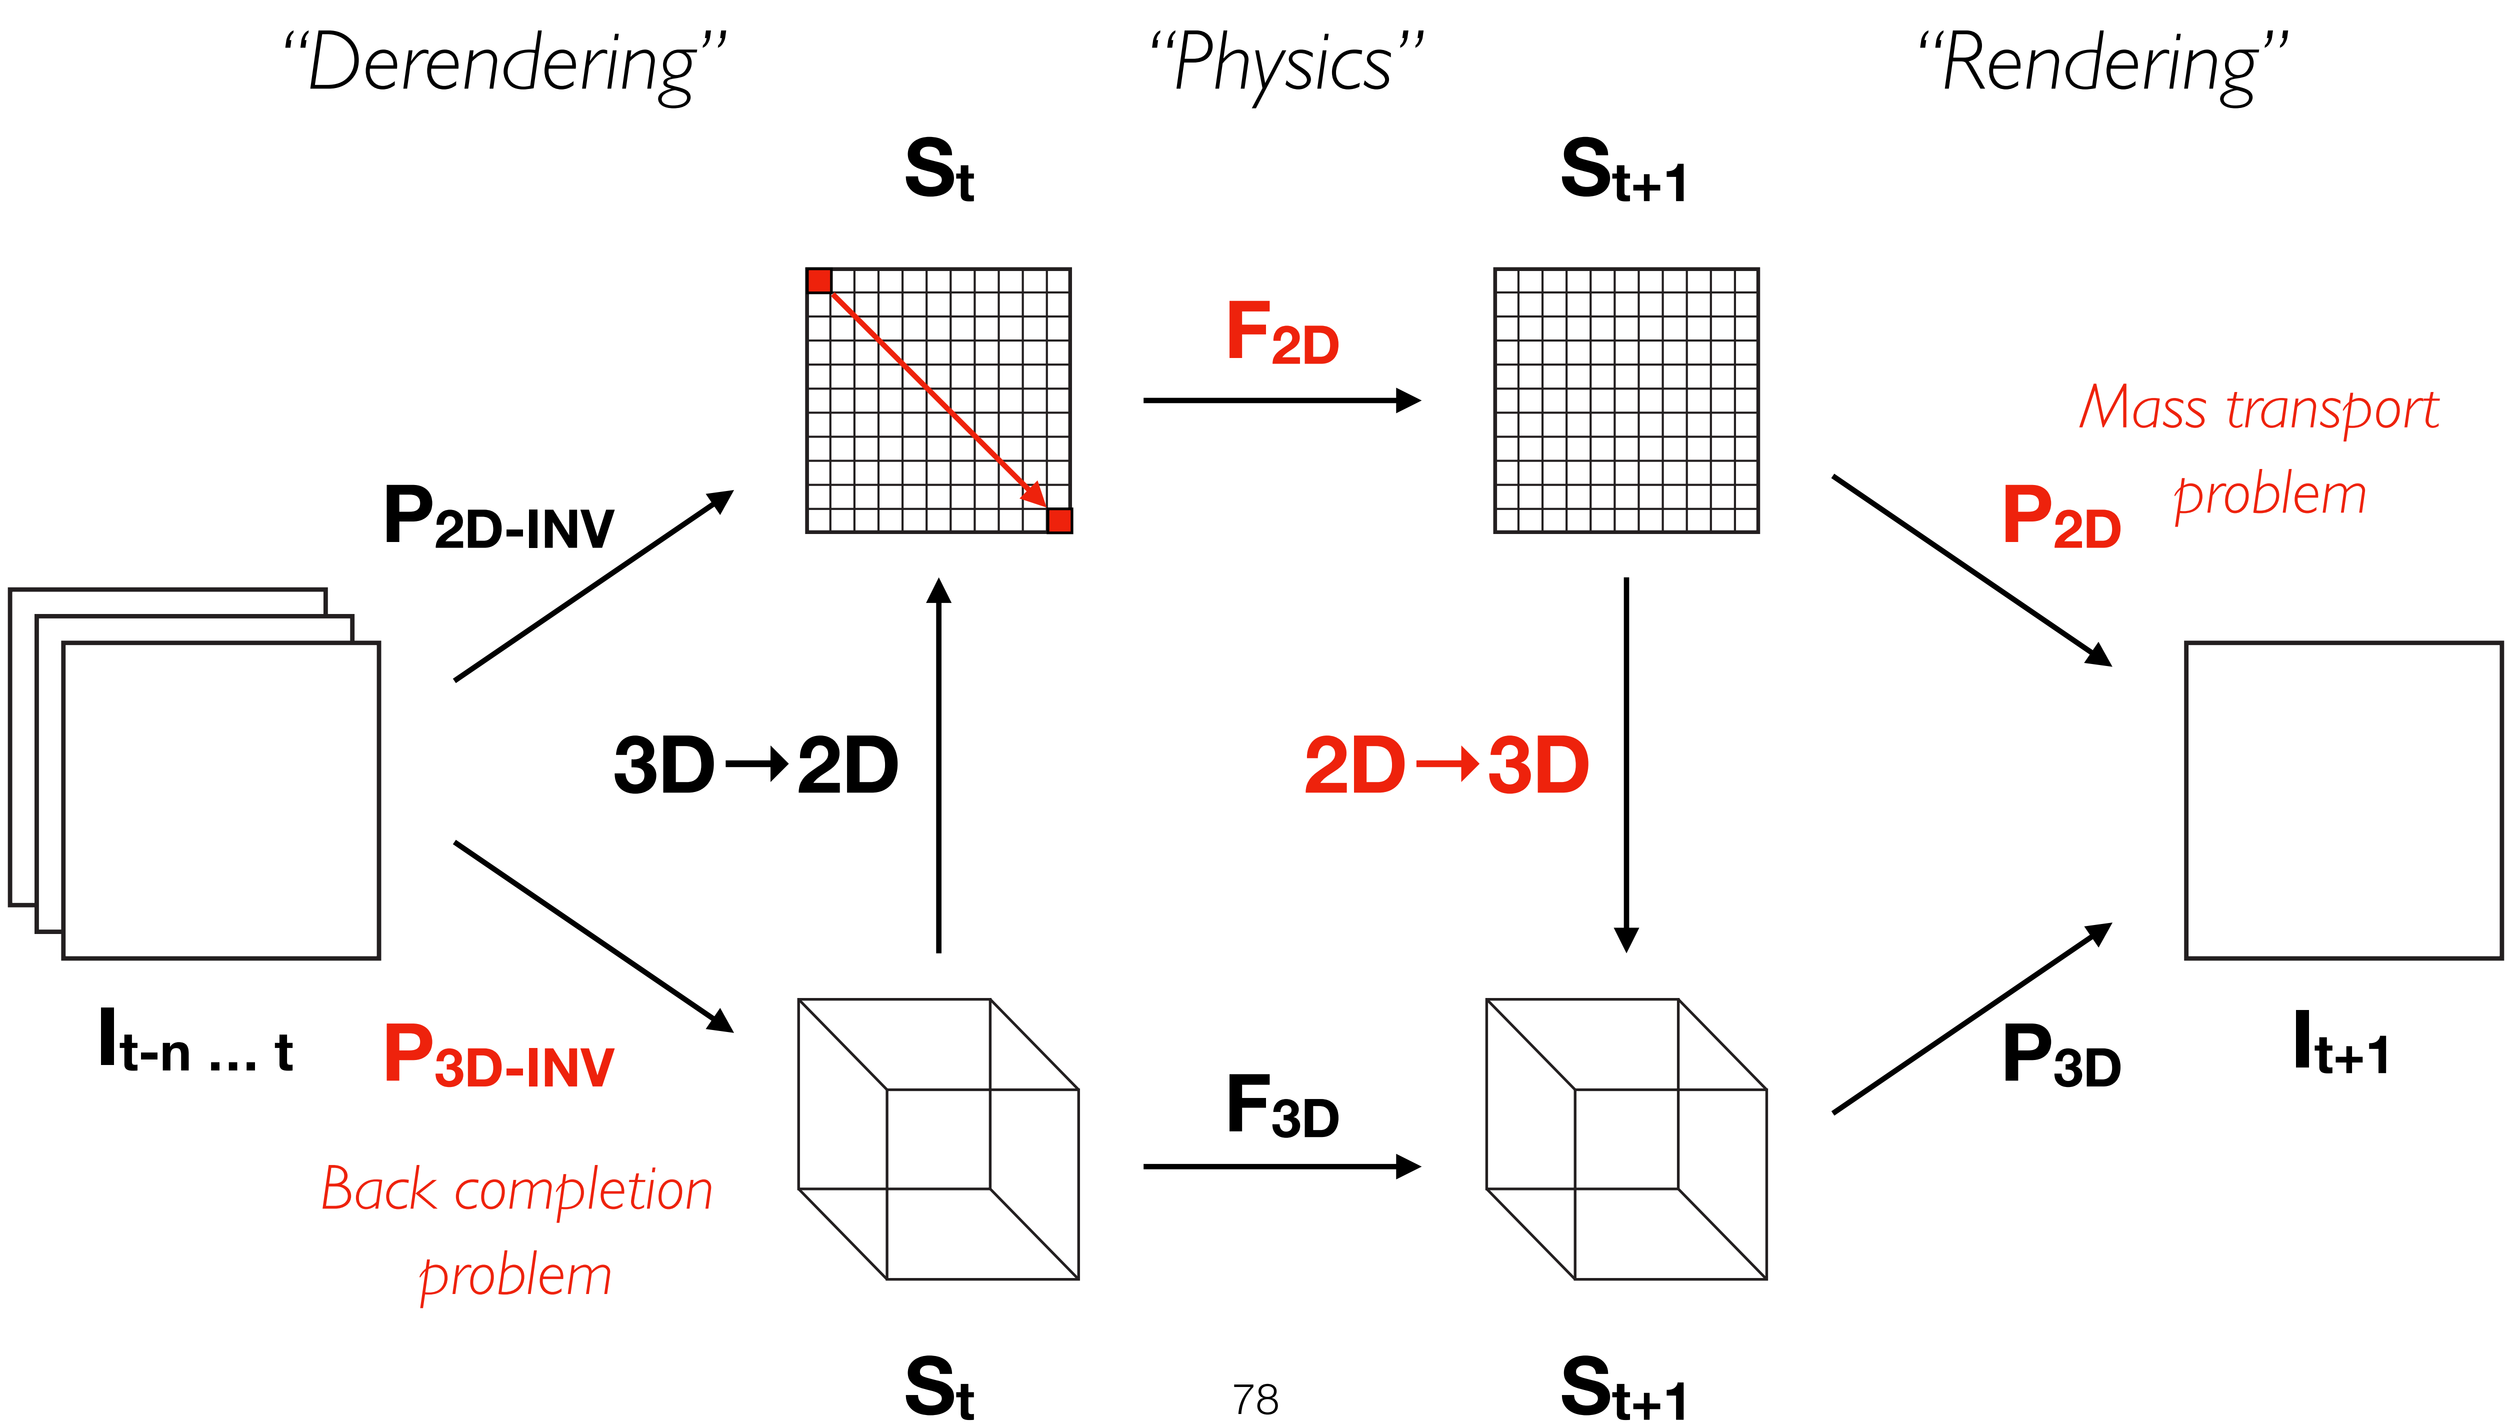

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early, perhaps by 3 months.

Alternative to spatially-uniform priors are **graph-based** priors

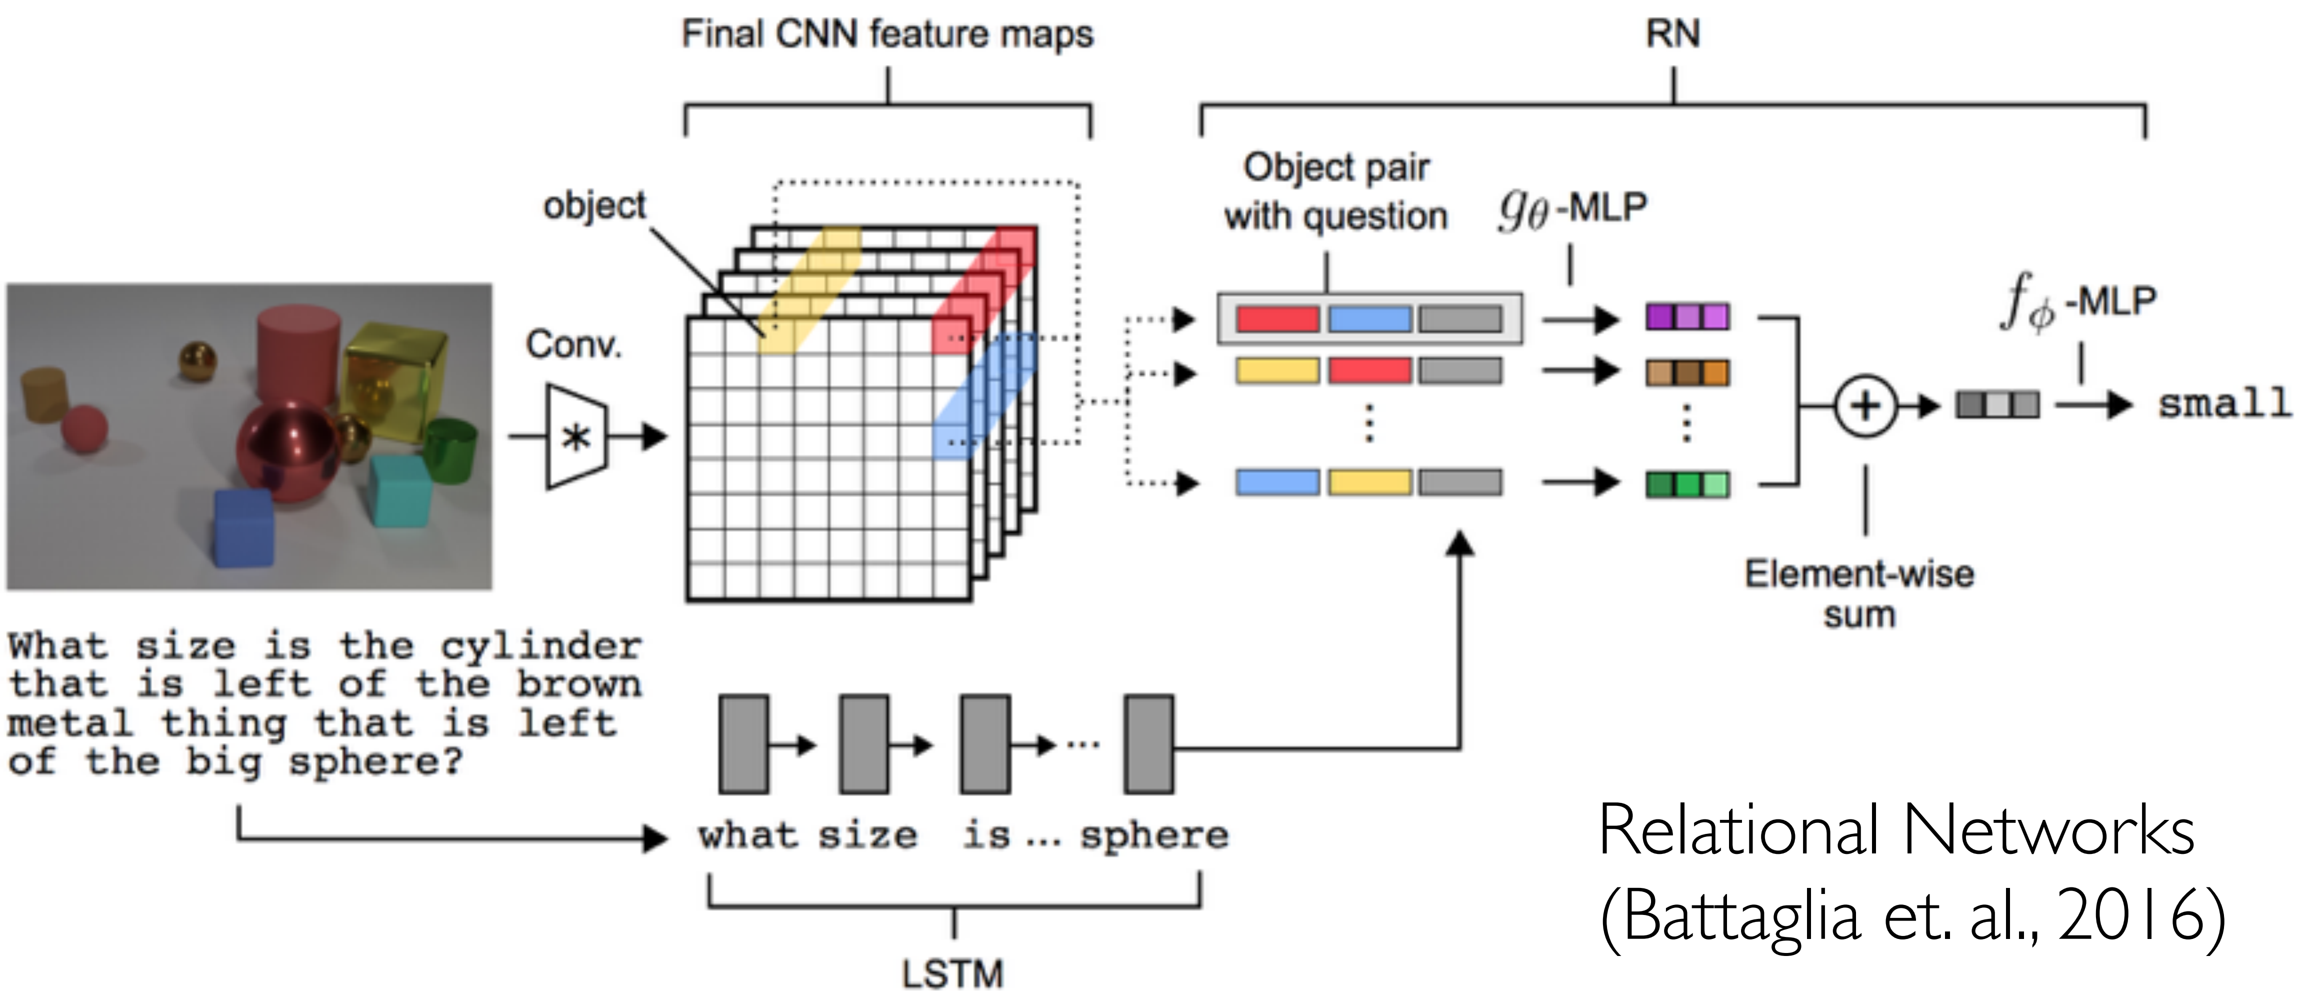

Relational Networks
(Battaglia et. al., 2016)

Intuitive Physics as Underlying Goal

Liz Spelke

Experimental results with infants: **object permanence** present very early, perhaps by 3 months.

Alternative to spatially-uniform priors are **graph-based** priors

… still local and convolutional, just on the graph.

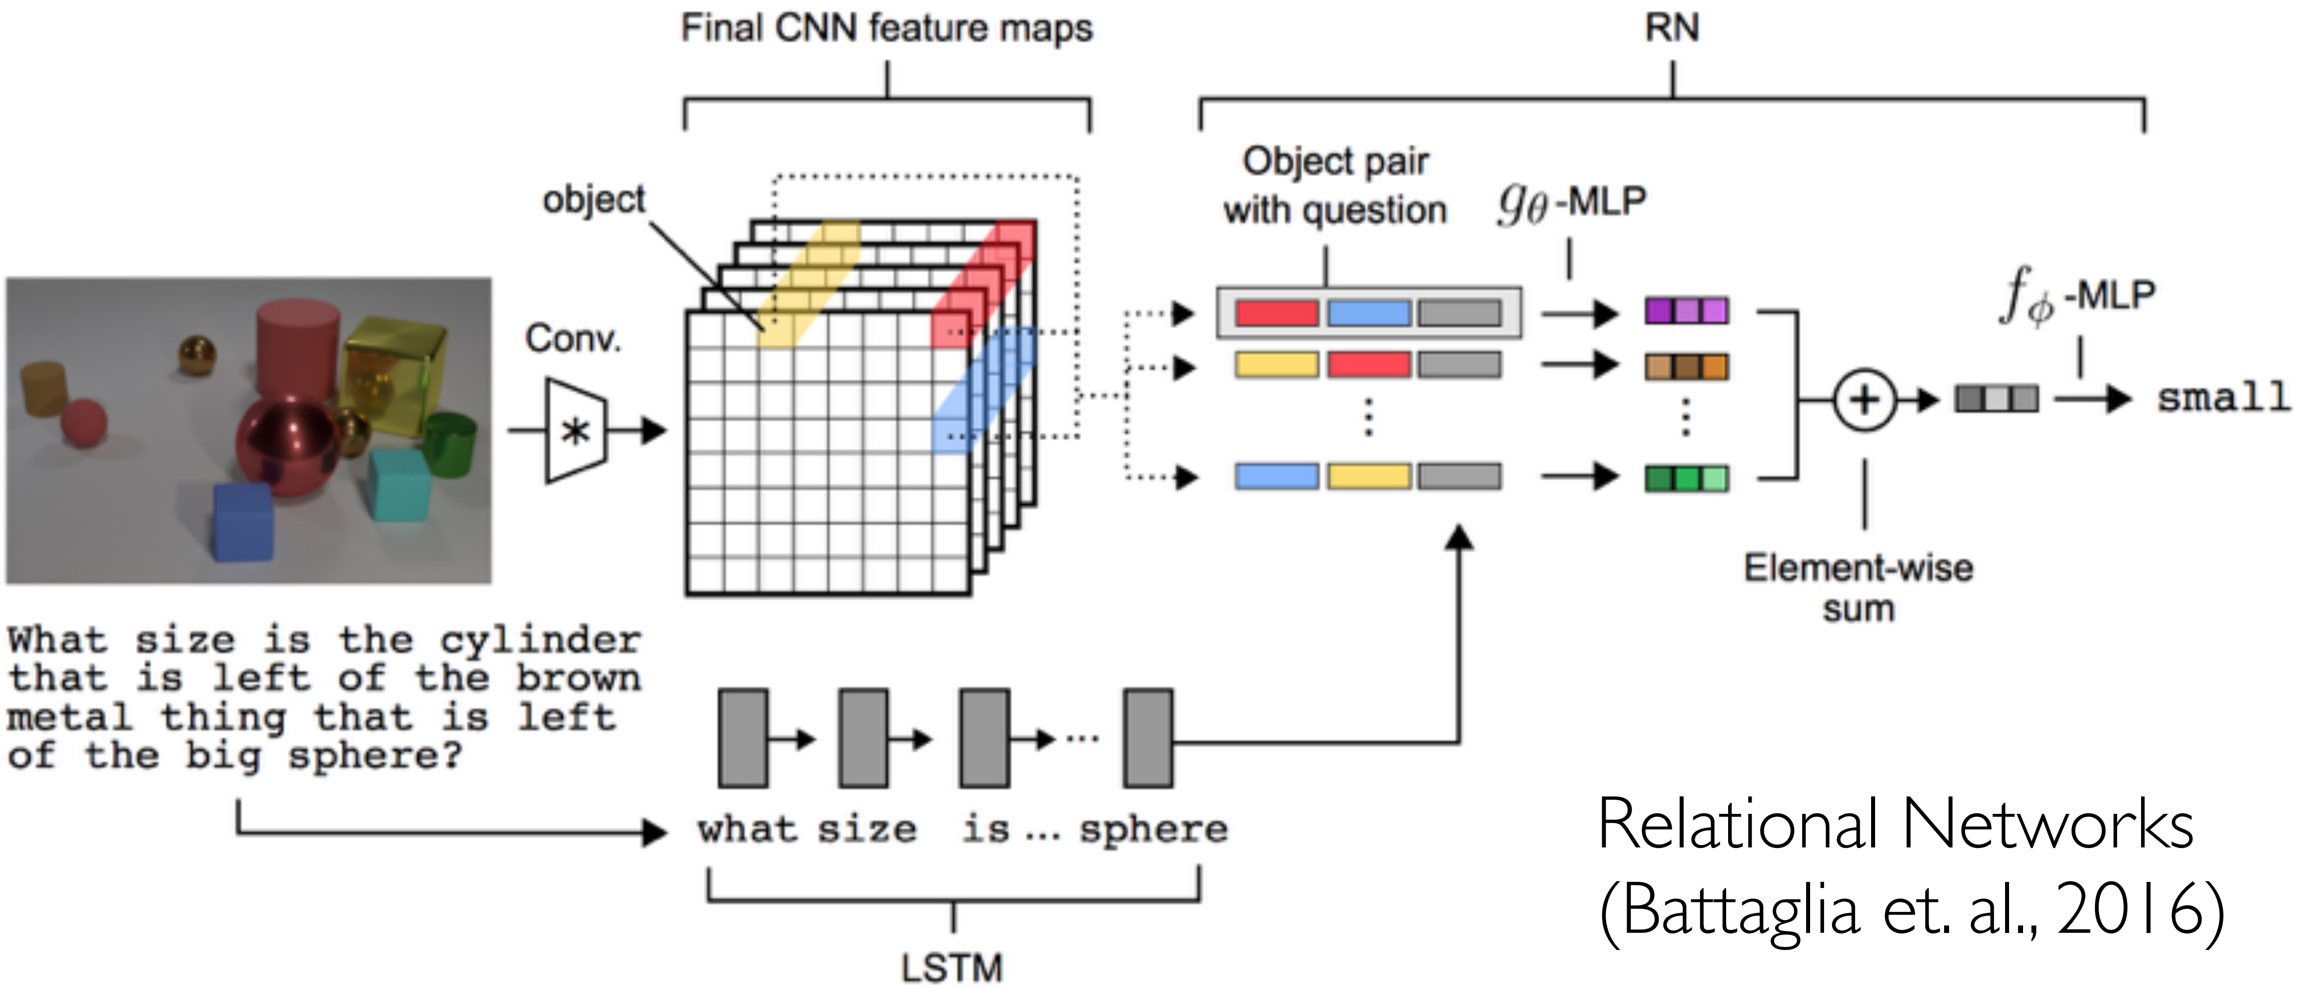

Relational Networks
(Battaglia et. al., 2016)

Relational Networks
(Battaglia et. al., 2016)

