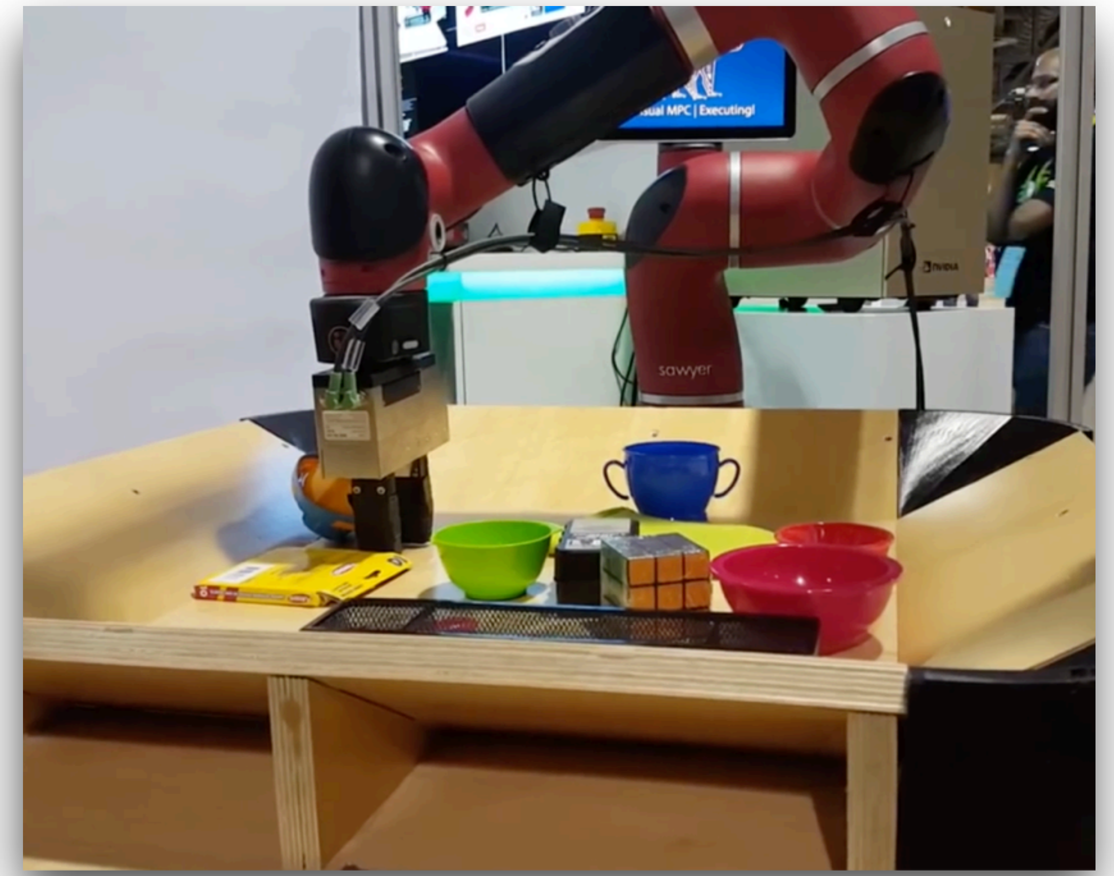

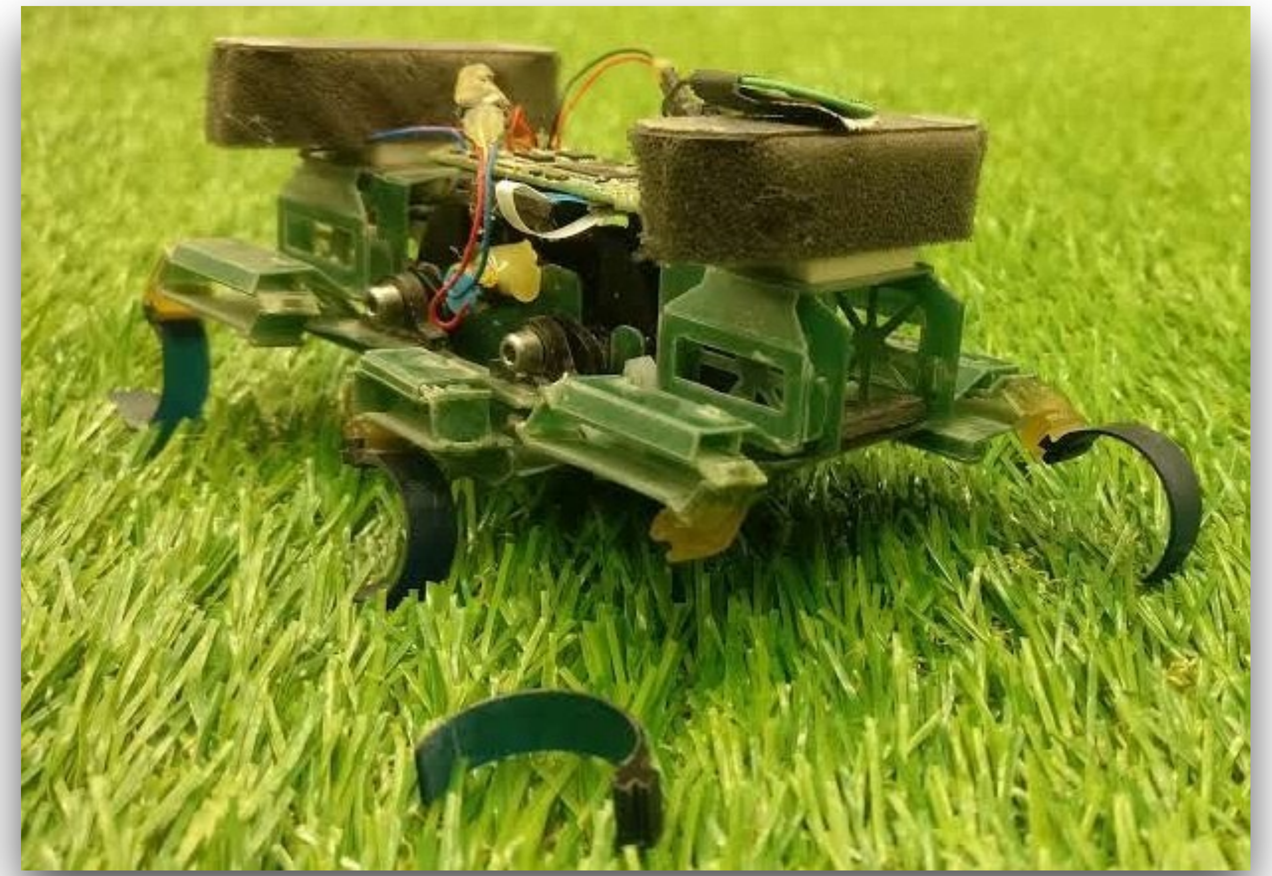

An agent that can do many things
(by modeling the world)

Chelsea Finn

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH
UC Berkeley

Google Brain

Stanford

Impressive Feats in AI

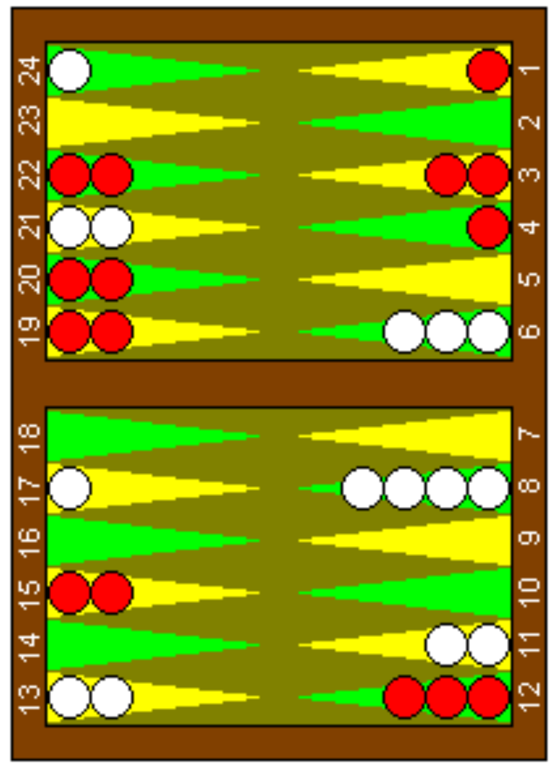

TD Gammon

Watson

helicopter acrobatics

machine translation

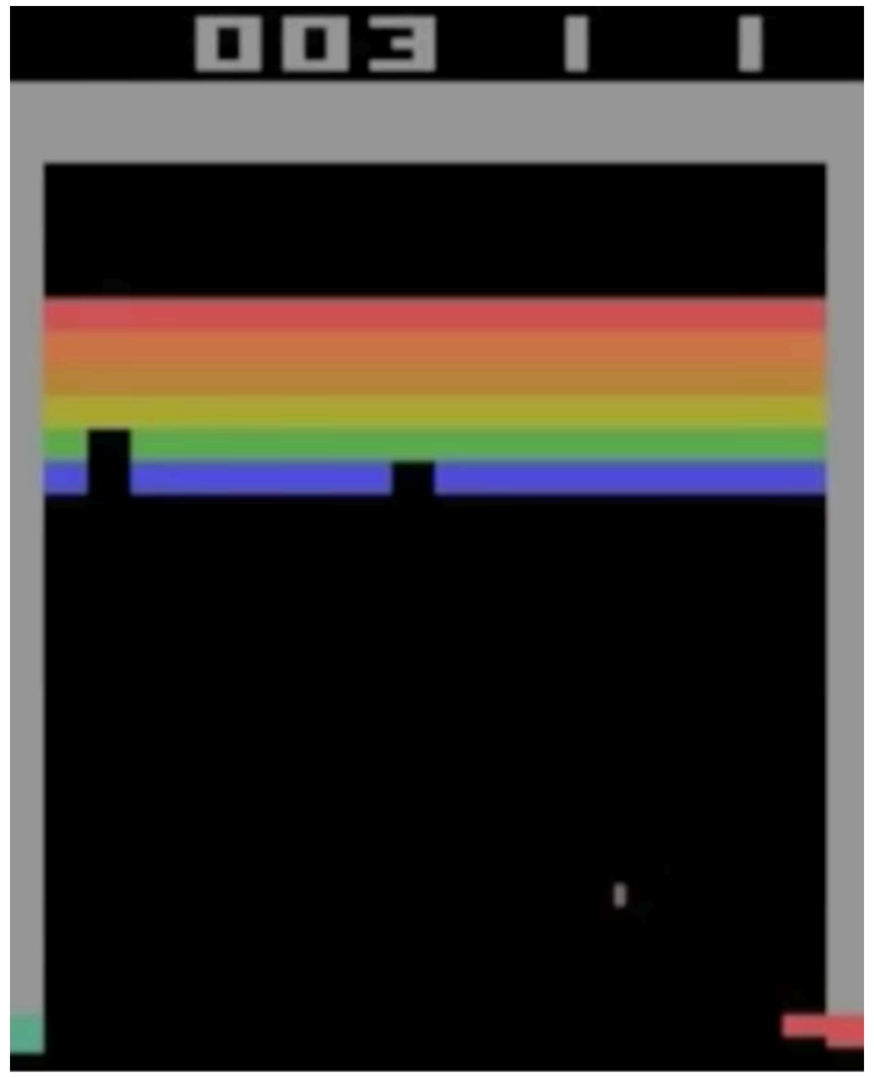

DQN

AlphaGo

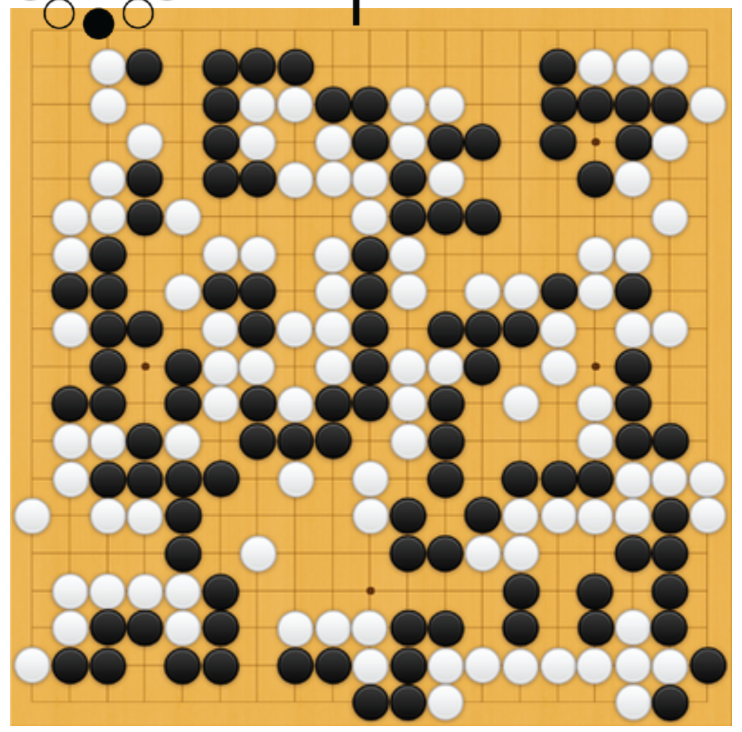

Why are these impressive?

They perform a **complex task** very well, sometimes even better than a human.

"specialists"

What is equally important:
but not impressive
(on the surface)

Generality: ability to perform many tasks

How can we build *generalists*?

Simple, yet **general**, manipulation skills are beyond the scope of current methods.

It turns out — the **simpler**, but **broader** capabilities are **really hard**.
(Moravec's Paradox)

This talk: can we do the **unimpressive** things?

Can we build an agent that can do **many tasks**?

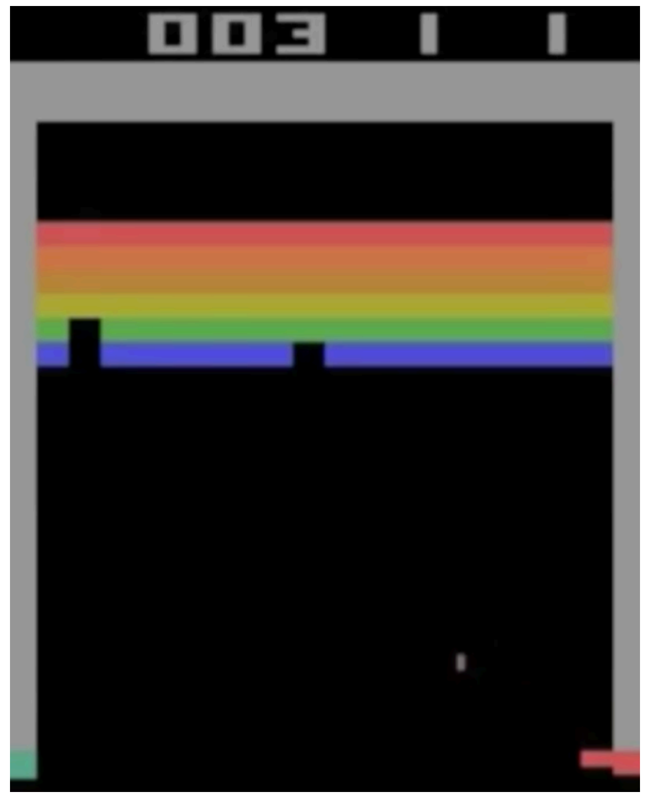

learning a **policy** in
a **closed universe**

learn **general-purpose** model
+
plan with model **for many tasks**

model-based control

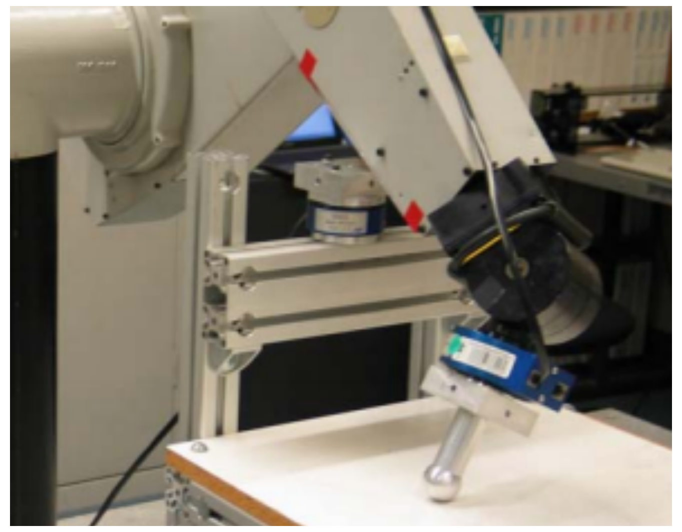

Petrovskaya, Park, Khatib '07

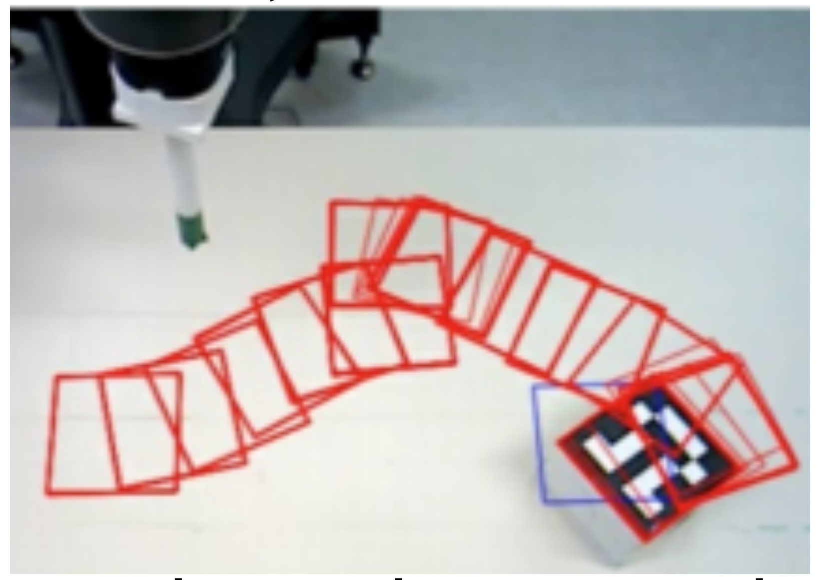

Arruda, Mathew, Kopicki,
Mistry, Azad, Wyatt '17

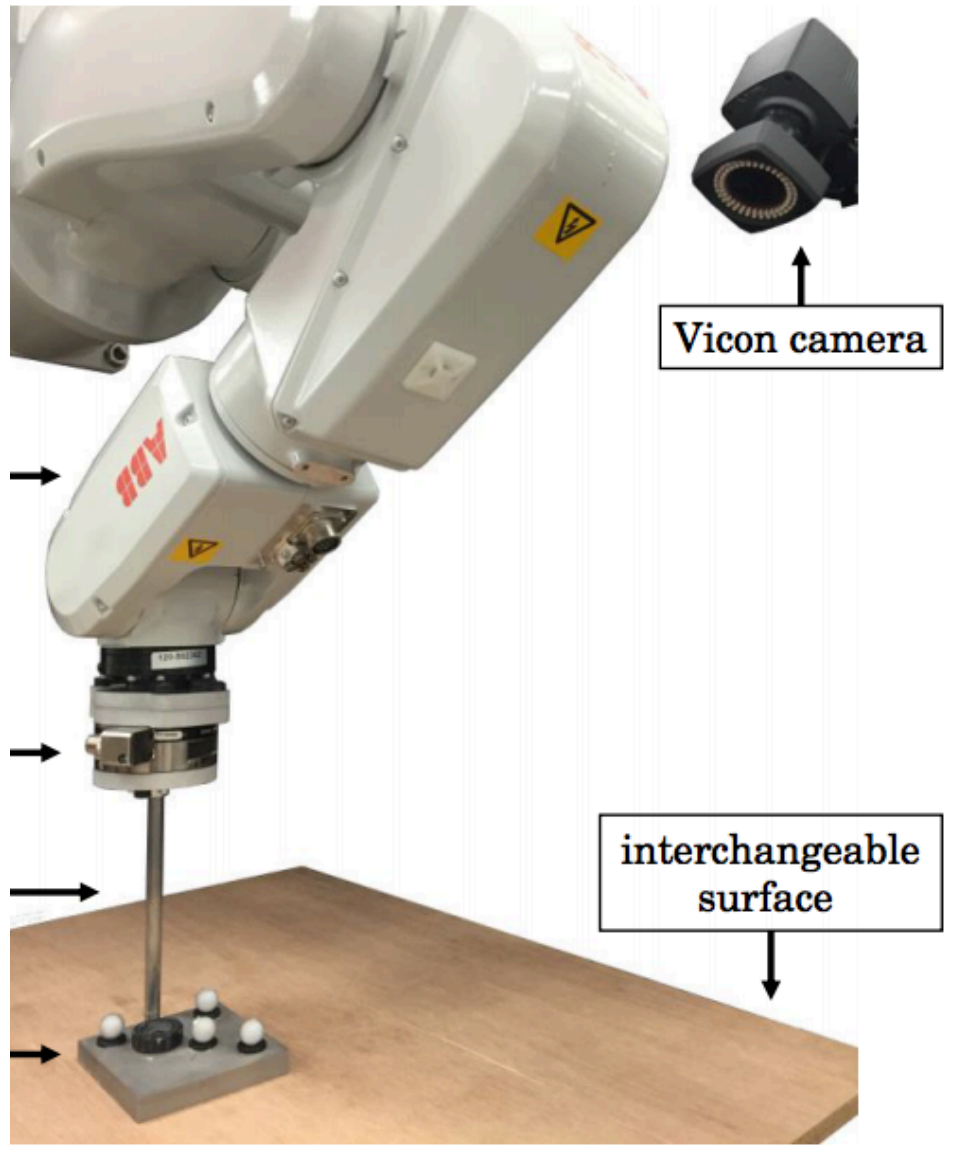

Yu, Bauza, Fazeli, Rodriguez '17

from **pixel observations**, with **limited supervision**, in the **physical world**

Can we build an agent that can do **many tasks**?

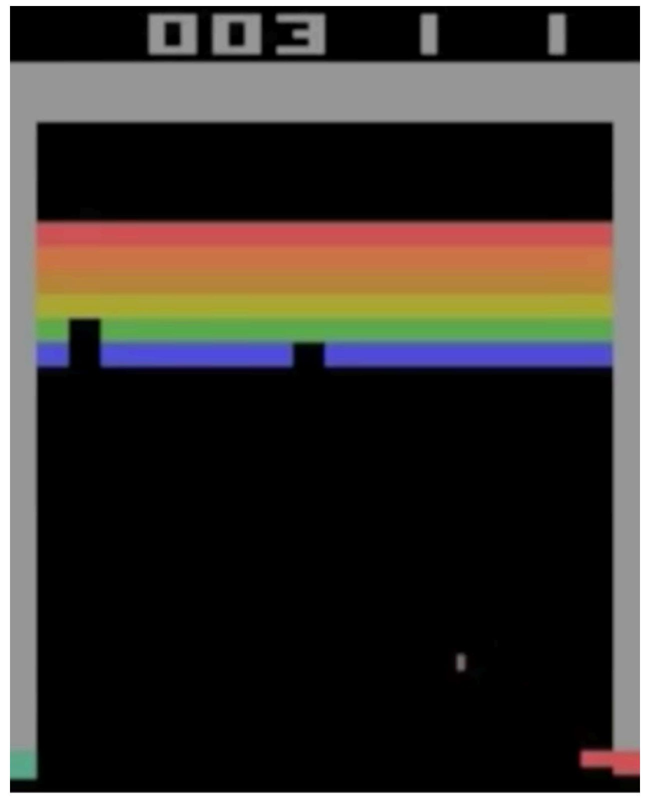

learning a **policy** in a **closed universe**

learn **general-purpose** model
+
plan with model **for many tasks**

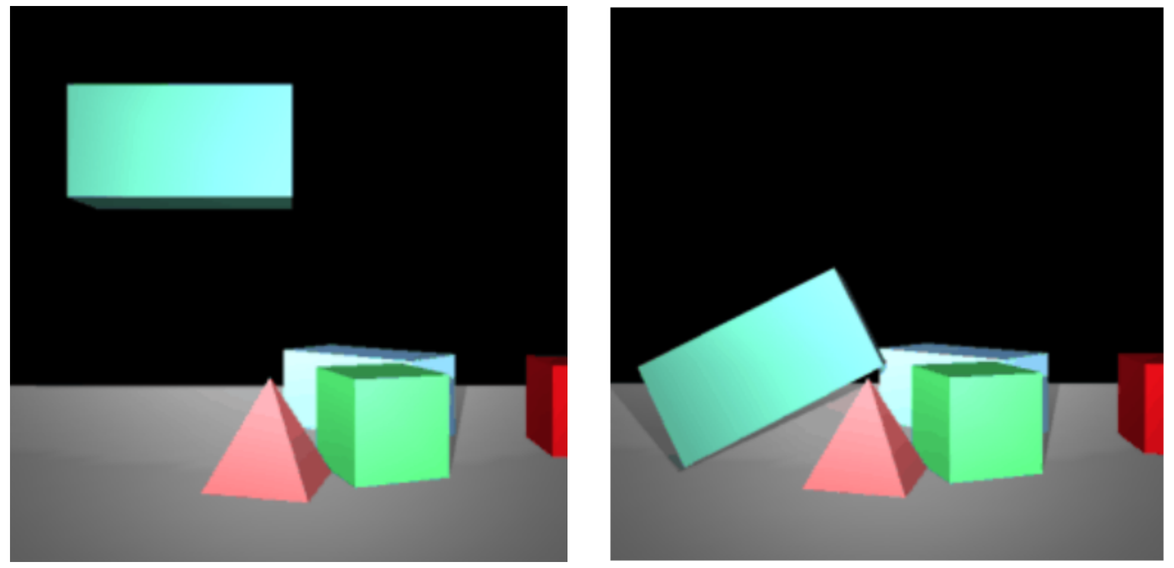

structured latent space model for long-horizon tasks

modeling **diverse**, **open-world** environments

from **pixel observations**, with **limited supervision**, in the **physical world**

