
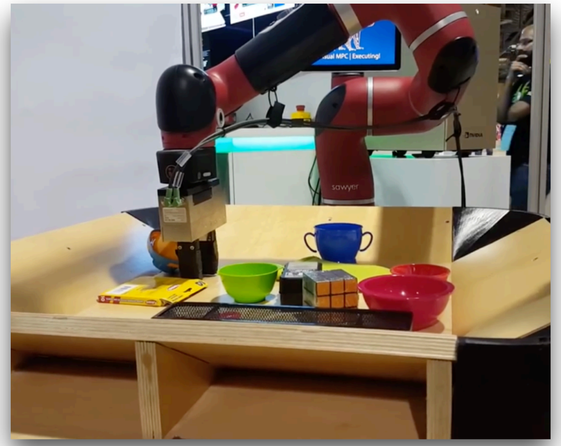

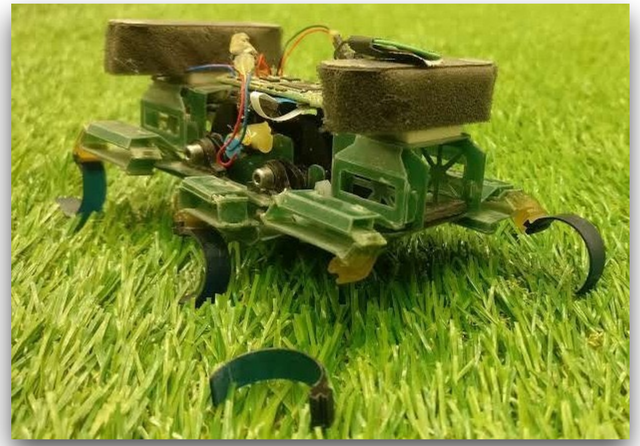

Learning Compound Tasks through Interaction and Observation

Chelsea Finn

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH
UC Berkeley

Google Brain

Stanford

Impressive Feats in AI

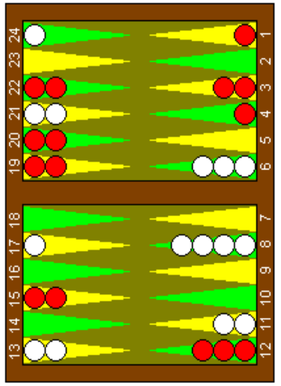

TD Gammon

Watson

helicopter acrobatics

machine translation

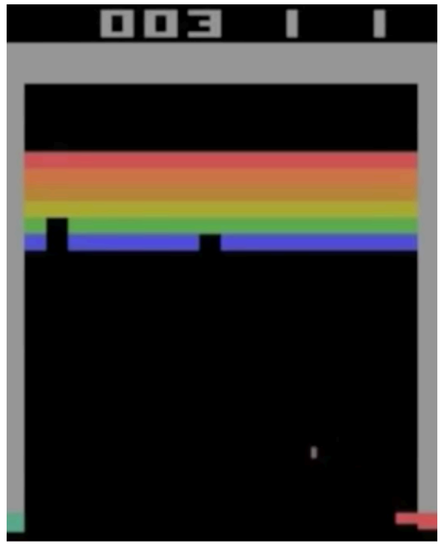

DQN

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 a robot that can do **many tasks**?

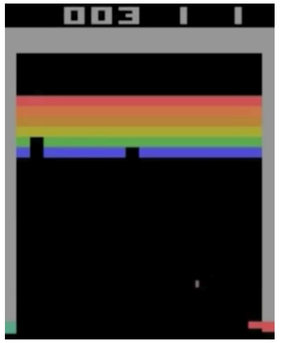

learning a **policy** in
a **closed universe**

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

model-based control

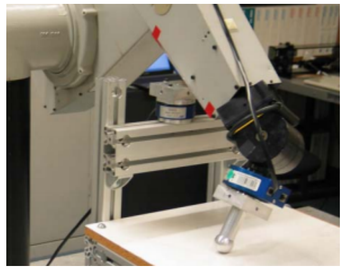

Petrovskaya, Park, Khatib '07

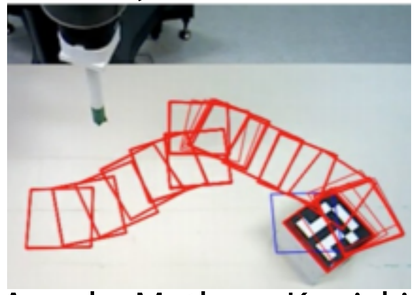

Arruda, Mathew, Kopicki,
Mistry, Azad, Wyatt '17

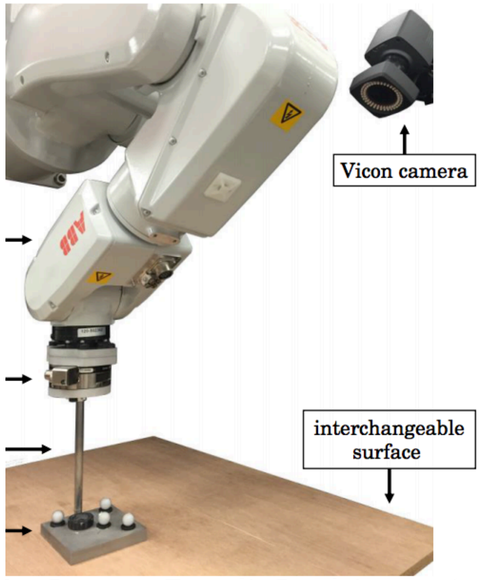

Yu, Bauza, Fazeli, Rodriguez '17

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

Can we build a robot that can do **many tasks**?

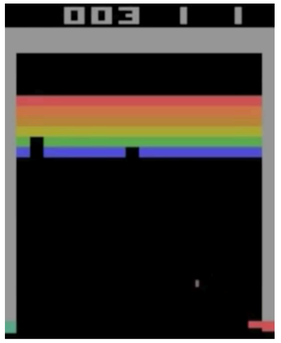

learning a **policy** in
a **closed universe**

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

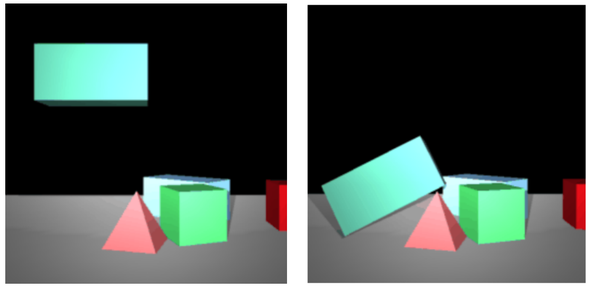

structured latent space
model for **long-horizon tasks**

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

long-horizon tasks in **diverse**,
open-world environments

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

1. Collect **diverse** interactions

Greater diversity —> more general-purpose model

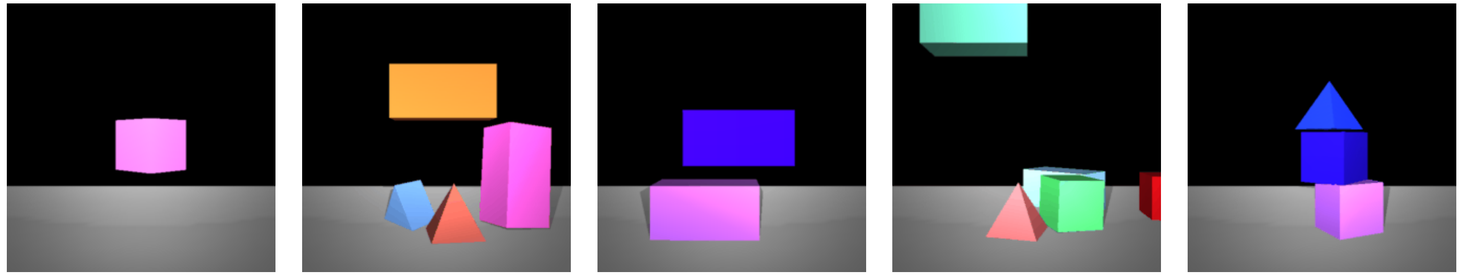

2. Learn **structured** representation & model

Structure —> long-horizon reasoning

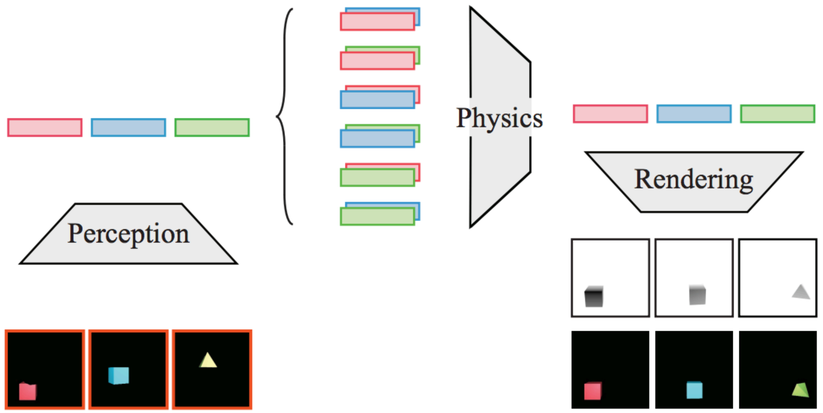

3. **Plan** using model

Online planning —> many tasks

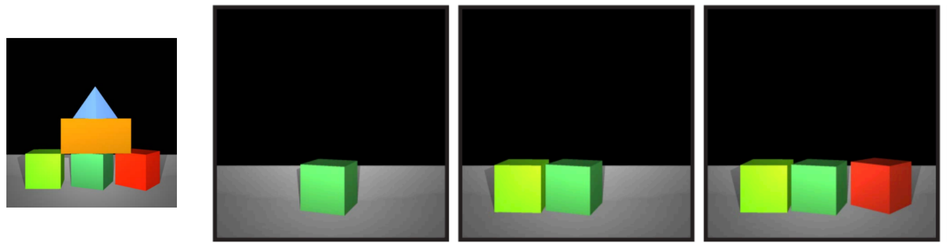

Goal: be able to build any tower

Michael Janner Jiajun Wu

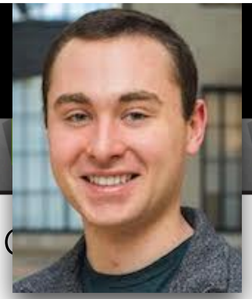

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR '19

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Follow up work: remove this assumption in Chang et al. '19

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

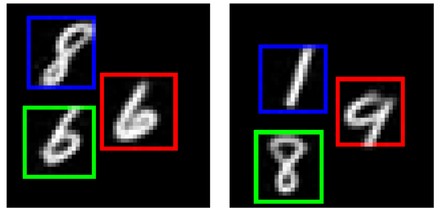

simple, 2D scenes

Wu et al. '17

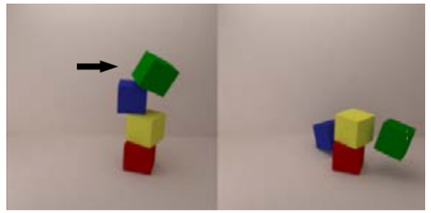

full supervision of
object properties

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR '19

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Follow up work: remove this assumption in Chang et al.'19