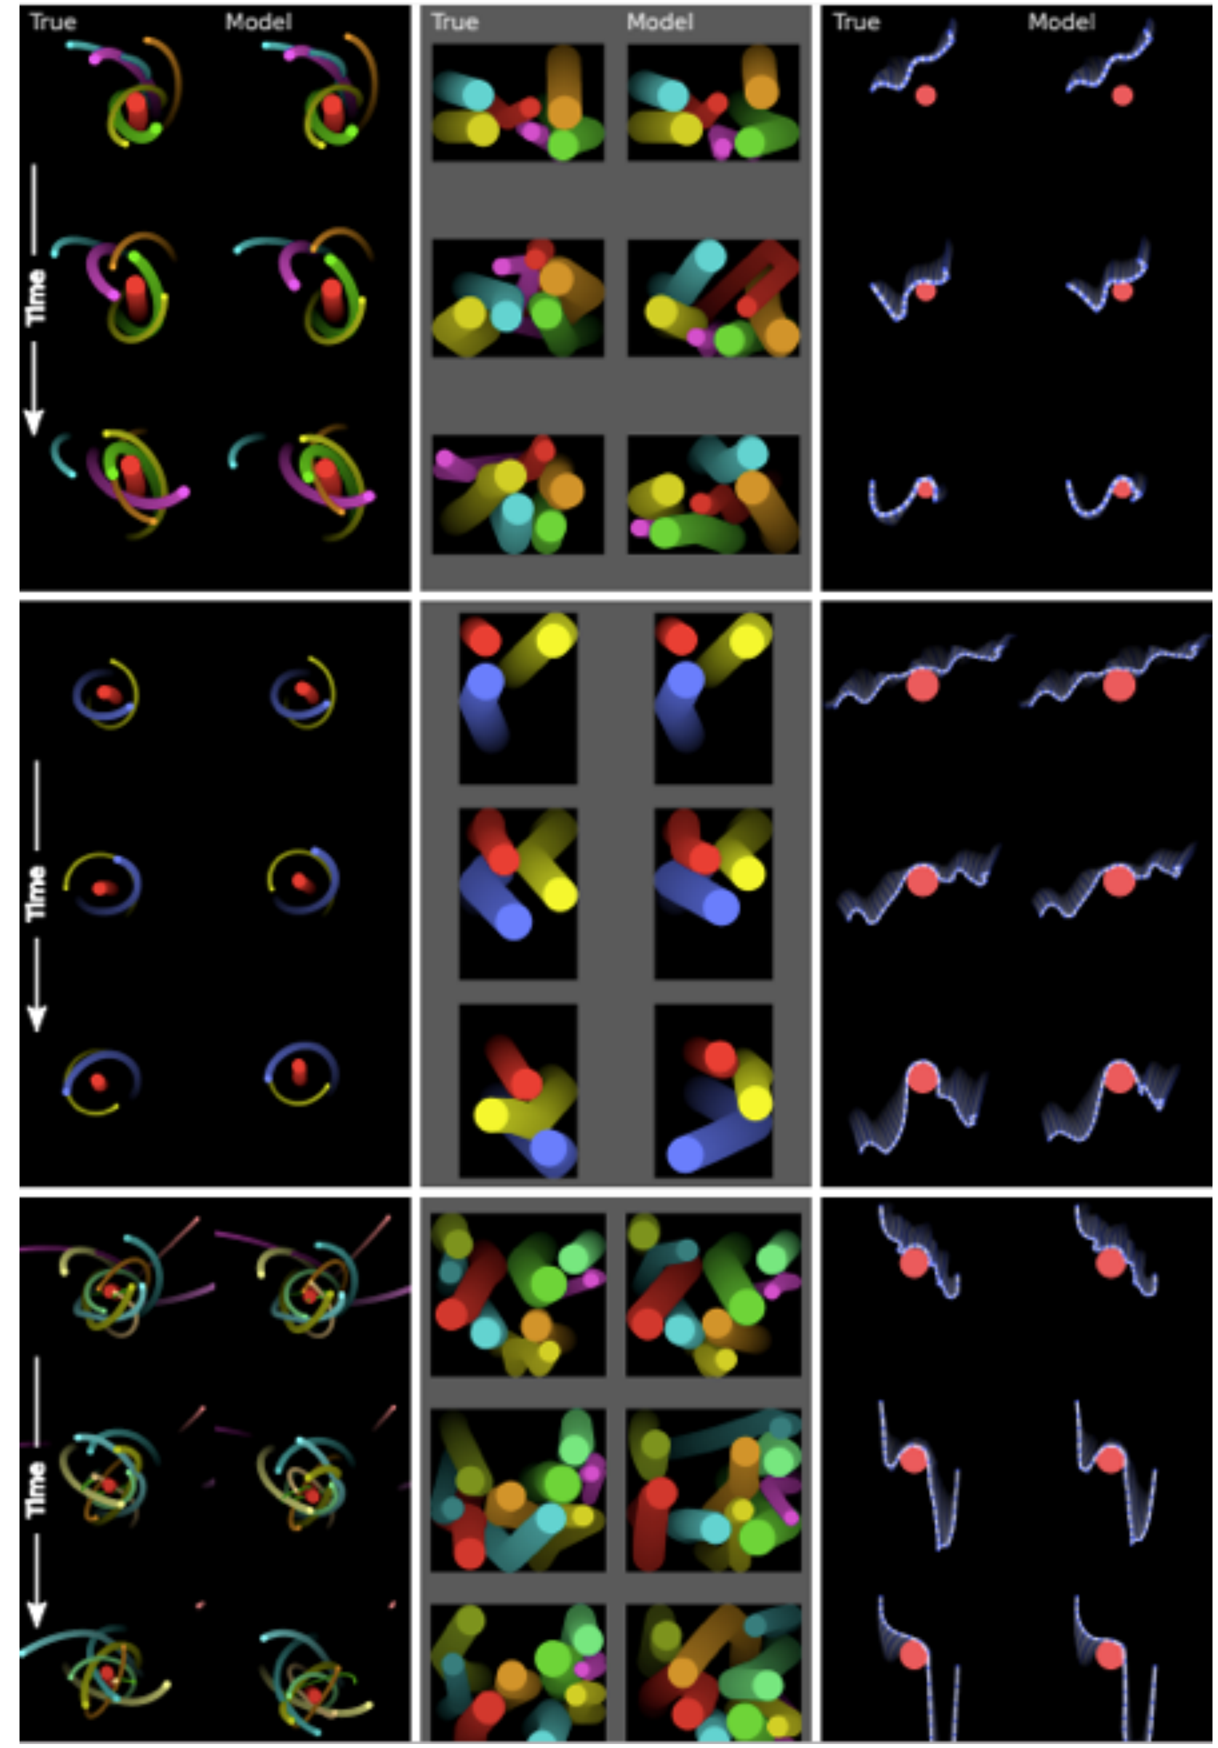

Neural Physics Engine
(Chang et. al., 2016)

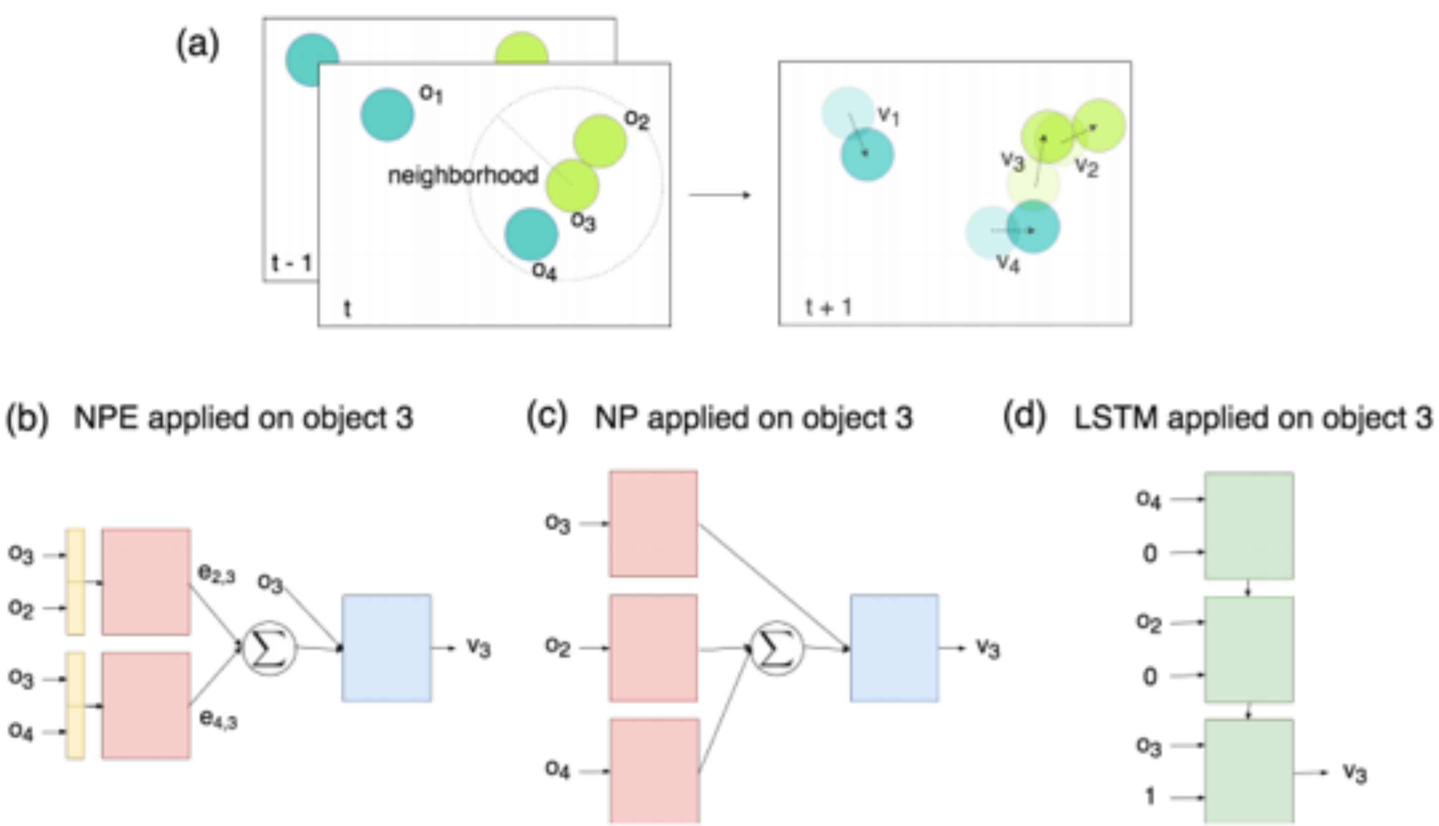

Intuitive Physics as Underlying Goal

Complex Scenes

Complex Materials

Describe objects through complex graphs:

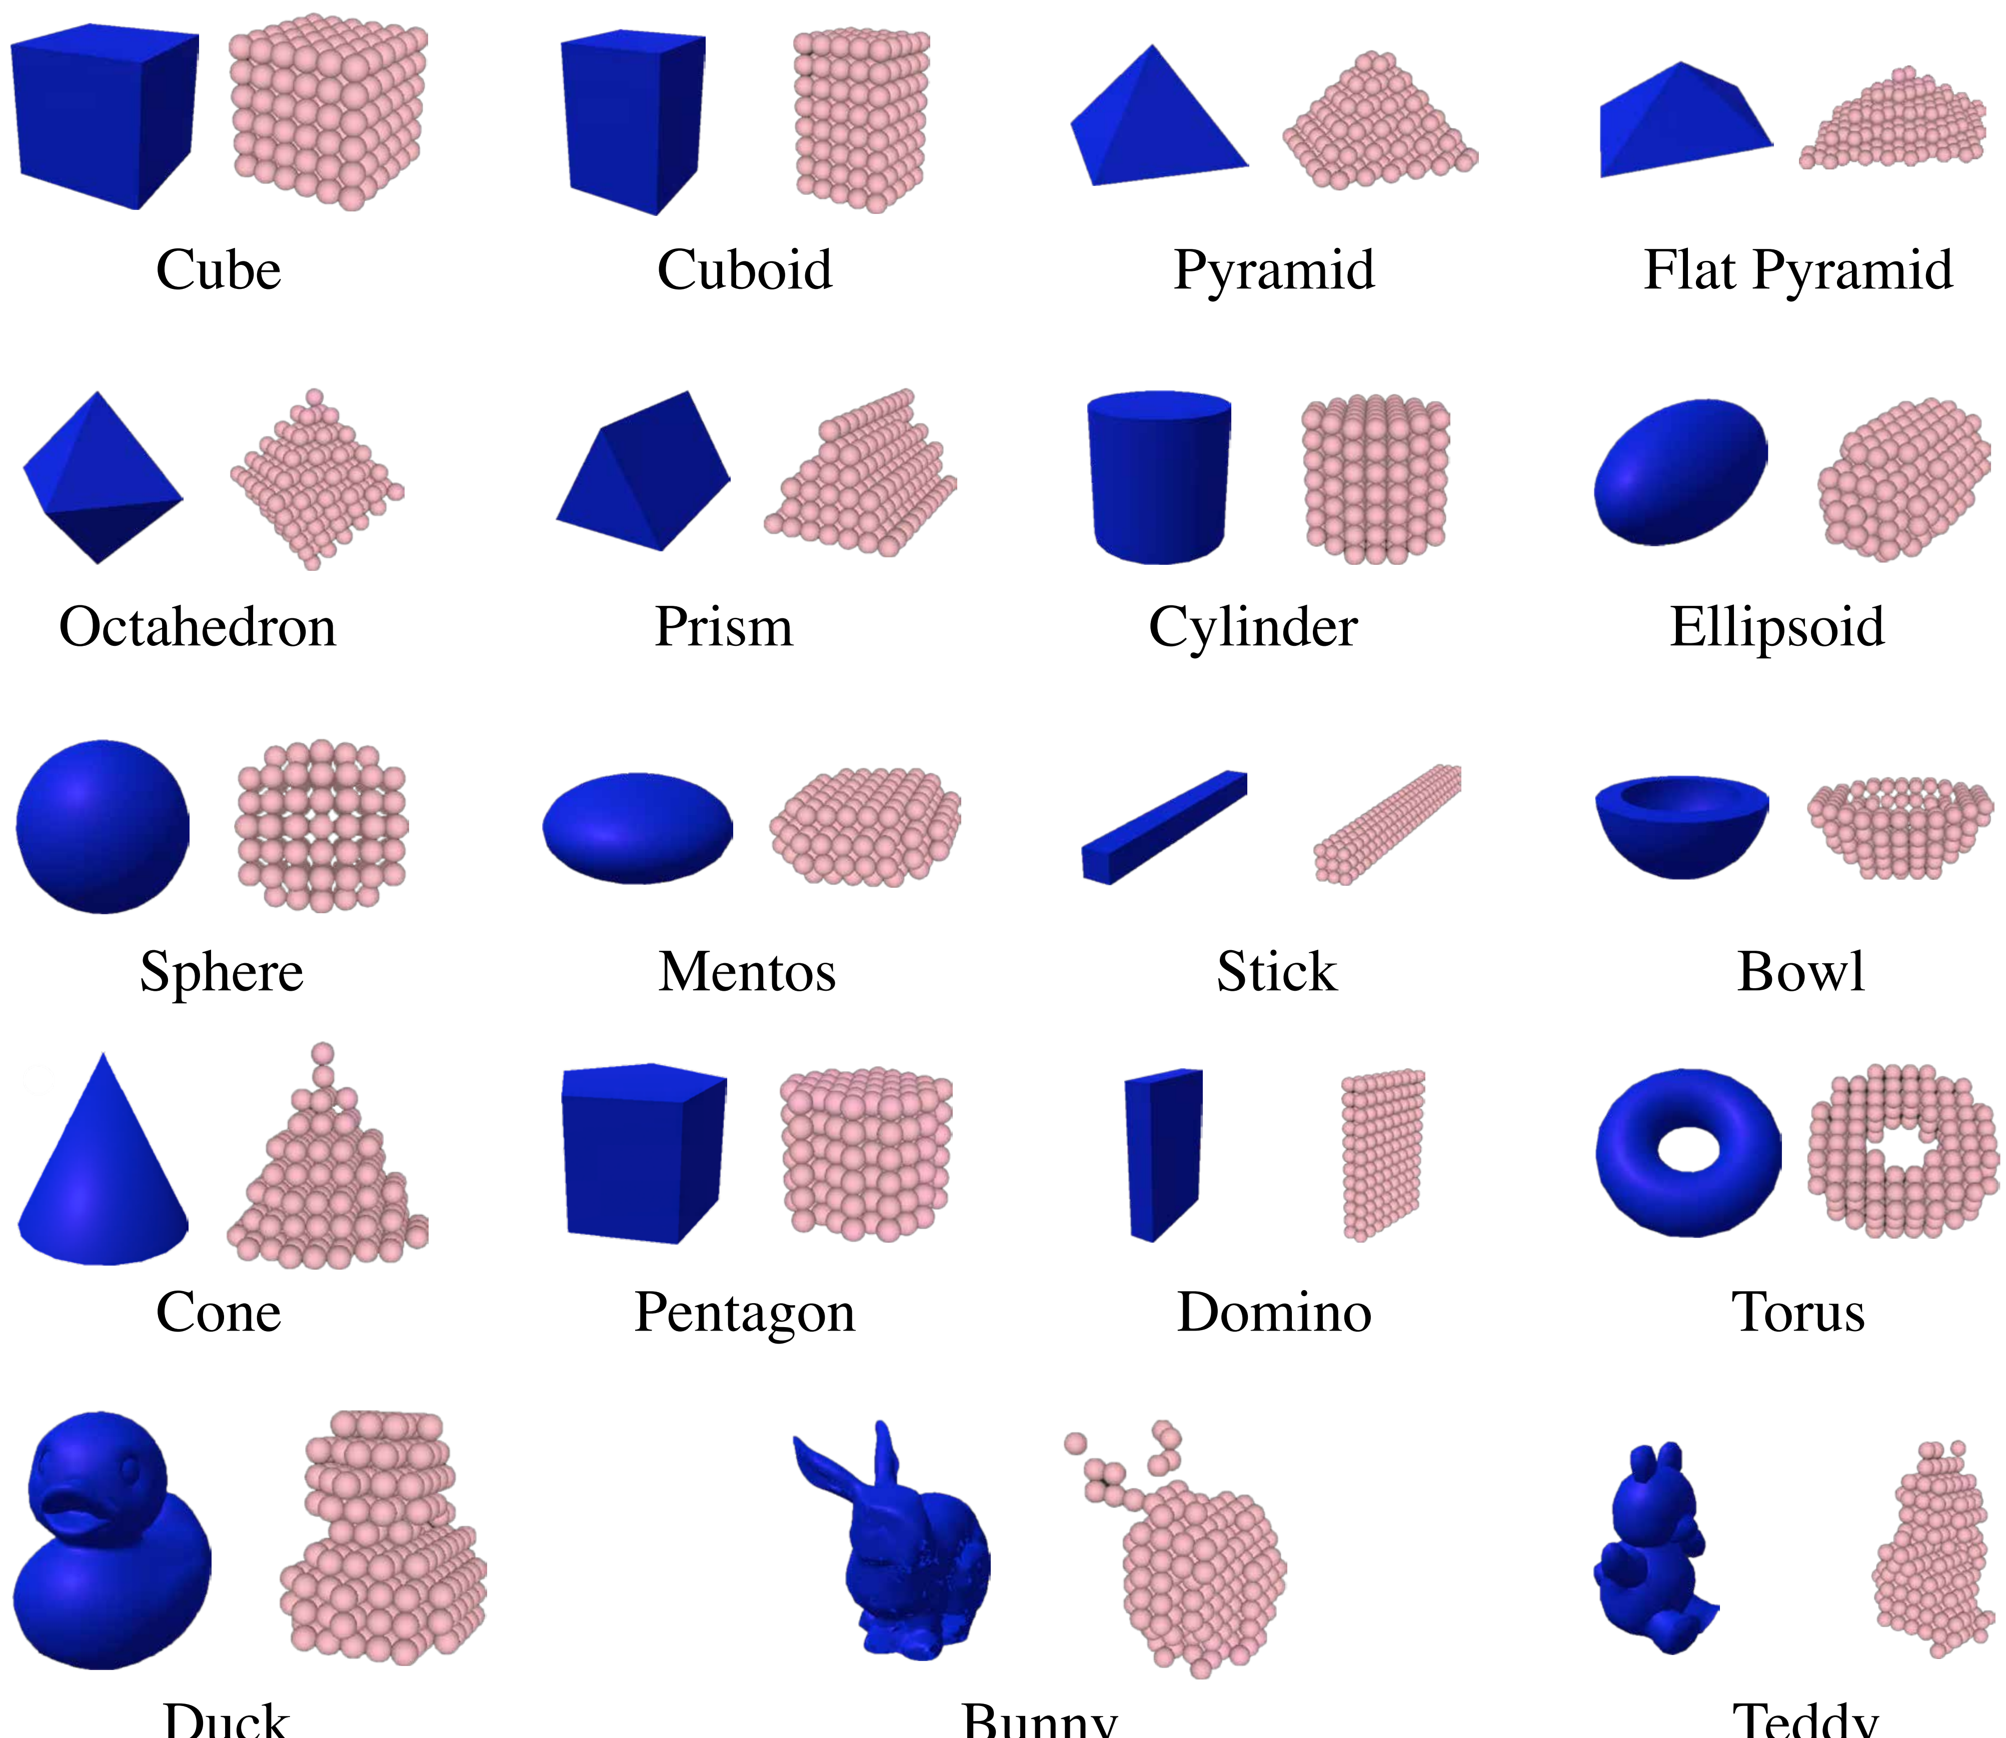

In fact, describe whole scenes.

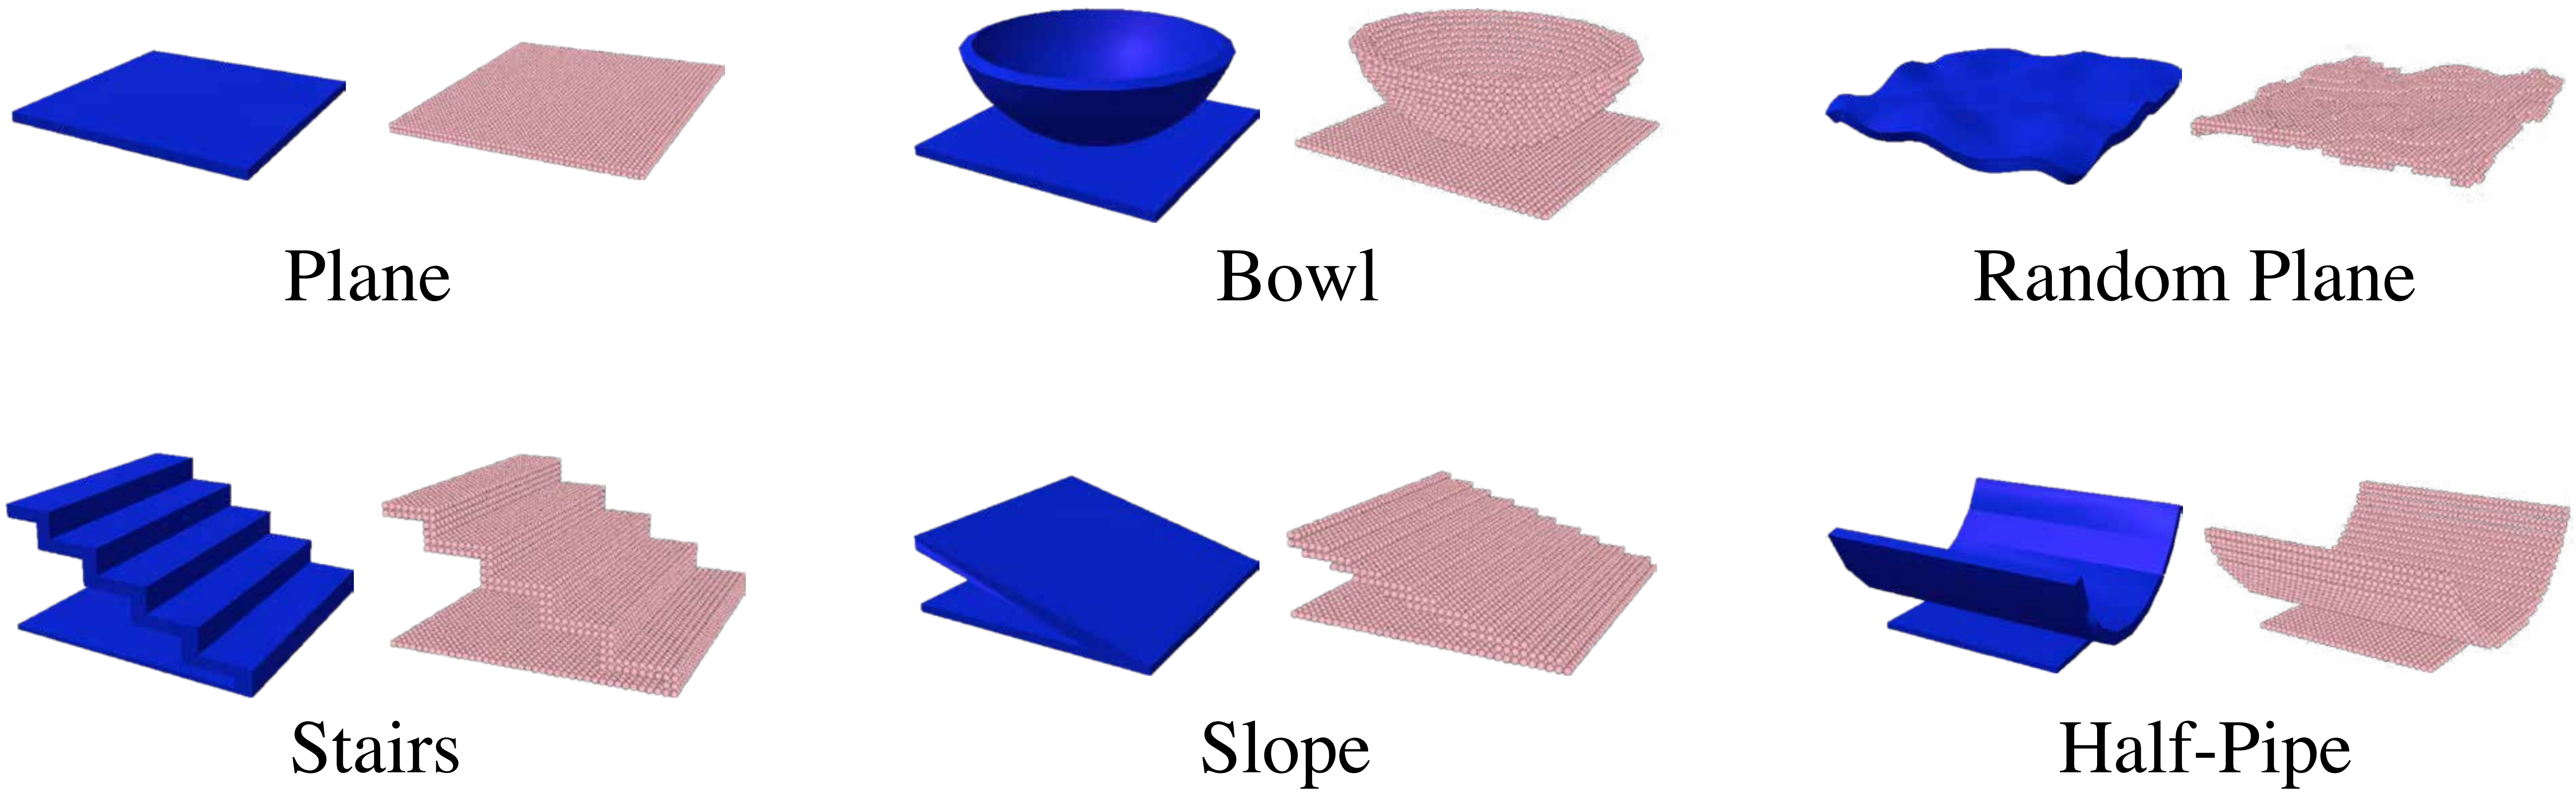

In fact, describe whole scenes.

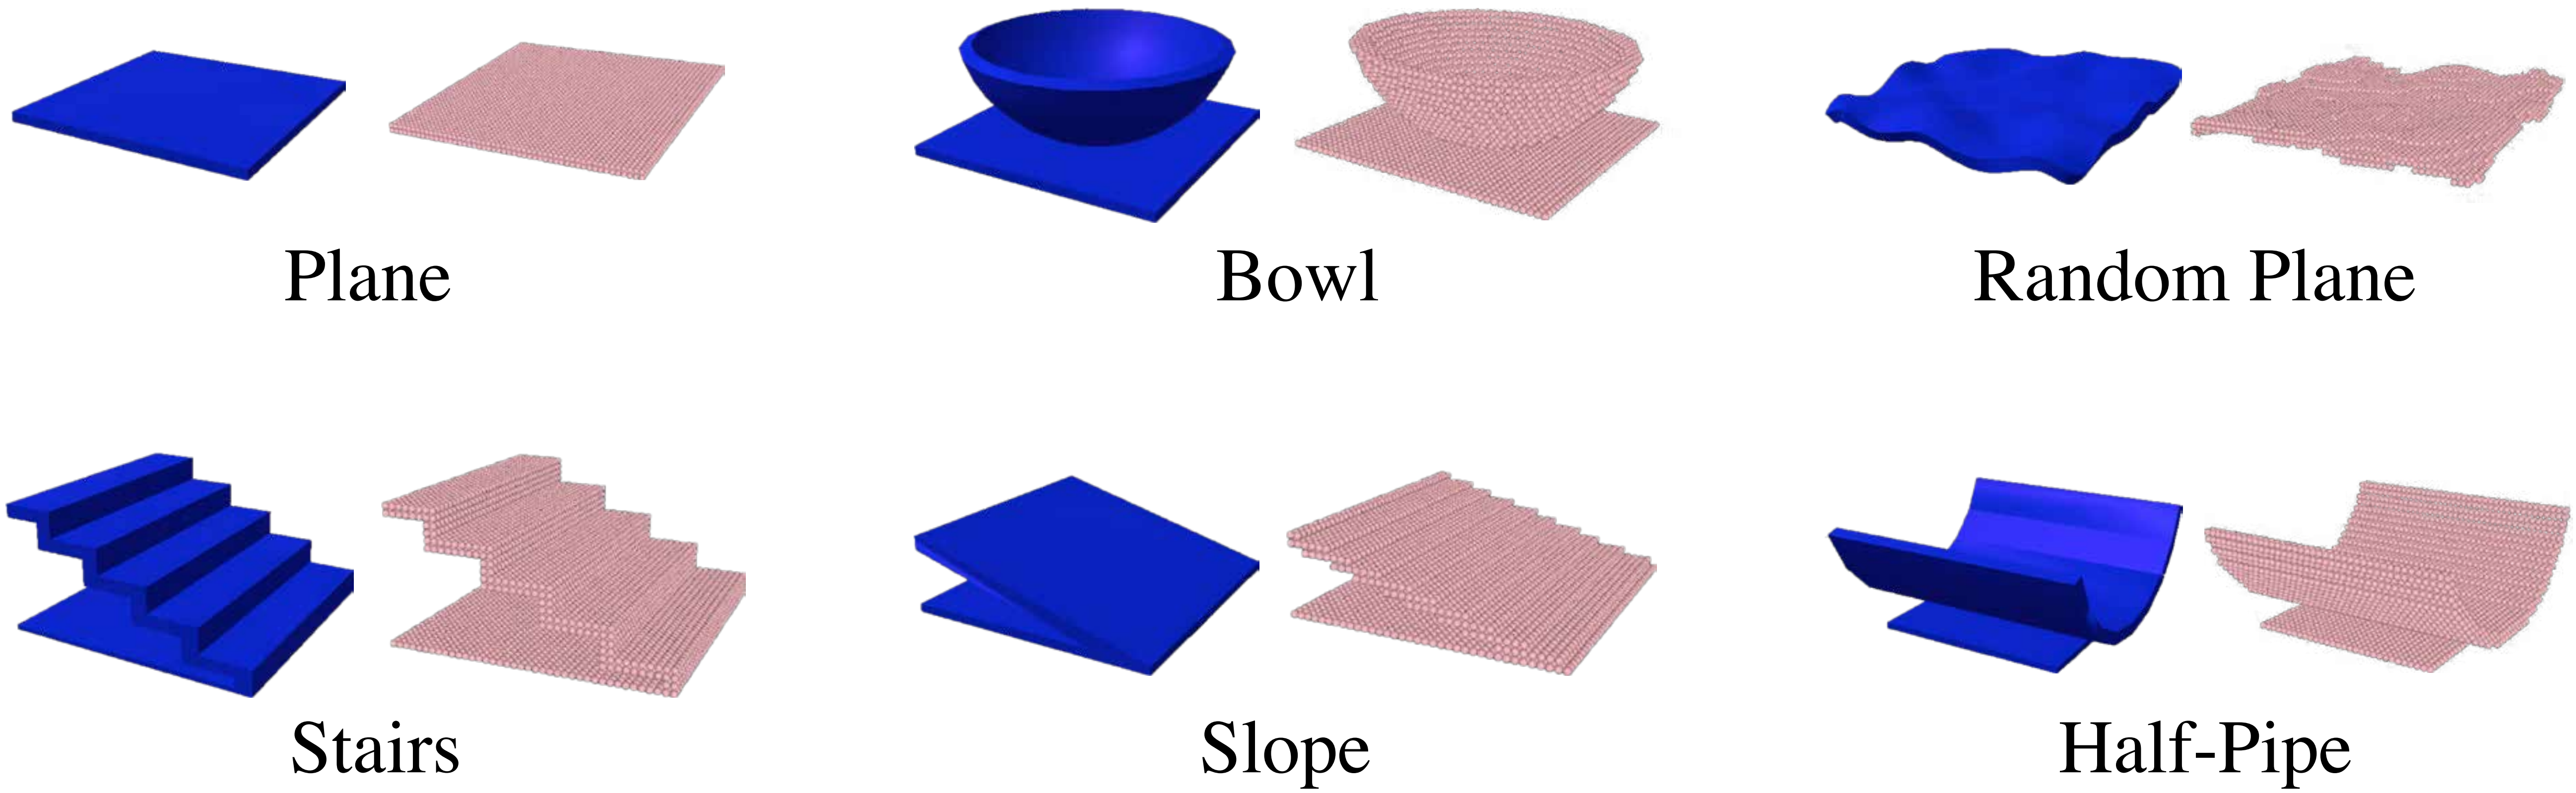

$$G = \langle N, E \rangle$$ scene graph

In fact, describe whole scenes.

