Generalizable Autonomy in Robot Manipulation

Animesh Garg | UNIVERSITY OF TORONTO VECTOR INSTITUTE NVIDIA

Generalizable Autonomy in Robot Manipulation

Vacuuming

Sweeping/Mopping

Cooking

Laundry

Generalizable Autonomy in Robot Manipulation

…

Vacuuming

Sweeping/Mopping

Cooking

Laundry

…

Diversity:
New Scenes,
Tools,…

Complexity:
Long-term
Settings

Generalizable Autonomy in Robot Manipulation

Vision: Build Intelligent Robotic Companions towards Human Enrichment and Augmentation

Generalizable Autonomy in Robot Manipulation

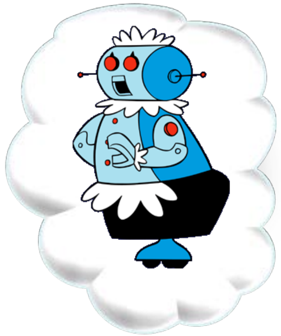

1956 Dartmouth AI Project

1956

Generalizable Autonomy in Robot Manipulation

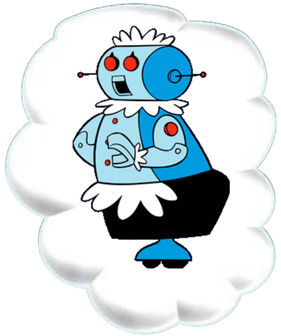

Dartmouth AI Meeting

UNIMATE
1st Industrial robot

1956 ‘61 1968

Generalizable Autonomy in Robot Manipulation

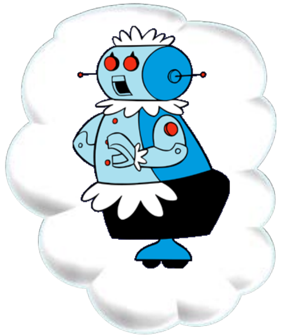

Dartmouth AI Me

UNIM
1st Indus

1956 ‘61 1968 2013

Generalizable Autonomy in Robot Manipulation

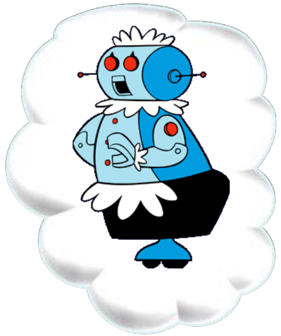

Dartmouth AI Meet

UNIMA
1st Industri

P

1956 ‘61 1968 2013 2018

Generalizable Autonomy in Robot Manipulation

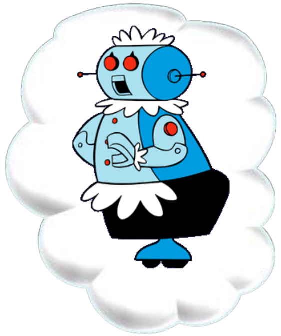

Dartmouth AI Me

UNIM
1st Indus

1956 '61 1968 2013 2018 2019

Generalizable Autonomy in Robot Manipulation

Then

Now

How to Generalize to Unstructured Scenarios?

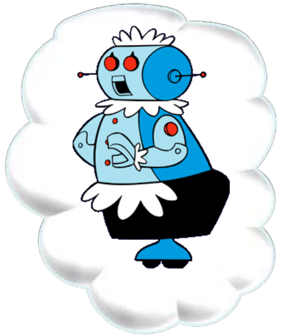

Dartmouth AI Meeting

UNIMATE
1st Industrial robot

Personal Robot Assistant

Atlas Robot Humanoid

Staged Repeated Tasks

Structured Environments

1956 ‘61 1968 2013 2018 2019

Generalizable Autonomy in Robot Manipulation

Then

Now

How to Generalize to Unstructured Scenarios?

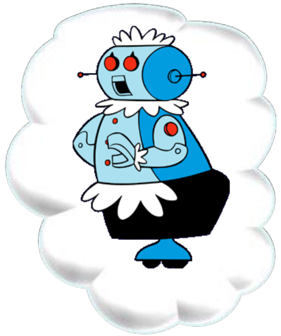

Manufacturing/Retail

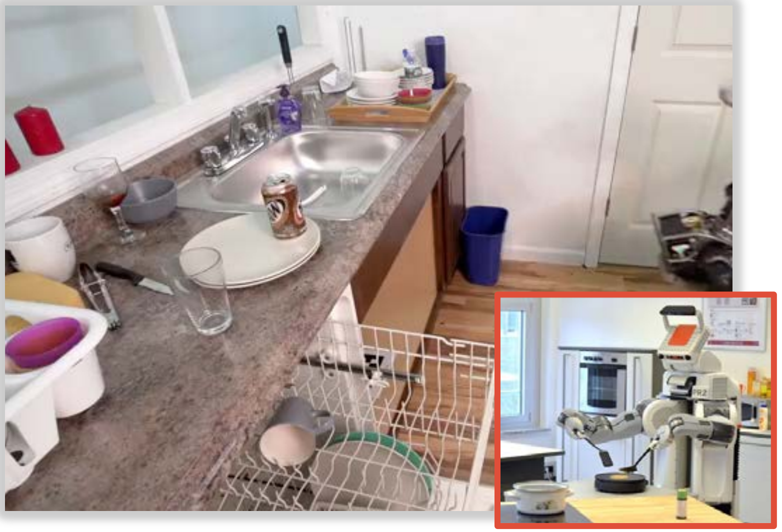

Personal/Service

Healthcare/Medicine

Generalizable Autonomy in Robot Manipulation

Vision: Build Intelligent Robotic Companions

Approach: Learning with Structured Inductive Bias and Priors

Demonstration

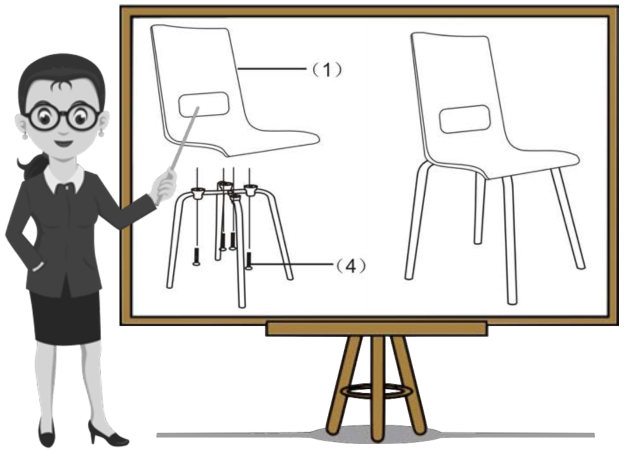

Instructional Input
(Teleoperation, Video, Language)

Task Imitation

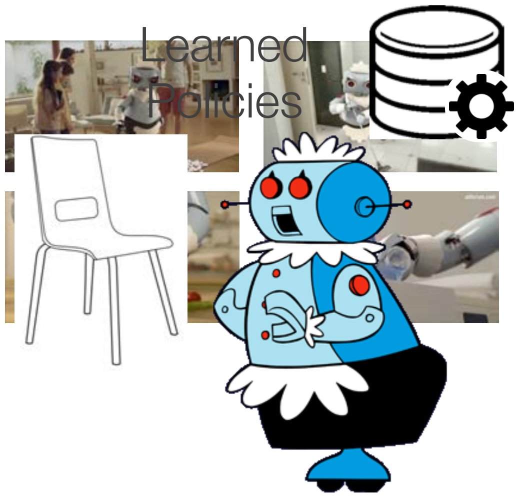

Learn to do the task in
Same Environment

Generalization

New Task Variations
in Novel Environments

Layers of Imitation

Movement Skills | Skill Sequencing | Semantic Purpose | Task Specification

Control | **Planning** | **Perception**

Cheng et al. Sci.Rob. 2019

Generalizable Autonomy in Robot Manipulation

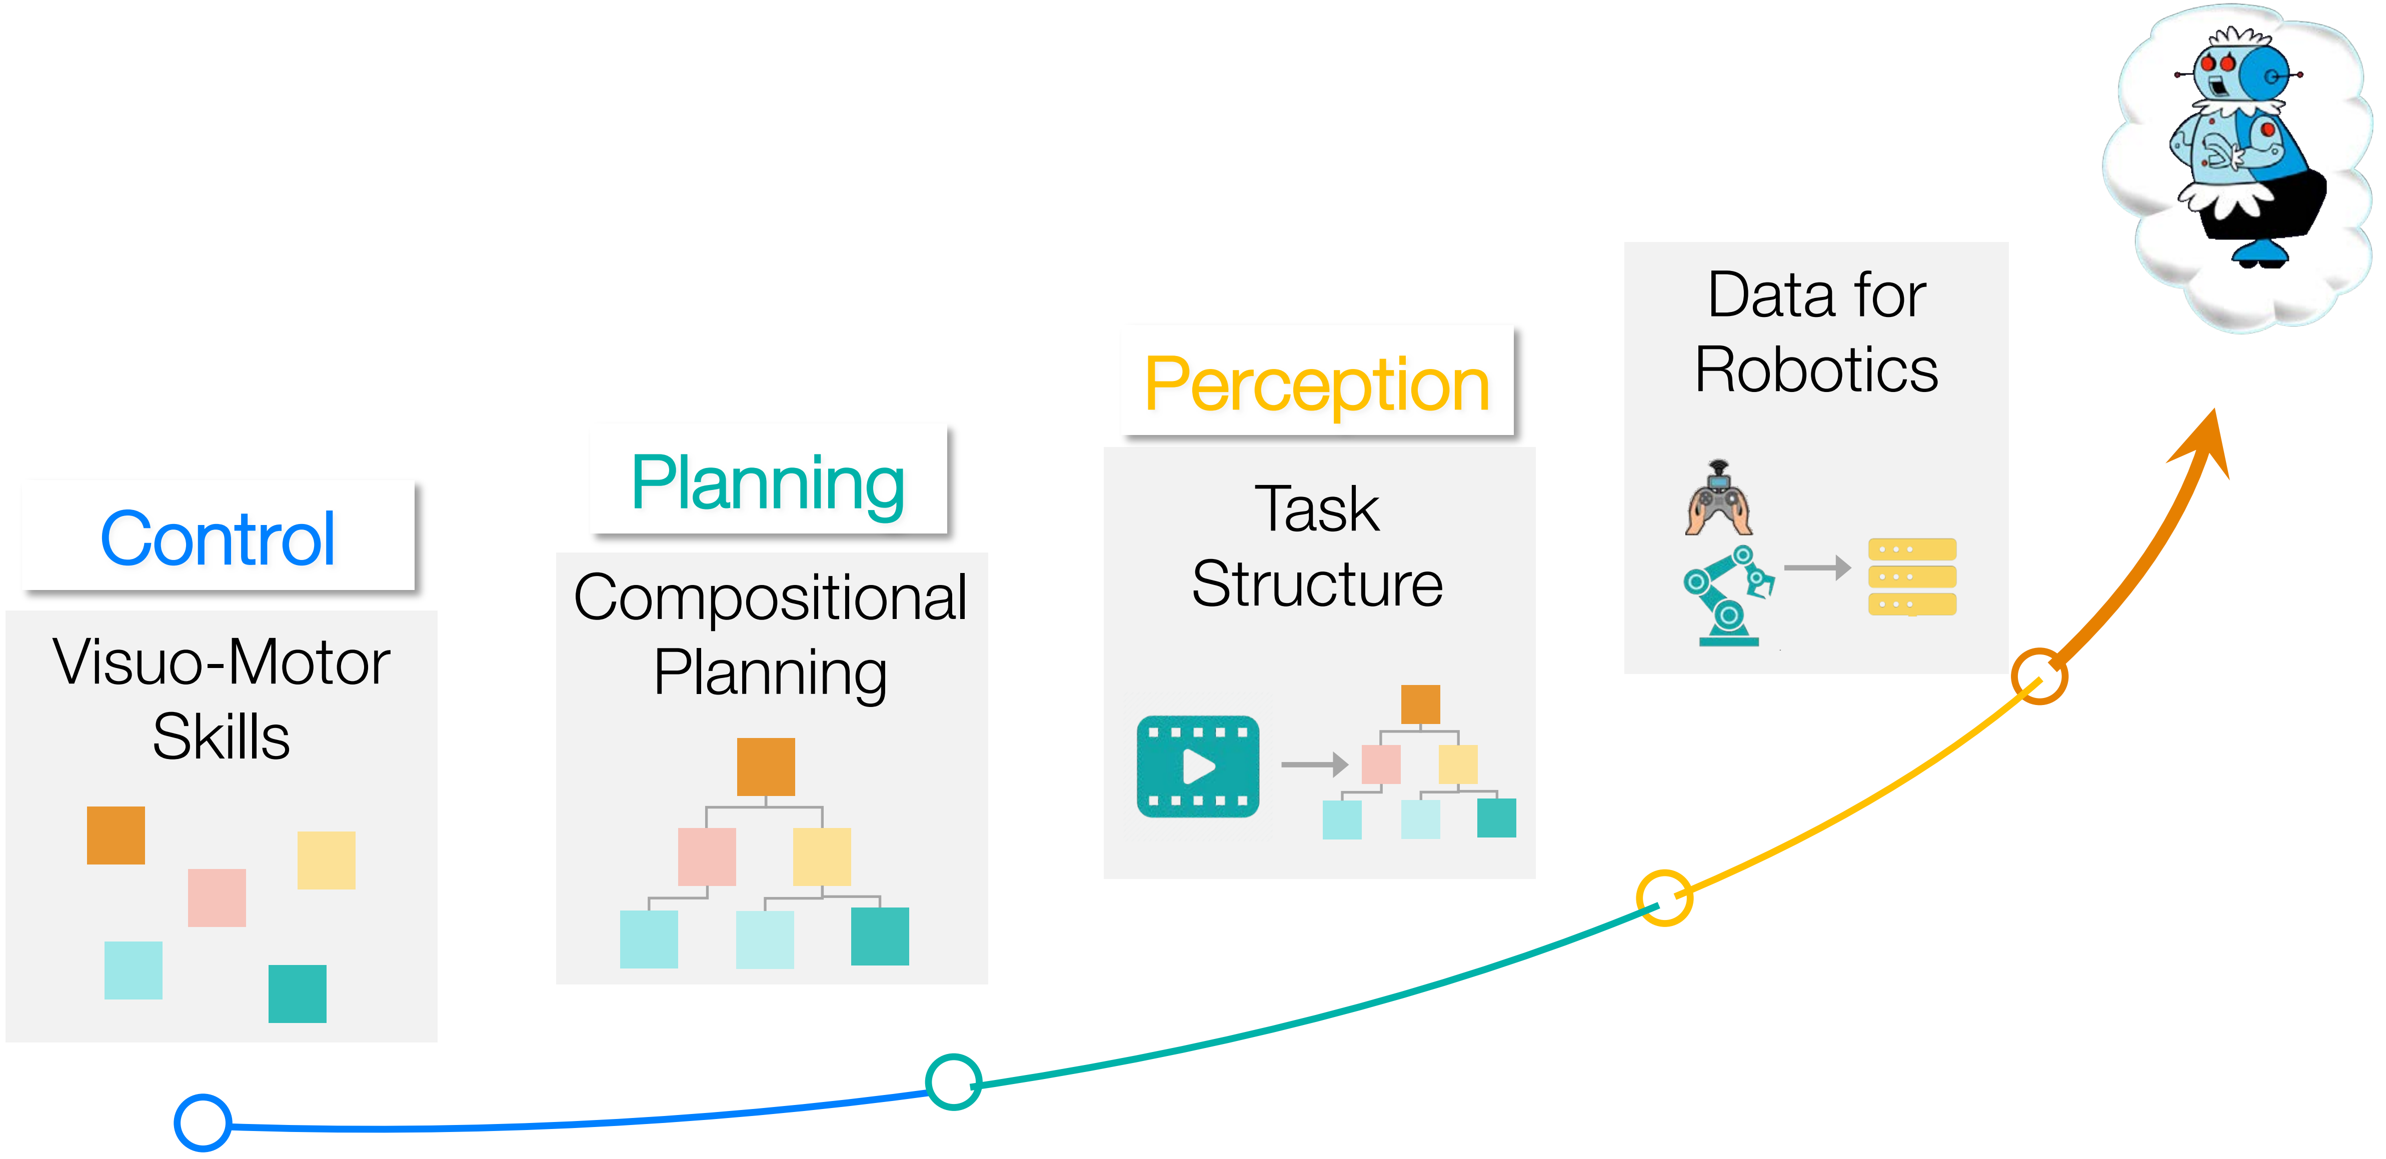

Generalizable Autonomy in Robot Manipulation

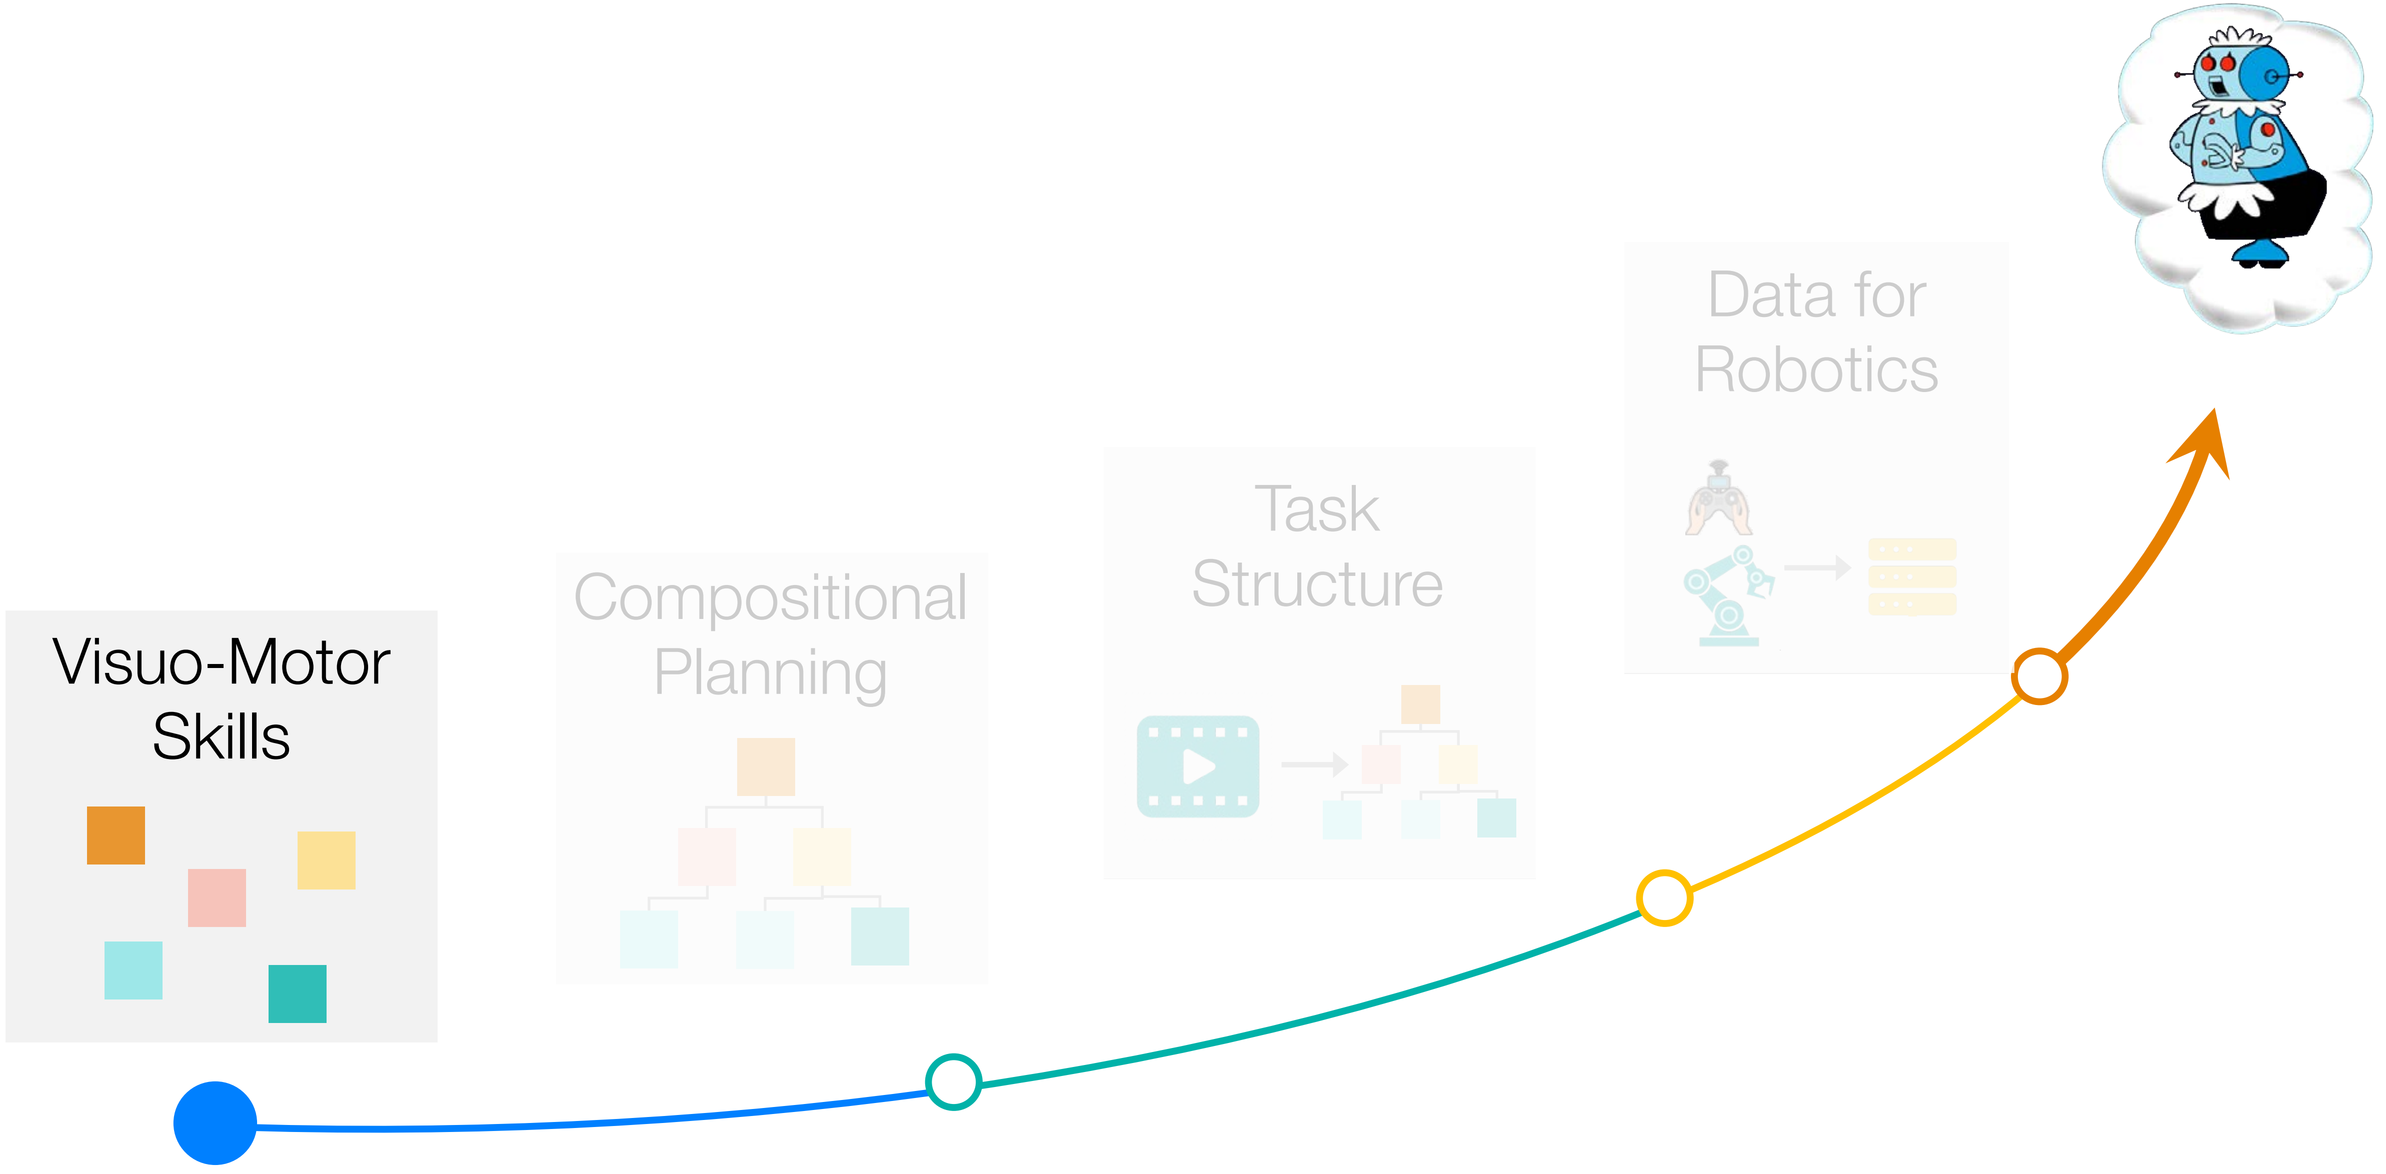

Visuo-Motor Skills

Challenge: Algorithmic frameworks to learn a diversity of skills

Approach: Close the Visuo-Motor Loop with Learning based Control

• • •

Vacuuming

Sweeping/Mopping

Cooking

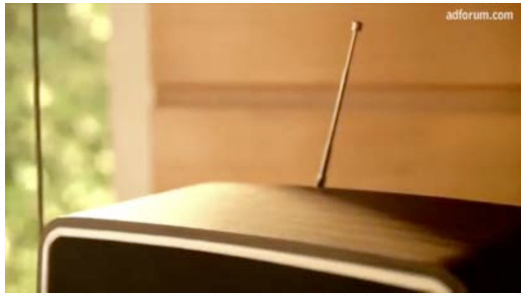

Cleaning

• • •

Visuo-Motor Skills: Generalization

Cleaning

Skills: Surface Wiping

Hard Stains – Push Harder?

Different Surfaces – Be Gentle?

Generalization

Visuo-Motor Skills: Current Paradigm

Model Based Task (Operational) Space Control

Actual State: Image, Force, Joint Enc.

Desired State: x_d

Robot Model Parameters: M, J

Action: τ

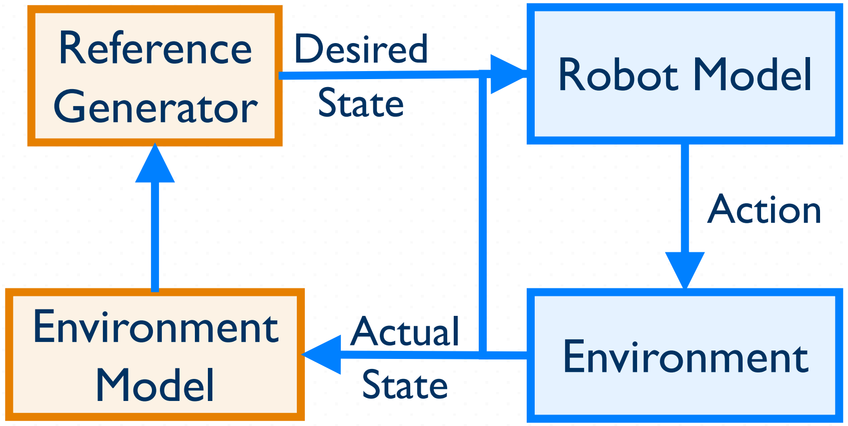

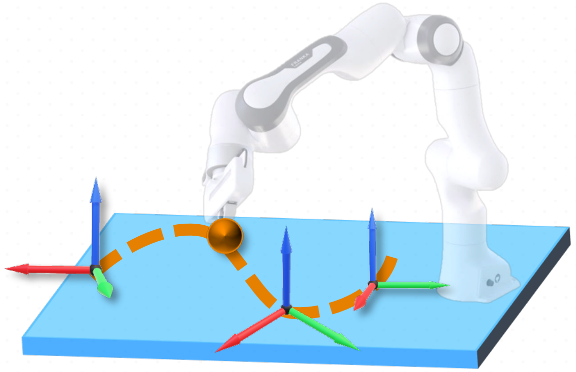

Robot Model

$$\ddot{x}_{ref} = K_p(x_d - x) + K_v(\dot{x}_d - \dot{x}) + \ddot{x}_d$$

$$M(q,\dot{q}) + C(q,\dot{q}) + G(q) + \varepsilon(q,\dot{q}) = \tau$$

$$\tau = J^T\,(JM^{-1}J^T)^{-1}(\ddot{x}_{ref} - \dot{J}\dot{q} + JM^{-1}F)$$

+ Leveraged Robot Model
+ Compliant Control

Environment Model + Reference Generator

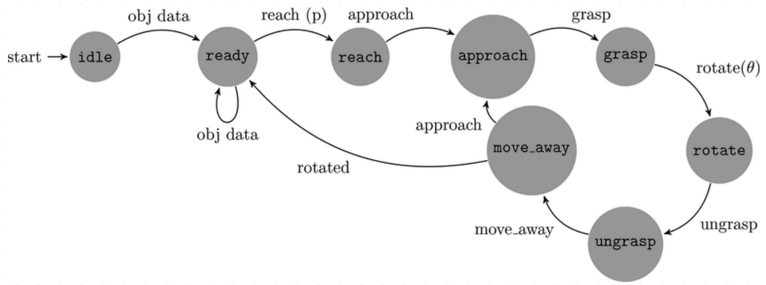

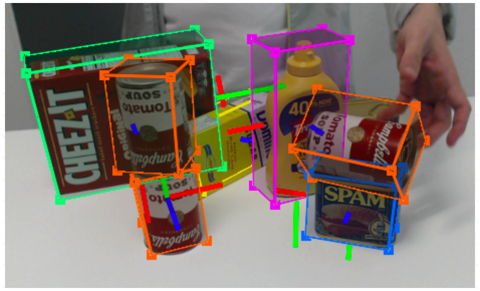

- Needs Environment (Task) Model

- Task Dependent State

- Explicit State Estimation

[Khatib.1987], [Bruyninckx et al. , '96], [Schaal '03], [Rodriguez et al. 2012], [Vijaykumar et al. '11],[Ijspreet et al. '13]][Li, Billard, et al. '14], [Lee, Abbeel et al. '15]

Visuo-Motor Skills: Current Paradigm

Deep Reinforcement Learning

State: Image

Policy π

Action: Torque

τ

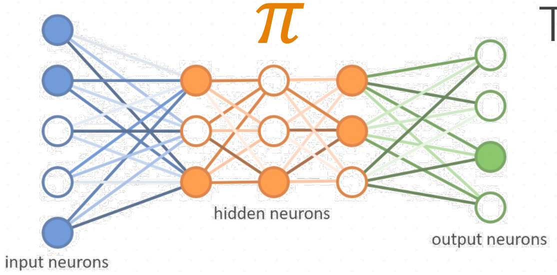

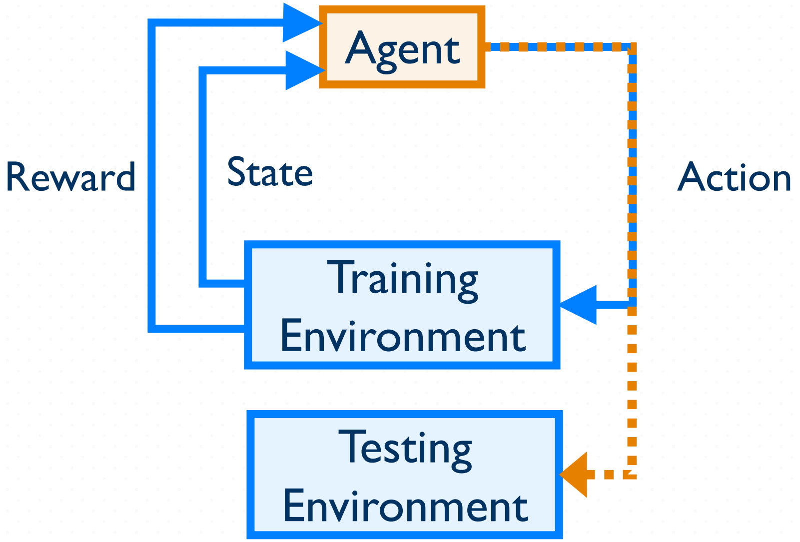

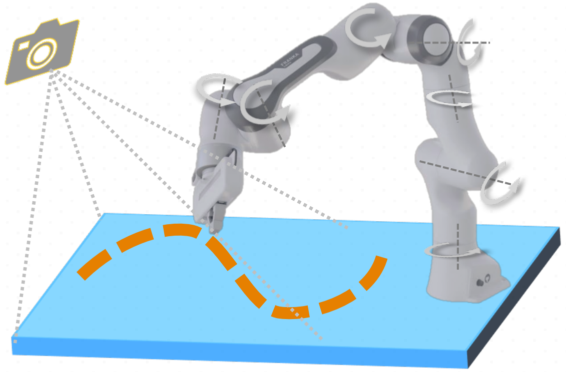

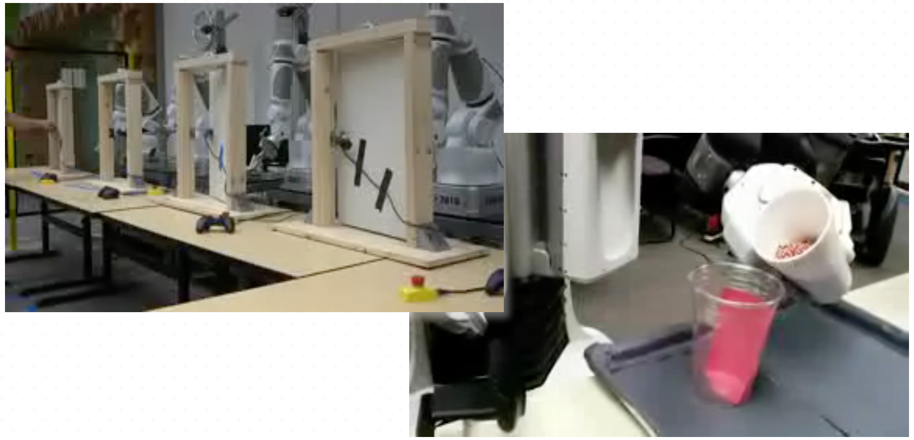

+ Model Free: No Environment Model
+ State is Image

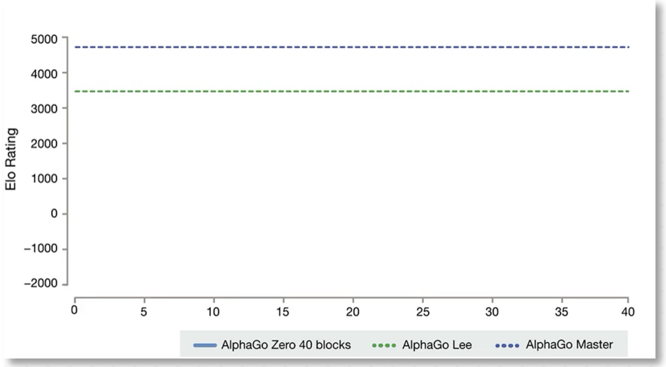

- Sample Inefficient
- Learn robot model (implicitly)

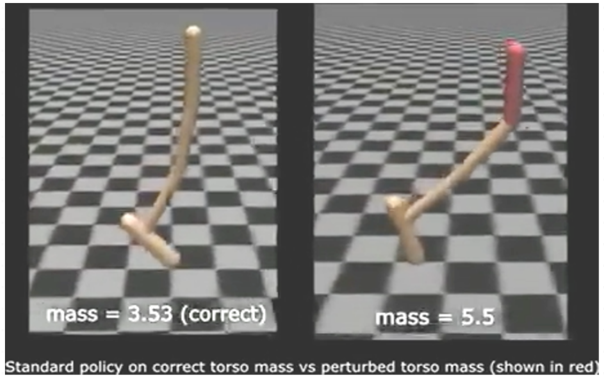

- If Training $\not\approx$ Testing: Policy Fails!

[Agrawal et al., '16], [Levine et al., '16], [Peng et al., '17], [Gu et al., '16], [Chebotar et al., '17], [Yahya et al., '16], [James et al., '17], [Popov et al., '17]

Visuo-Motor Skills: Our Approach

RL with Variable Impedance Task-Space

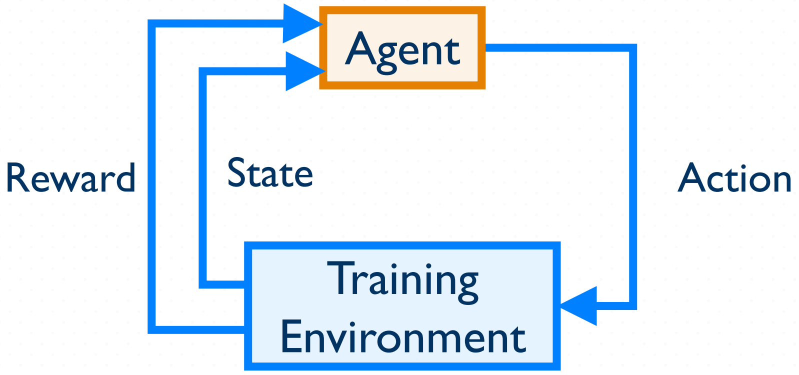

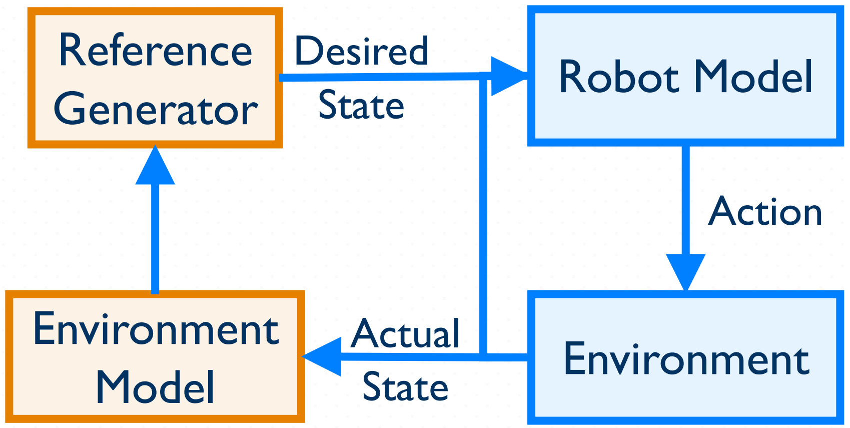

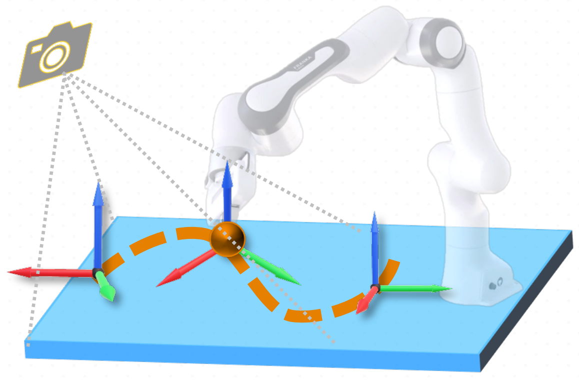

Visuo-Motor Skills: Our Approach

RL with Variable Impedance Task-Space

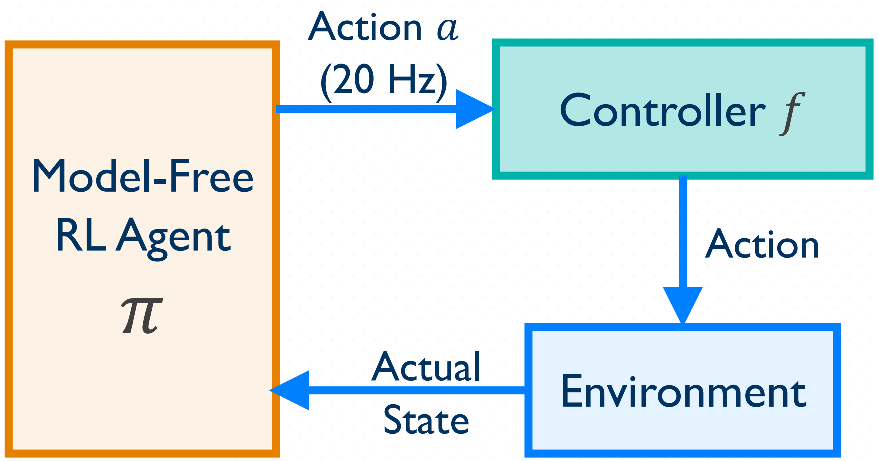

Reference Generator (learned)

$$\tau = f(\pi(o_t))$$

Robot Model (Deterministic)

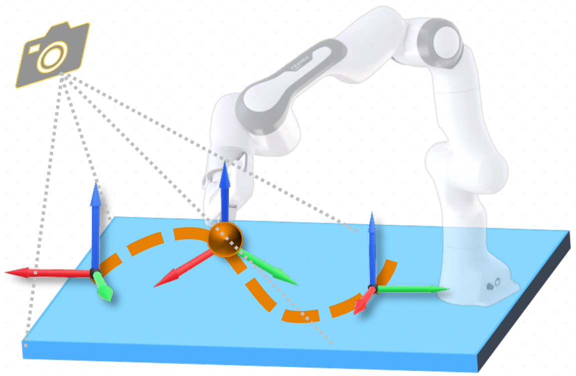

Visuo-Motor Skills: Our Approach

RL with Variable Impedance Task-Space

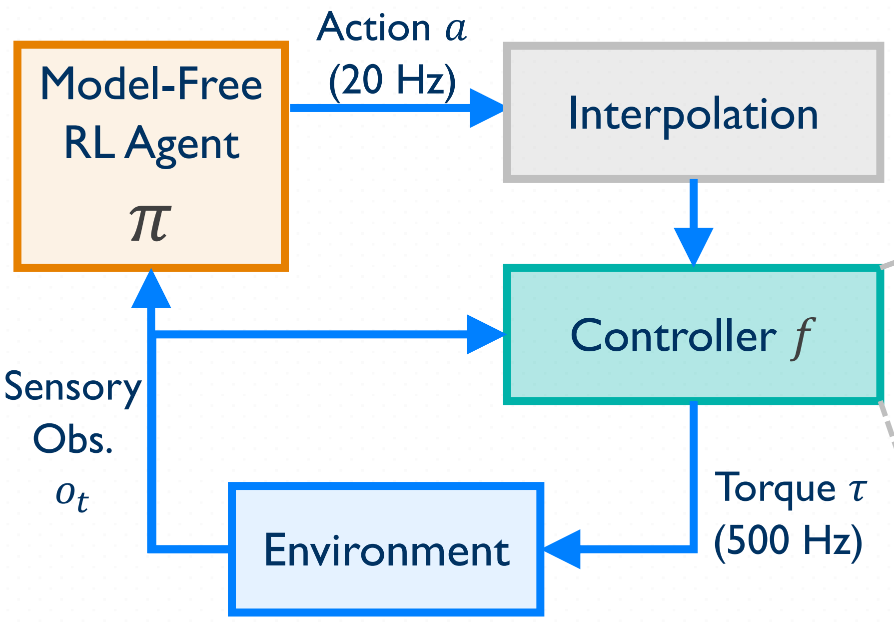

$$\tau = f(\pi(o_t))$$

$$\pi(o_t) = a: [\underbrace{\boxed{x_d, \dot{x}_d}}_{\text{Pose and Velocity}}, \underbrace{\boxed{K_p, K_v}}_{\text{Impedance Gains}}]$$

$$\tau = \underbrace{\boxed{f}}_{\text{Deterministic Position-Velocity Control Jacobian } J \text{ and Inertia } M}(x_d, \dot{x}_d, K_p, K_v)$$

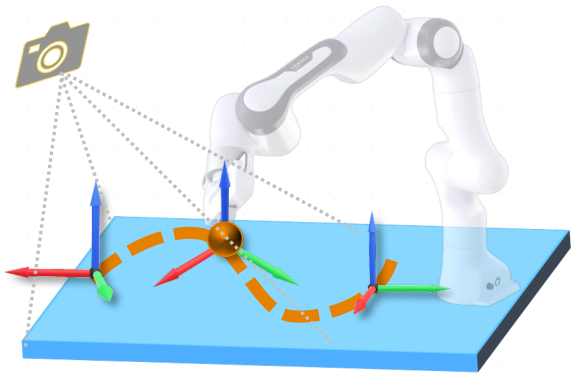

+ Model Free: No Environment Model
+ State is Image

+ Leverages Robot Model
+ Compliant Control

+ Sample Efficient
+ Transferable

Visuo-Motor Skills: Action Representation

Surface Wiping

Input: Image (48x48)

Minimize the number of Dirty Tiles

Maintain Contact with the Table

Don't push with more than Robot Payload

Reward: $\lambda_1 \sum(dirt\_on\_table) + \lambda_2(distance\_to\_table) - \lambda_3 \mathbb{I}\,(F \geq 40N)$