1. Collect **diverse** interactions

Greater diversity —> more general-purpose model

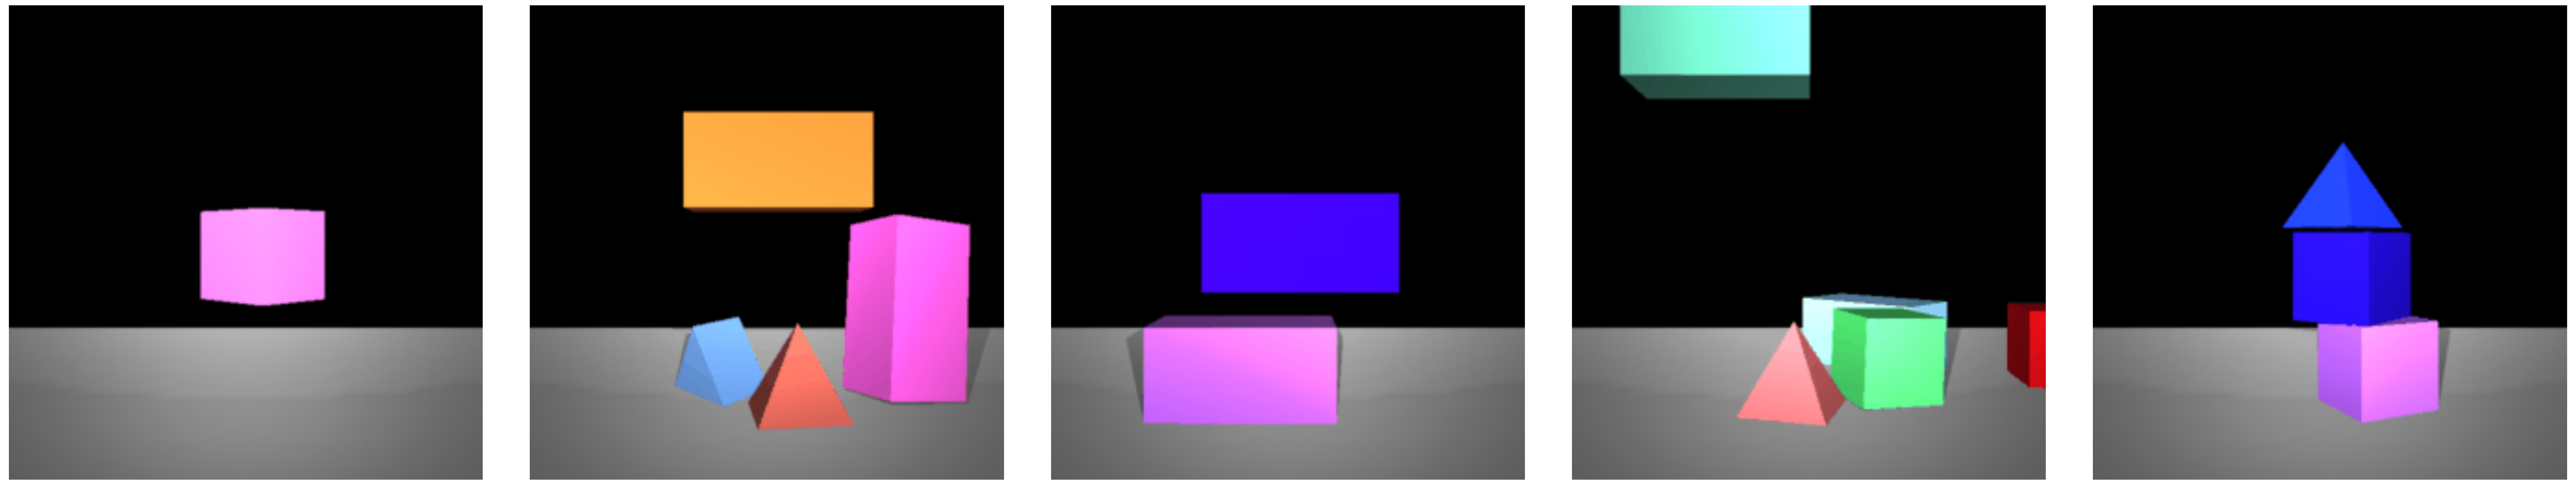

2. Learn **structured** representation & model

Structure —> long-horizon reasoning

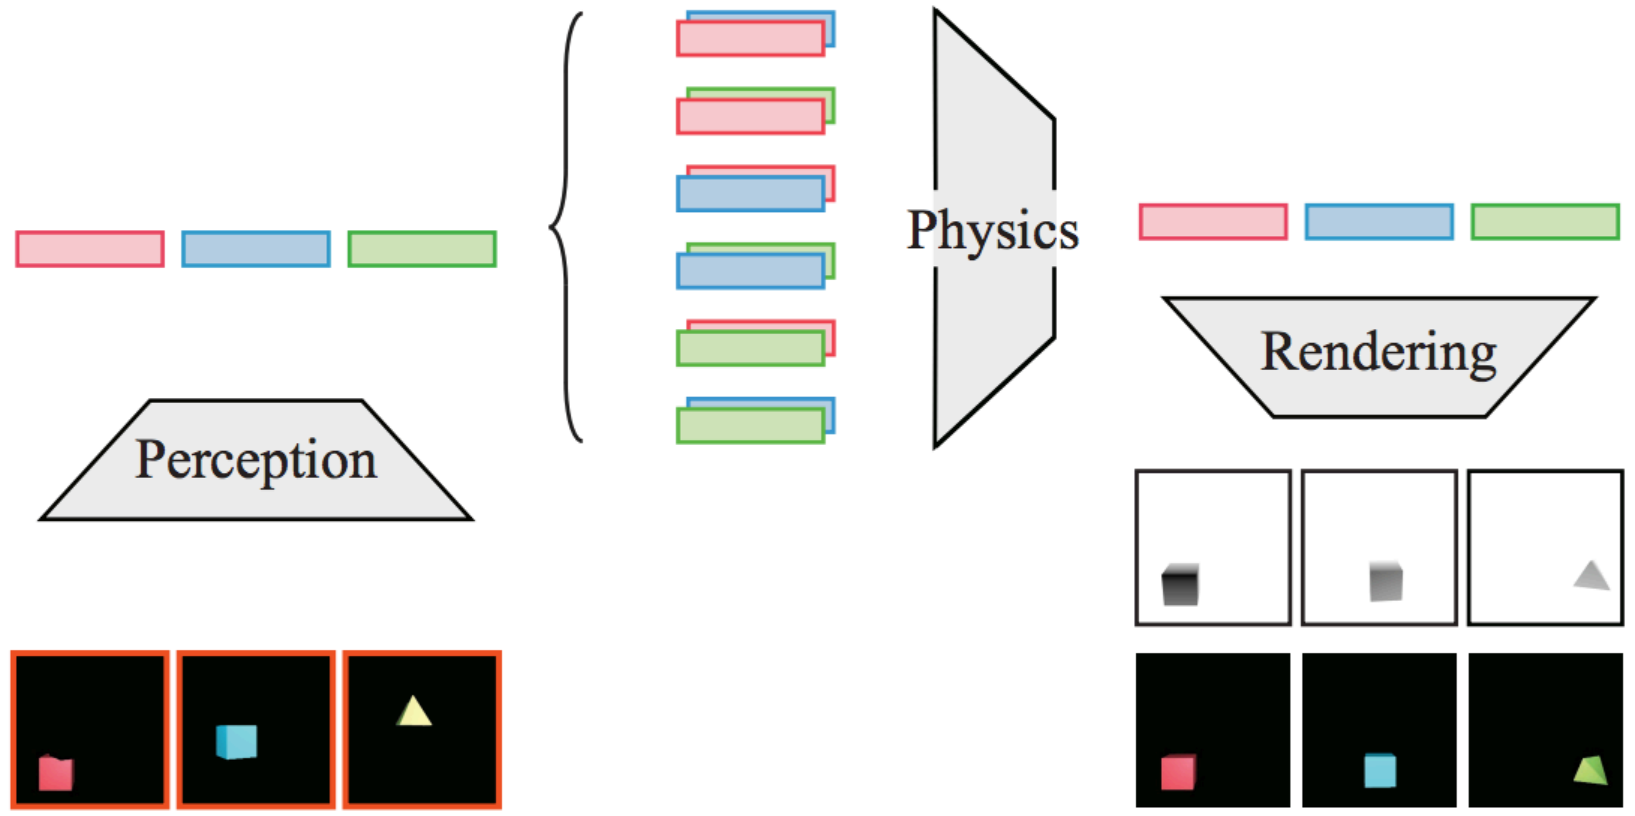

3. **Plan** using model

Online planning —> many tasks

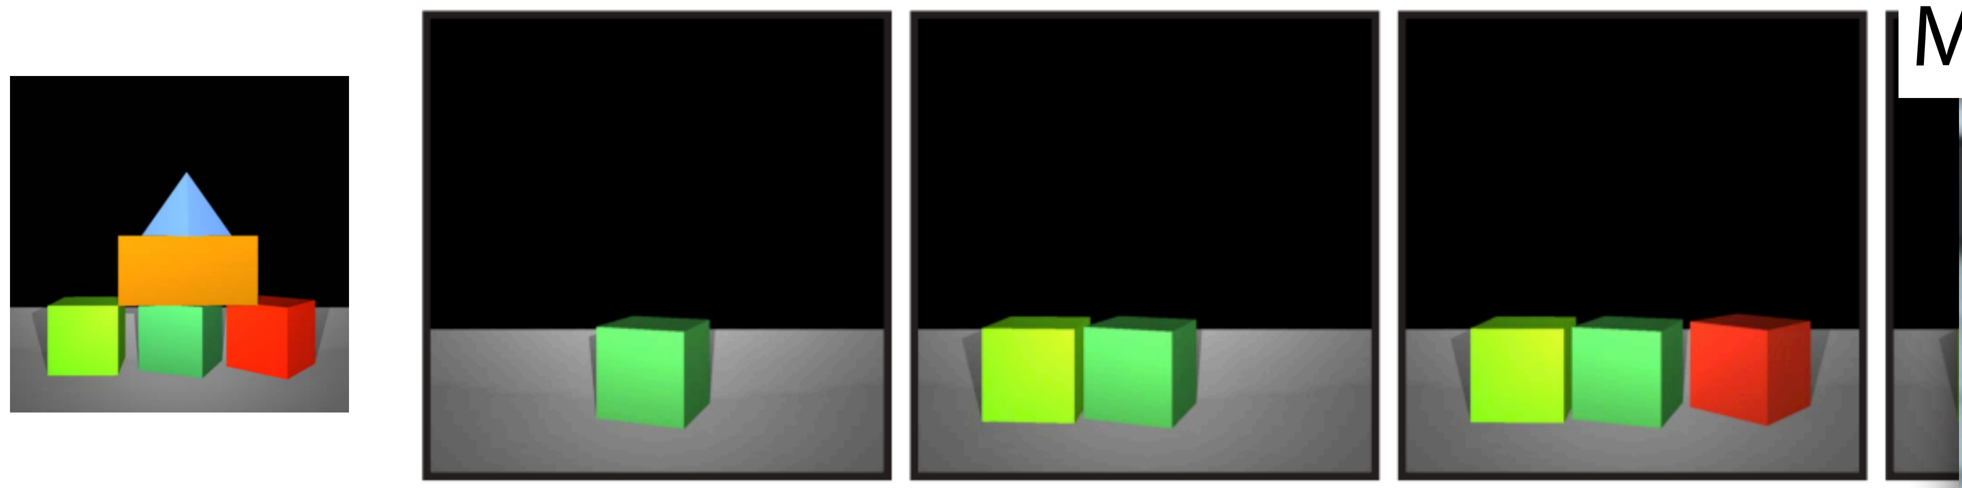

Goal: be able to build any tower c

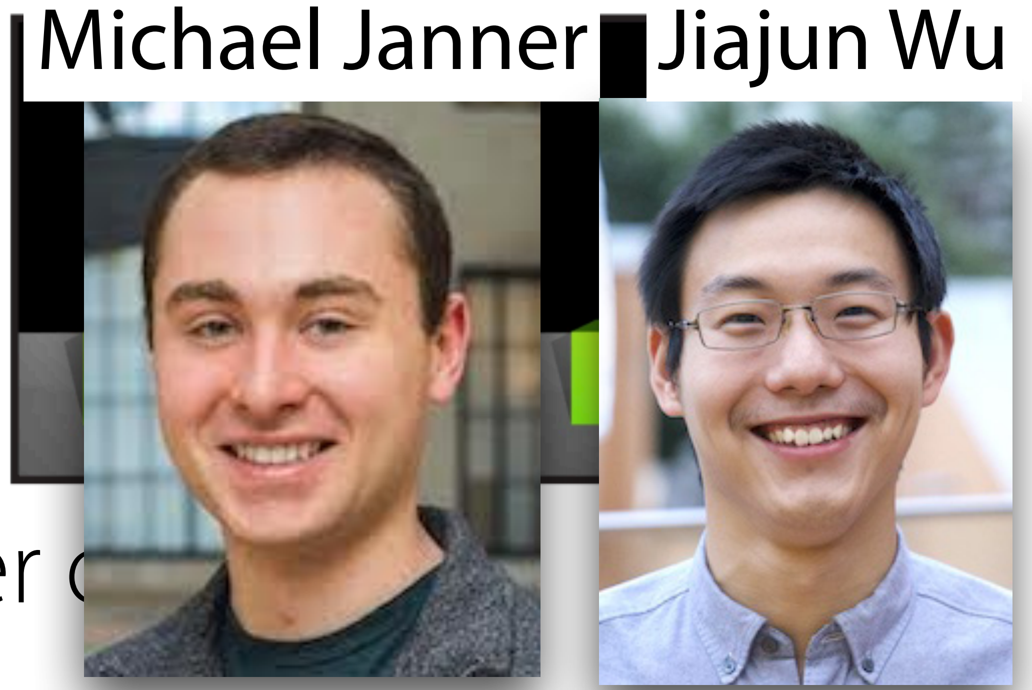

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Eslami et al. '16,
Kosiorek et al. '18

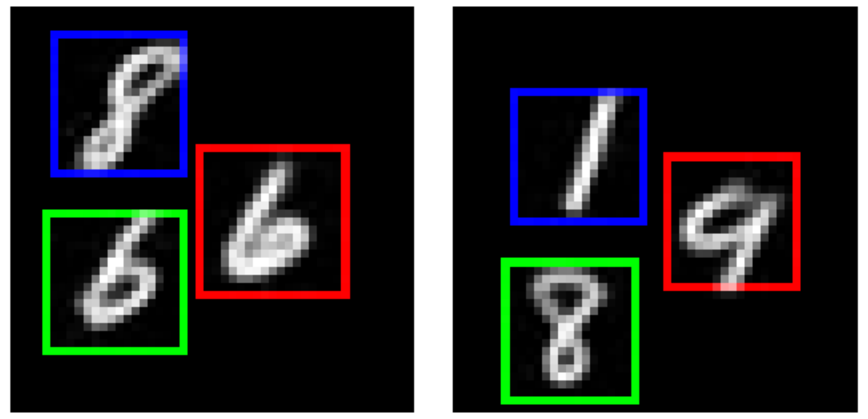

simple, 2D scenes

Wu et al. '17

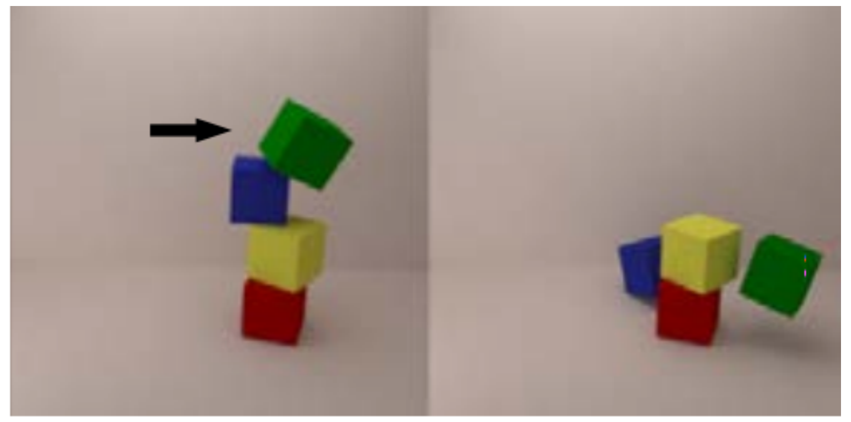

full supervision of
object properties

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

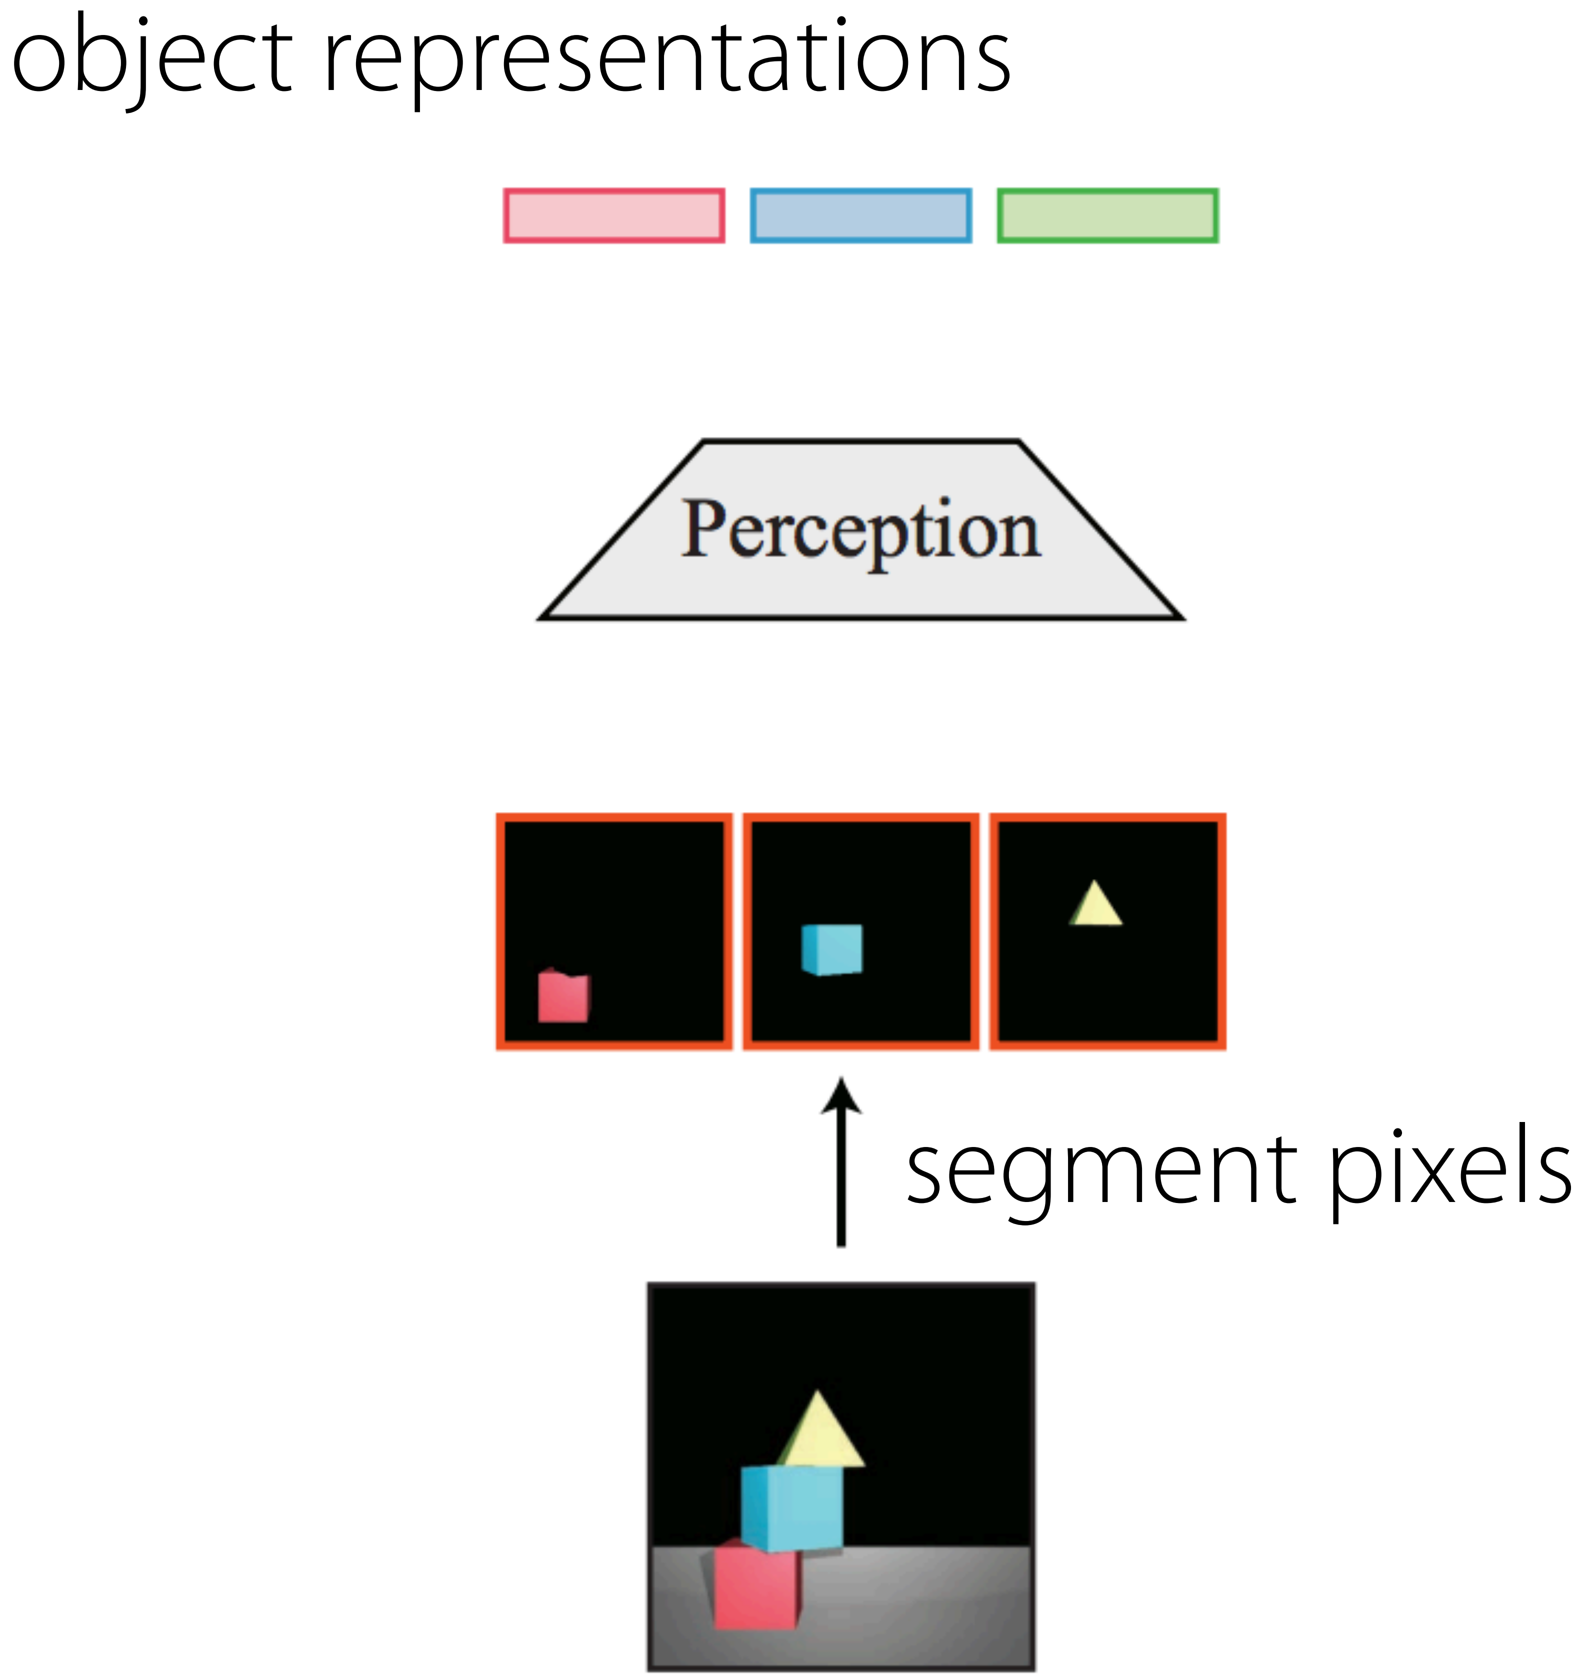

Eslami et al. '16,
Kosiorek et al. '18

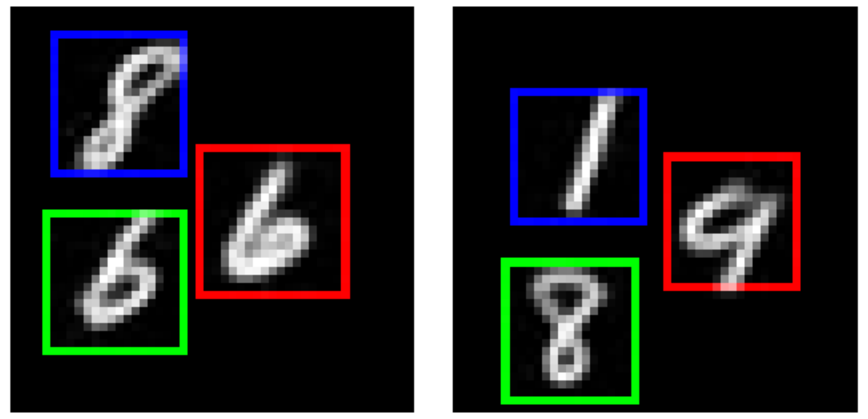

simple, 2D scenes

Wu et al. '17

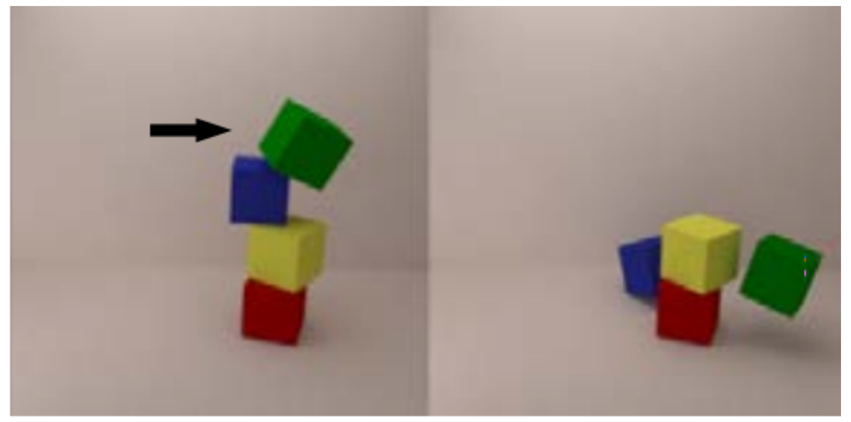

full supervision of
object properties

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Eslami et al. '16,
Kosiorek et al. '18