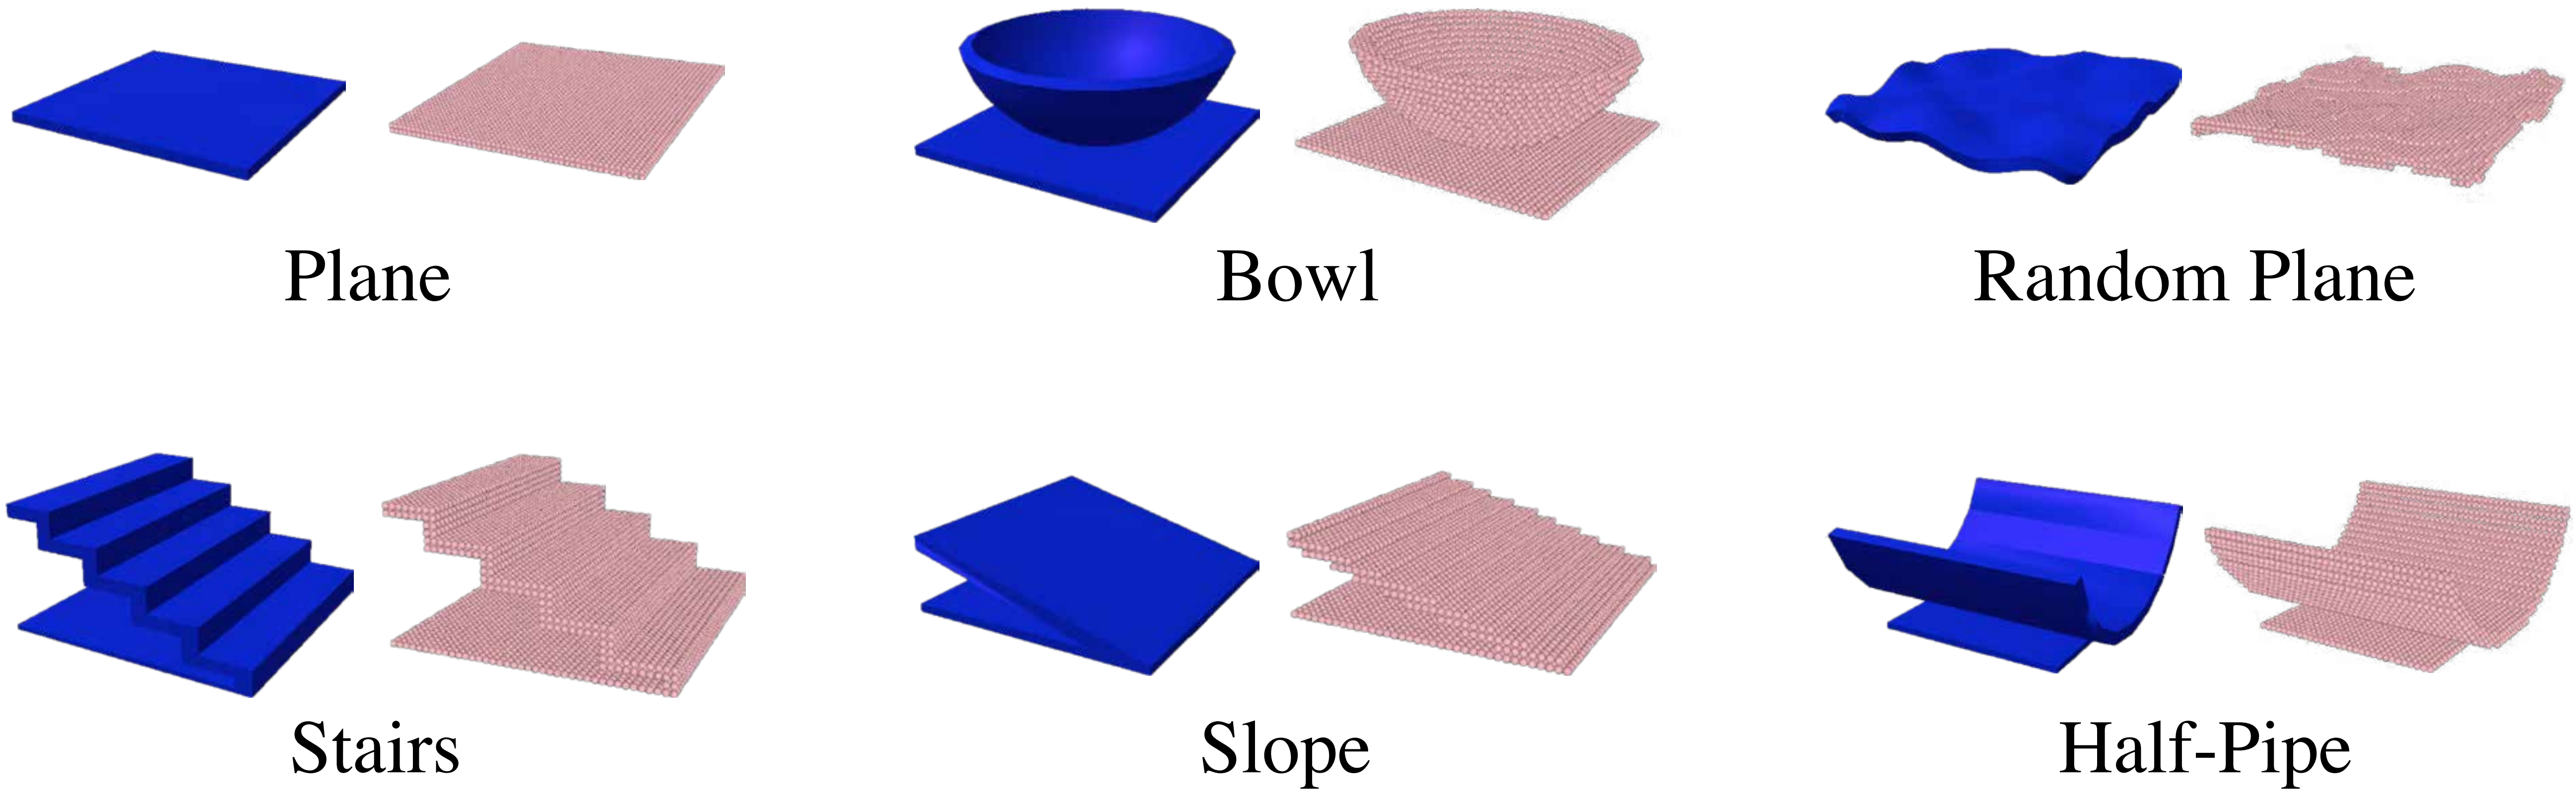

$$G = \langle N, E \rangle$$ scene graph

N = nodes corresponding to particles comprising objects

E = edges corresponding to relationships between particles

In fact, describe whole scenes.

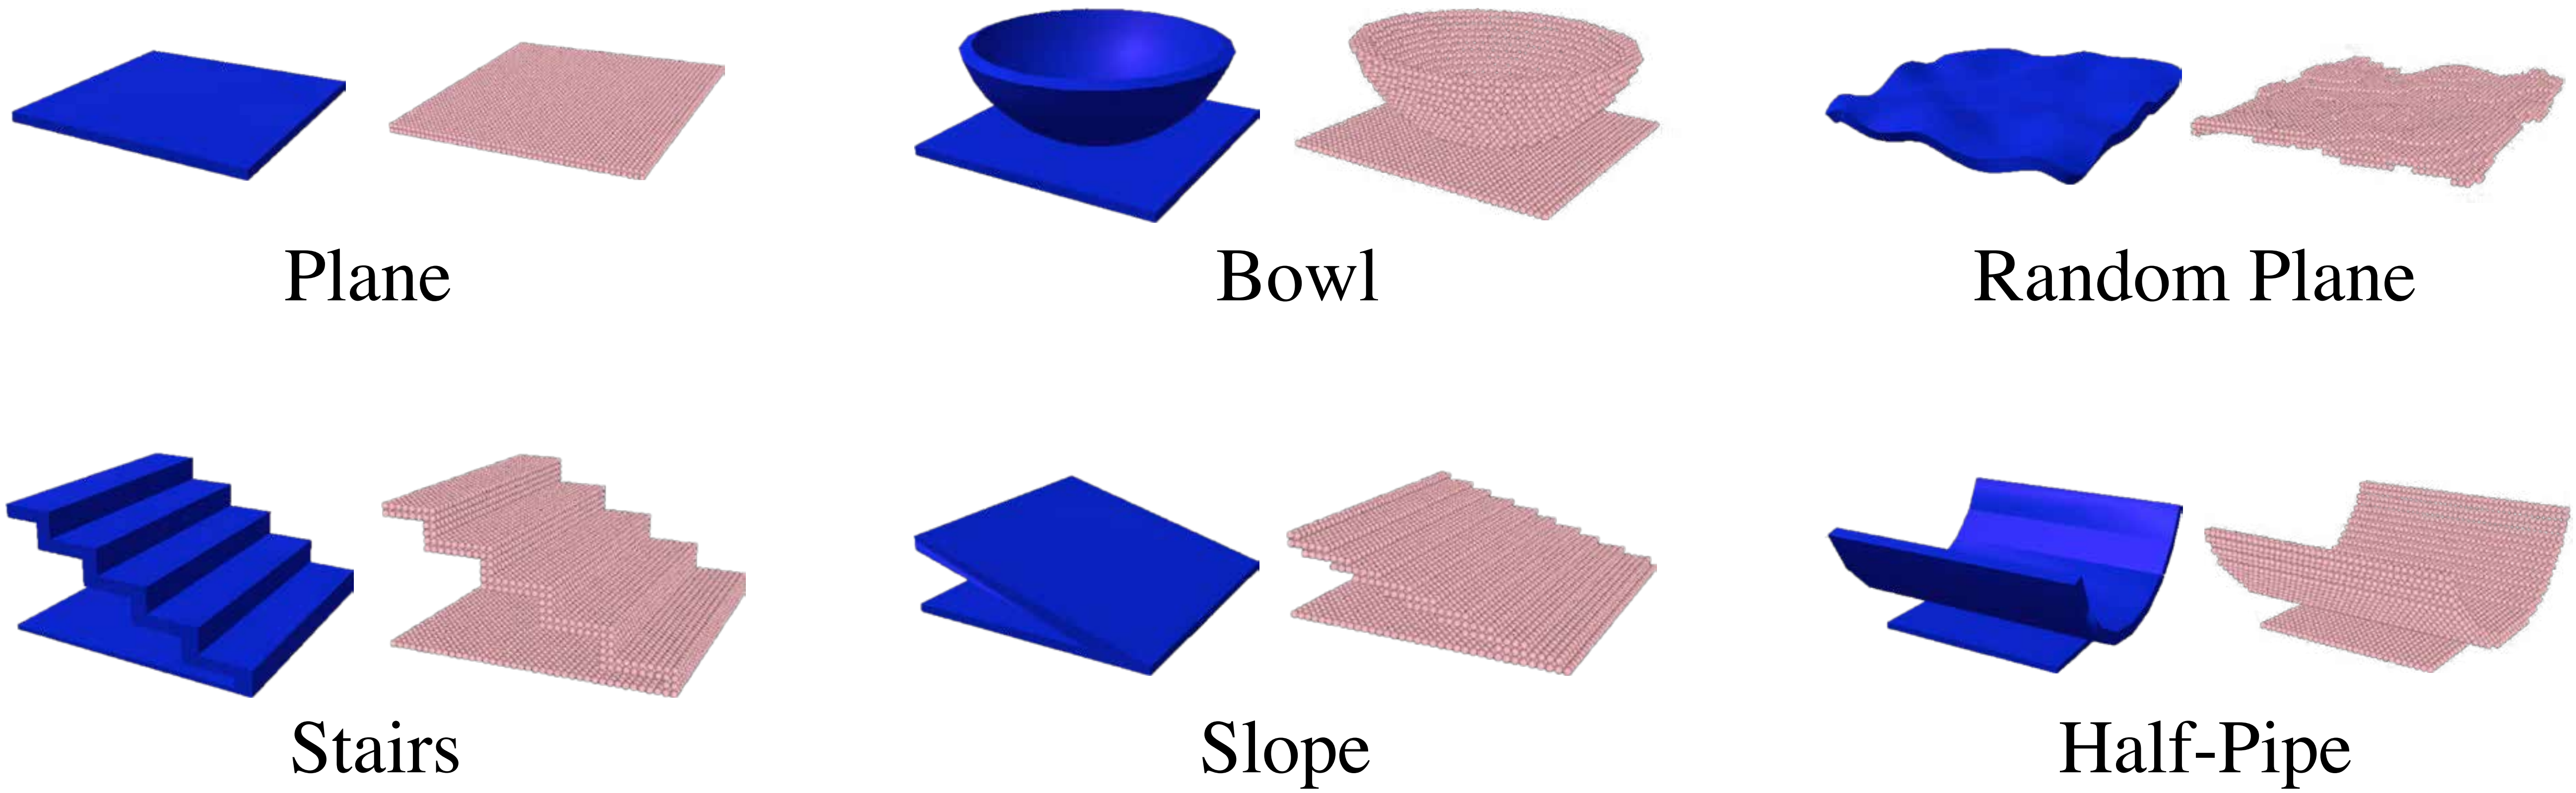

$$G = \langle N, E \rangle$$ scene graph

N = nodes corresponding to particles comprising objects

E = edges corresponding to relationships between particles

edges are labelled by vector capturing bond characteristics

Intuitive Physics as Underlying Goal

Of course, humans don't think about all the particles at once all the time.

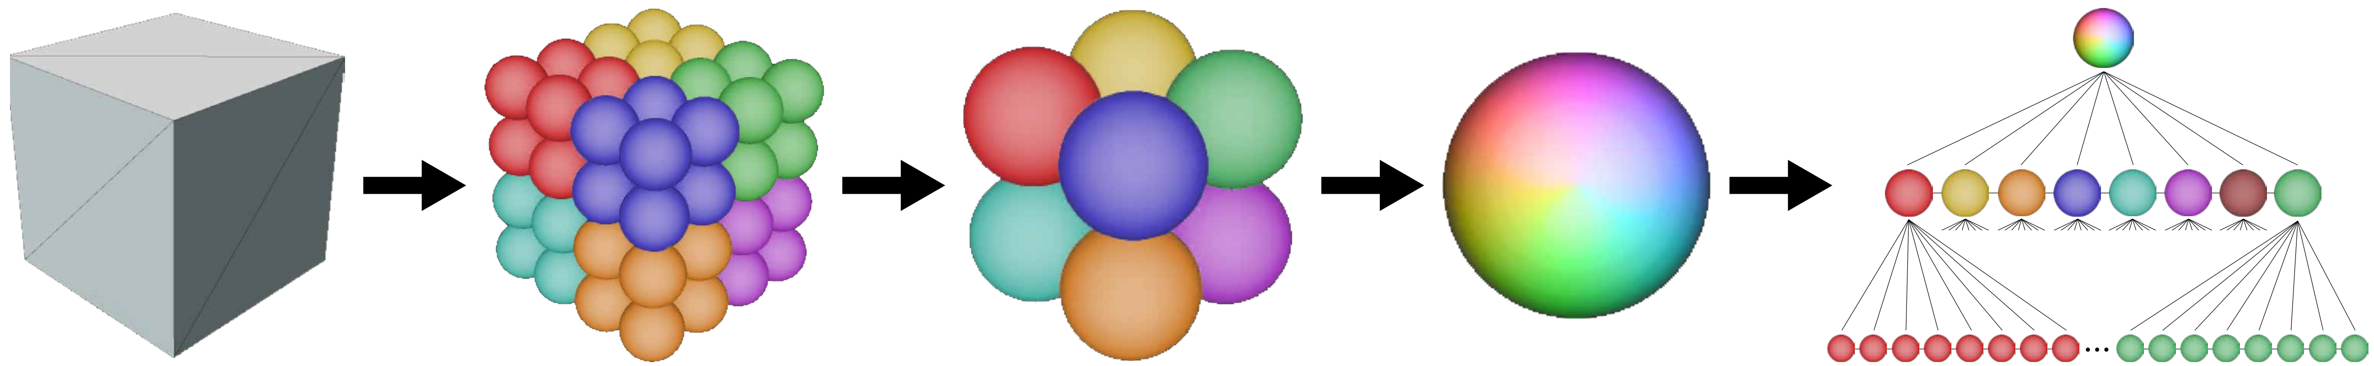

Intuitive Physics as Underlying Goal

Of course, humans don't think about all the particles at once all the time.

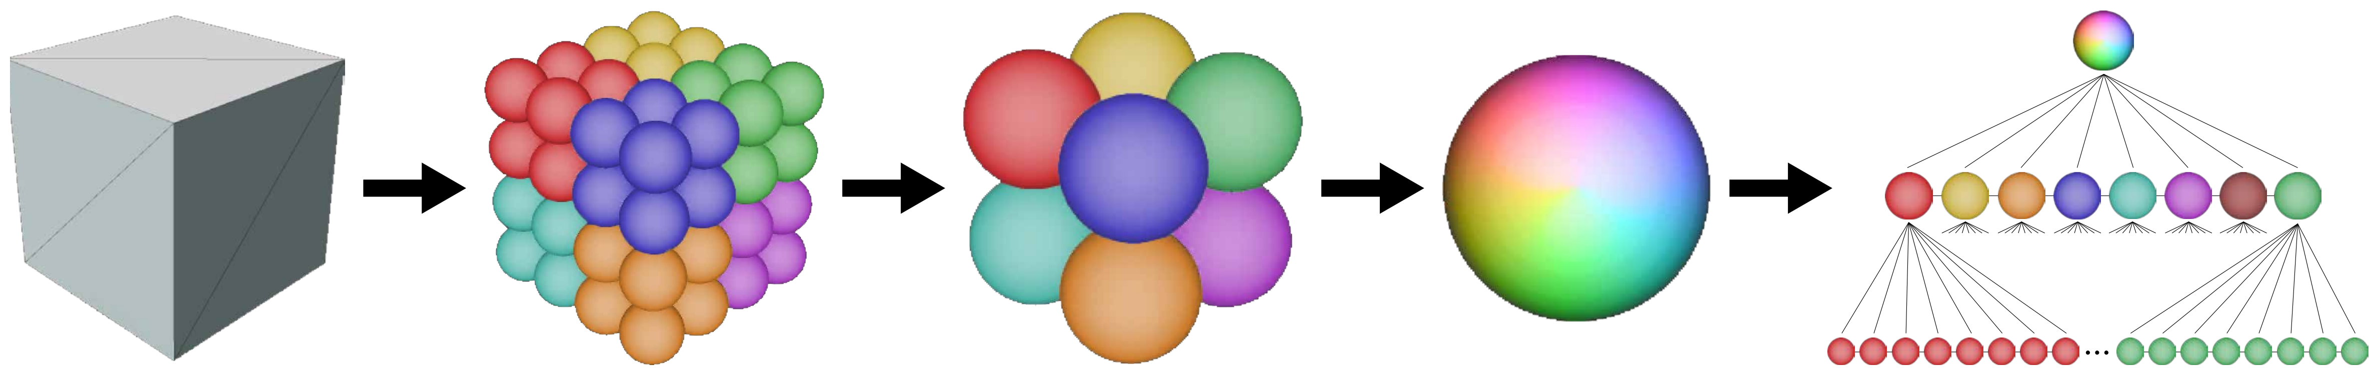

$$G \longmapsto G_H$$

G_H = dynamic "hierarchicalization" of underlying scene graph

(right now computed via k-means)

Intuitive Physics as Underlying Goal

Of course, humans don't think about all the particles at once all the time.

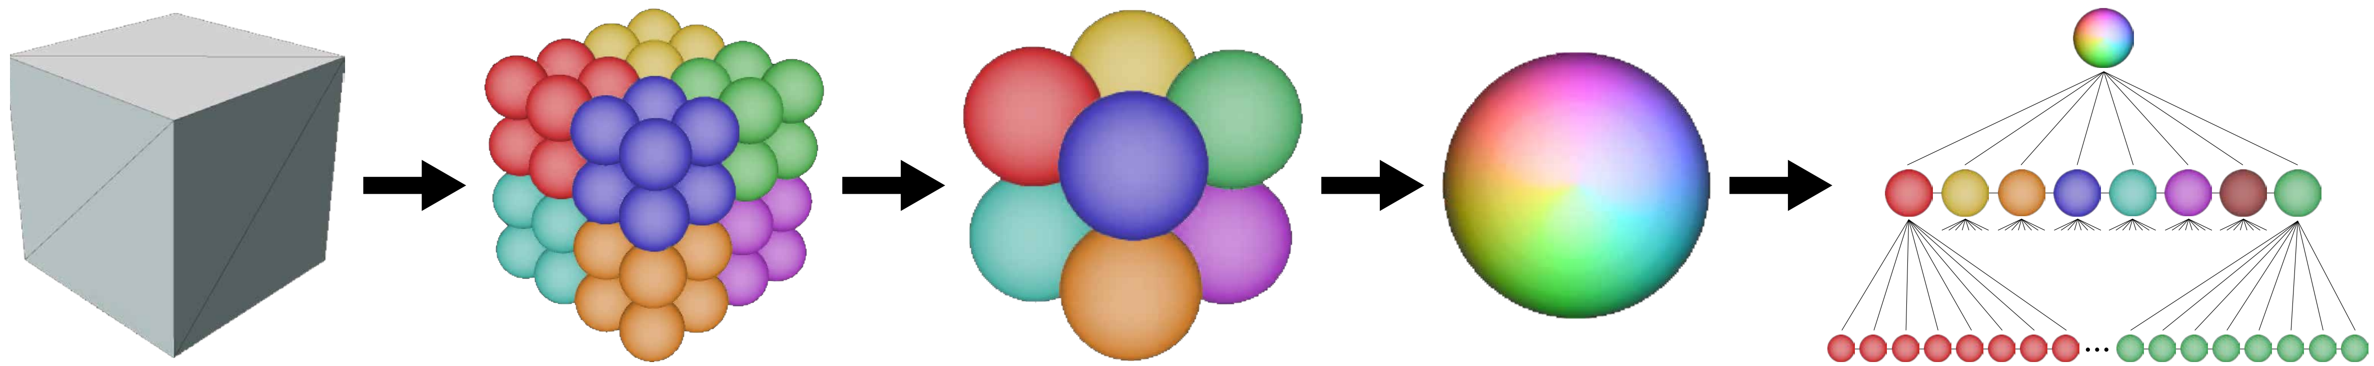

$$G \longmapsto G_H$$

G_H = dynamic "hierarchicalization" of underlying scene graph

(right now computed via k-means)

graph convolution → hierarchical graph convolution

Intuitive Physics as Underlying Goal

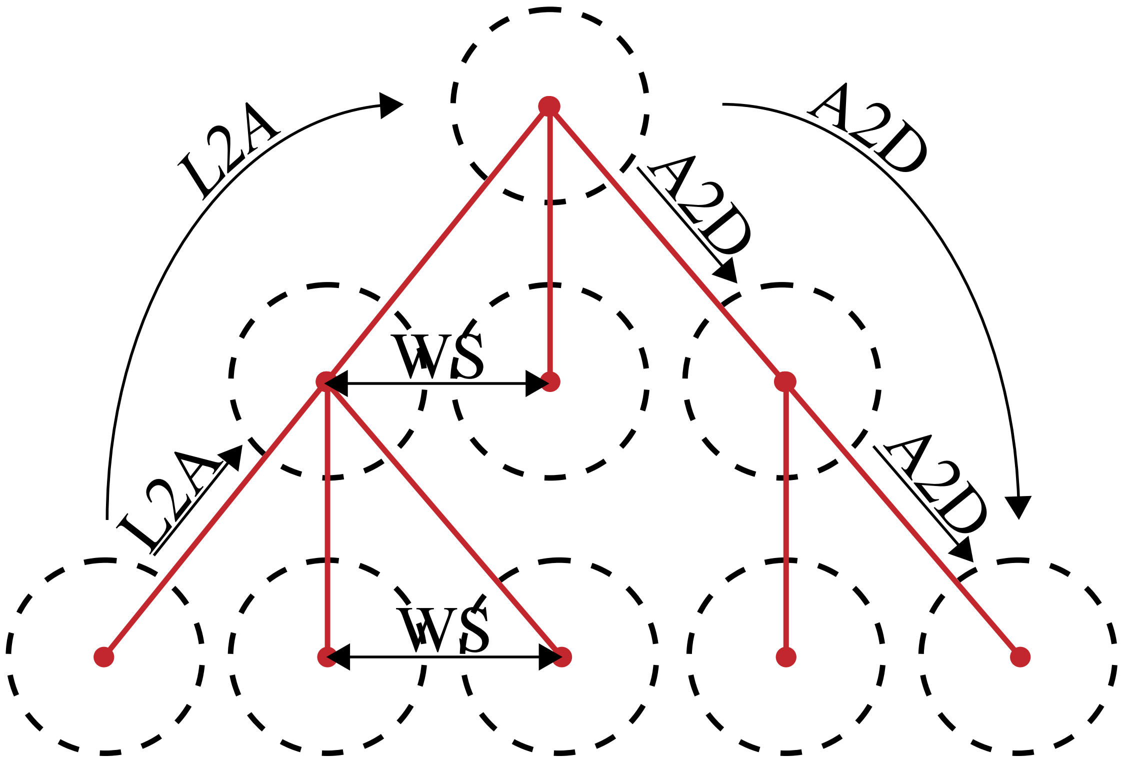

ϕ^{L2A} graph conv. leaves to ancestors

ϕ^{WS} graph conv. with siblings

ϕ^{A2D} graph conv. ancestors to descendants

Intuitive Physics as Underlying Goal

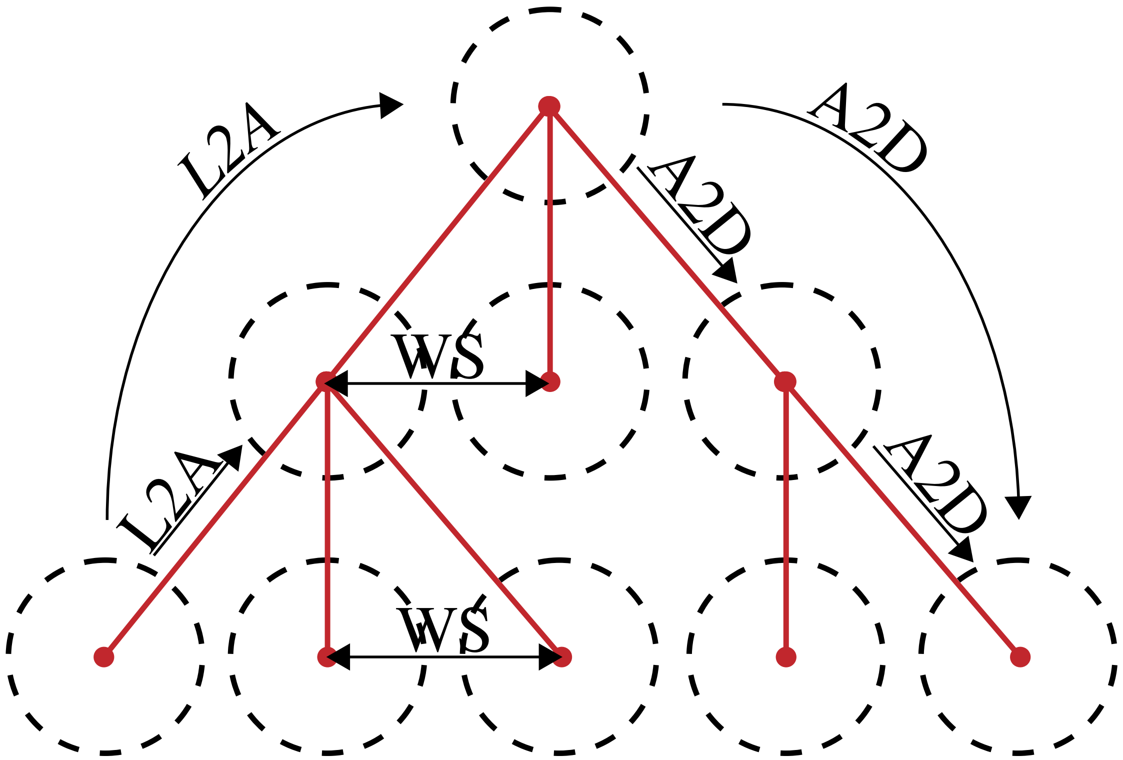

ϕ^{L2A} graph conv. leaves to ancestors

ϕ^{WS} graph conv. with siblings

ϕ^{A2D} graph conv. ancestors to descendants

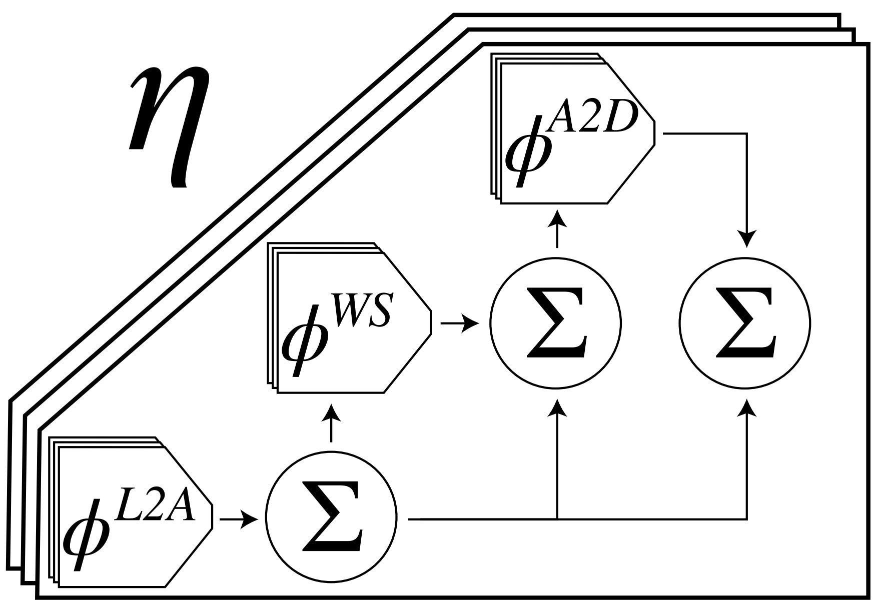

η module composing these three operations from one up-down cycle, adding physical effects