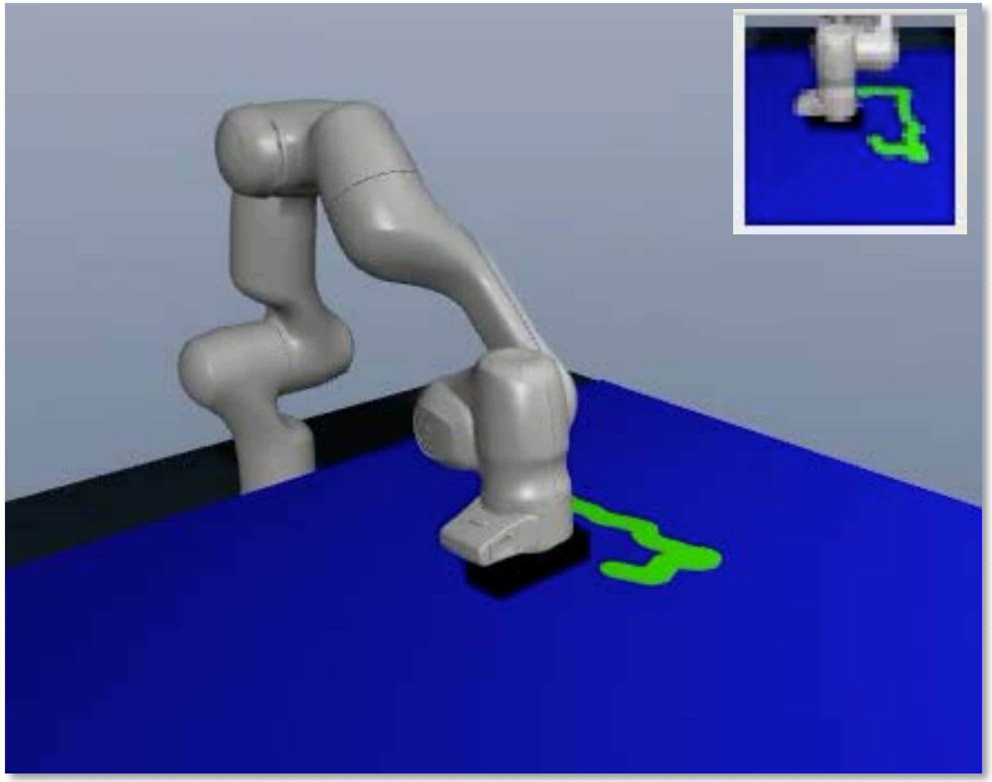

Trained Policy Rollout (Ours)

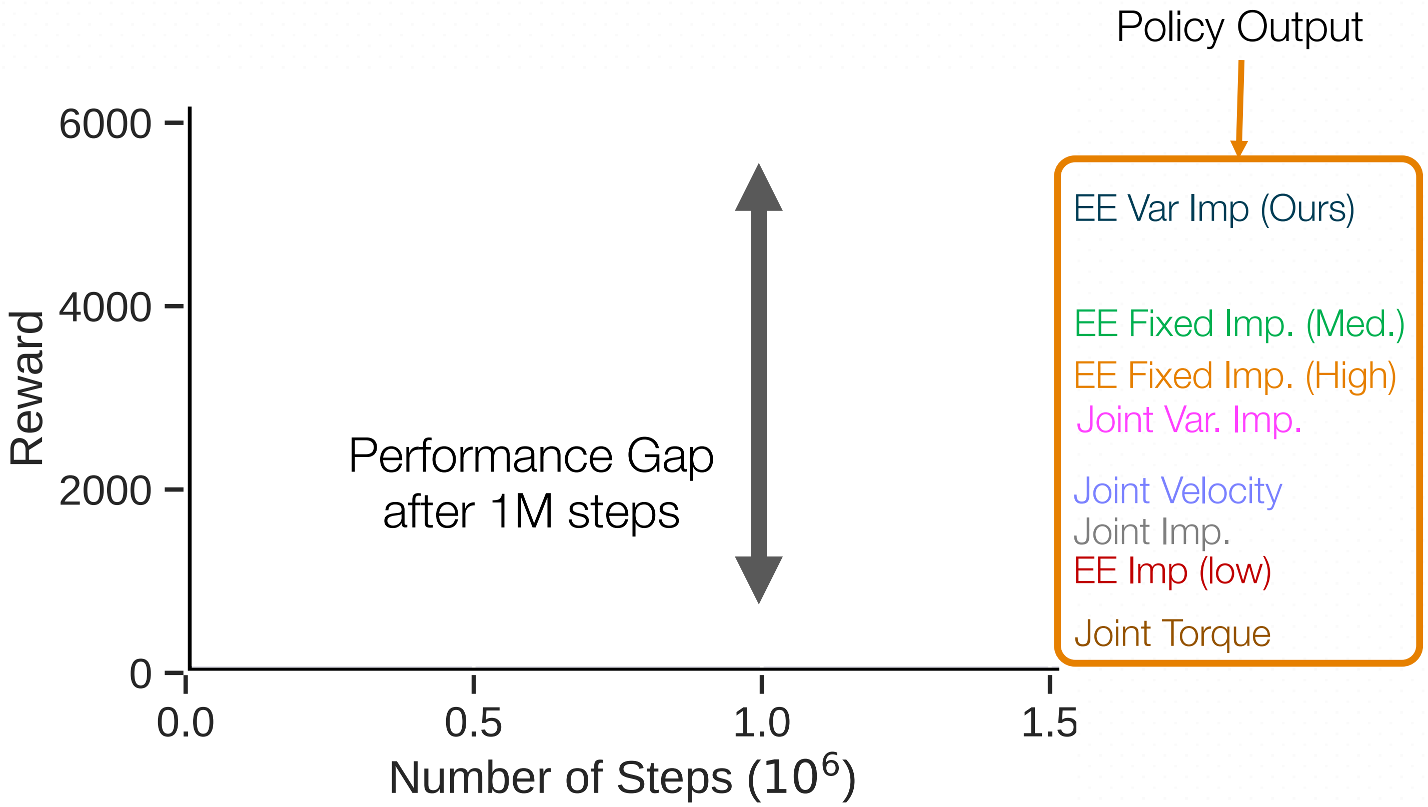

Evaluation During Training (PPO)

Visuo-Motor Skills: Action Representation

$$\tau = f_{Sim}(\pi(o_t))$$

$$\tau = f_{Real}(\pi(o_t))$$

Success 80% (10 Trials)

Generalizable Autonomy in Robot Manipulation

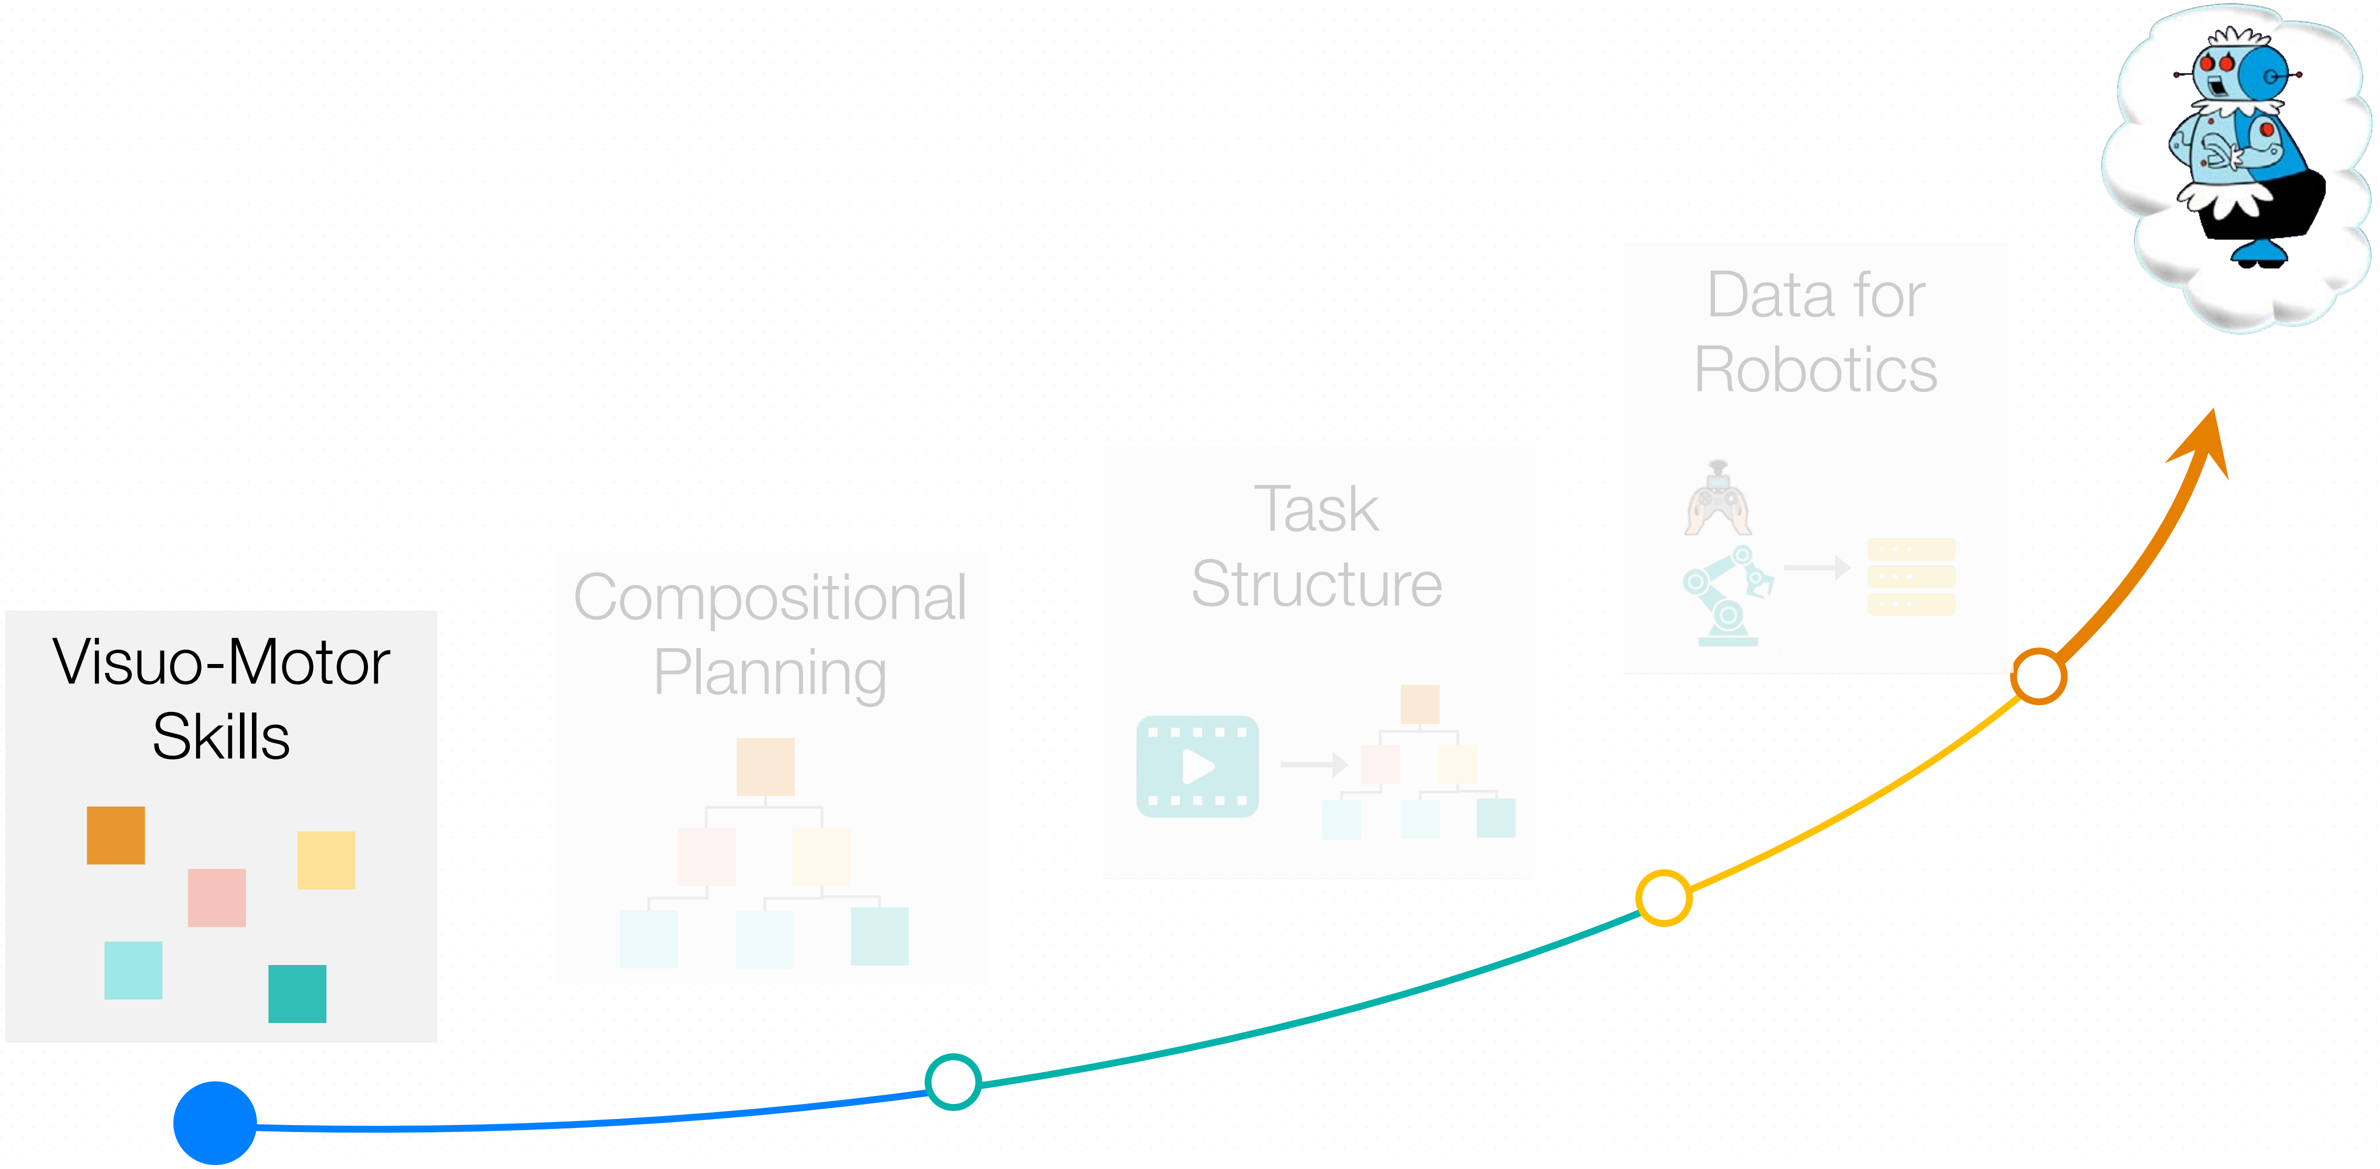

Skills: Imitation from Heuristics

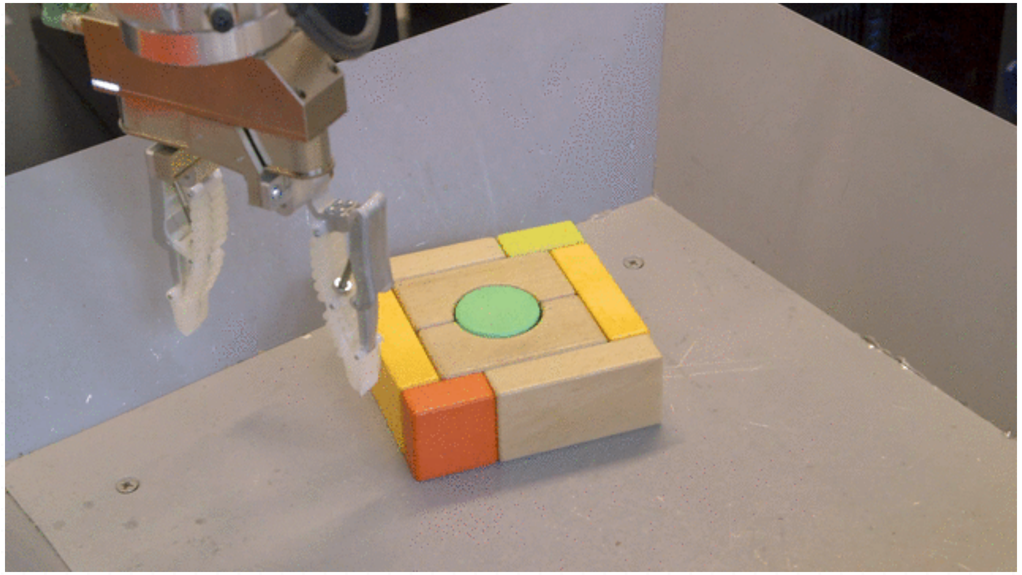

Promise of Deep RL
closed loop-control with images

…albeit, with a lot of training

[Kalashnikov et al (2018). Levine et al. (2016), Pinto et al. (2016), Kalashnikov et al. (2018), Yu et al. (2016), Haarnoja et al. (2018), Lee et al. (2019), Vecerik et al. (2017)]

Skills: Heuristics often beat RL

RL struggles with structured, multi-step skills

Even simple heuristics beat RL

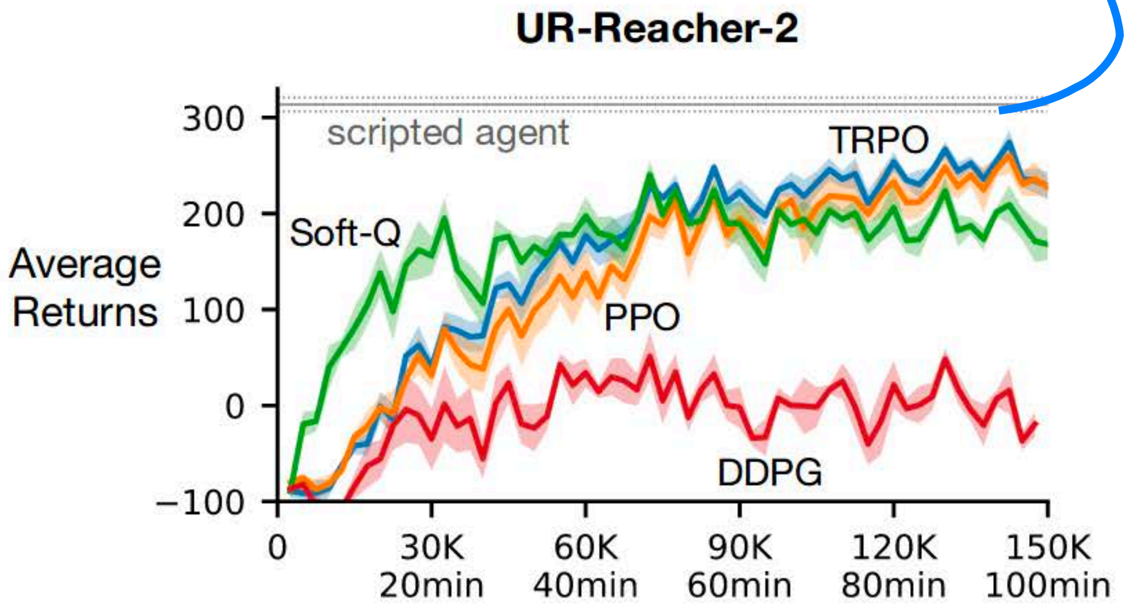

IROS 2019, Mahmood et al. 2018

Skills: Exploration without Guidance

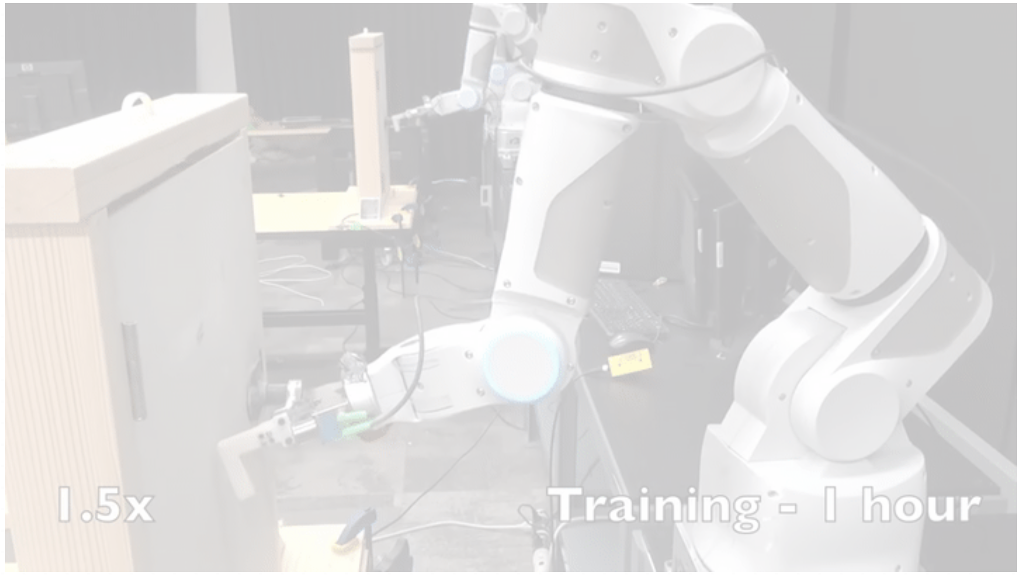

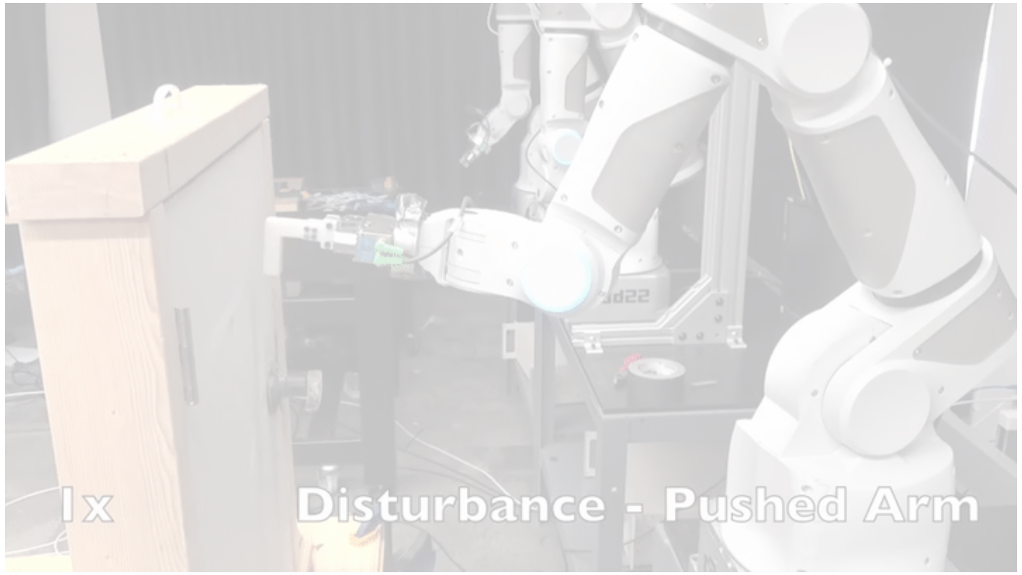

Random Exploration is slow

…even when first steps are obvious

Can Human Intuition Guide Exploration?

Gu et al. 2017

Skills: Imitation from Heuristics

Human
RL Task

Teachers

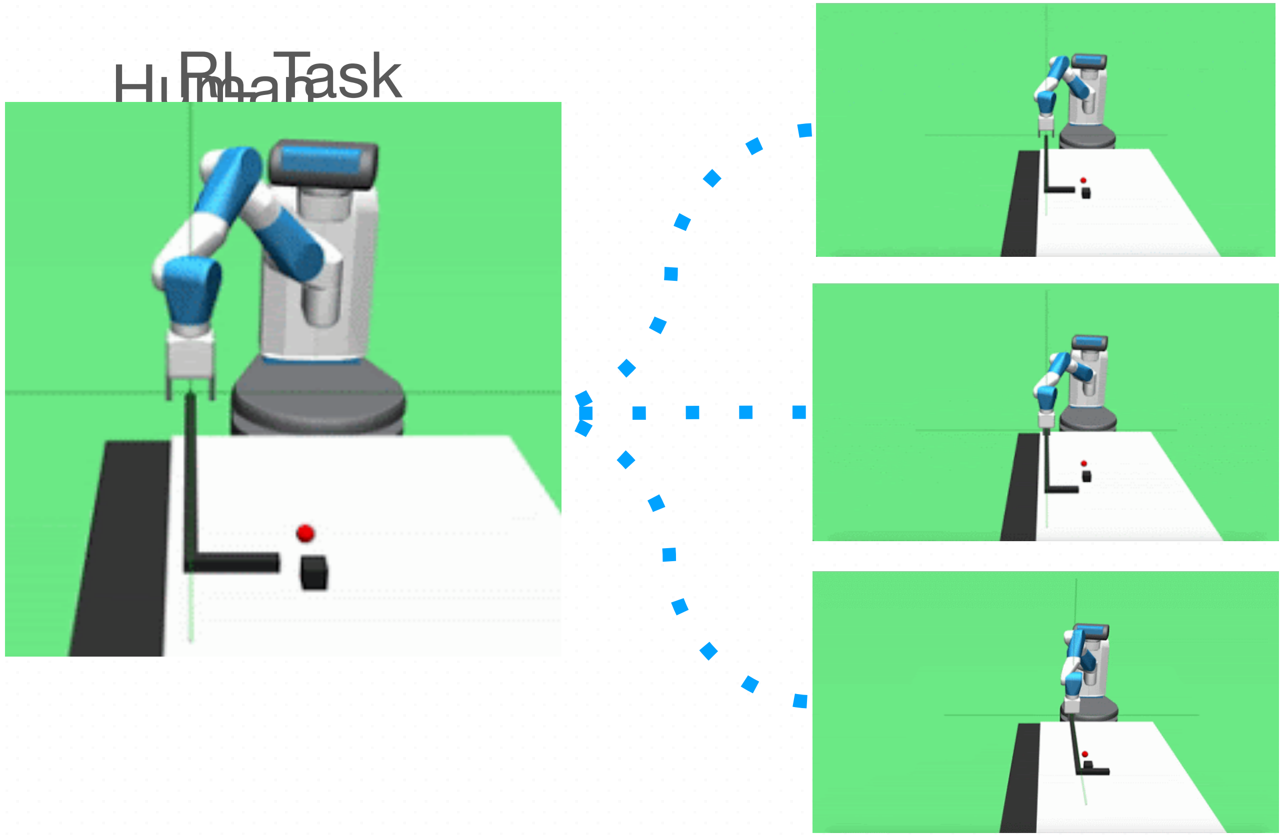

Intuition

Implement Useful Skills
...but not full solution

Teachers

Black-box controllers
solving parts of the task

Skills: Imitation from Heuristics

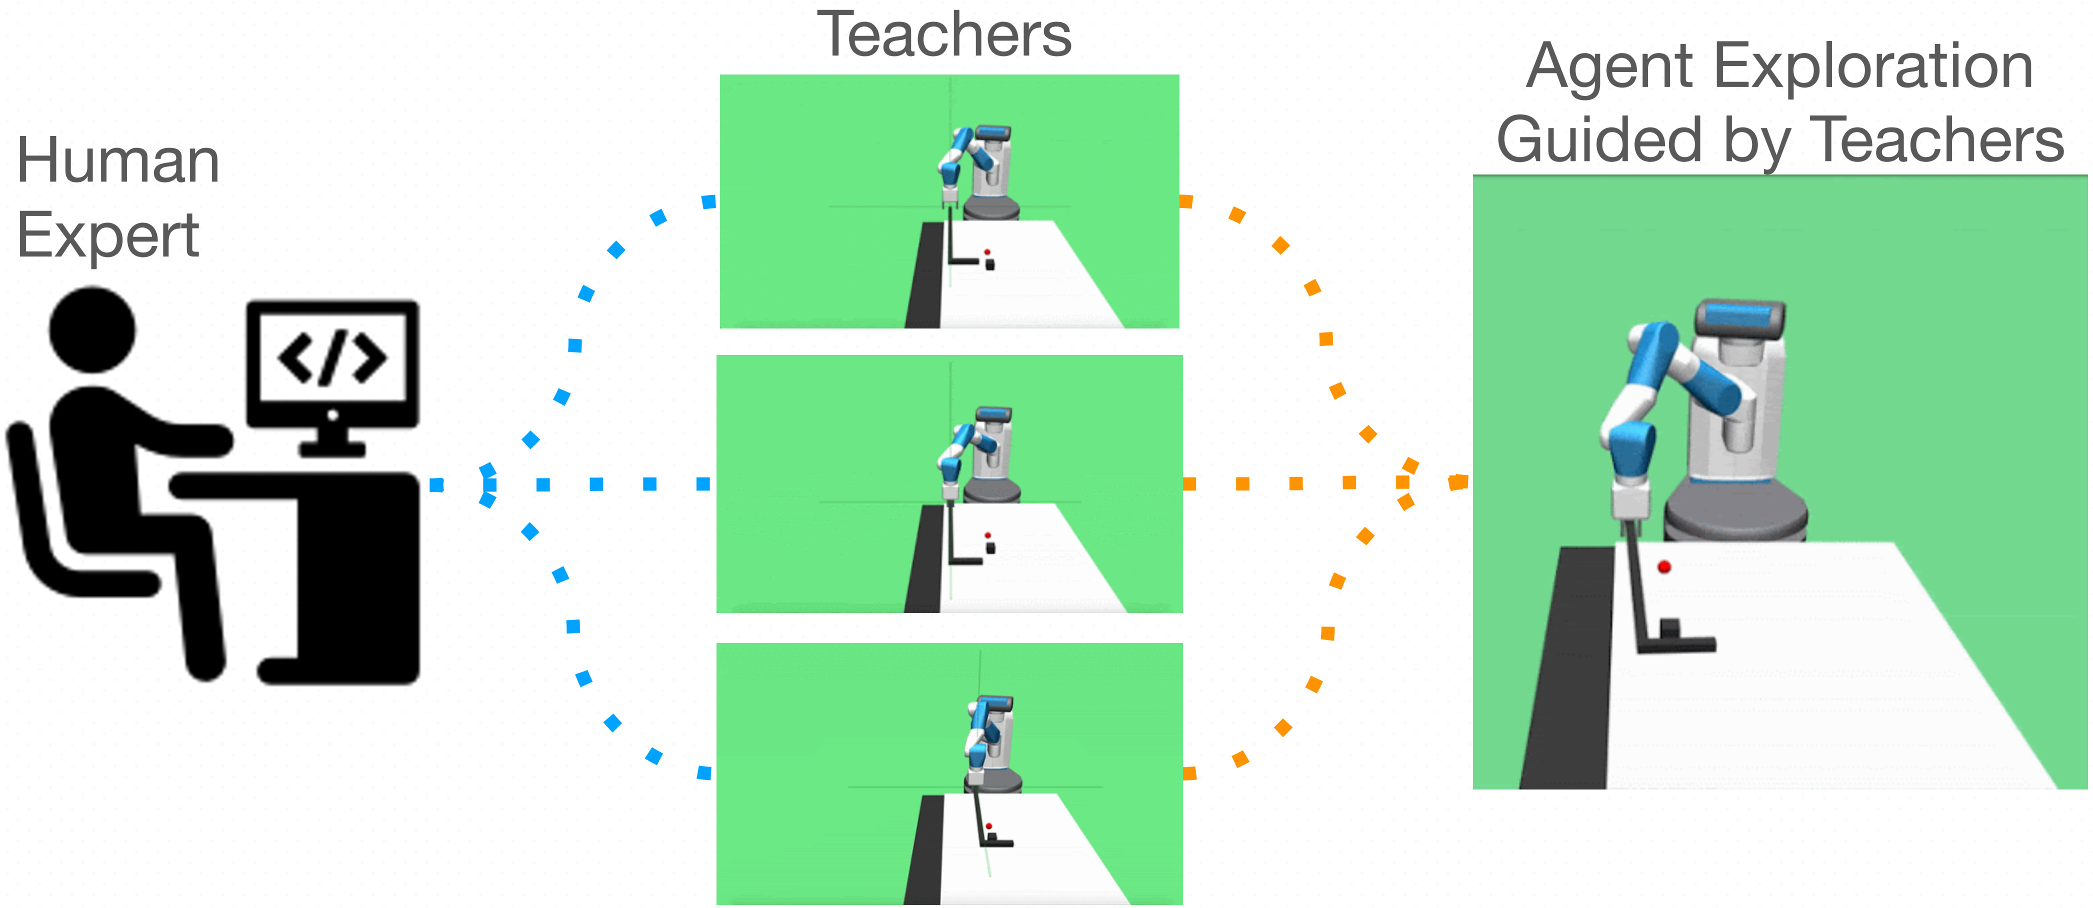

Goals: A) **faster** agent training B) **optimal** test-time agent performance

Skills: Imitation from Heuristics

Naive action choice might not work well!

Partial

Contradictory

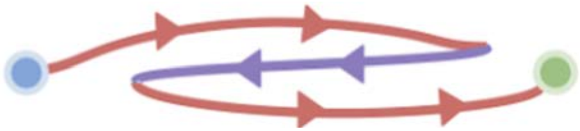

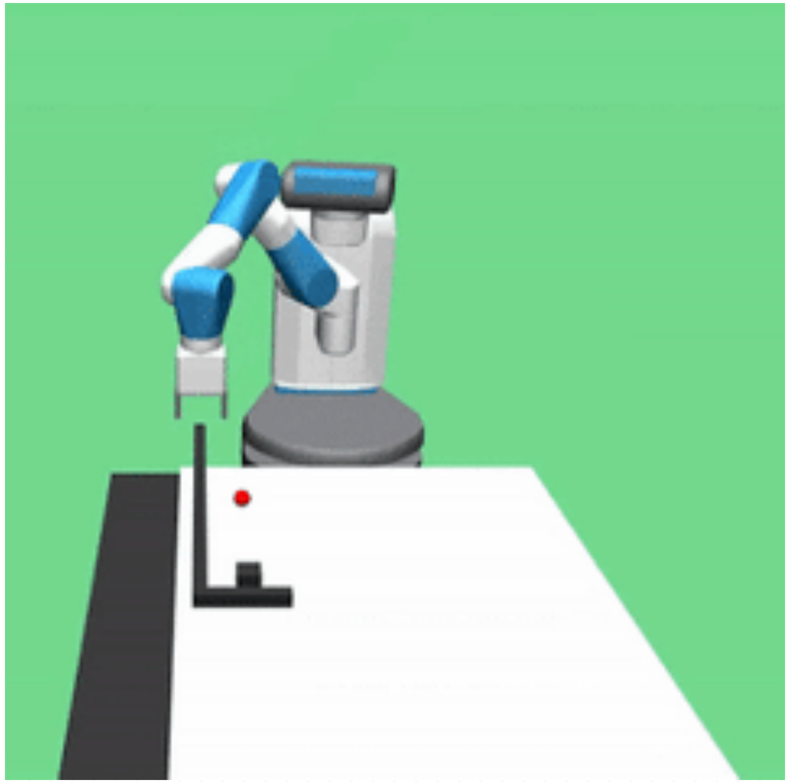

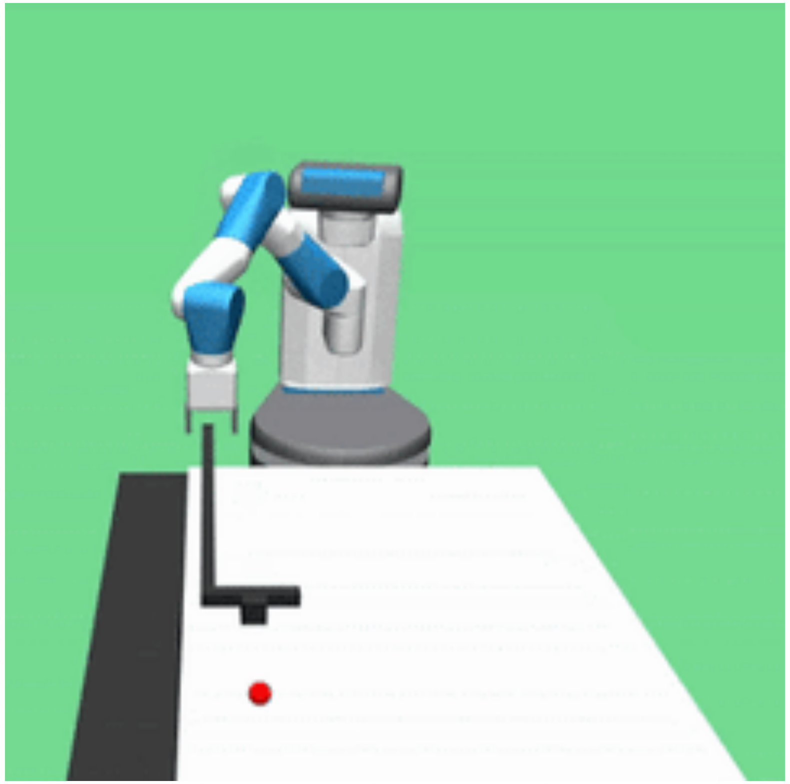

Off-Policy RL: DDPG Review

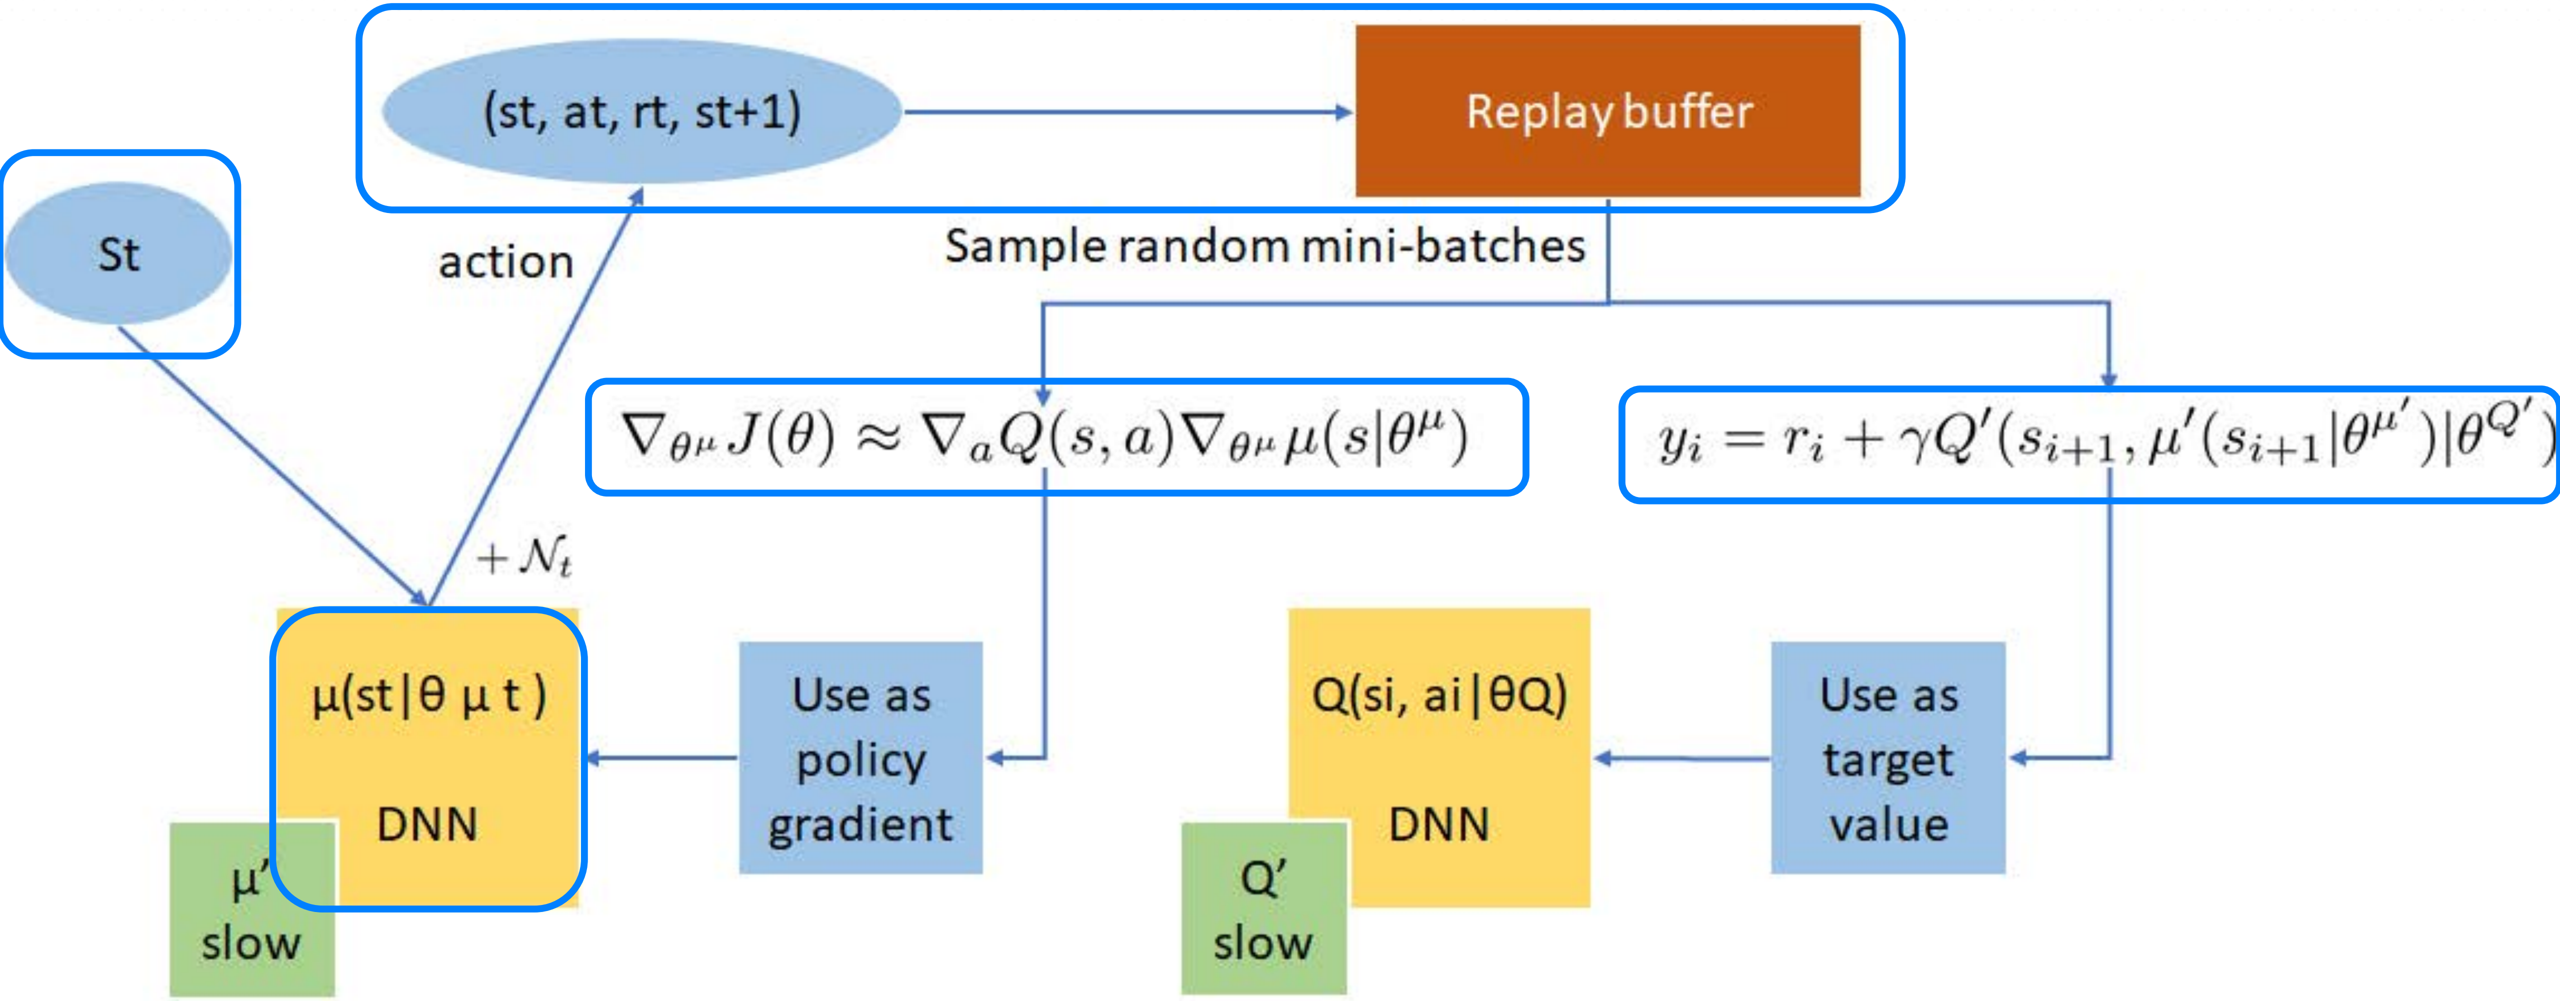

Lillicrap et al. 2015

AC-Teach: Actor-Critic with Teachers

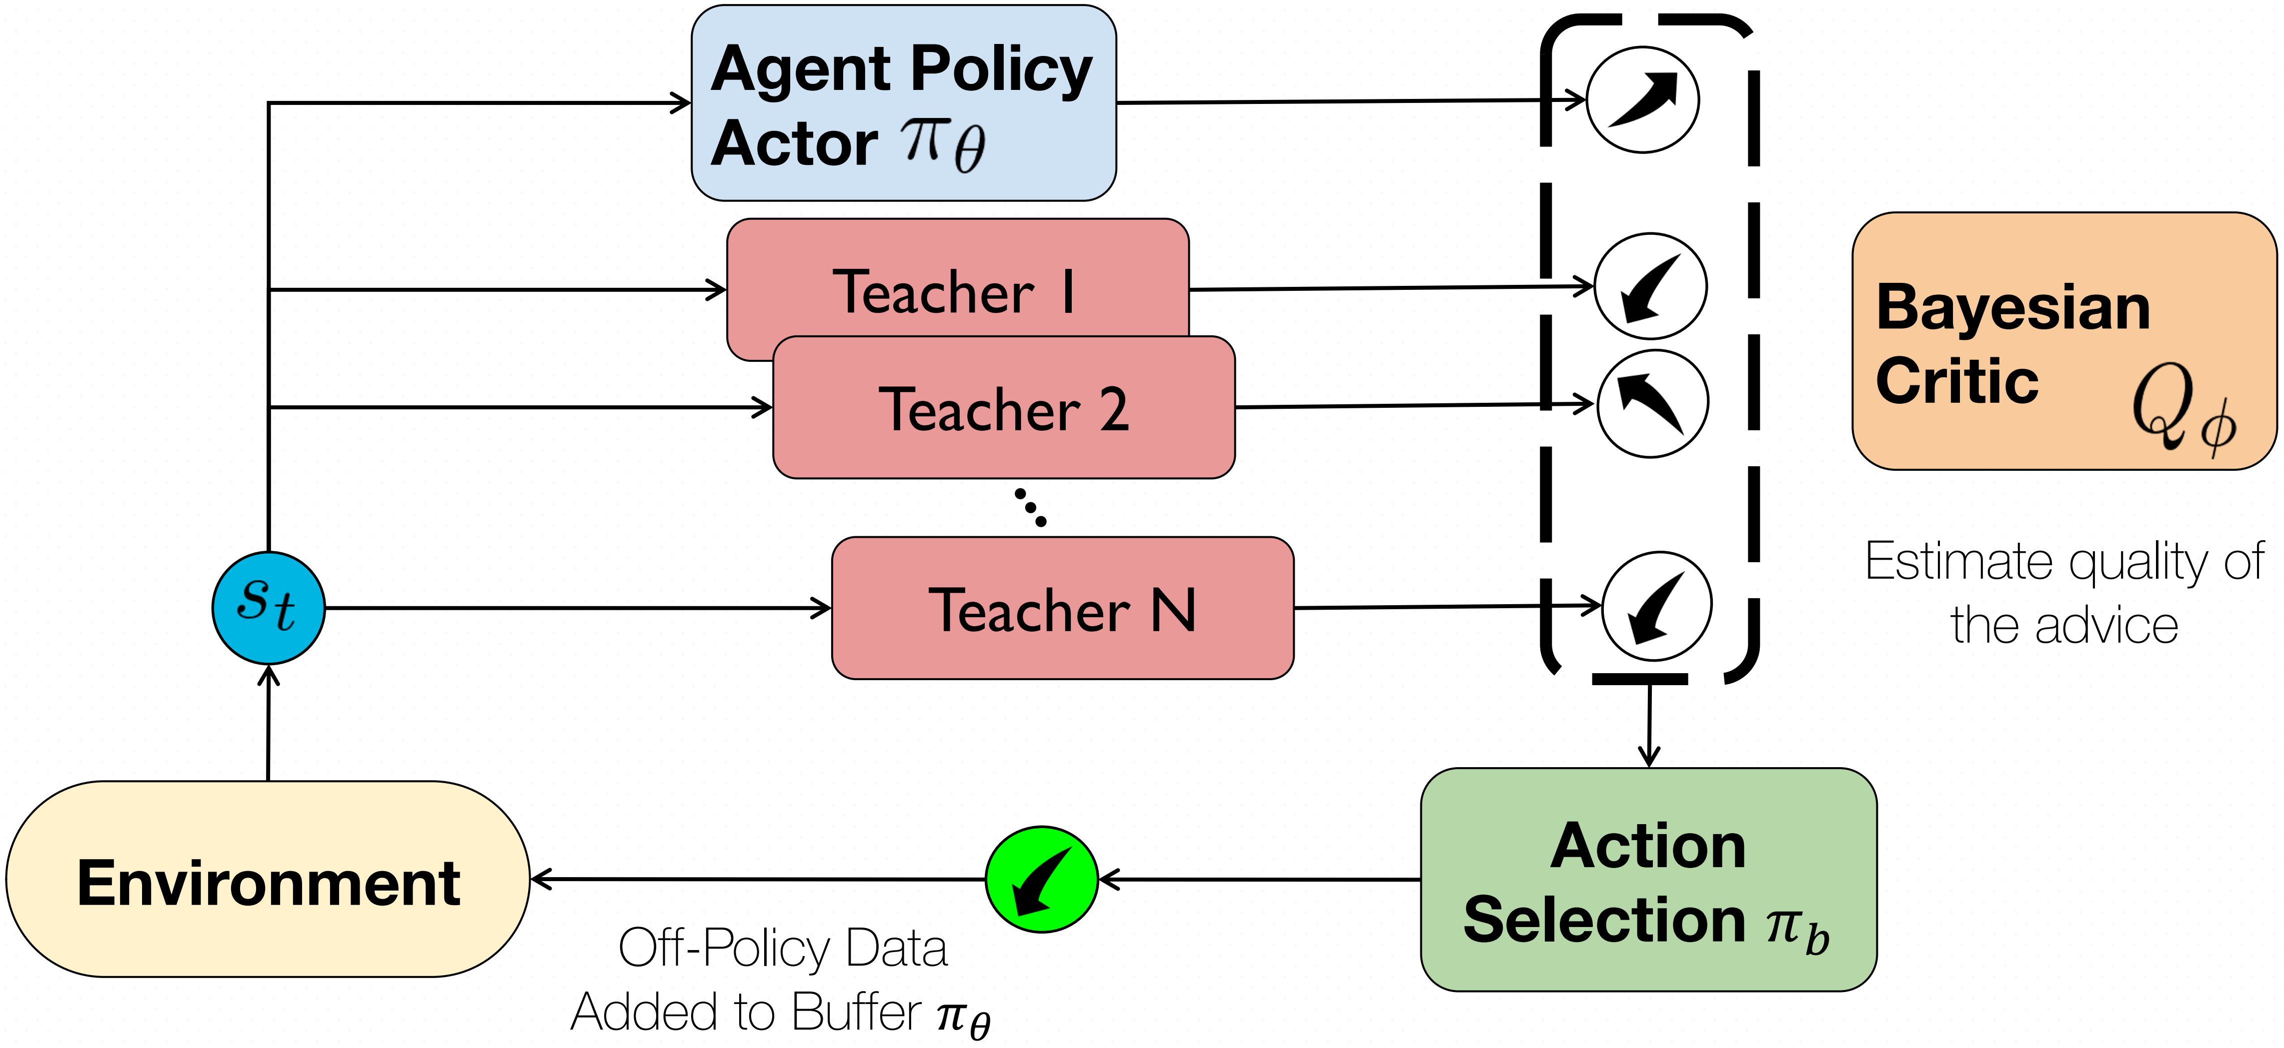

Experiments

Task:

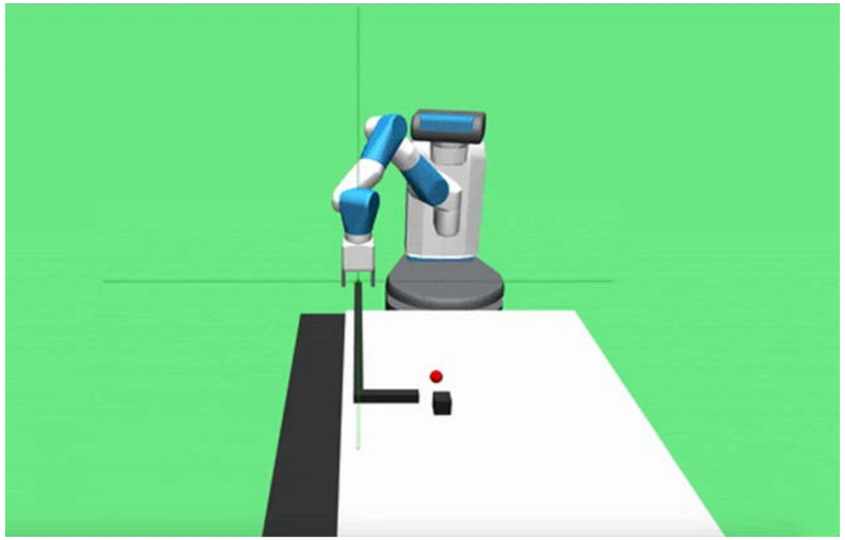

Teachers:

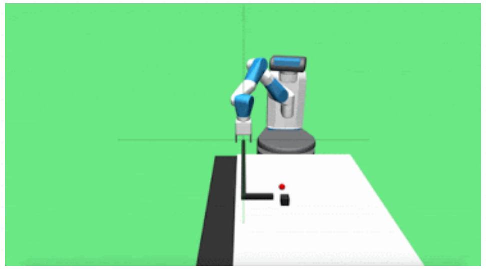

grab hook

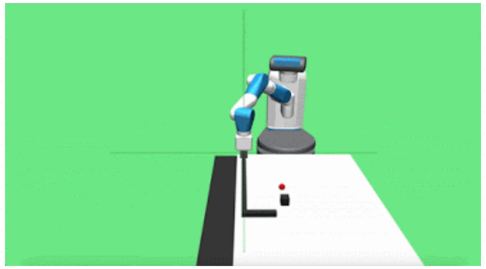

position hook

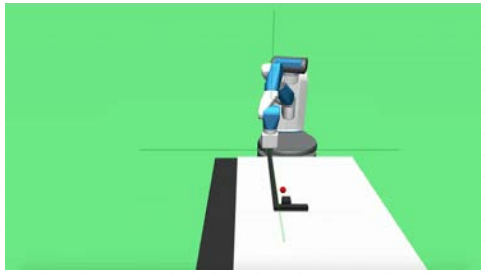

pull

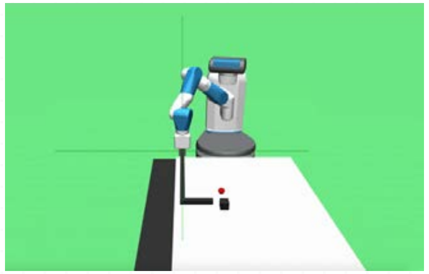

push

Results

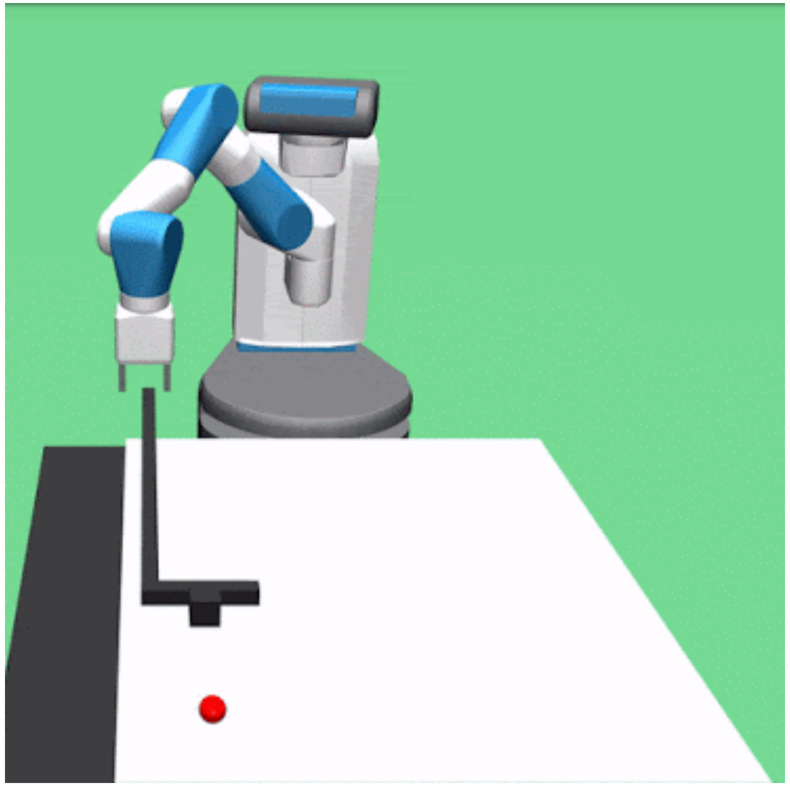

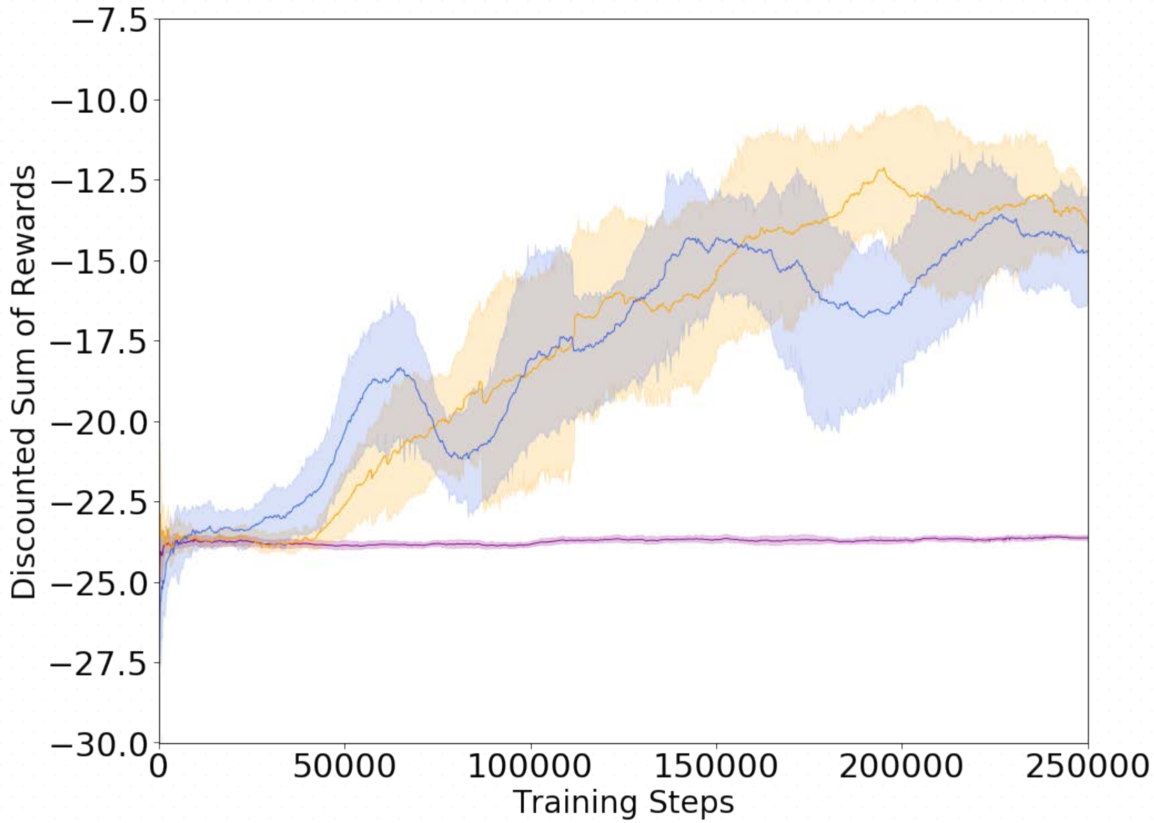

AC-Teach is able to leverage a single teacher well

Results

B-DDPG + AC-Teach (ours) B-DDPG + DQN B-DDPG (no teachers)

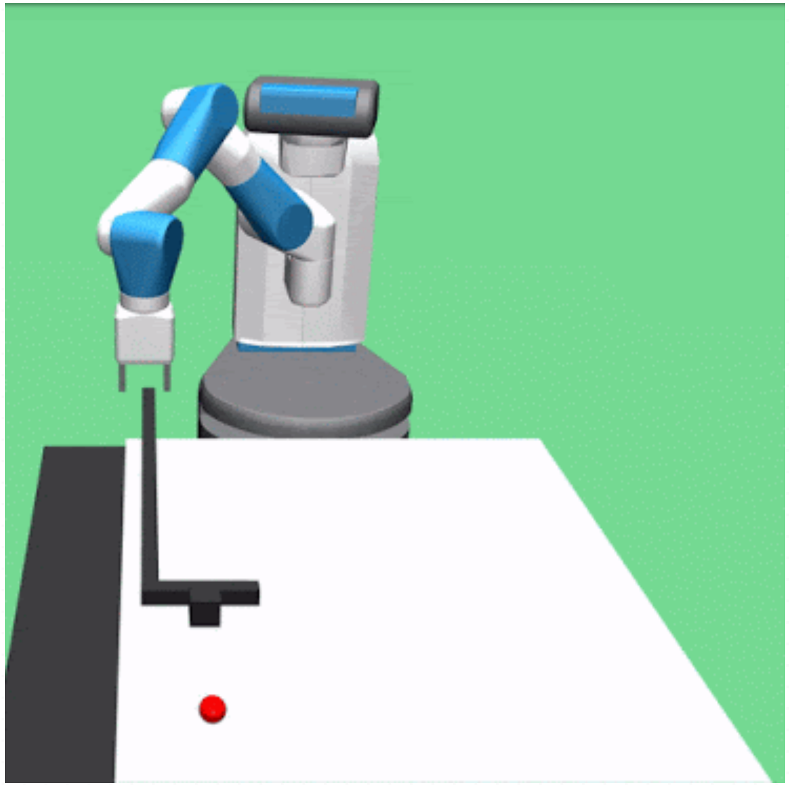

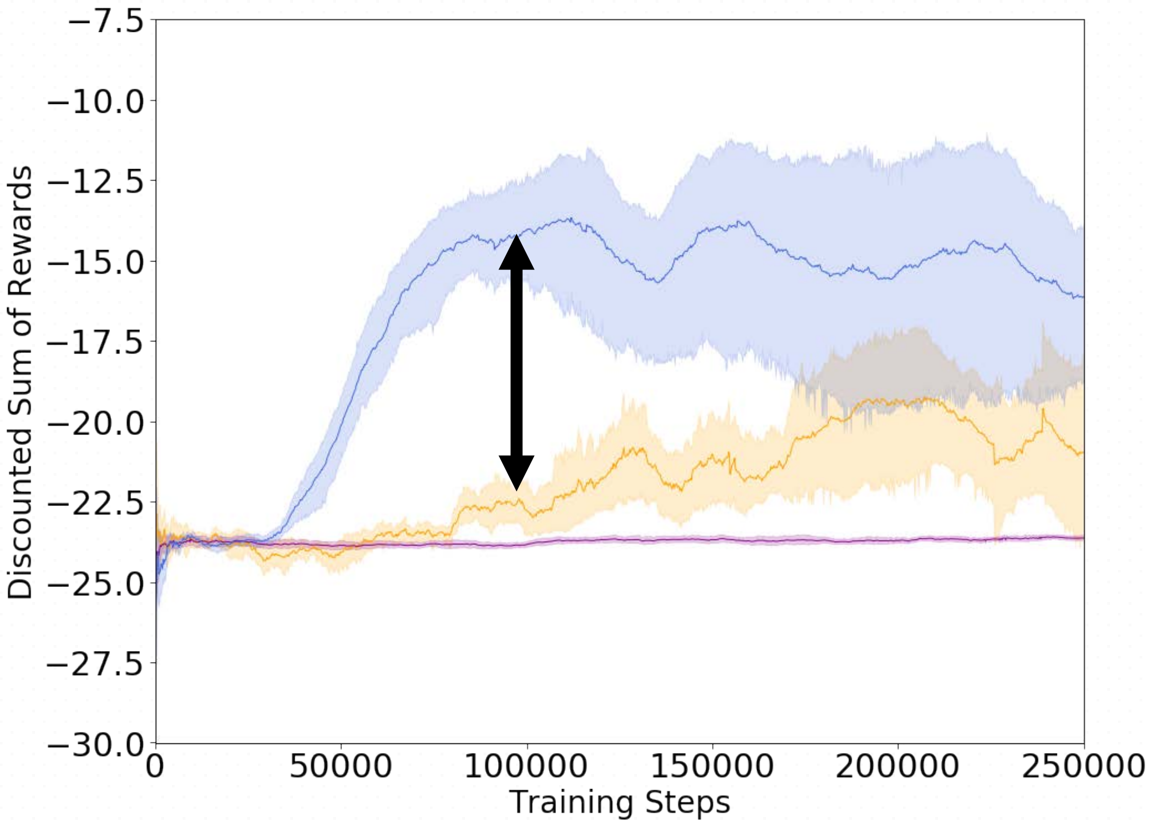

AC-Teach speeds up training given multiple teachers

Results

B-DDPG + AC-Teach (ours) — B-DDPG + DQN — B-DDPG (no teachers)

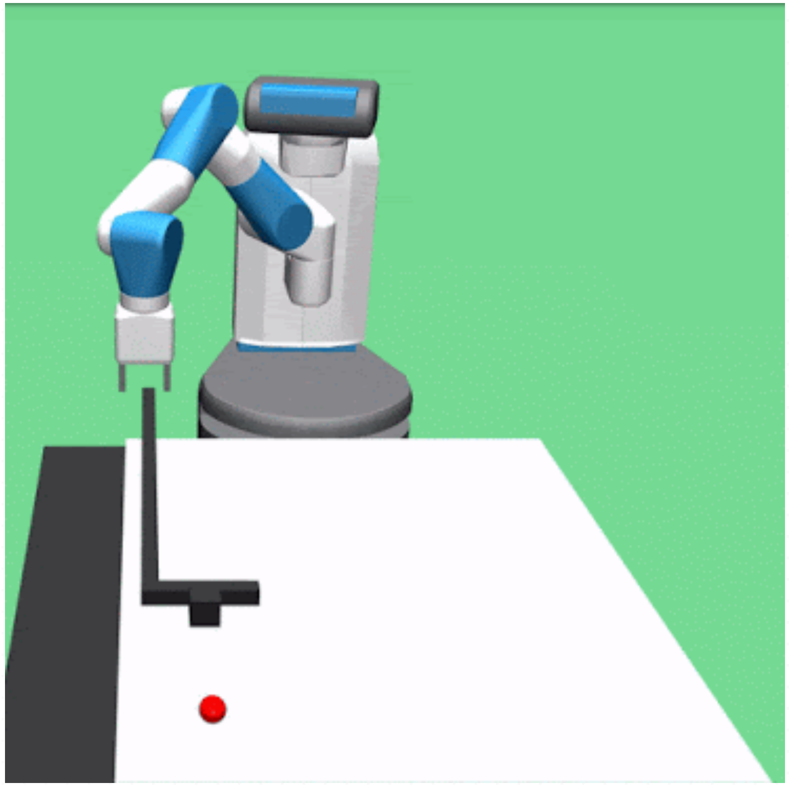

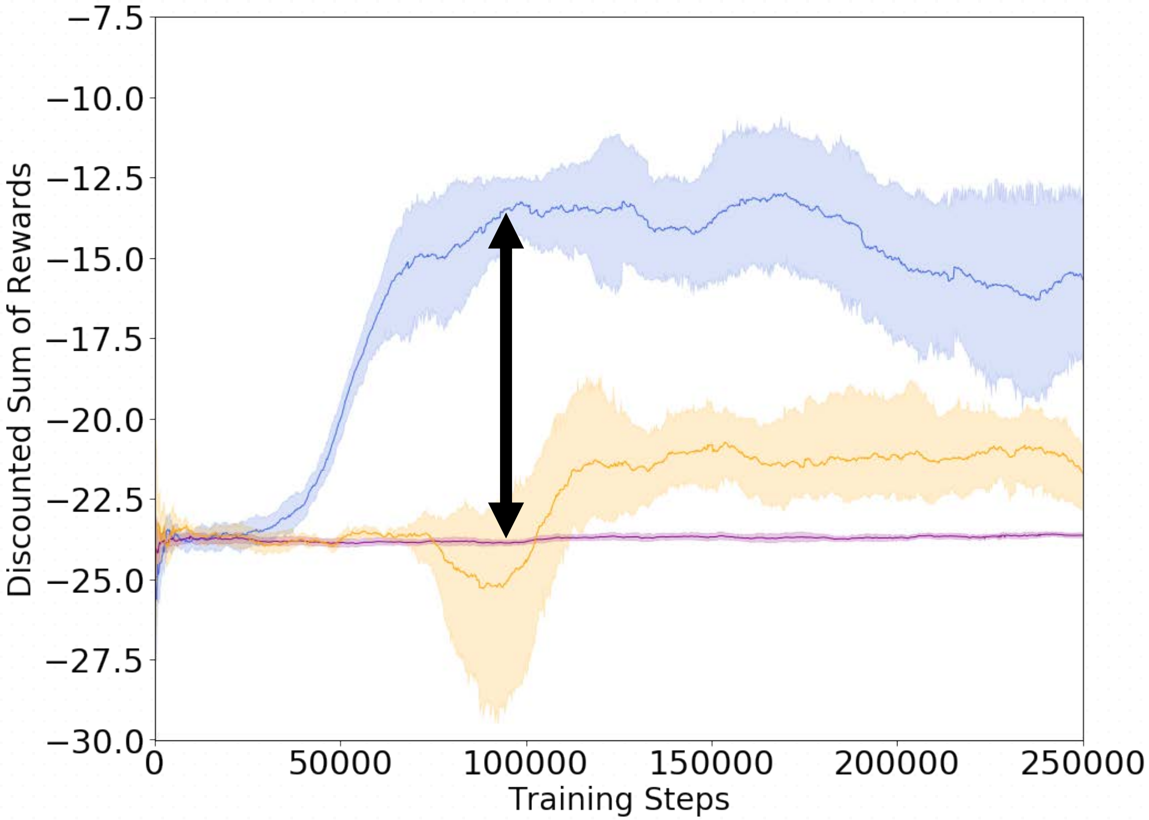

AC-Teach has agent learn behaviors not in teacher set

AC-Teach: CoRL 2019

Visuo-Motor Skills

Grasping
Pushing
Picking
Wiping
Open door

IROS 2019

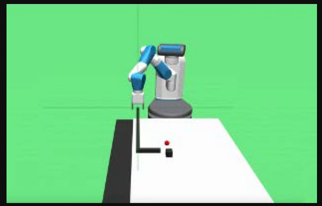

CoRL 2019

Action Representations and Weak-Supervision provide structure to enable learning efficiency and generalization

Generalizable Autonomy in Robot Manipulation

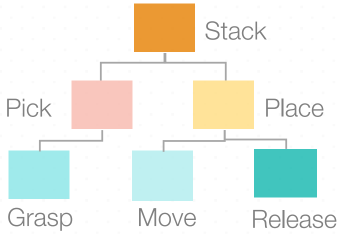

Object Arrangement

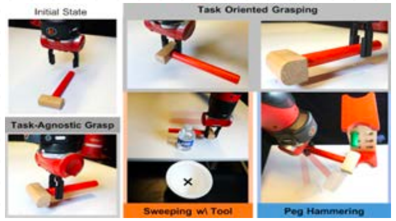

RSS 2018, IJRR 2019

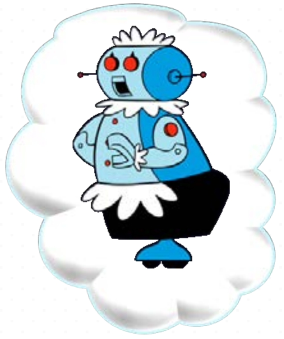

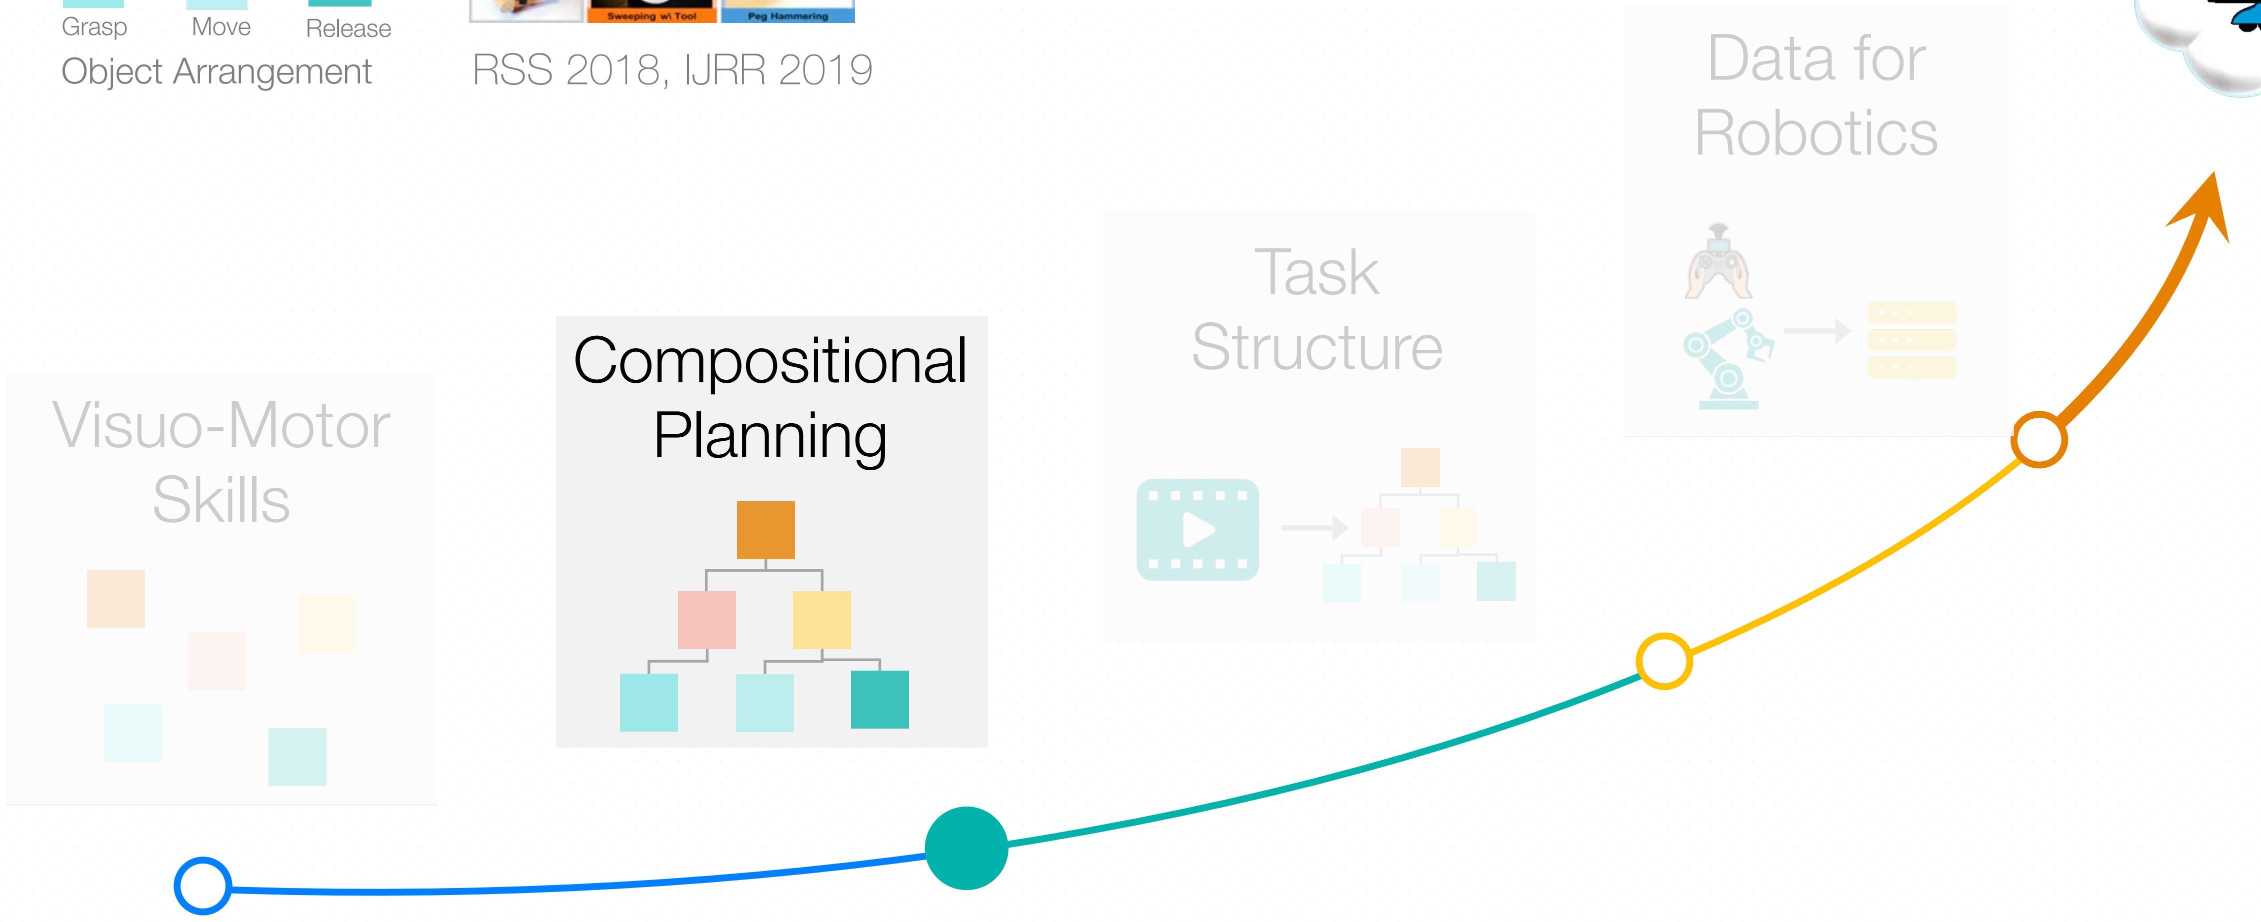

Sequential Skills

Skills: Surface Wiping


Skills: Tool Use

Primitive Skills

- Grasping
- Pushing
- Picking
- Wiping
- Open door

Sequential Skills

- Hammering (with unknown objects)
- Cutting (with new knife)
- Sweeping (with new broom)

Sequential Skills: Manipulation with Tools

Tool-Use

Initial State

Unknown Object

Task-Agnostic Grasping[1]

Optimizes for Grasp Success Only

Suboptimal for Task!

Task-Oriented Grasping

Sweeping

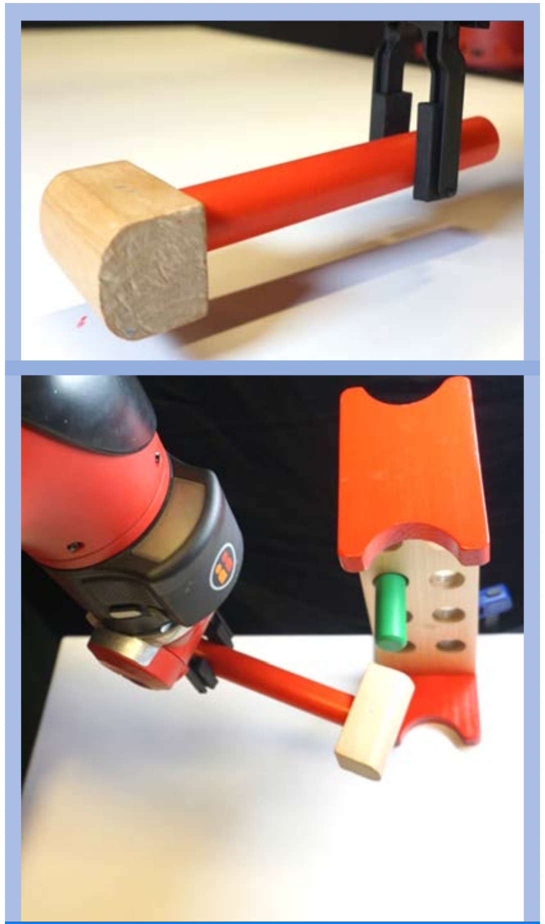

Hammering

[1] Pinto et al. '16, Levine et al. '16, Mahler et al. '18, Kalashnikov et al. '18

Visuo-Motor Skills: Task-Oriented Grasping

Input =

Output =

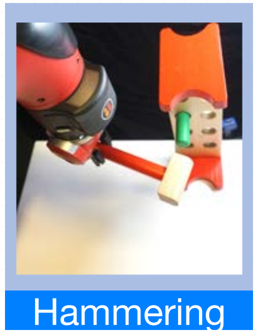

Depth Image o

Task T

Grasp Parameters $g = (g_x, g_y, g_z, g_\phi)$

Policy a $= \pi(o, g)$

Task Success

Grasp Success

$$g^*, \pi^* = \underset{g,\pi}{\mathrm{argmax}}\ Q_T^\pi(o, g)$$

Score Function

$$Q_T^\pi(o, g) = P_\pi(S_T = 1 \mid S_G = 1, o, g)\, P(S_G = 1 \mid o, g)$$

$$Q_T^\pi(o, g) = Q_{T|G}^\pi(o, g)\, Q_G(o, g)$$

Task Success

Task Conditioned Grasp Success

Grasp Success

RSS 2018, IJRR 2019

Visuo-Motor Skills: Task-Oriented Grasping

Input: Depth Image (64x64)

CNN Model	Output	Loss Function
grasping model	$Q_G(o,g)$	$\mathcal{L}(S_G, Q_G(o,g;\theta_1))$
task-oriented grasping model	$Q^{\pi}_{T\mid G}(o,g)$	$\mathbb{I}(S_G = 1)\mathcal{L}(S_T, Q^{\pi}_{T\mid G}(o,g;\theta_2))$
Manipulation	$\pi(a \mid o,g;\theta_3)$	$\nabla_{\theta_3} \log \pi(a \mid o,g;\theta_3)$

Training

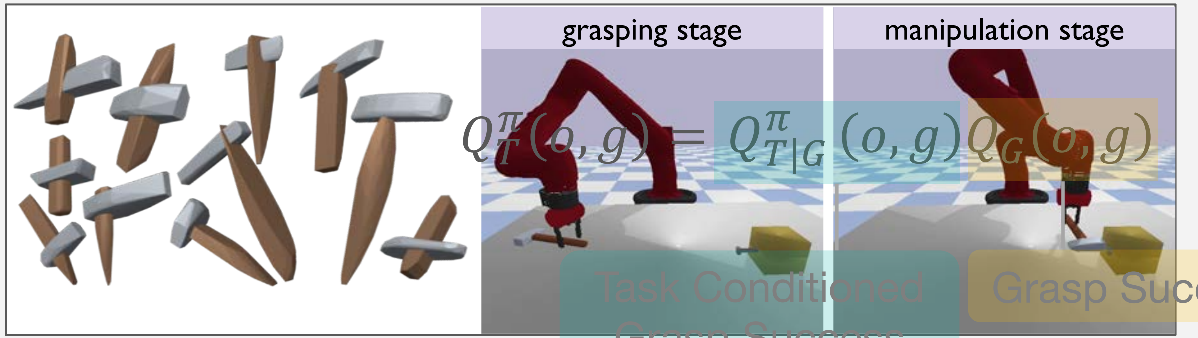

$$Q^{\pi}_{T}(o,g) = Q^{\pi}_{T\mid G}(o,g) Q_G(o,g)$$

Task Conditioned Grasp Success

Grasp Success

18K diverse training objects

large-scale data collection (1.5m Episodes)