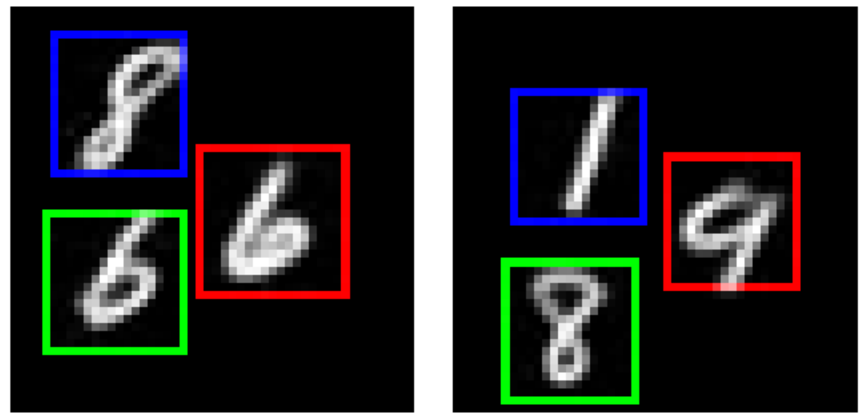

simple, 2D scenes

Wu et al. '17

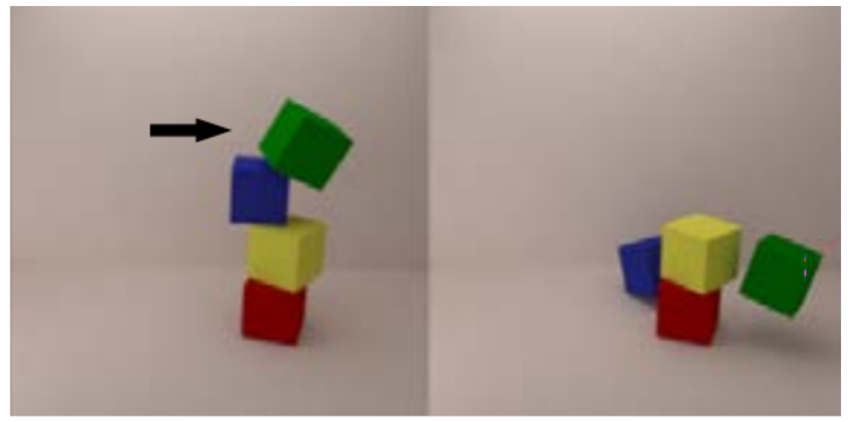

full supervision of
object properties

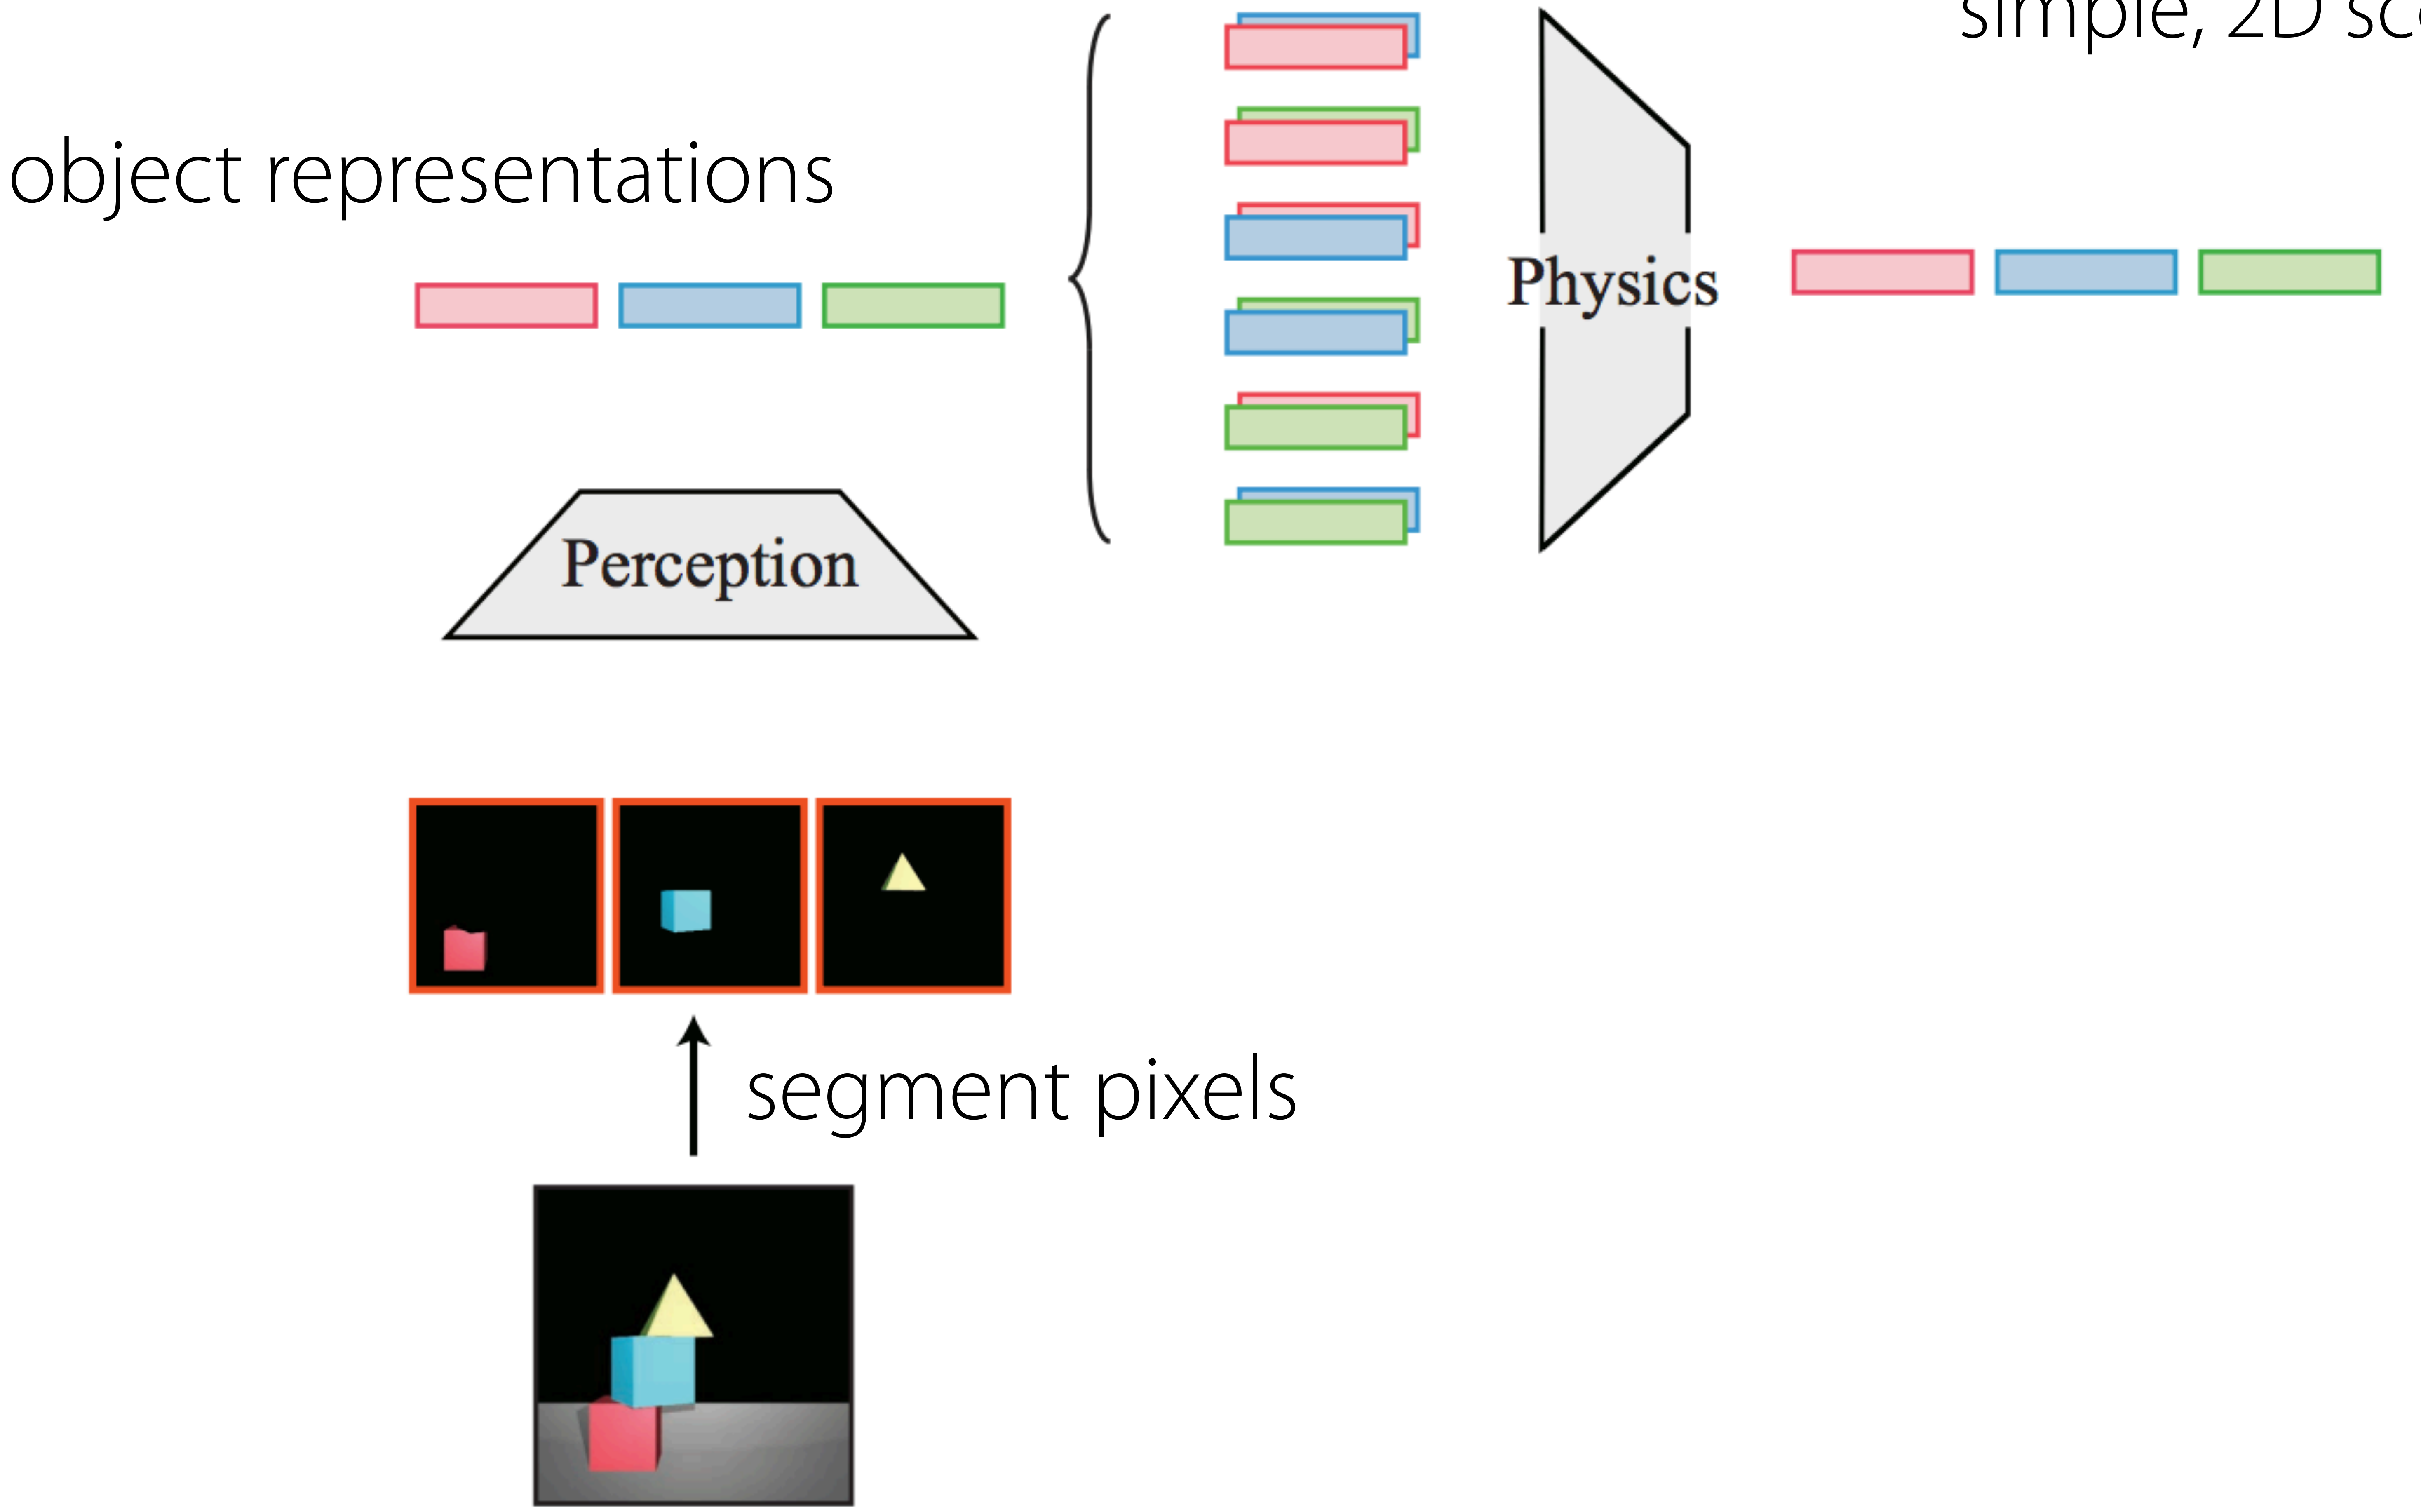

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Eslami et al. '16,
Kosiorek et al. '18

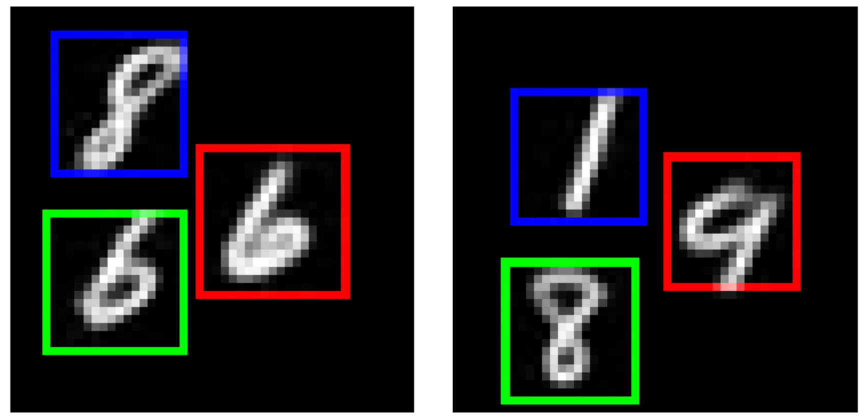

simple, 2D scenes

Wu et al. '17

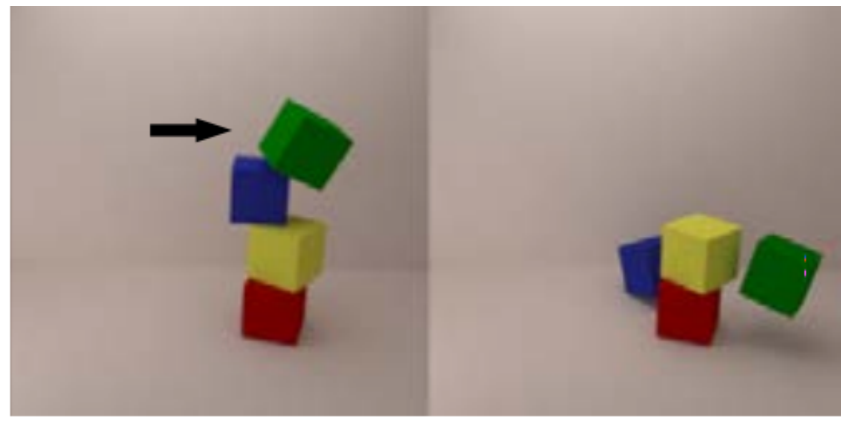

full supervision of object properties

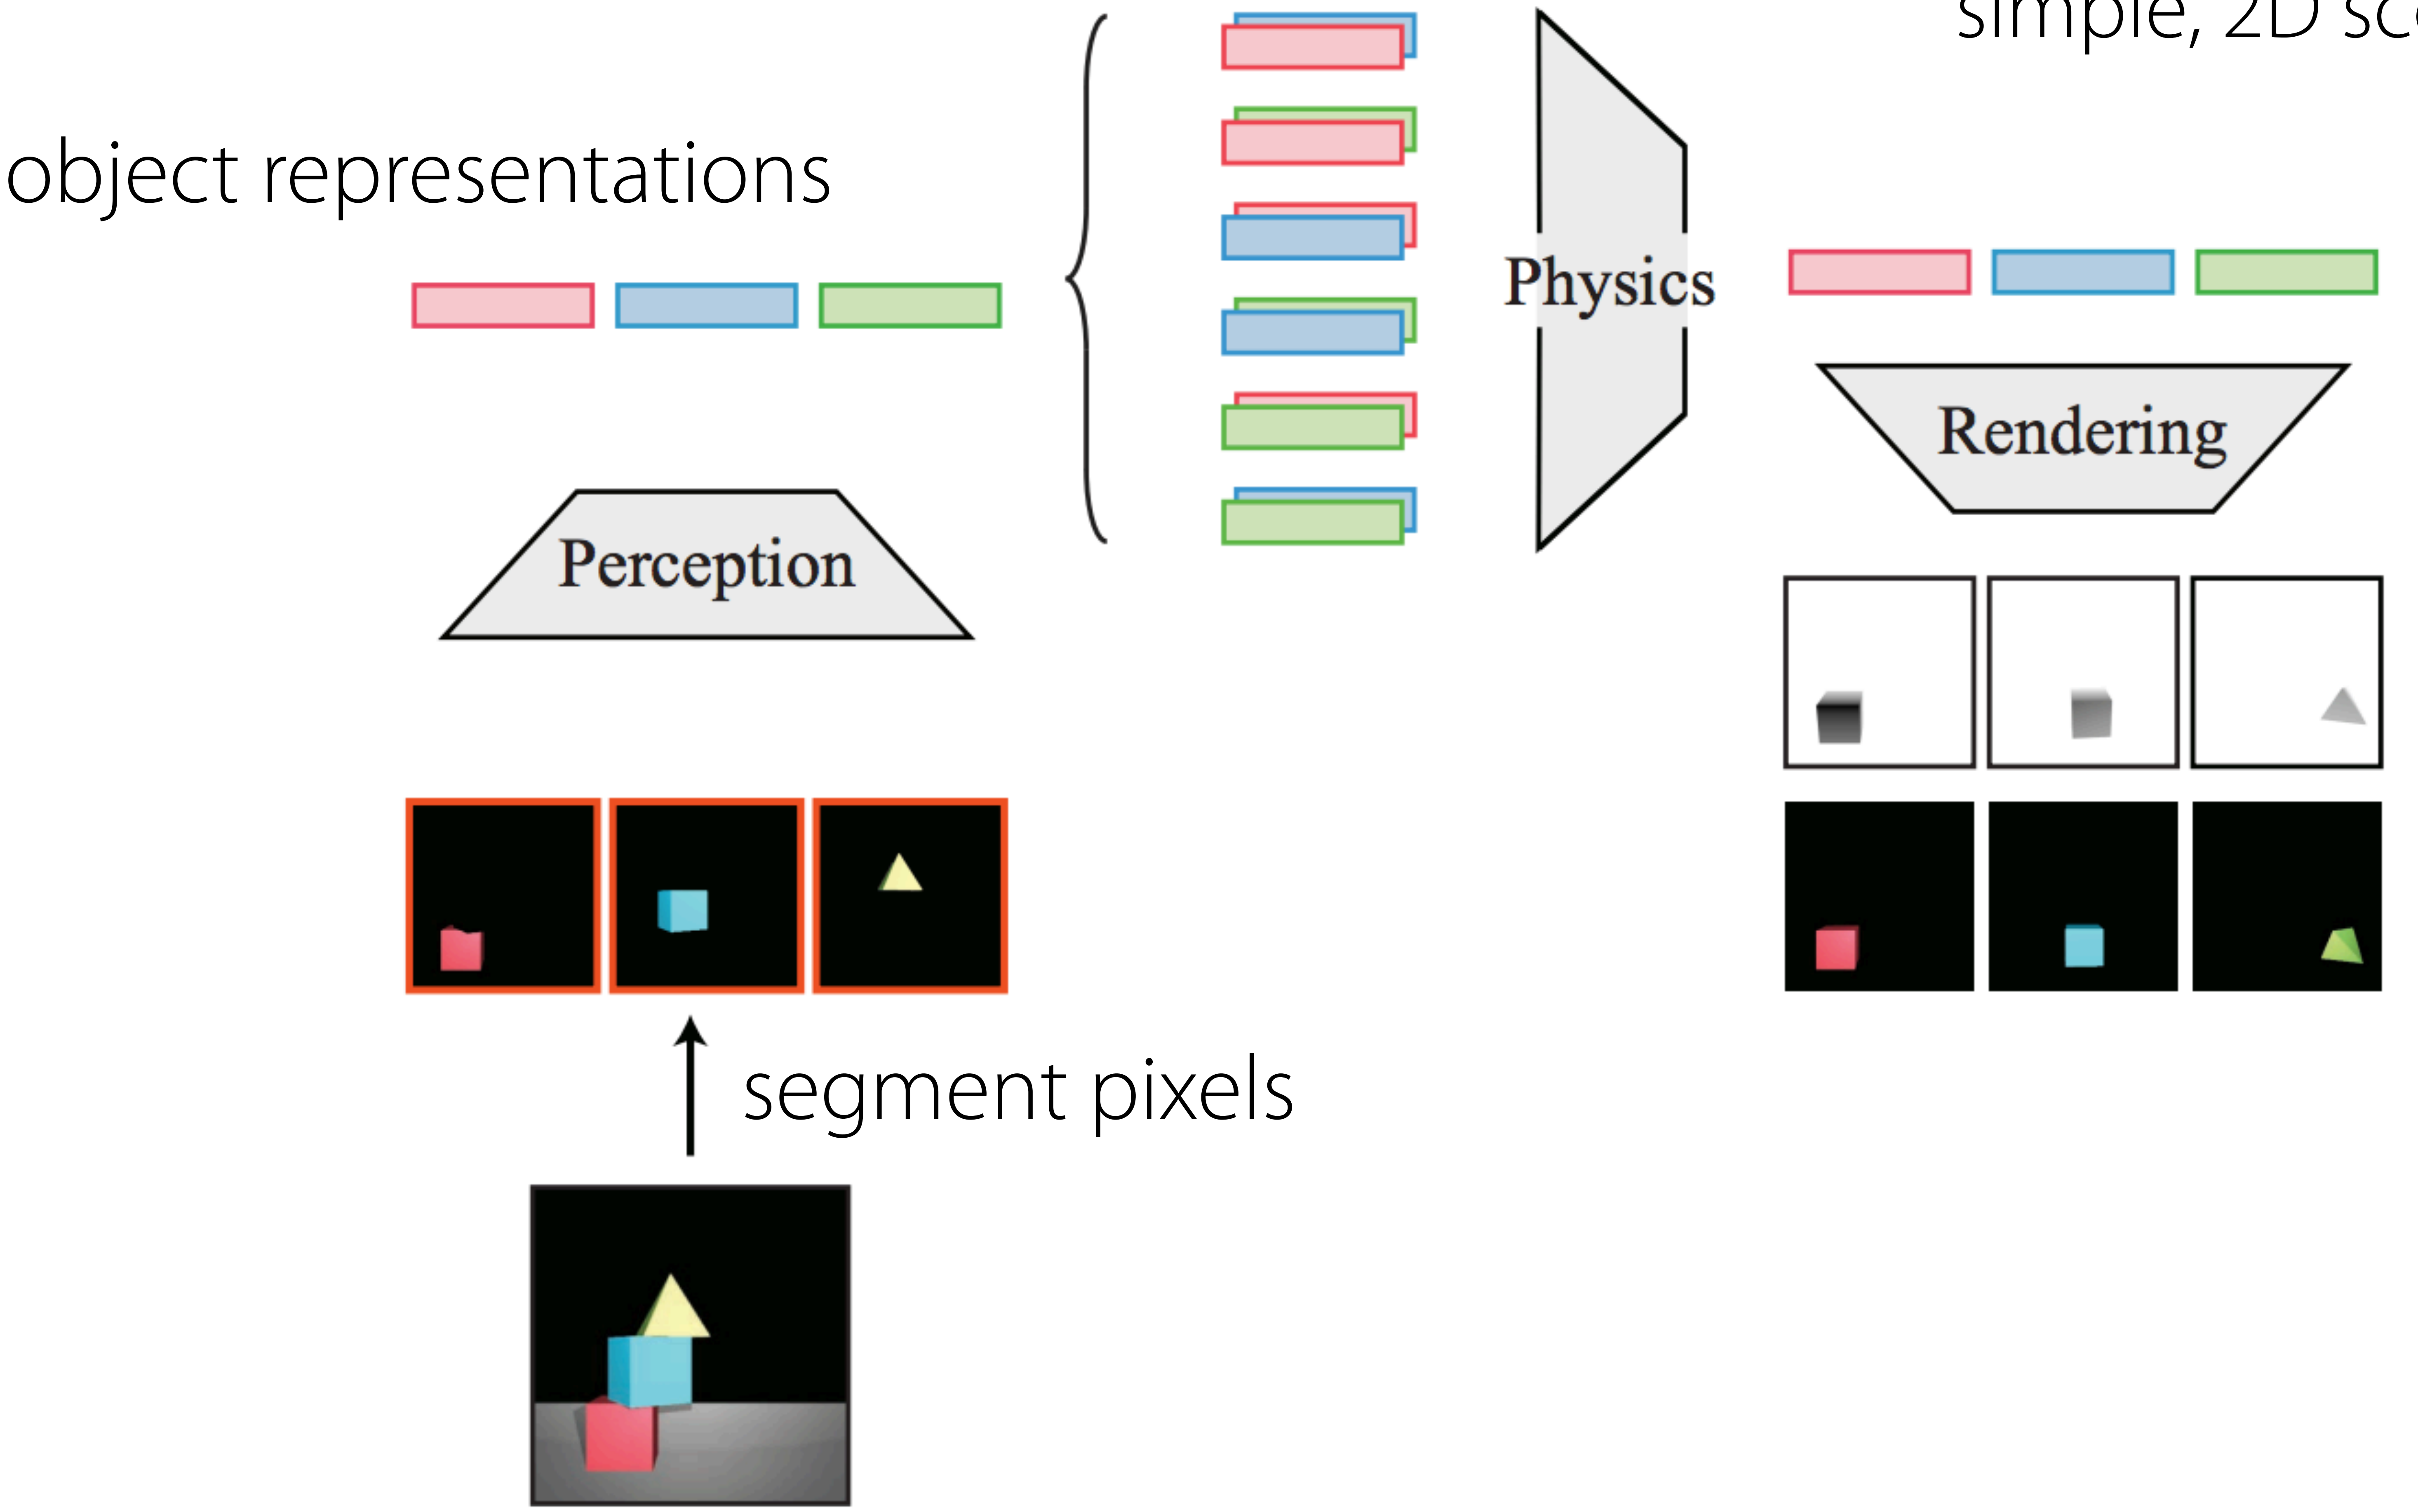

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Eslami et al. '16,
Kosiorek et al. '18