Intuitive Physics as Underlying Goal

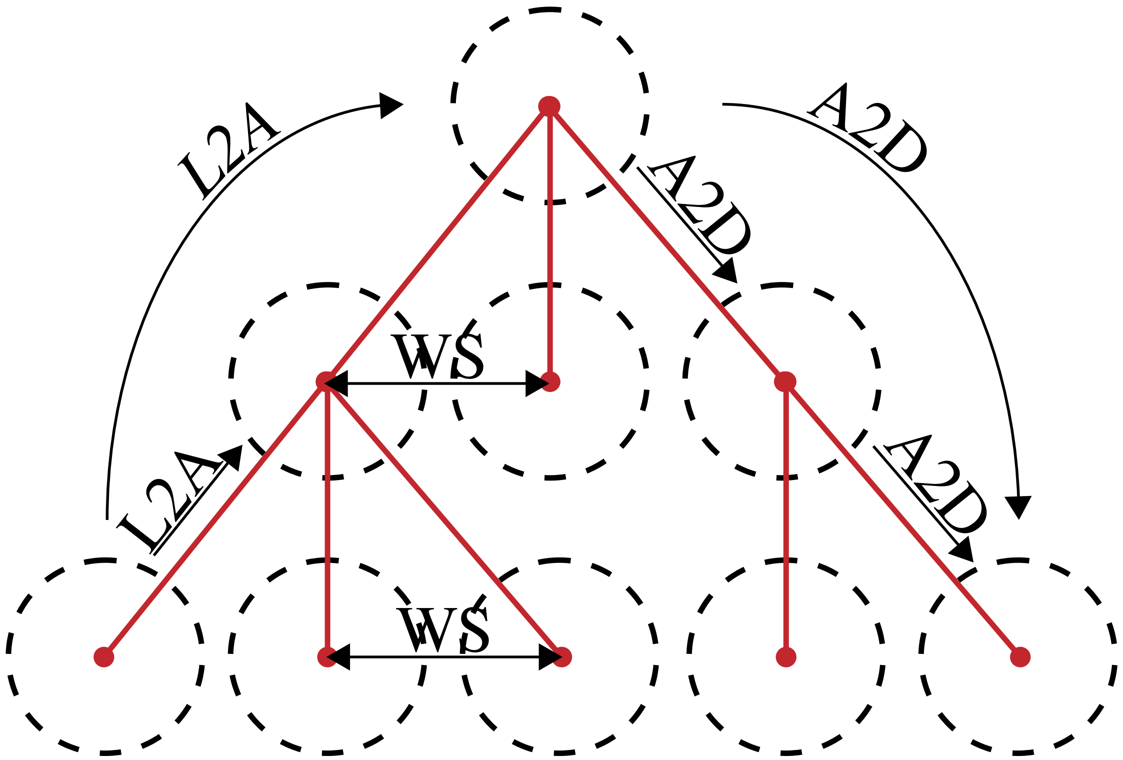

ϕ^{L2A} graph conv. leaves to ancestors

ϕ^{WS} graph conv. with siblings

ϕ^{A2D} graph conv. ancestors to descendants

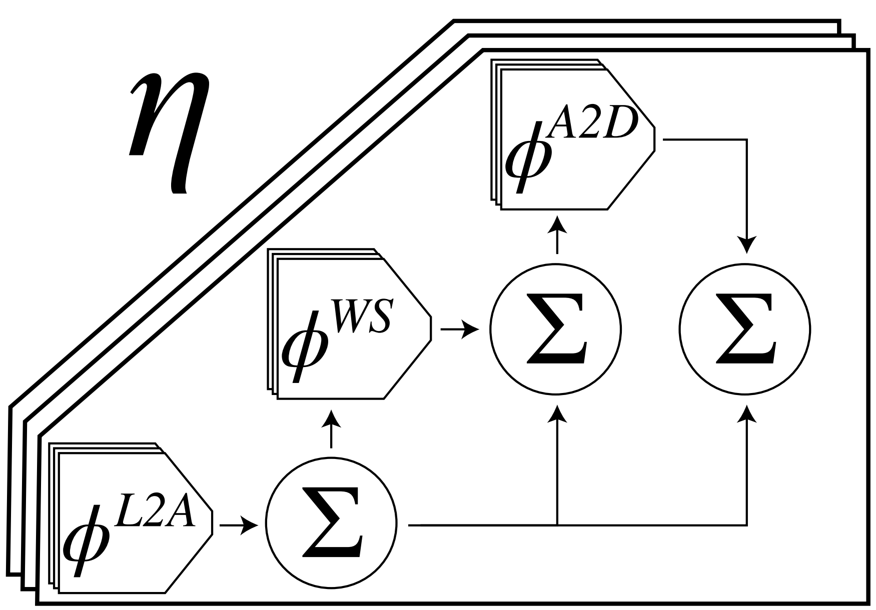

η module composing these three operations from one up-down cycle, adding physical effects

Hierarchical graph convolution propagates interactions efficiently

Intuitive Physics as Underlying Goal

Hierarchical Relational Network (HRN):

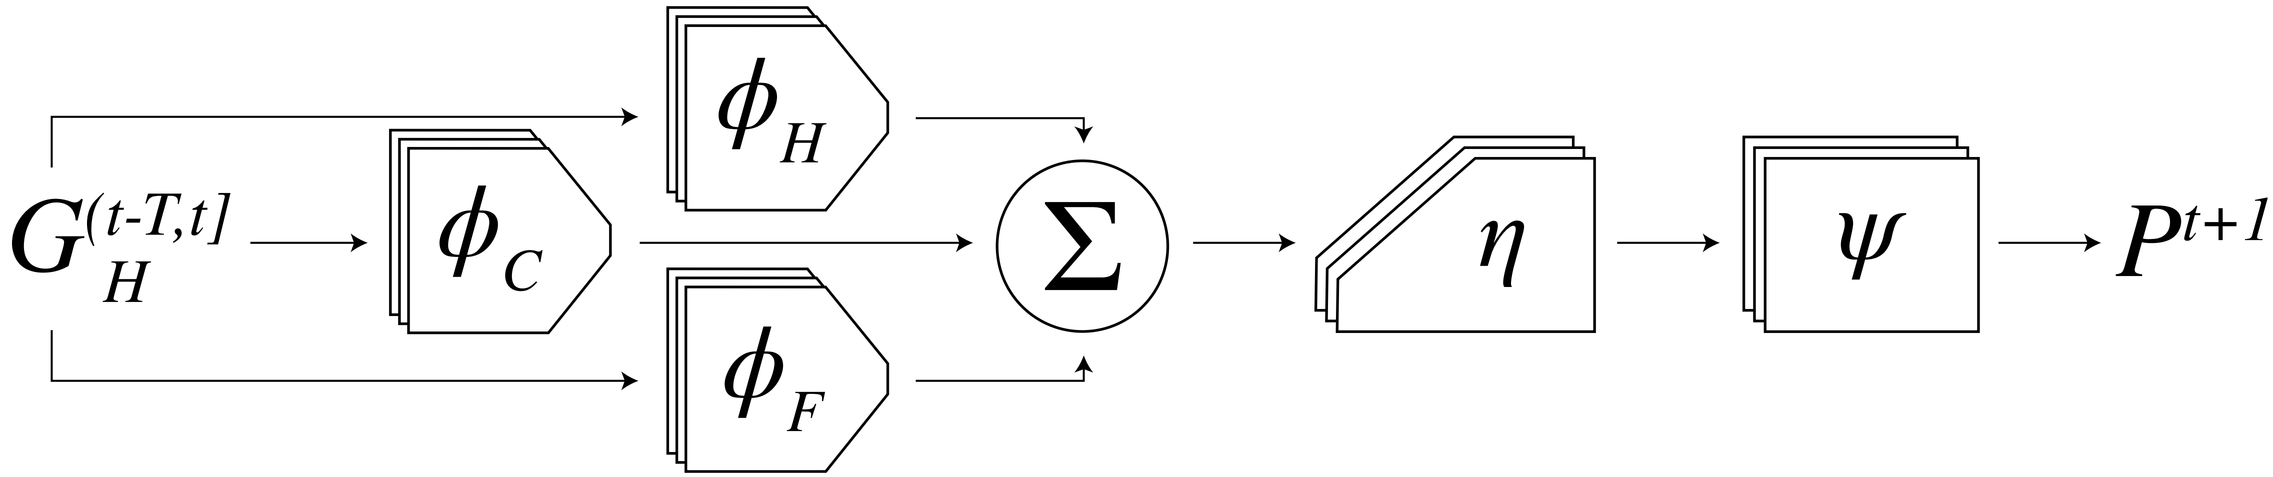

…generates momentum updates (***P***) from hierarchical graph state (***G***).

Intuitive Physics as Underlying Goal

Hierarchical Relational Network (HRN):

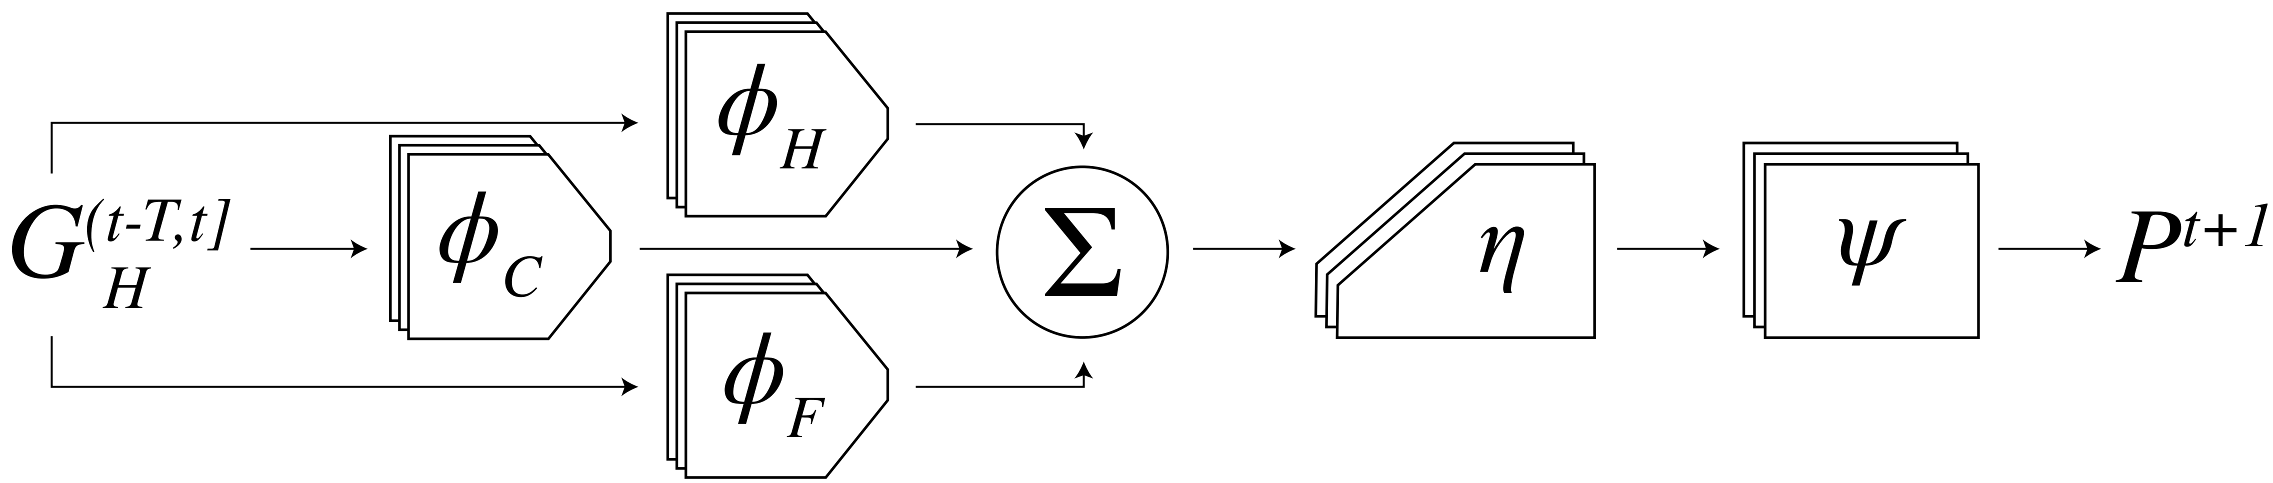

…generates momentum updates (***P***) from hierarchical graph state (***G***).

Network learns to interpret graph structure (including meaning of material-vector edge labels)…