RSS 2018, IJRR 2019

Visuo-Motor Skills: Task-Oriented Grasping

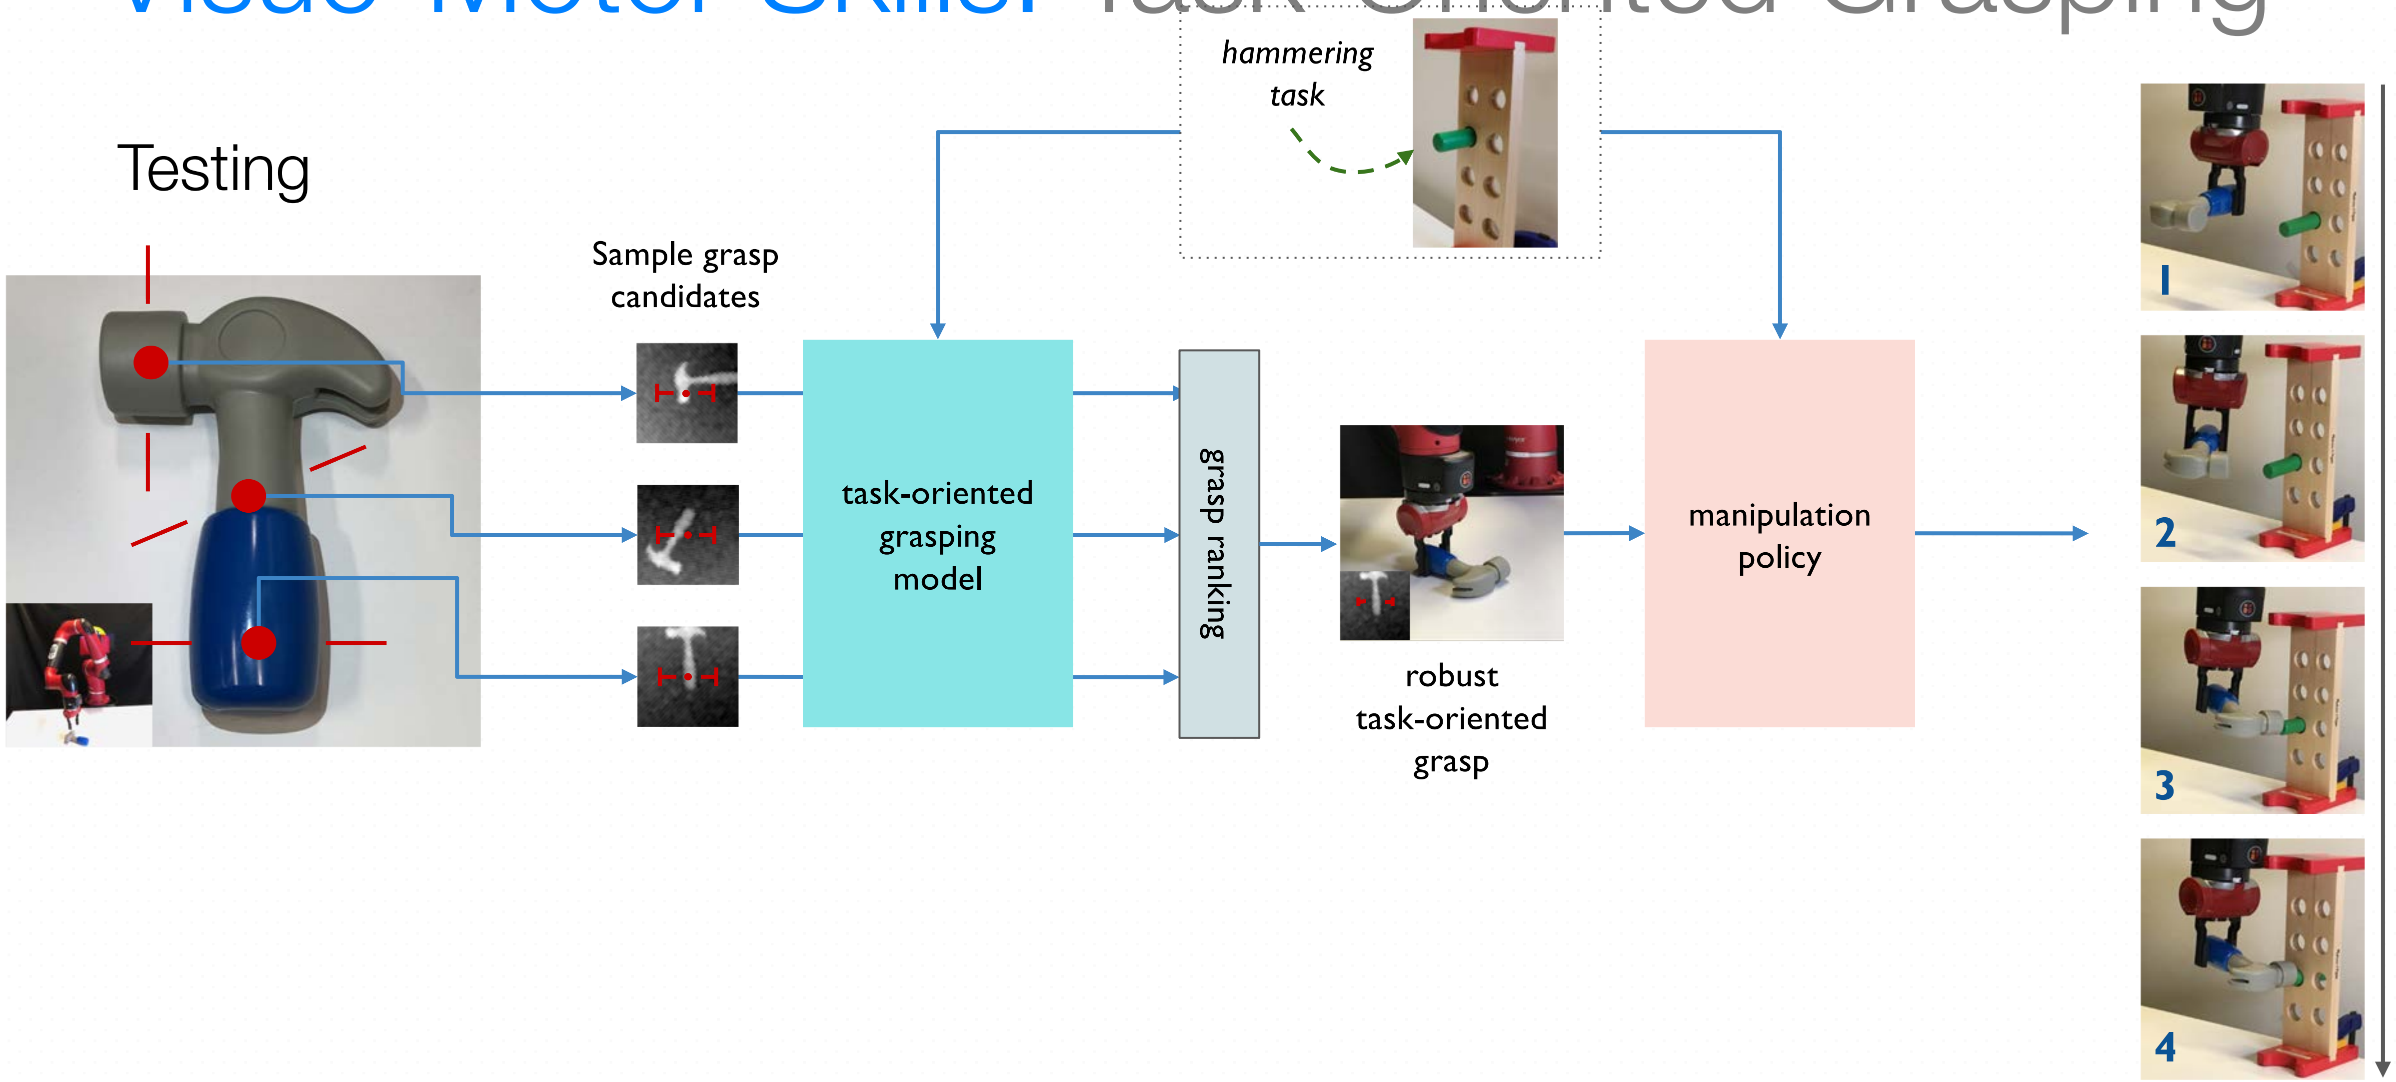

Sequential Skills: Task-Oriented Grasping

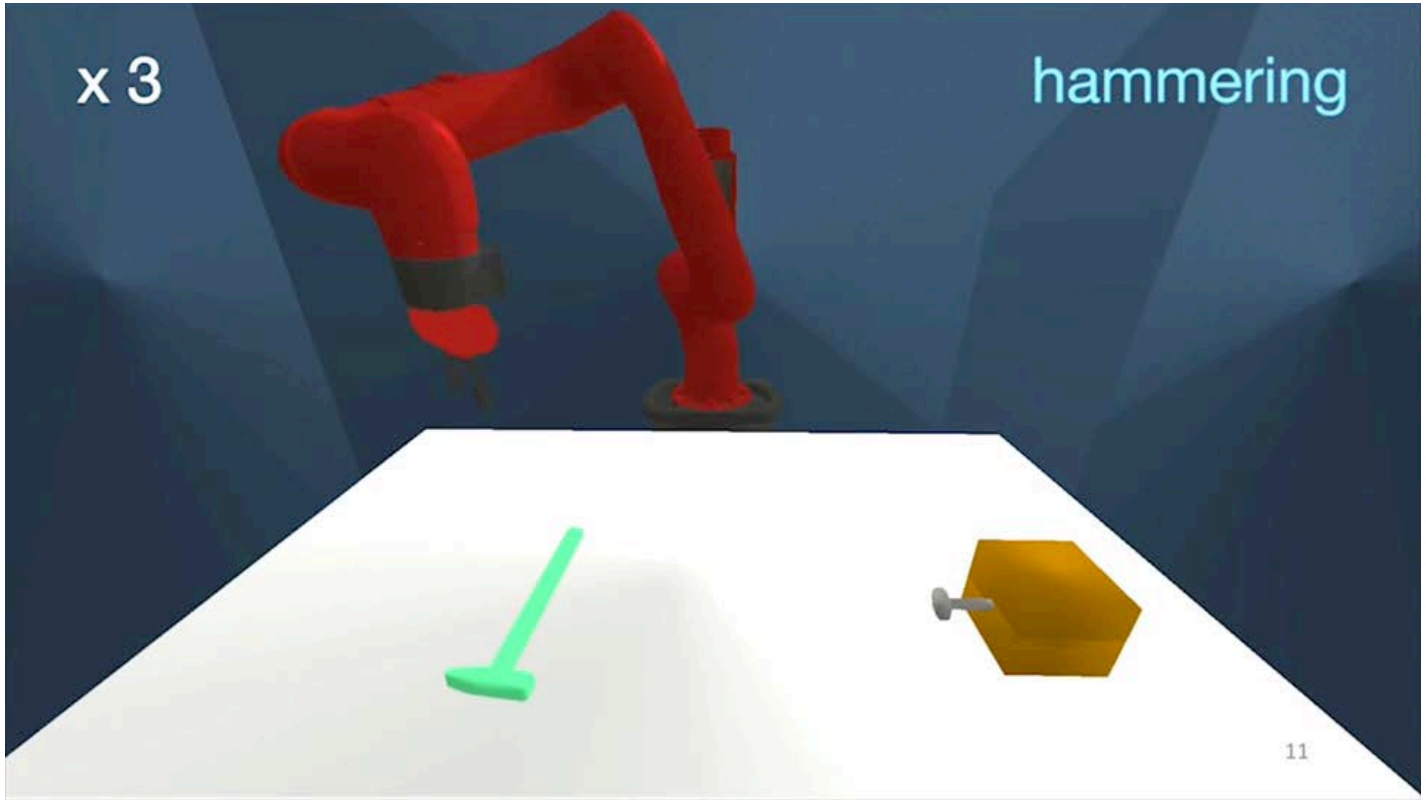

Trained Policy Rollout (Ours)
Unseen Test Objects

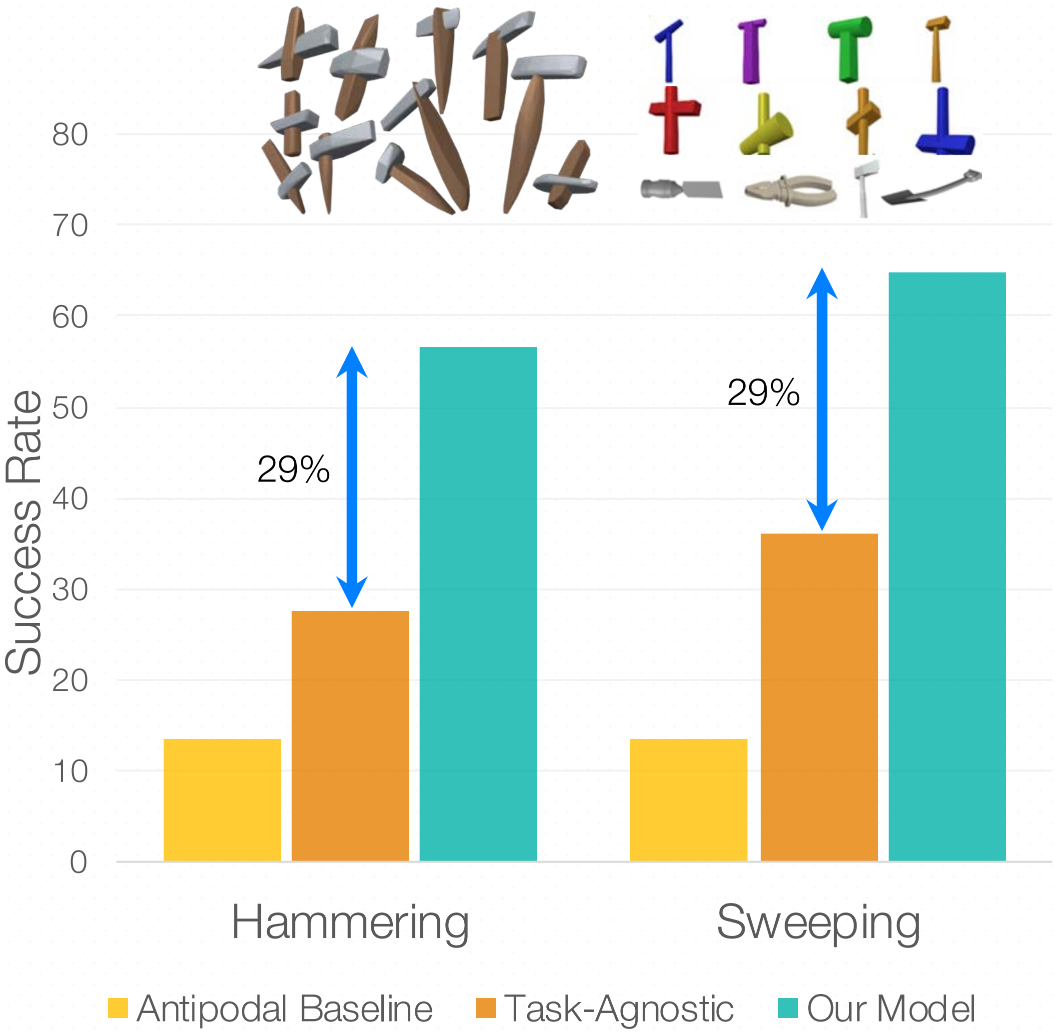

Two-Stage Optimization
Joint Optimization

RSS 2018, IJRR 2019

Sequential Skills: Task-Oriented Grasping

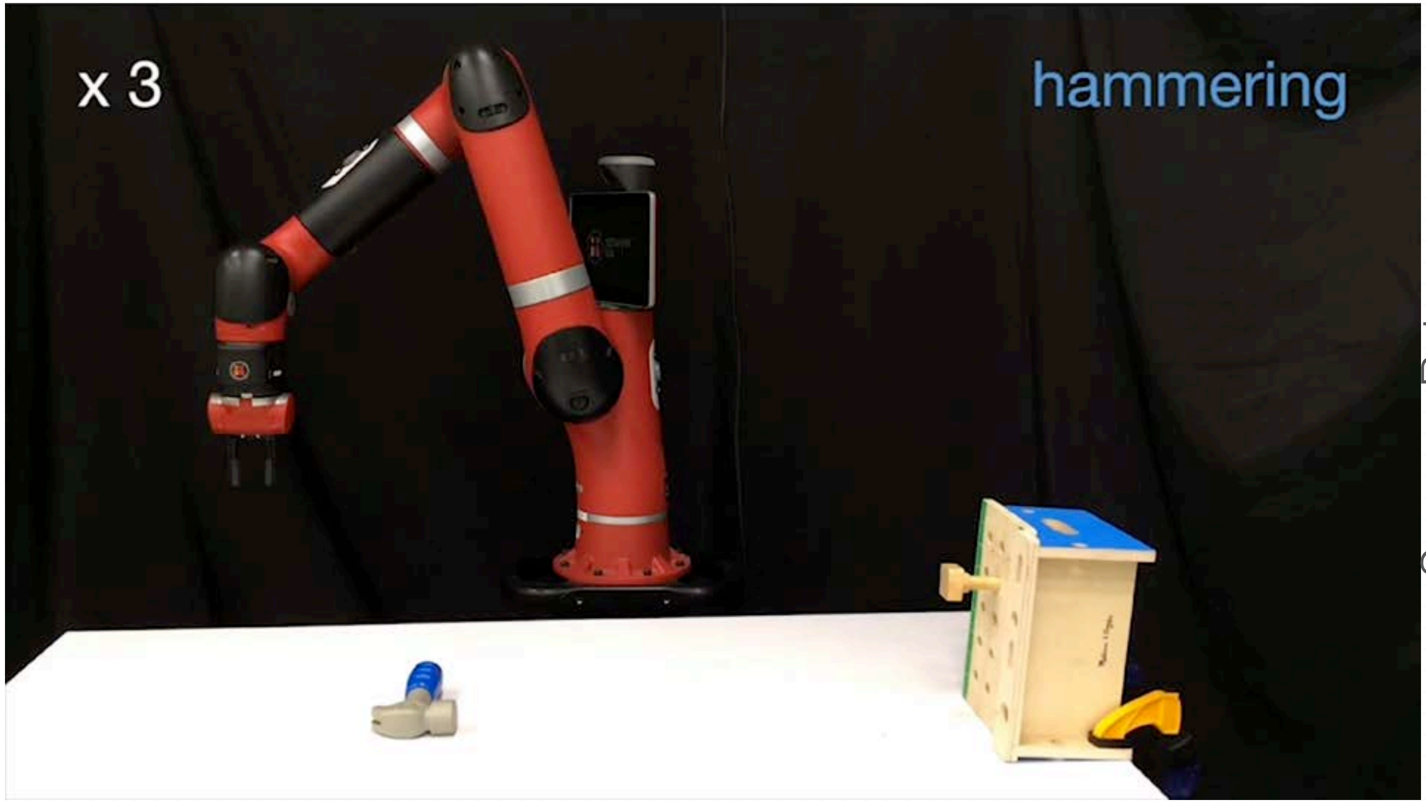

Trained Policy Rollout (Ours)
Unseen Test Objects

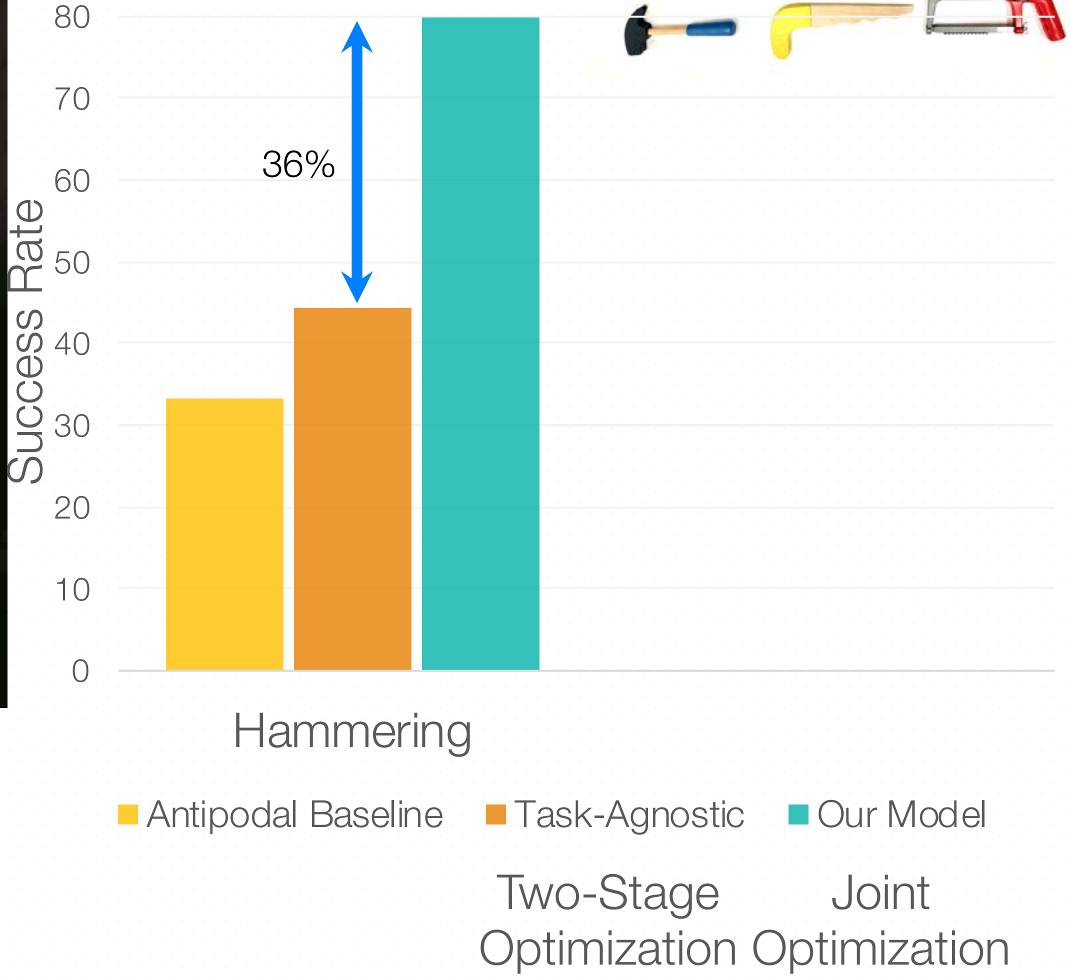

RSS 2018, IJRR 2019

Sequential Skills: Task-Oriented Grasping

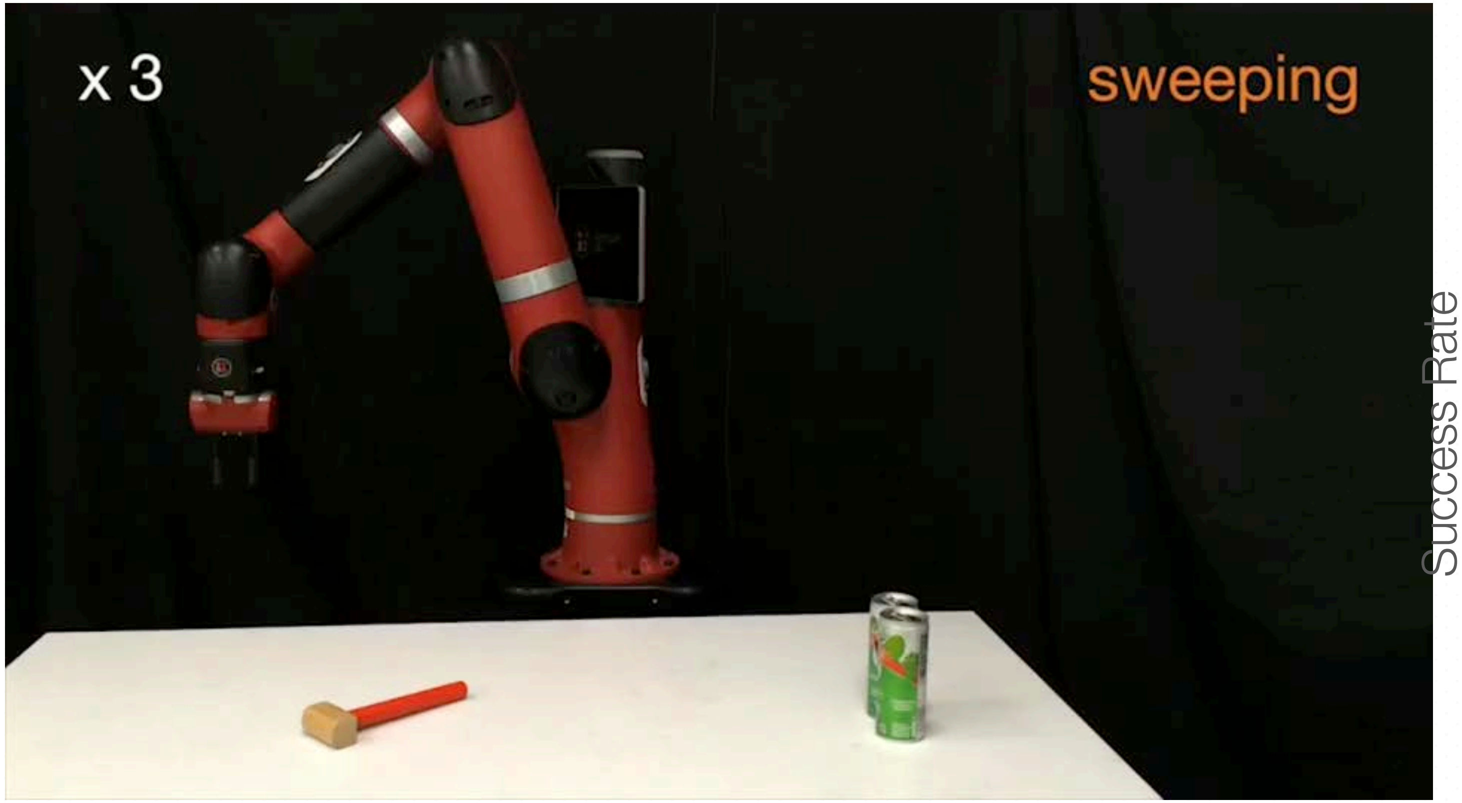

Trained Policy Rollout (Ours)
Unseen Test Objects

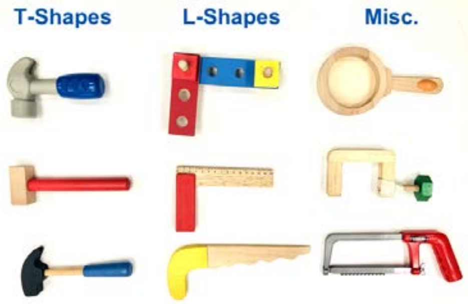

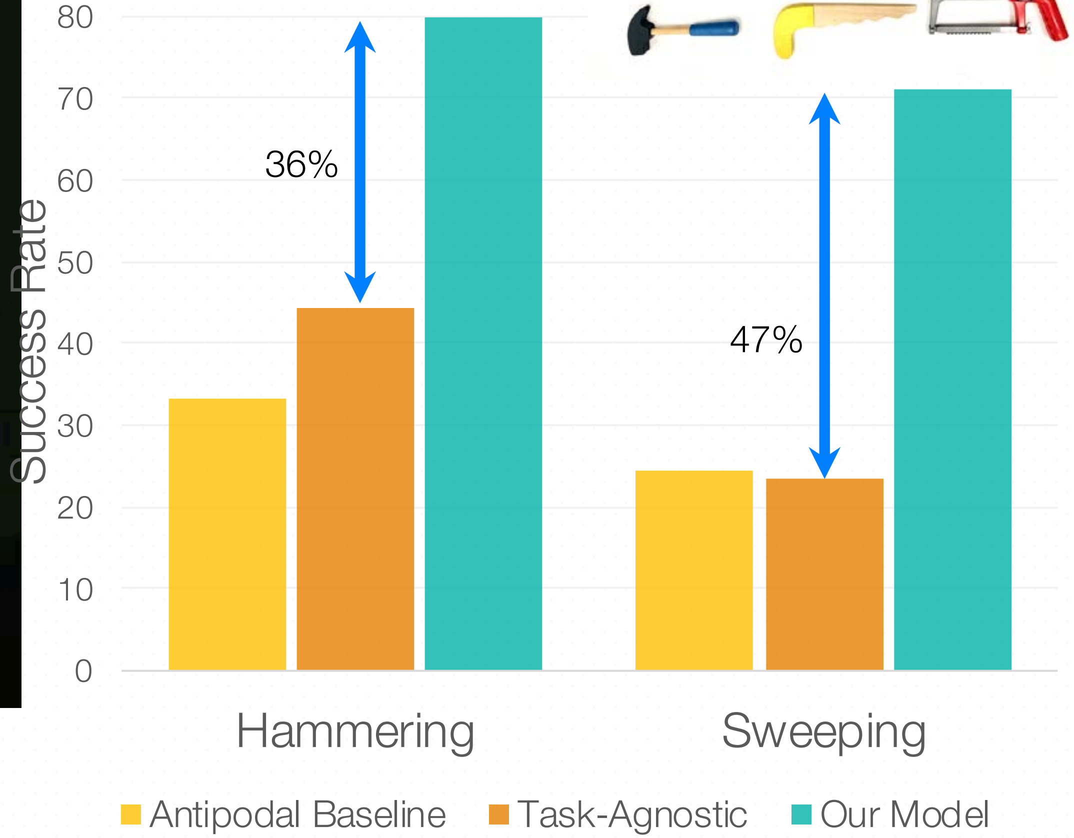

Two-Stage Optimization
Joint Optimization

RSS 2018, IJRR 2019

Sequential Skills

Skills: Surface Wiping

Skills: Tool Use

Primitive Skills

Grasping

Pushing

Picking

Wiping

Open door

Sequential Skills

Hammering (with unknown objects)

Cutting (with new knife)

Sweeping (with new broom)

Sequential Skills: Multi-Step Reasoning

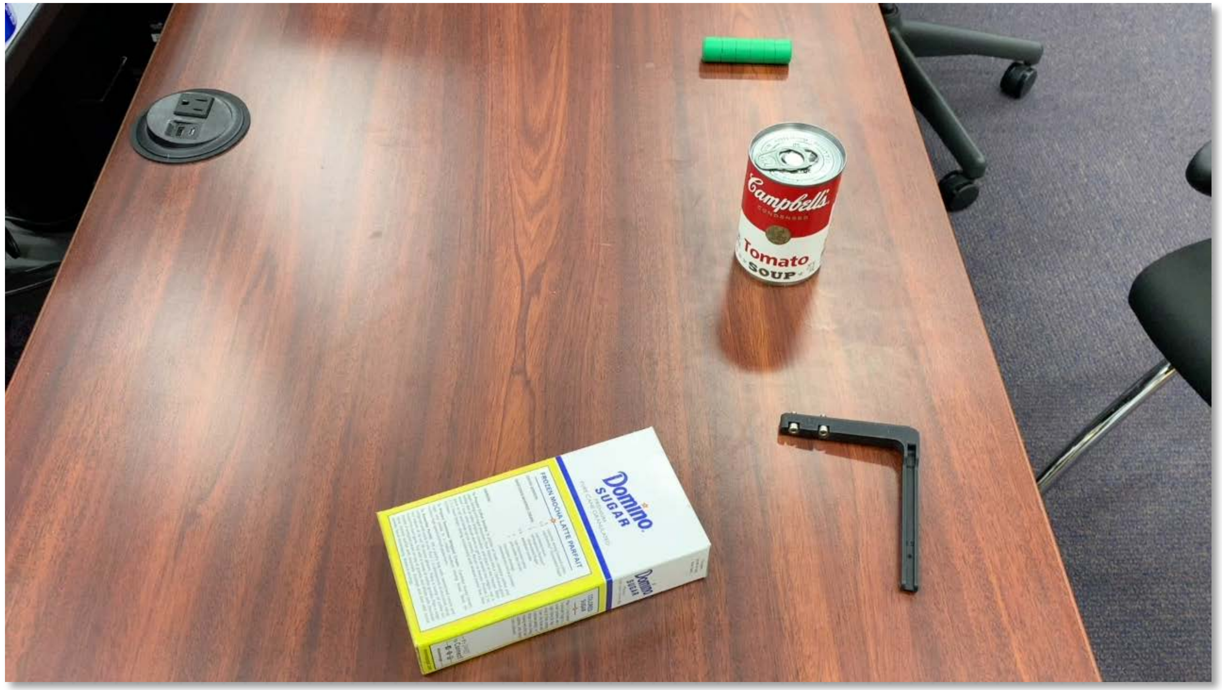

Skills: Multi-Step Reasoning

Generalization

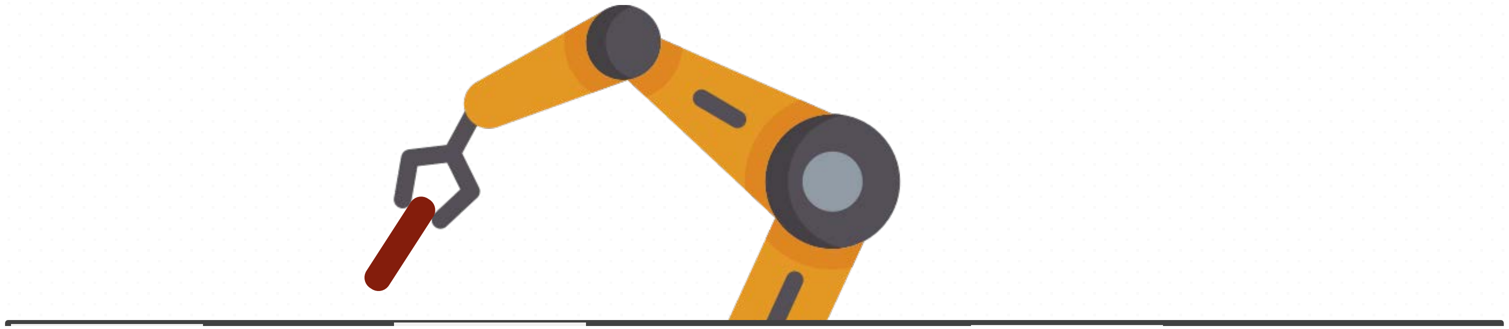

Can we learn multi-step reasoning in robotics under physical and semantic constraints

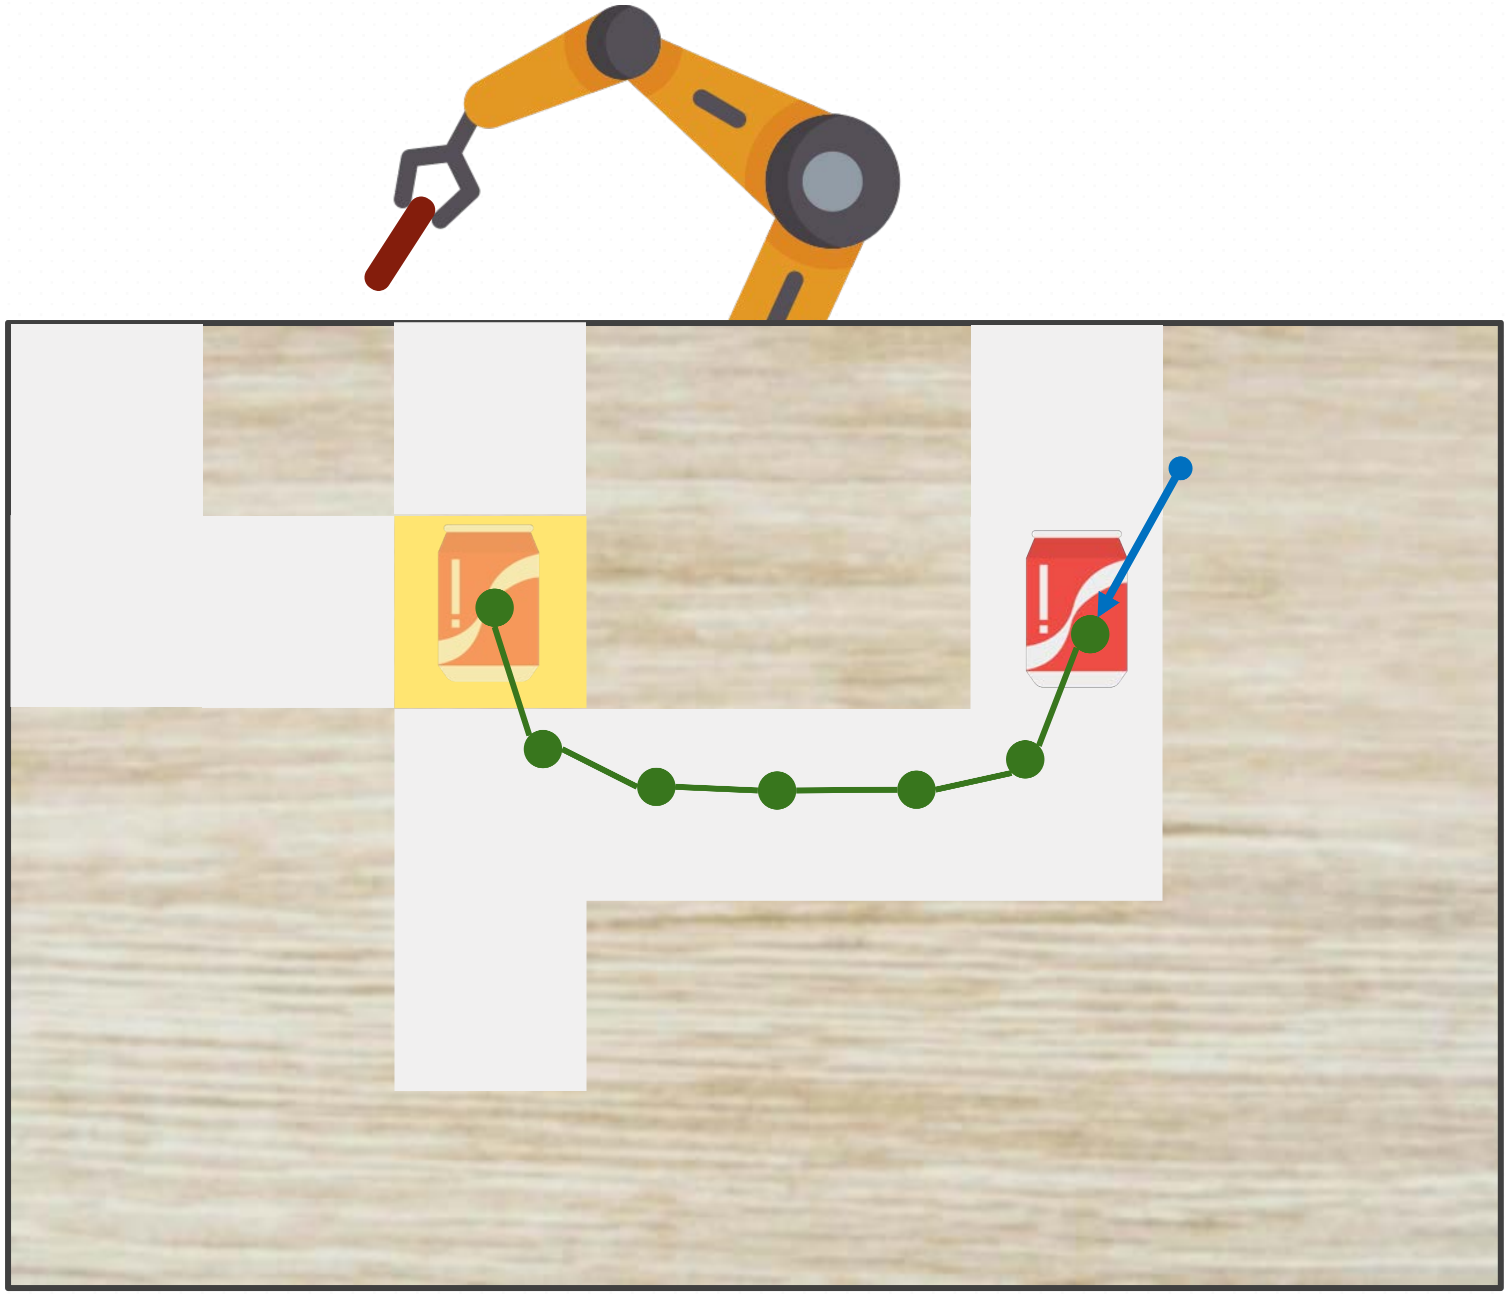

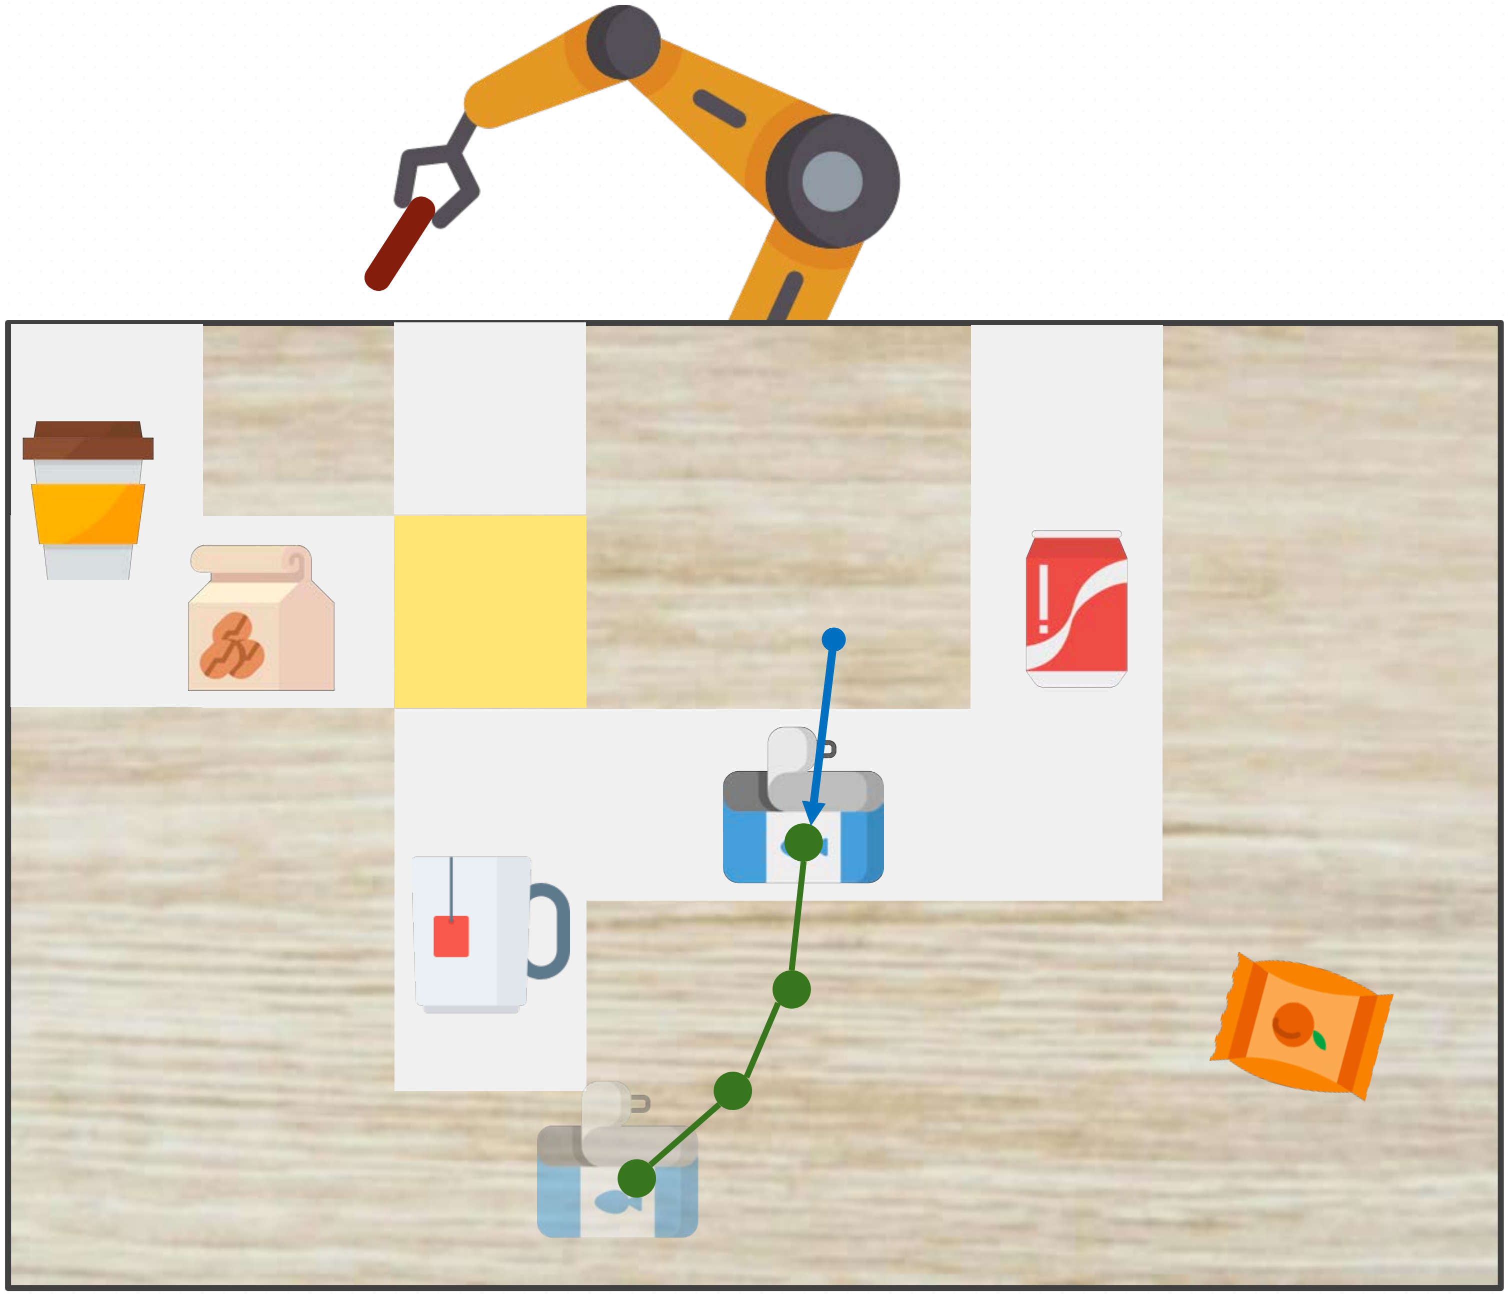

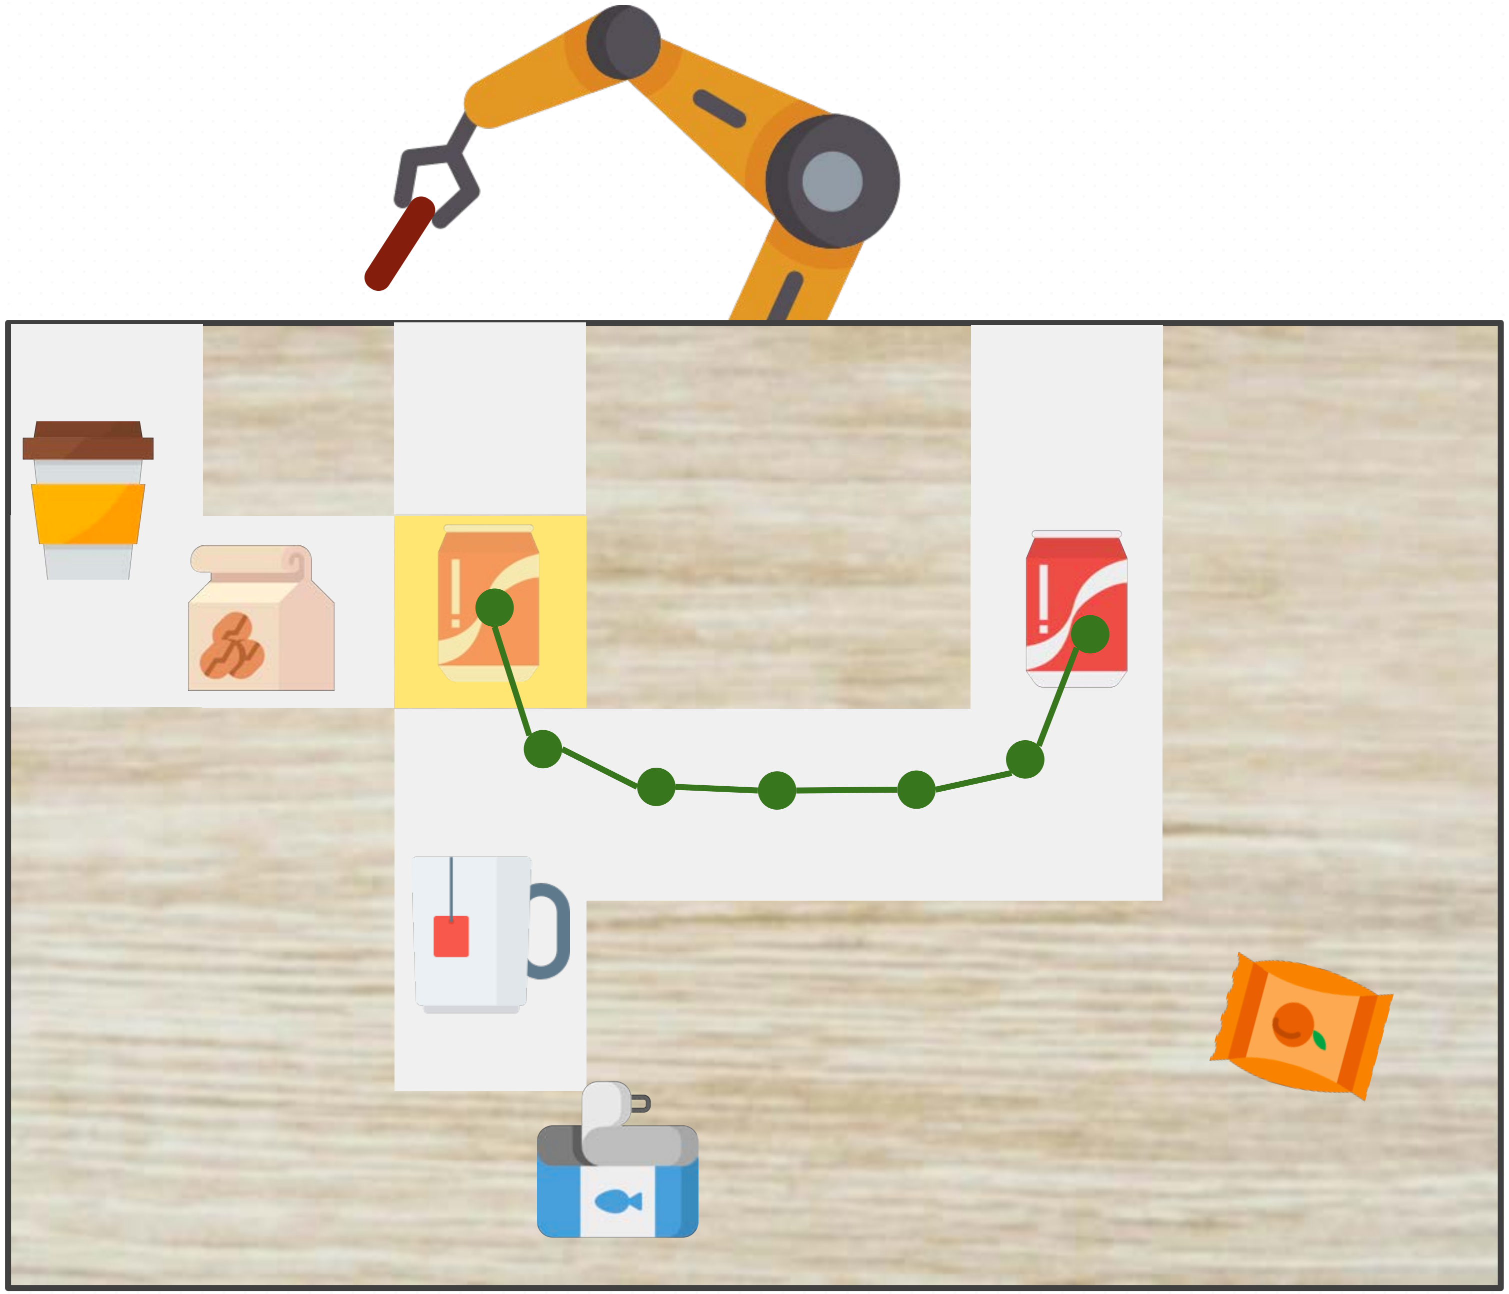

Model-based learning

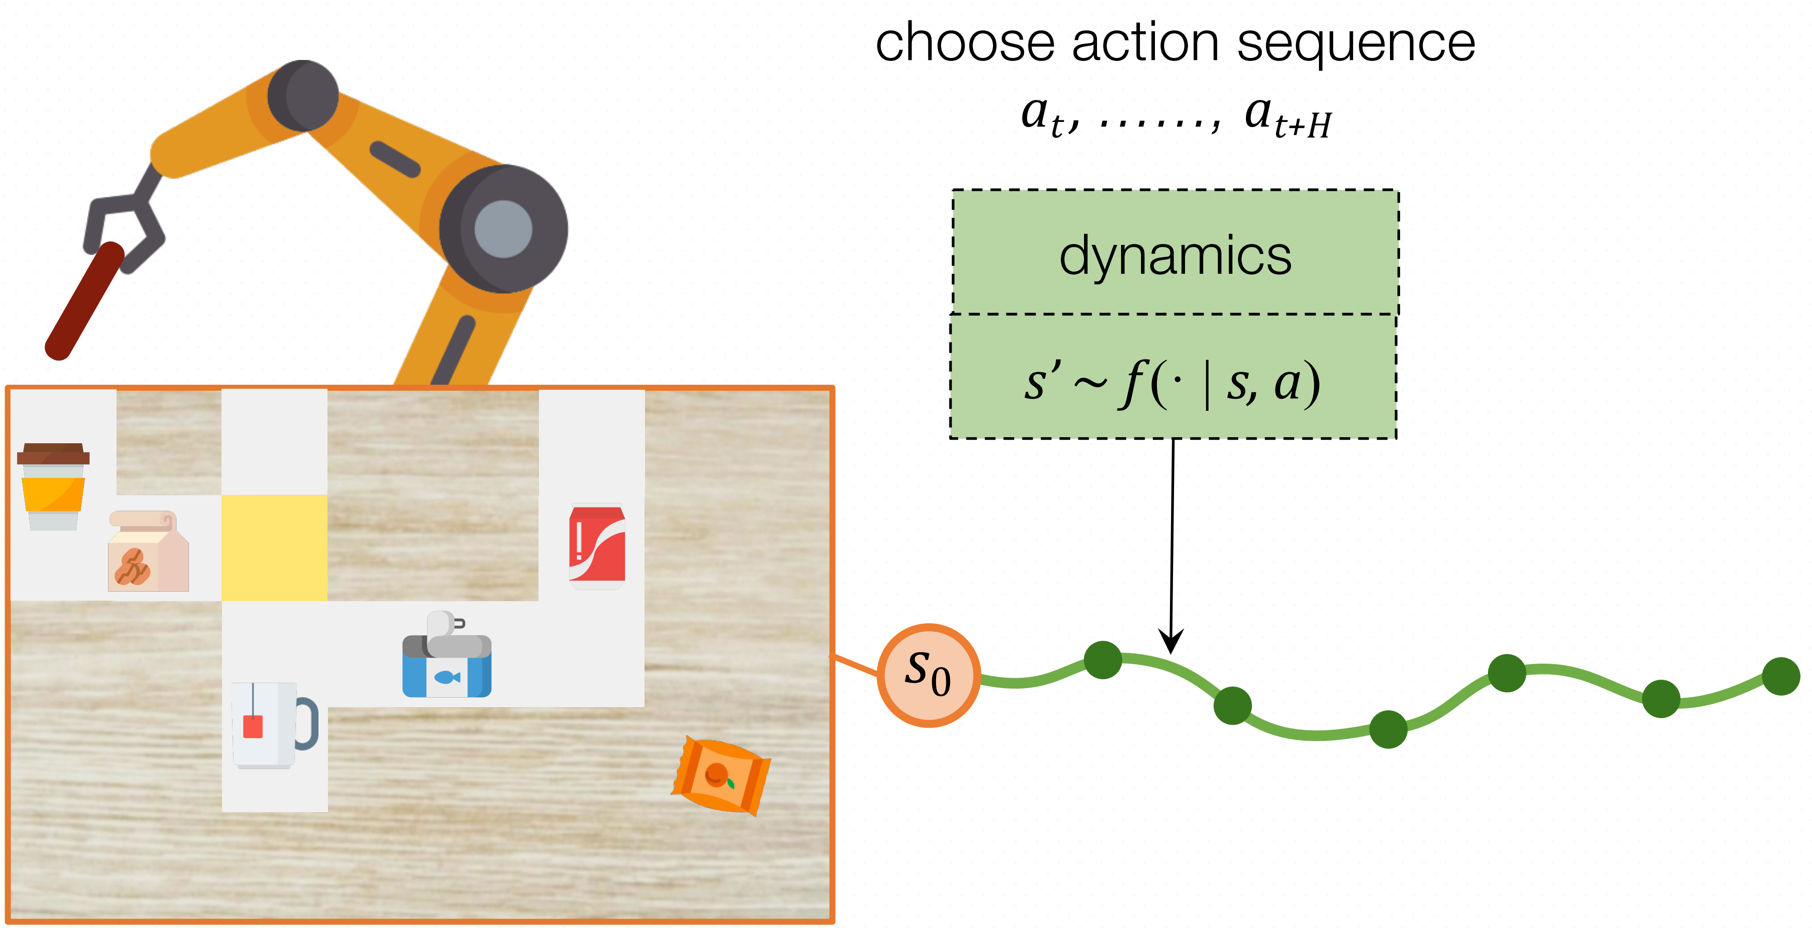

[Deisenroth et al, RSS'07], [Guo et al, NeurIPS'14], [Watter et al, NeurIPS'15], [Finn et al, ICRA'17], ……