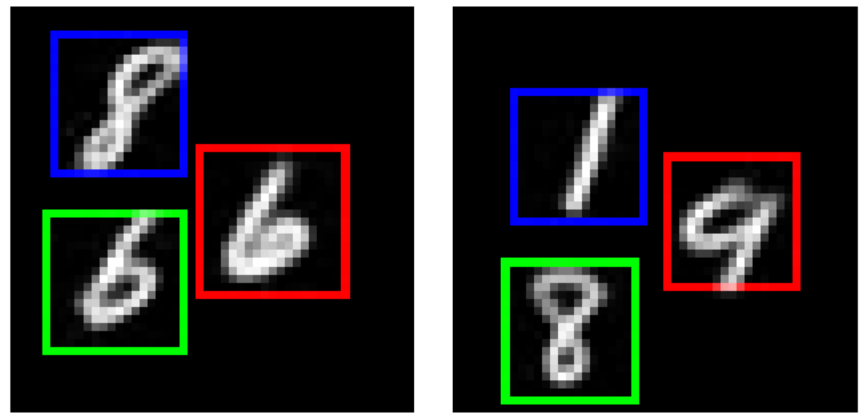

simple, 2D scenes

Wu et al. '17

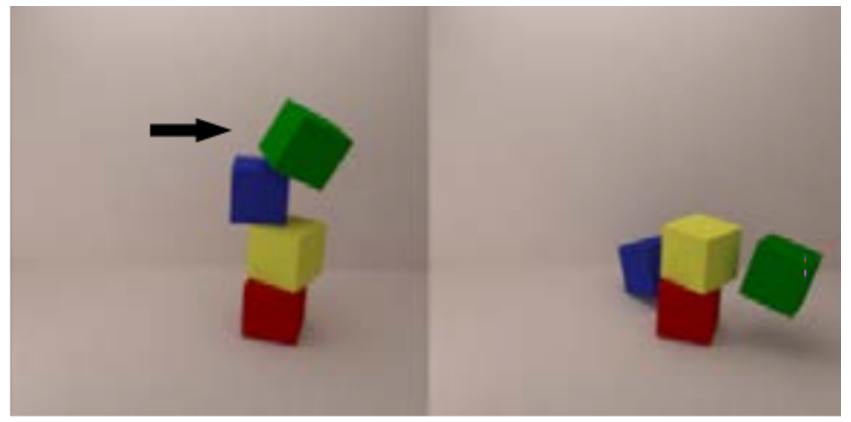

full supervision of object properties

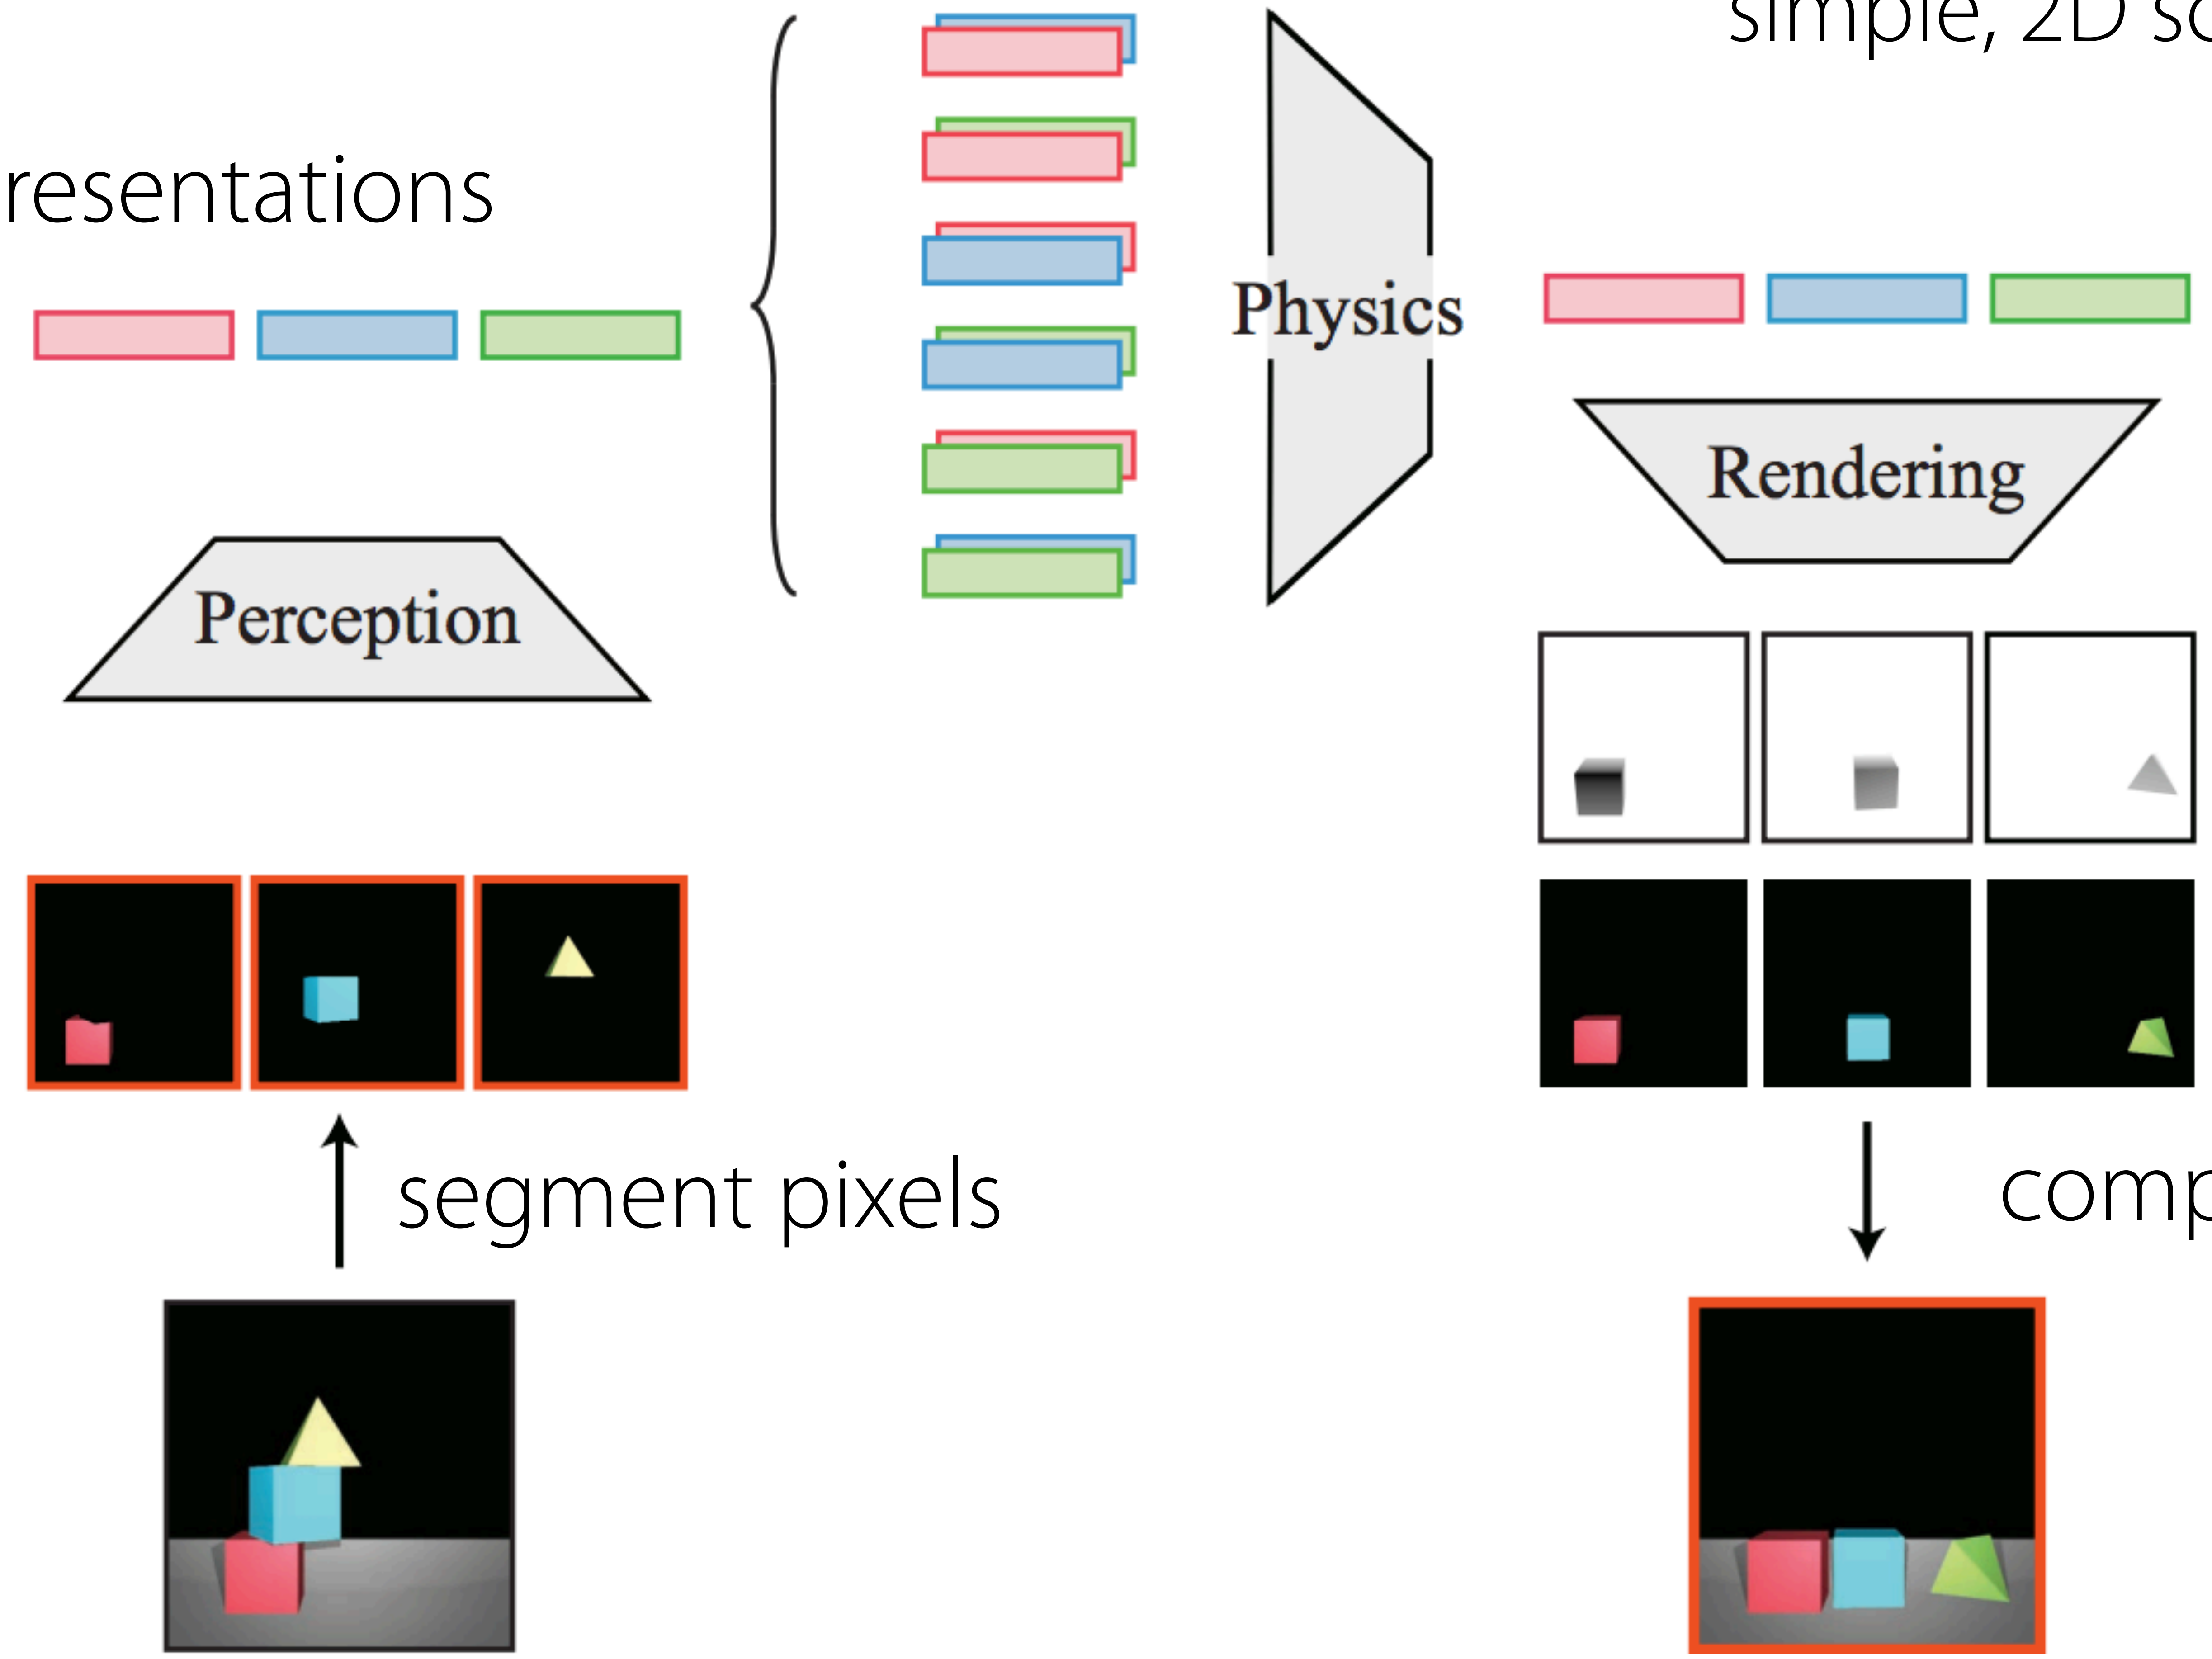

All modules trained with **reconstruction loss** (L_2+L_{VGG})

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Plan using model

goal space: image of object configuration

action space: which object & where to drop

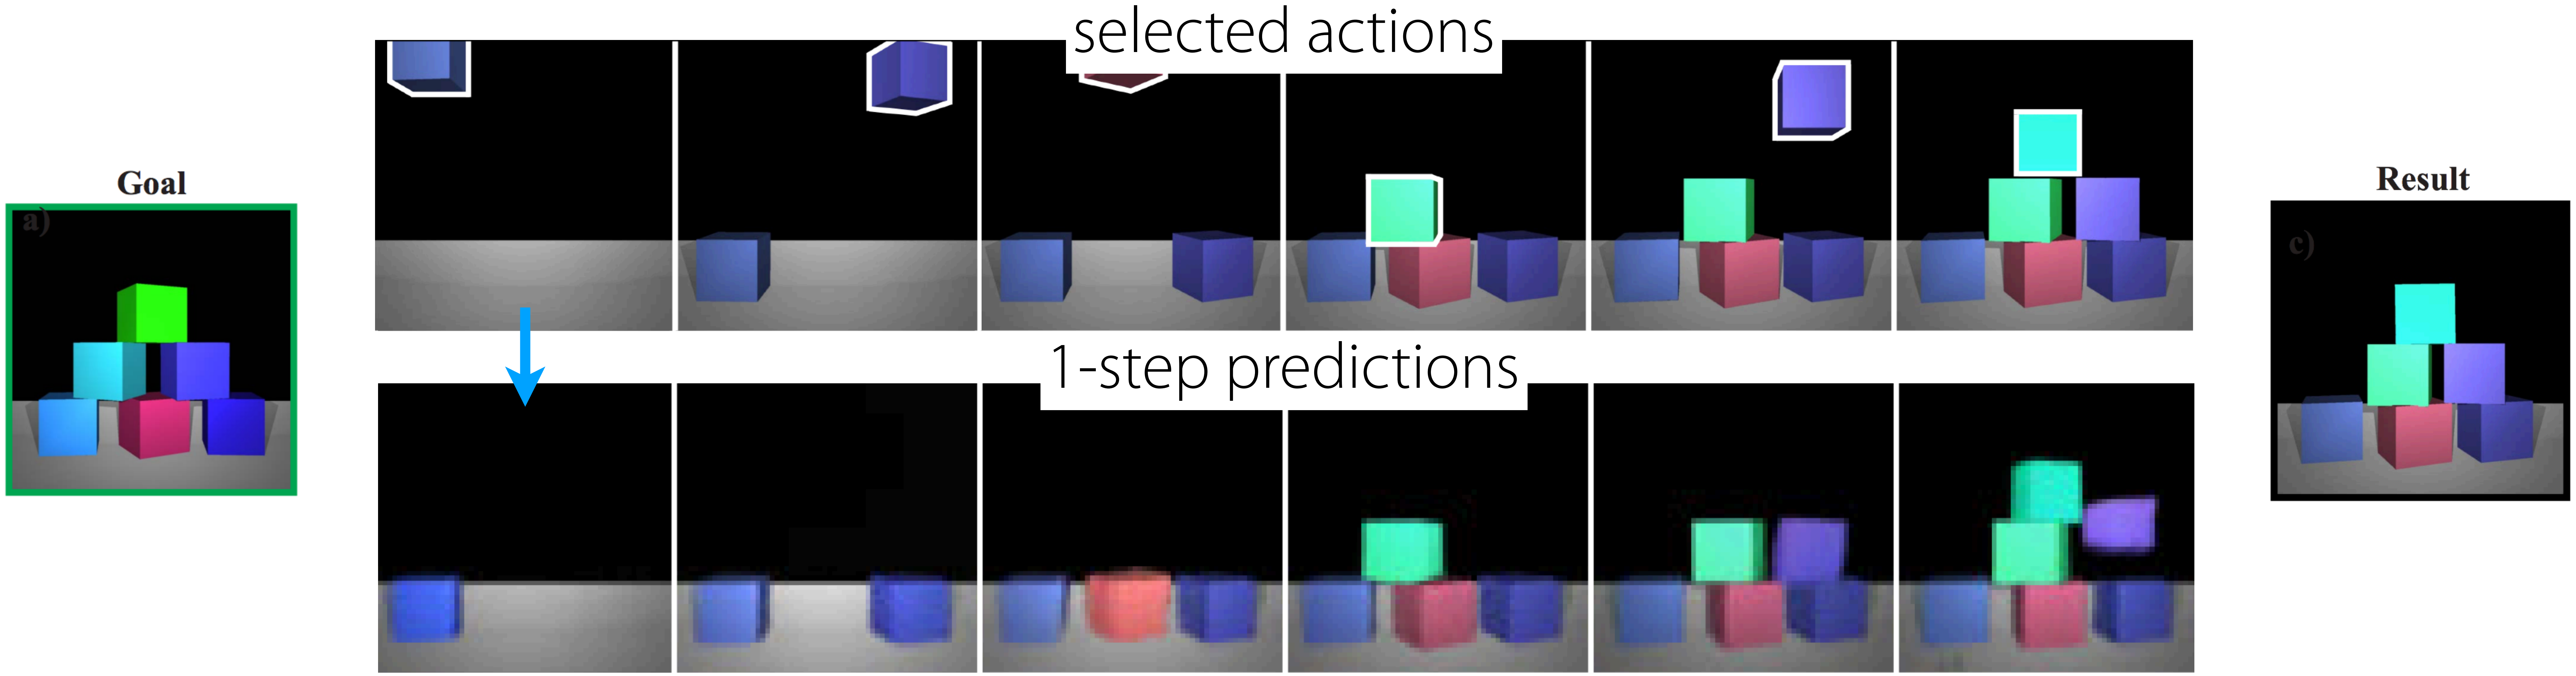

- **sampling-based**, **beam search** to plan action sequence
- evaluate action sequence based on **distance** in **latent space & pixel space**
- **replan** after each action

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Real world performance with single mode

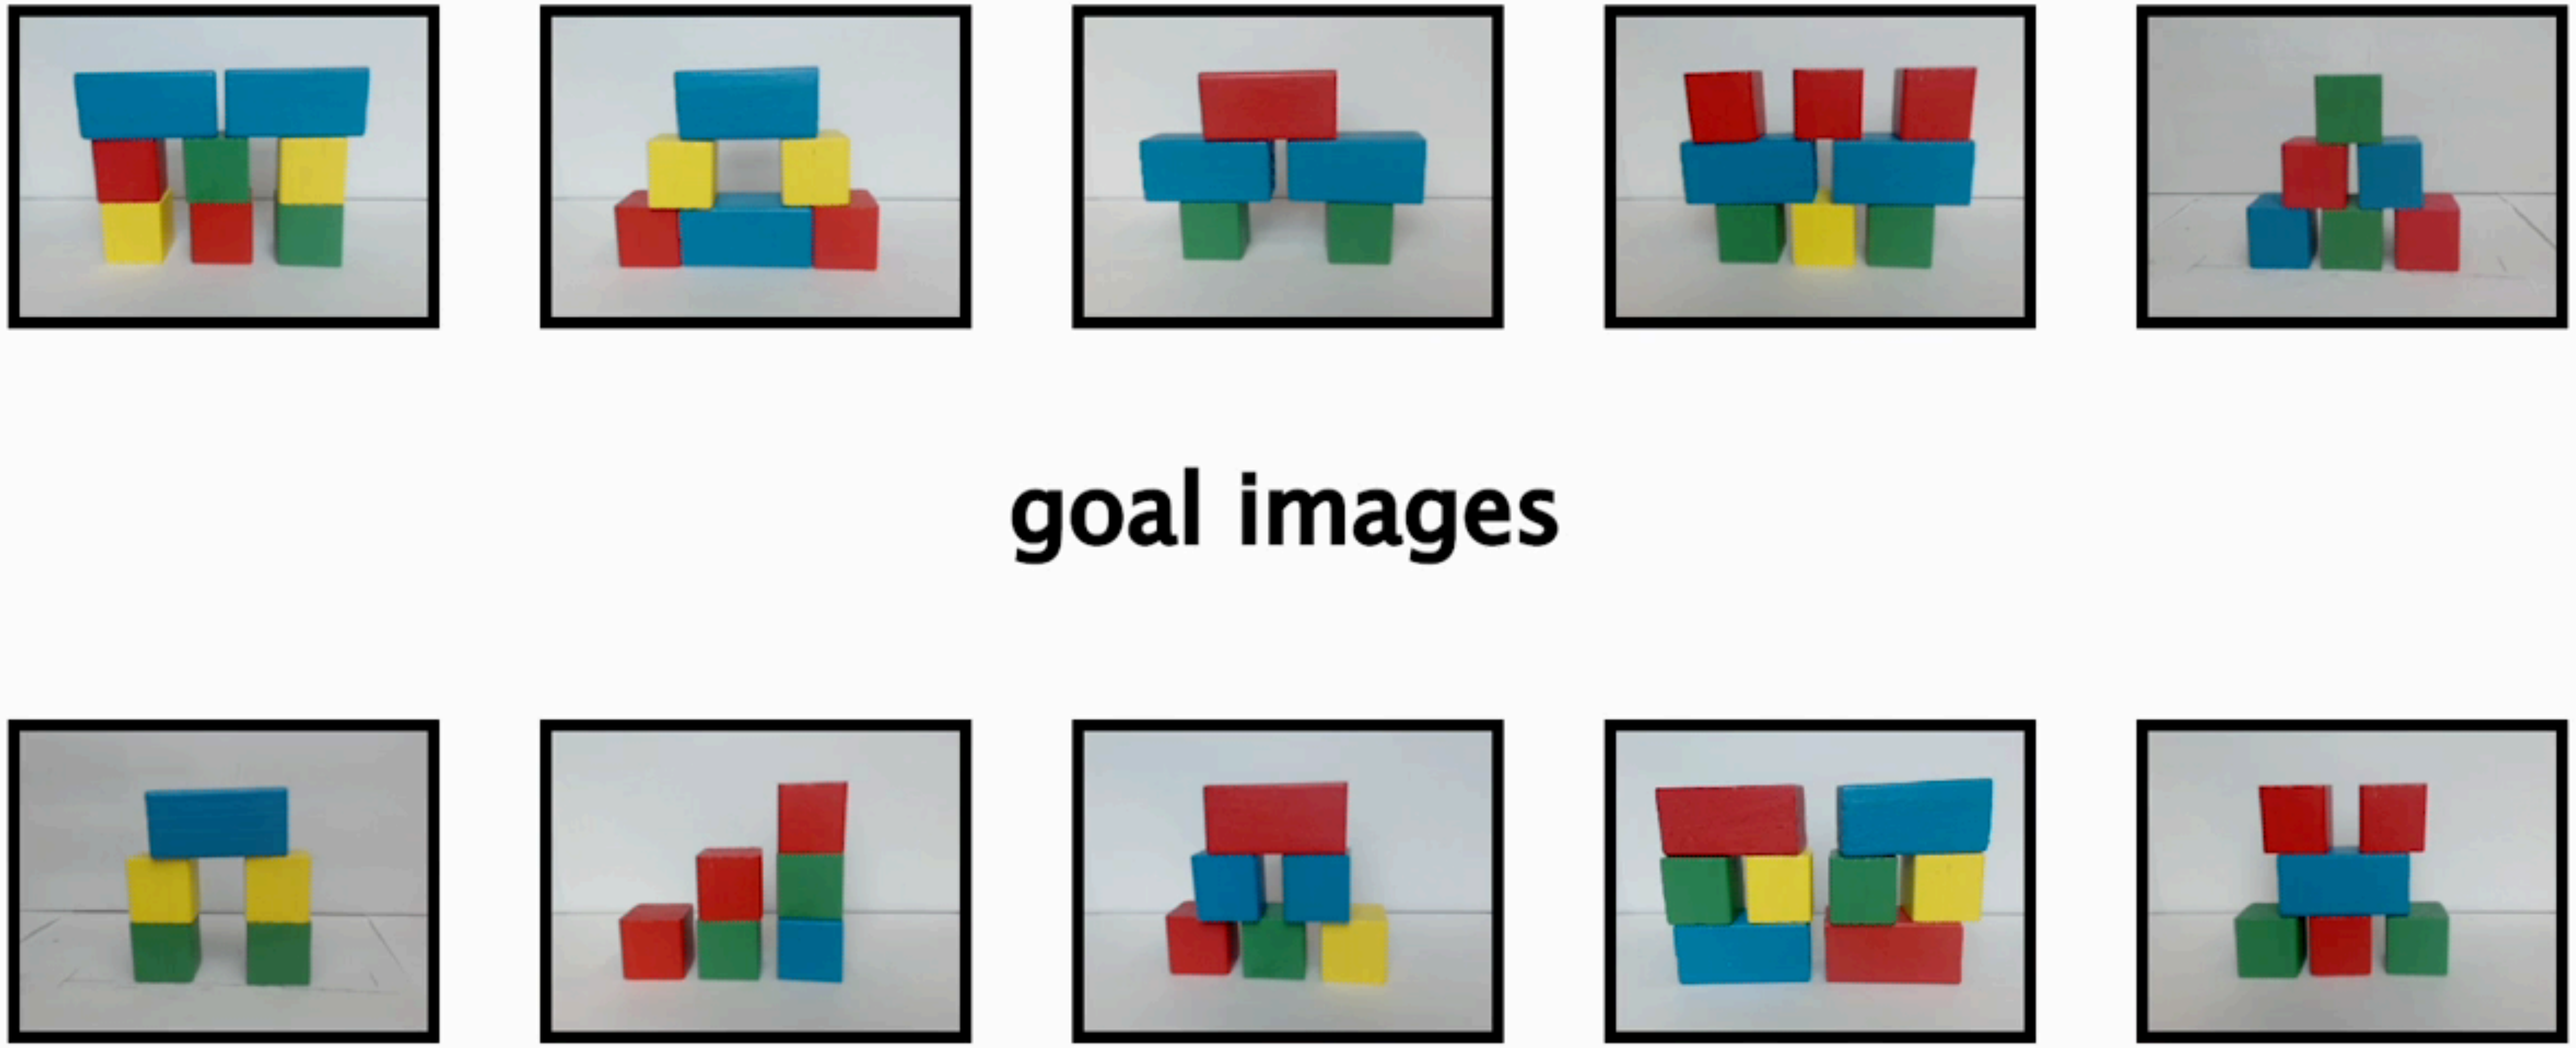

Takeaways

Learning model on diverse interactions → achieve **many** tasks

Structured latent space → achieve **complex**, **long-horizon** tasks

Janner, Levine, Freeman, Tenenbaum, Finn, Wu '18

Can we build an agent that can do **many tasks**?