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

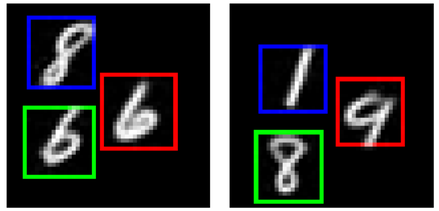

simple, 2D scenes

Wu et al.'17

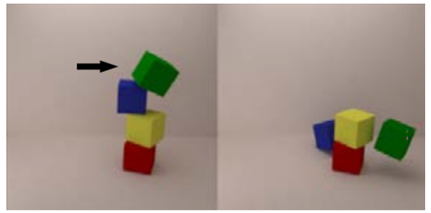

full supervision of
object properties

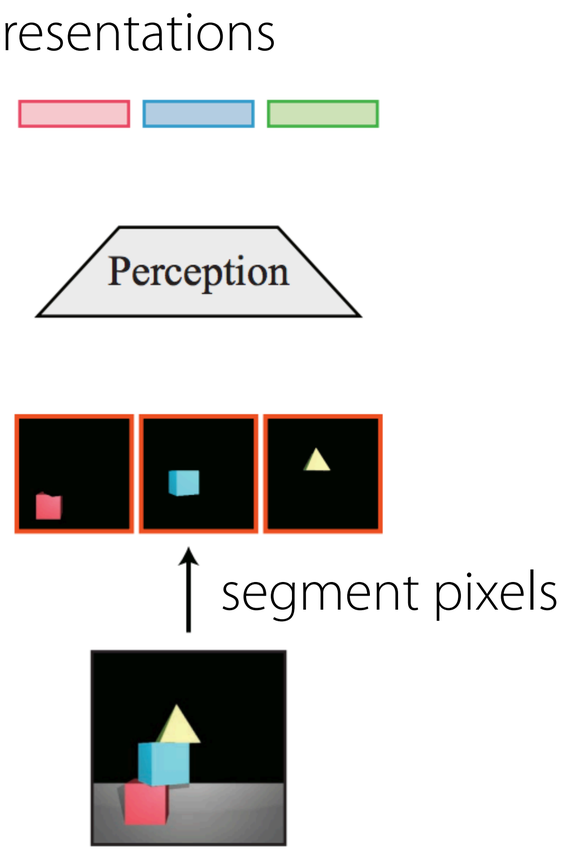

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR'19

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Follow up work: remove this assumption in Chang et al.'19

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

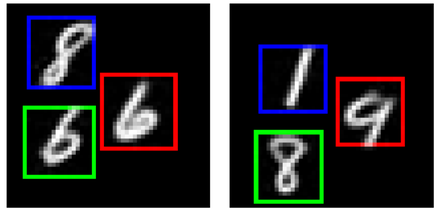

simple, 2D scenes

Wu et al.'17

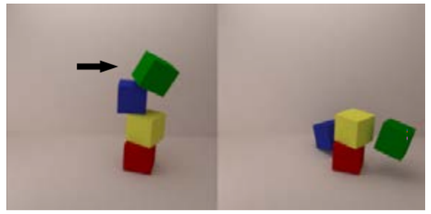

full supervision of
object properties

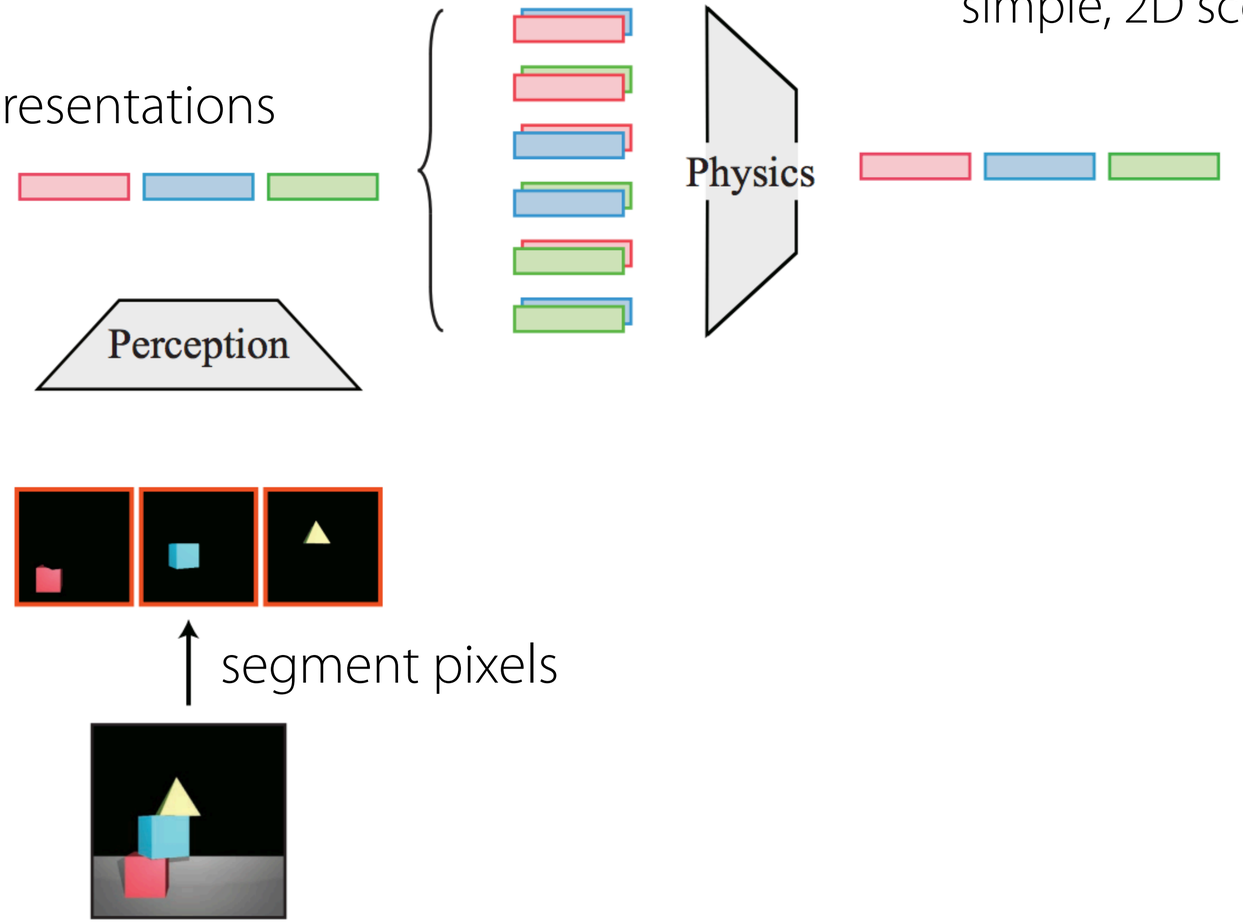

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR'19

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Follow up work: remove this assumption in Chang et al.'19