Model-based learning

?

data ↑
learning ↑

[Deisenroth et al. RSS’07]

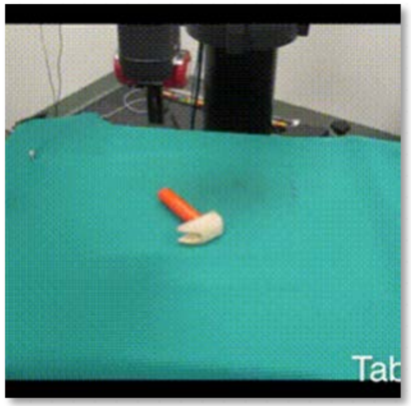

[Agrawal et al. ICRA’16]

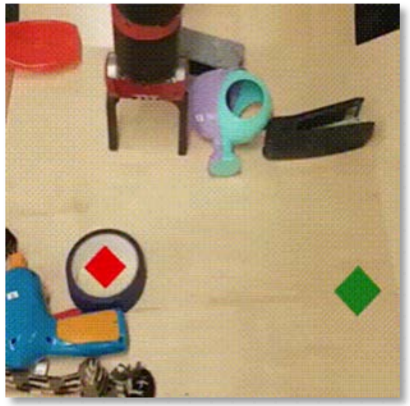

[Ebert et al. CoRL’17]

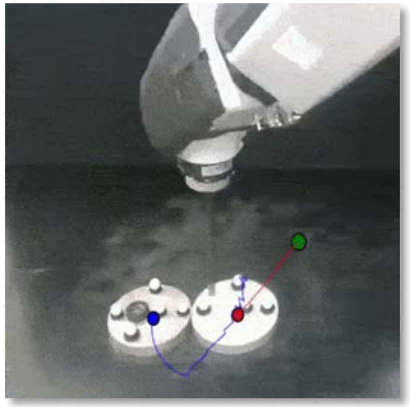

[Janer et al. ICRA’19]