Deformable cone bouncing off a flat floor

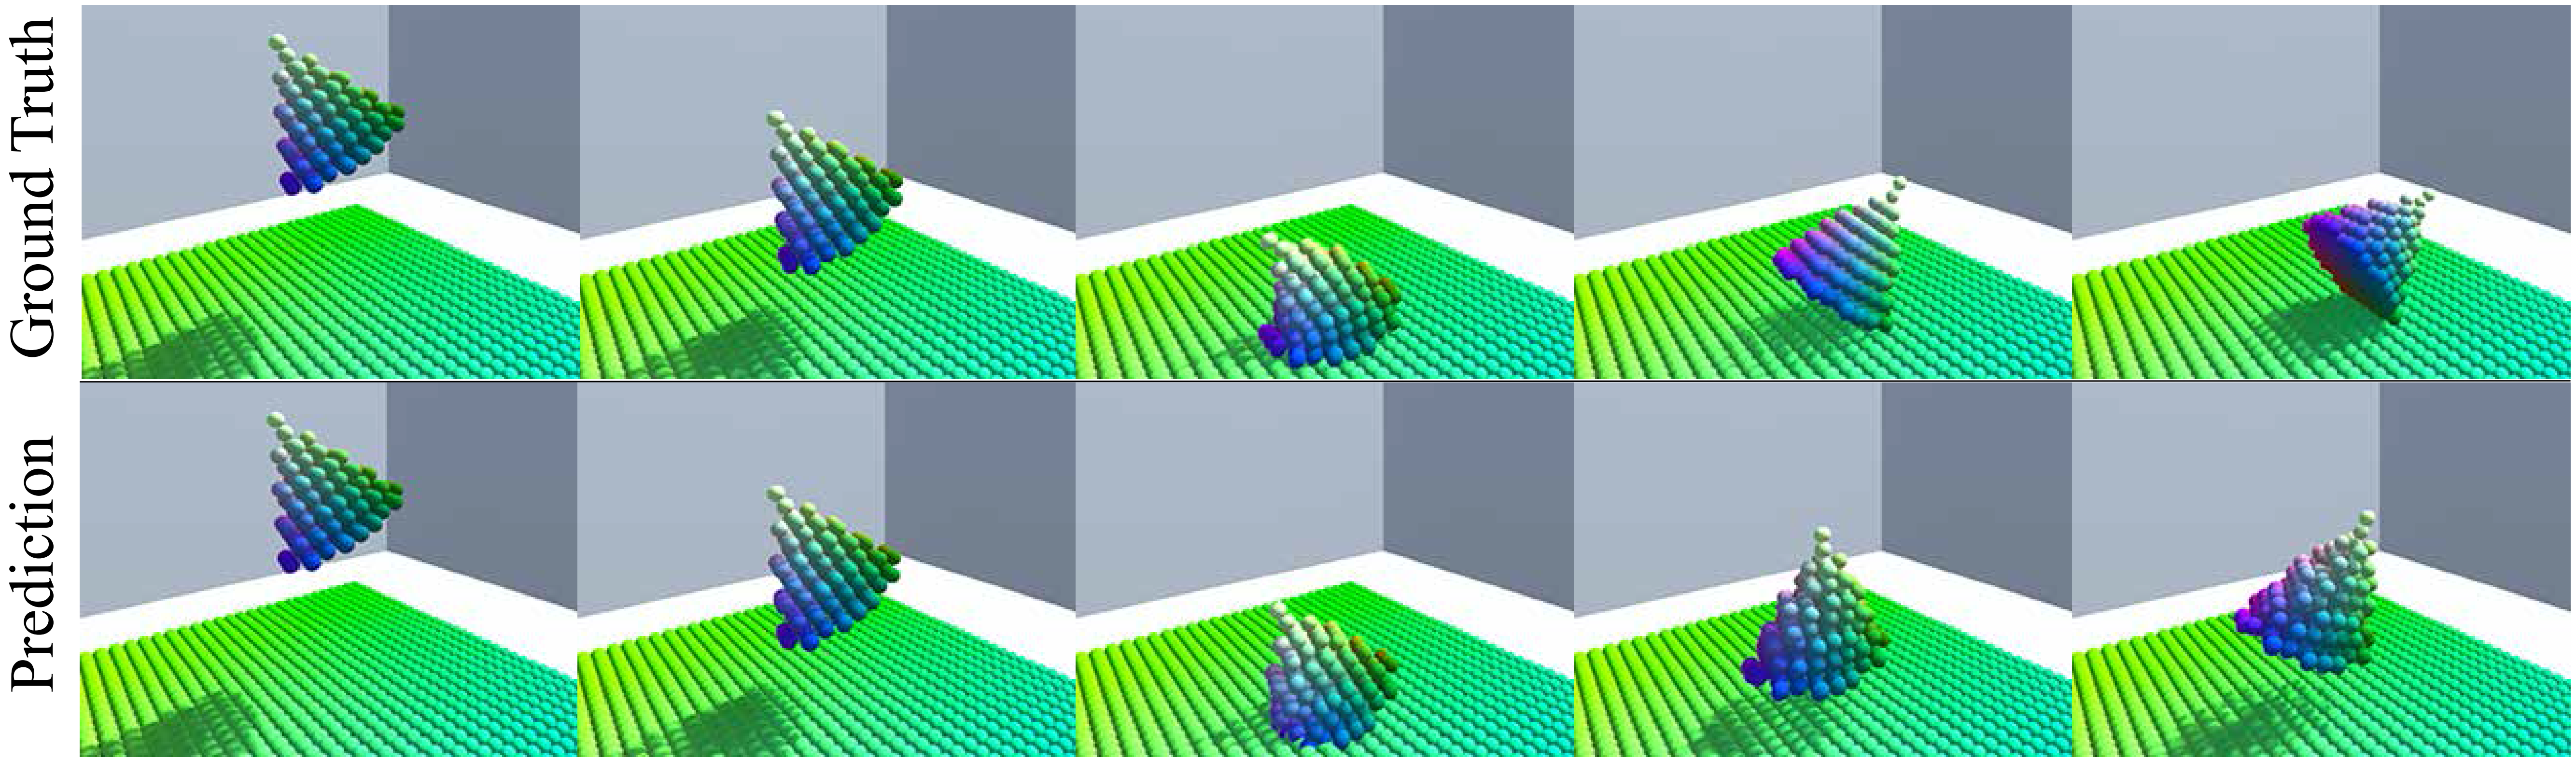

Deformable cone bouncing off a flat floor

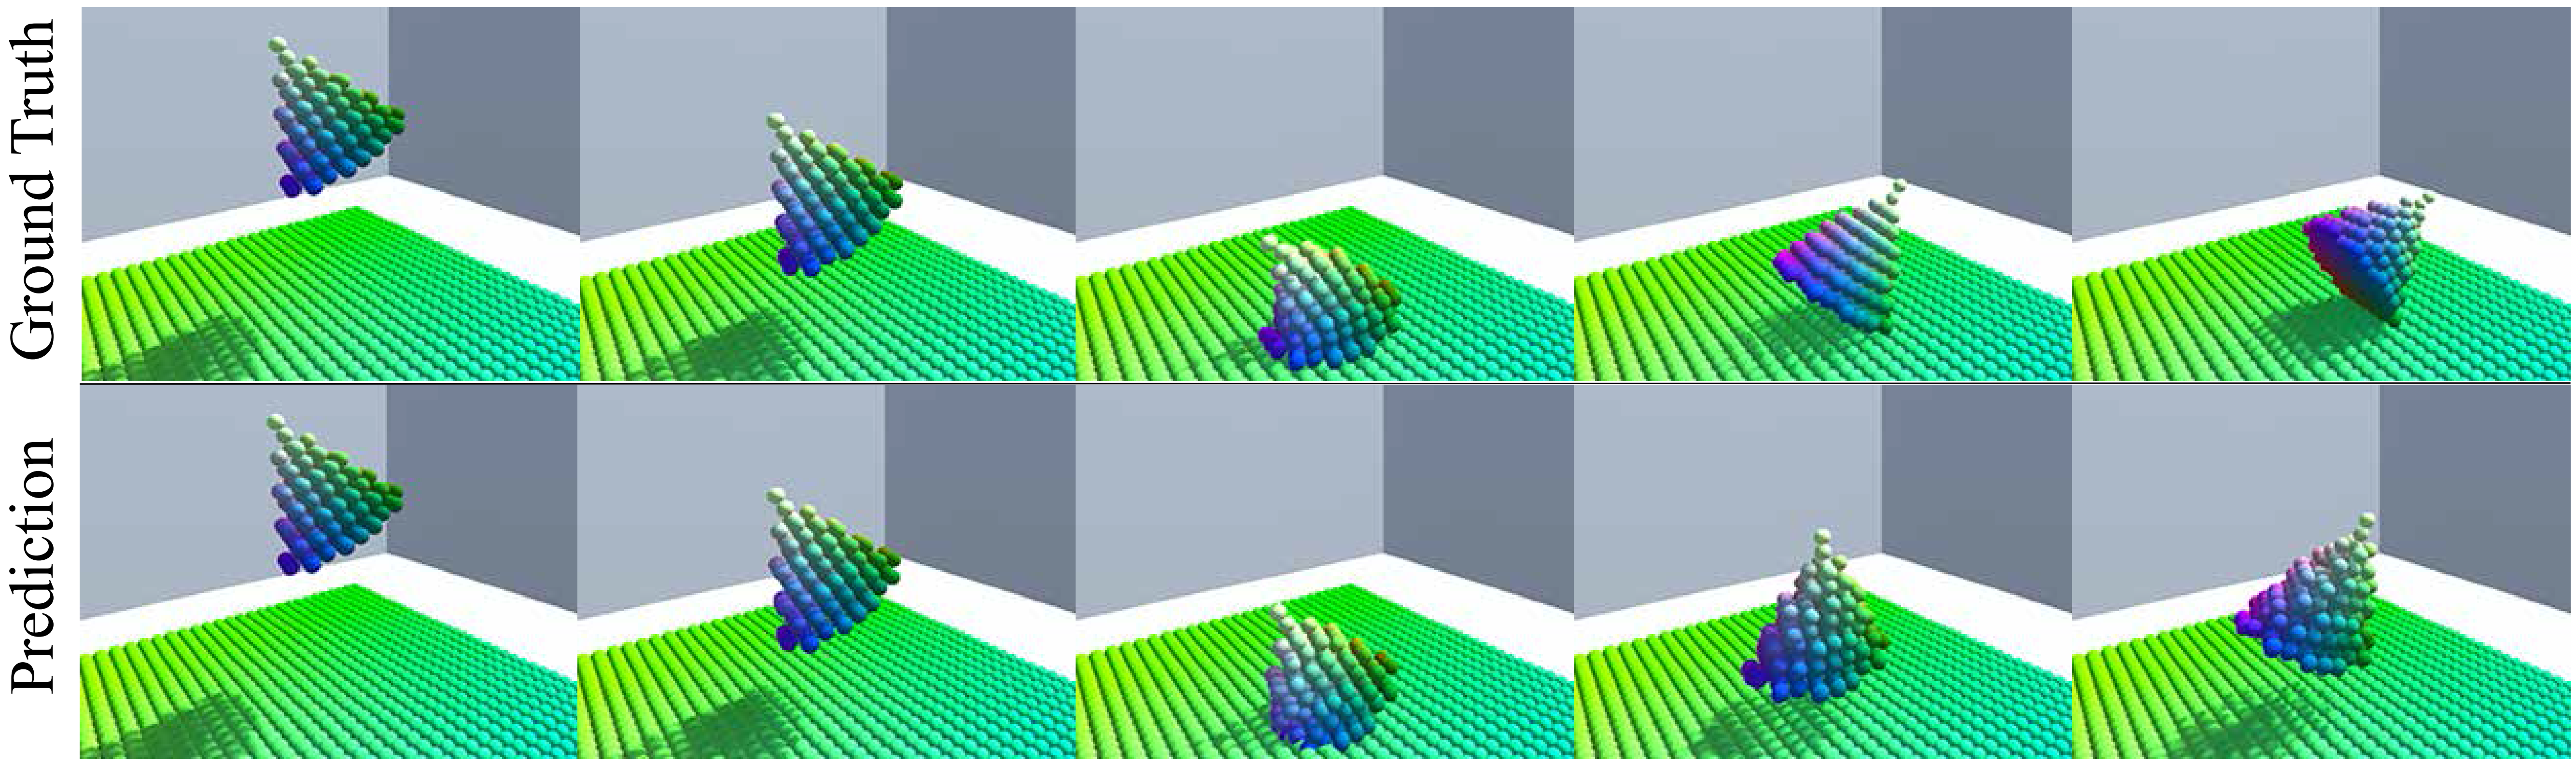

Stanford bunny

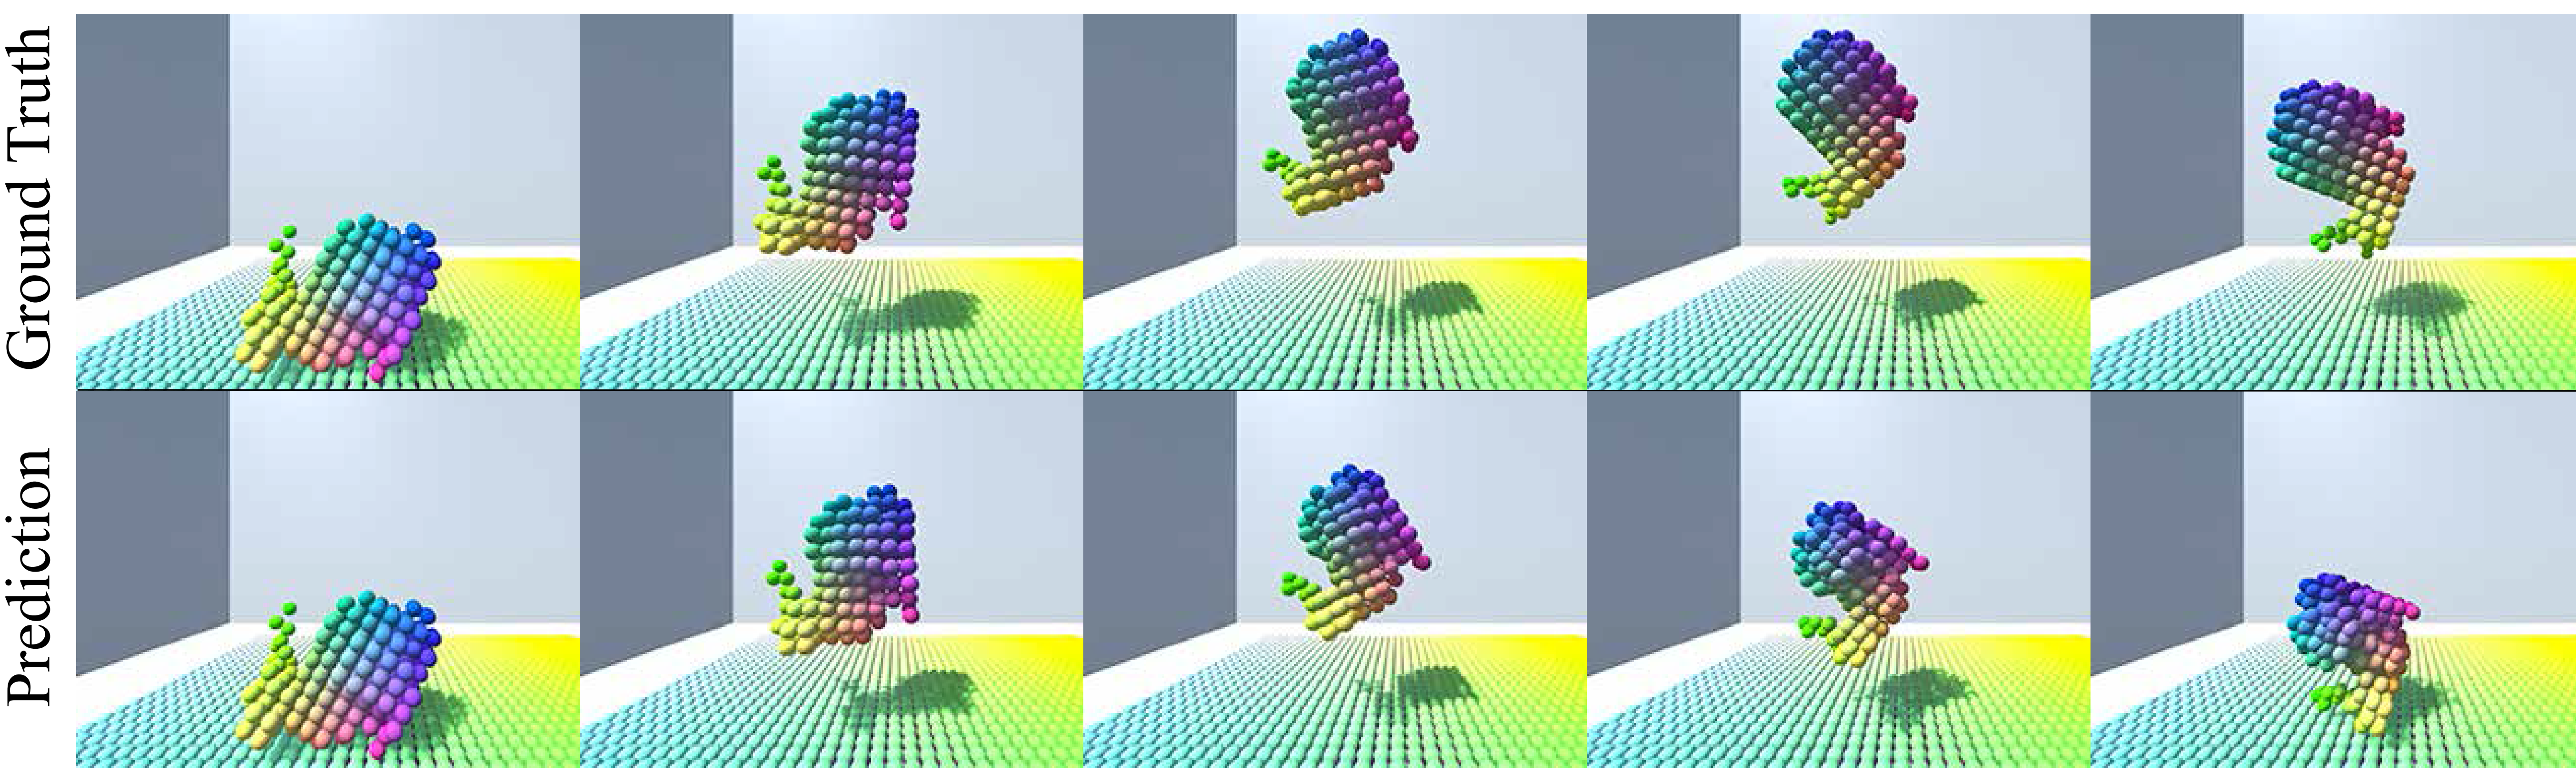

Deformable box bouncing off an incline

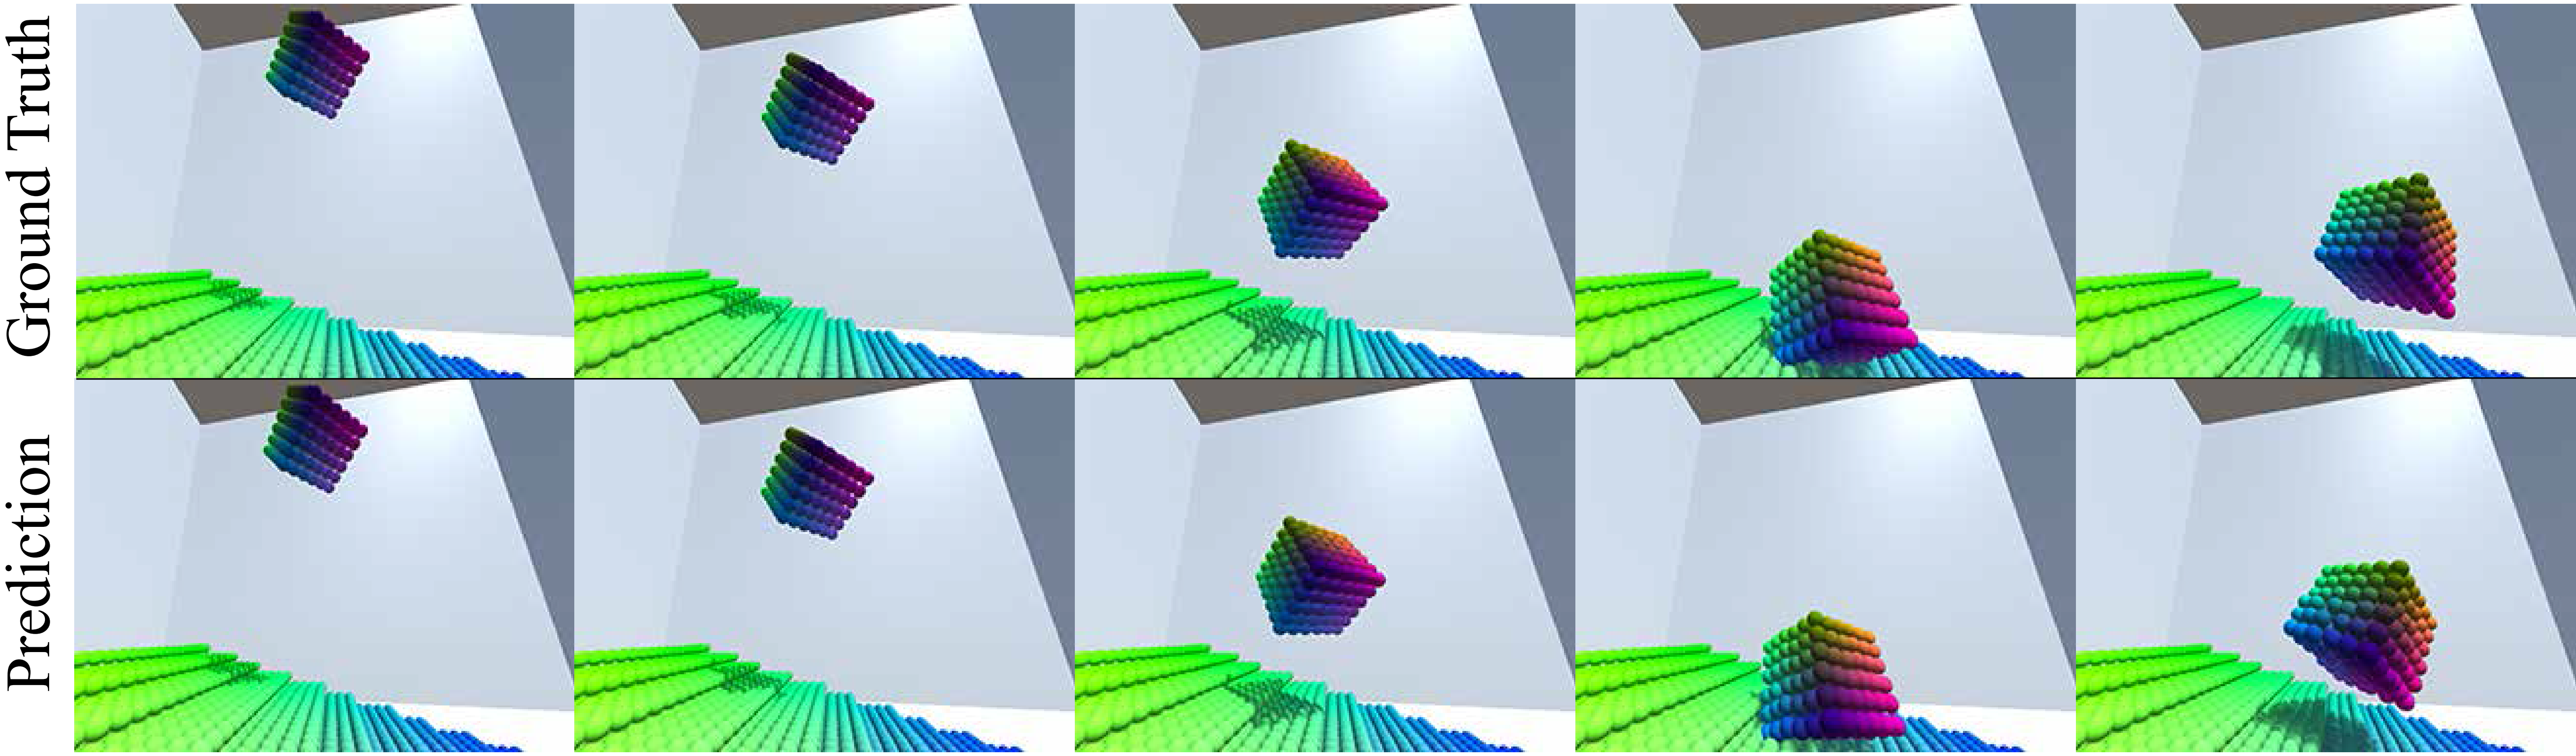

Deformable box bouncing off an incline

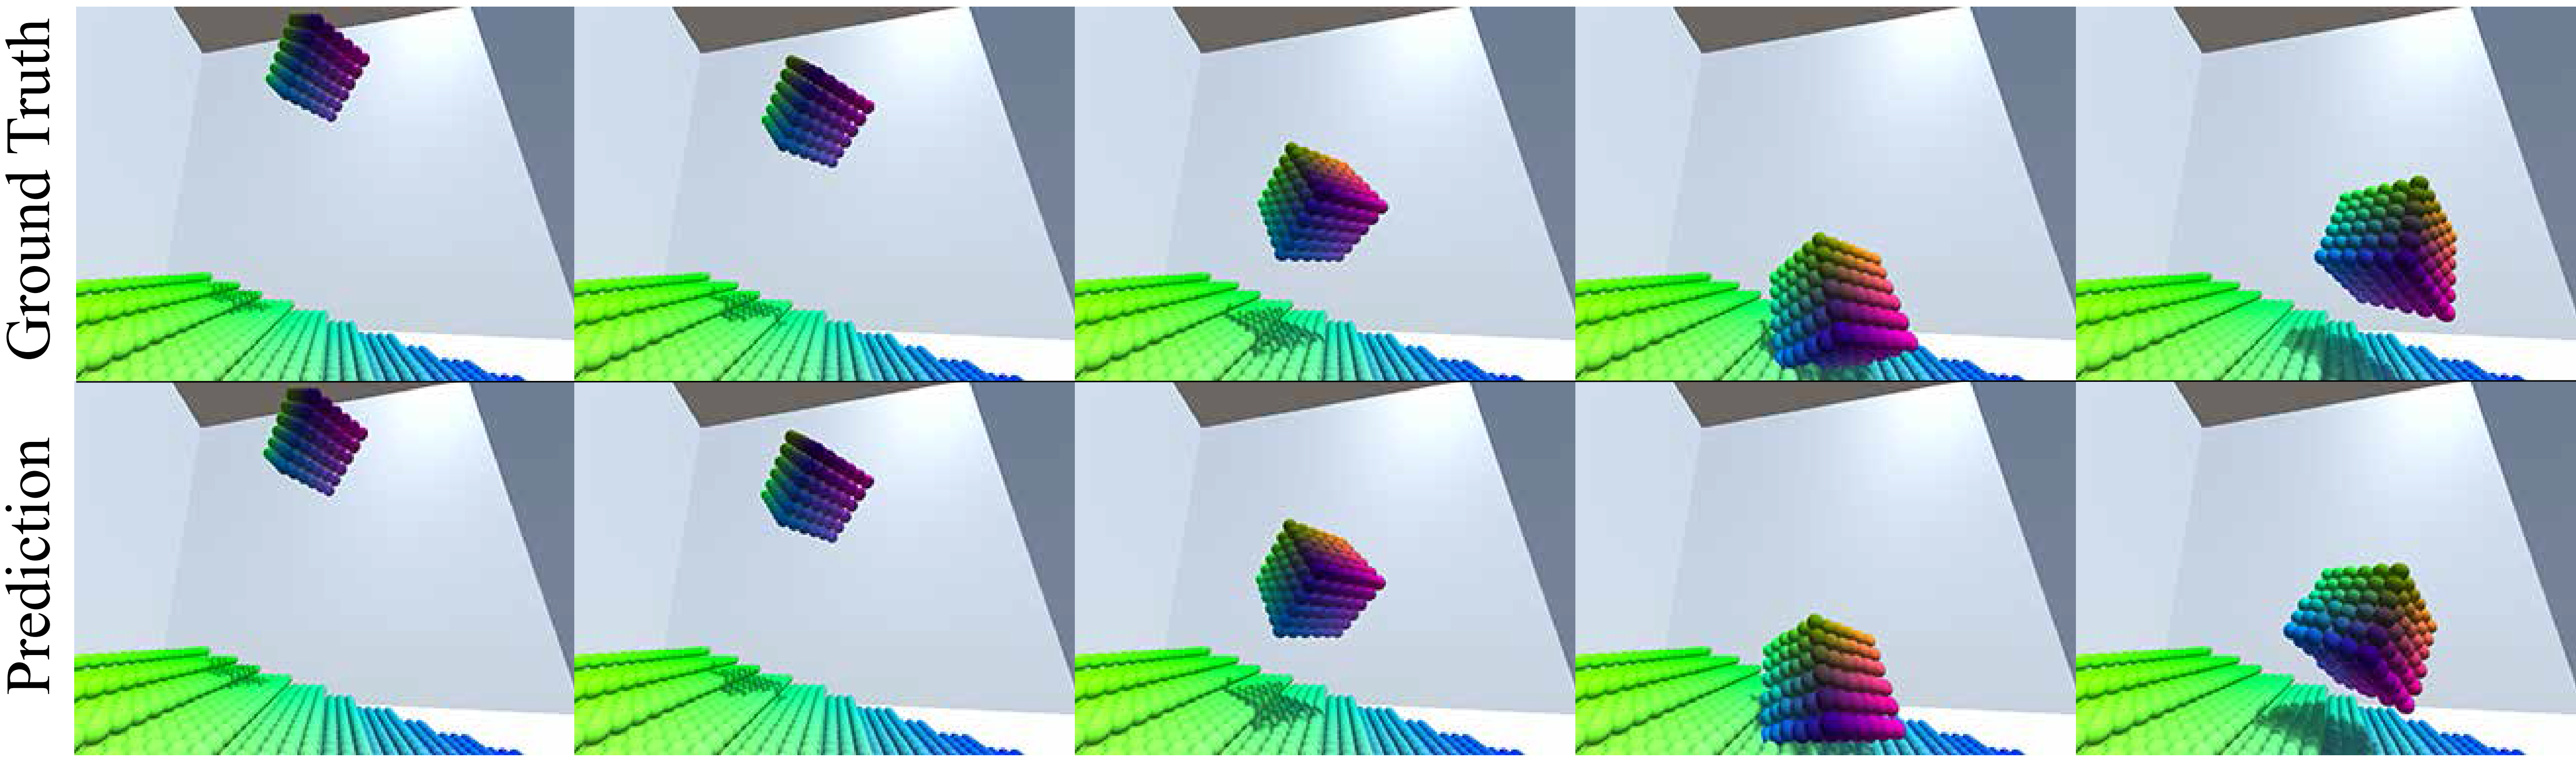

Multiple rigid objects colliding

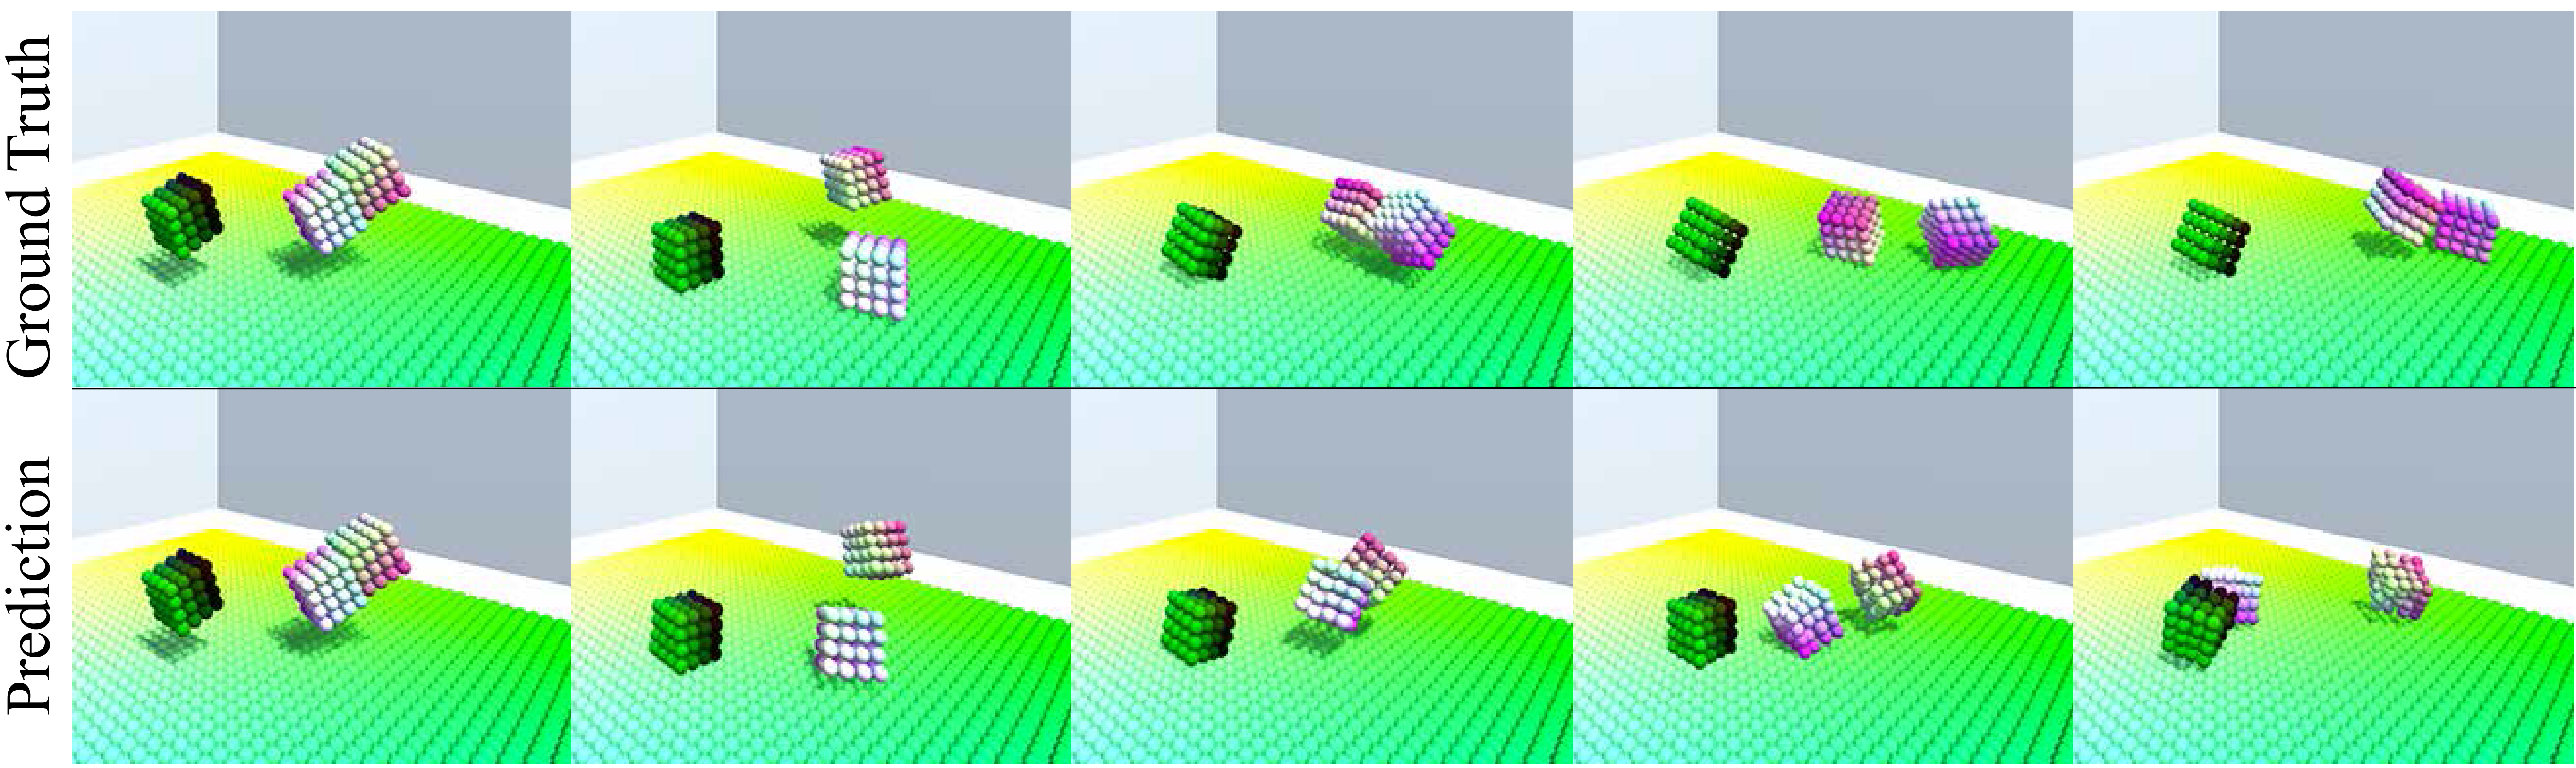

rigid sphere rolling out of rigid bowl