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

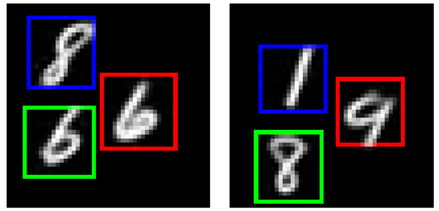

simple, 2D scenes

Wu et al.'17

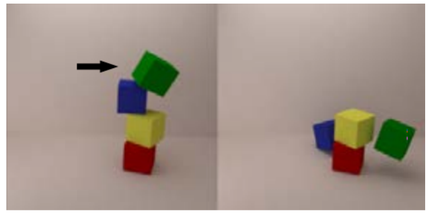

full supervision of
object properties

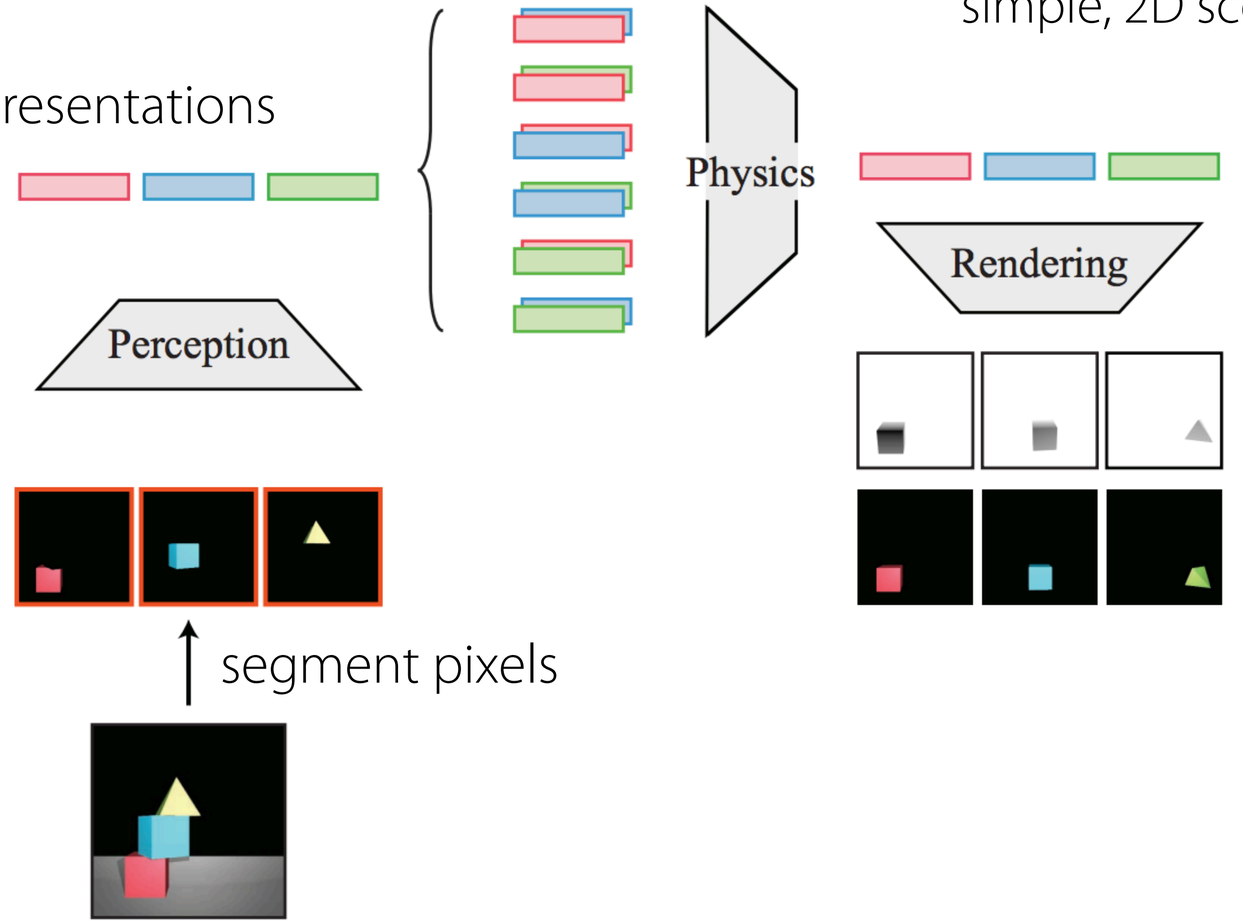

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR'19

Learn **structured** representation & model

object-centric model

Assume: object segmentation masks for individual frames

Follow up work: remove this assumption in Chang et al.'19

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

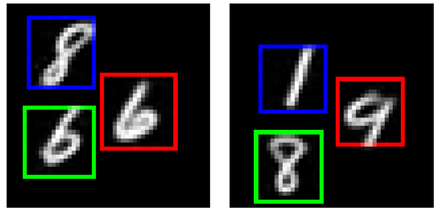

simple, 2D scenes

Wu et al.'17

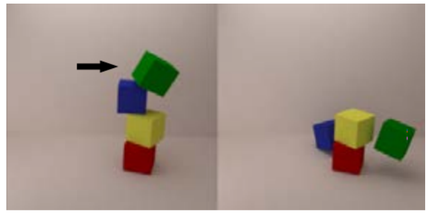

full supervision of
object properties

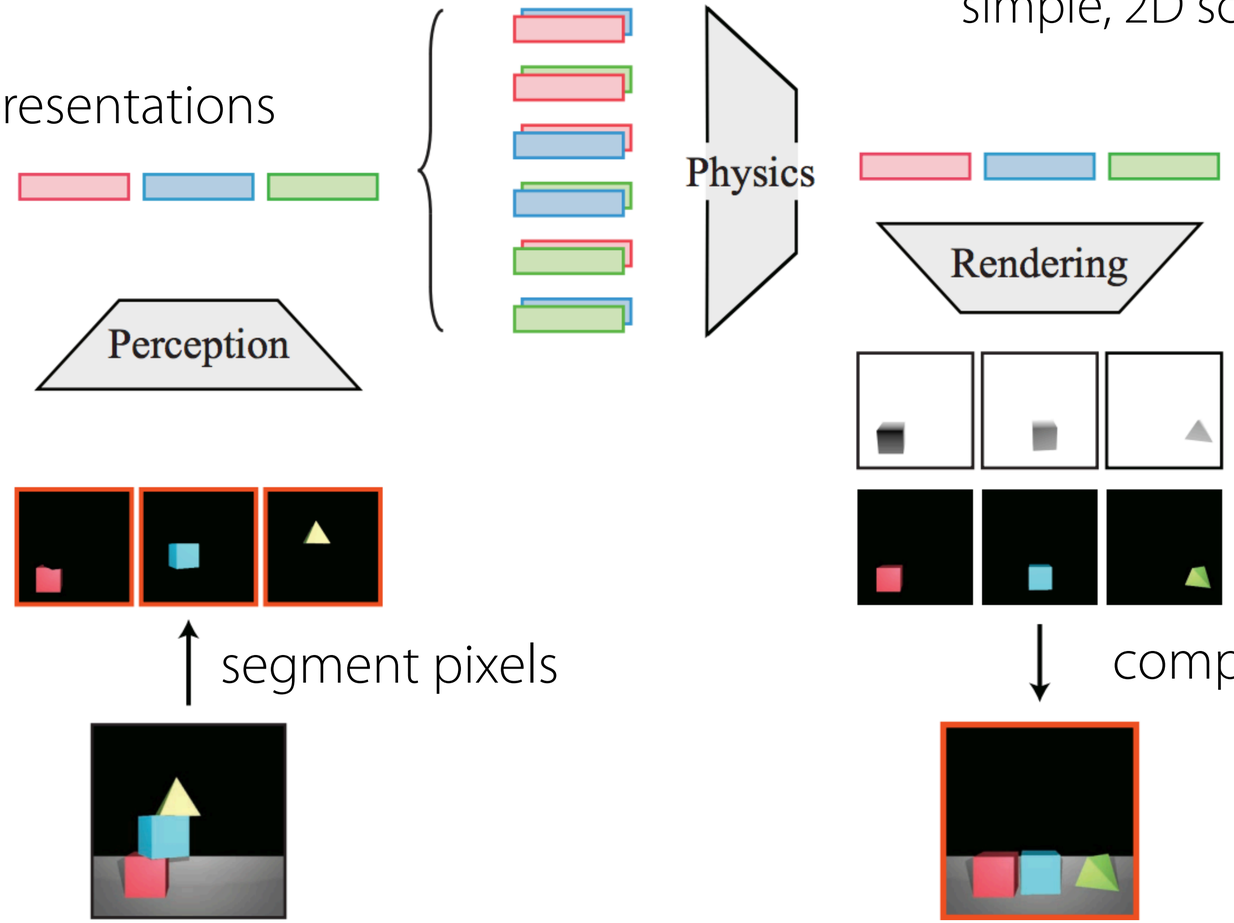

All modules trained with **reconstruction loss** (L_2+L_{VGG})

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR'19

Plan using model

goal space: image of object configuration

action space: which object & where to drop

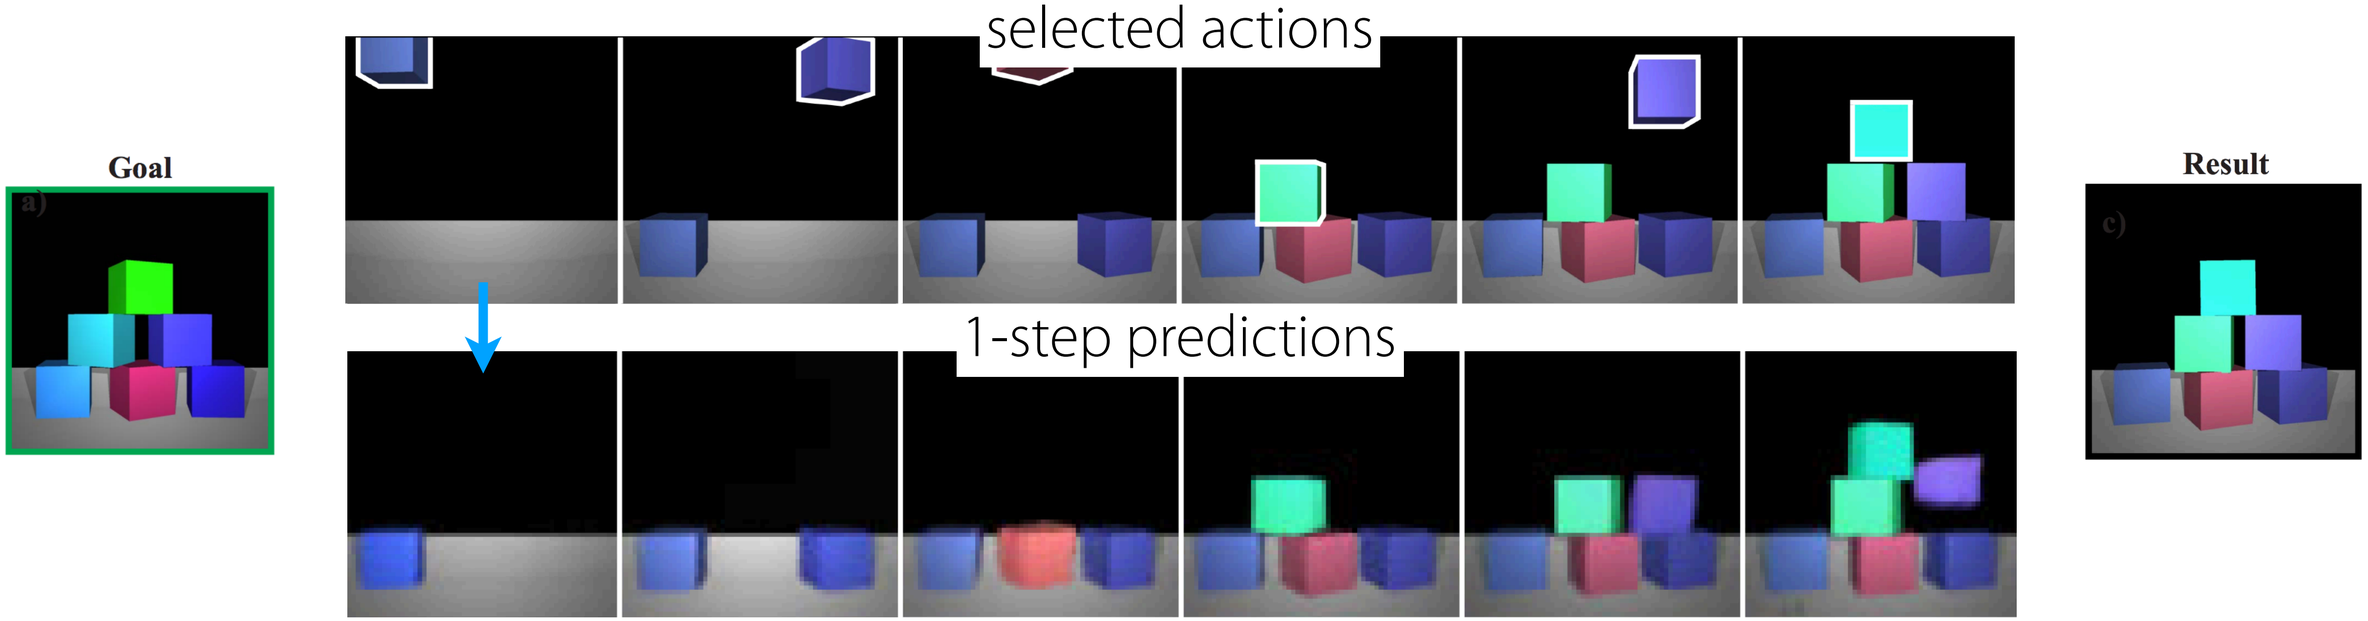

- **sampling-based**, **beam search** to plan action sequence
- evaluate action sequence based on **distance** in **latent space & pixel space**
- **replan** after each action