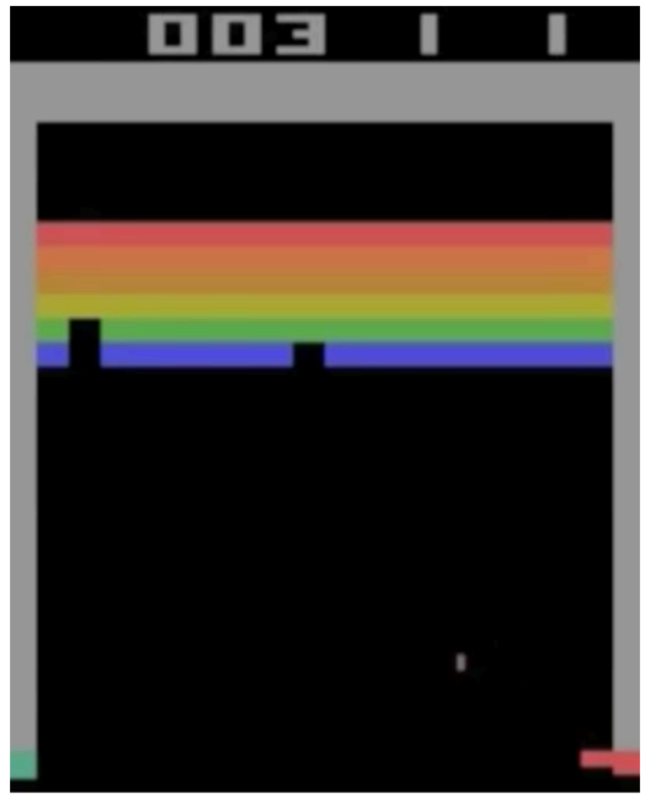

learning a **policy** in a **closed universe**

learn **general-purpose** model
+
plan with model **for many tasks**

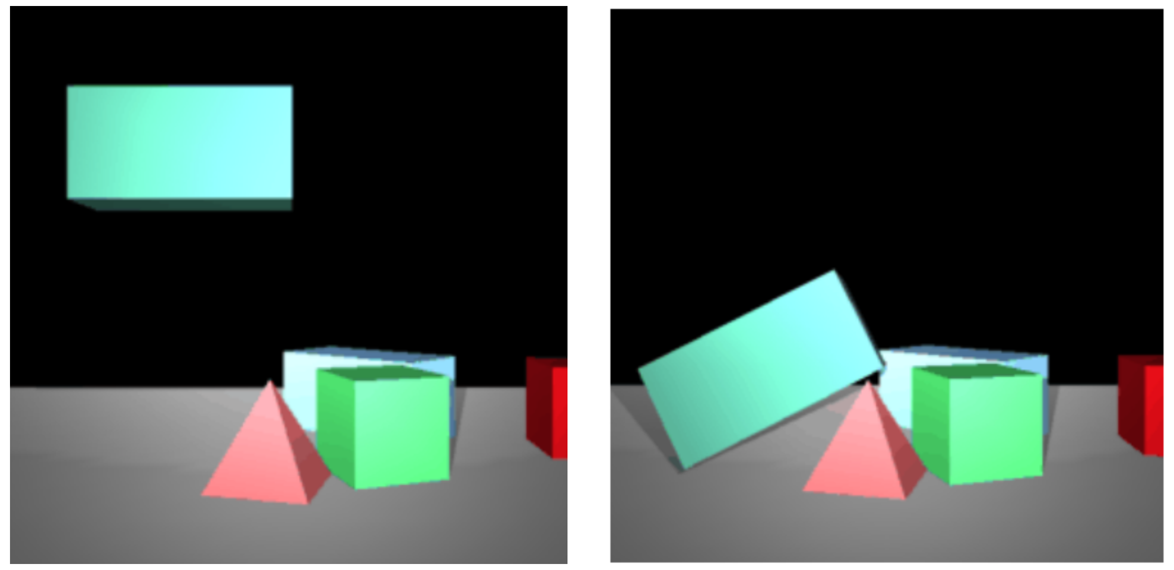

structured latent space model for long-horizon tasks

modeling **diverse**, **open-world** environments

from **pixel observations**, with **limited supervision**, in the **physical world**

Diverse Open-World Environments

self-supervised robot learning

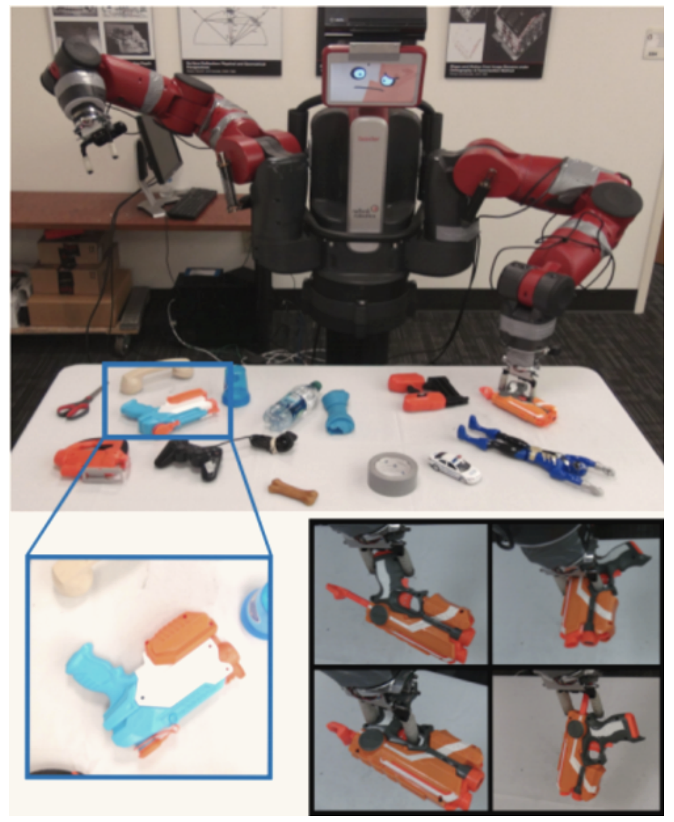

Pinto & Gupta '16

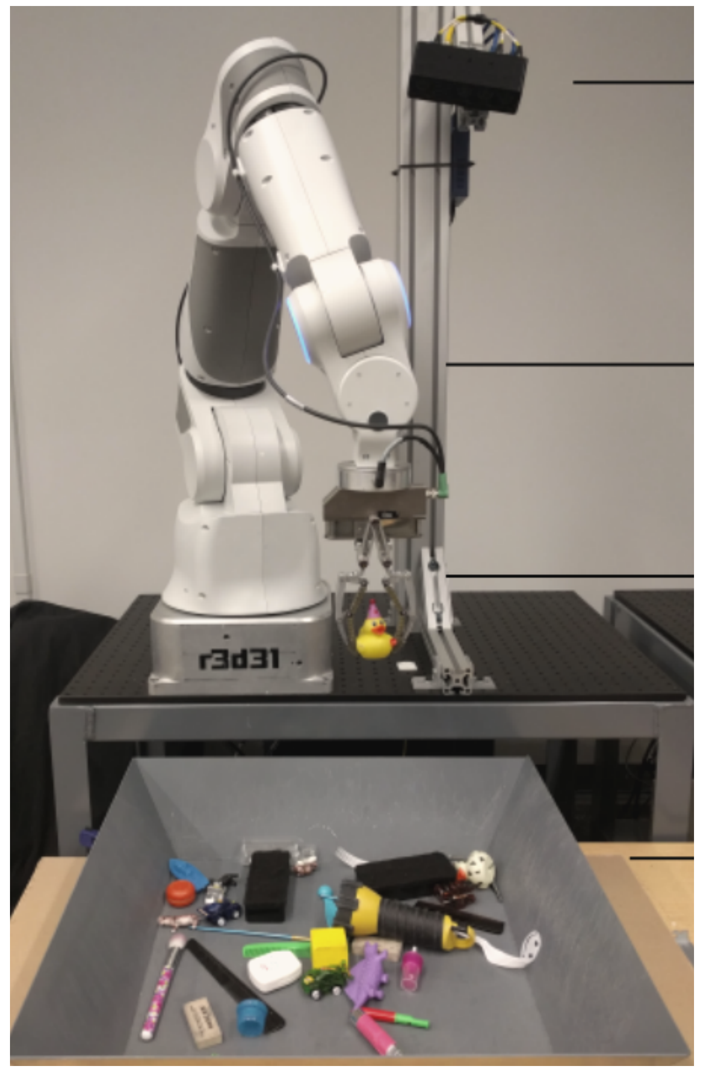

Levine, Pastor, Krizhevsky, Quillen '16

Nair*, Chen*, Agrawal*, Isola, Abbeel, Malik, Levine '17

Our goal: generalize to **novel objects** and, also to **many tasks**

(by learning a general-purpose model)

Overall approach: Collect data, learn model, plan to achieve many tasks

Collect **diverse** data in a **scalable** way

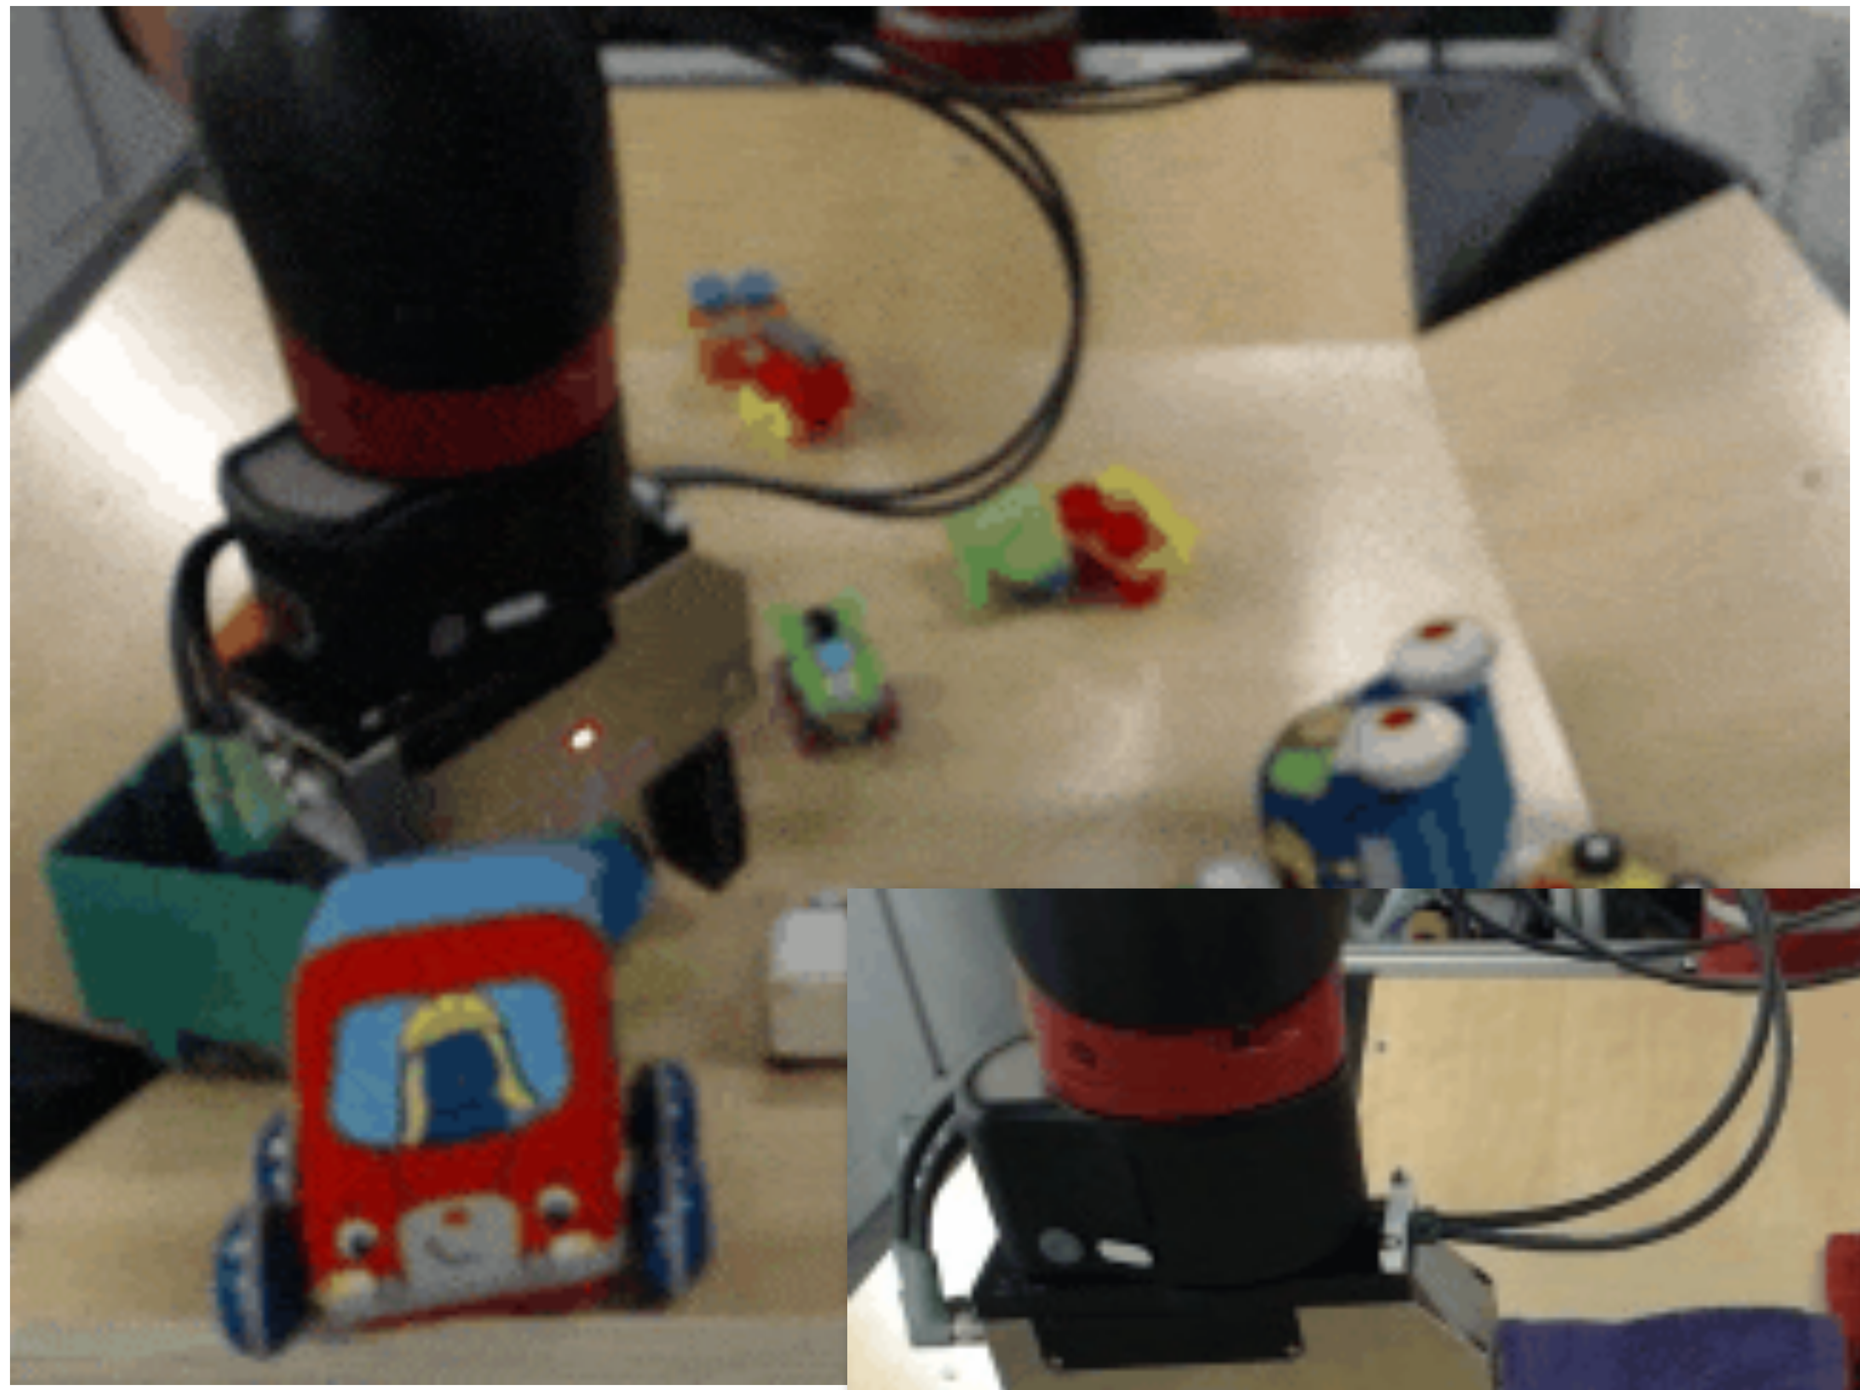

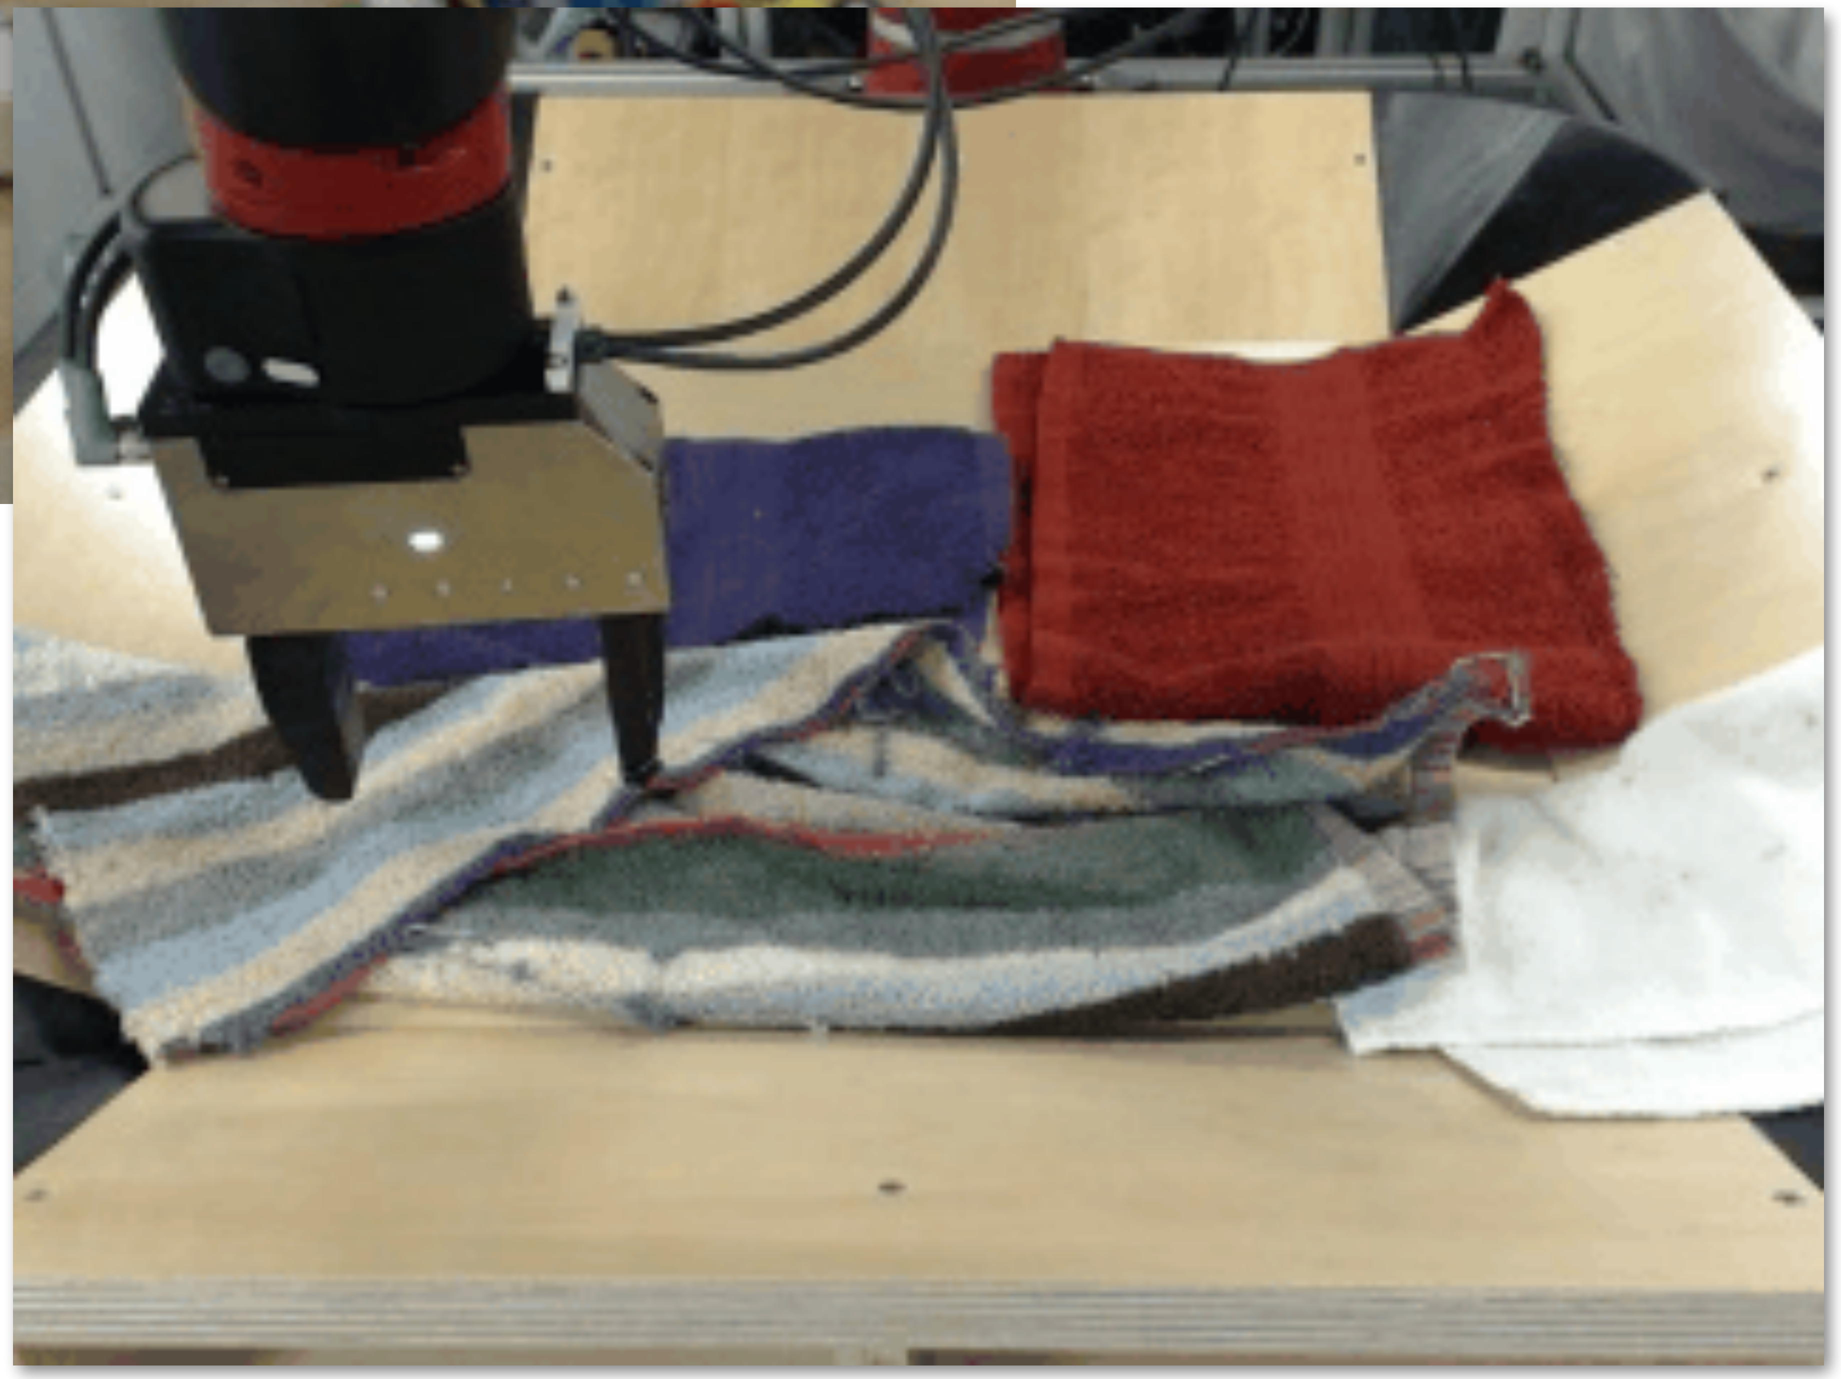

In contrast to policy learning: no notions of progress or success!