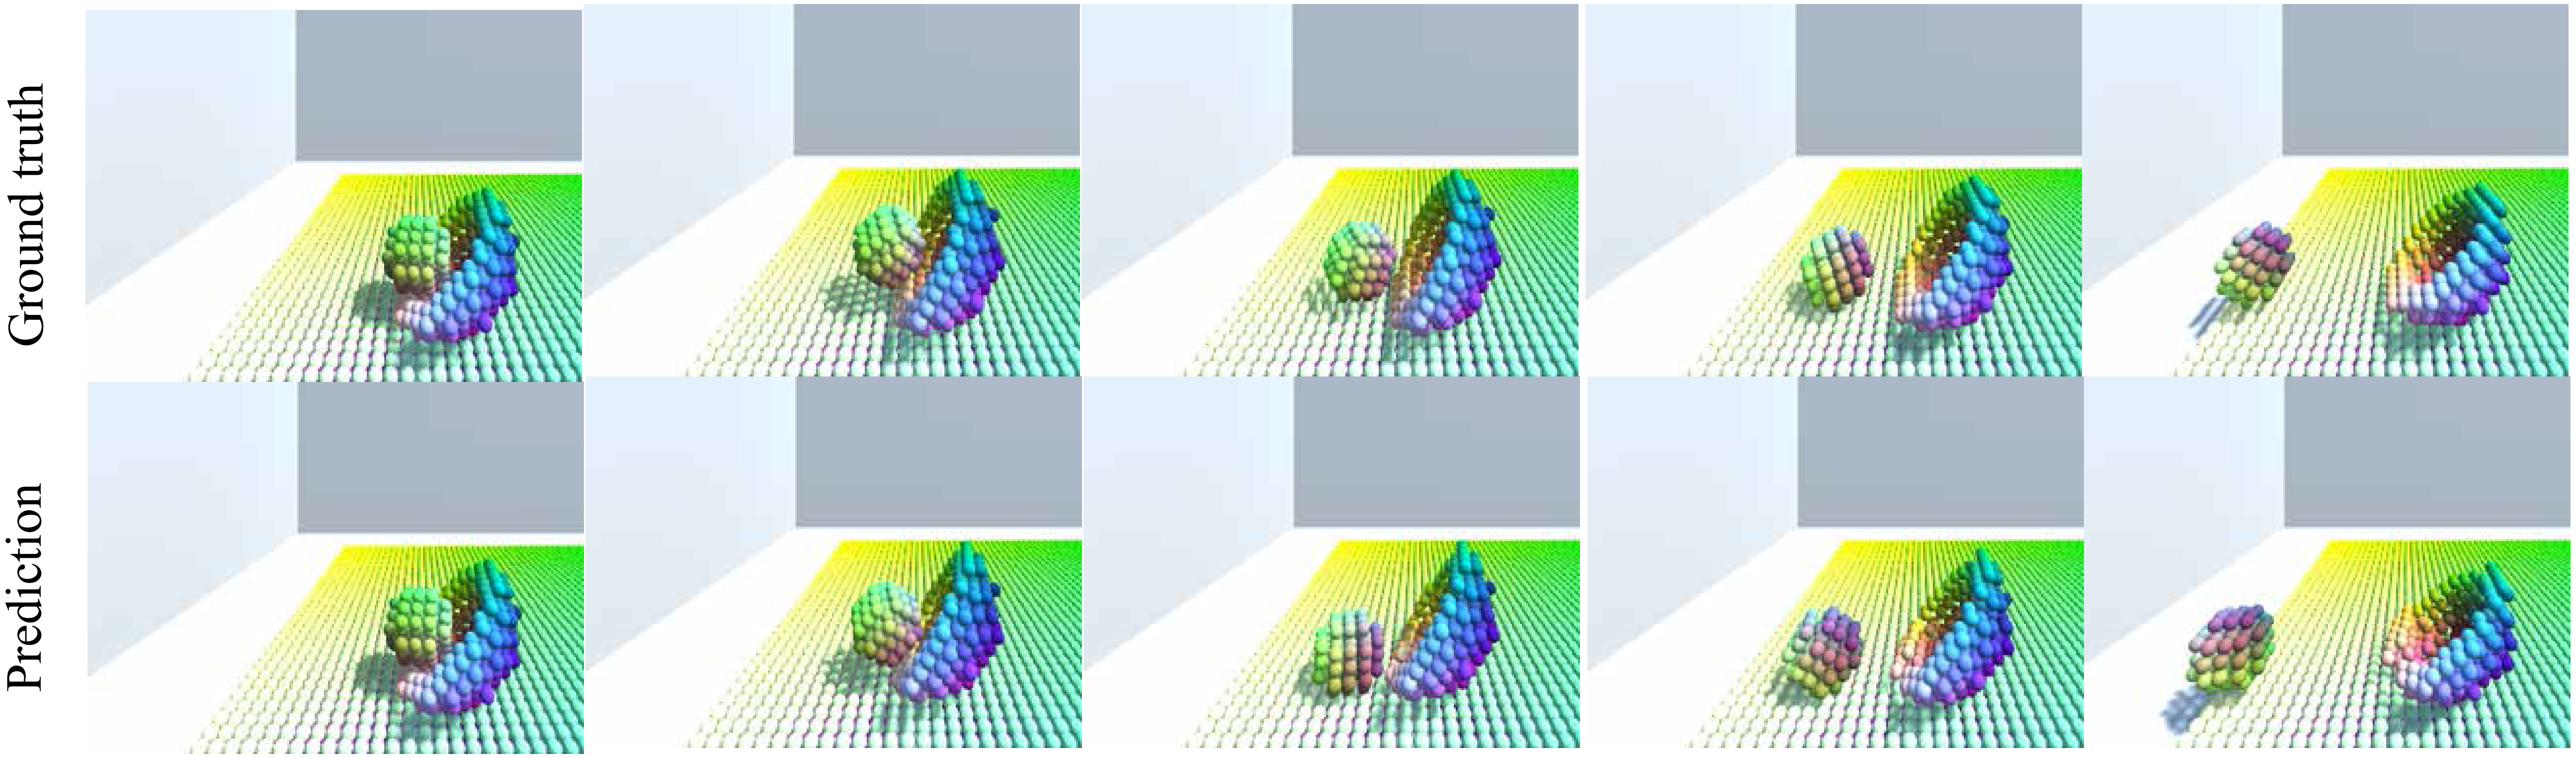

rigid sphere rolling out of rigid bowl

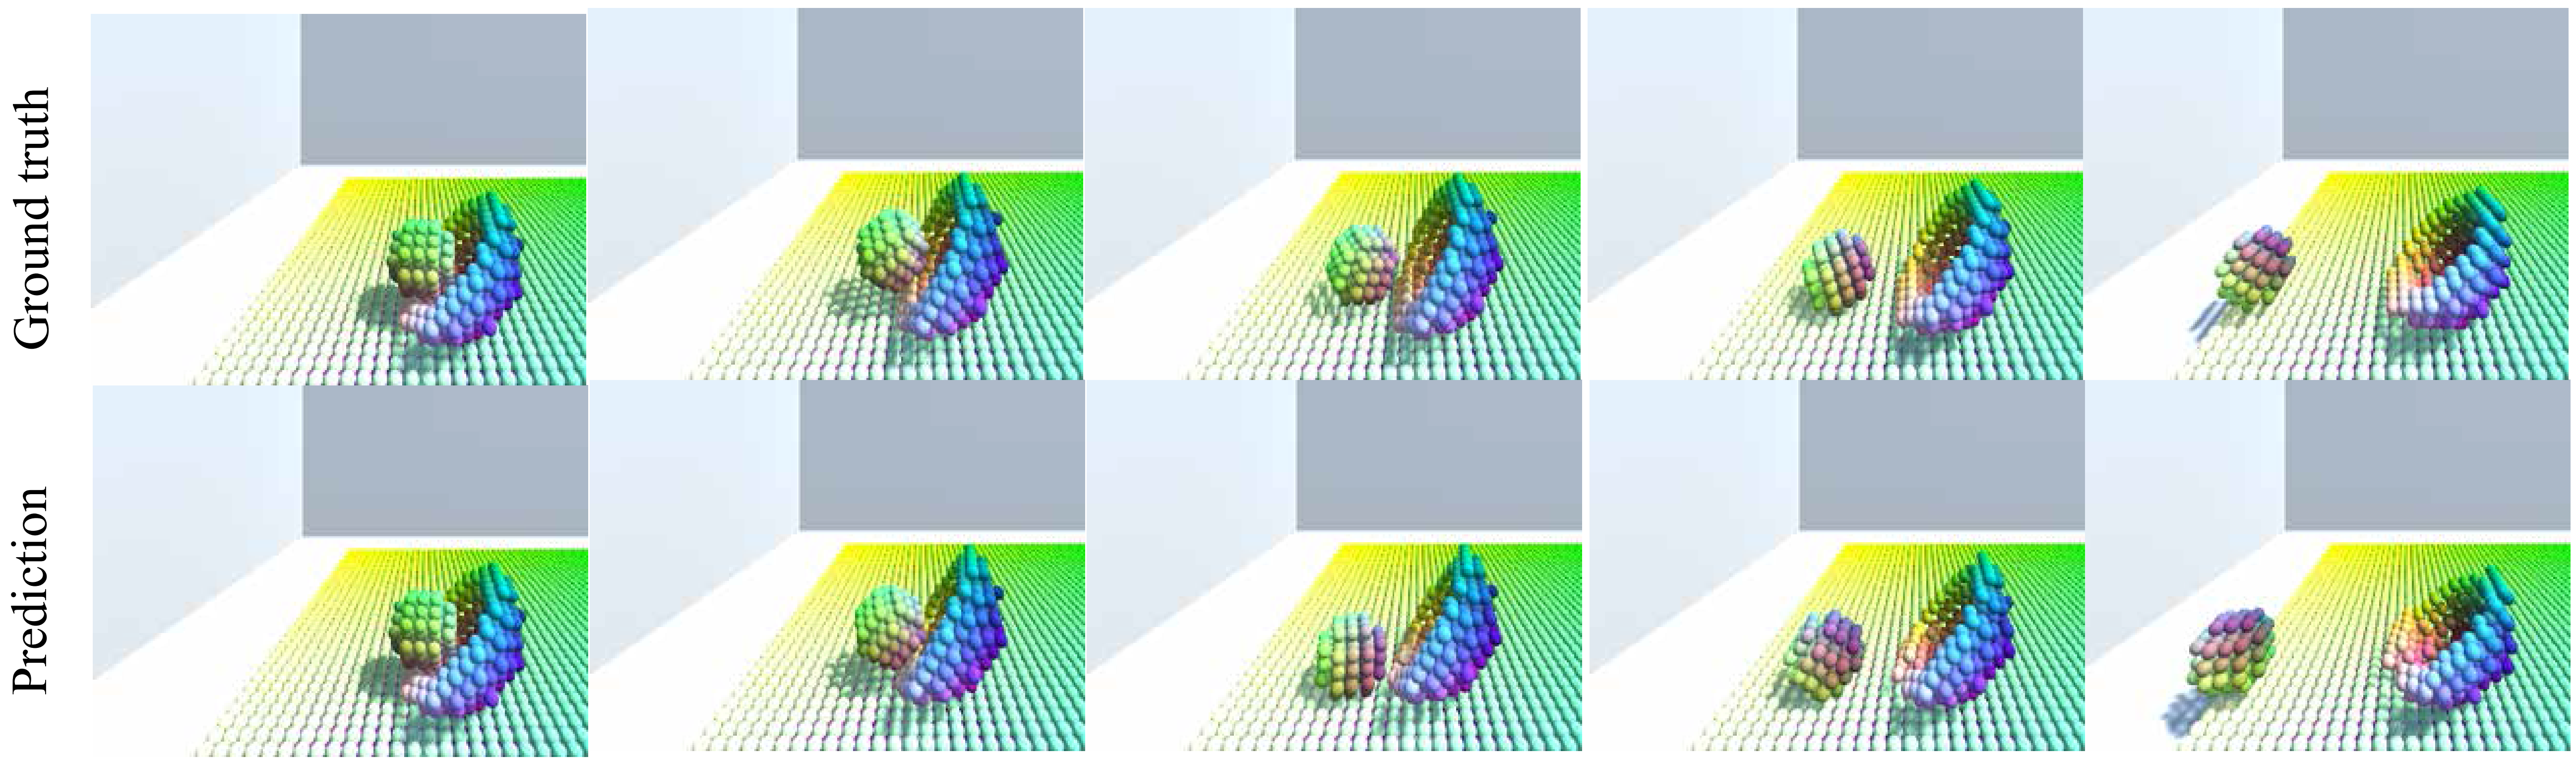

floppy teddybear bouncing off floor and recovering

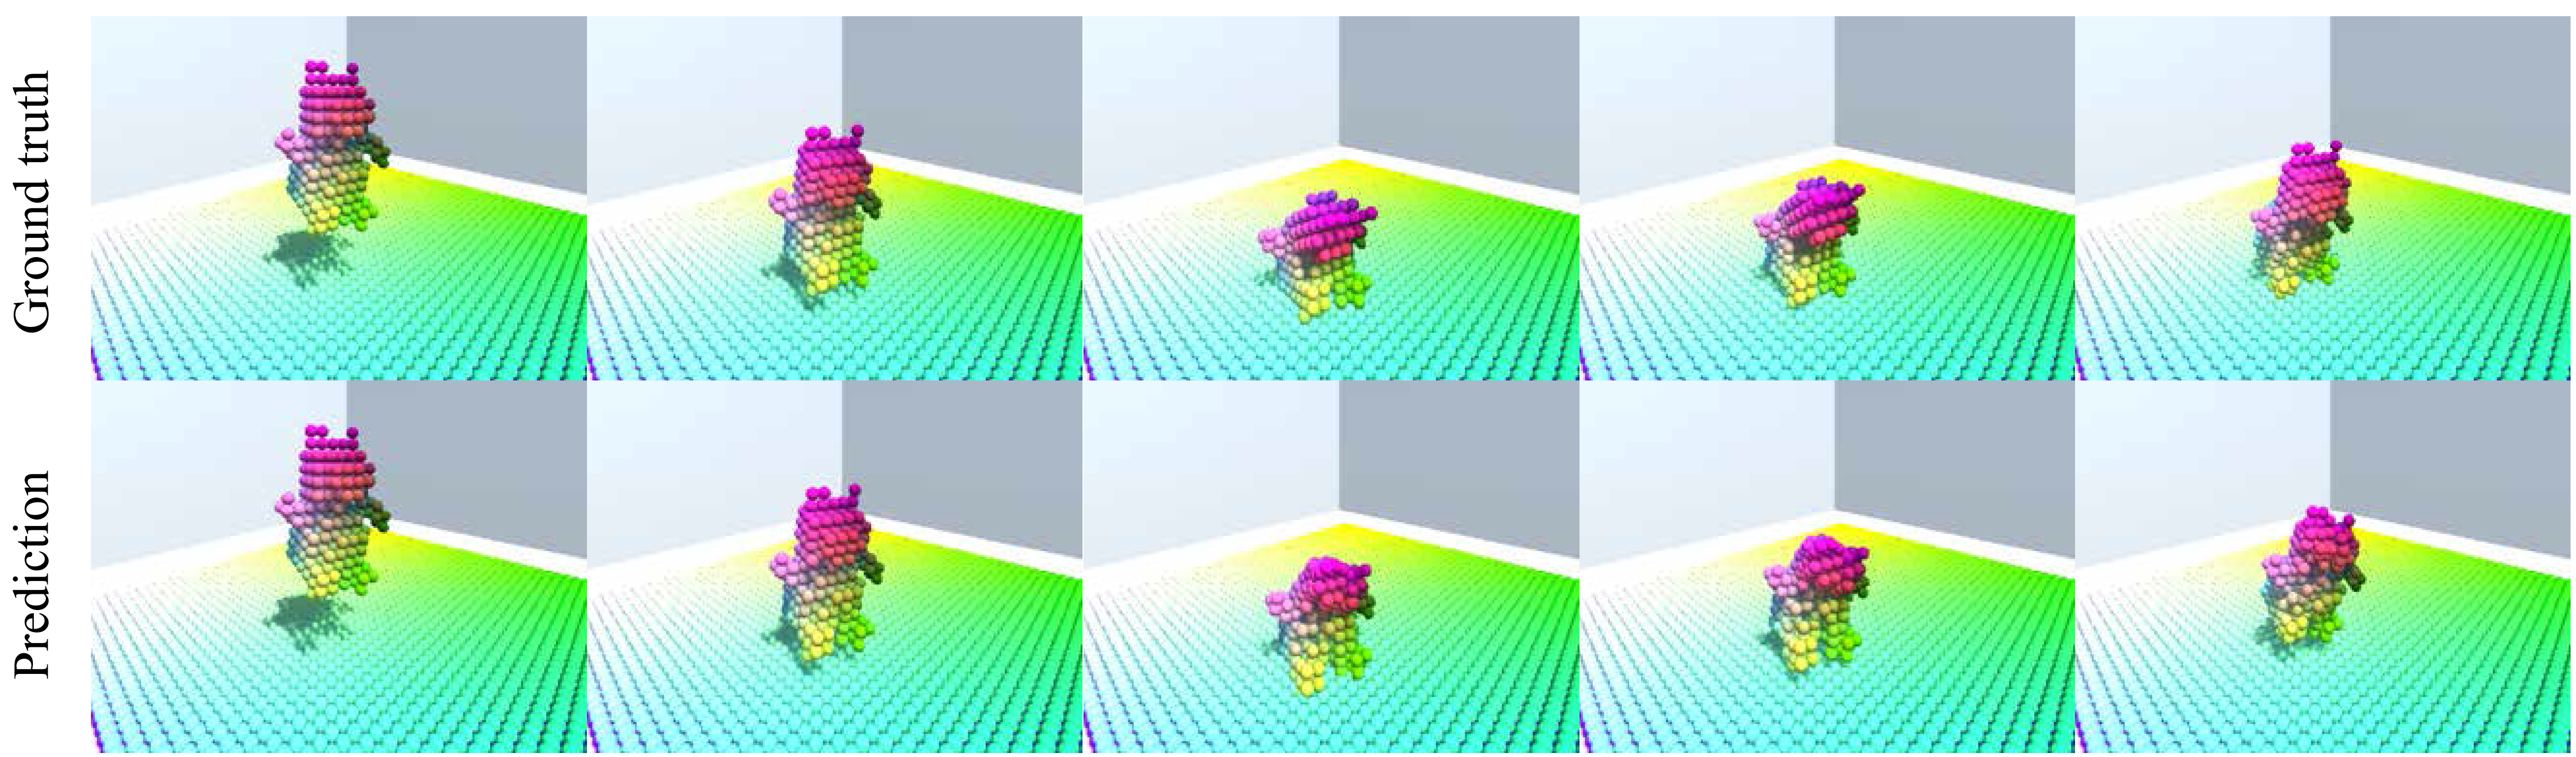

knocking over an unstable block tower

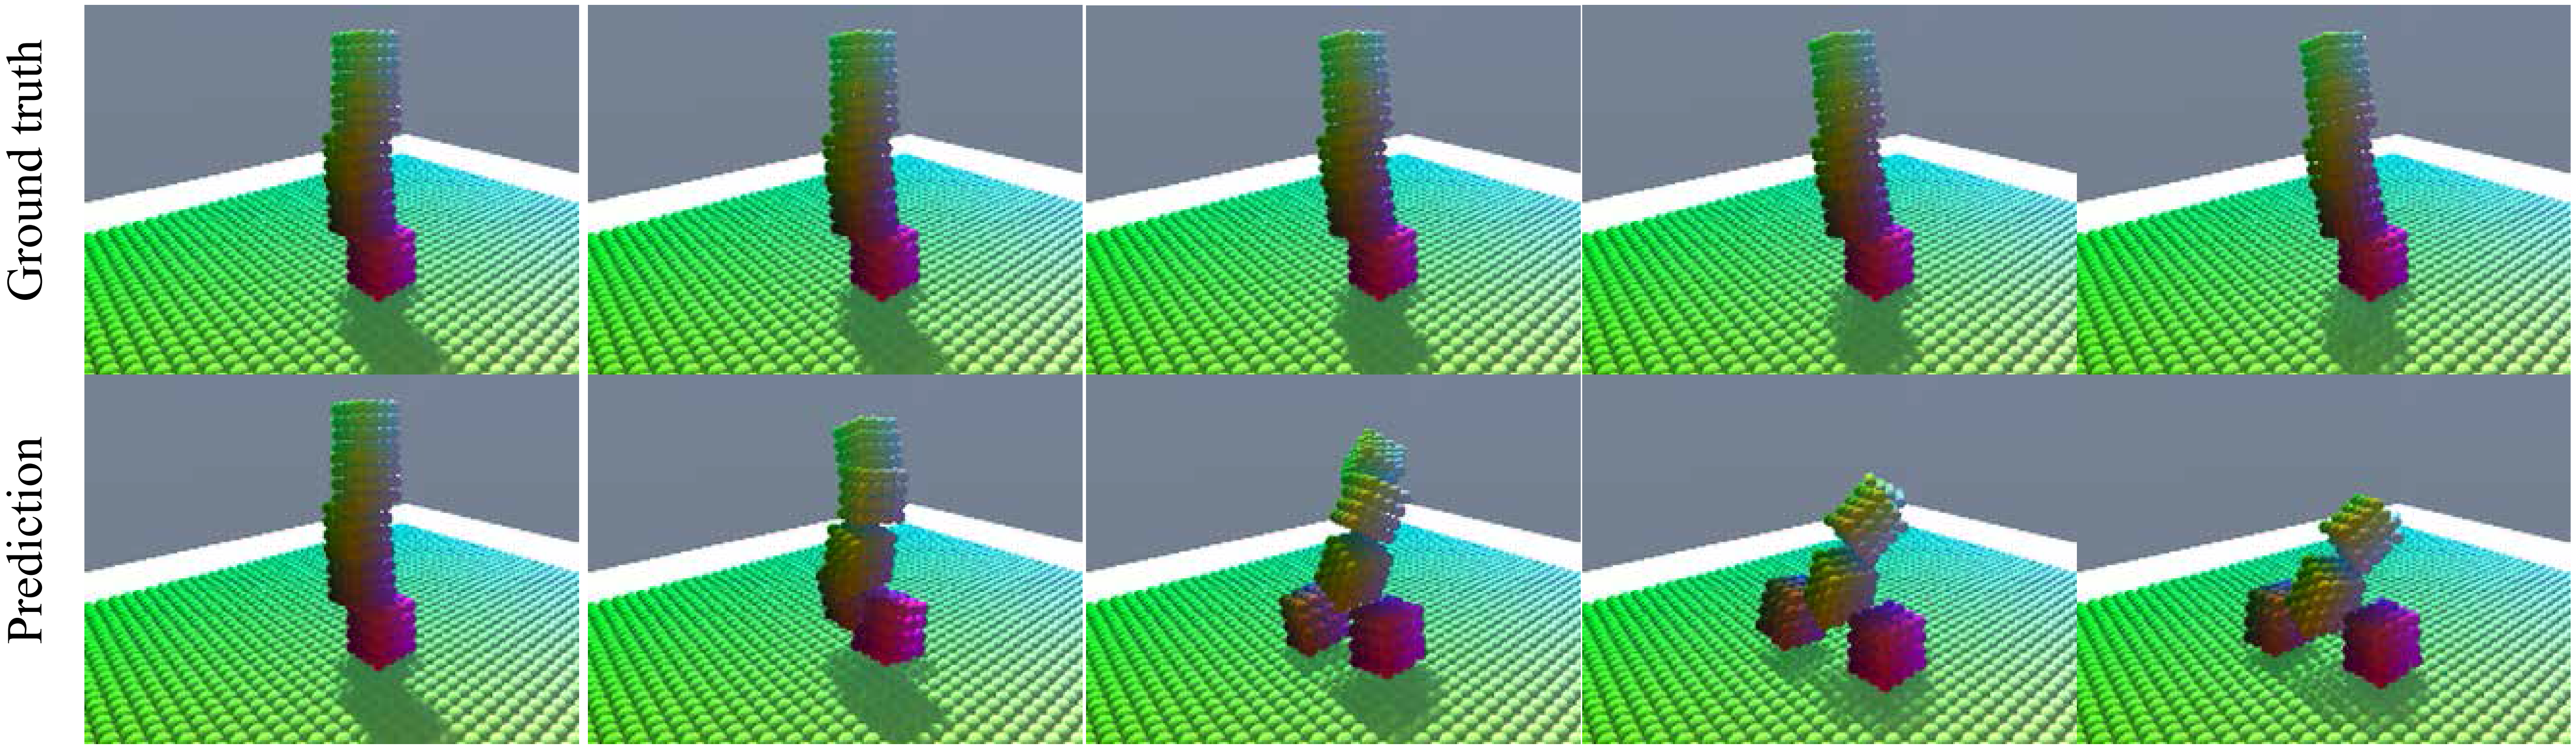

in GT the tower does fall, but prediction falls too fast . . .

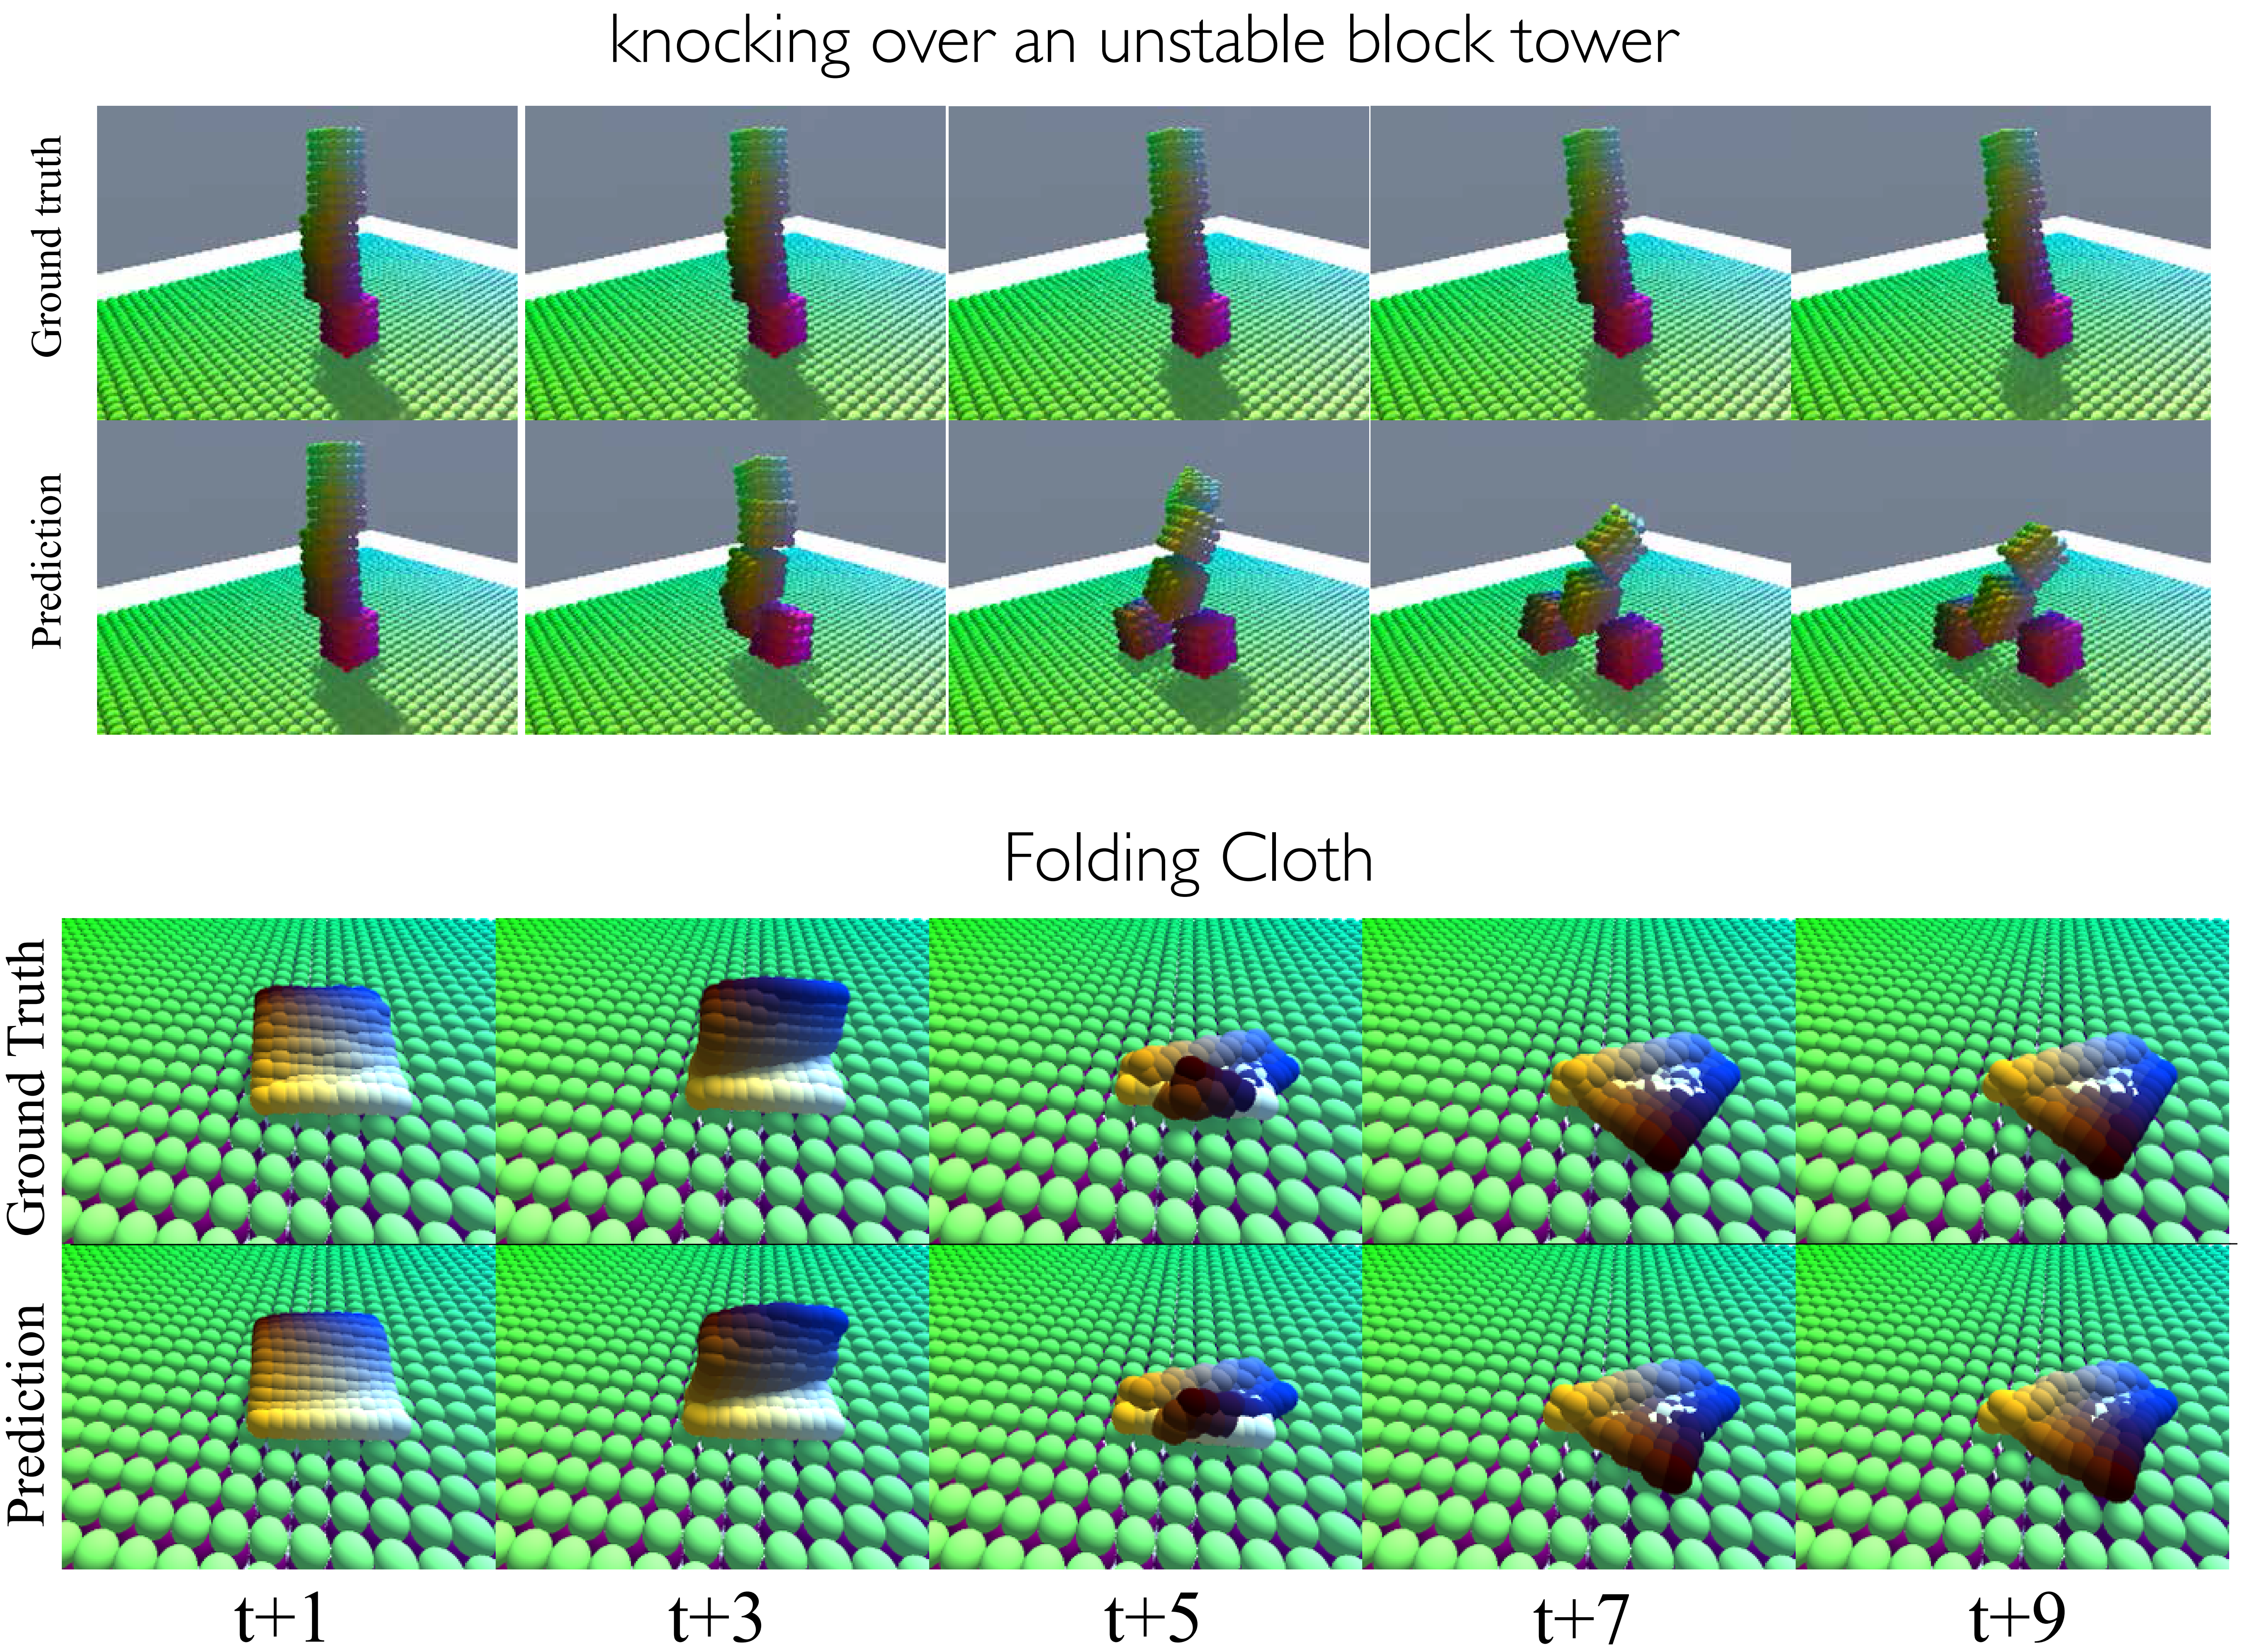
knocking over an unstable block tower
Ground truth
Prediction
Folding Cloth
Ground Truth
Prediction
t+1
t+3
t+5
t+7
t+9