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR'19

Real world performance with single mode

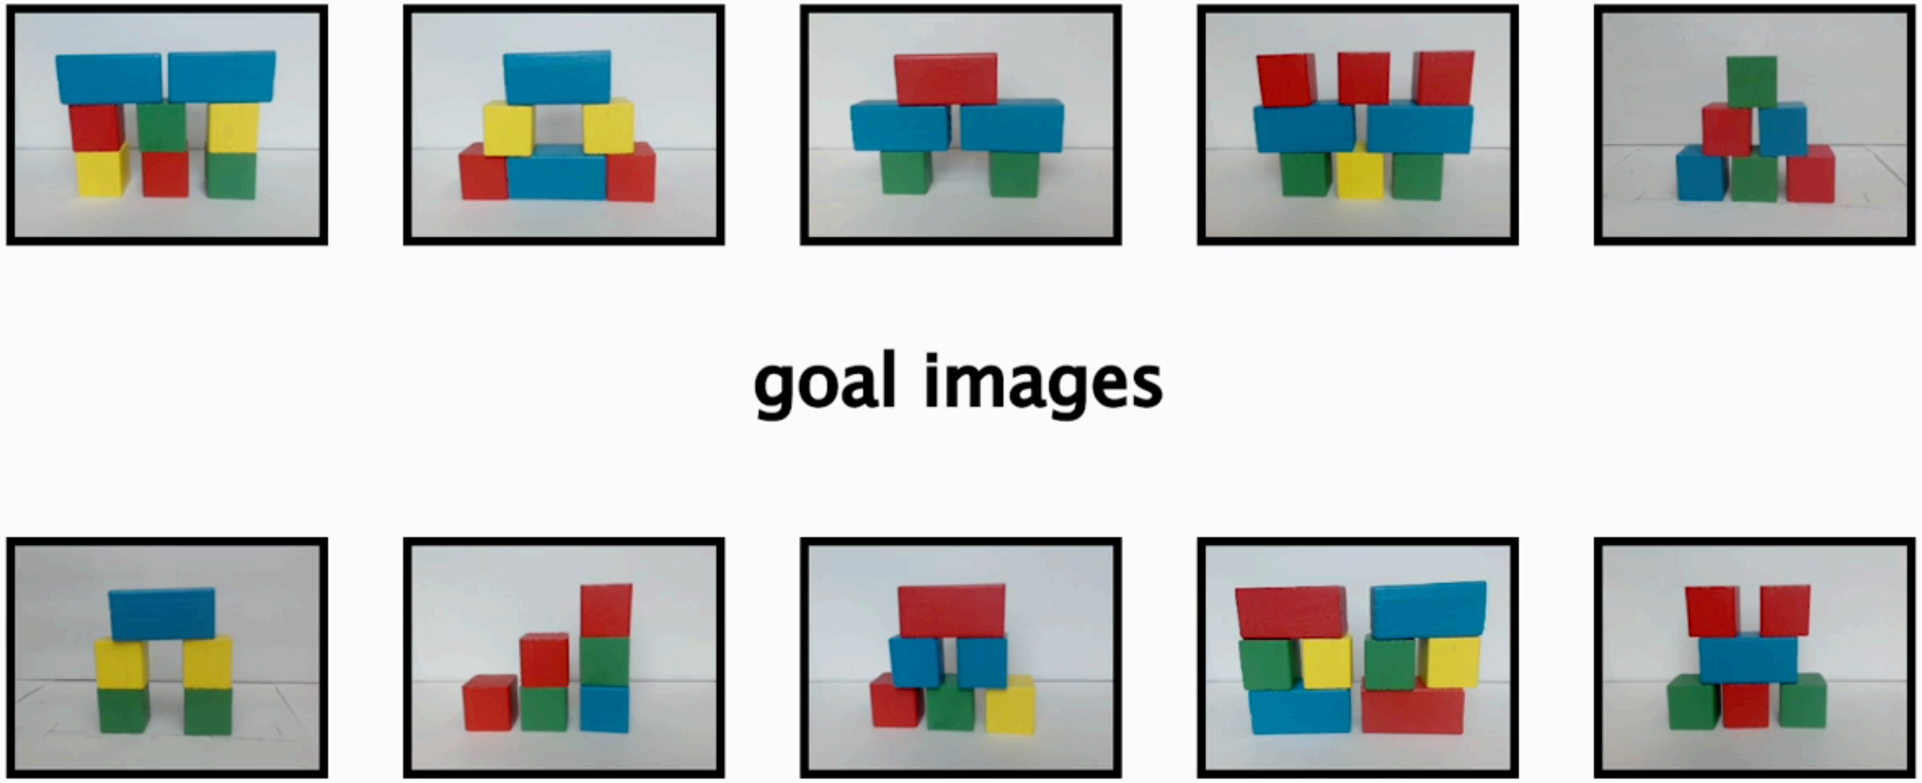

Takeaways

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

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

Janner, Levine, Freeman, Tenenbaum, Finn, Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. ICLR '19

Can we build a robot that can do **many tasks**?

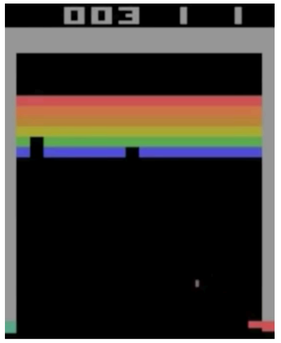

learning a **policy** in
a **closed universe**

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

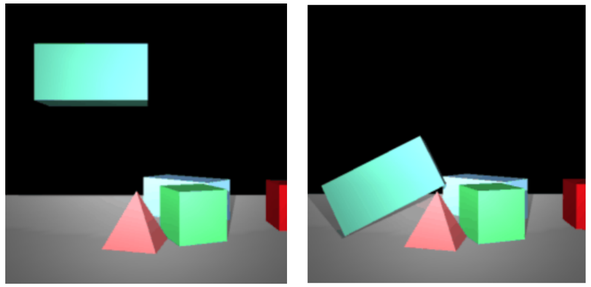

structured latent space
model for **long-horizon tasks**

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

long-horizon tasks in **diverse**,
open-world environments

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

Can we build a robot that can do **many tasks**?

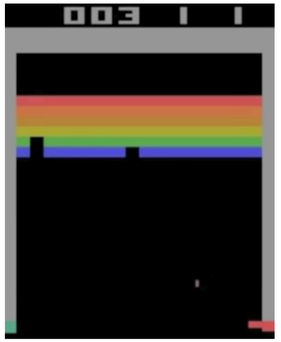

learning a **policy** in
a **closed universe**

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

structured latent space
model for **long-horizon tasks**

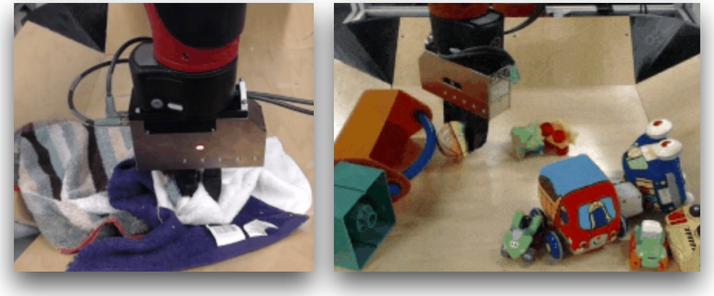

modeling **diverse**, **open-world**
environments

long-horizon tasks in **diverse**,
open-world environments

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

Diverse Open-World Environments

self-supervised robot learning

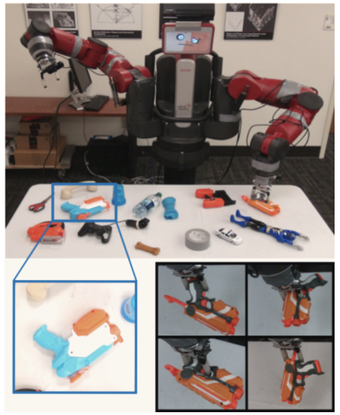

Pinto & Gupta '16

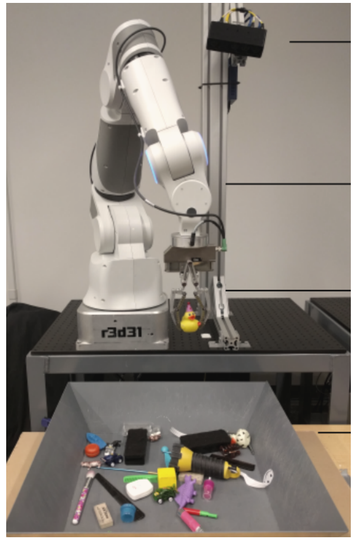

Levine, Pastor, Krizhevsky, Quillen '16

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 data

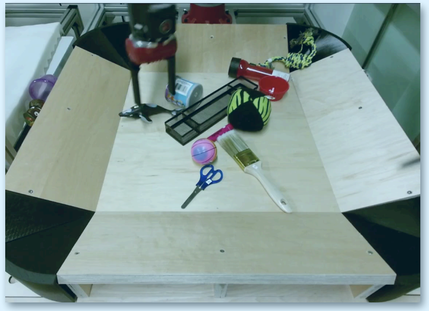

Collect **diverse** data in a **scalable** way

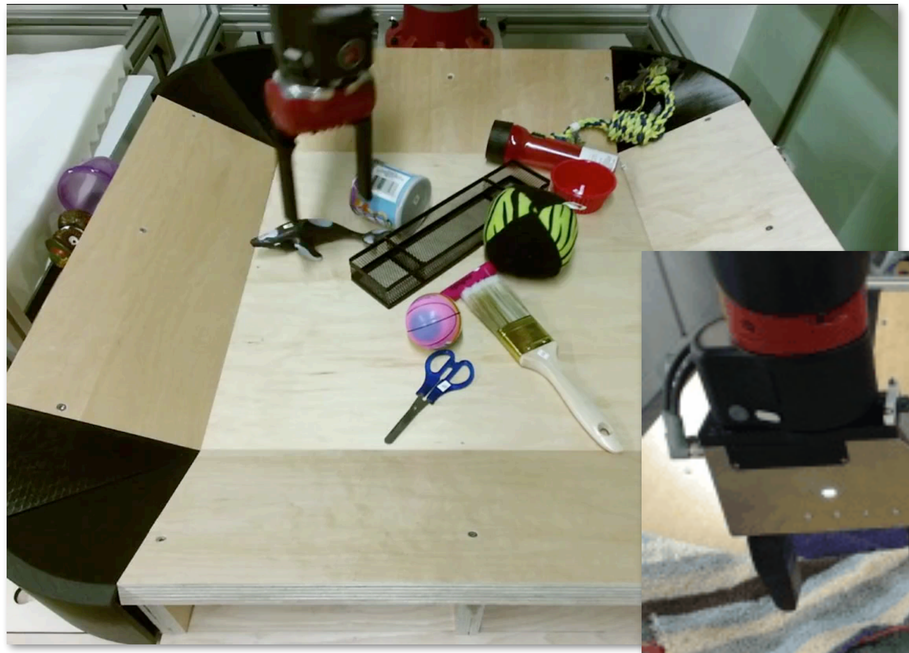

In contrast to policy learning: no notions of progress or success!