Learn to predict

$$I_t, a_{t:t+H} \longrightarrow I_{t:t+H}$$

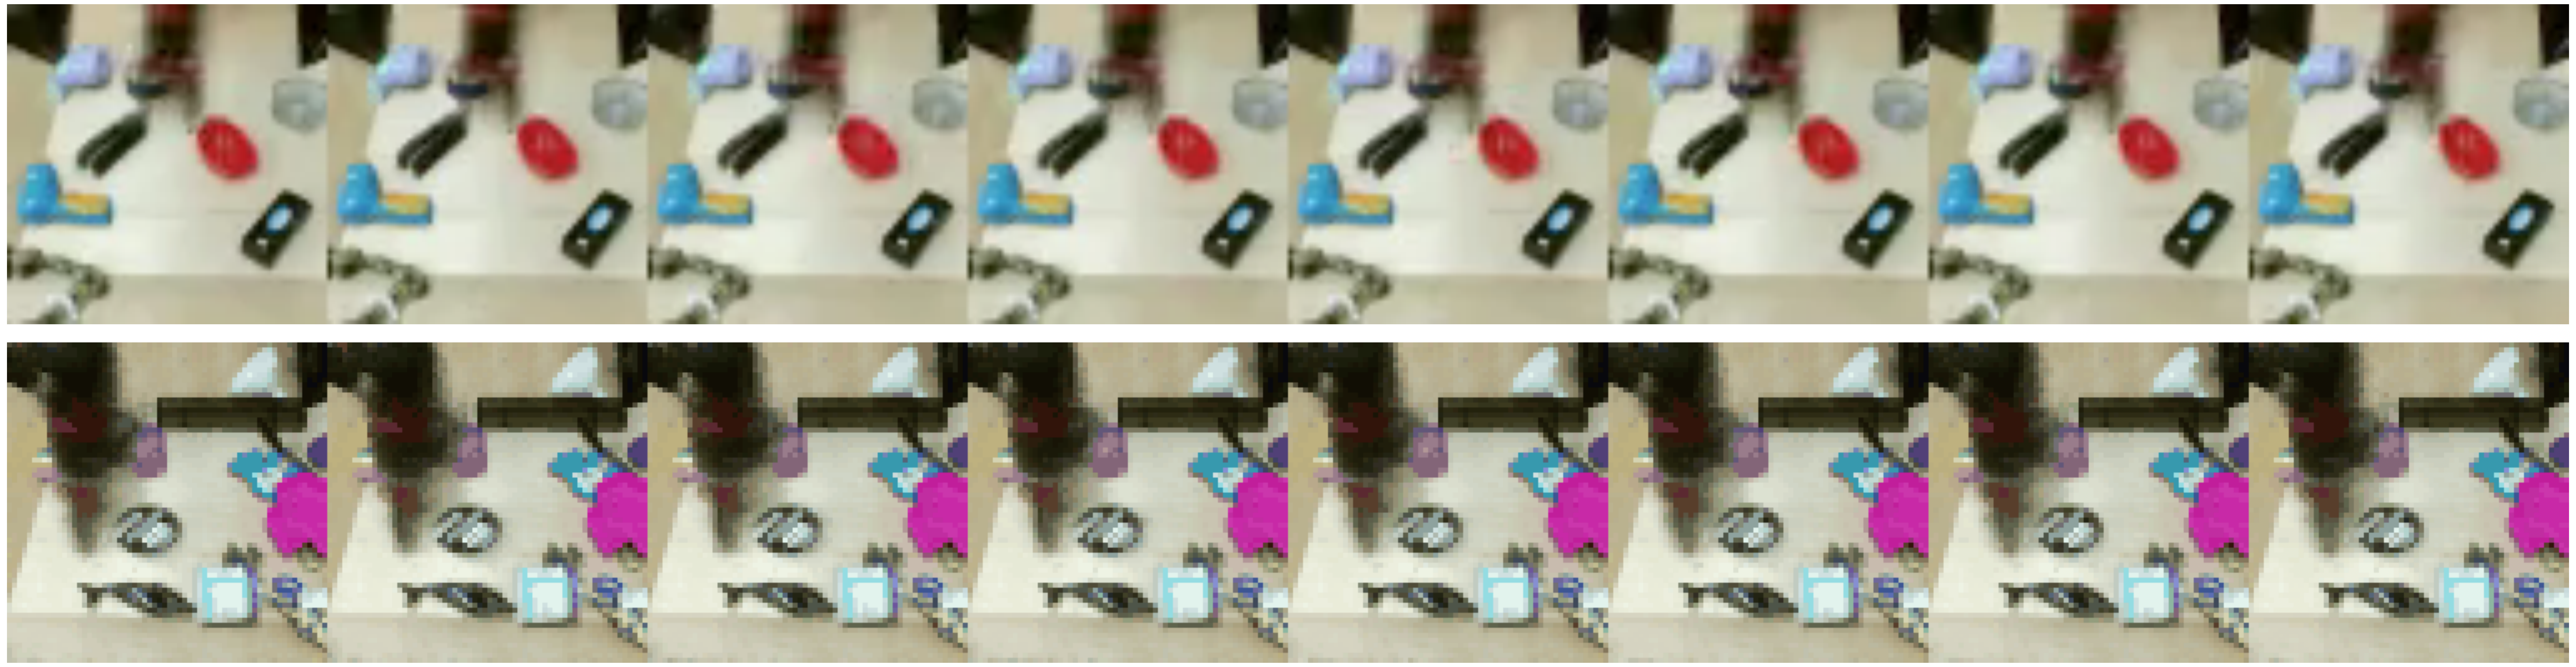

Contrast to:

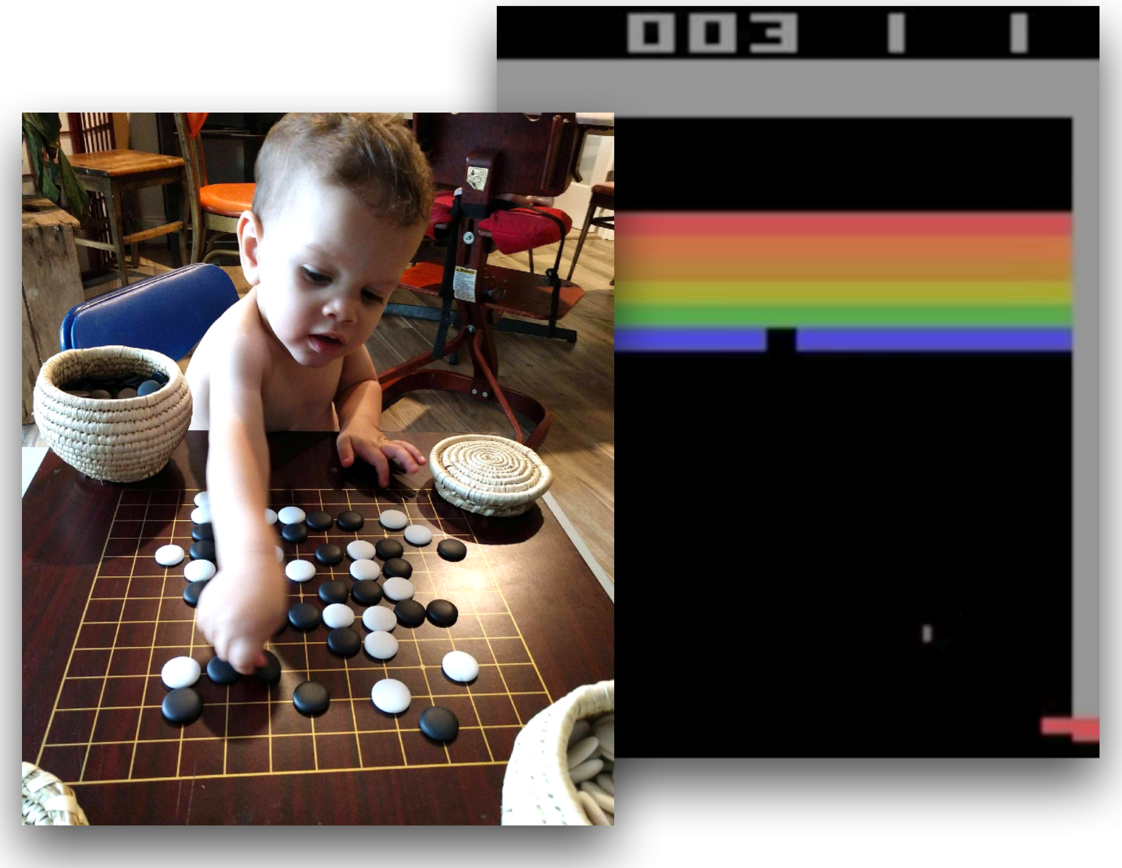

Models capture **general-purpose** knowledge about the world

Use **all** of the available supervision signal.

Also: No assumptions about task **representations**.

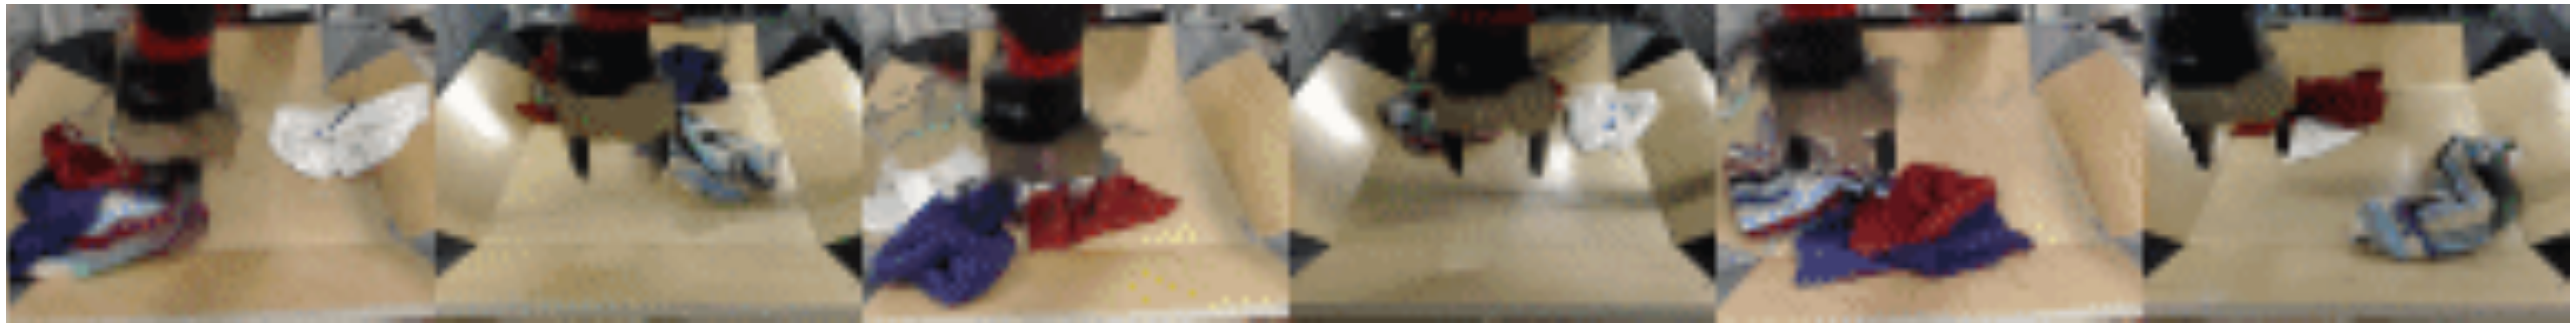

Are these models useful?

How can we use these models to plan?

(to achieve many human-specified goals)

Planning with Visual Foresight

1. Consider potential action sequences
2. Predict the future for each action sequence
3. Pick best future & execute corresponding action
4. Repeat 1-3 to replan in real time

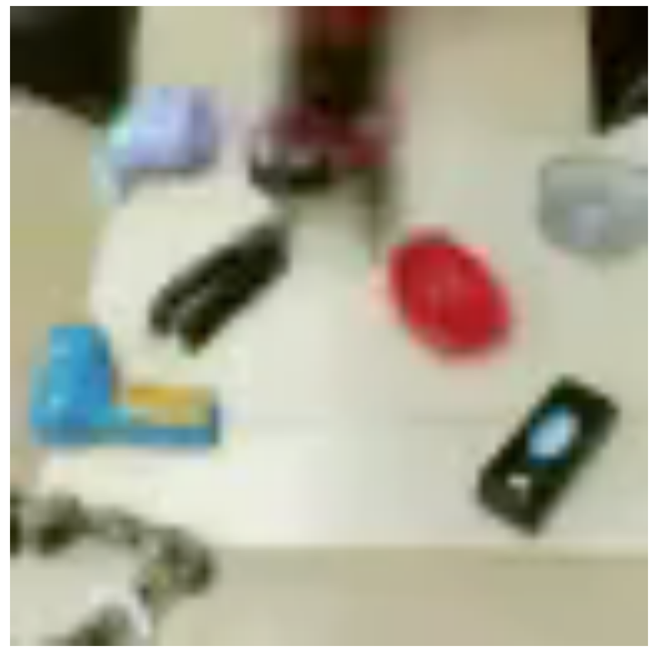

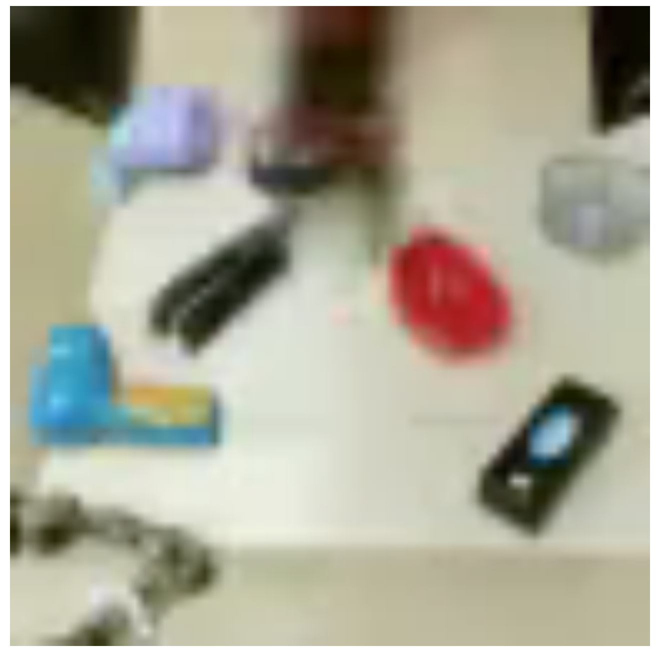

visual "model-predictive control" (MPC)

How to predict video?

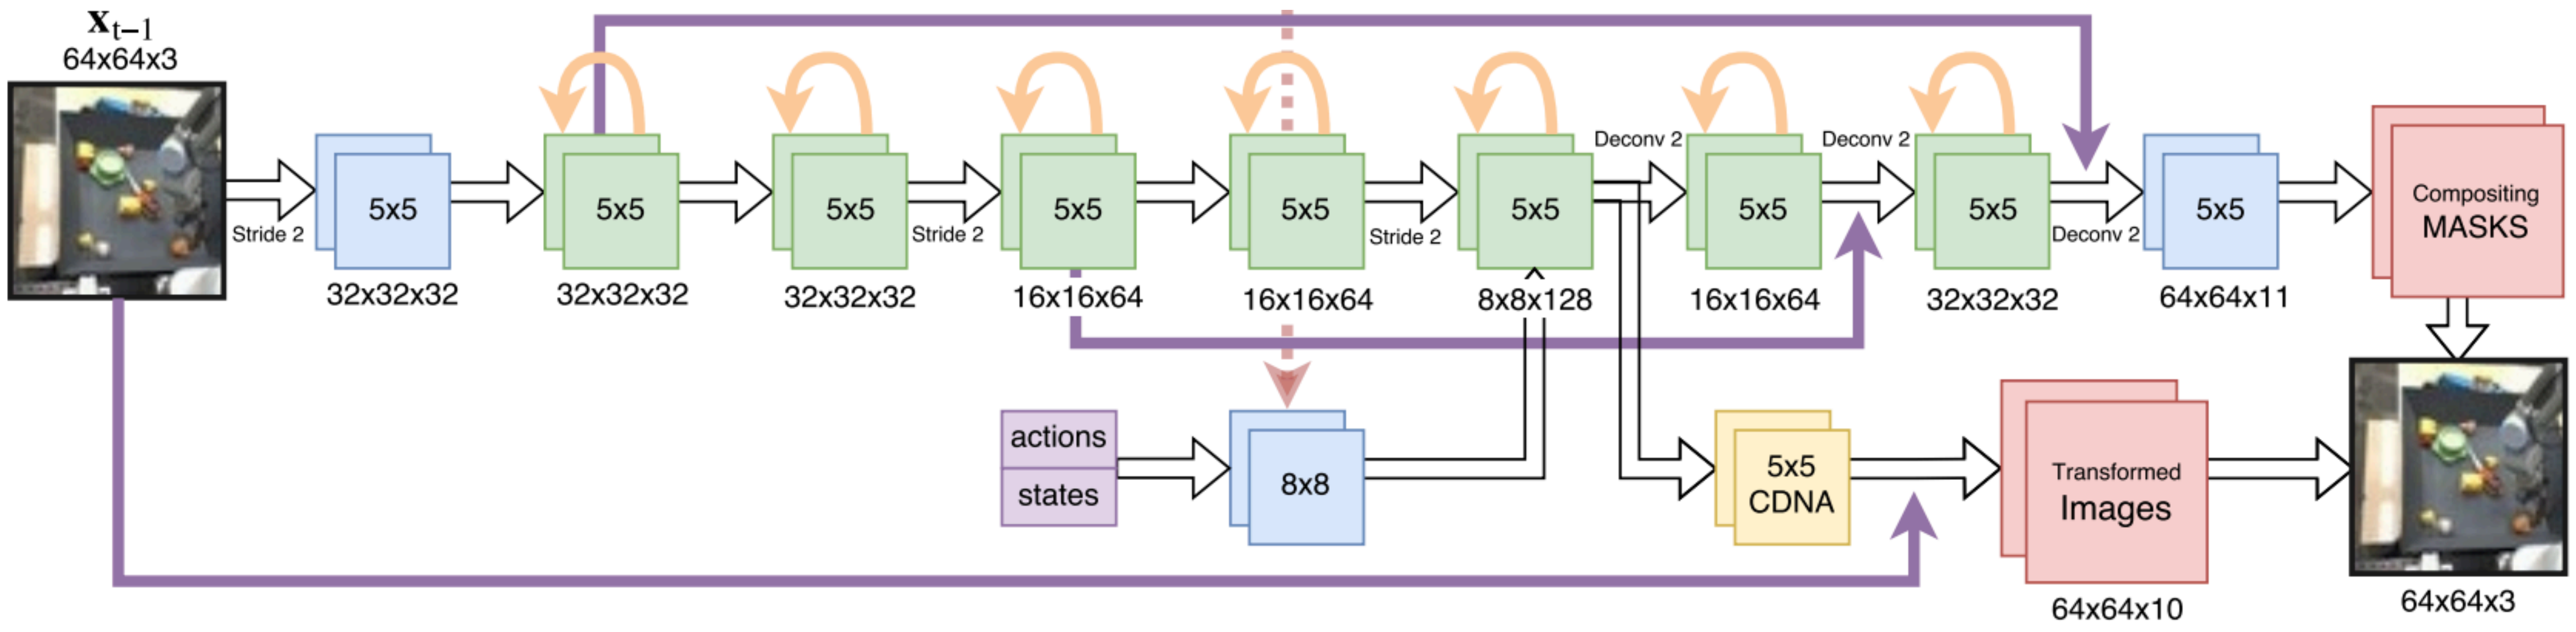

- deep recurrent network
- multi-frame prediction
- action-conditioned
- explicitly model motion

Finn, Goodfellow, Levine NIPS '16
Finn & Levine ICRA '17

Which future is the best one?

Human specifies a goal by:

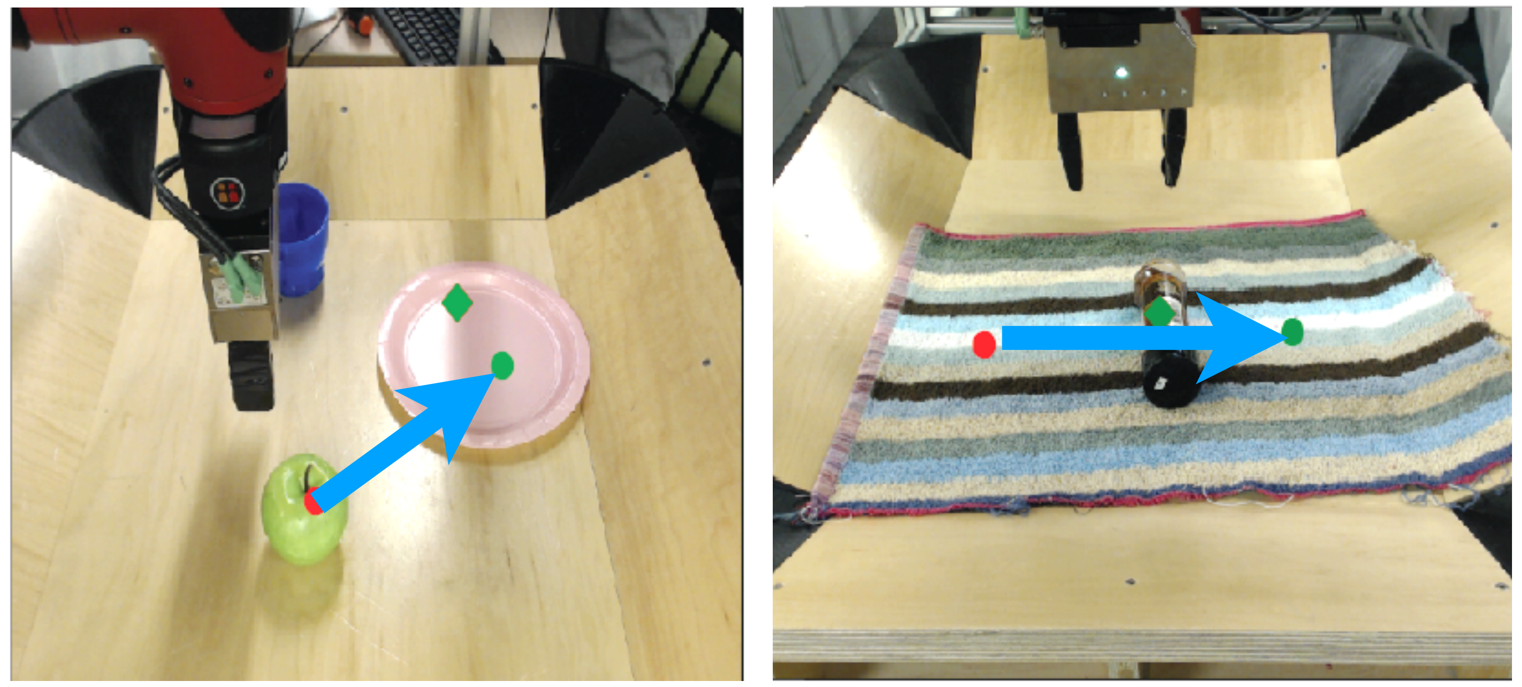

Selecting where
pixels should move.

Providing an image
of the goal.

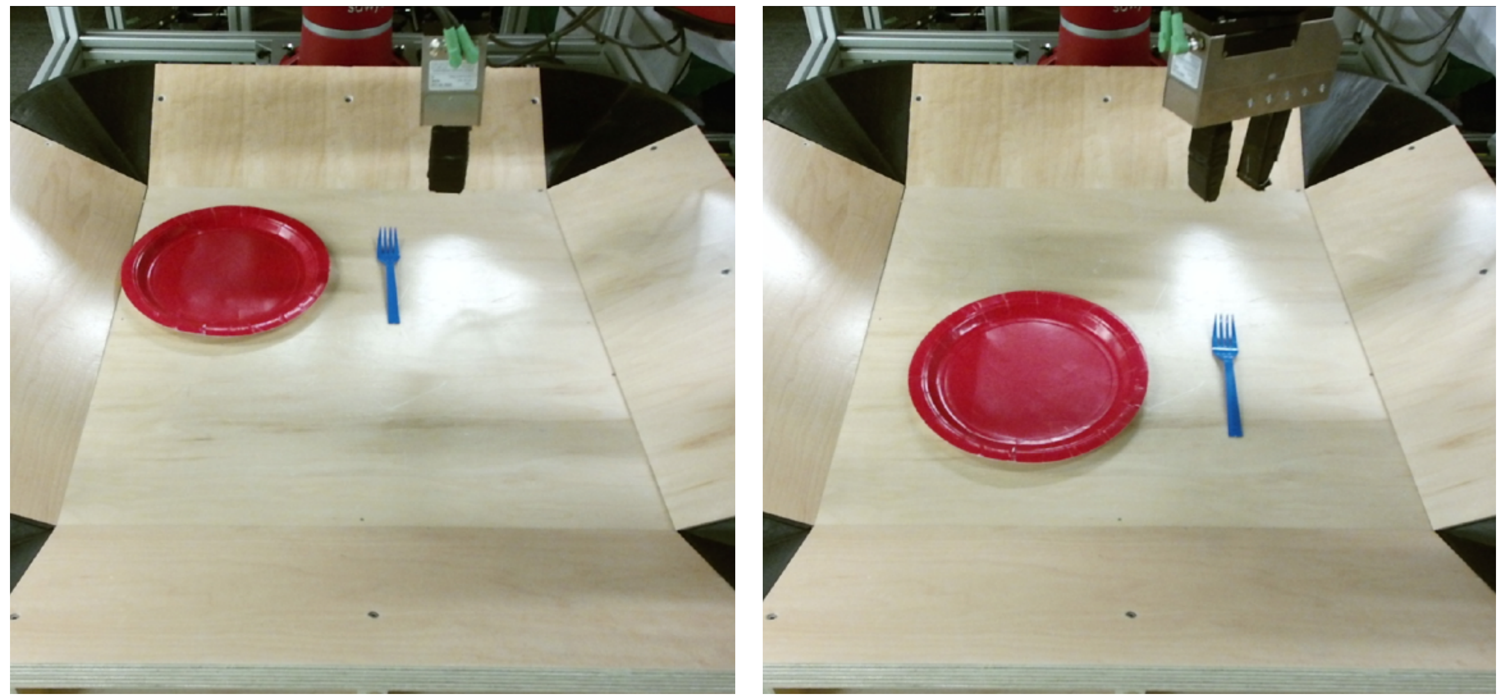

Providing a few
examples of success.