CAVIN: Hierarchical planning in learned latent spaces

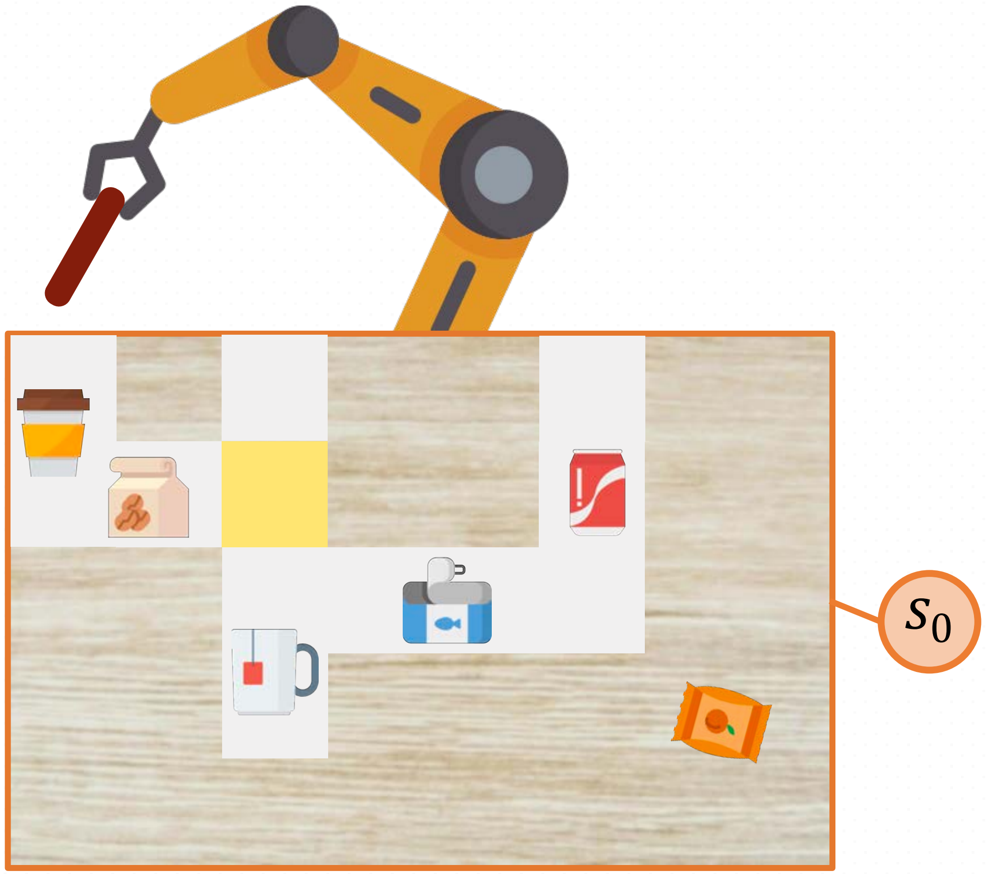

CAVIN Planner

effect code c

motion code z

Leverage Hierarchical Abstraction in Action Space

Without Hierarchical Supervision

CAVIN: Hierarchical planning in learned latent spaces

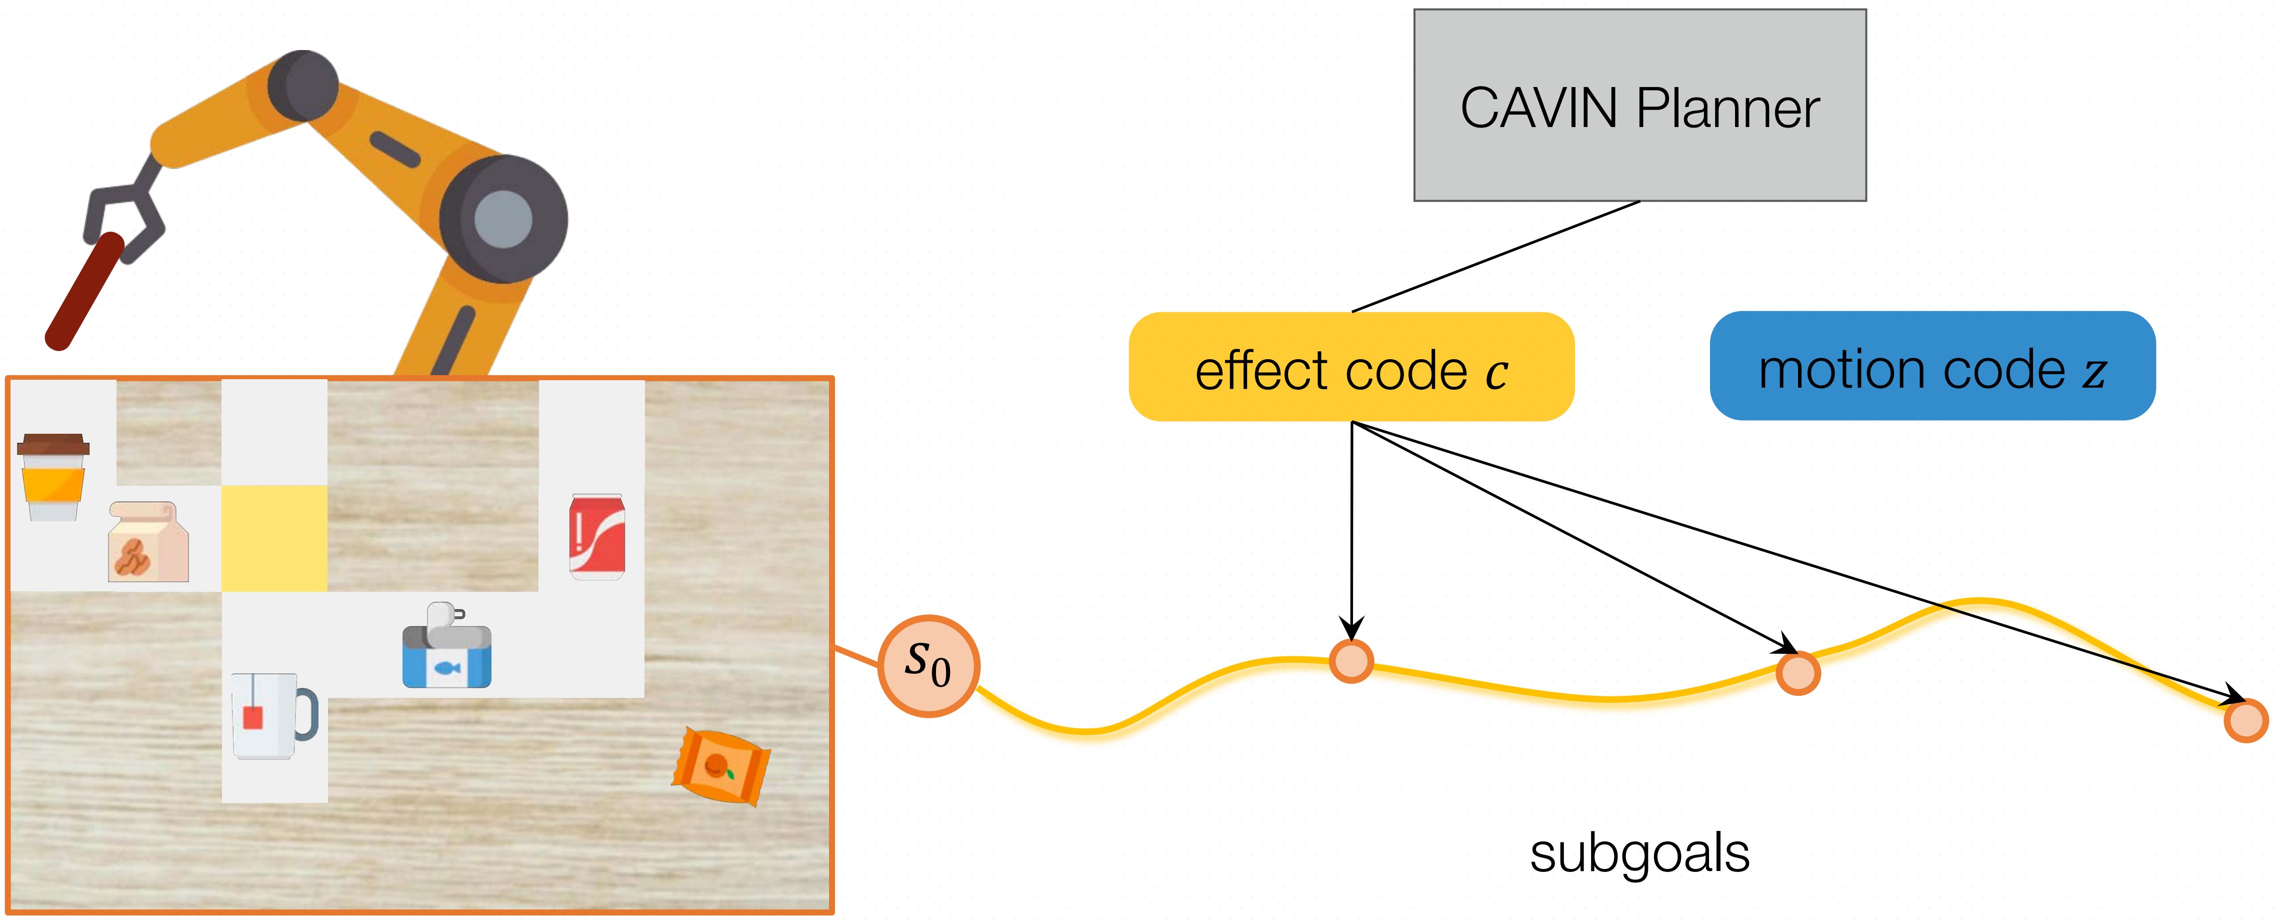

CAVIN: Hierarchical planning in learned latent spaces

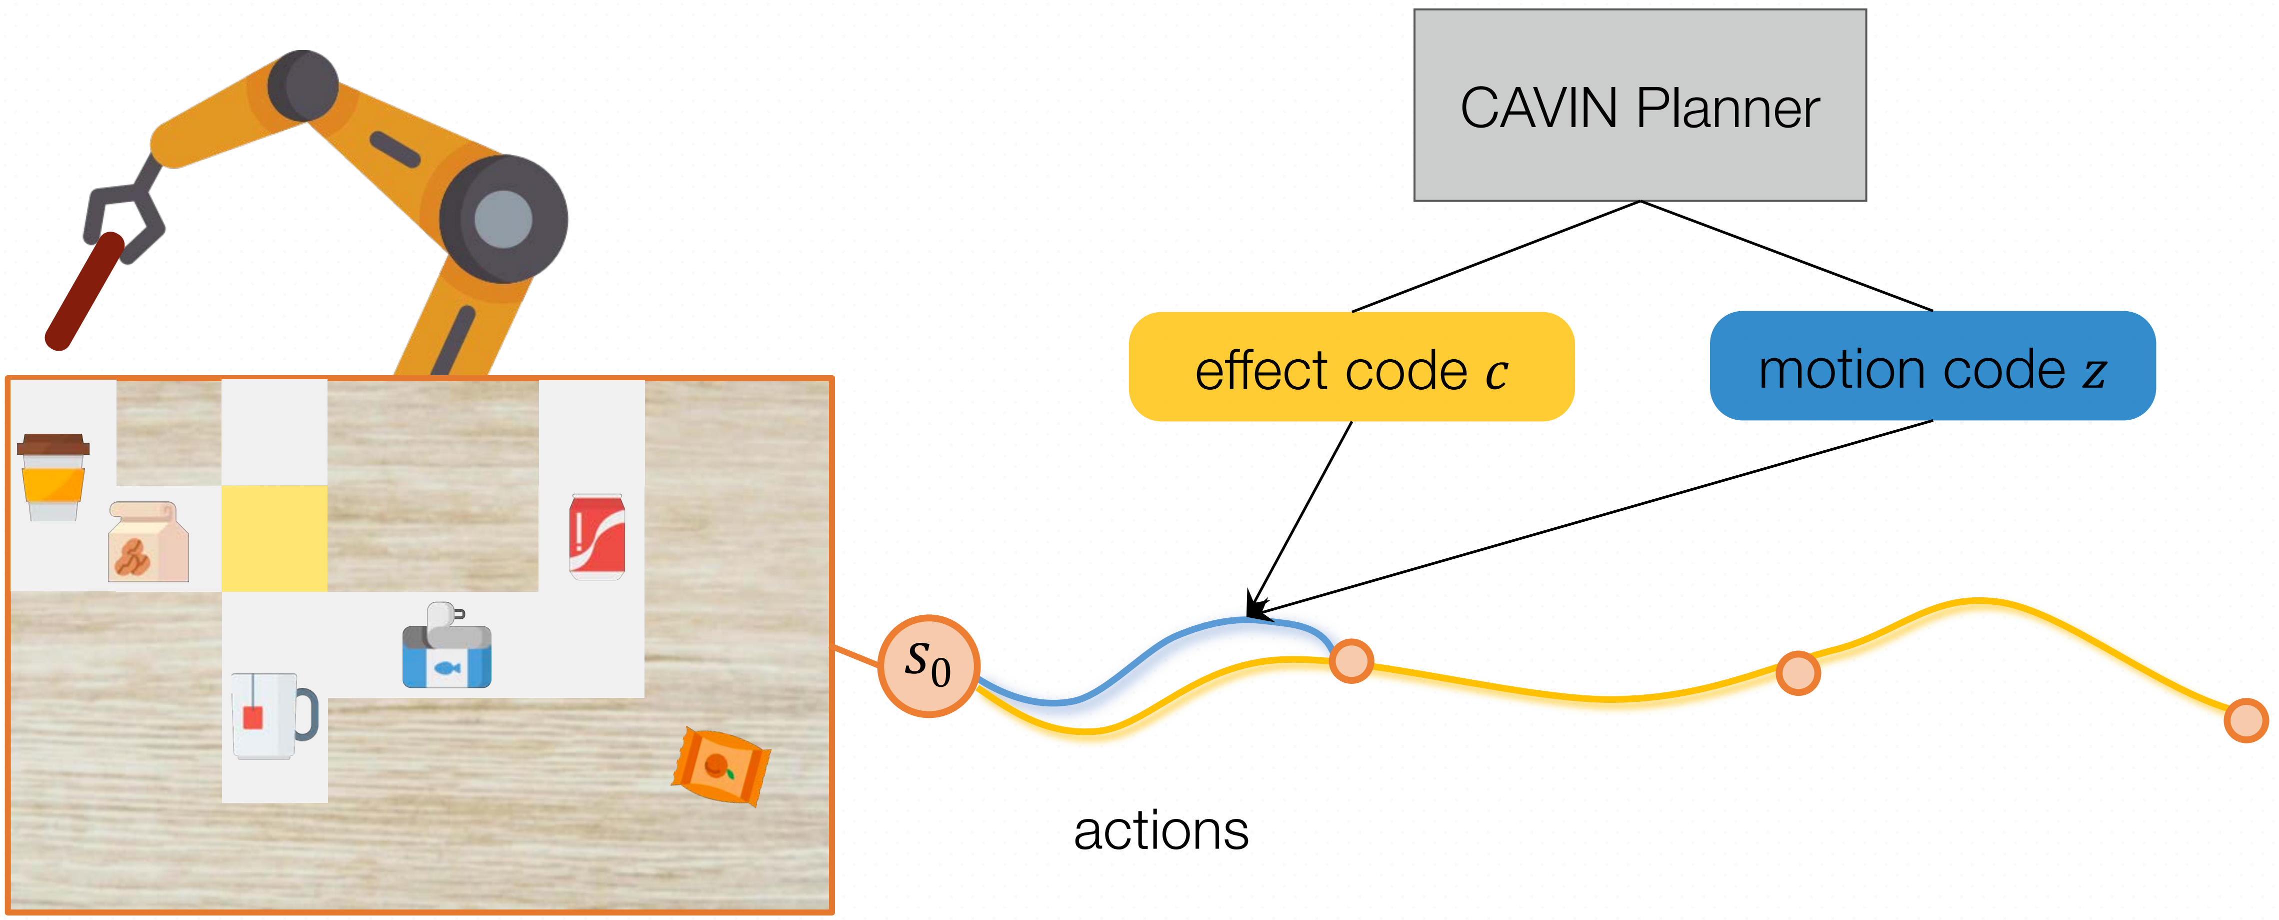

CAVIN: Hierarchical planning in learned latent spaces

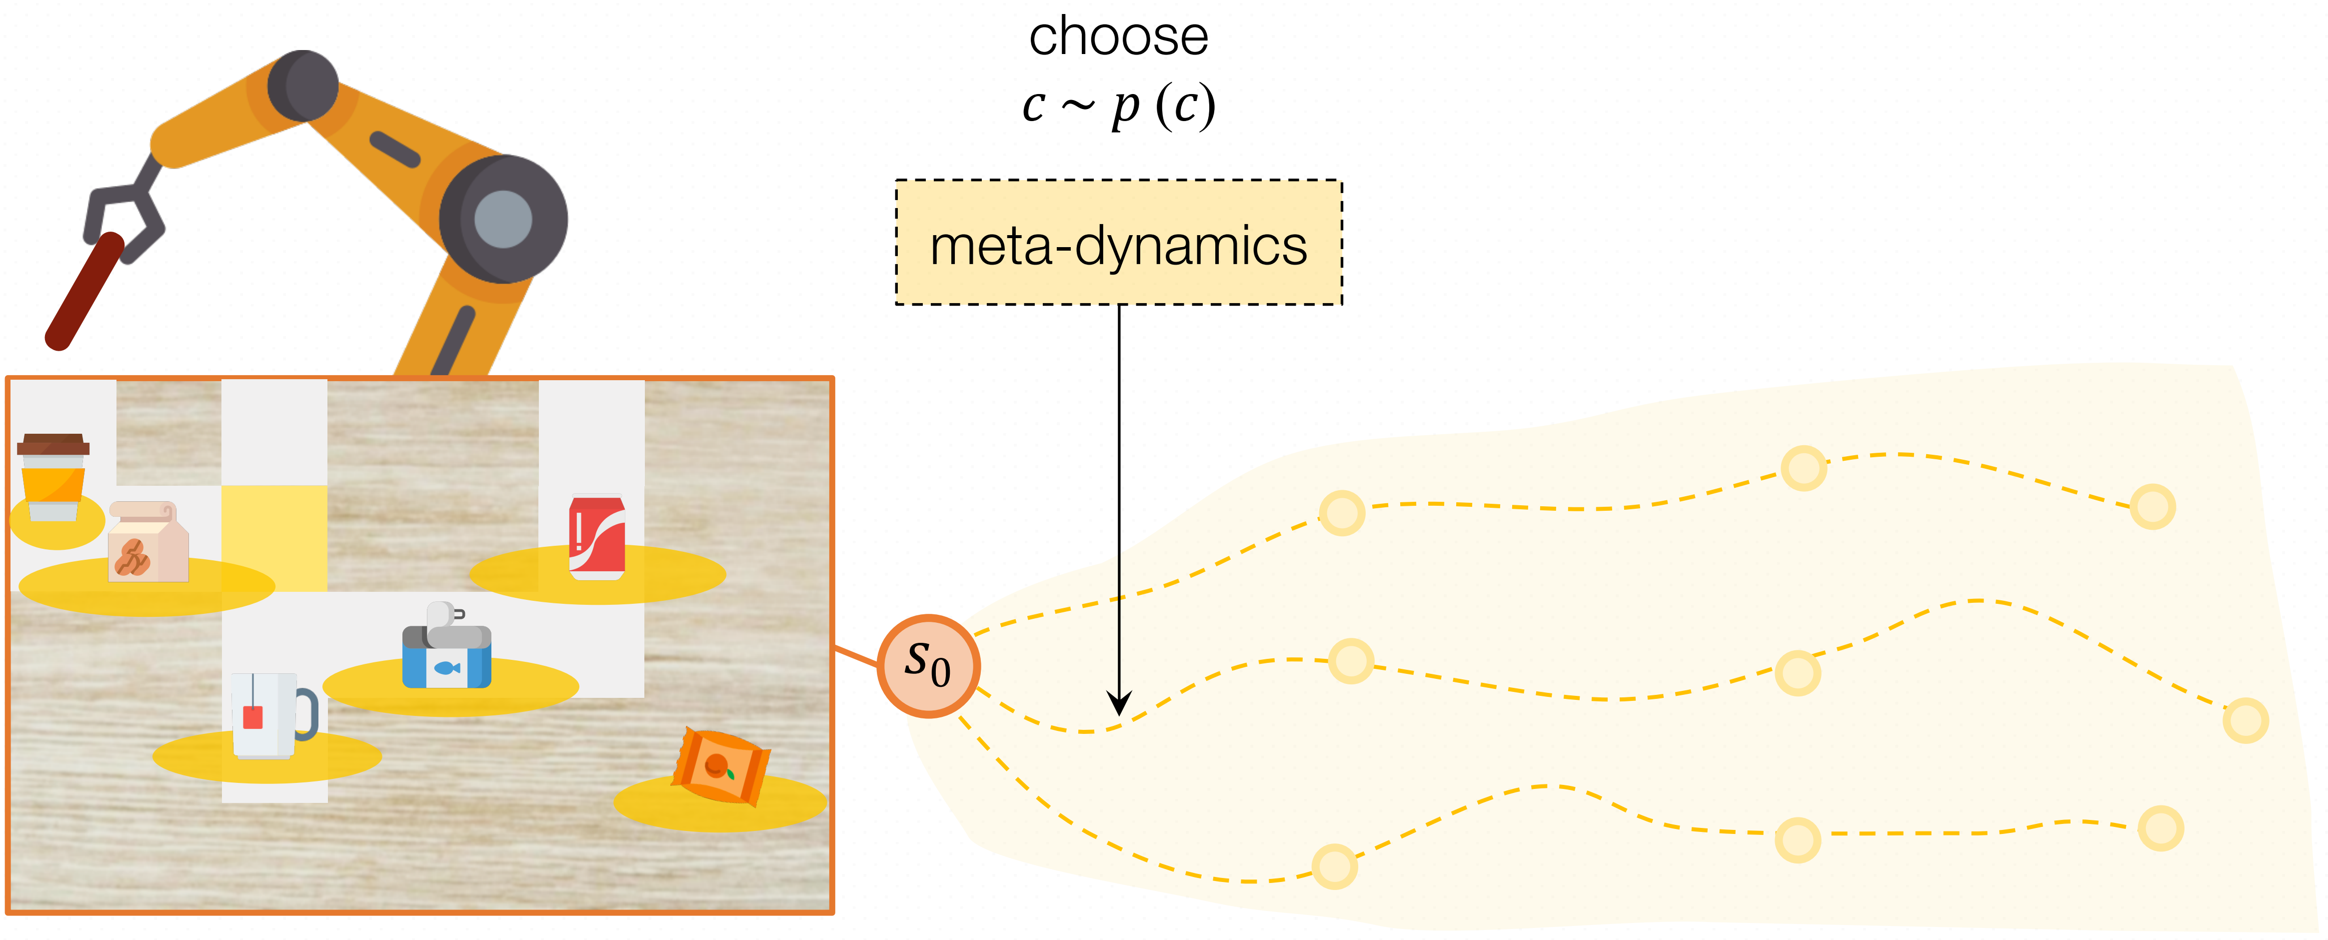

CAVIN: Hierarchical planning in learned latent spaces

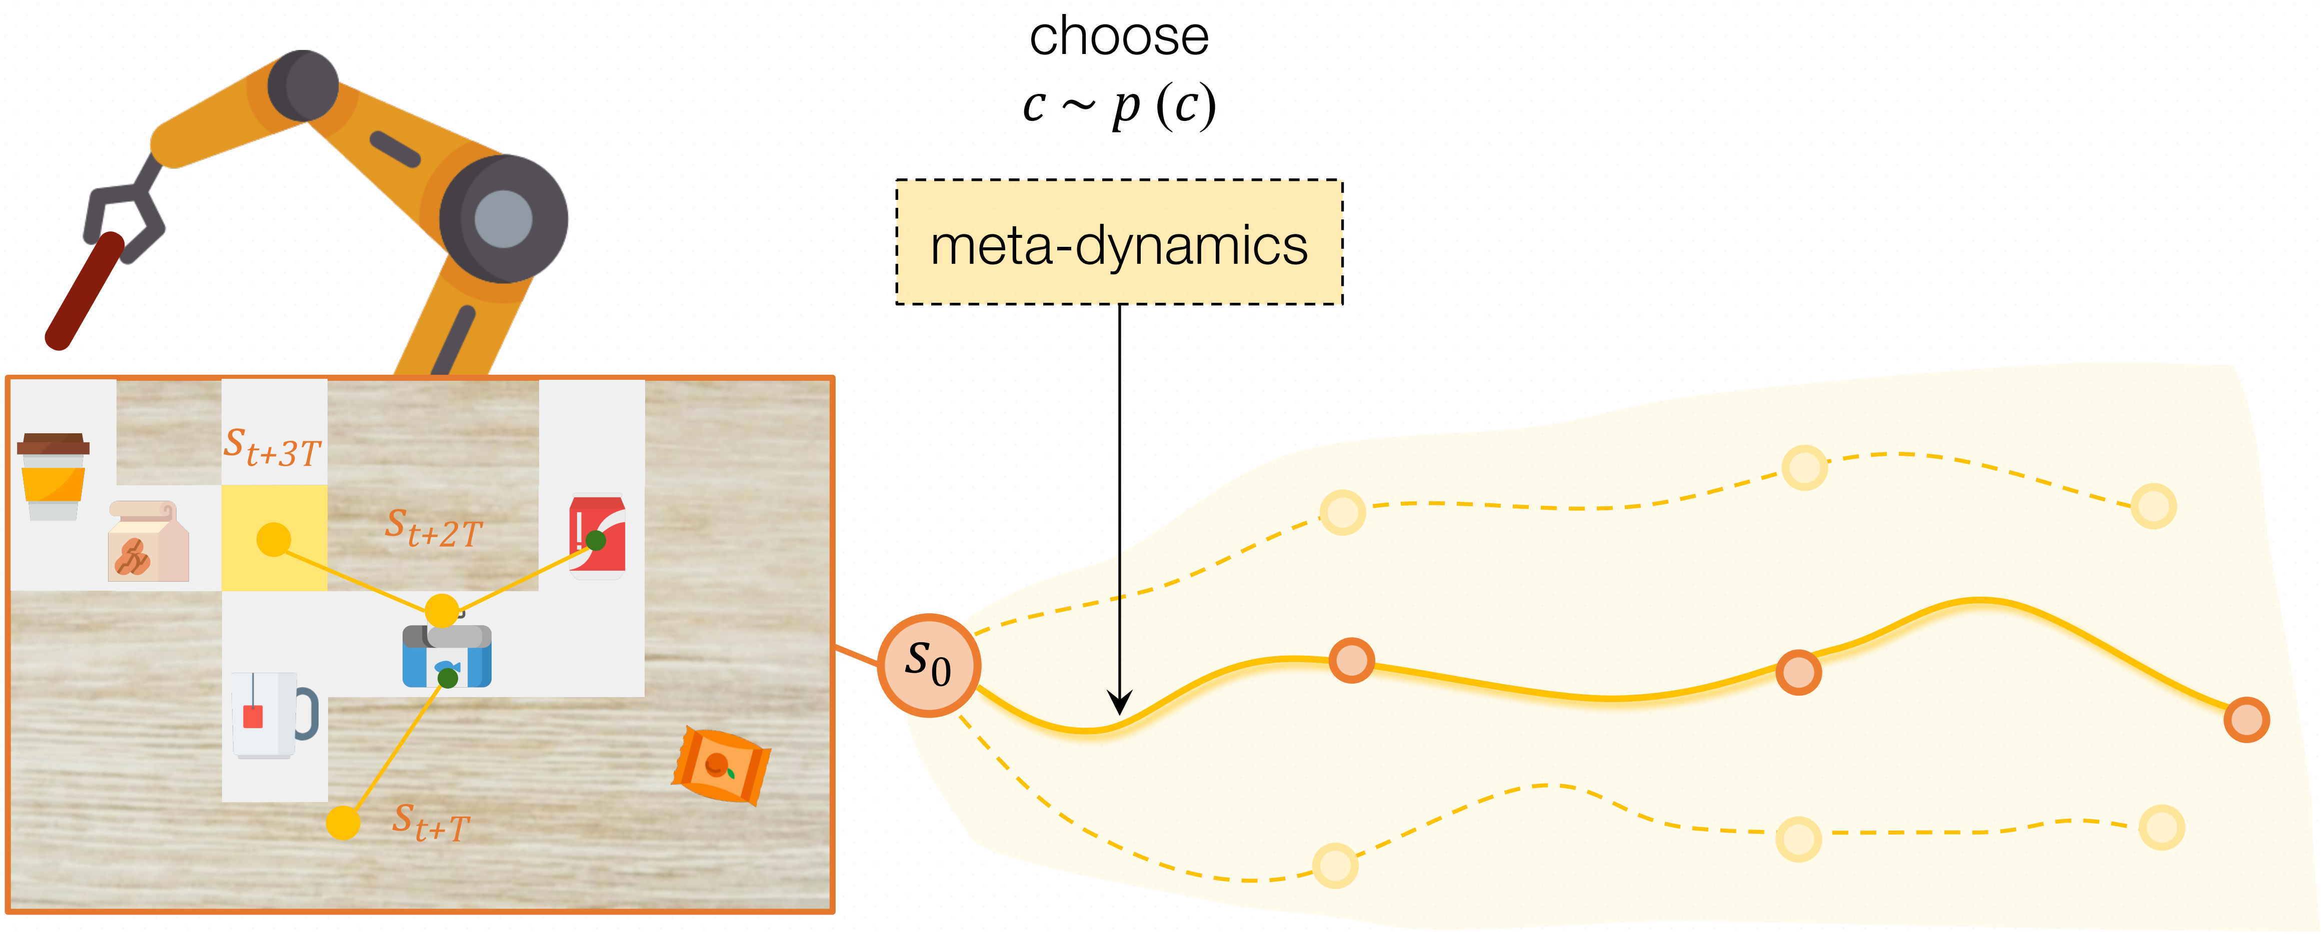

Hierarchical planning in learned latent spaces

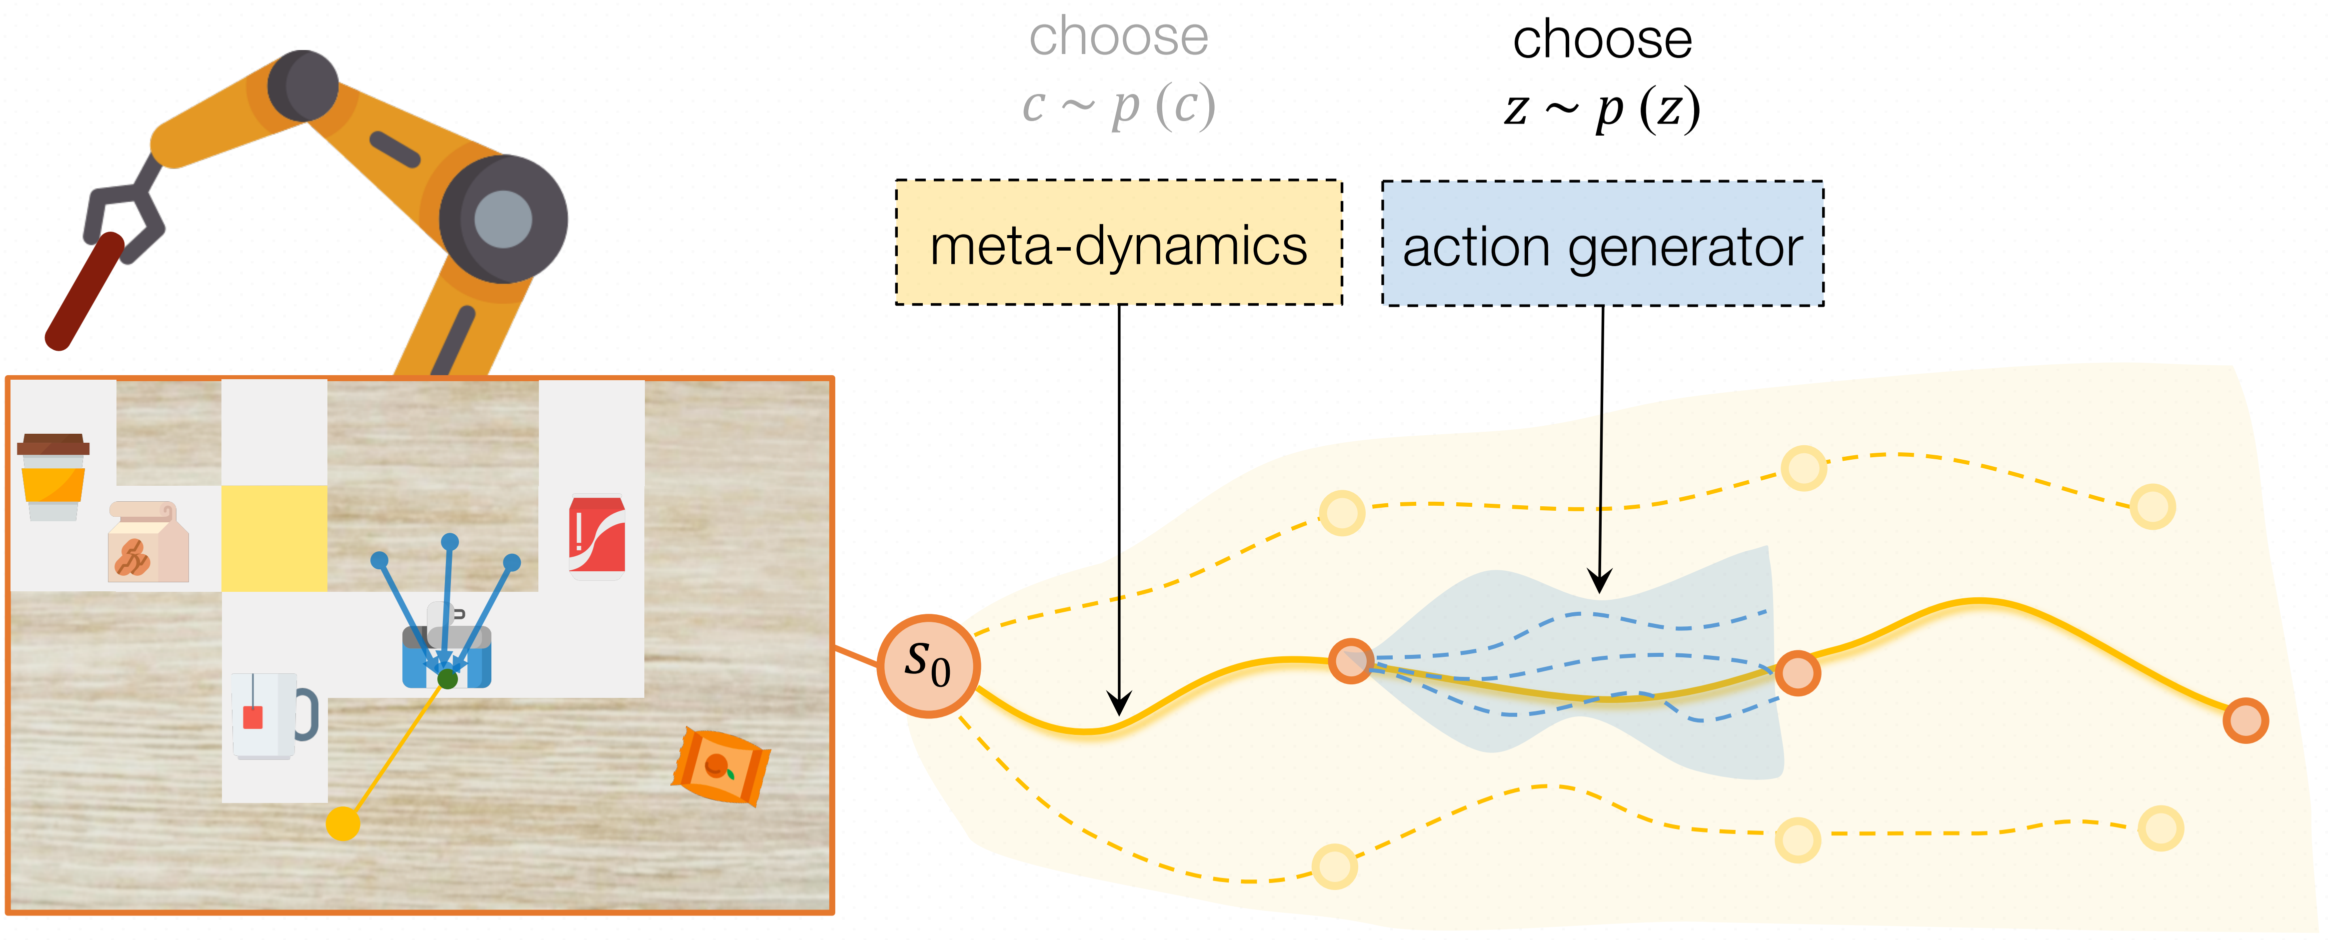

CAVIN: Hierarchical planning in learned latent spaces

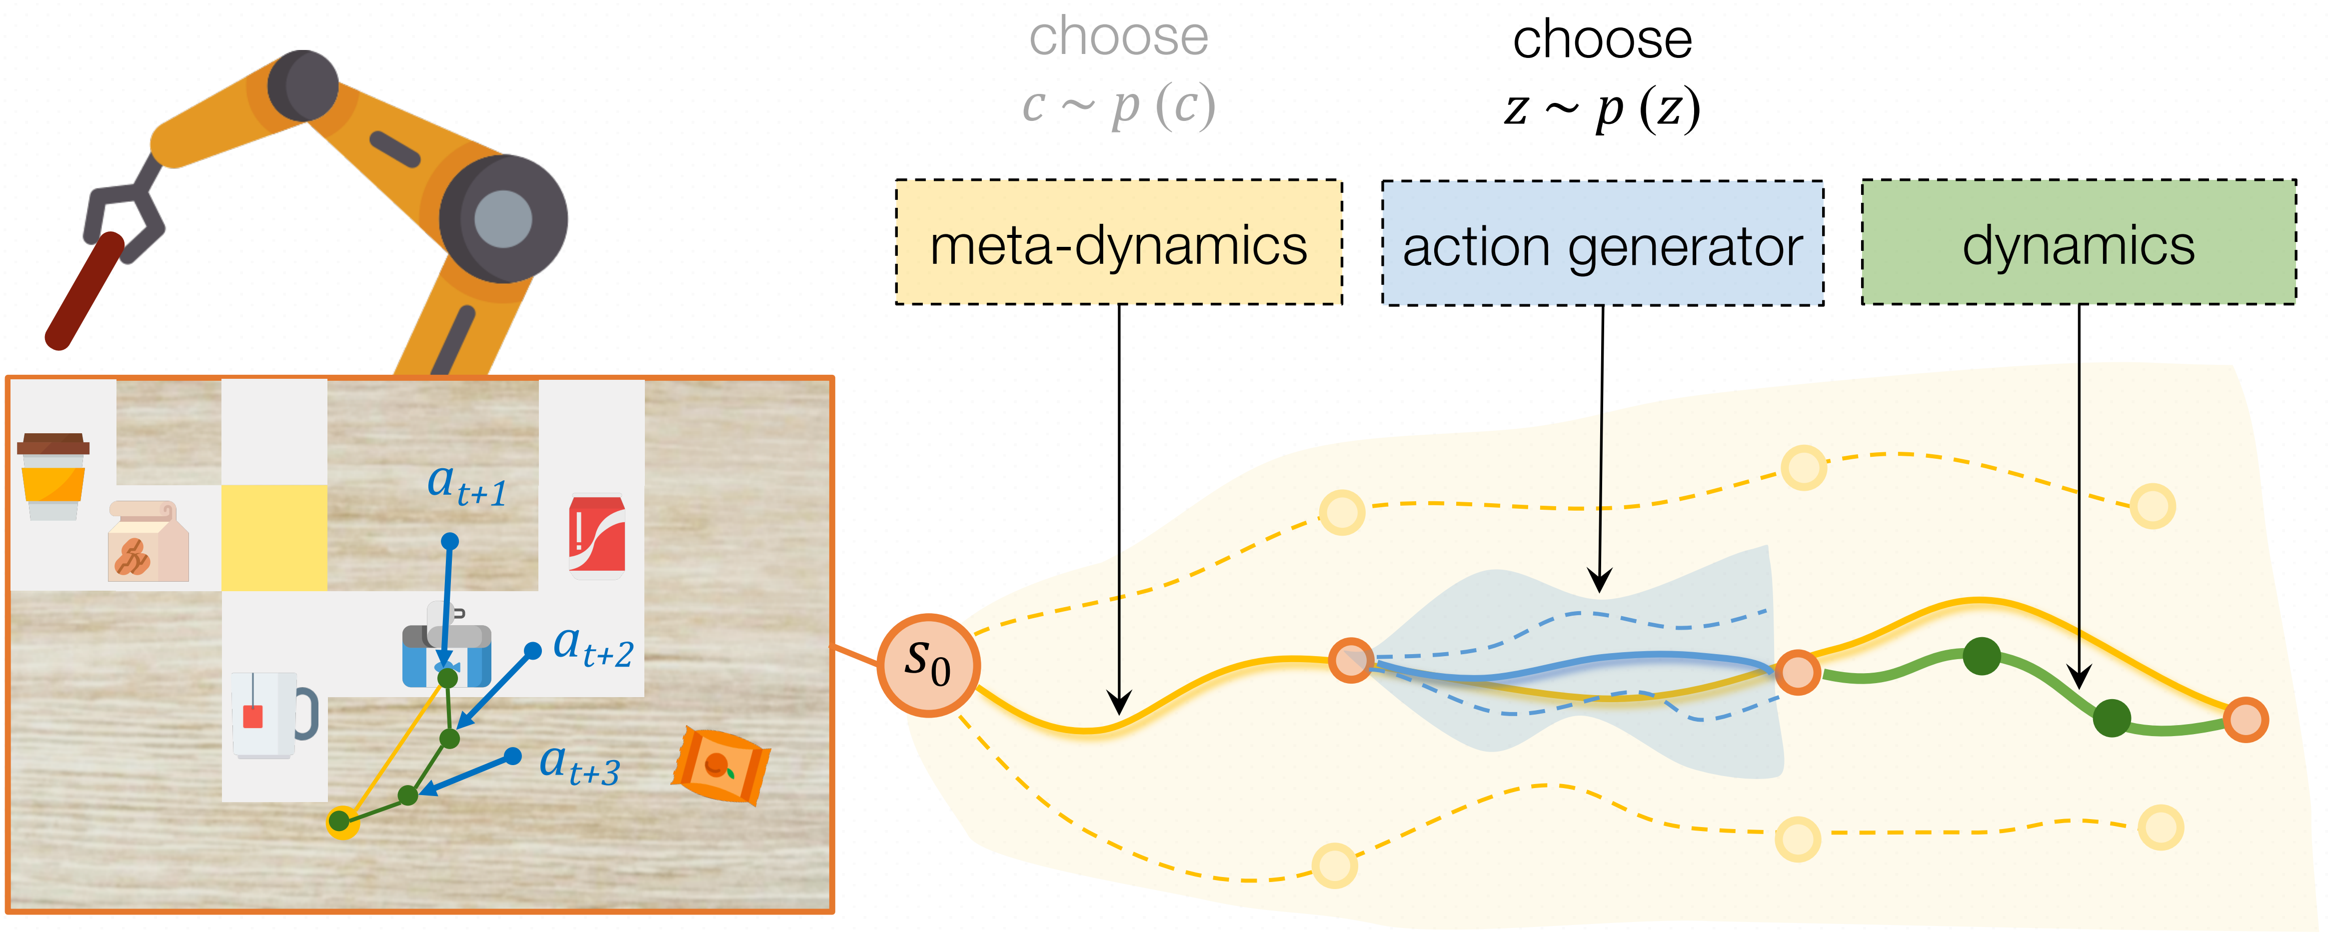

Learning with cascaded variational inference

task-agnostic interaction

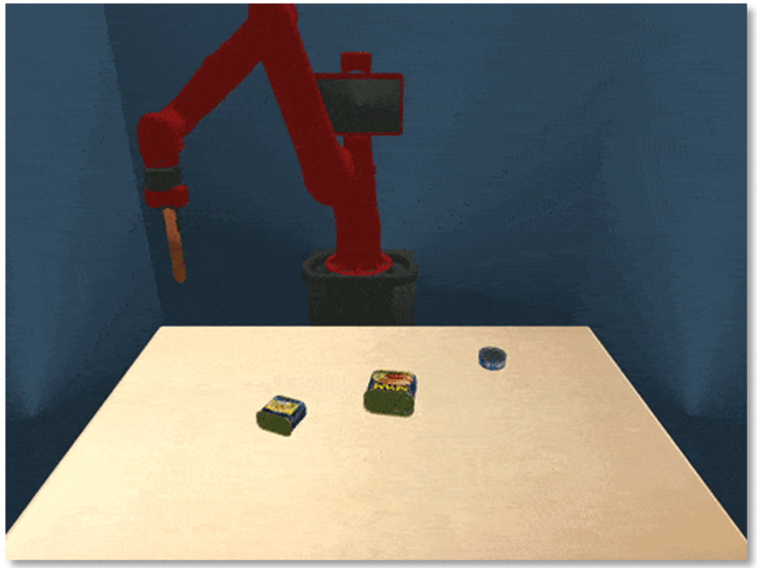

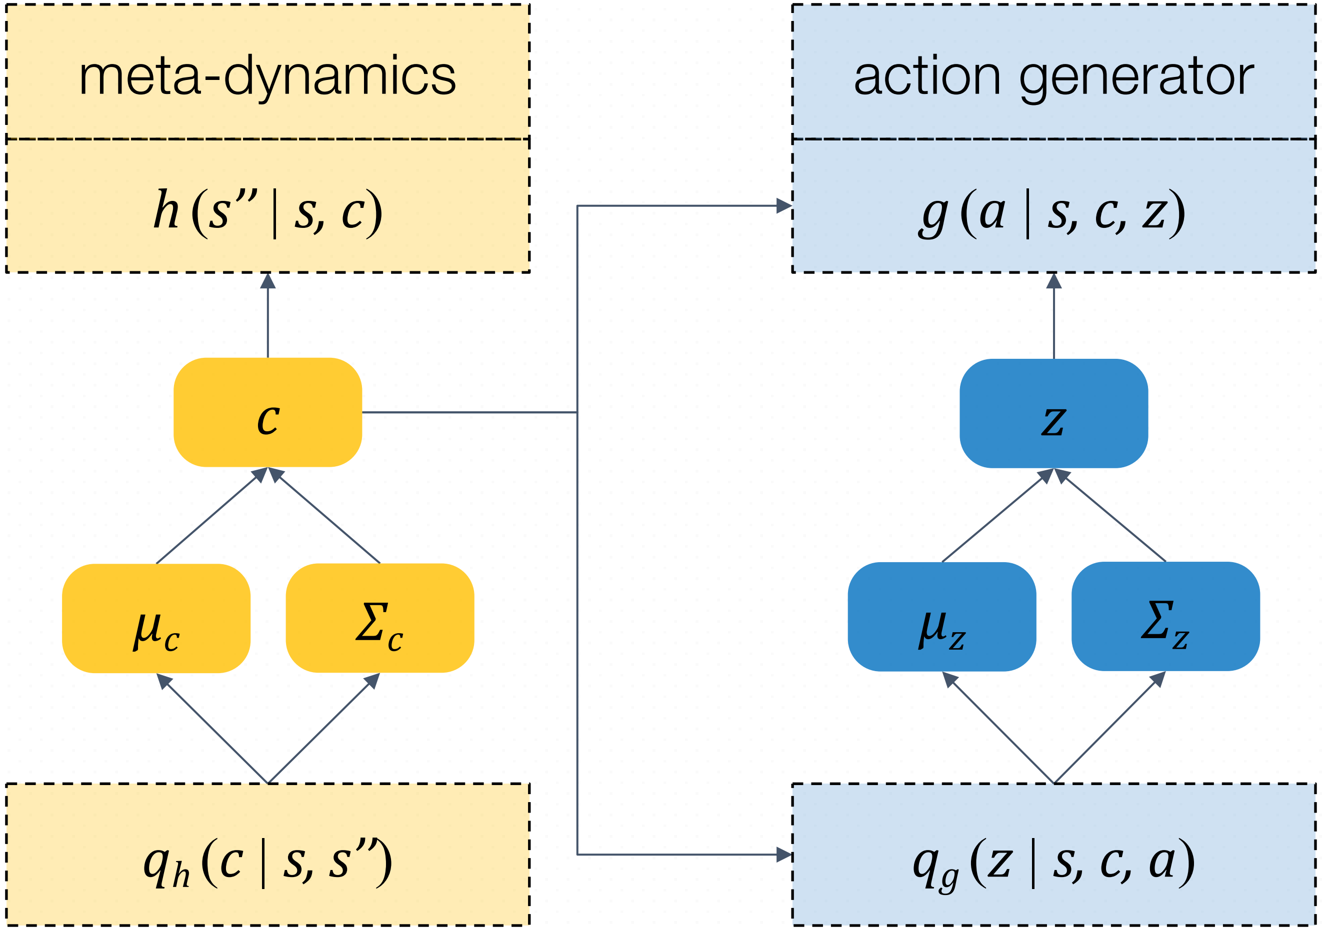

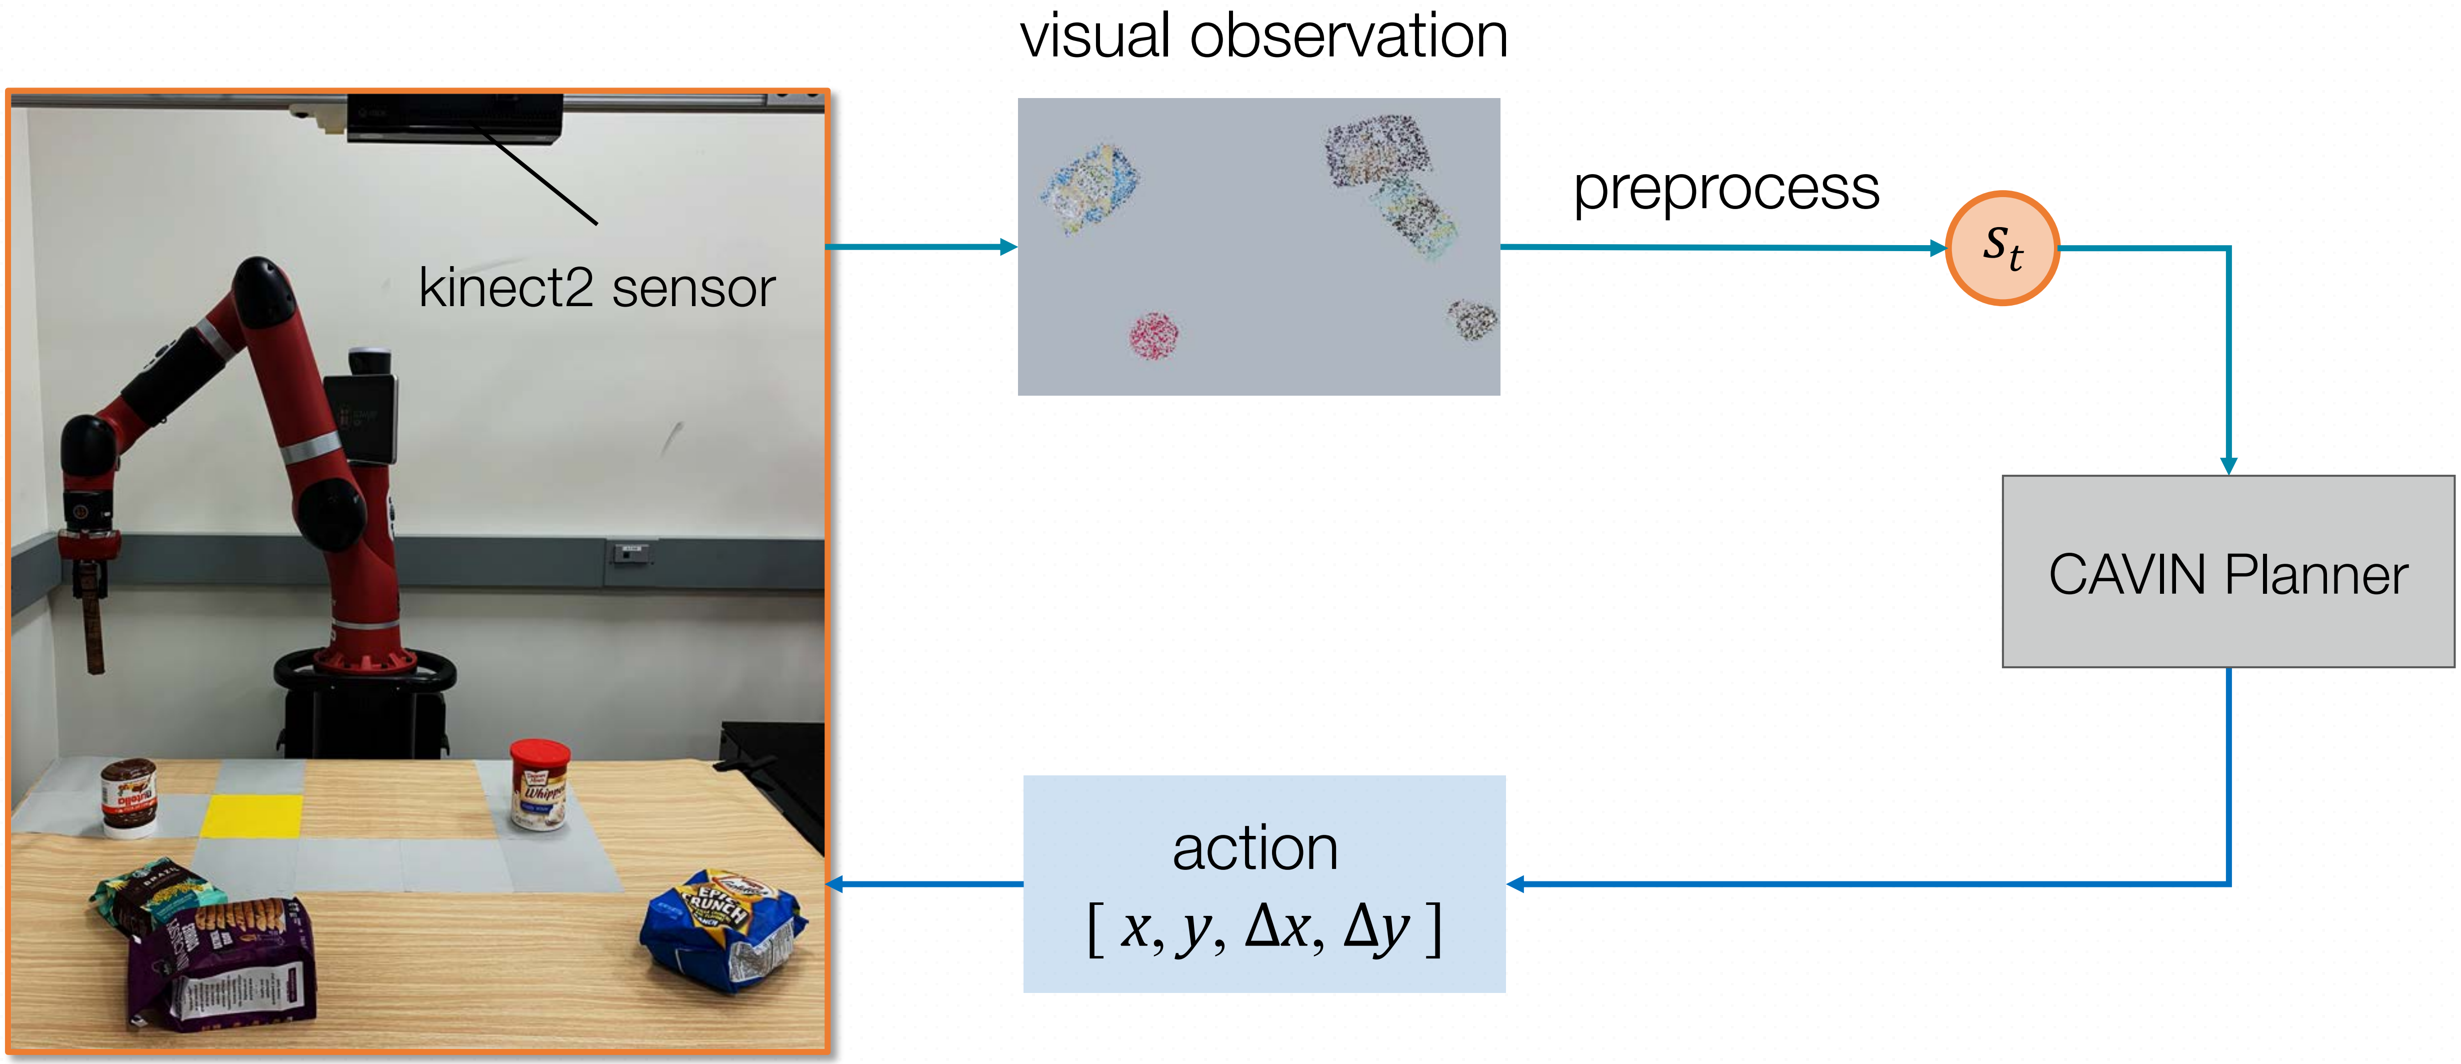

kinect2 sensor
visual observation
preprocess
s_t
CAVIN Planner
action
$[\, x, y, \Delta x, \Delta y \,]$

Tasks

clearing

Clear all objects within the area of blue tiles.

insertion

Move the target to the goal without traversing red tiles.

crossing

Move the target to the goal across grey tiles


Real
Simulated

Quantitative Evaluation

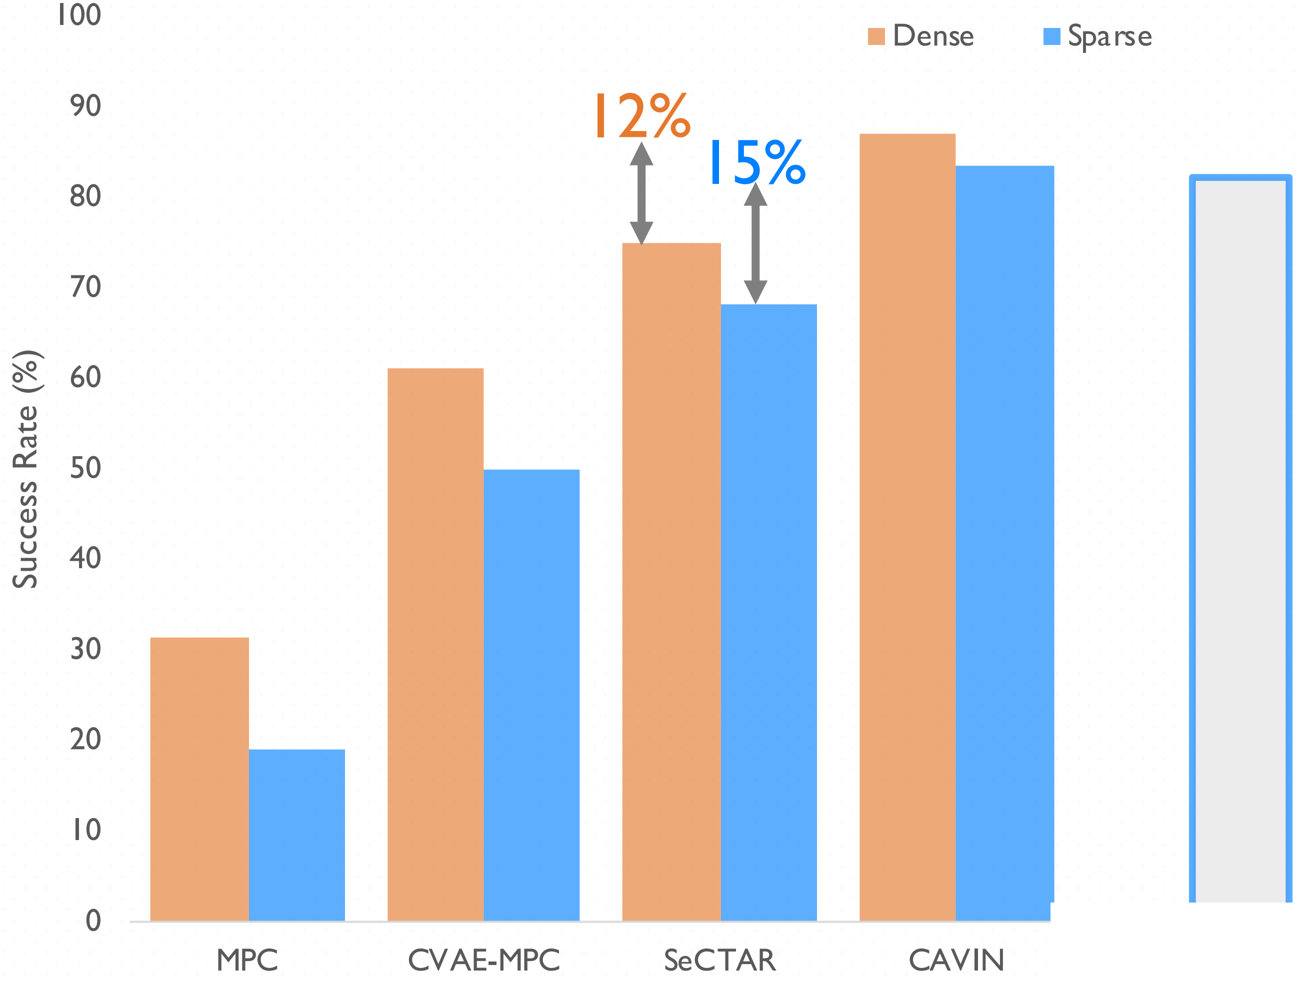

Hierarchical Latent space dyn.

↓

Better performance with sparse reward signal

Averaged over 3 Tasks
with 1000 test instances each

MPC (Guo et al. '14, Agrawal et al. '16, Finn et al. 17); CVAE-MPC (Ichter et al. 18), SeCTAR (Co-Reyes et al '18)

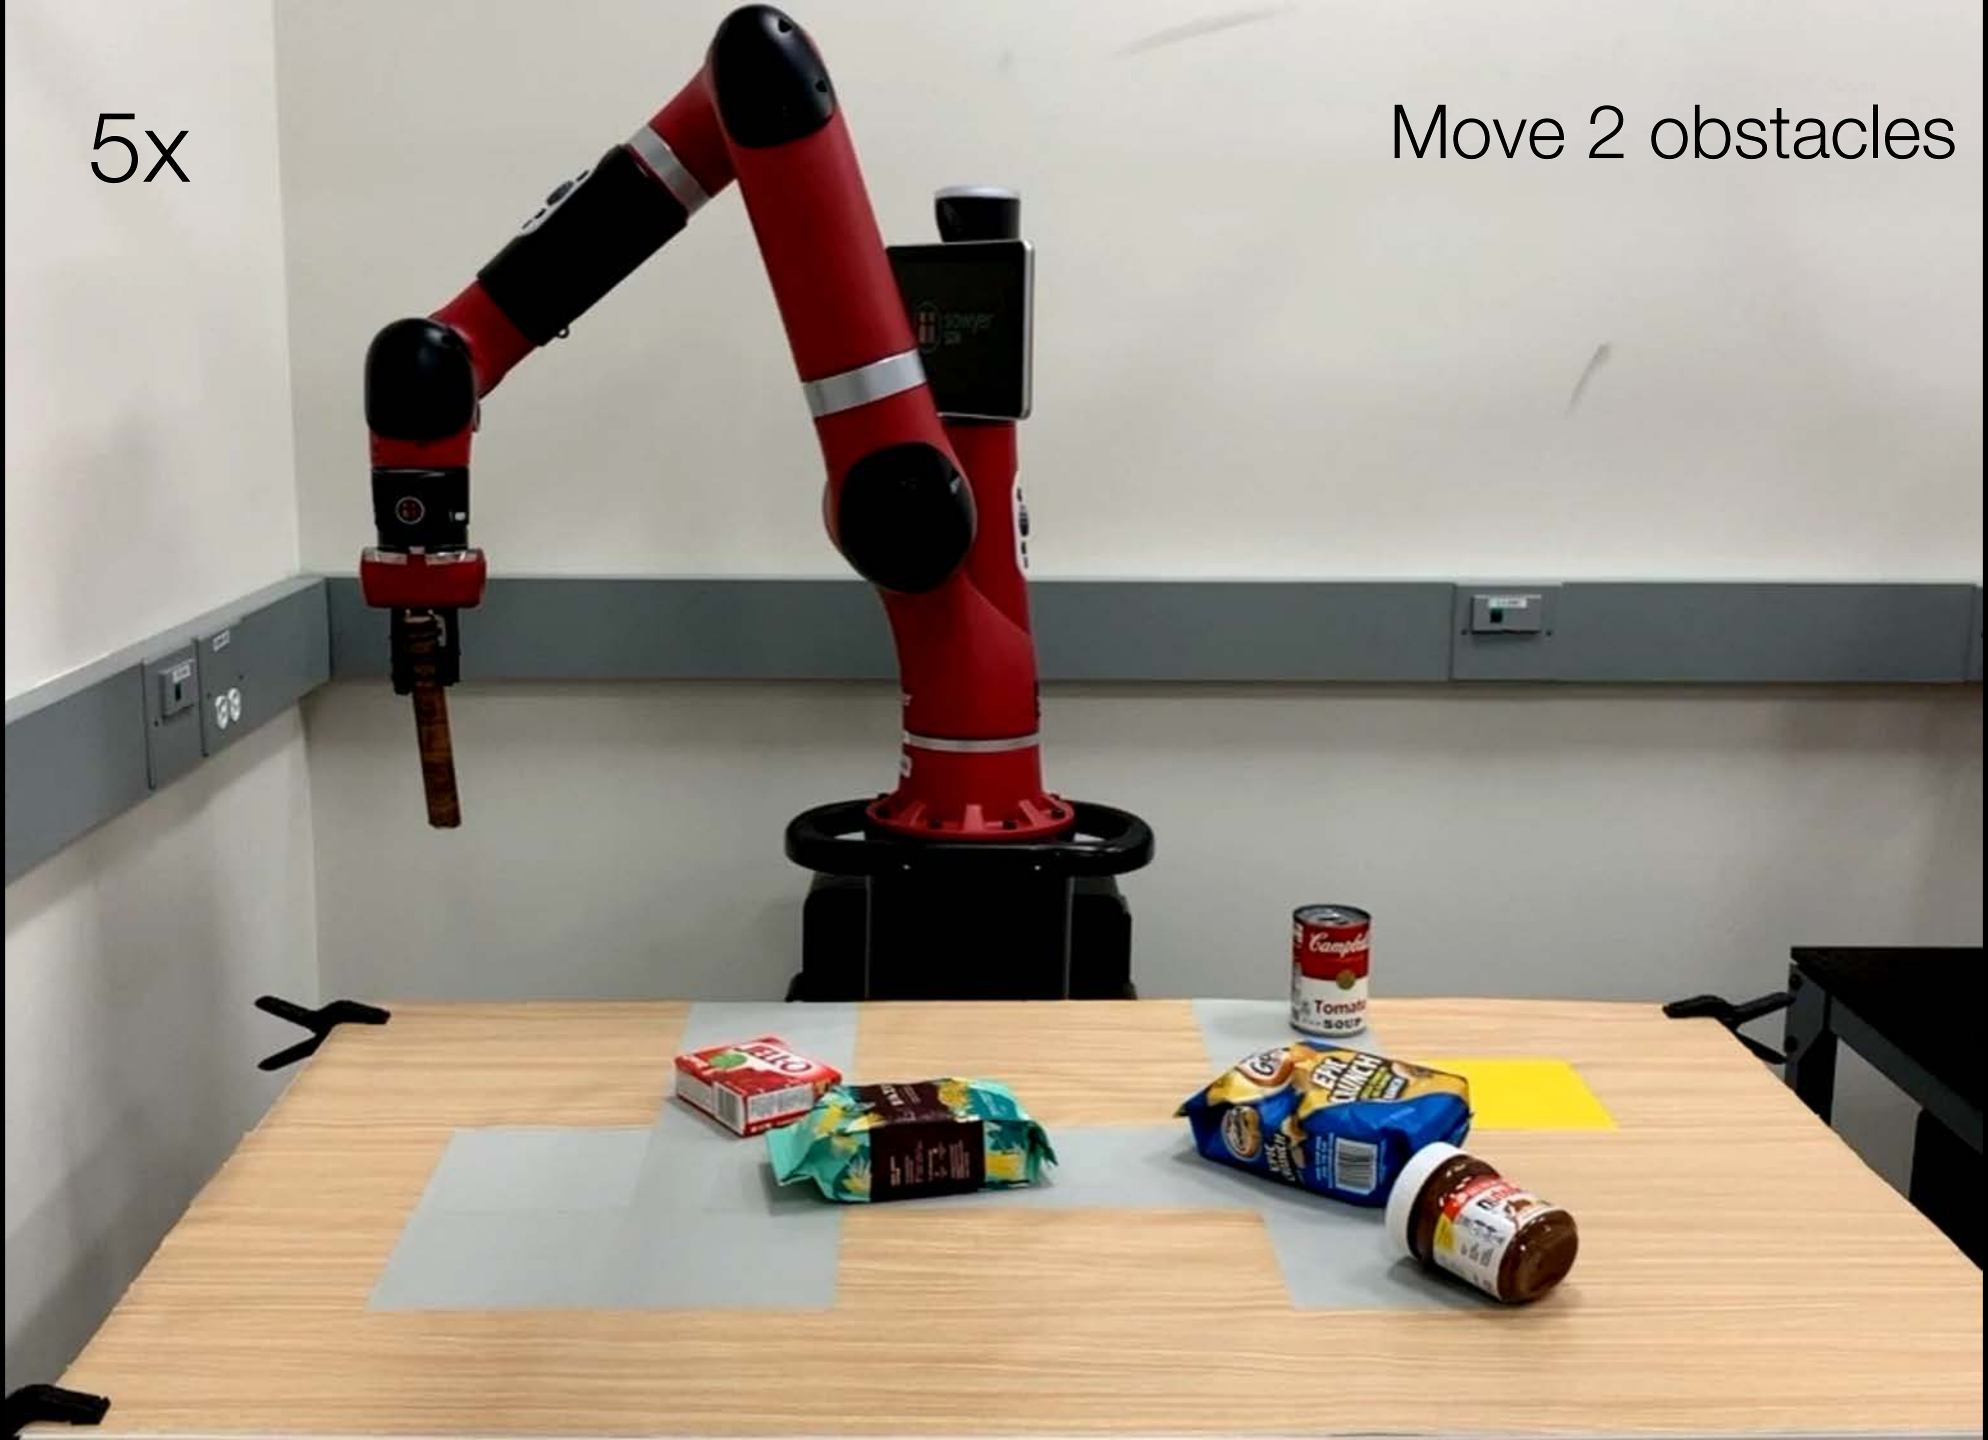
5x
Move 2 obstacles


5x
Get around
nutella

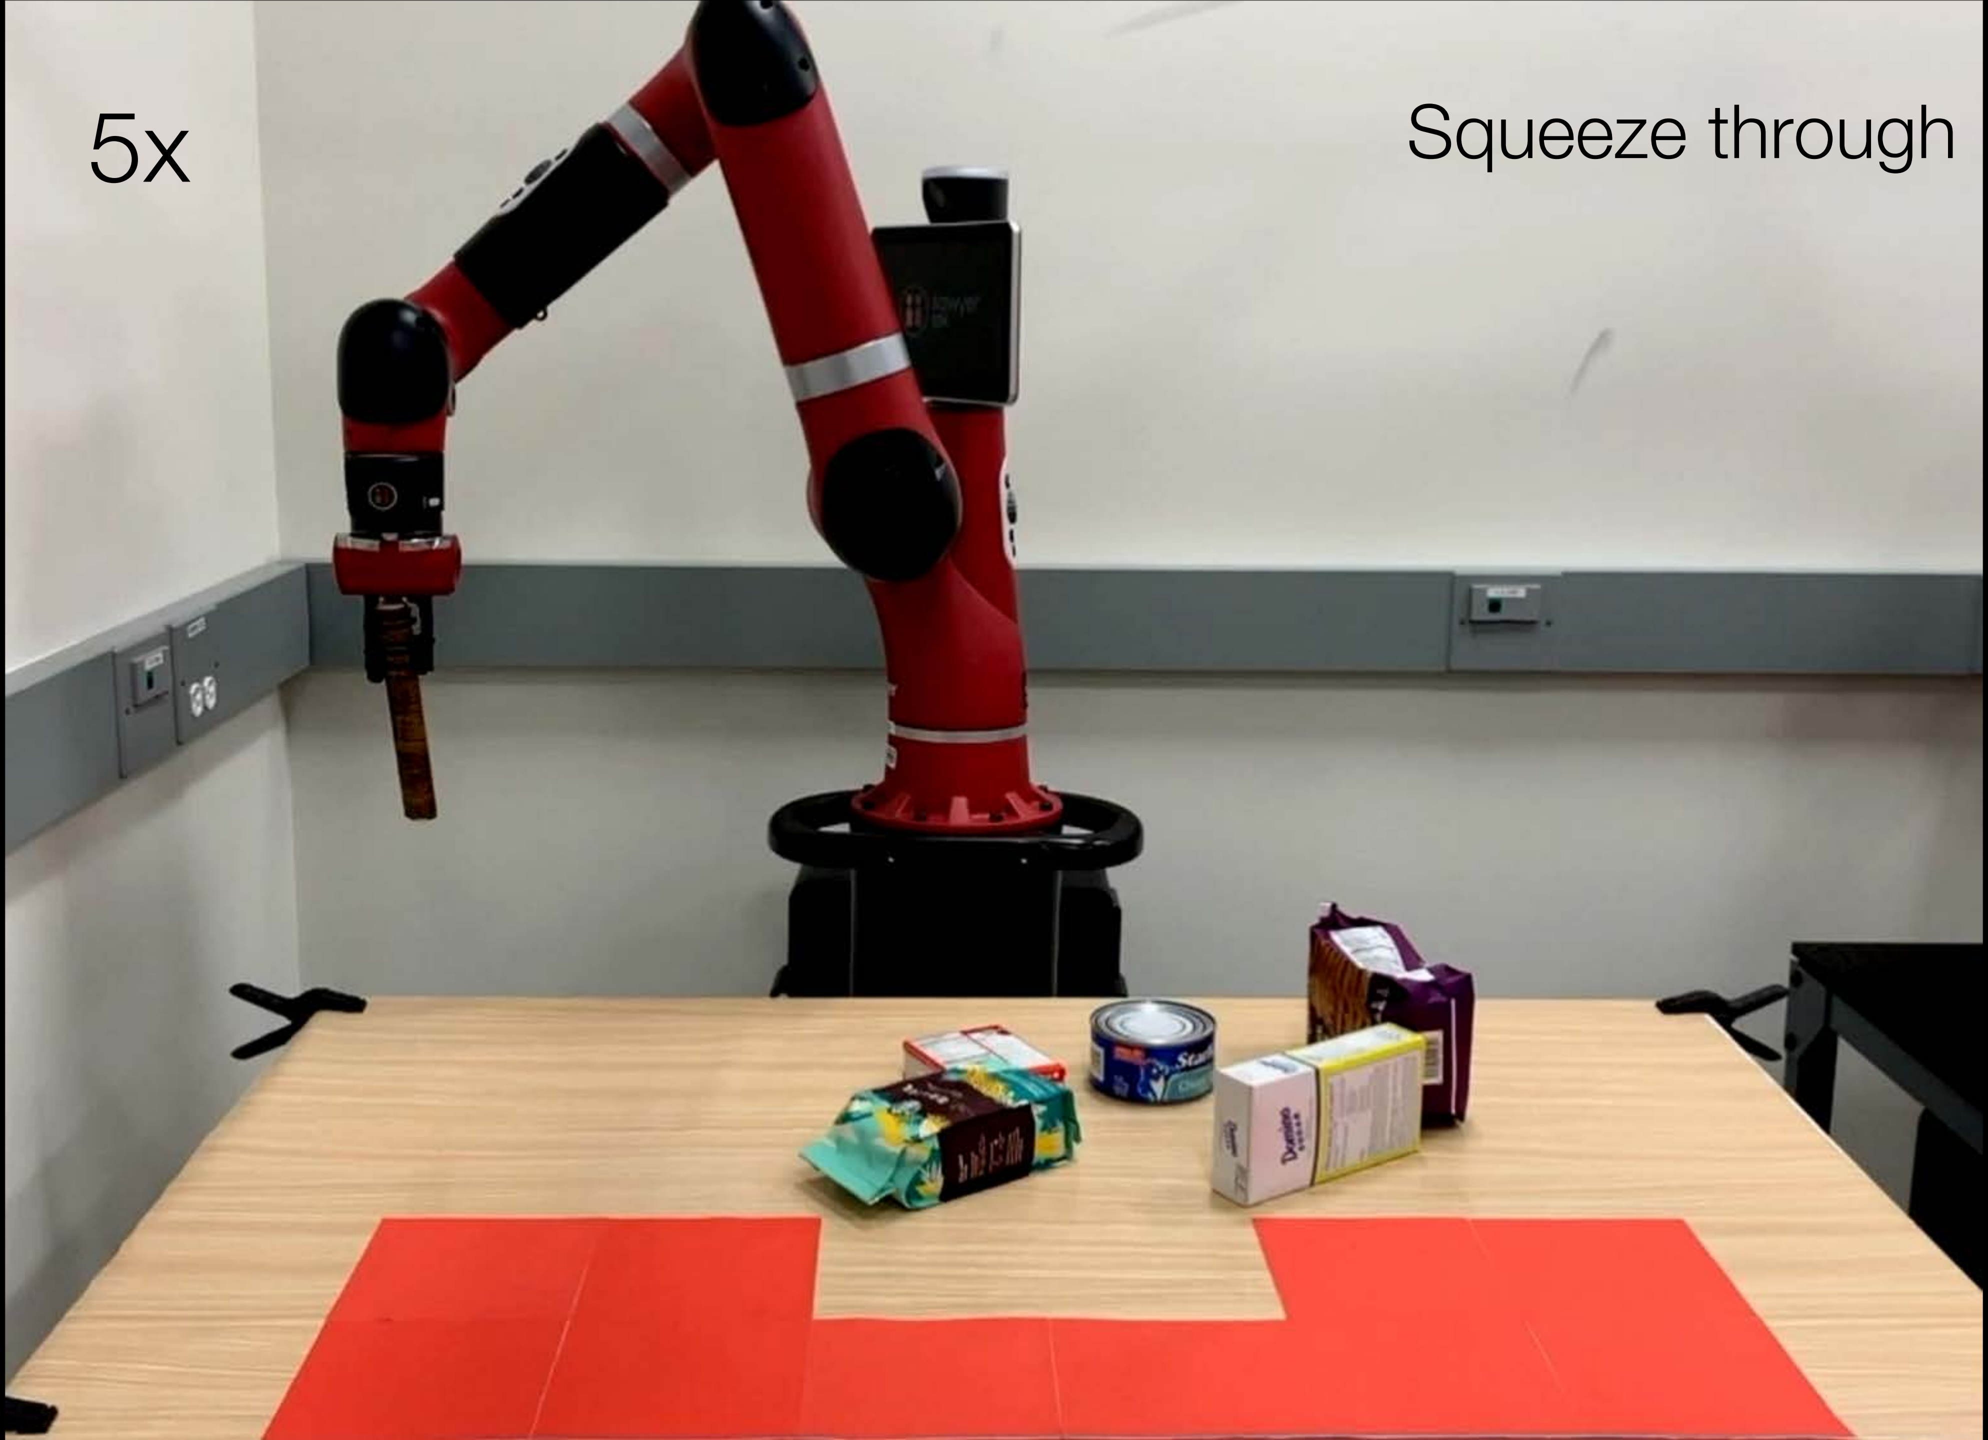
5x
Squeeze through

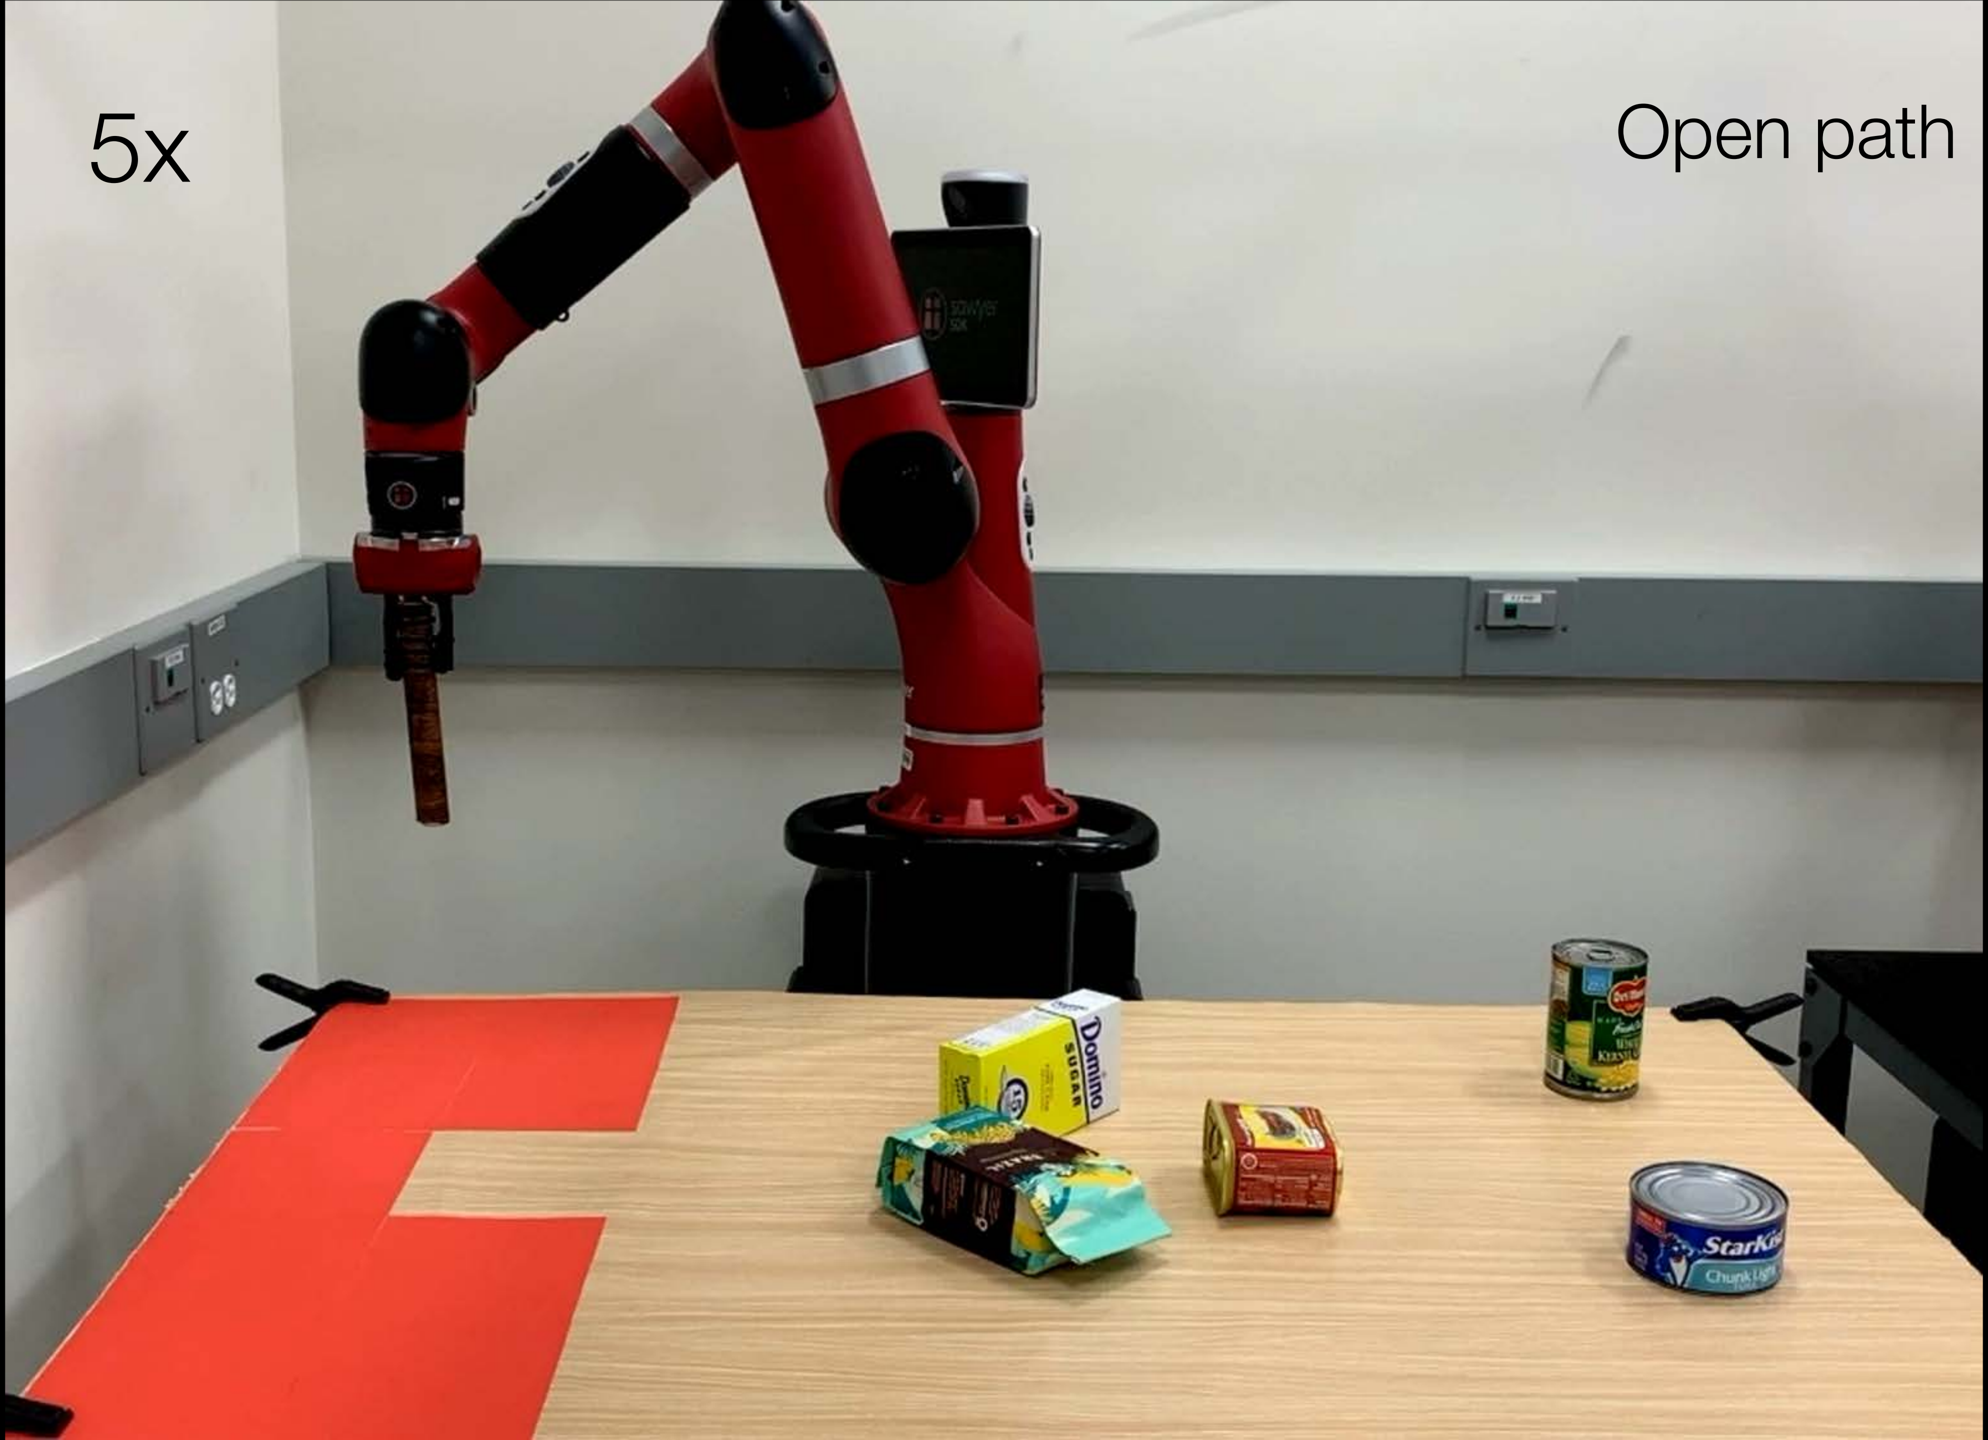
5x
Open path
Domino
SUGAR
StarKist
Chunk Light