Collect data

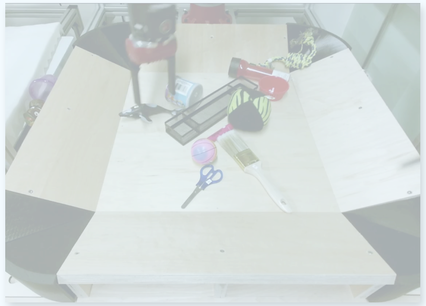

Learn to predict

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

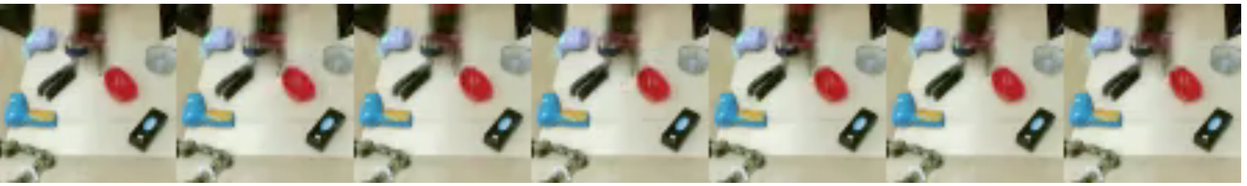

Contrast to:

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

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

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

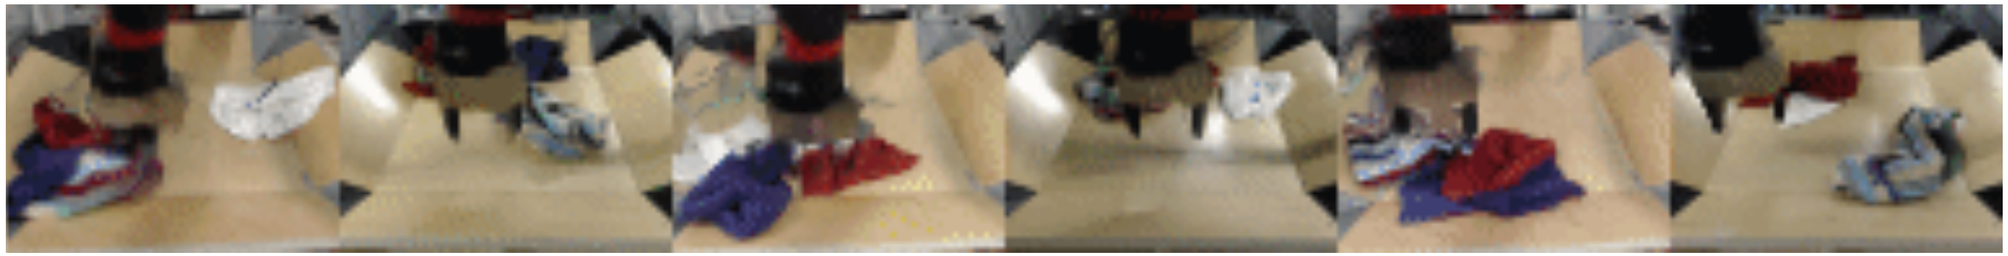

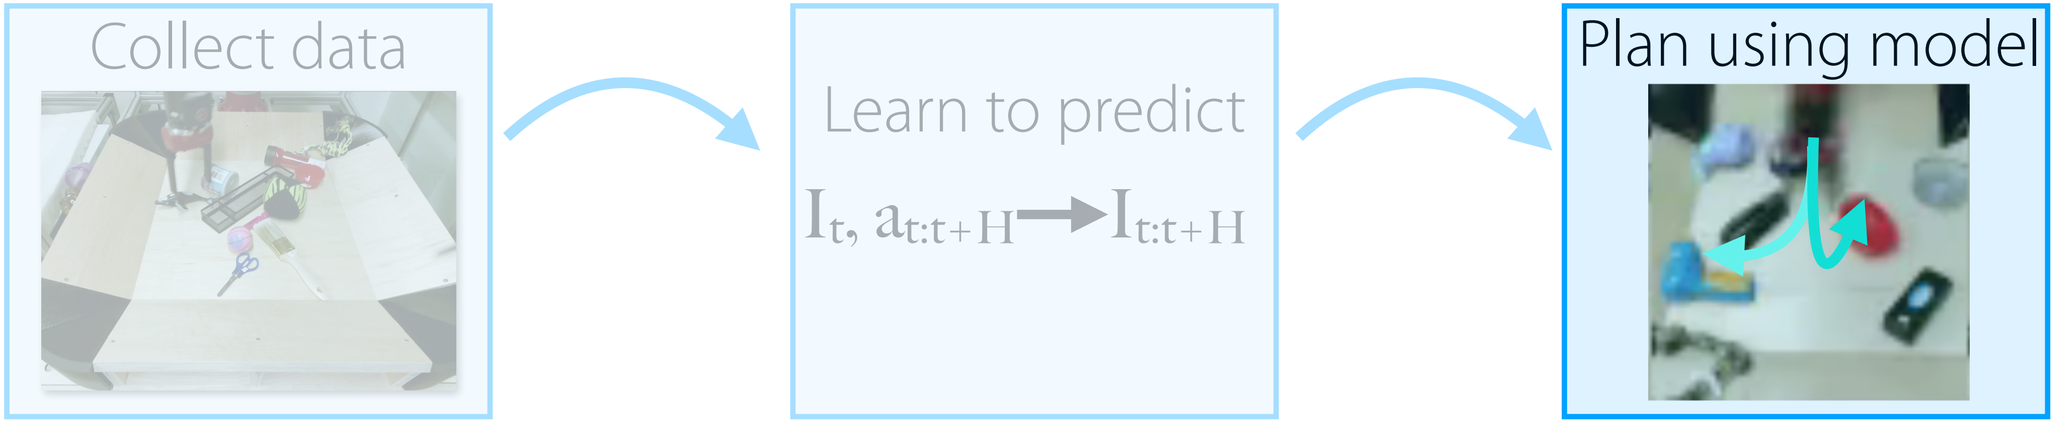

Planning with **visual foresight**:

- sampling-based optimization over actions
- replan action sequence at each time step

visual MPC

Finn & Levine. Deep Visual Foresight for Planning Robot Motion. ICRA '17

How to predict video?

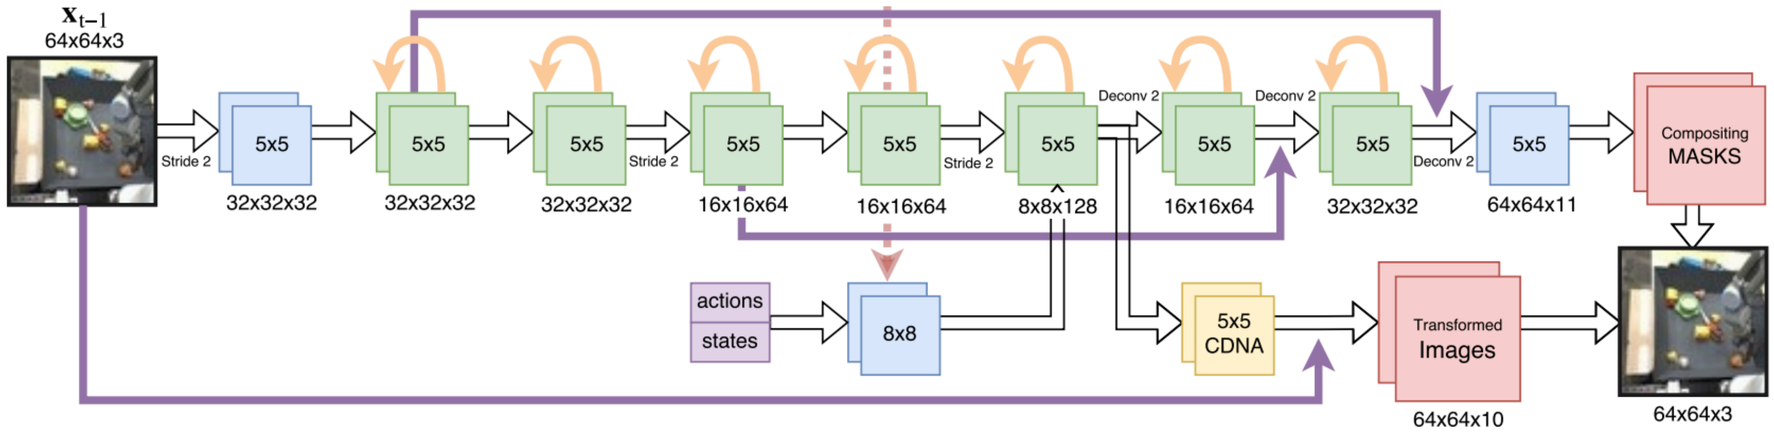

- deep recurrent network
- multi-frame prediction
- action-conditioned
- explicitly model motion

Finn, Goodfellow, Levine NIPS '16

Finn & Levine ICRA '17

Which future is the best one?

Human specifies a goal by:

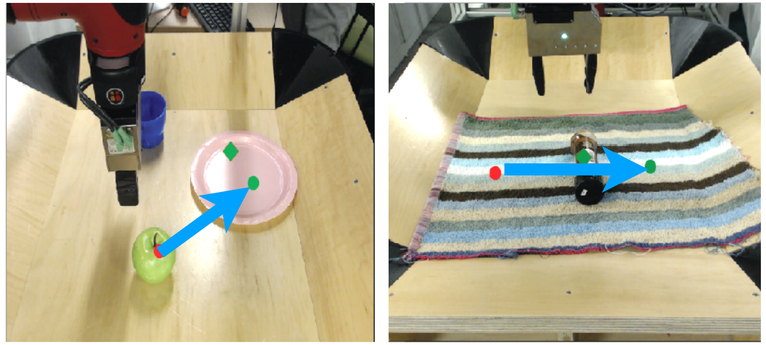

Selecting where pixels should move.

Providing an image of the goal.

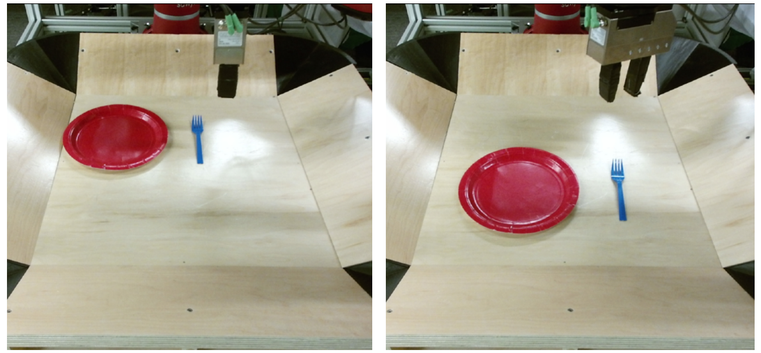

Providing a few examples of success.