knocking over an unstable block tower

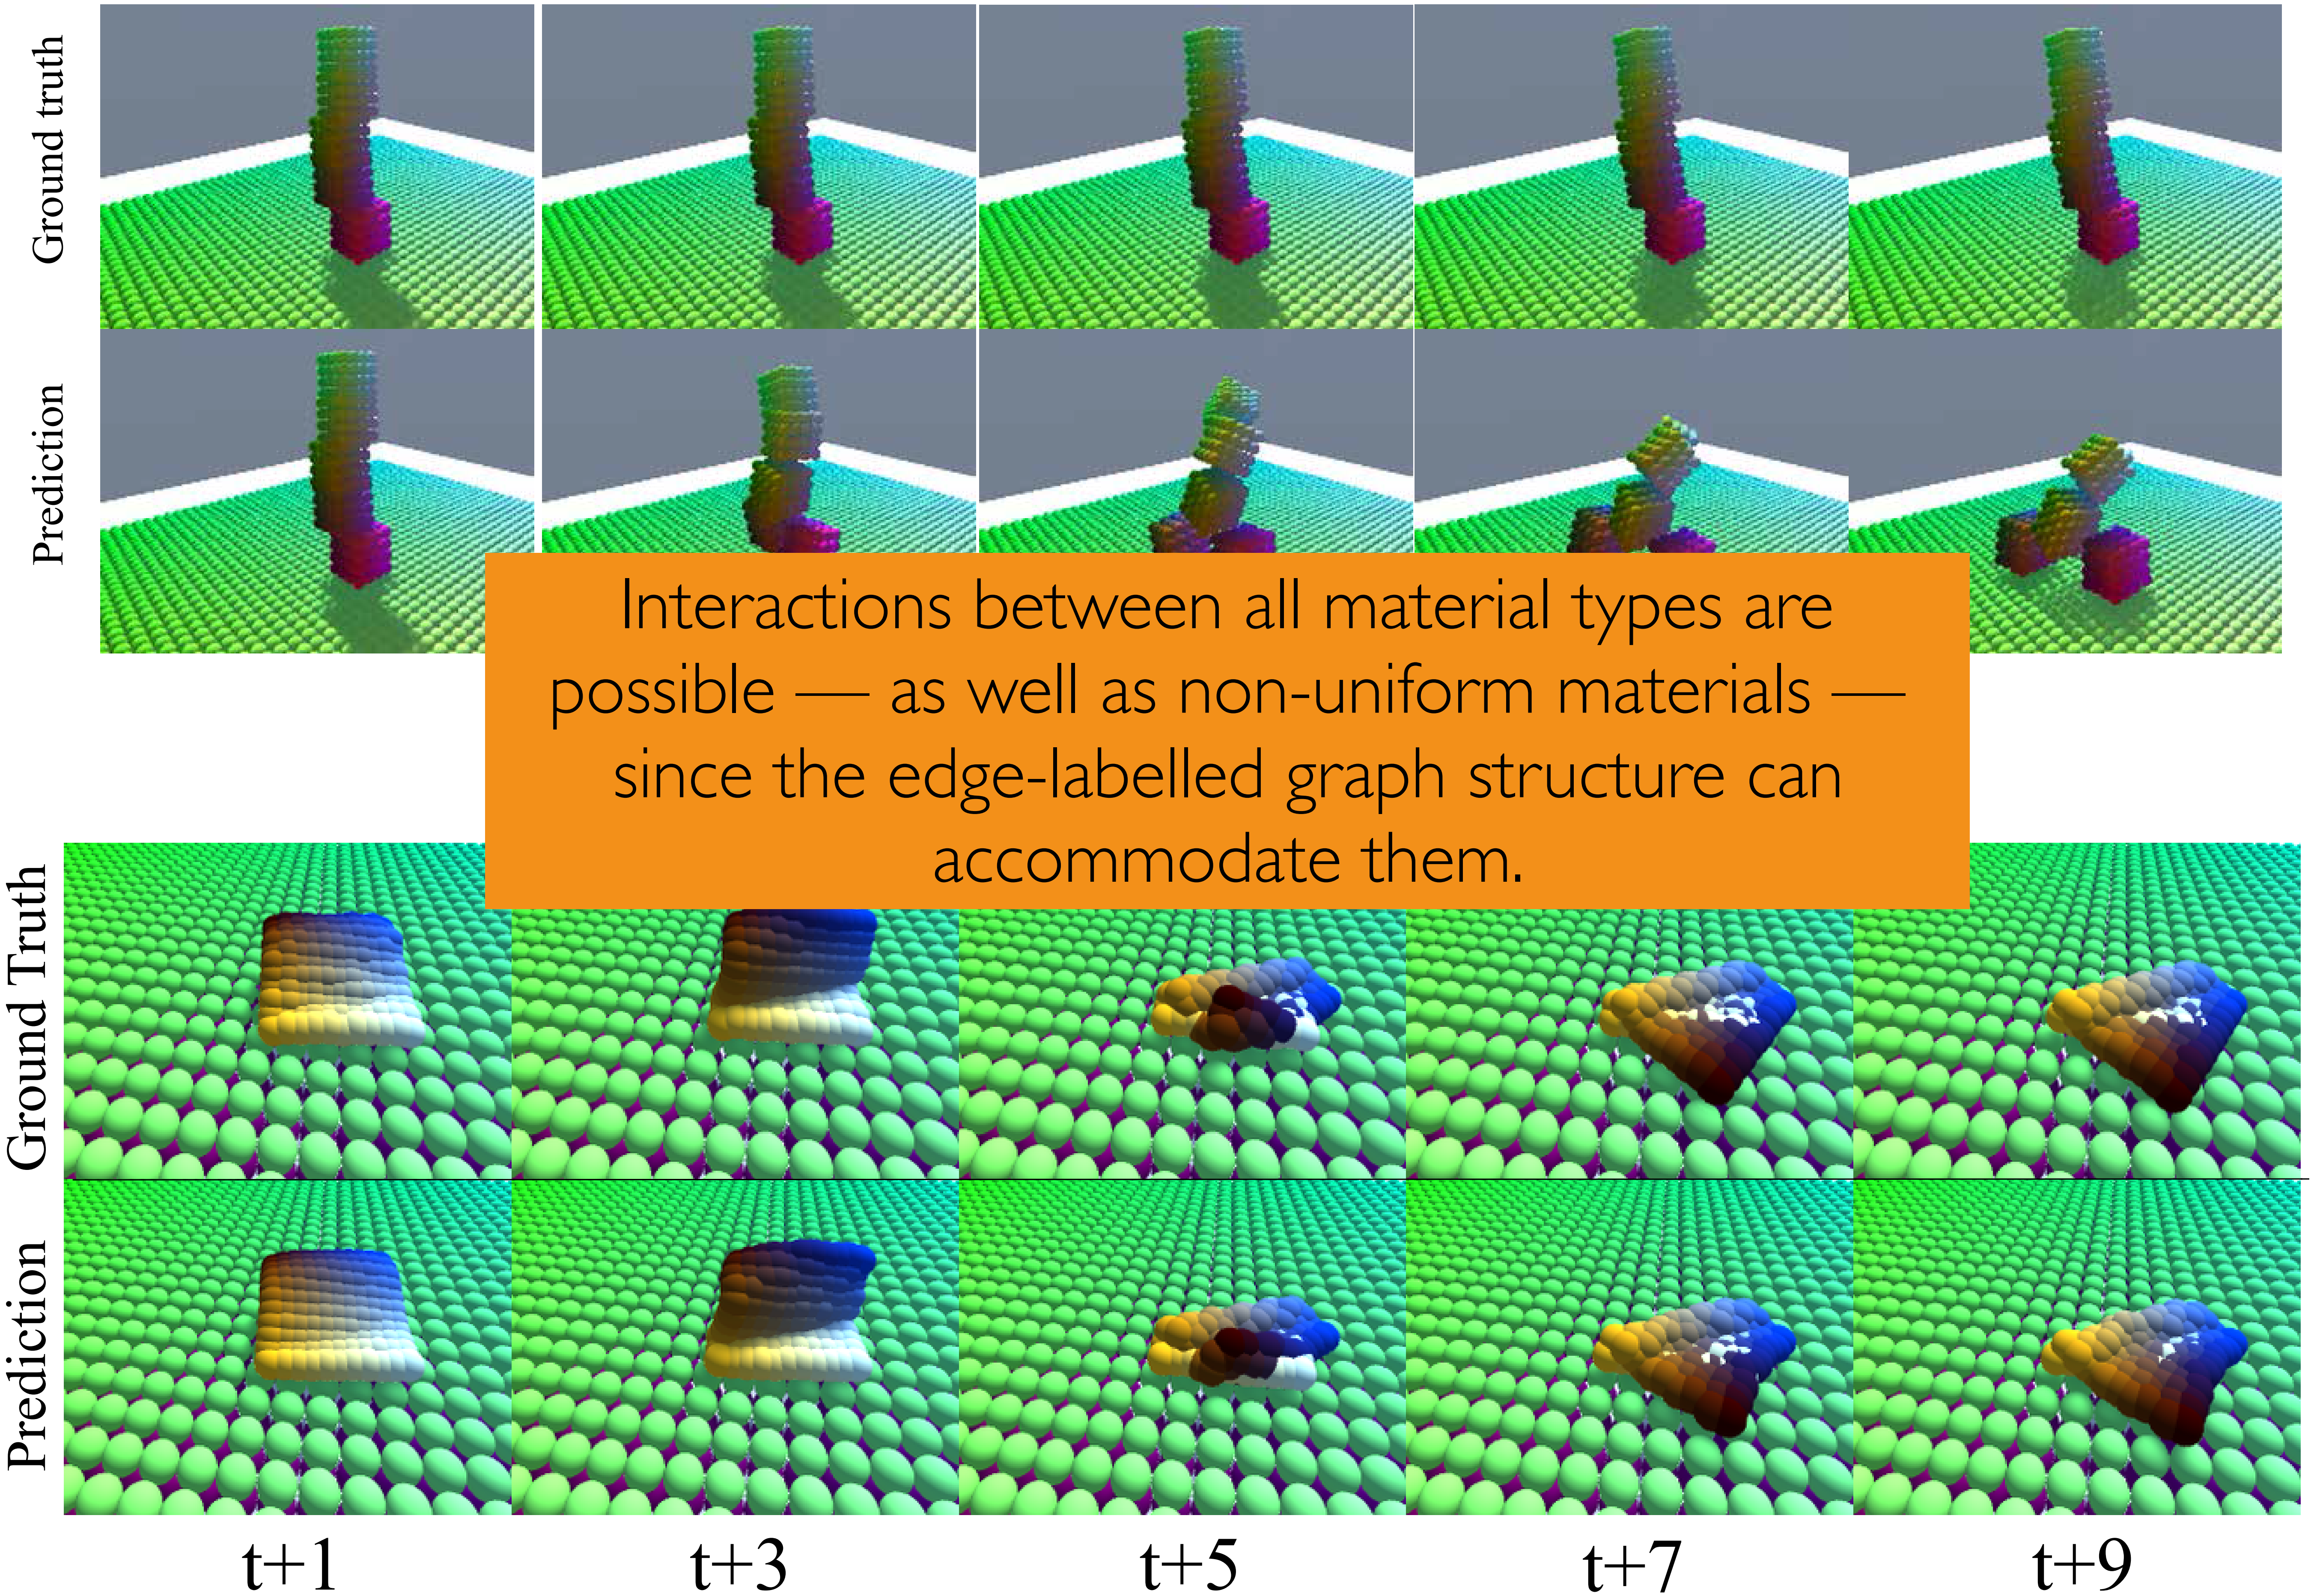

Interactions between all material types are possible — as well as non-uniform materials — since the edge-labelled graph structure can accommodate them.

Challenges:

slinky

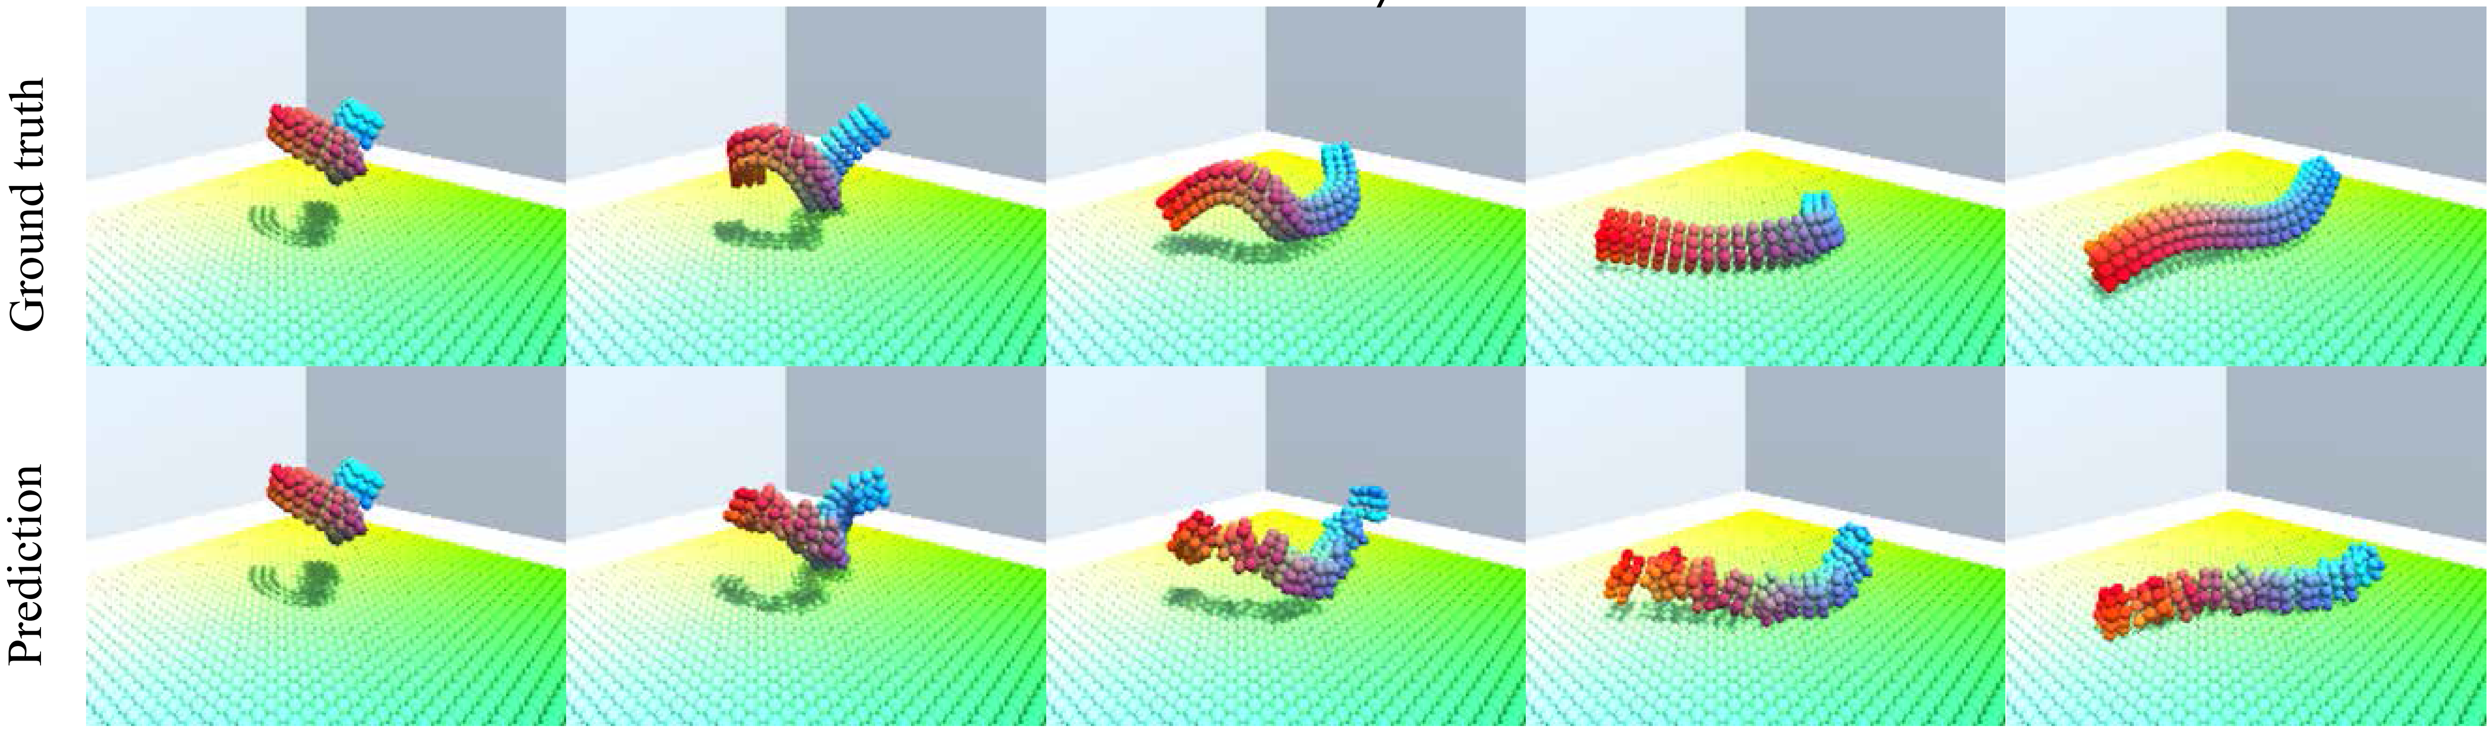

shape is not preserved super well over long rollouts…

Challenges:

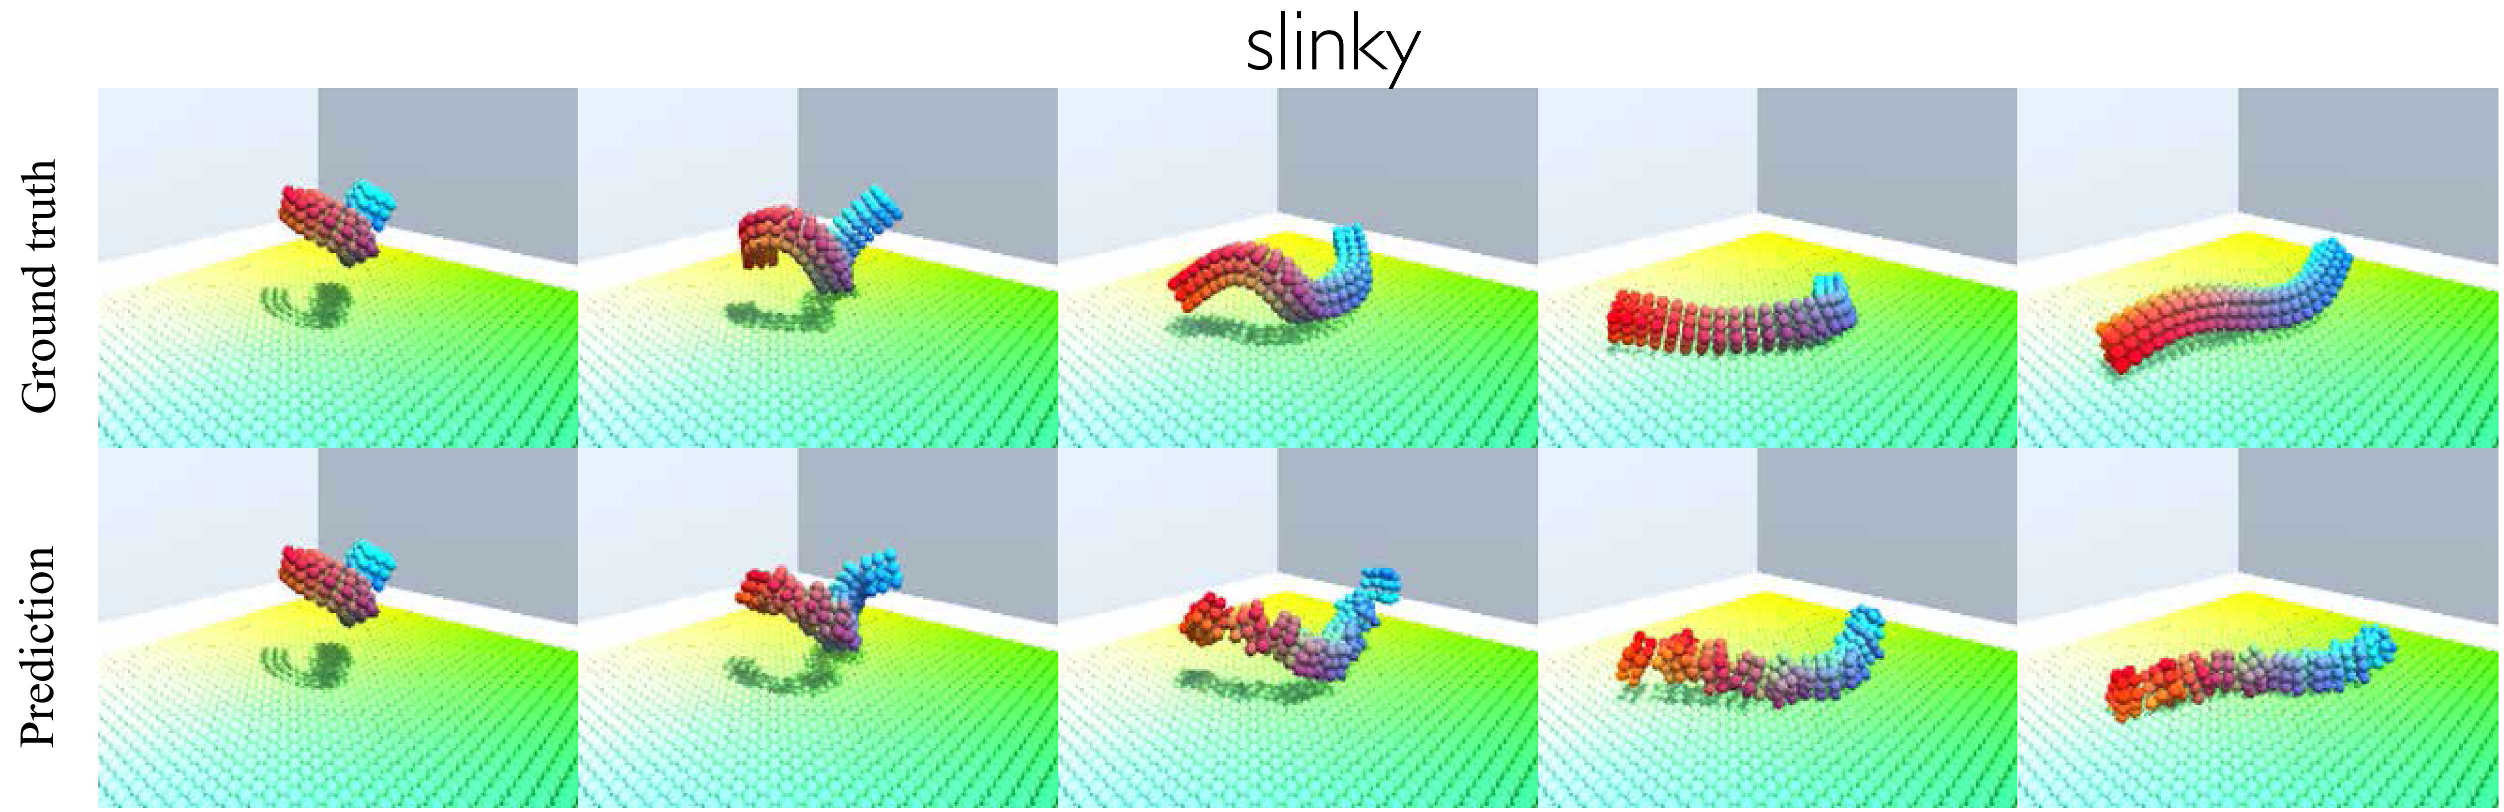

shape is not preserved super well over long rollouts…

Easy to impose simple shape conversation rules — in a "per material" way. (e.g. rigid different than cloth different than soft-body)

Challenges:

slinky

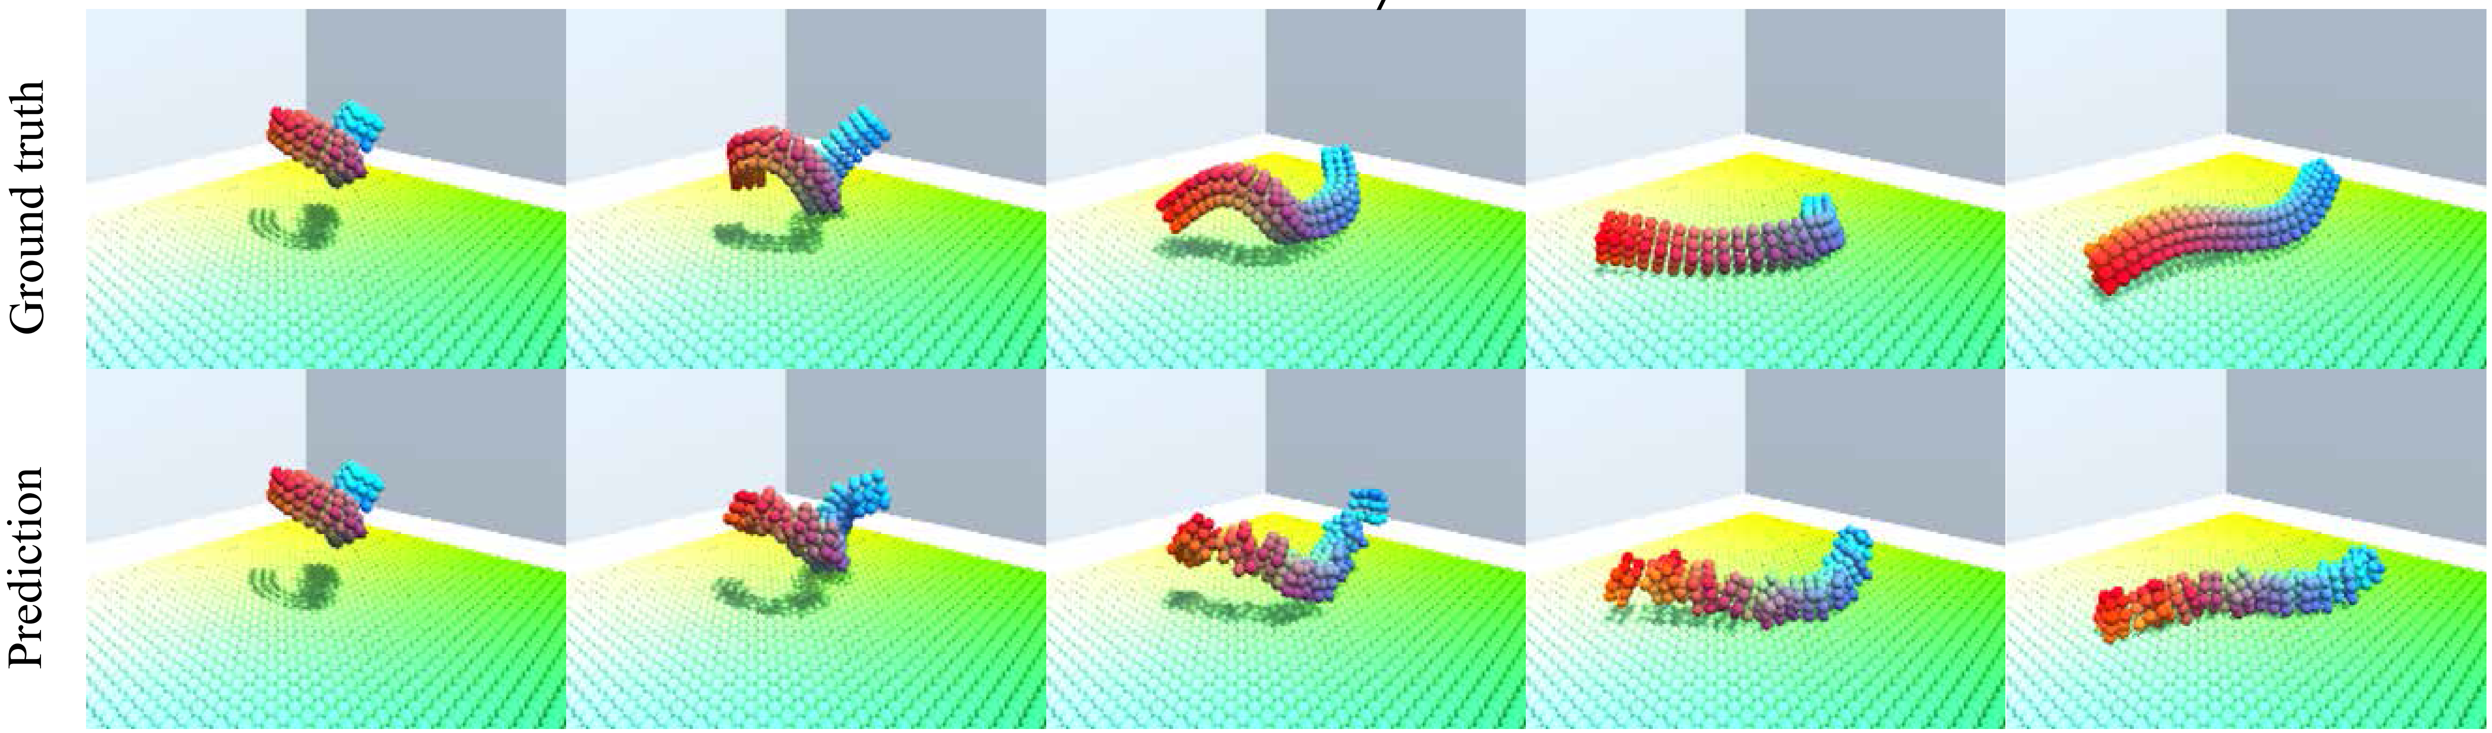

shape is not preserved super well over long rollouts…

Easy to impose simple shape conversation rules — in a "per material" way. (e.g. rigid different than cloth different than soft-body)

…less easy to understand how to do this in material-agnostic way.