Compositional Planning

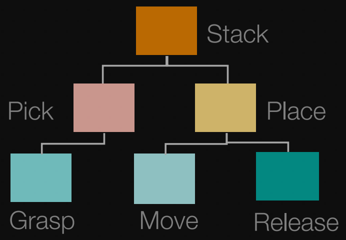

Object Arrangement

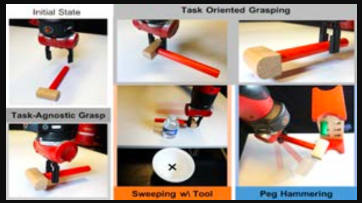

RSS 2018,
IJRR 2019

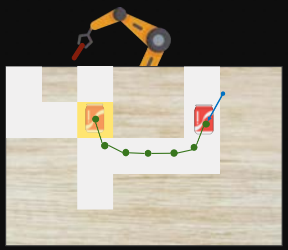

CoRL 2019 (oral)

Self-Supervision and Structured Latent Variable Models
lead to good representations that generalize

Generalizable Autonomy in Robot Manipulation

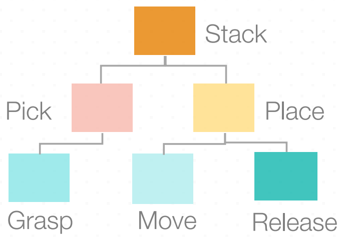

Object Arrangement

ICRA 2018

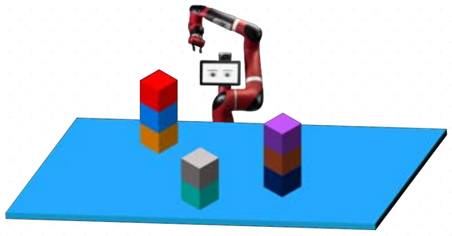

CVPR 2019 (oral)

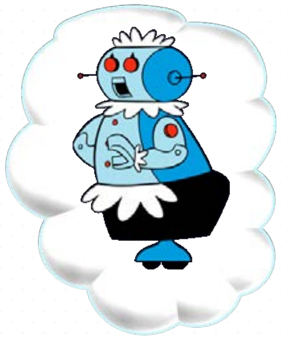

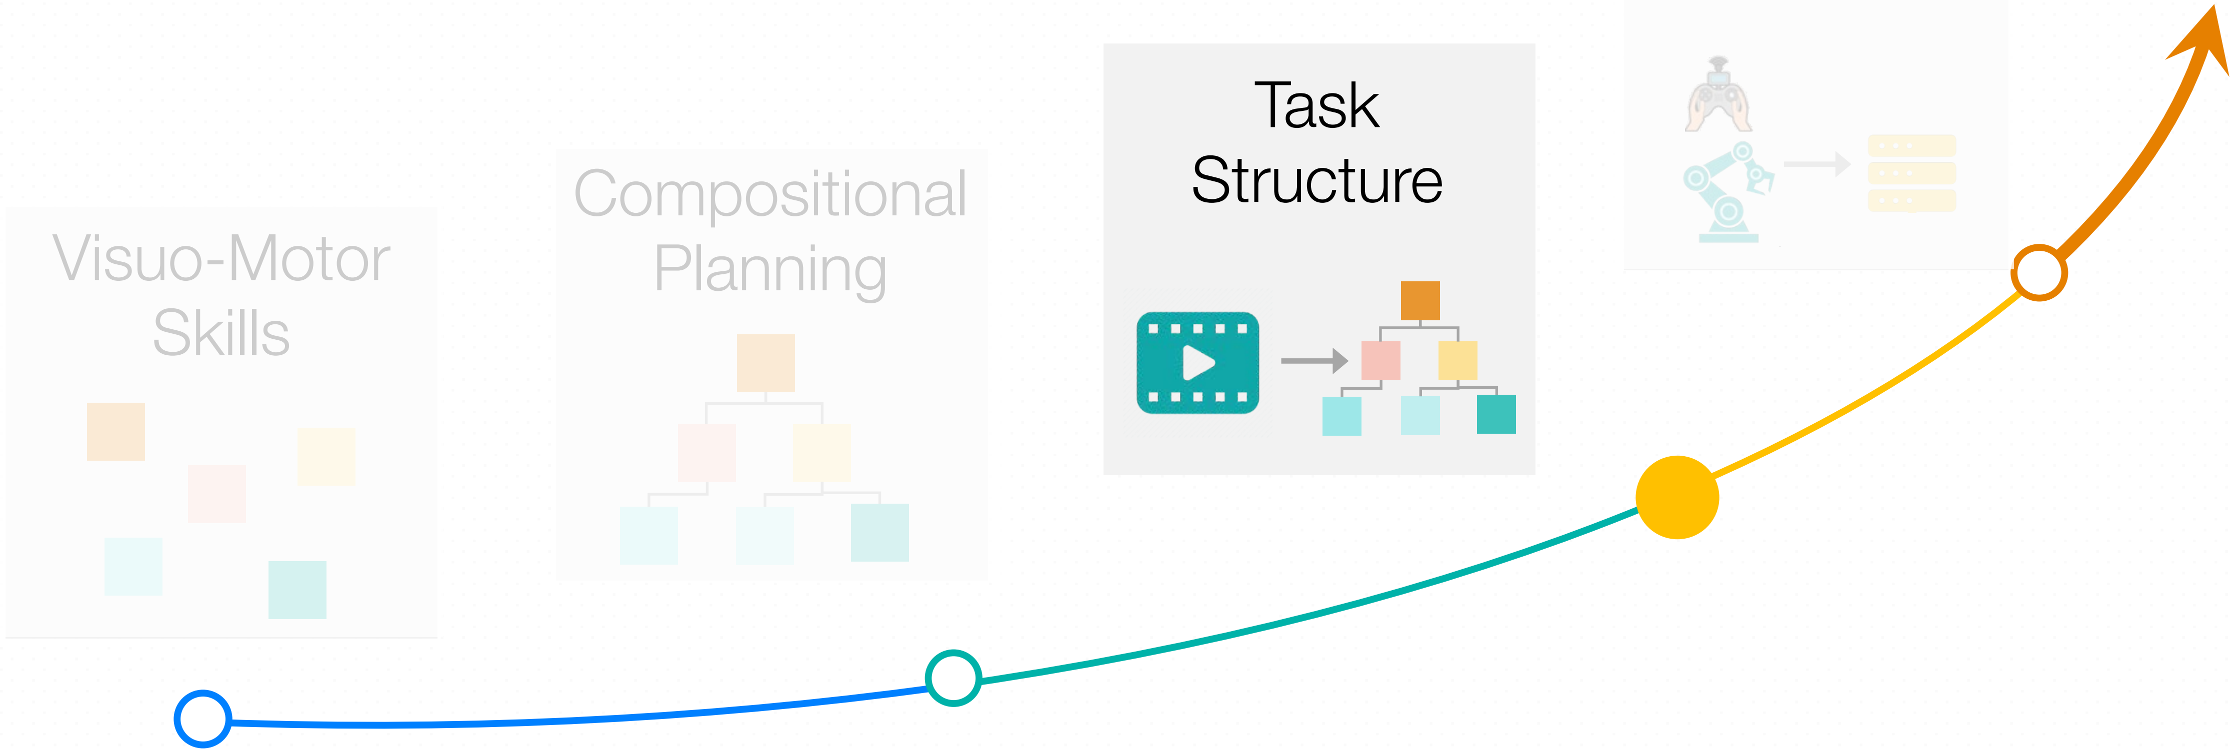

Complex Task Structure

Visuo-Motor Skills

Complex Task Structure

Visuo-Motor Skills

Complex Task Structure

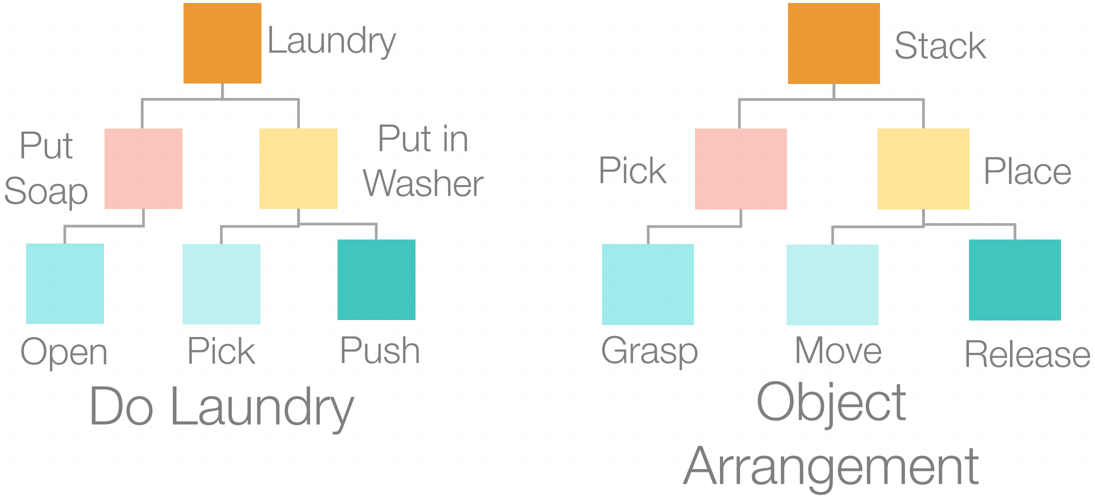

Compositional Planning: Current Paradigm

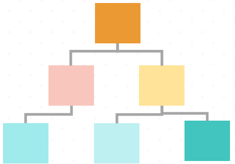

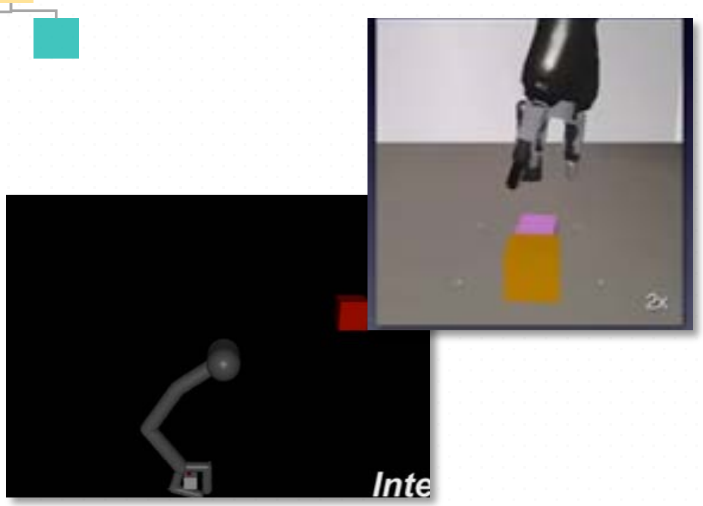

Schulman et al 2013

Kober and Peters '09

Pastor et. al. '09

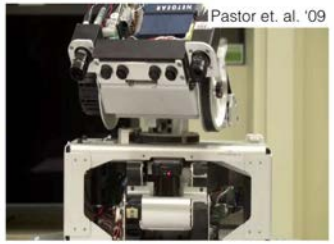

Desired

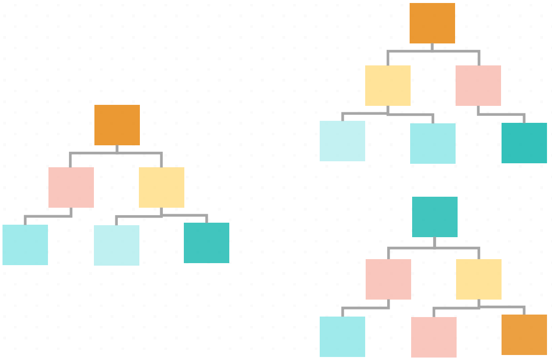

Train ≠ Test

Reinforcement Learning

- Sample Inefficient
- Multi-step Structured Tasks
- Needs non-trivial Reward Shaping

Imitation Learning

- Task Segmentation is non-trivial
- Multi-modality of Search Space
- Fixed Permutation of Primitives

Meta Imitation Learning

- New Task Structures
- Few-Shot performance
- Input State as Video

RL: [Schaal 1997], [Chebotar et al., '17], [Yahya et al., '16], [James et al., '17], [Popov et al., '17], [Zhu et al. 18], [Hausman et al. 18]
Imitation: [Calinon et al 2008], [Argall et al 2009], [Kober, Peters, et al. 09], [Pastor et al, 09], [Schulman et al. 2013], [Kroemer et al, 15], [Garg et al 2017]

Compositional Planning: Challenge

Task Domain

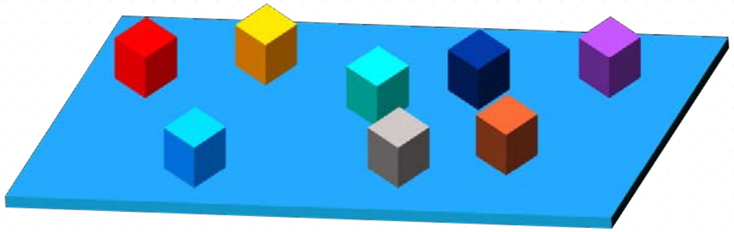

Instructional Demos

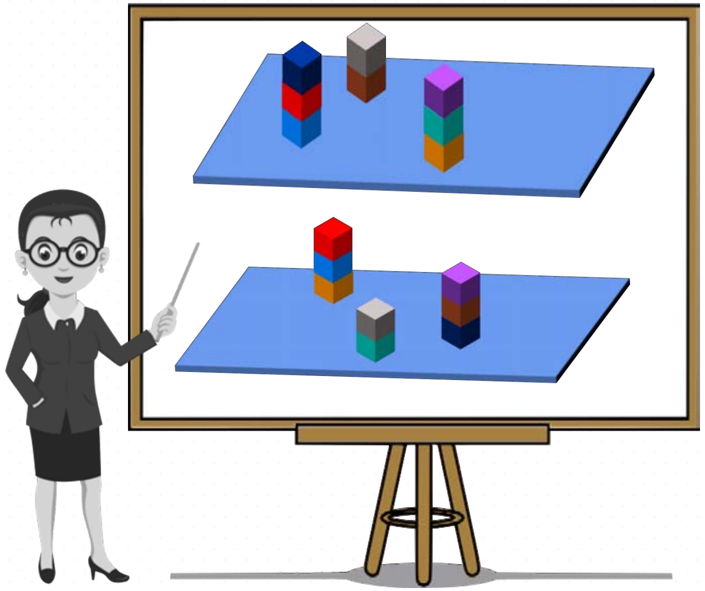

I. Learn Multiple Tasks in the Same Domain

Training Tasks

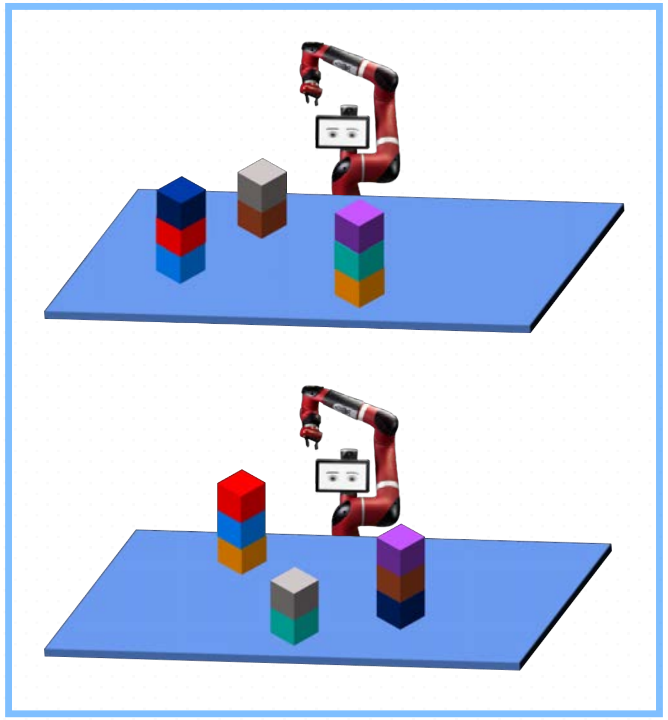

II. Generalize to New Tasks with a Single Demo

Test Task

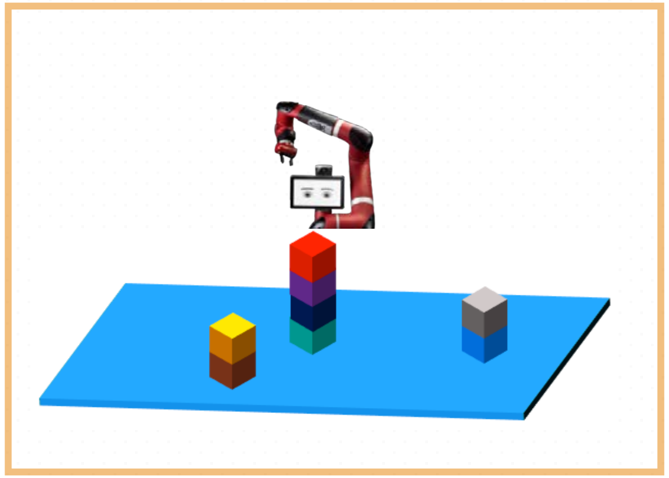

Compositional Planning

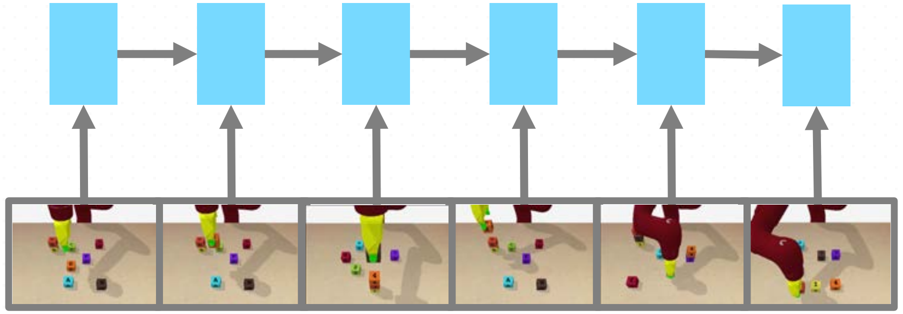

[Duan et al. 17; Finn et al. 2017;
Wang et al. 2017; Yu et al. 2018]

Models input demonstration
as a flat sequence

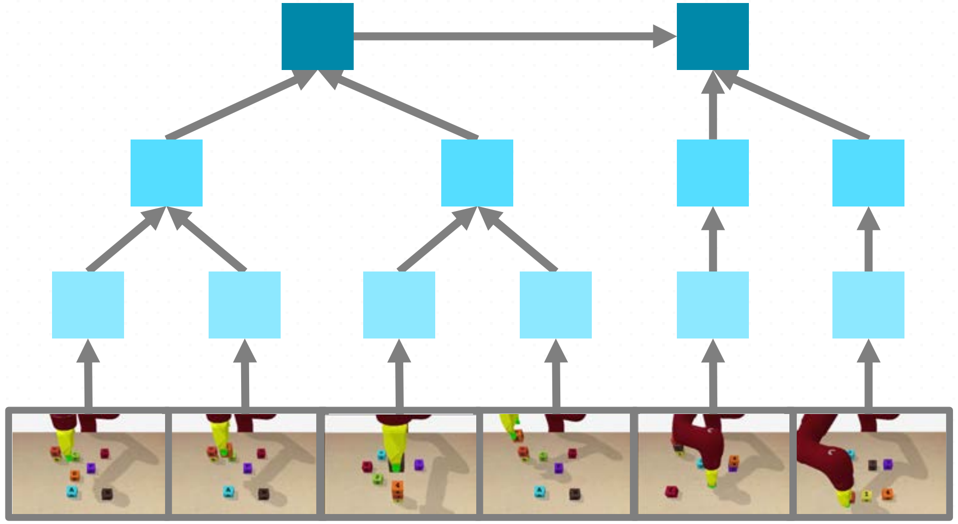

Our Method
[ICRA'18], [CVPR'19], [IROS'19]

Models input demonstration
as a Compositional Hierarchy

One Shot Imitation Learning from Videos

Compositional Planning: Task Programming

```
Block Stacking (...):
  while (done):
    pick_and_place (RED, BLUE):
      pick (RED):
        move_to (RED)
        Grasp (RED)
      <end> Pop
      place(BLUE):
        move_to (BLUE)
        Release (RED)
      <end> Pop
    <end> Pop
```

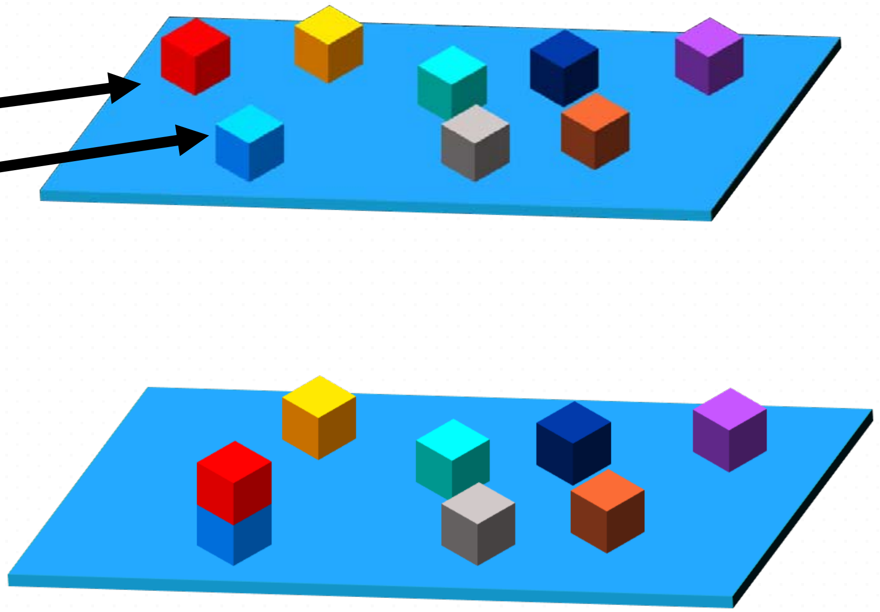

Task 1
Sub-task 1
Move Red-block on top of Blue

[ICRA'18, CVPR '19]

Compositional Planning: Task Programming

```
Block Stacking (...):                    Program 1
```

```
Block Stacking (...):                    Program 2
   while (done):
      pick_and_place (RED, BLUE):
         pick (RED):
            move_to (RED)
            Grasp (RED)
         <end> Pop
         place(BLUE):
            move_to (BLUE)
            Release (RED)
         <end> Pop
      <end> Pop
```

Training Task Structures

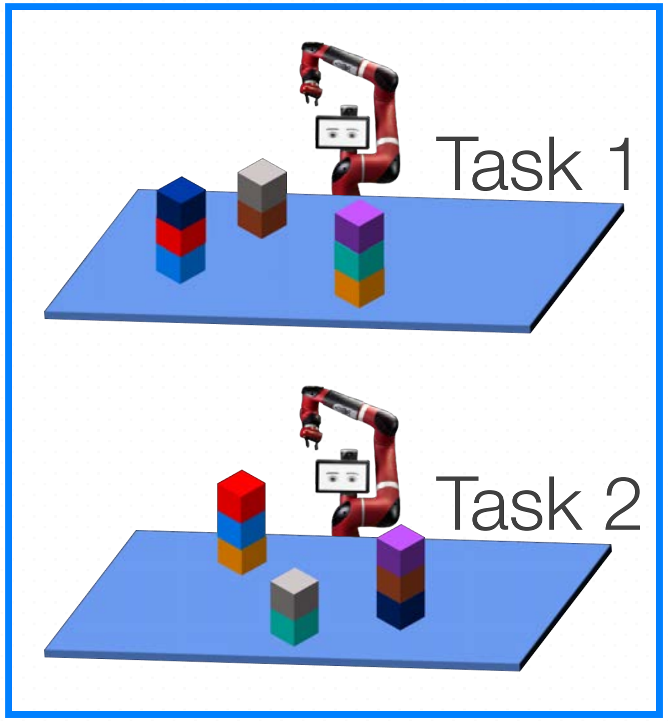

[ICRA'18, CVPR '19]

Neural Task Programming (NTP)

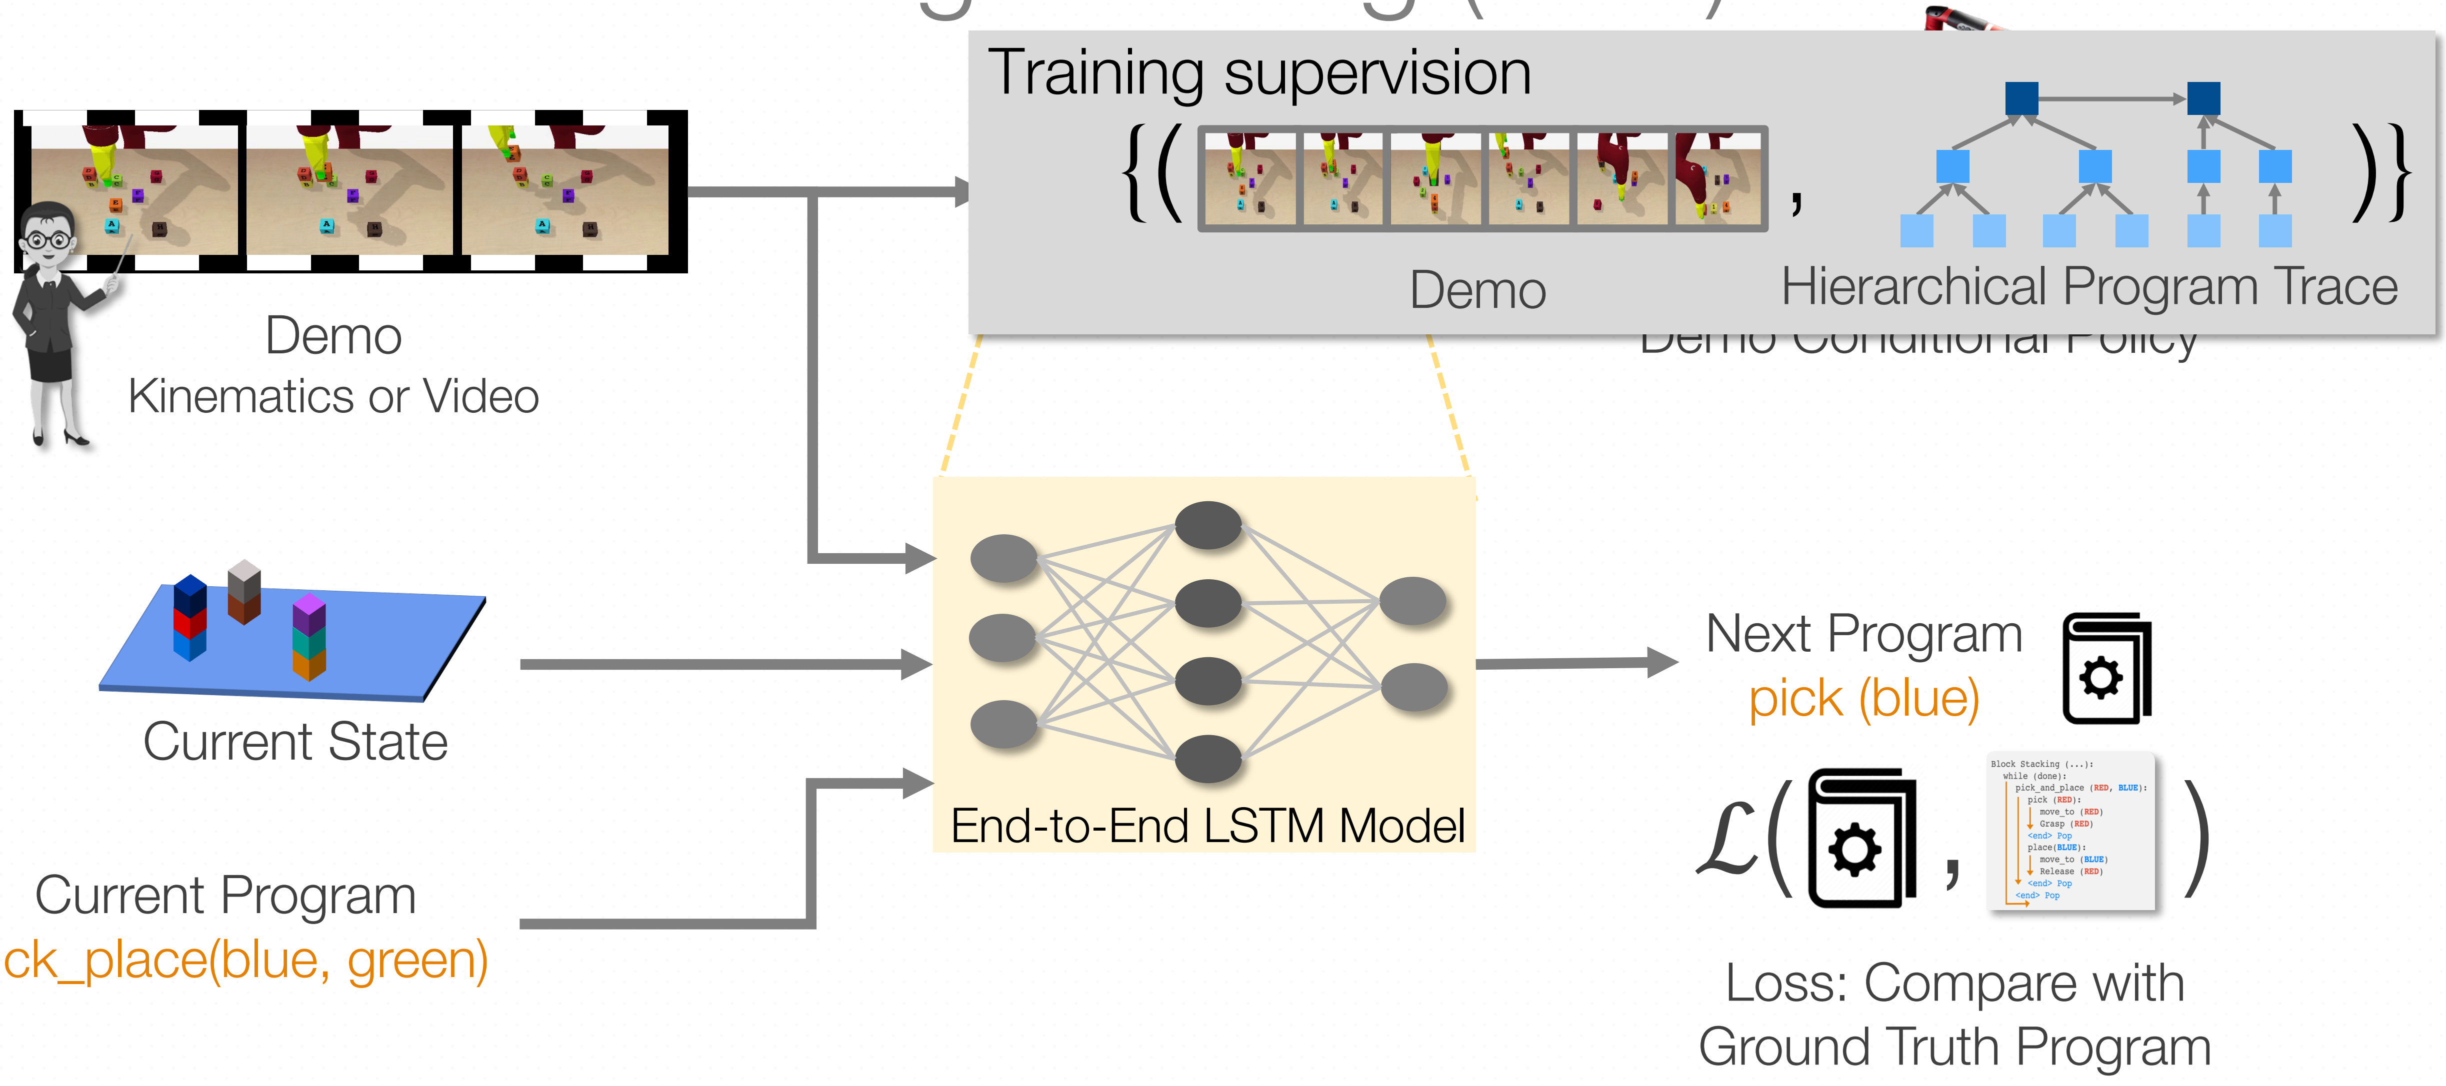

Hierarchical Policy Learning as Program Induction

[ICRA'18, CVPR '19]

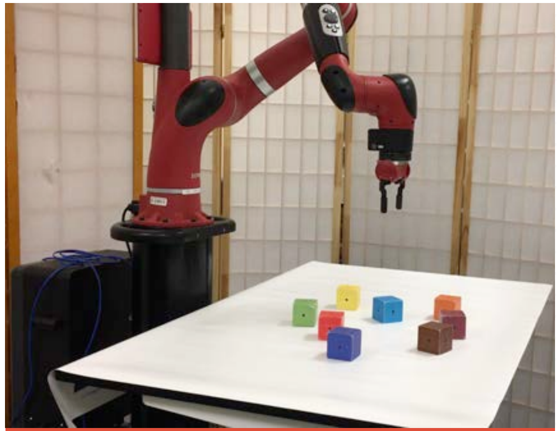

Move_to (Blue)

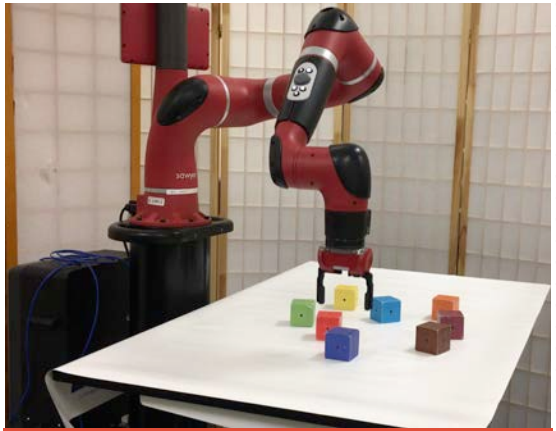

Grasp (Blue)

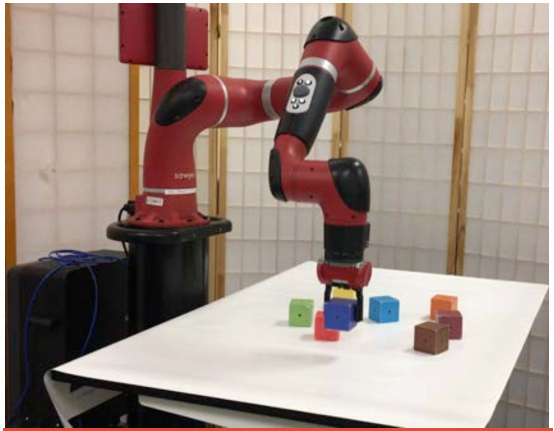

Move_to (Red)

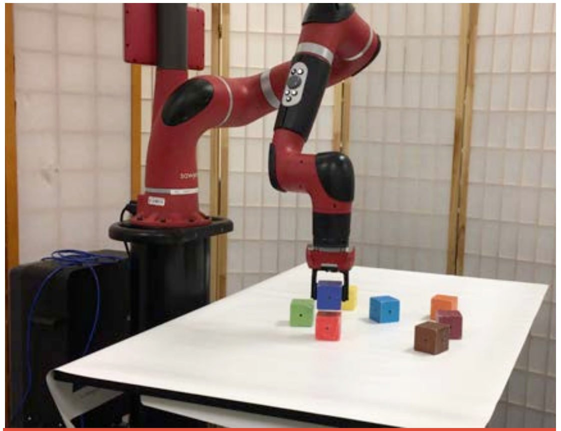

Release( )

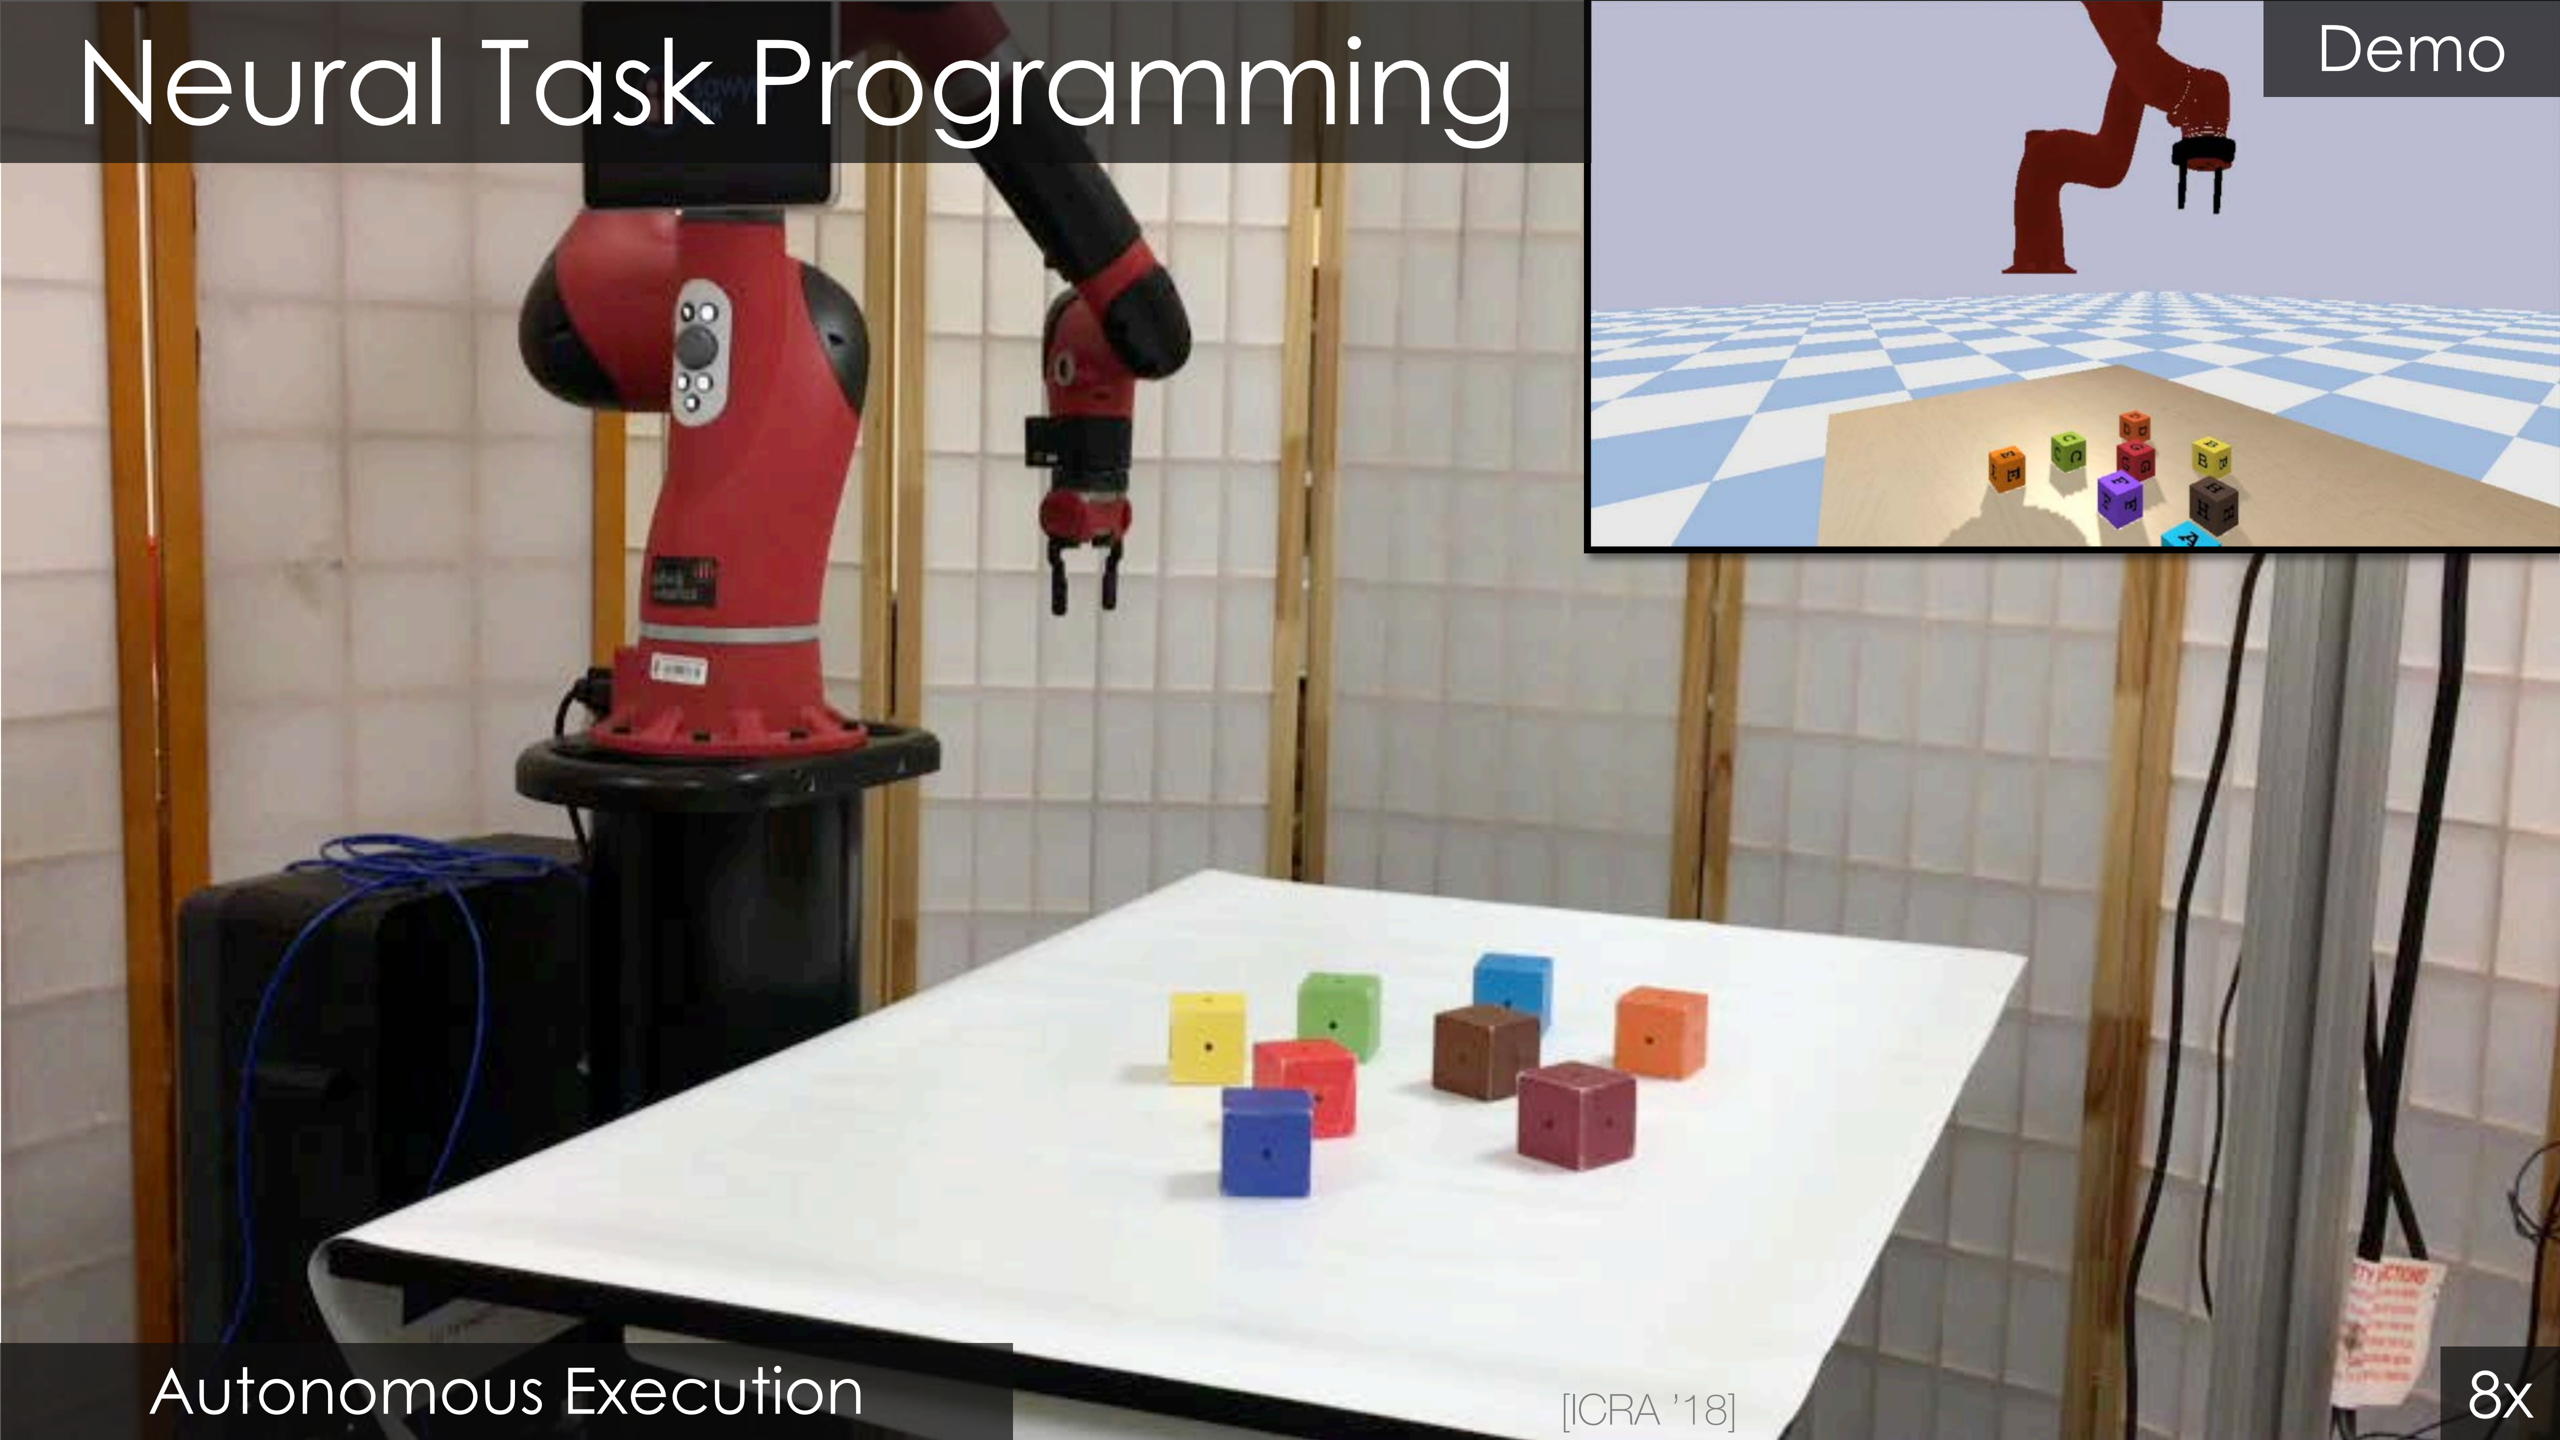
Neural Task Programming
Demo
Autonomous Execution
[ICRA '18]
8x