Finn & Levine ICRA '17

Ebert, Lee, Levine, Finn CoRL '18

Xie, Singh, Levine, Finn CoRL '18

How it works

Specify goal

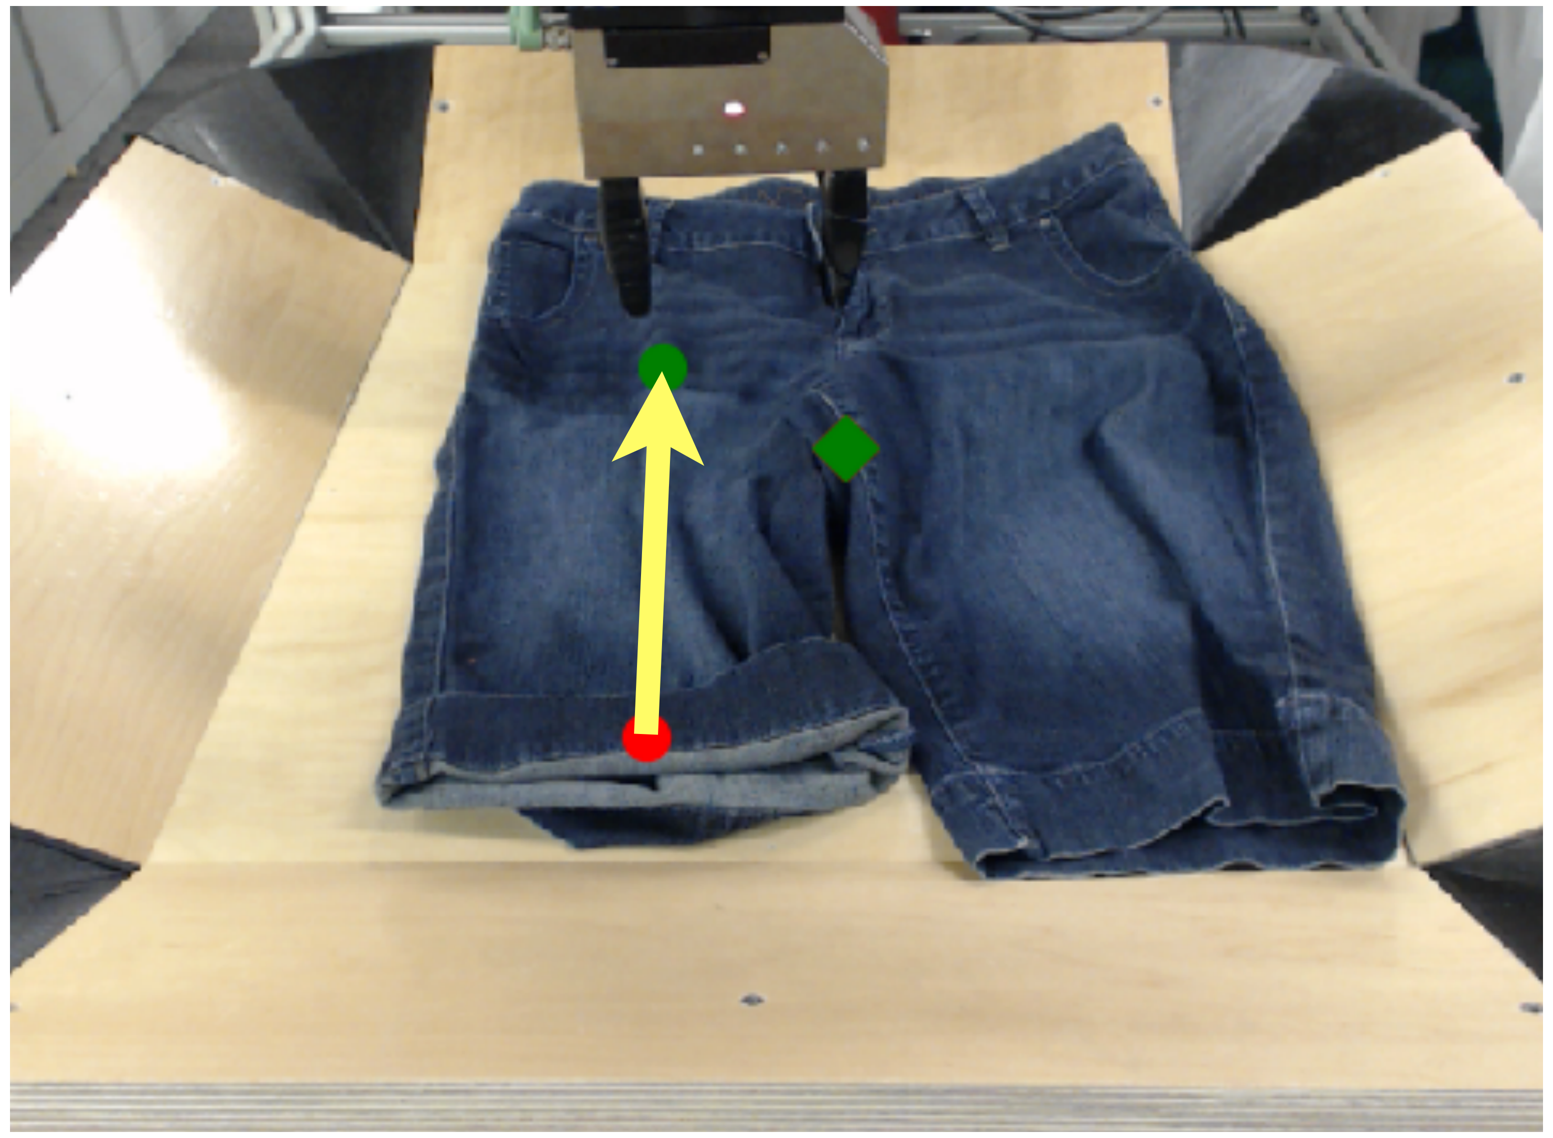

Visual MPC w.r.t. goal

Visual MPC execution

Frederik Ebert

Sudeep Dasari

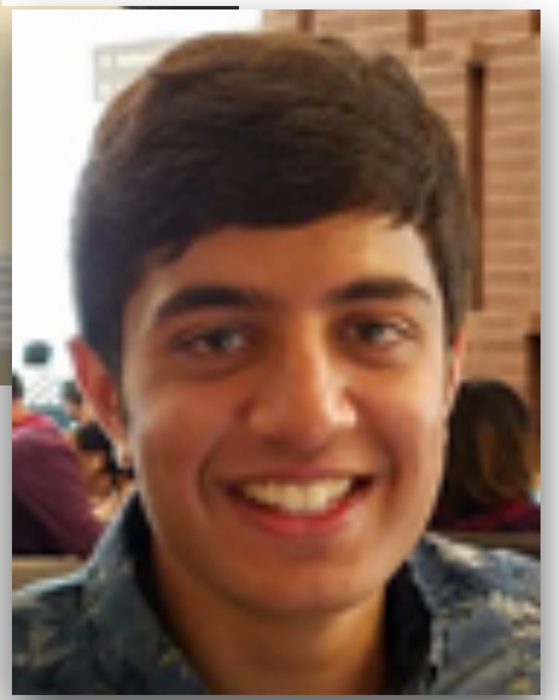

Ebert*, Finn*, Dasari, Xie, Lee, Levine. Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control.

How it works

Specify goal

(covering an object)

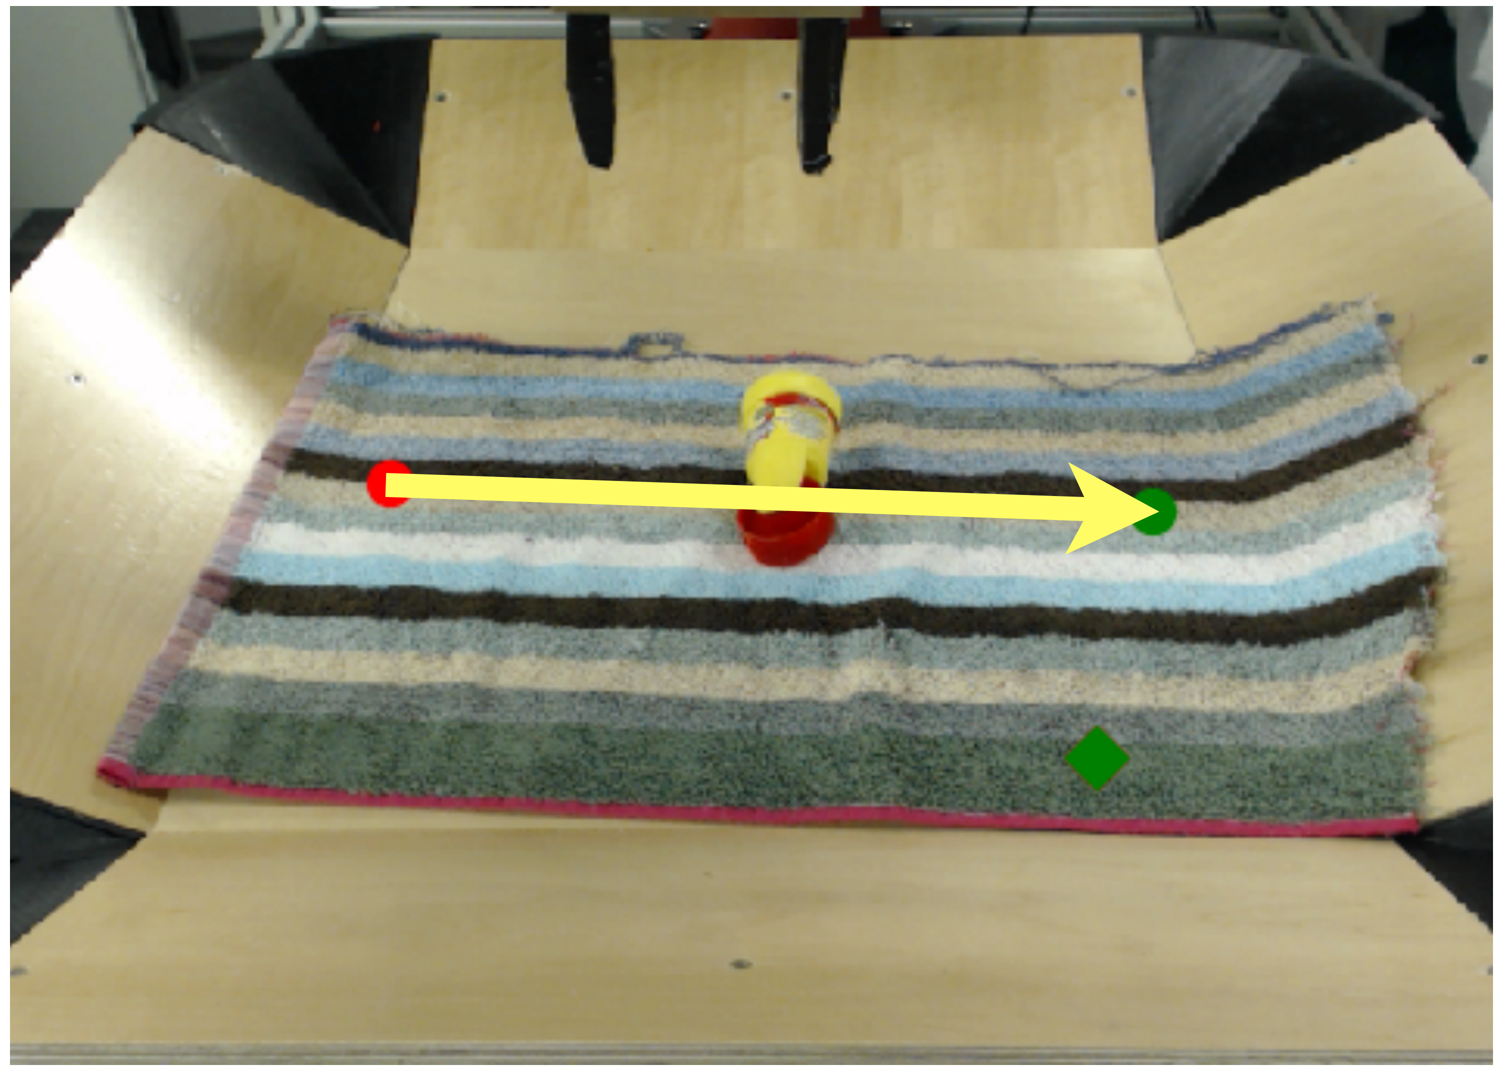

Visual MPC execution

Frederik Ebert

Sudeep Dasari

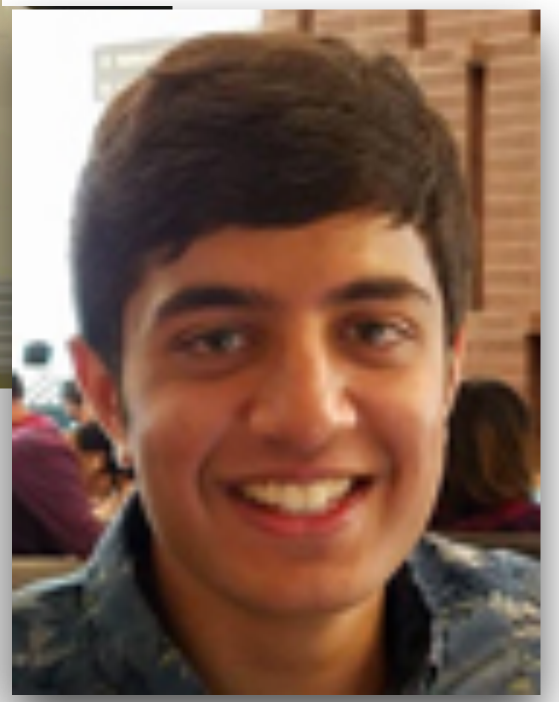

Ebert*, Finn*, Dasari, Xie, Lee, Levine. Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control.

How it works

Given 5 examples of success

infer goal classifier

visual MPC w.r.t. goal classifier

Visual MPC with learned objective

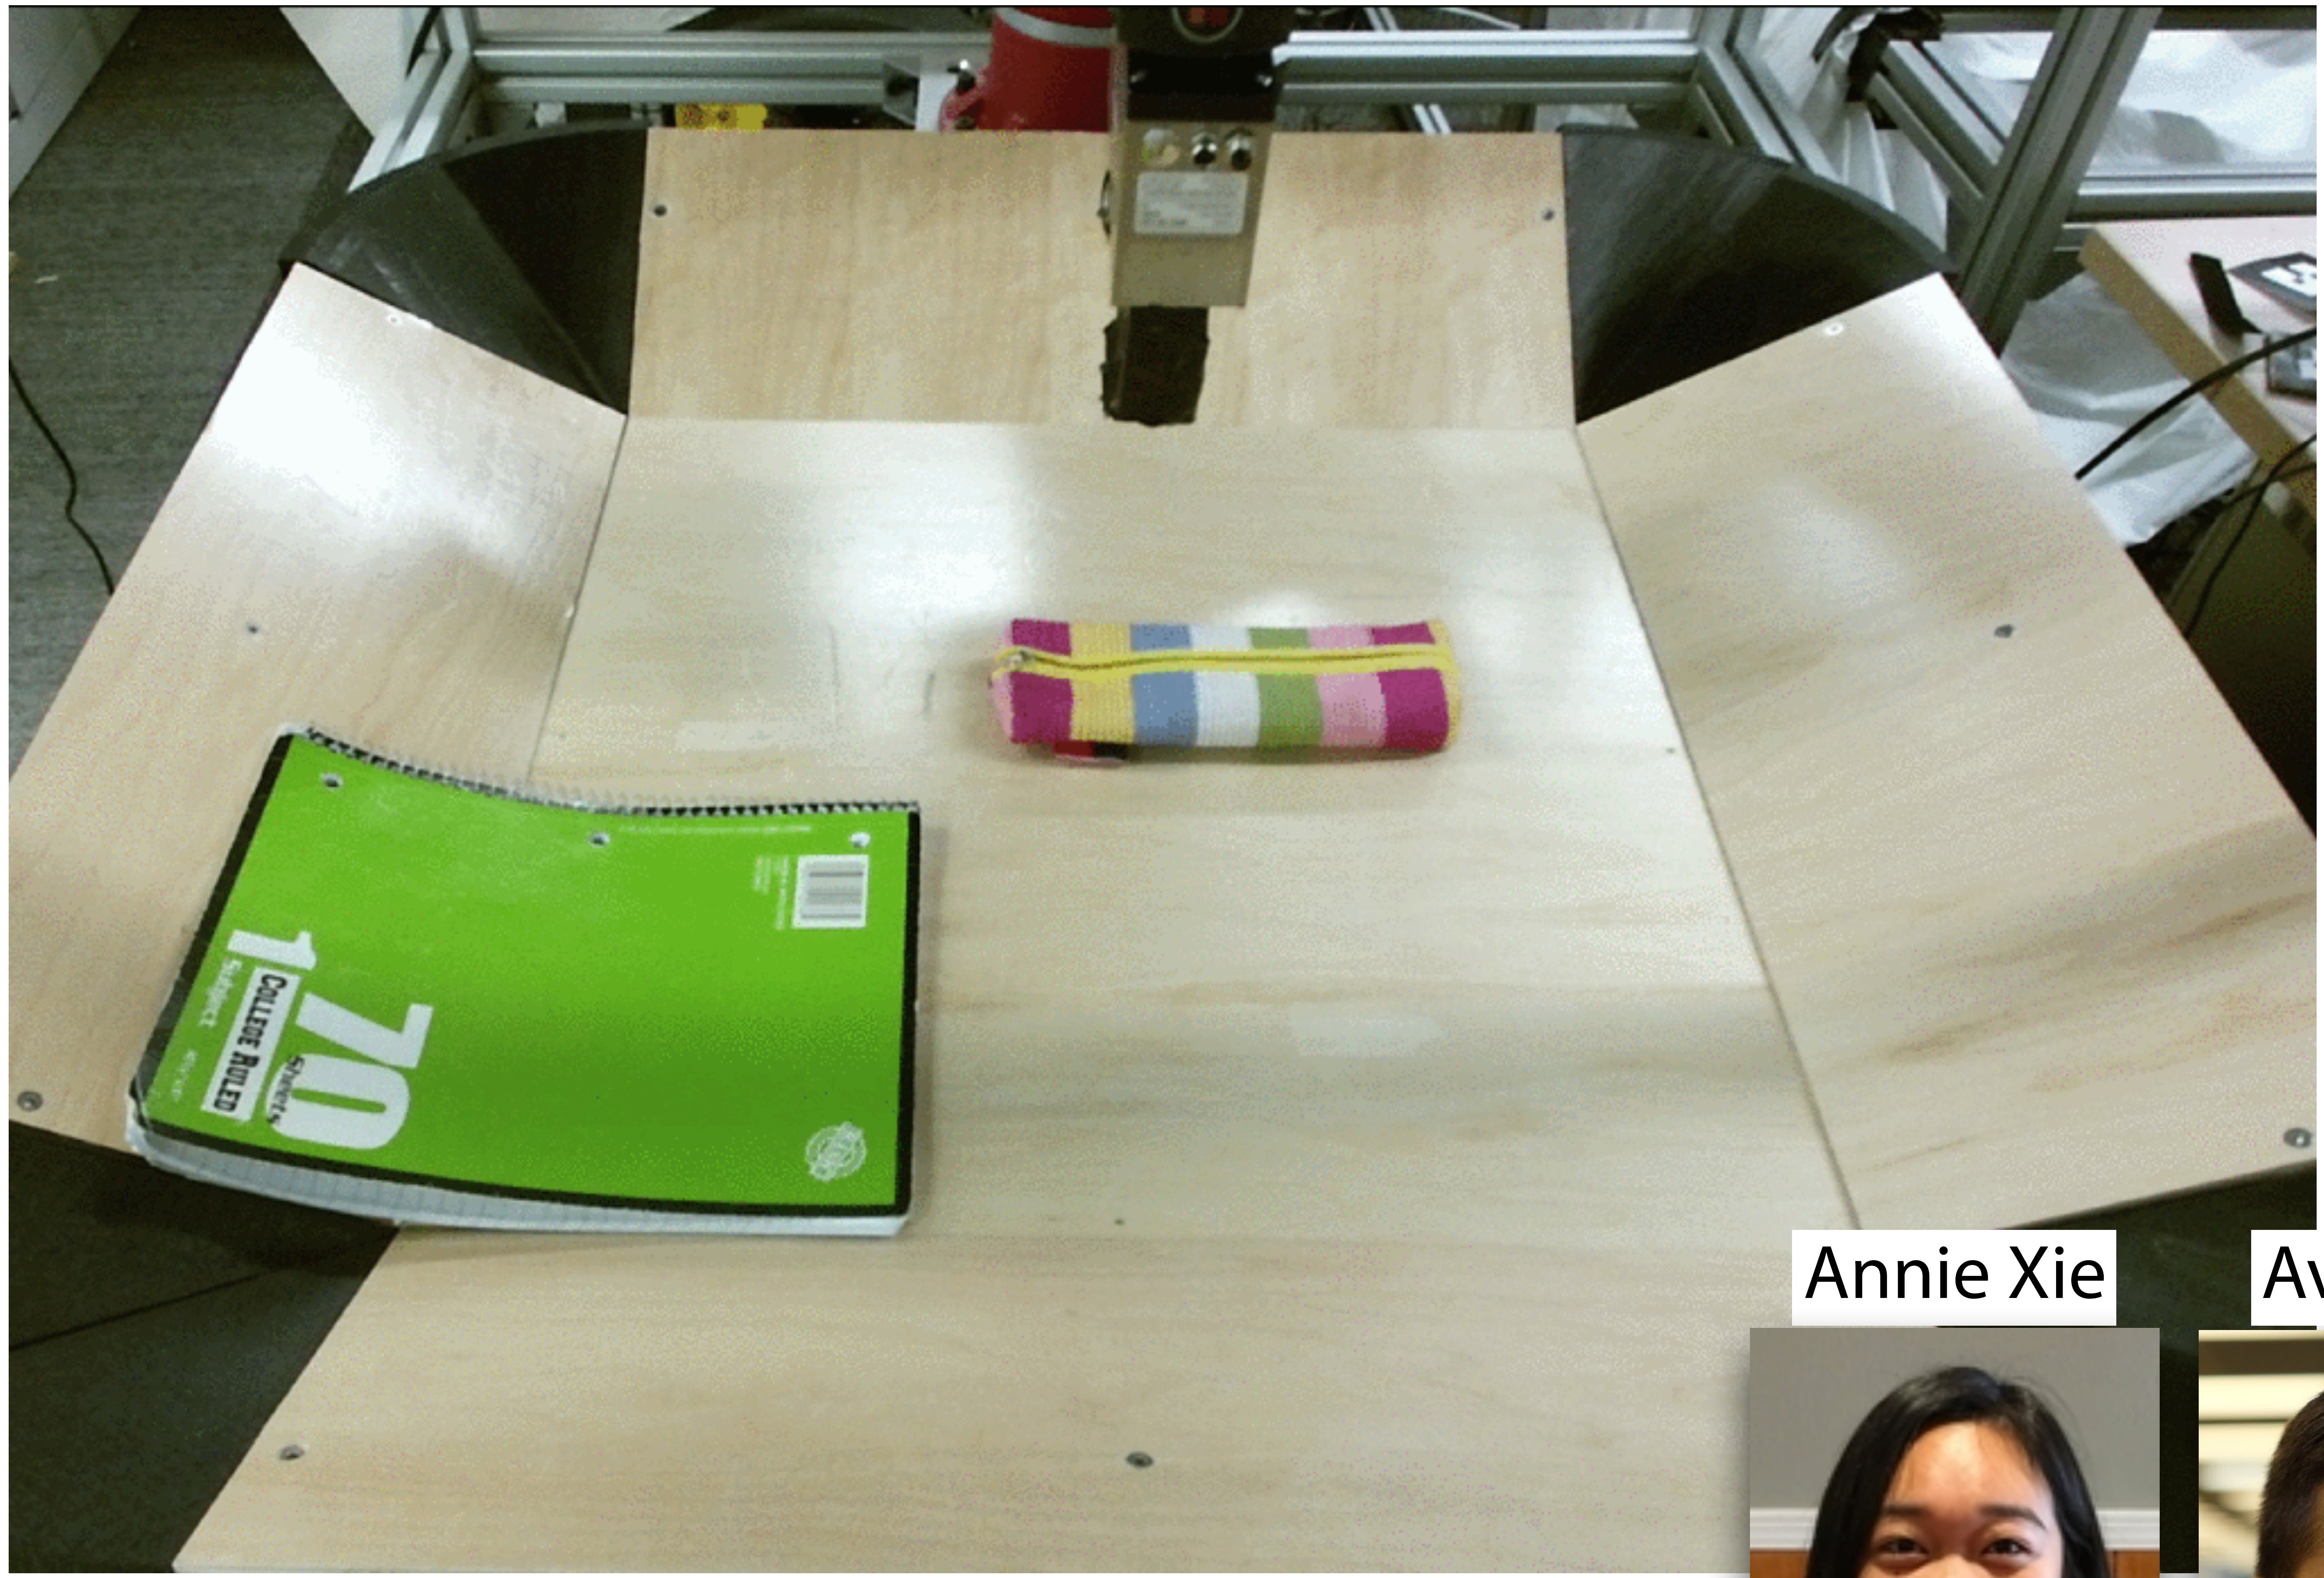

Annie Xie

Avi Singh

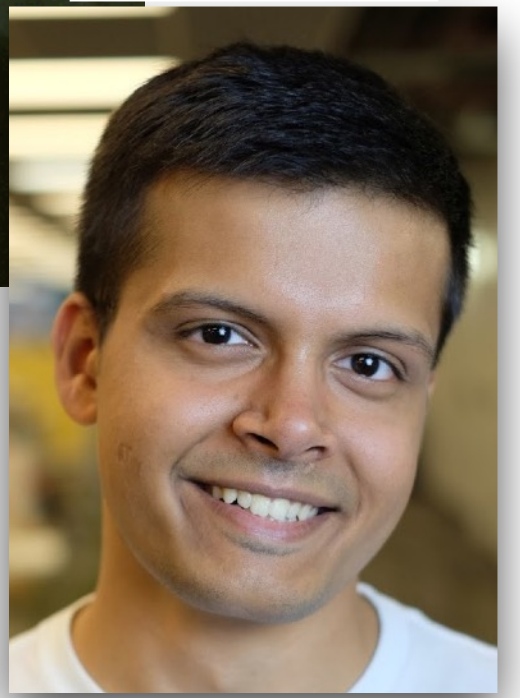

Xie, Singh, Levine, Finn. Few-Shot Goal Inference for Visuomotor Learning and Planning, CoRL 2018