Challenges:

Extracting the graph description from video.

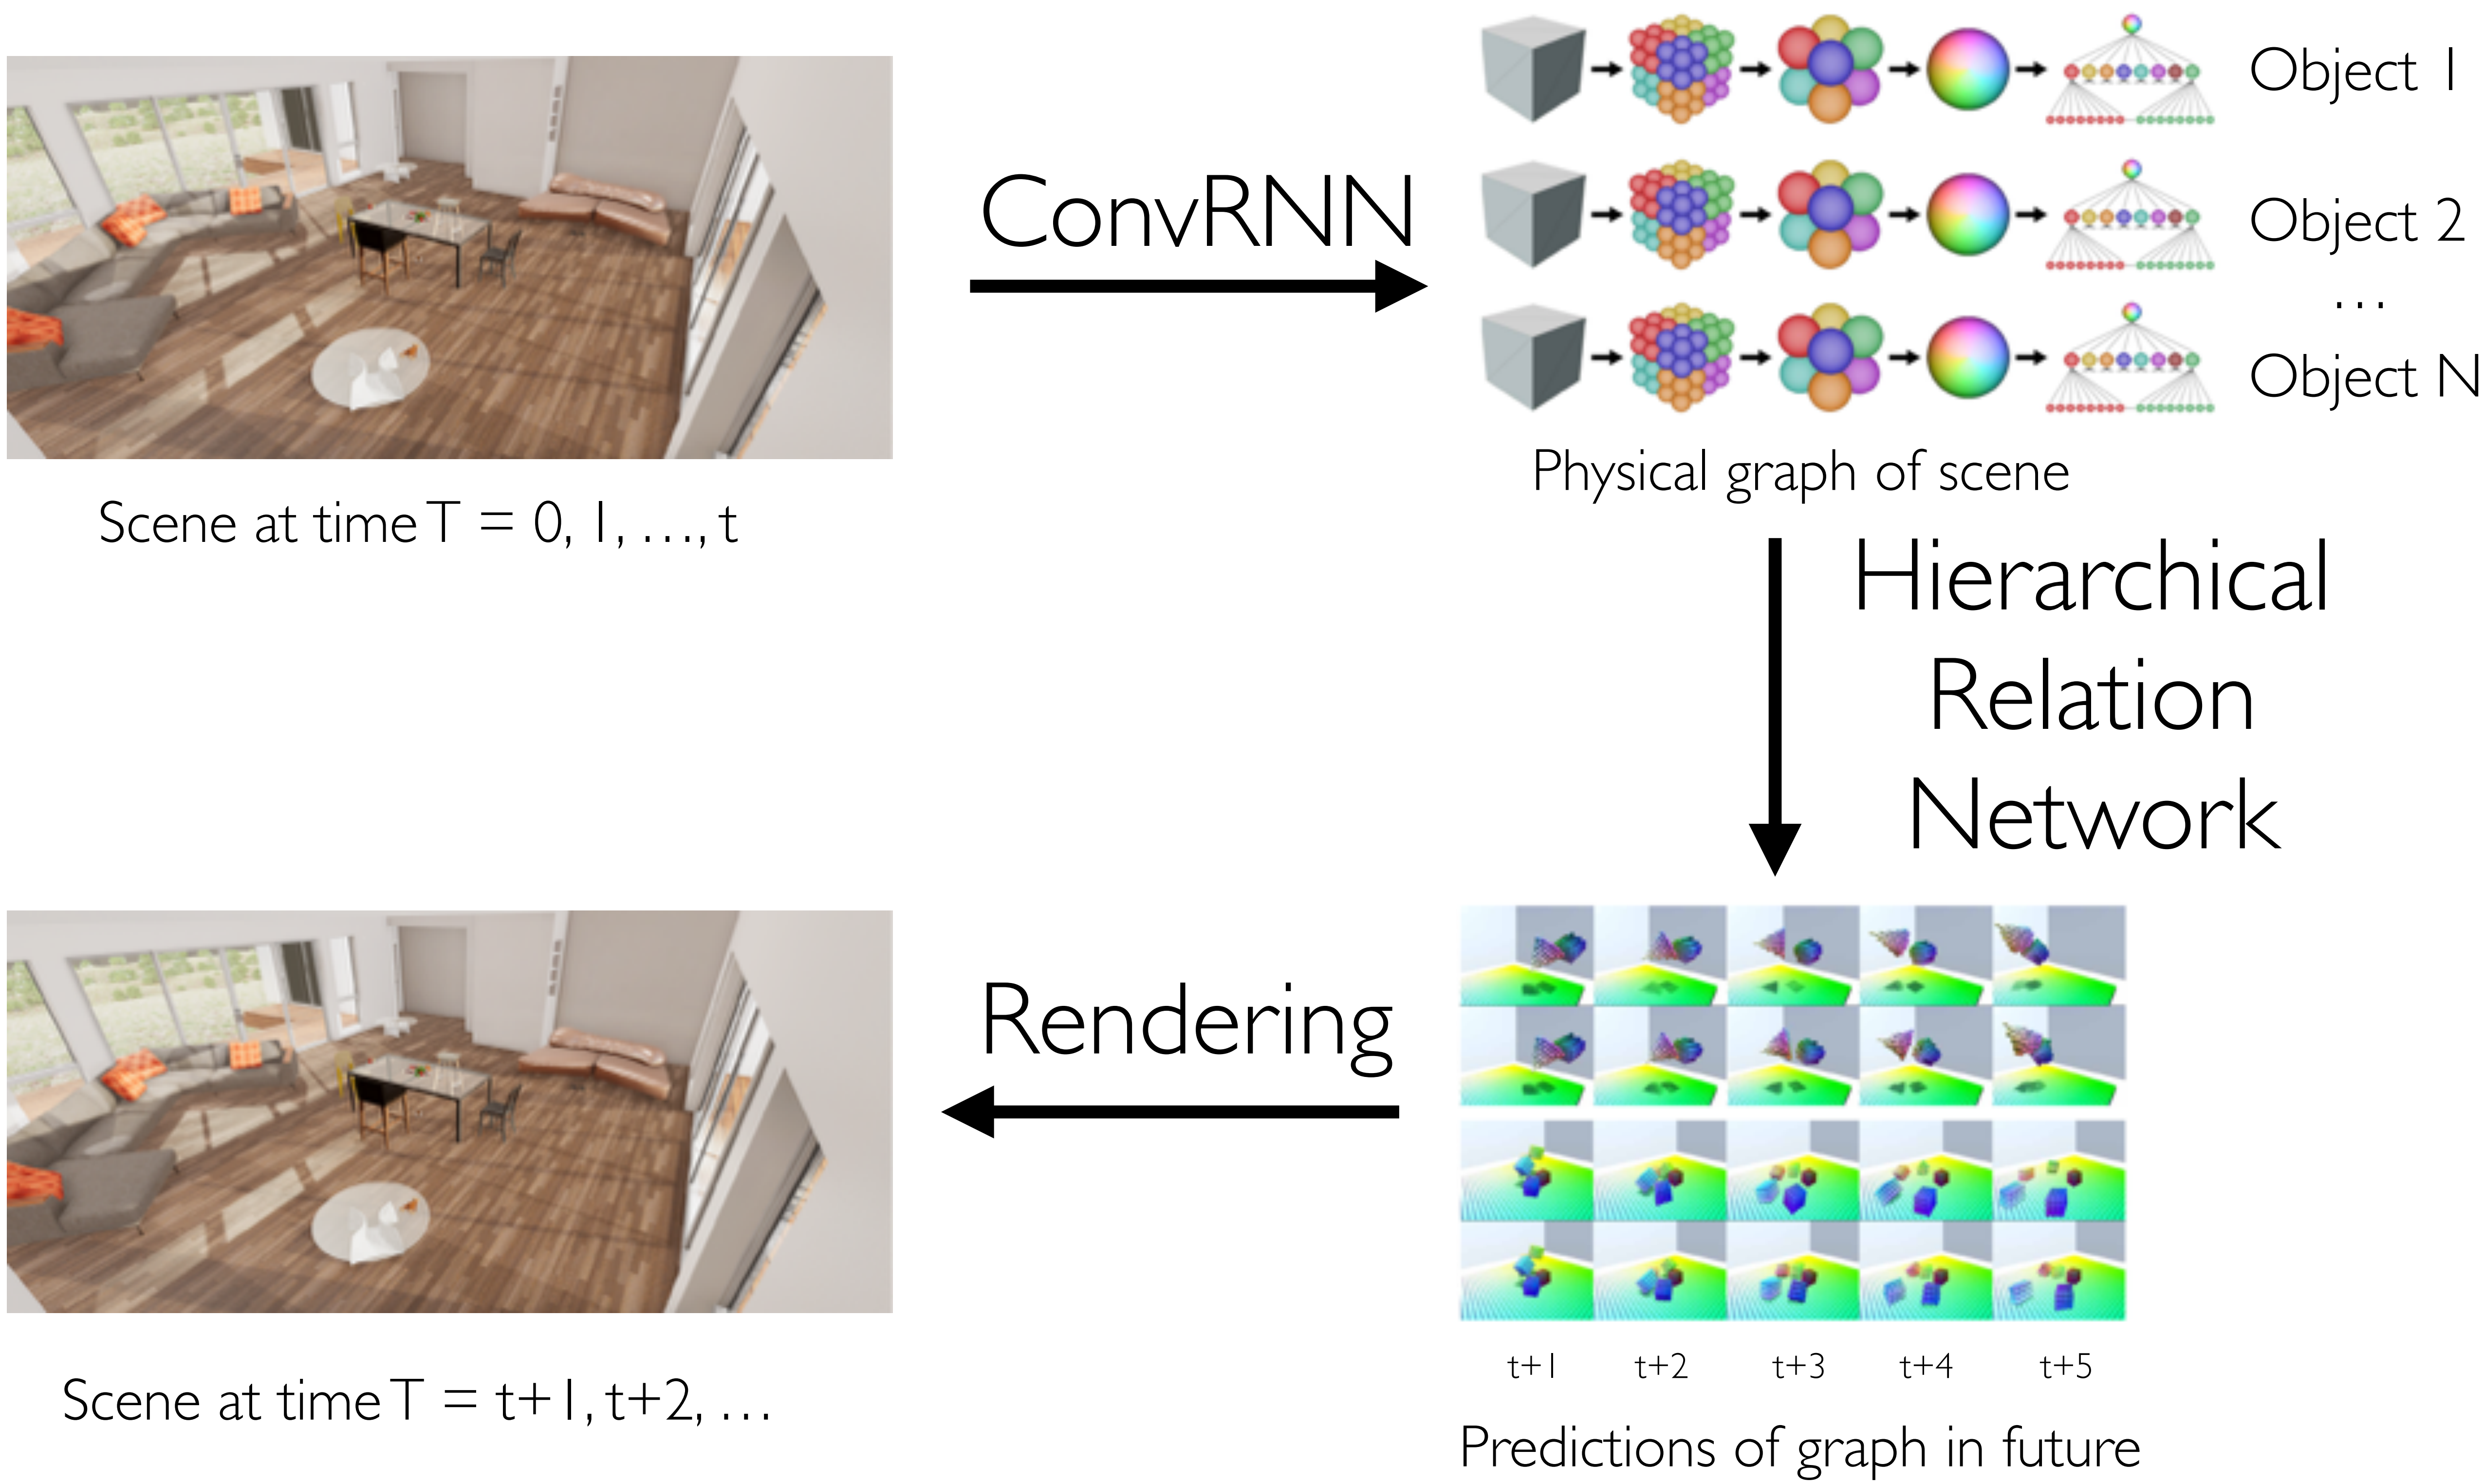

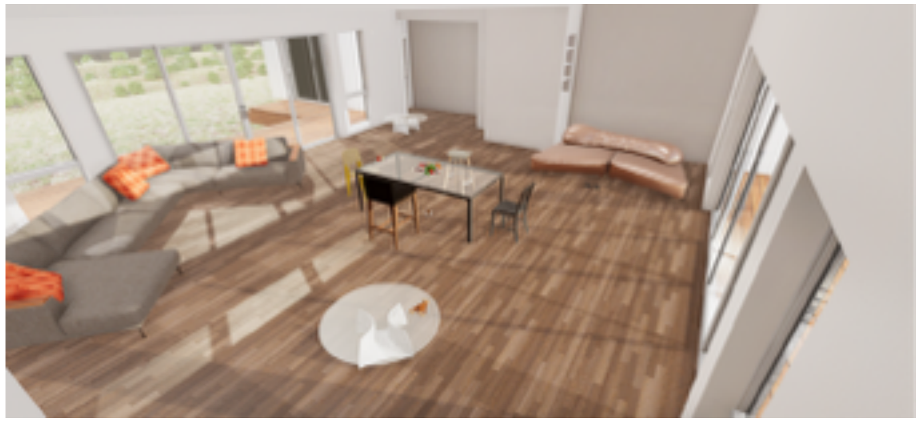

ConvRNN

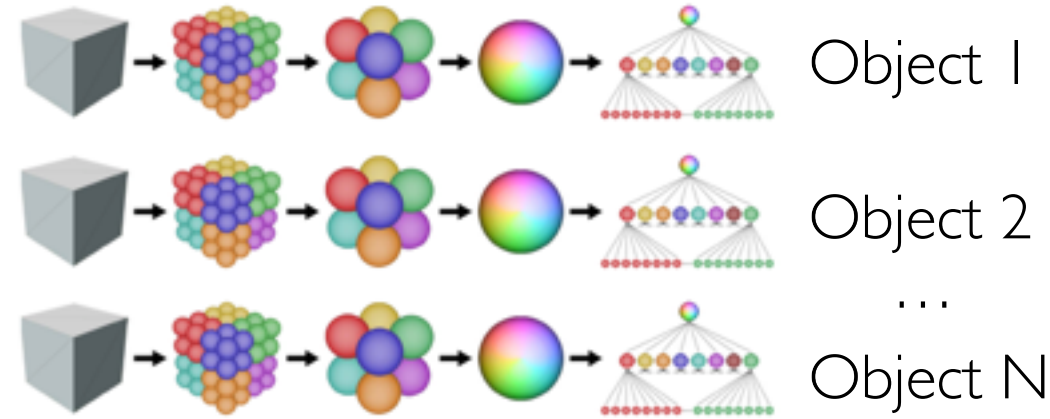

Physical graph of scene

Hierarchical Relation Network

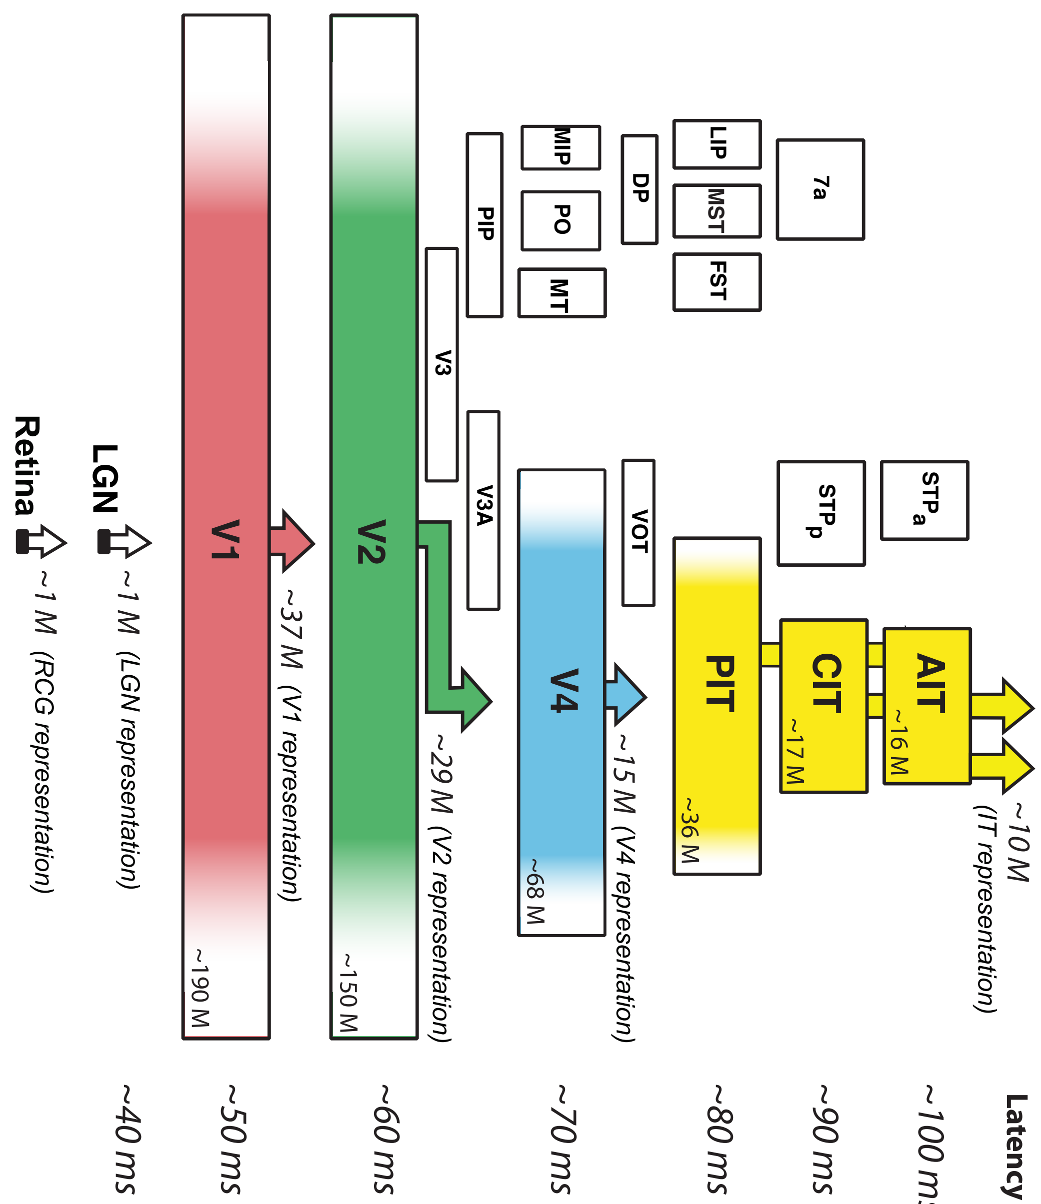

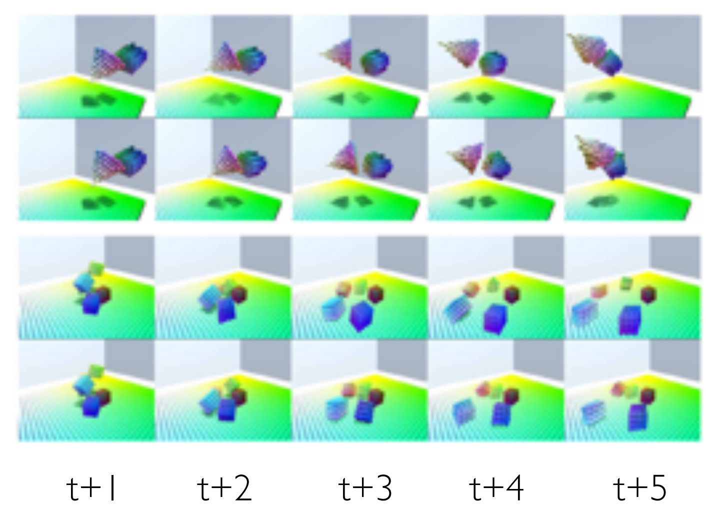

Predictions of graph in future

Human-centered feedback loop

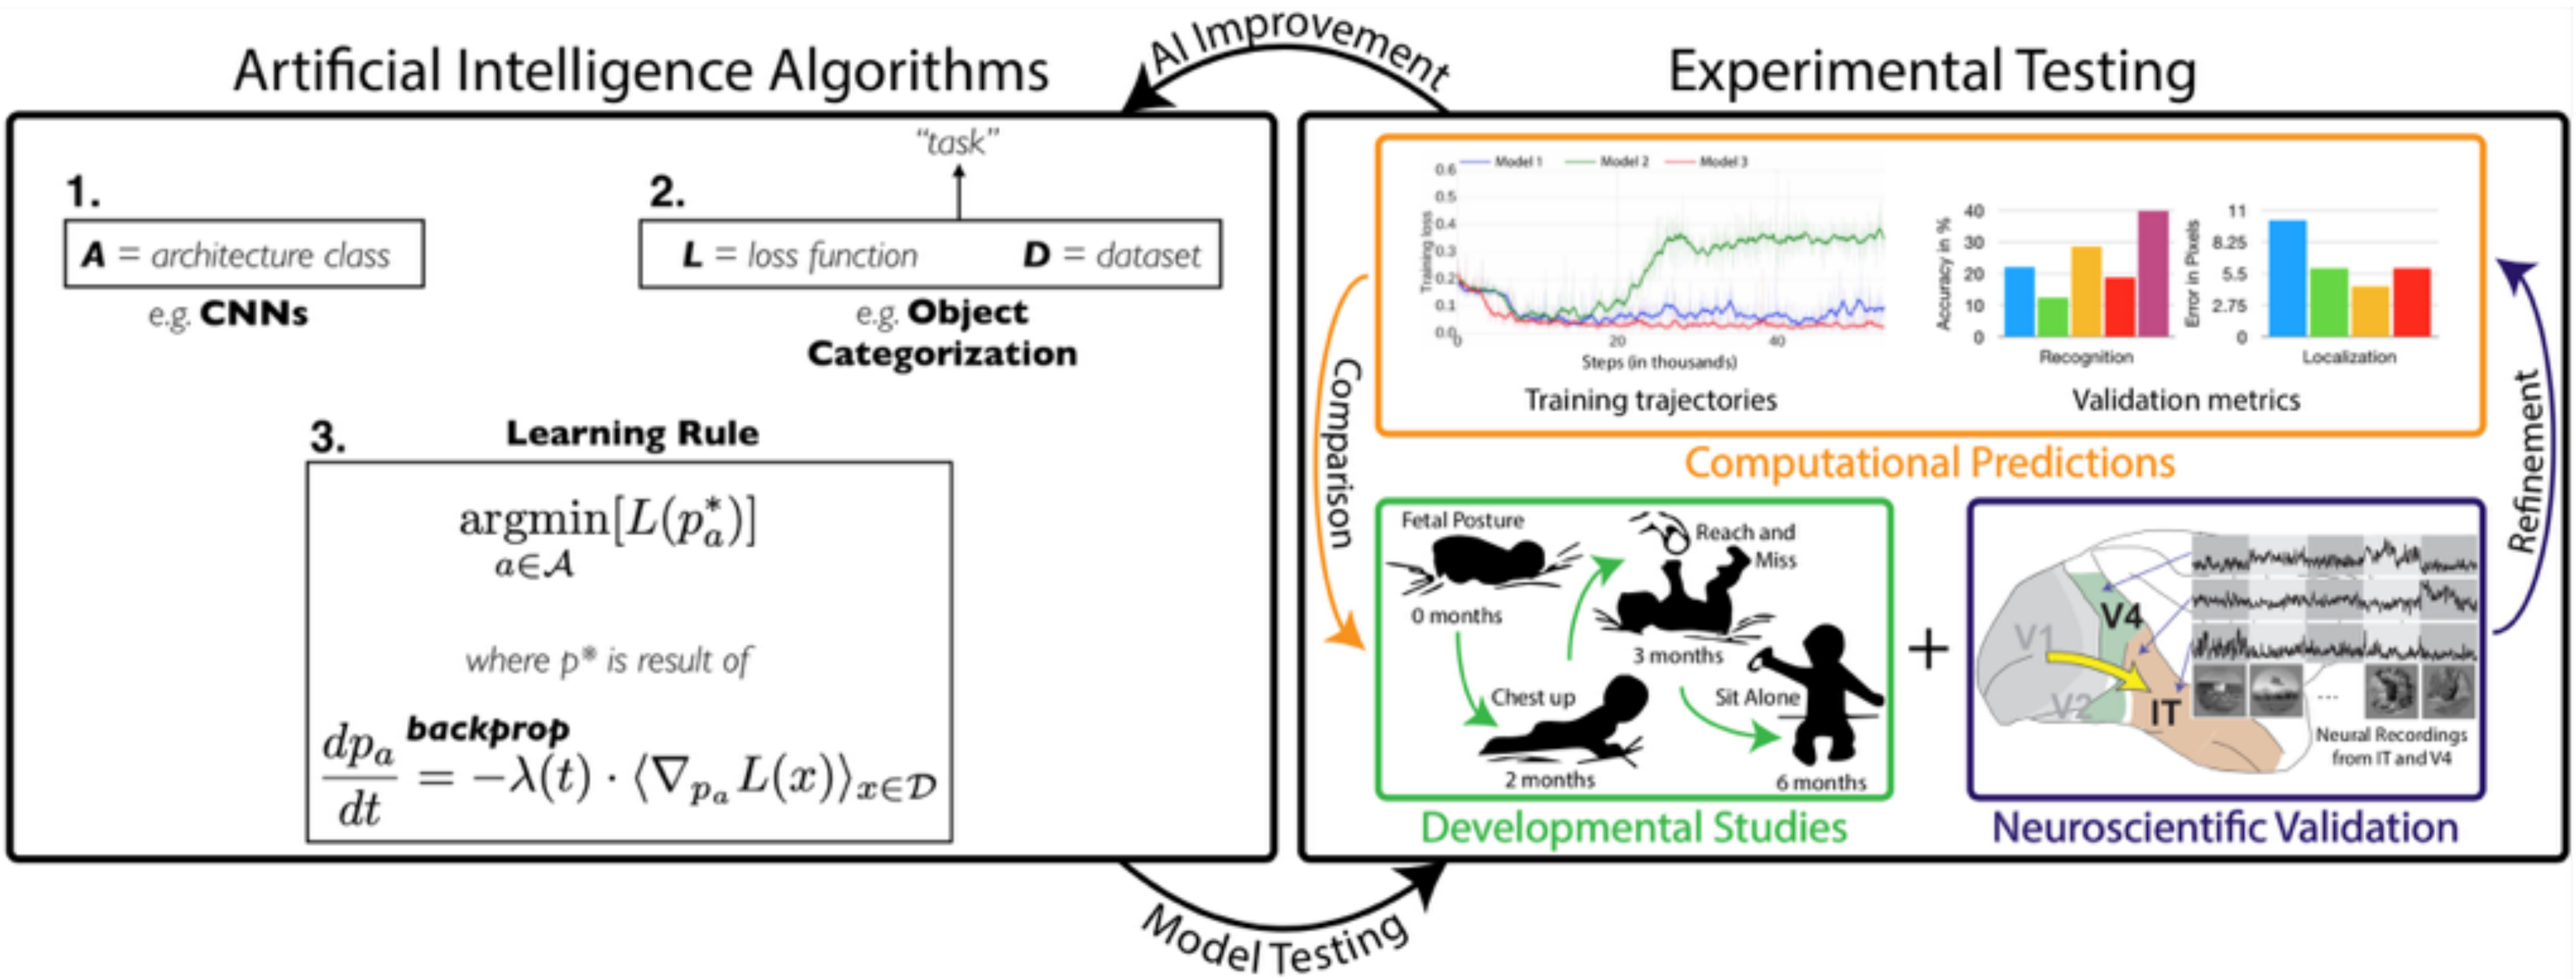

Thanks!