Finn & Levine ICRA '17
Ebert, Lee, Levine, Finn CoRL '18
Xie, Singh, Levine, Finn CoRL '18

How it works

Specify goal

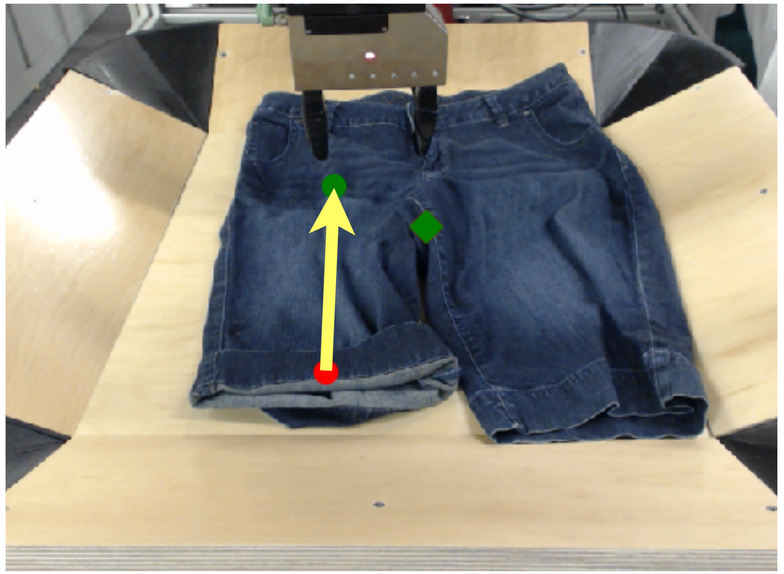

Visual MPC execution

Visual MPC w.r.t. goal

Frederik Ebert Sudeep Dasari

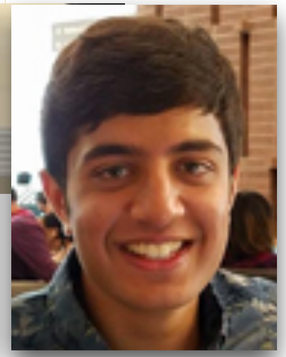

Ebert*, Finn*, Dasari, Xie, Lee, Levine. Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control.

How it works

Given 5 examples of success

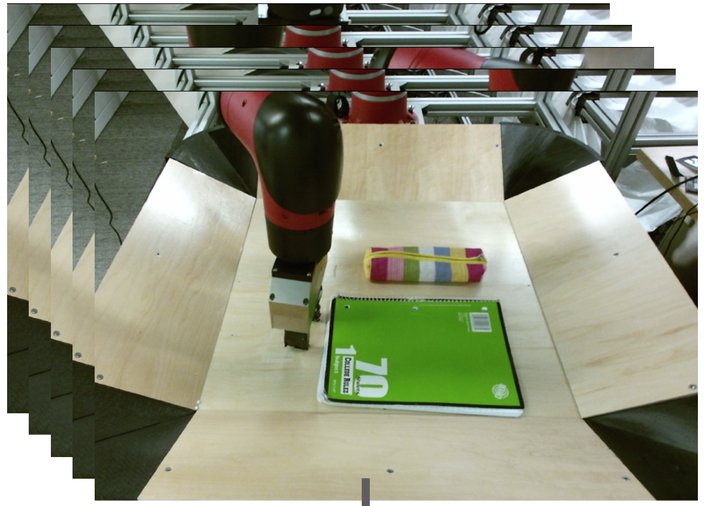

infer goal classifier

visual MPC w.r.t. goal classifier

Visual MPC with learned objective

Annie Xie

Avi Singh

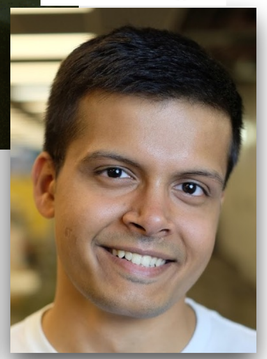

Xie, Singh, Levine, Finn. Few-Shot Goal Inference for Visuomotor Learning and Planning, CoRL 2018