Planning with a **single model** for many tasks

Video speed: 2x

Demo at NIPS 2017: Long Beach, CA

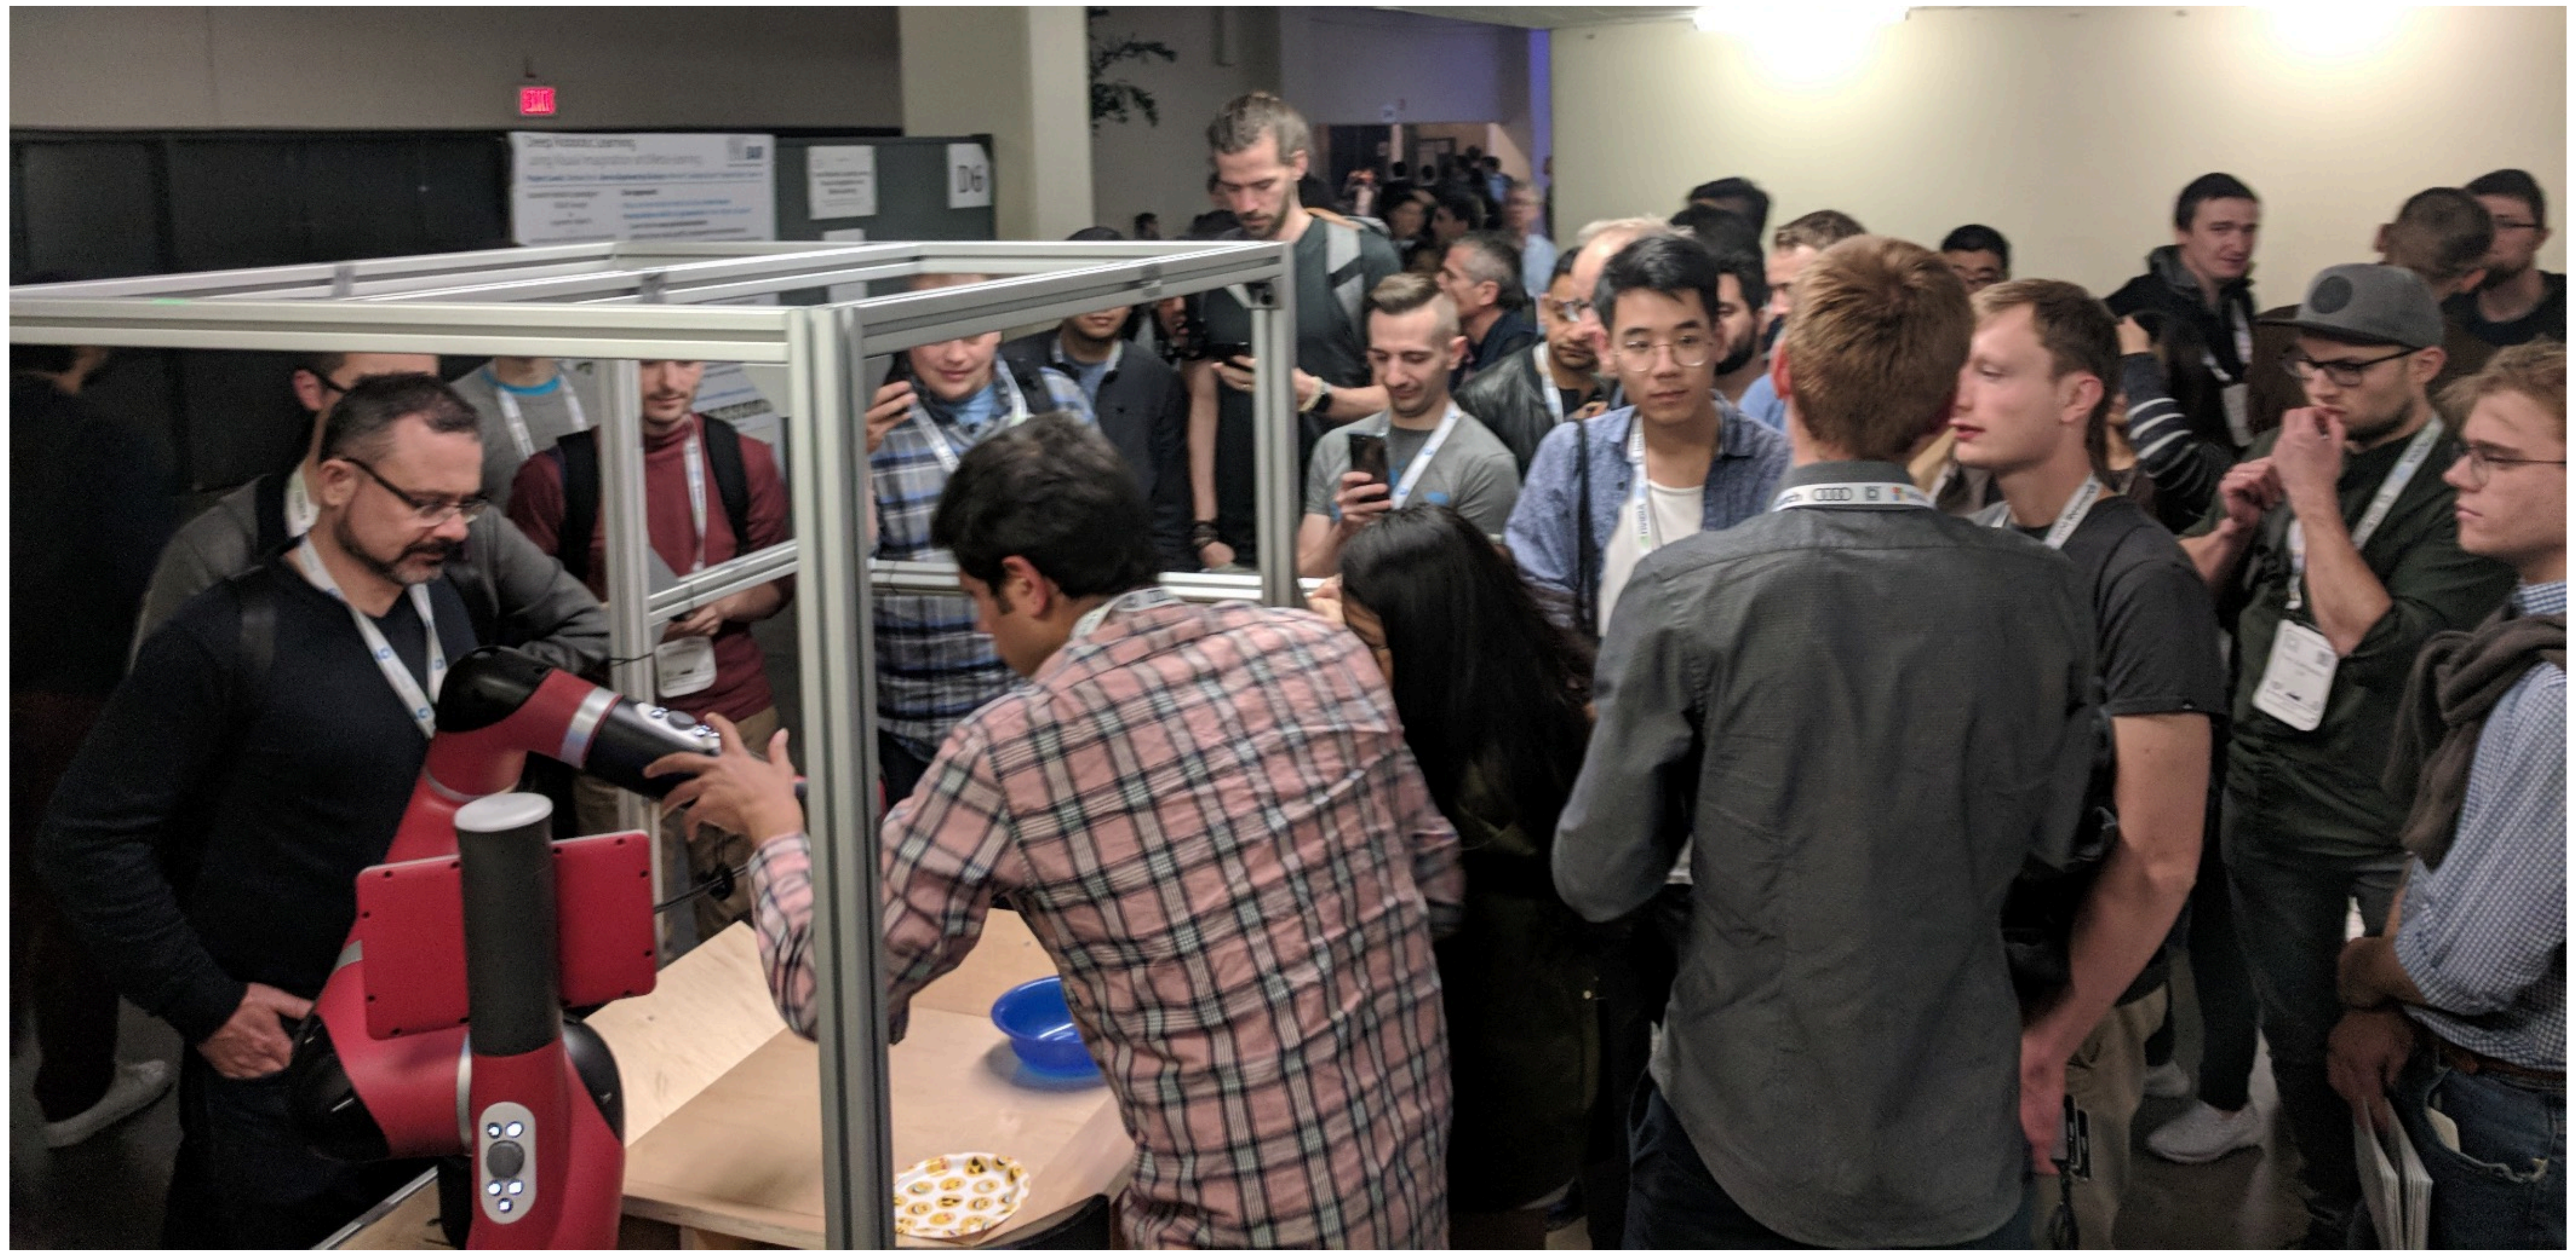

planning with visual models

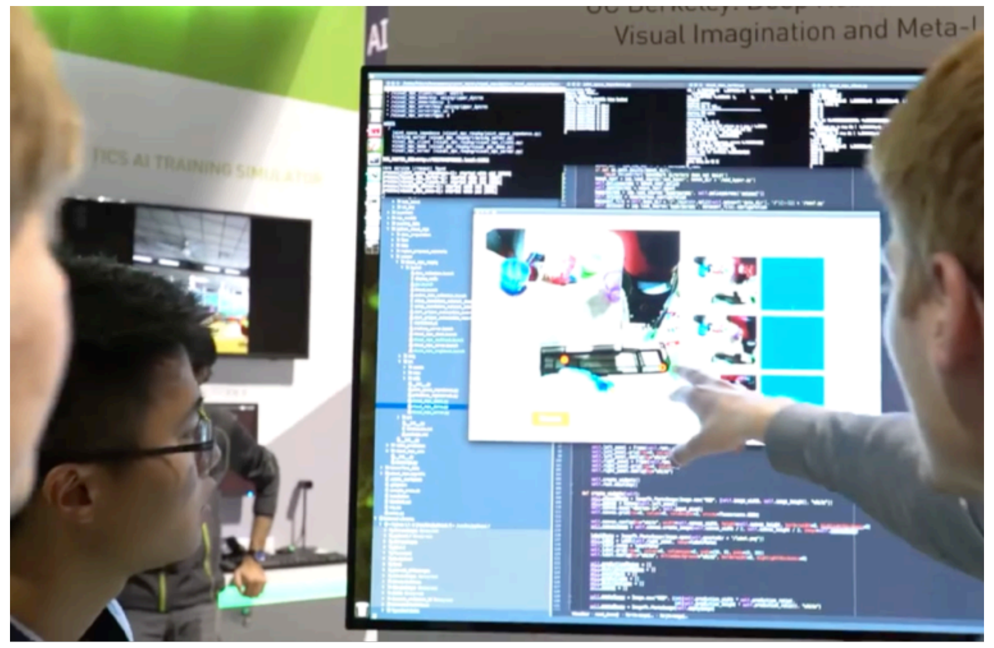

one-shot imitation

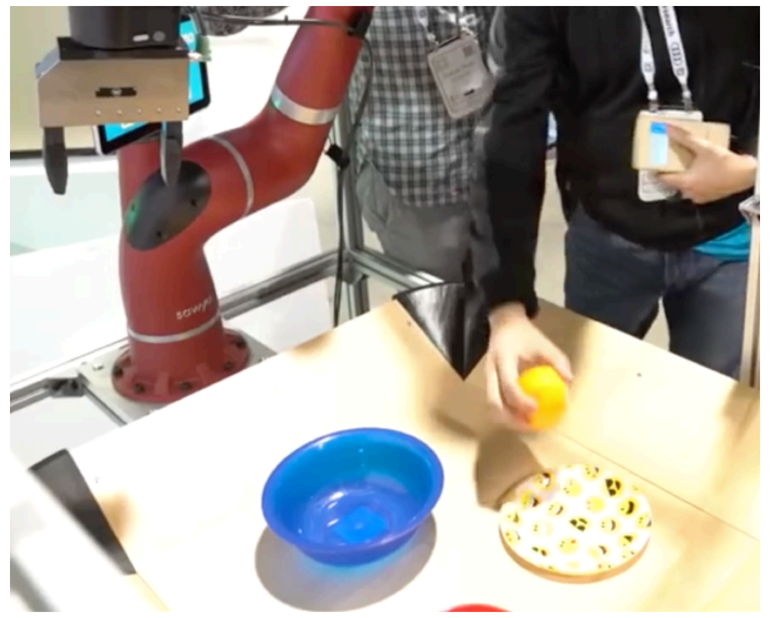

Demo at AI4ALL Outreach Camp

The students were unimpressed.

(but still had fun)

Takeaways

Can we build an agent that can do ***many tasks***?

from **pixel observations**, with **limited supervision**, in the **physical world**

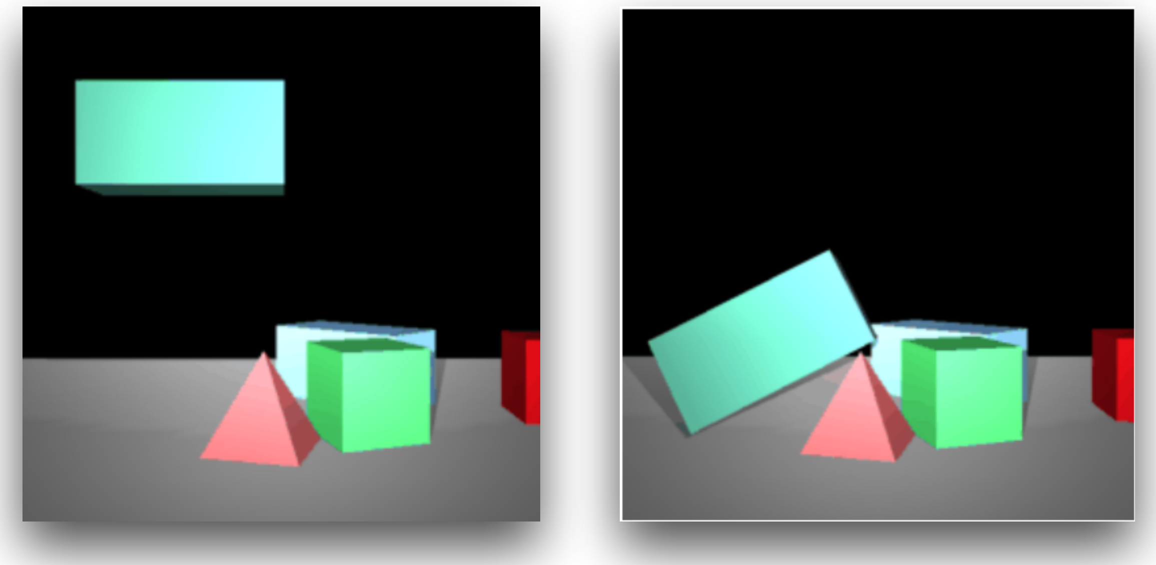

structured latent space

model for long-horizon tasks

+ complex, **long-horizon** tasks

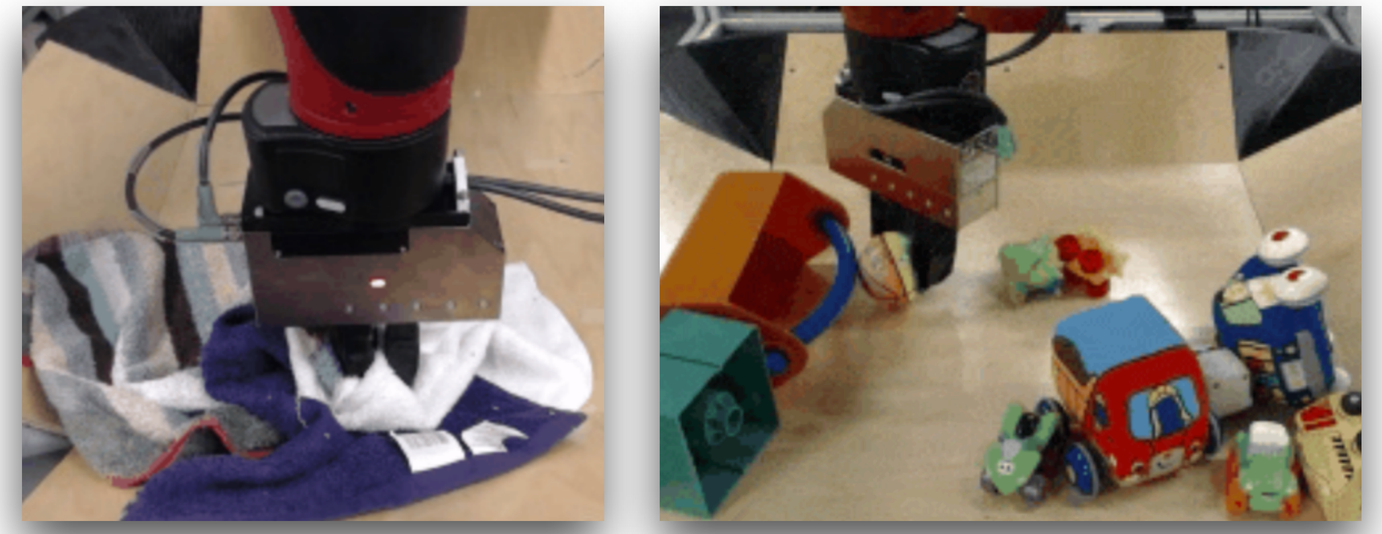

modeling **diverse**, **open-world** environments

+ significant **object diversity**
+ **minimal supervision**

Future work: best of both worlds?

Future work: How can we build better, more useful models of the world?

Can we model **uncertainty** over future observations?

More and more uncertainty over time.

Stochastic adversarial video prediction

Lee, Zhang, Ebert, Abbeel, Finn, Levine. 2018

Can we **adapt the model** with a small amount of experience?

Physical properties unknown until interaction.

Few-shot, online model adaptation

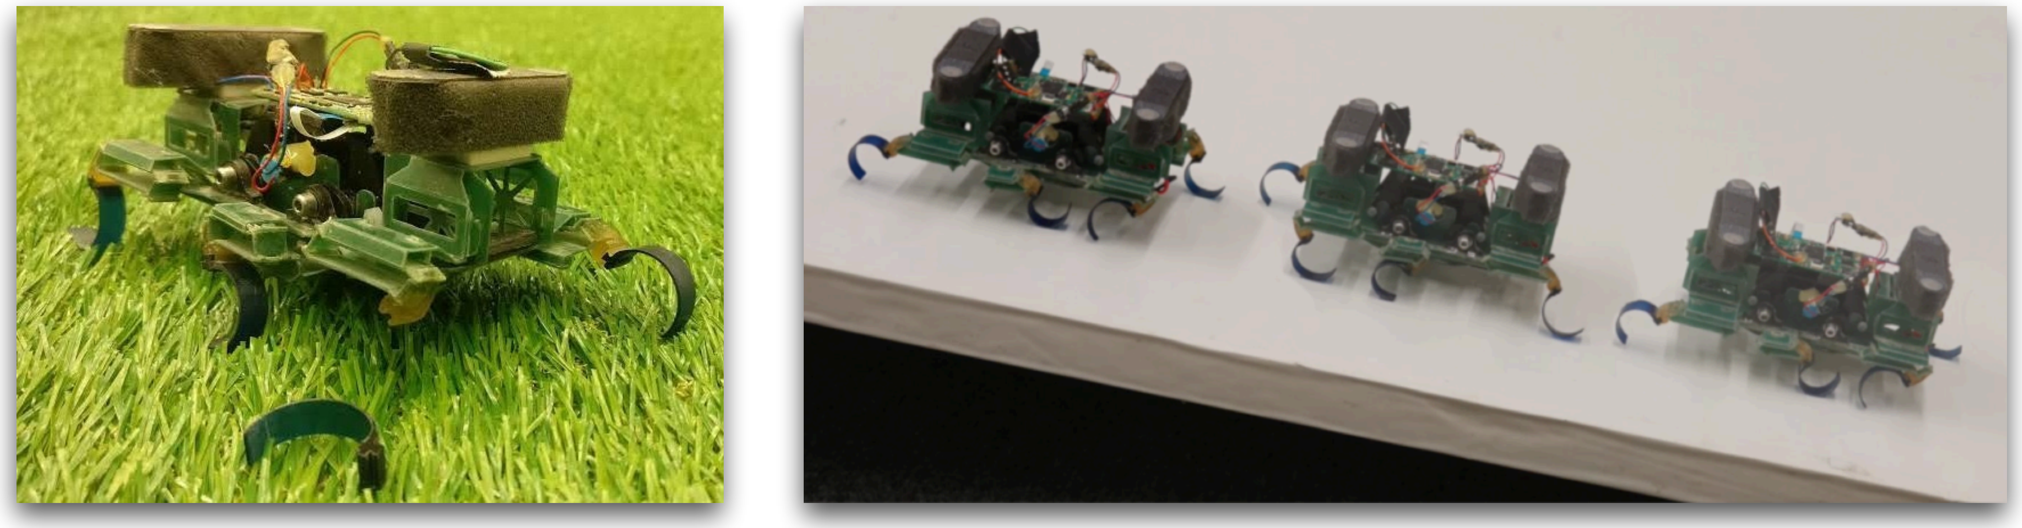

Nagabandi*, Clavera*, Liu, Fearing, Abbeel, Levine, Finn. 2018

How should we **model the reward**?

Agents need internal representation of the goal in the real world.

Goal inference from images

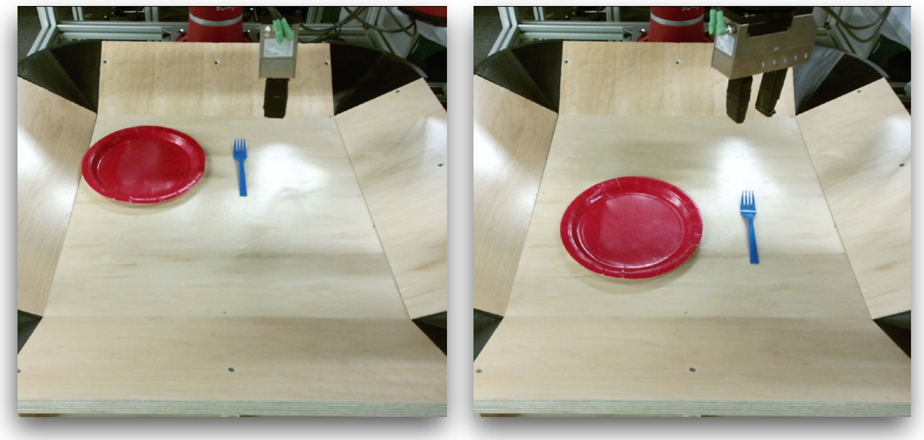

Xie, Singh, Levine, Finn. CoRL 2018

Collaborators

Frederik Ebert

Sudeep Dasari

Annie Xie

Avi Singh

Michael Janner

Sergey Levine

Pieter Abbeel

Bill Freeman

Josh Tenenbaum

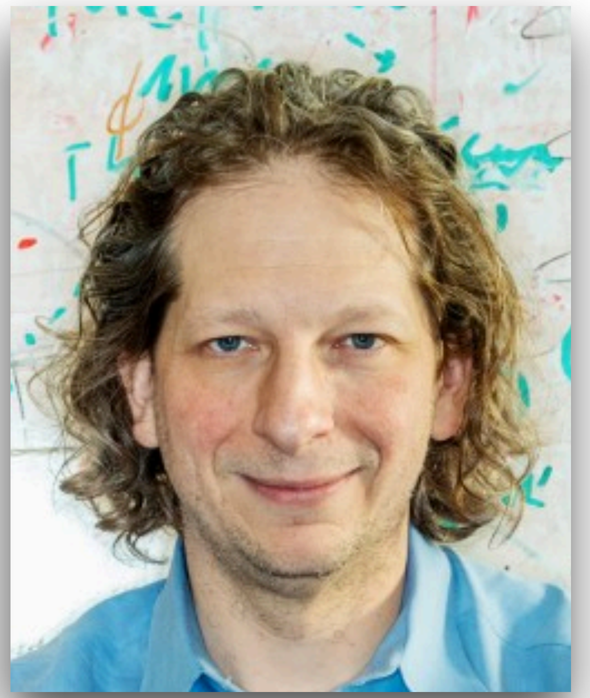

Jiajun Wu

Papers, data, and **code** linked at: people.eecs.berkeley.edu/~cbfinn

Questions?