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

Video speed: 2x

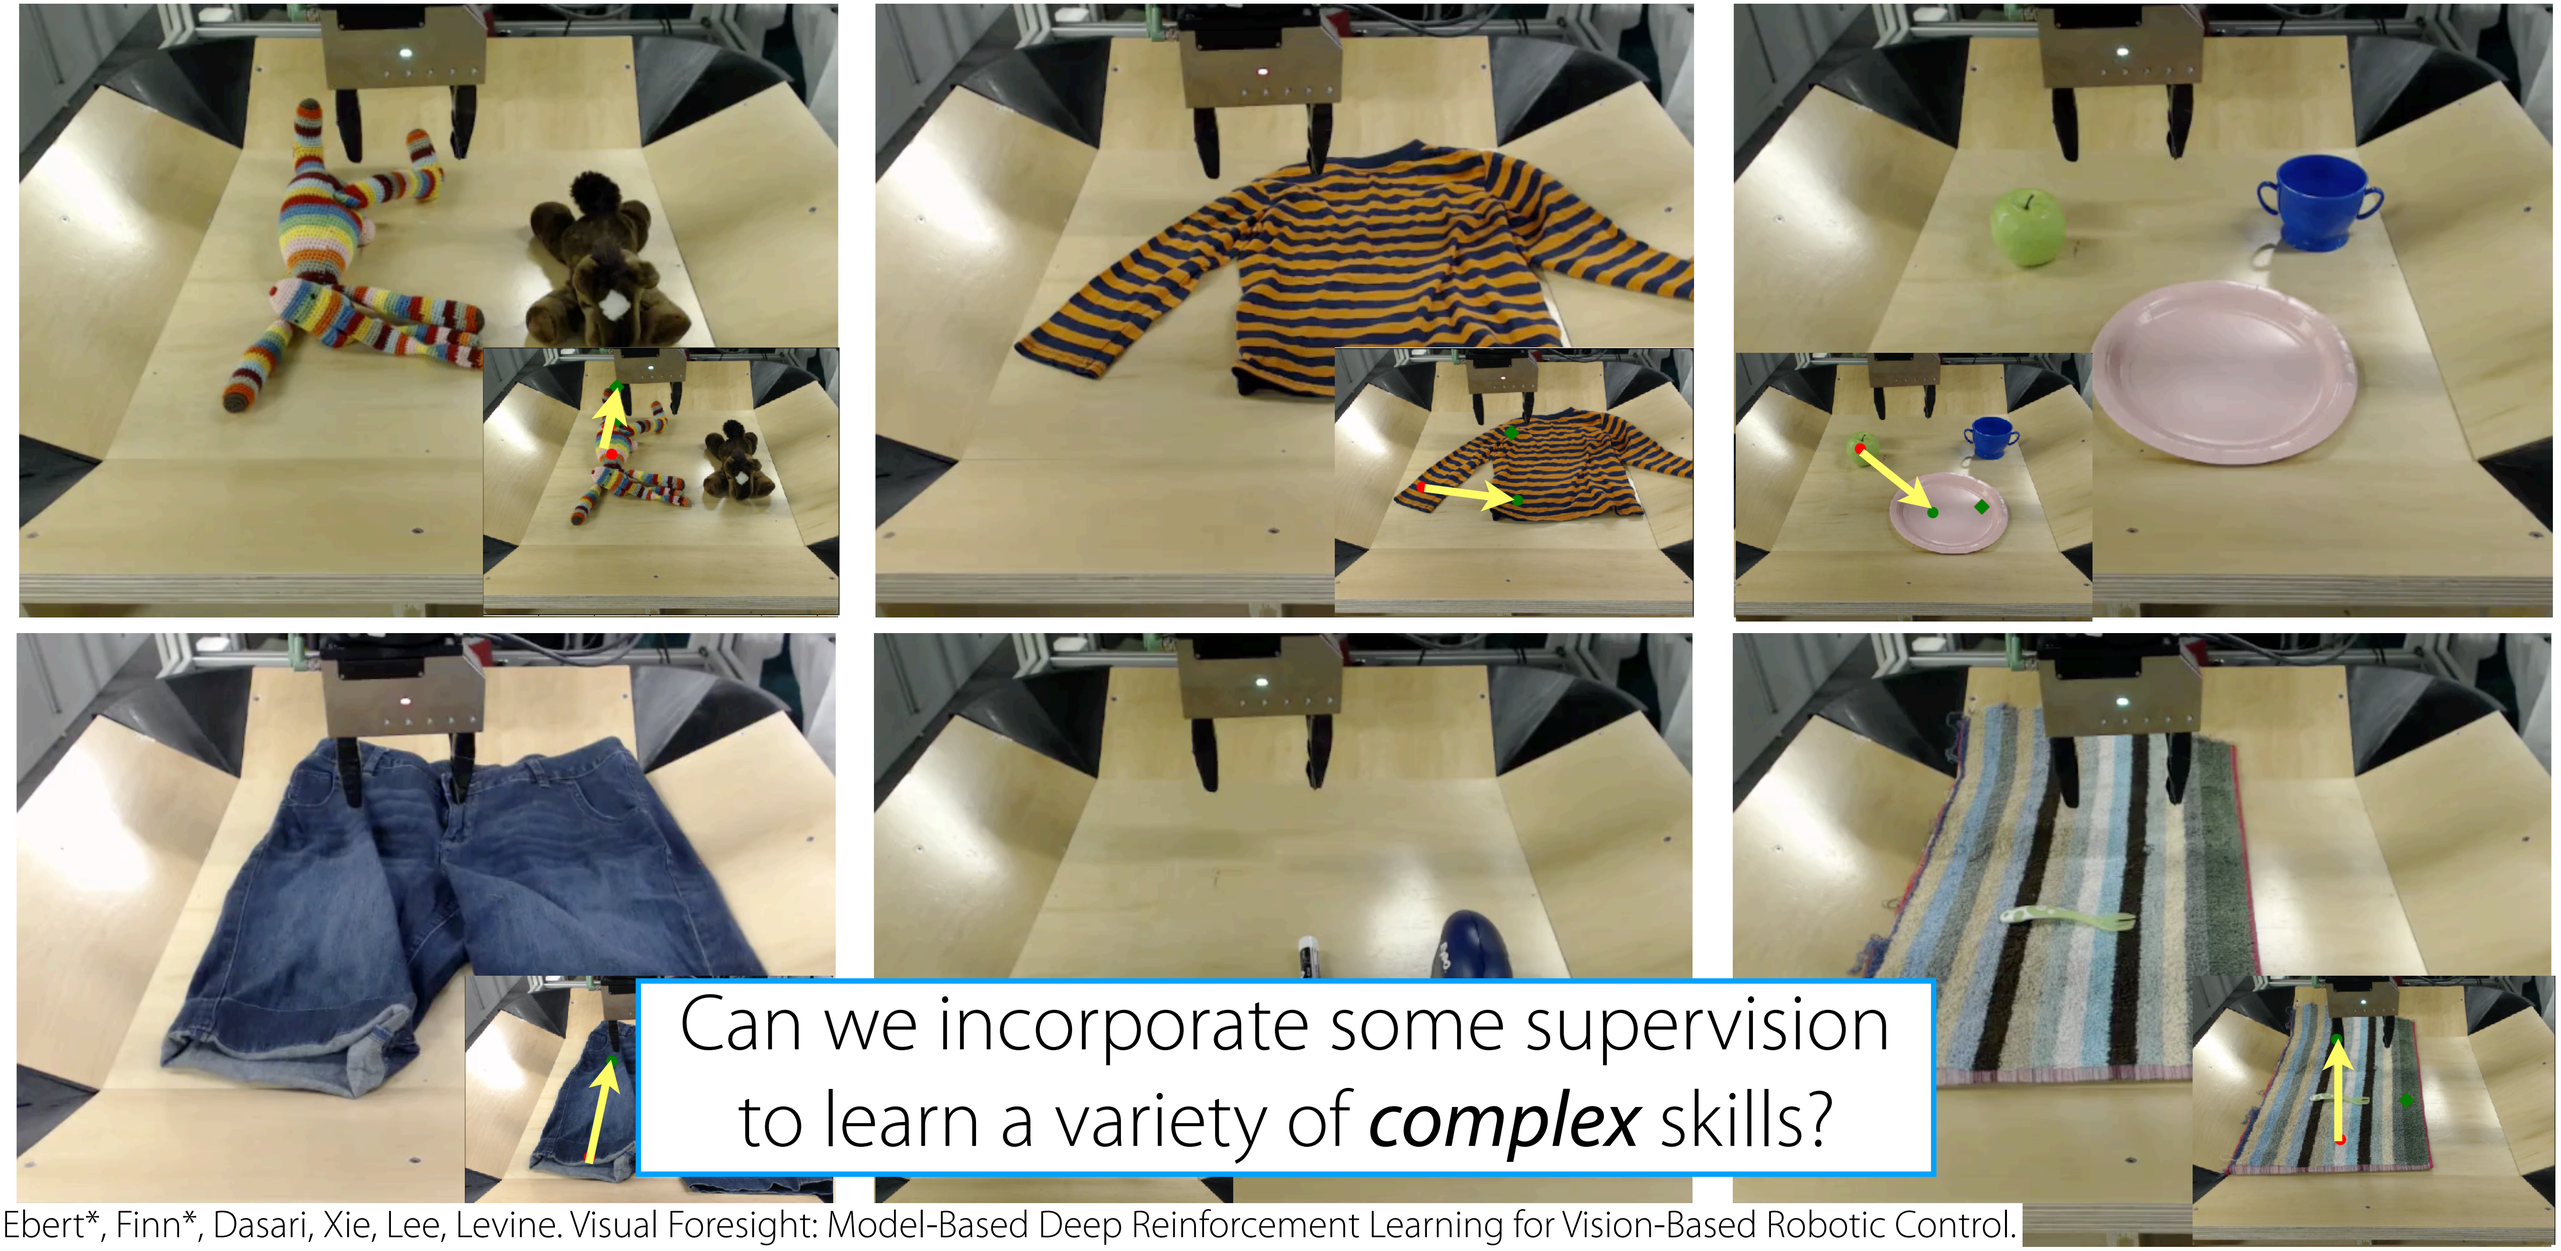

Can we incorporate some supervision to learn a variety of ***complex*** skills?

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

Can we build a robot that can do **many tasks**?

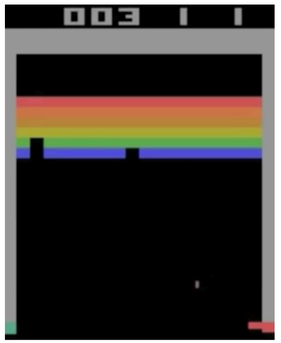

learning a **policy** in
a **closed universe**

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

structured latent space
model for **long-horizon tasks**

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

long-horizon tasks in **diverse**,
open-world environments

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

Collect data

Learn to predict

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

Plan using model

direct data collection

improve model

guide planning

Collect **diverse**, multi-task demonstrations

Fit model of $p(a_{t:t+H} \mid I_t)$ to the demonstration data.

Example multi-task demonstrations:

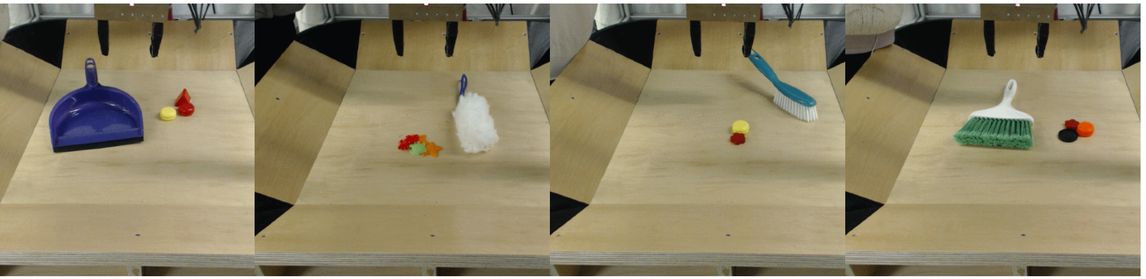

Samples from action proposal model:

Annie Xie

Frederik Ebert

Xie, Ebert, Levine, Finn. Improvisation through Physical Understanding, RSS '19

How it works

Specify goal

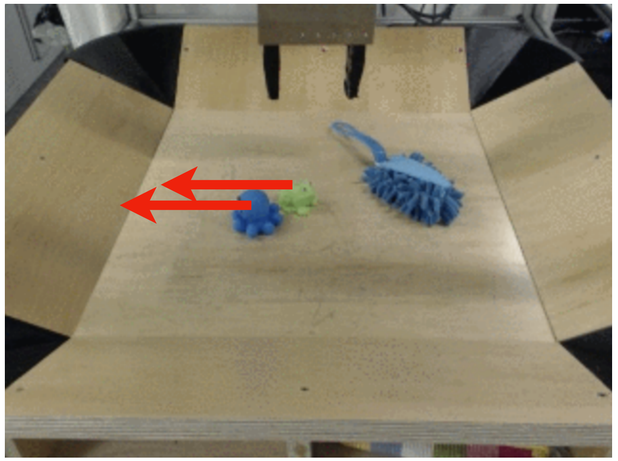

Guided visual planning w.r.t. goal

Executing actions

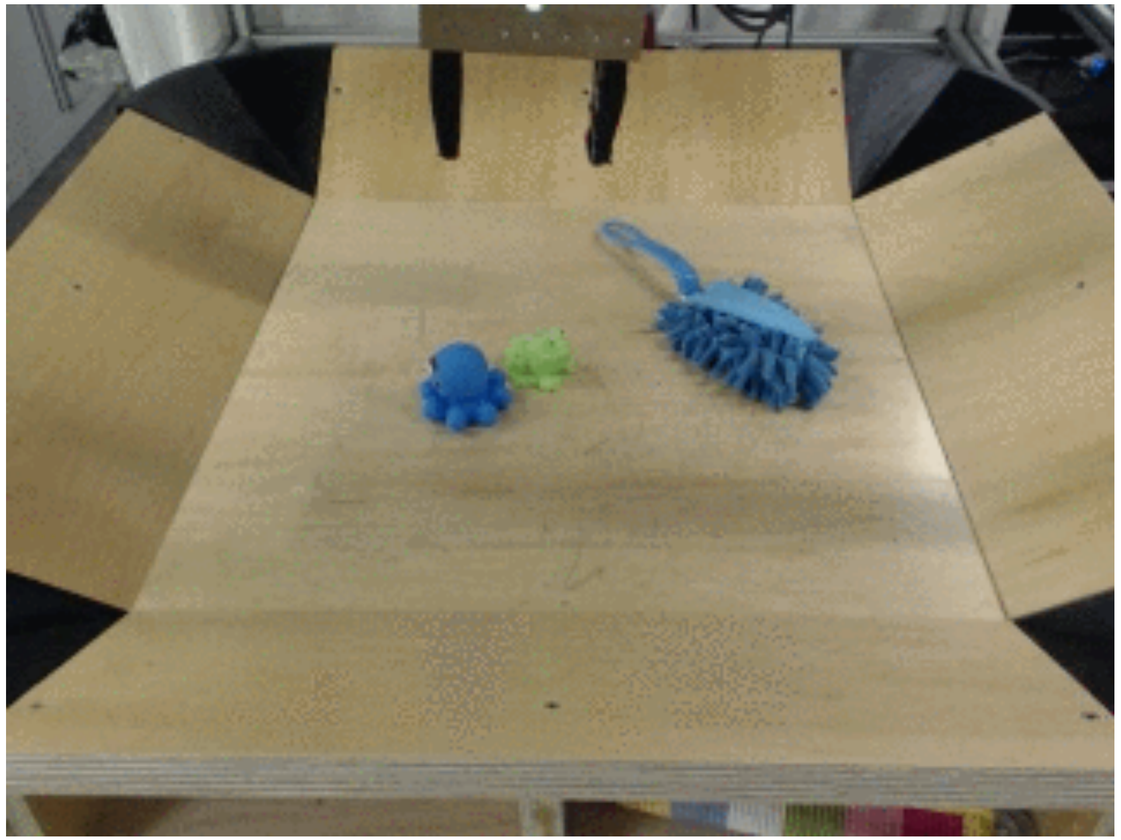

Xie, Ebert, Levine, Finn. Improvisation through Physical Understanding, RSS '19

Qualitative Experiments

solve new tasks

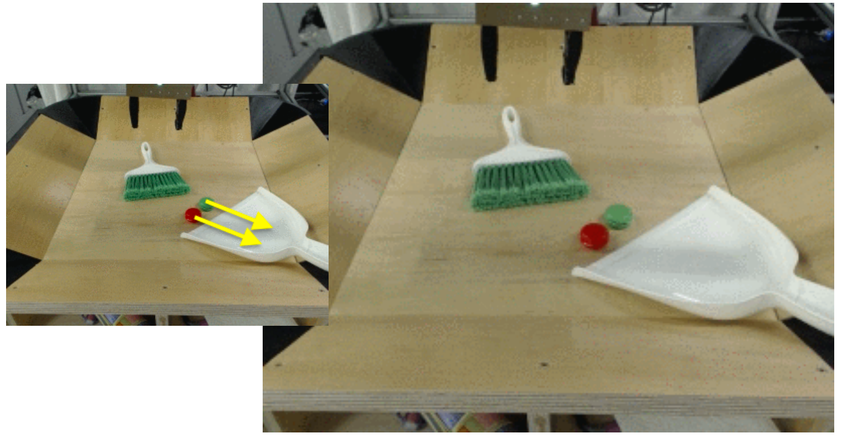

unseen tools

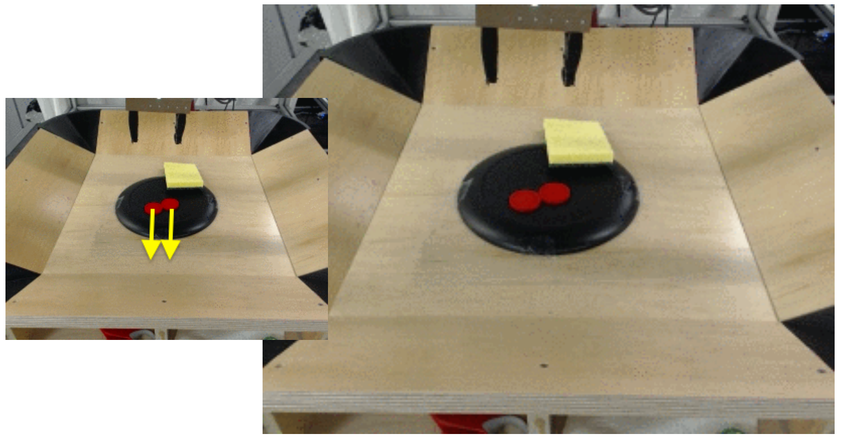

decide when to use a tool…

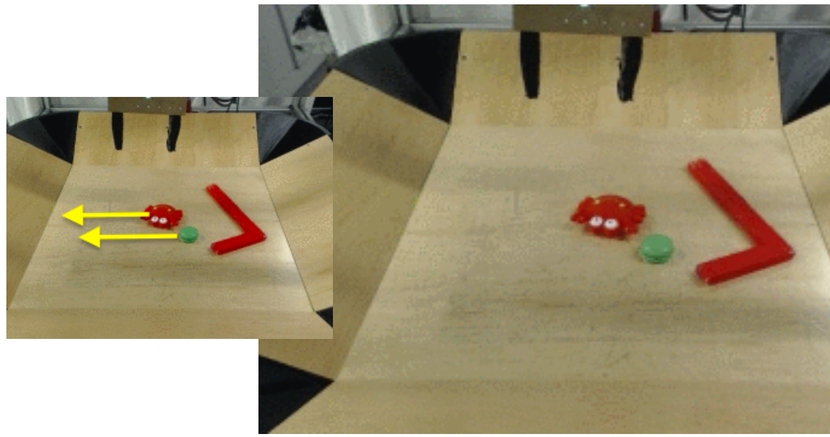

out-of-reach objects

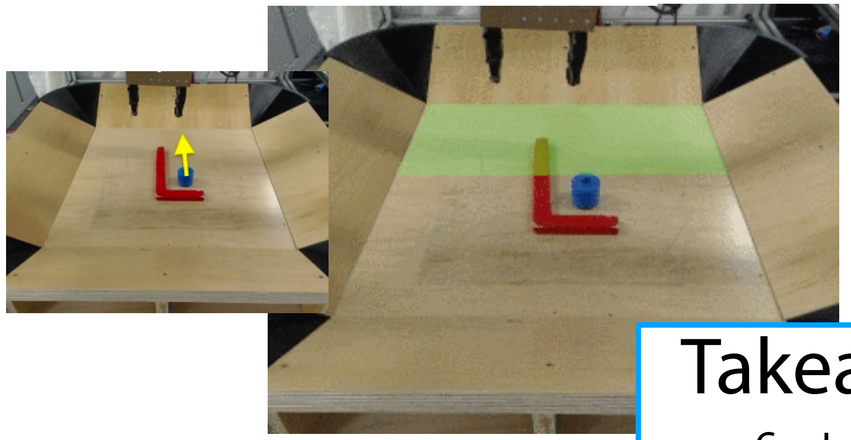

unseen *unconventional* tools

…and when not to

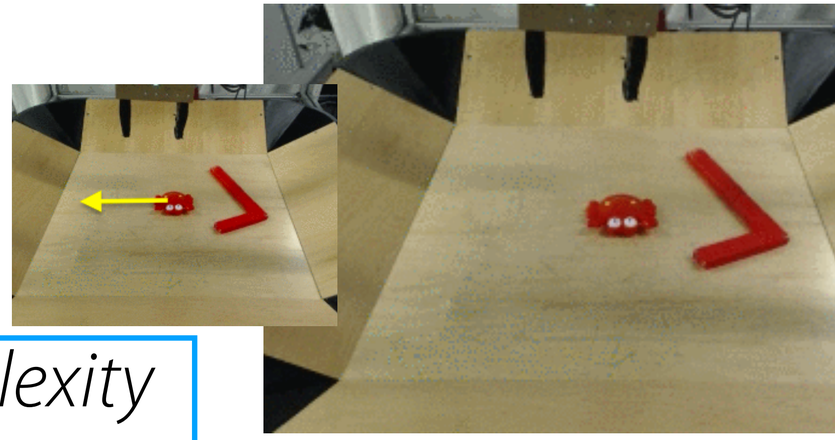

Takeaway: Achieve greater *complexity* of skills while maintaining *generality*.

Xie, Ebert, Levine, Finn. Improvisation through Physical Understanding, RSS '19

Demo at NIPS 2017: Long Beach, CA

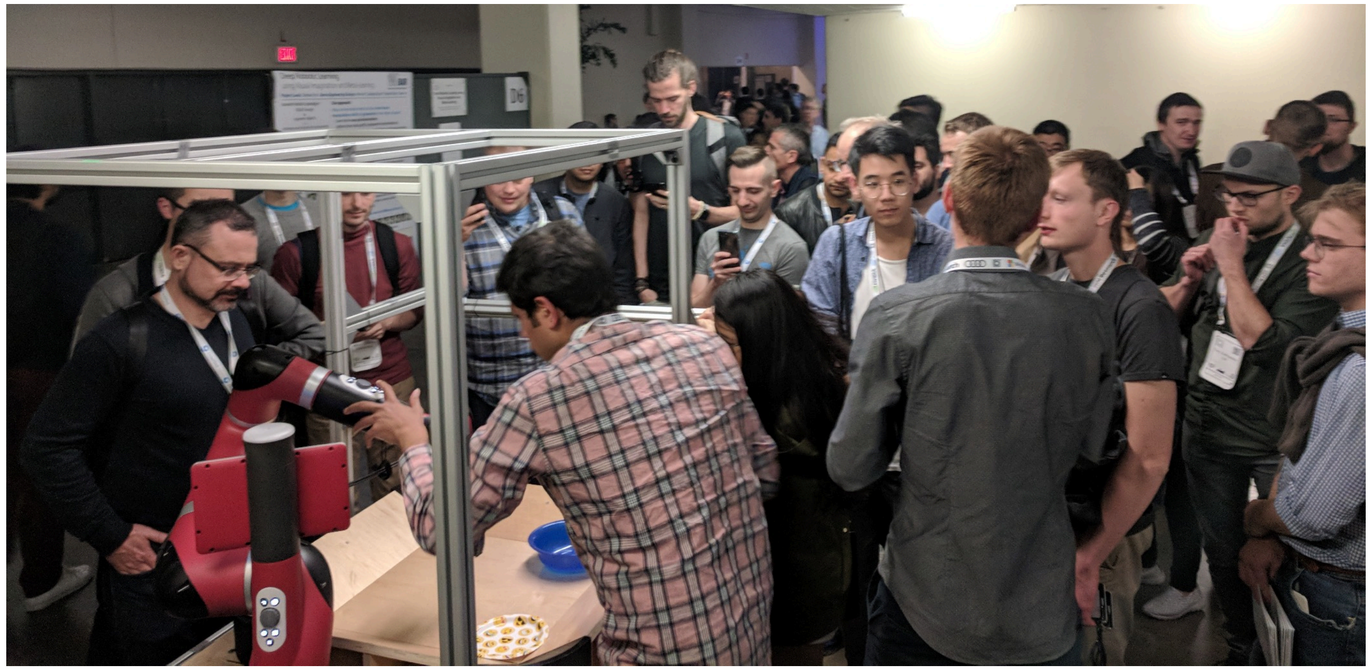

planning with visual models

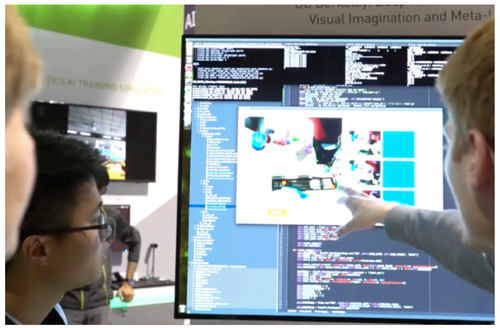

one-shot imitation

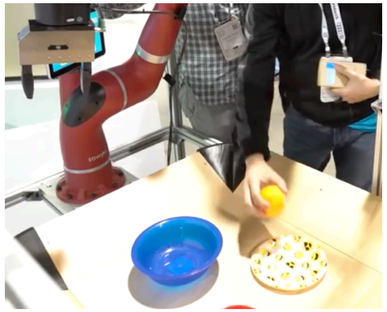

Demo at AI4ALL Outreach Camp

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**

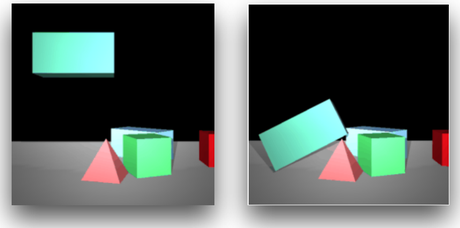

structured latent space model
for **long-horizon tasks**

+ complex, **long-horizon** tasks

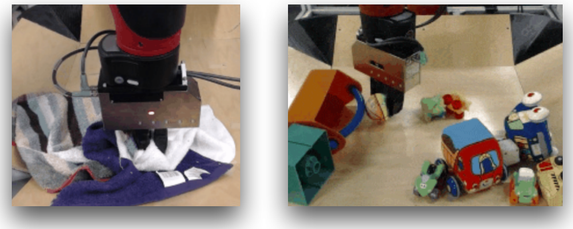

modeling **diverse**, **open-world** environments

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

long-horizon tasks in **diverse**, **open-world** environments

+ significant **object diversity**
+ **long-horizon tasks**

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

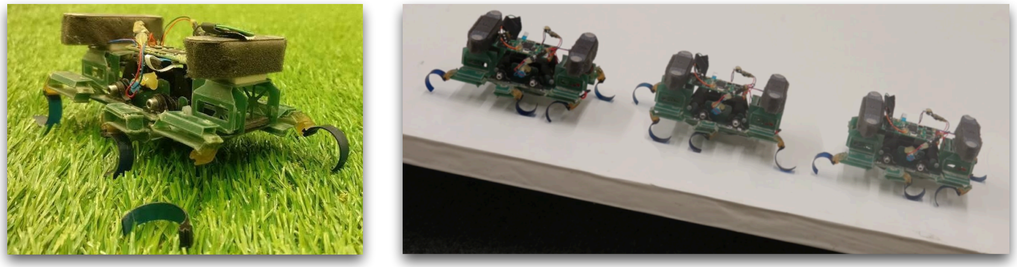

Nagabandi*, Clavera*, Liu, Fearing, Abbeel, Levine, Finn. 2018

How should we **model the reward**?

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

Goal inference from images

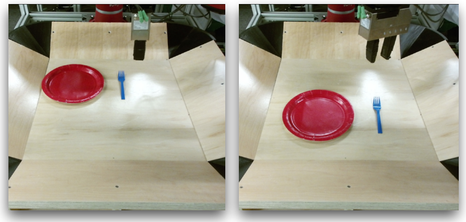

Xie, Singh, Levine, Finn. CoRL 2018

Collaborators & Students

Frederik Ebert

Sudeep Dasari

Annie Xie

Avi Singh

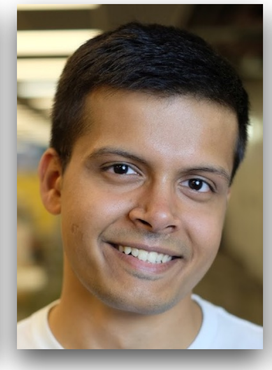

Michael Janner

Sergey Levine

Pieter Abbeel

Bill Freeman

Josh Tenenbaum

Jiajun Wu

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

Questions?