Recovery from Intermediate Failures

Output is not an Open Loop State Machine

Closed Loop Feedback Policy

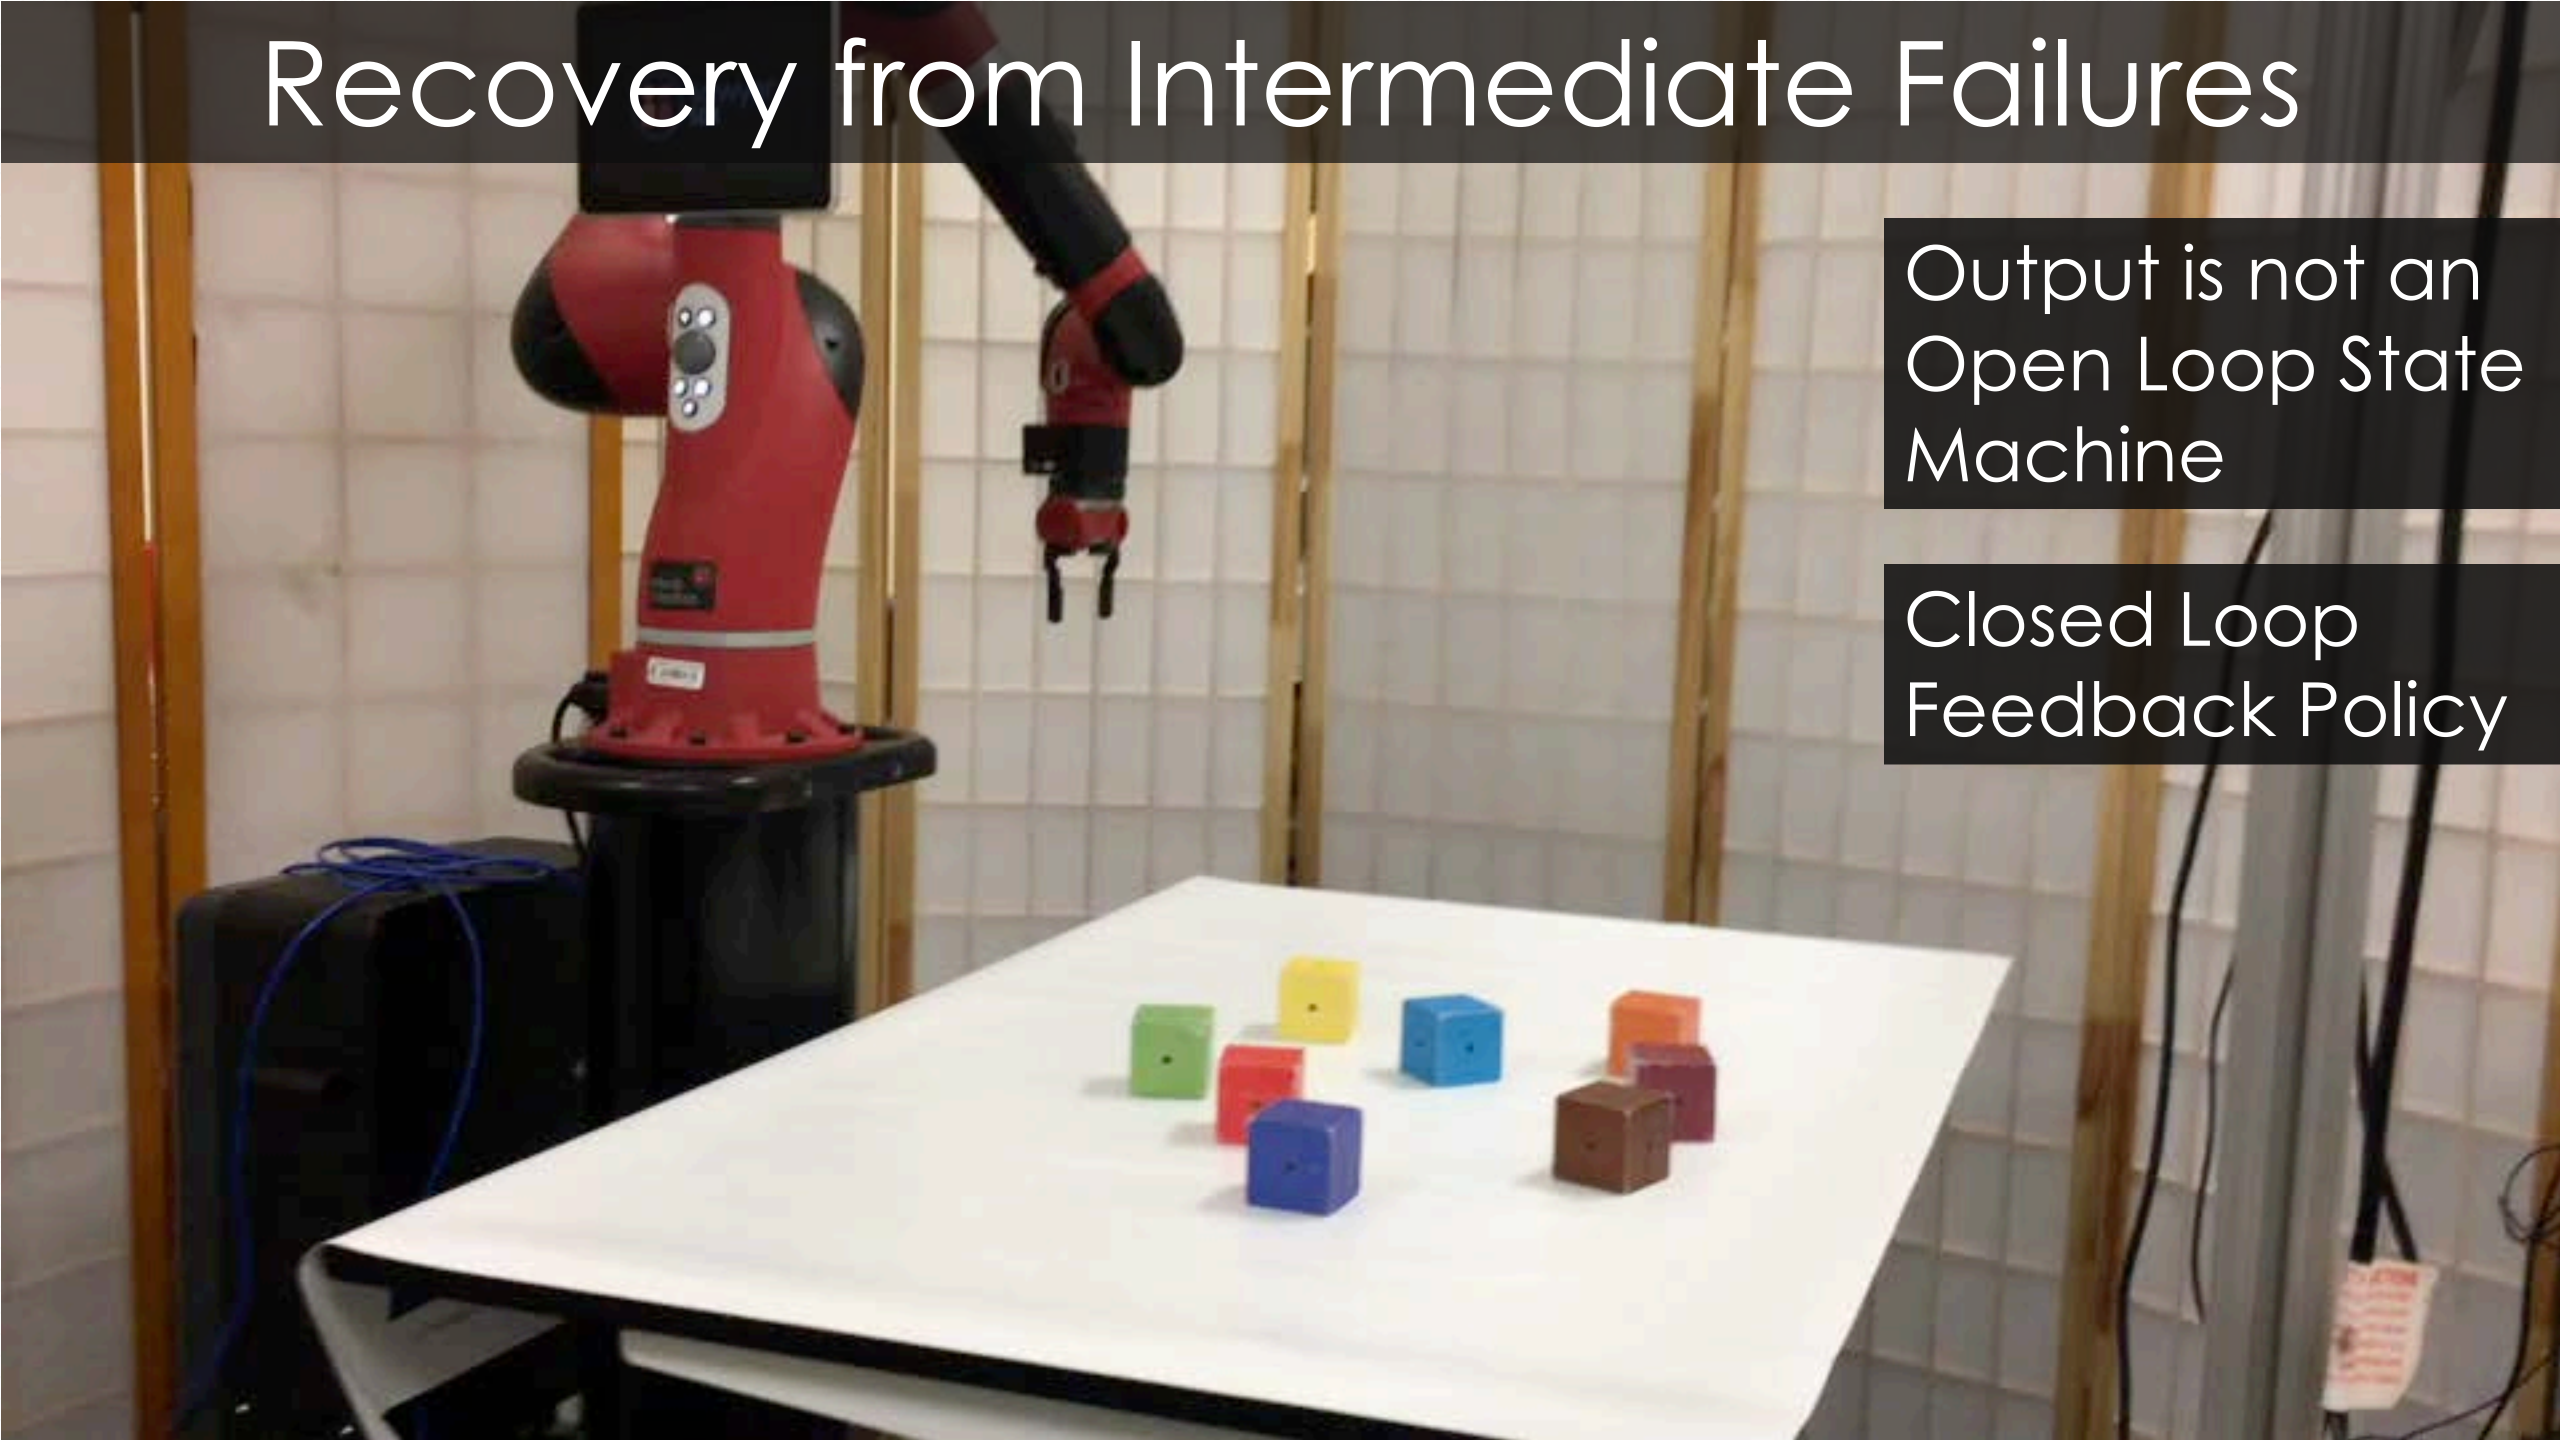

Neural Task Programming
Results

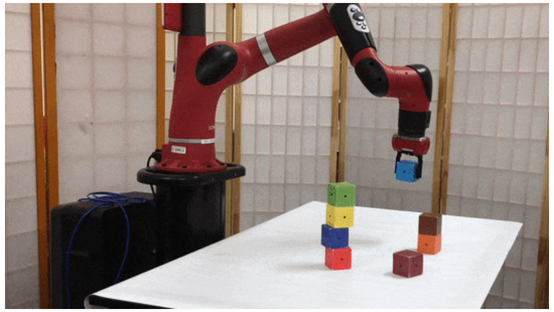

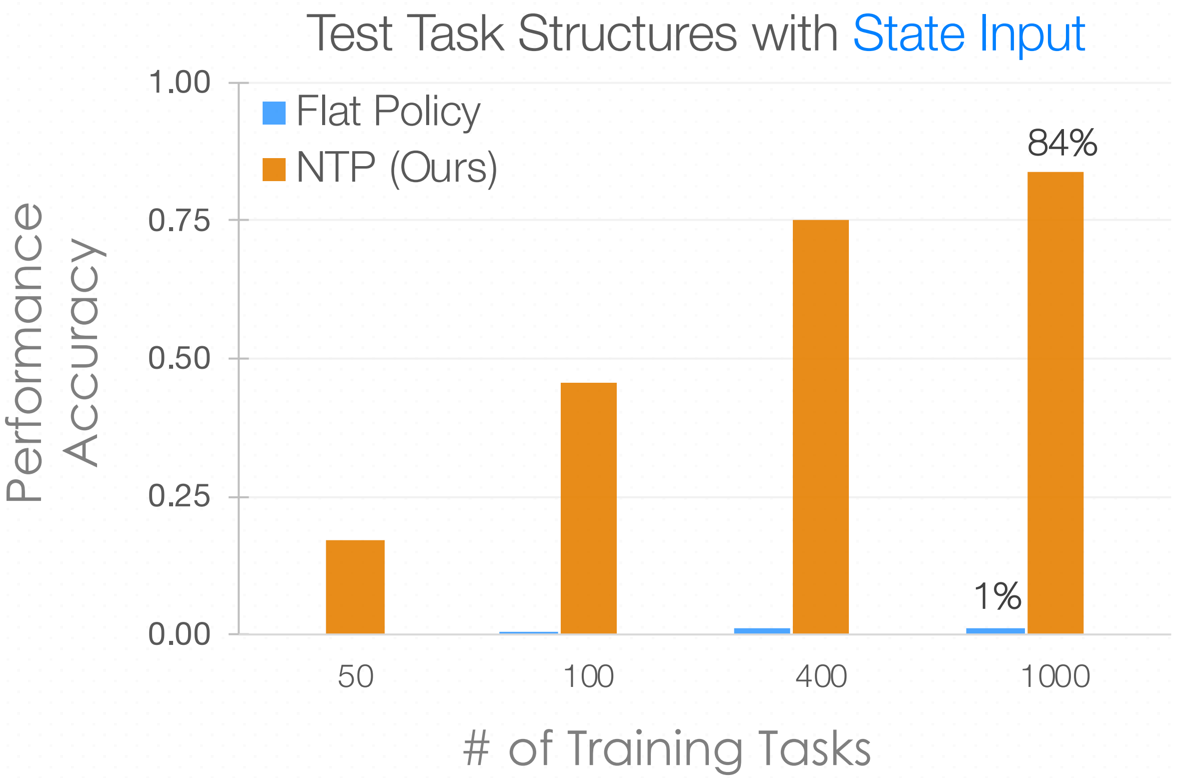

Pose Est. + Plan

E2E Plan

Better Generalization than Flat Policy + Works with Vision

[ICRA'18, CVPR '19]

Failure Modes

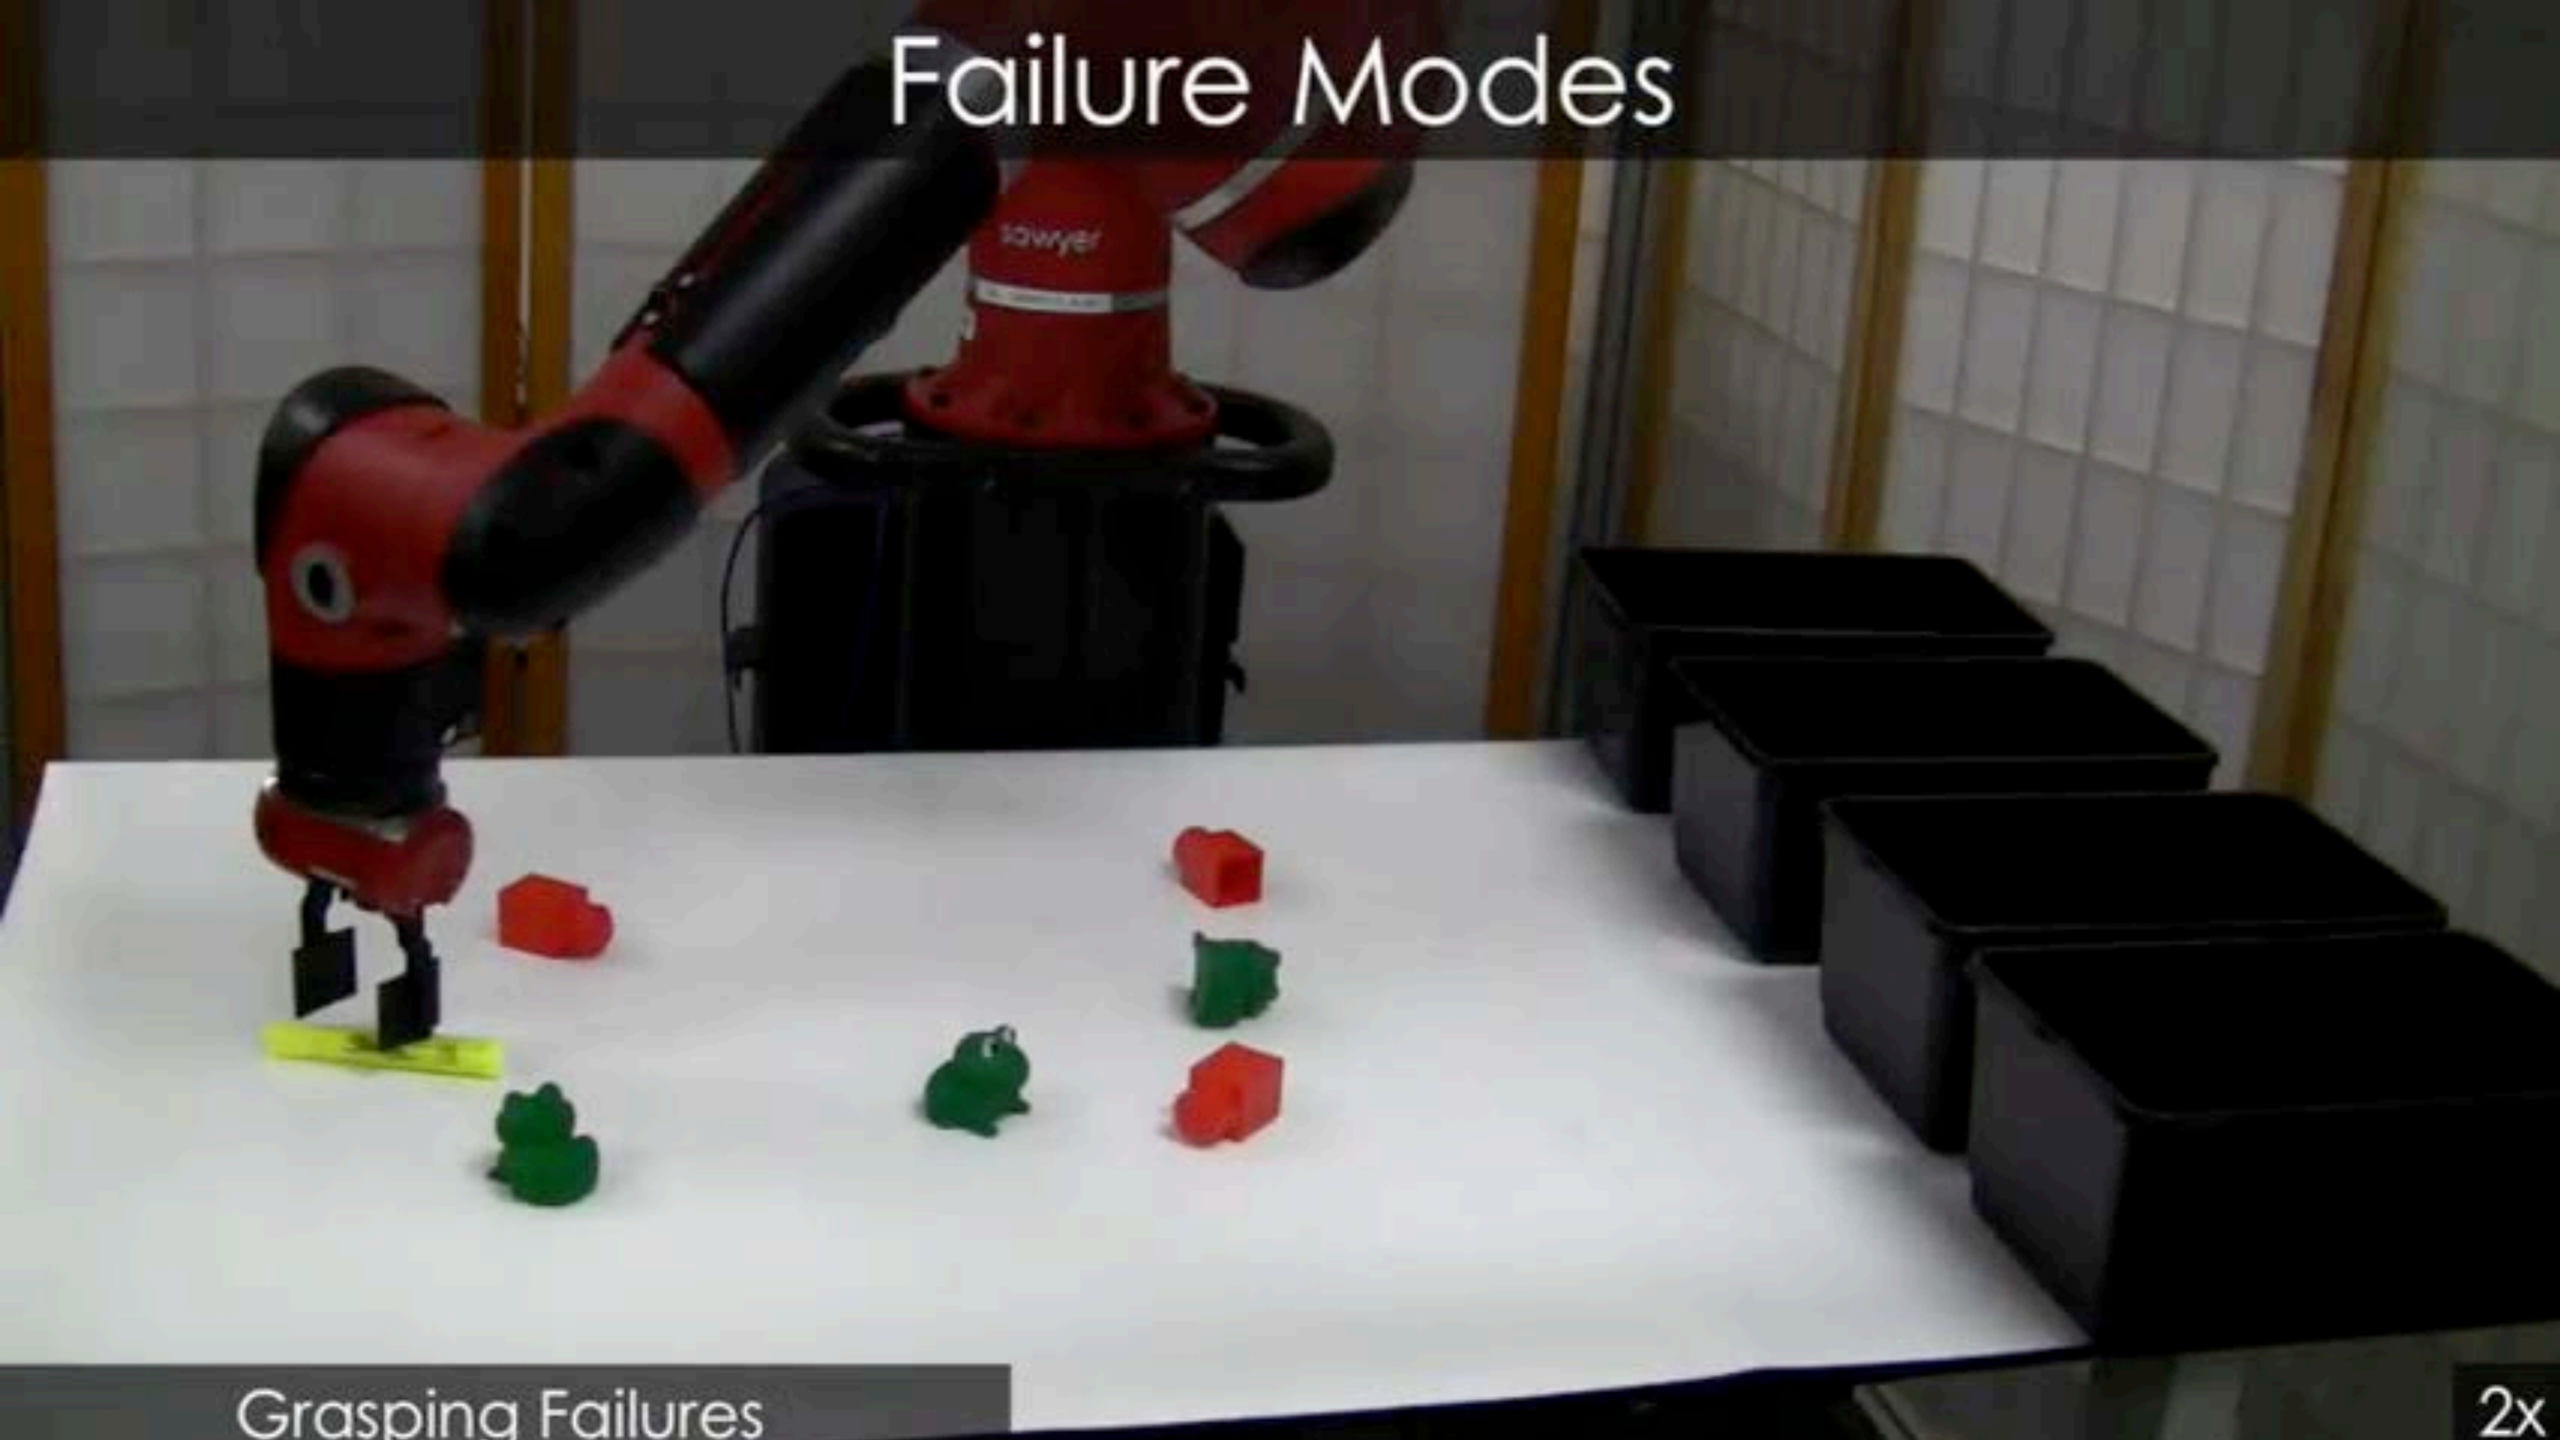

Grasping Failures

2x

Compositional Planning: Task Programming

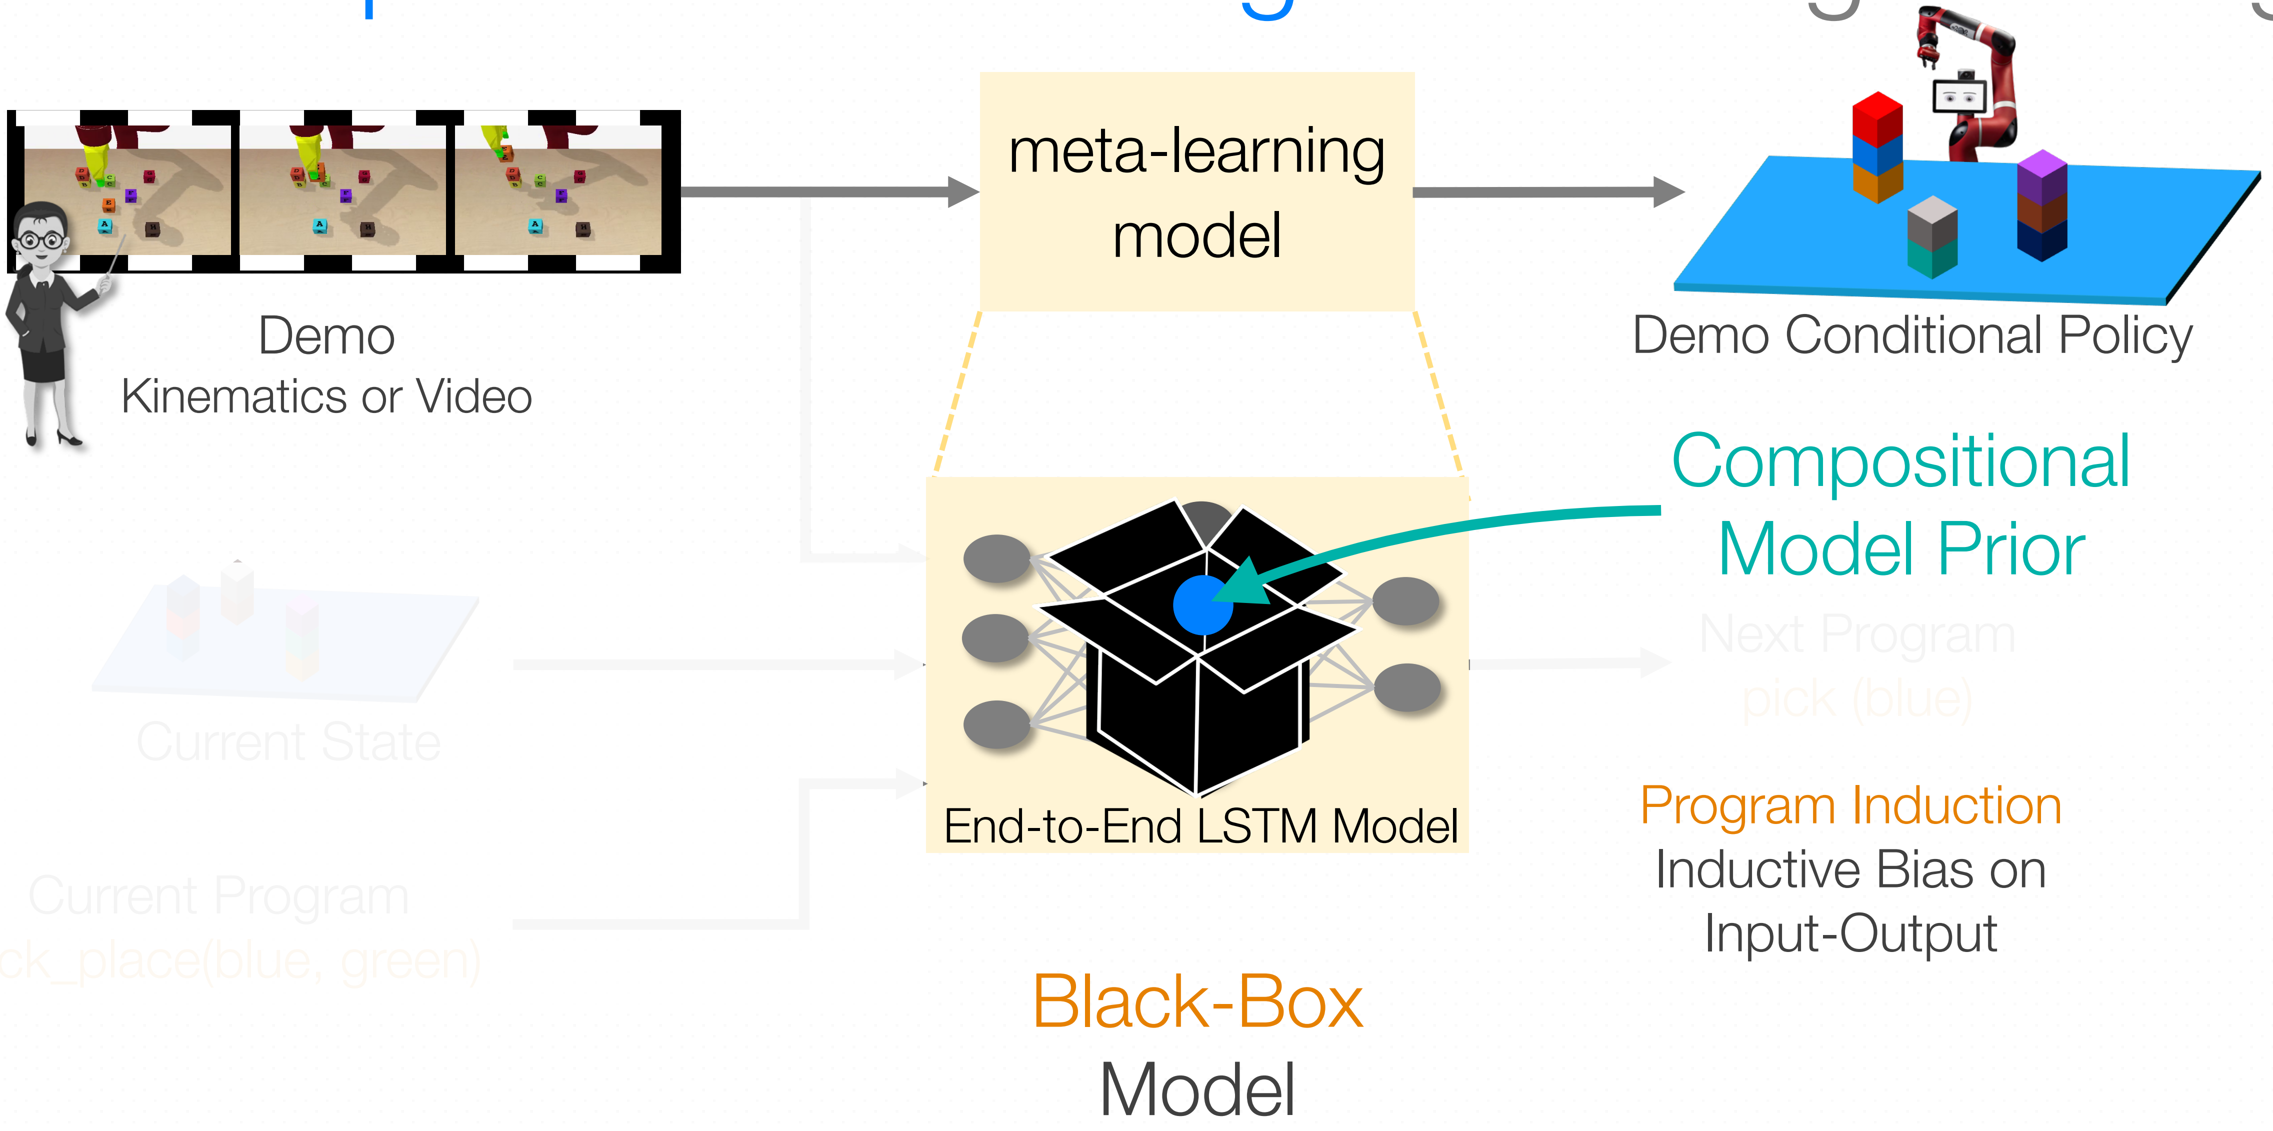

[ICRA'18, CVPR '19]

Compositional Planning: Task Graphs

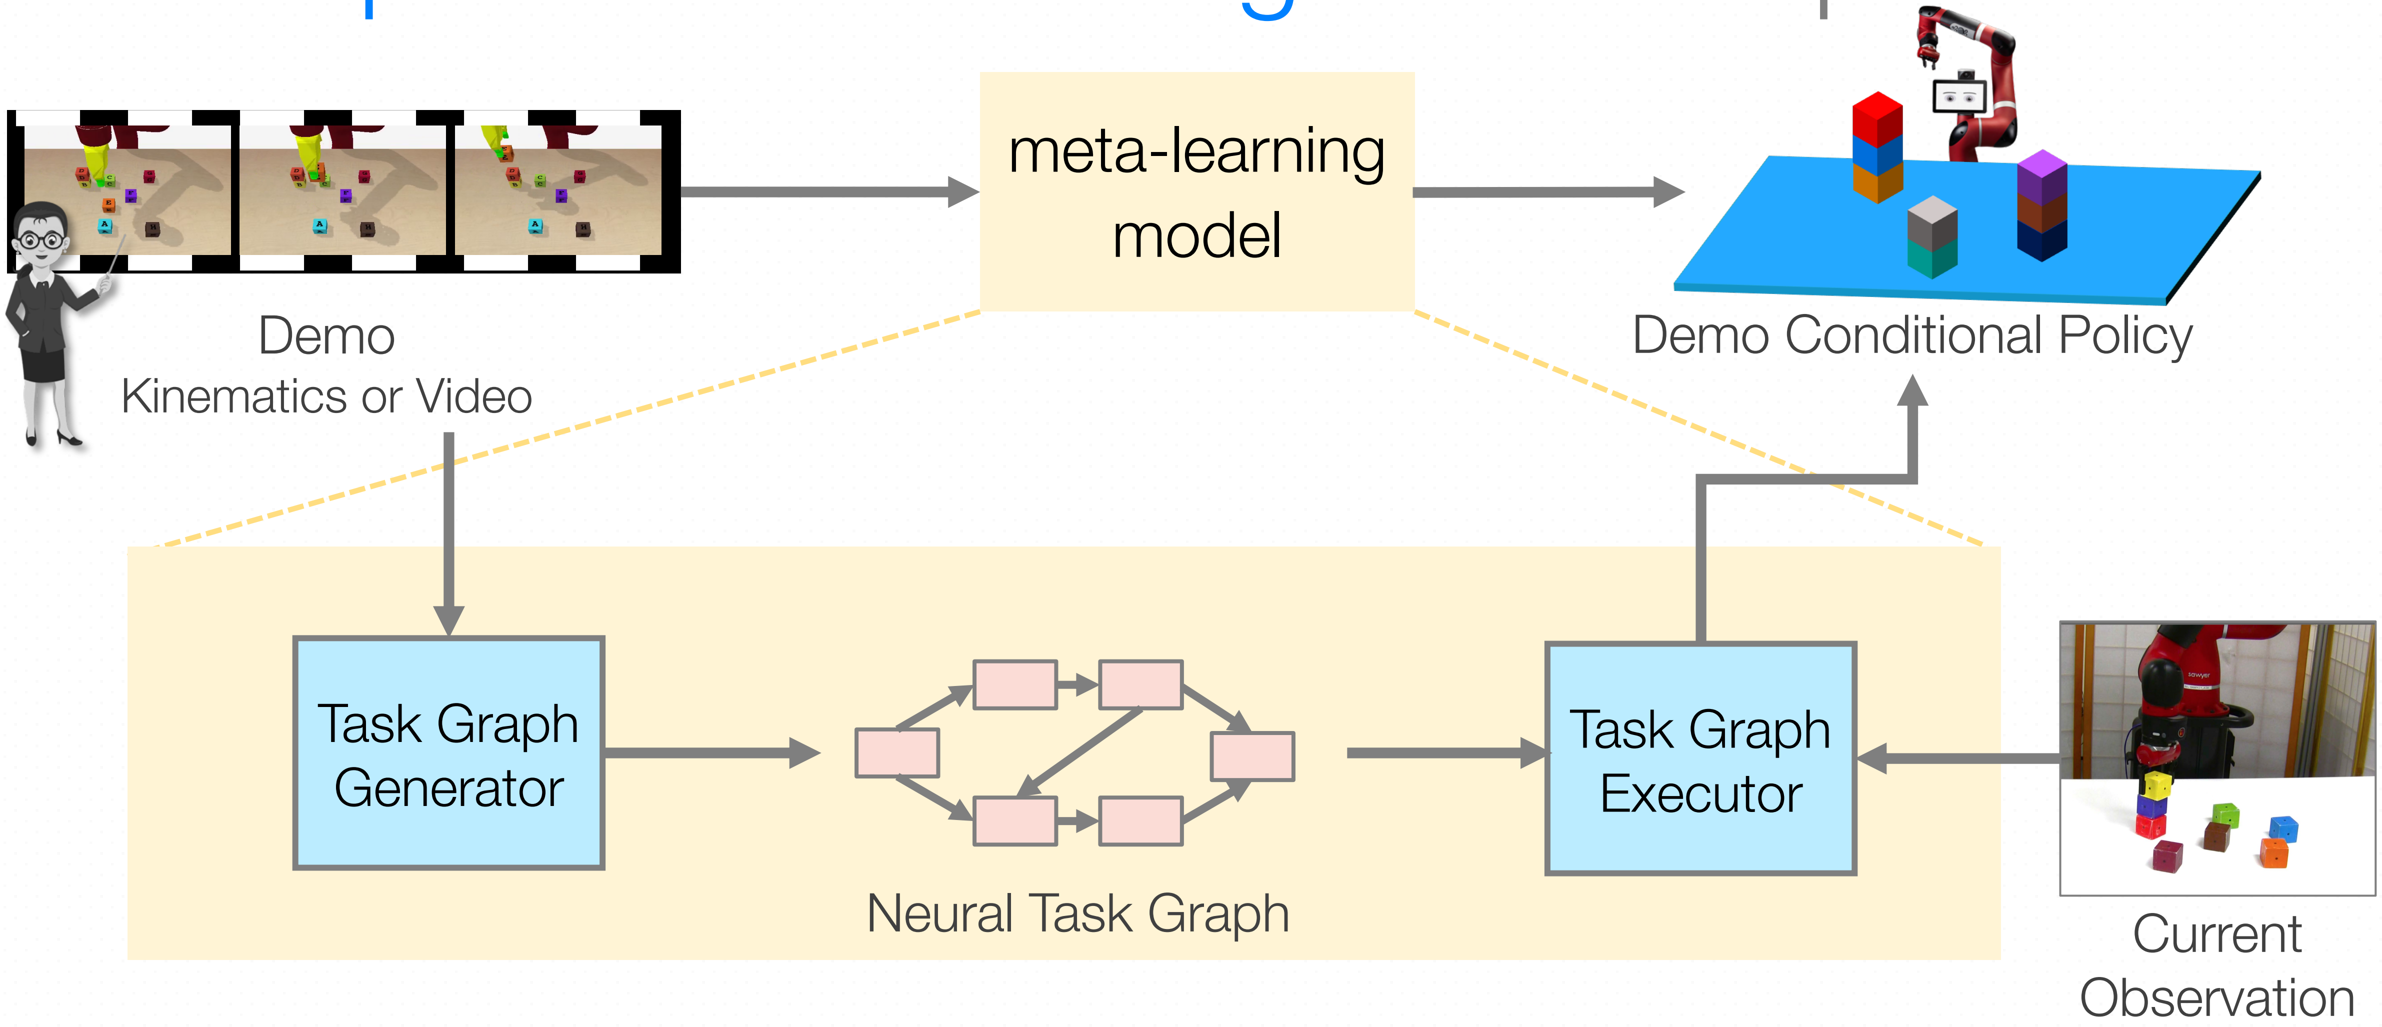

Hierarchical Policy Learning as Graph Induction

[ICRA'18, CVPR '19]

Neural Task Graphs (NTG)

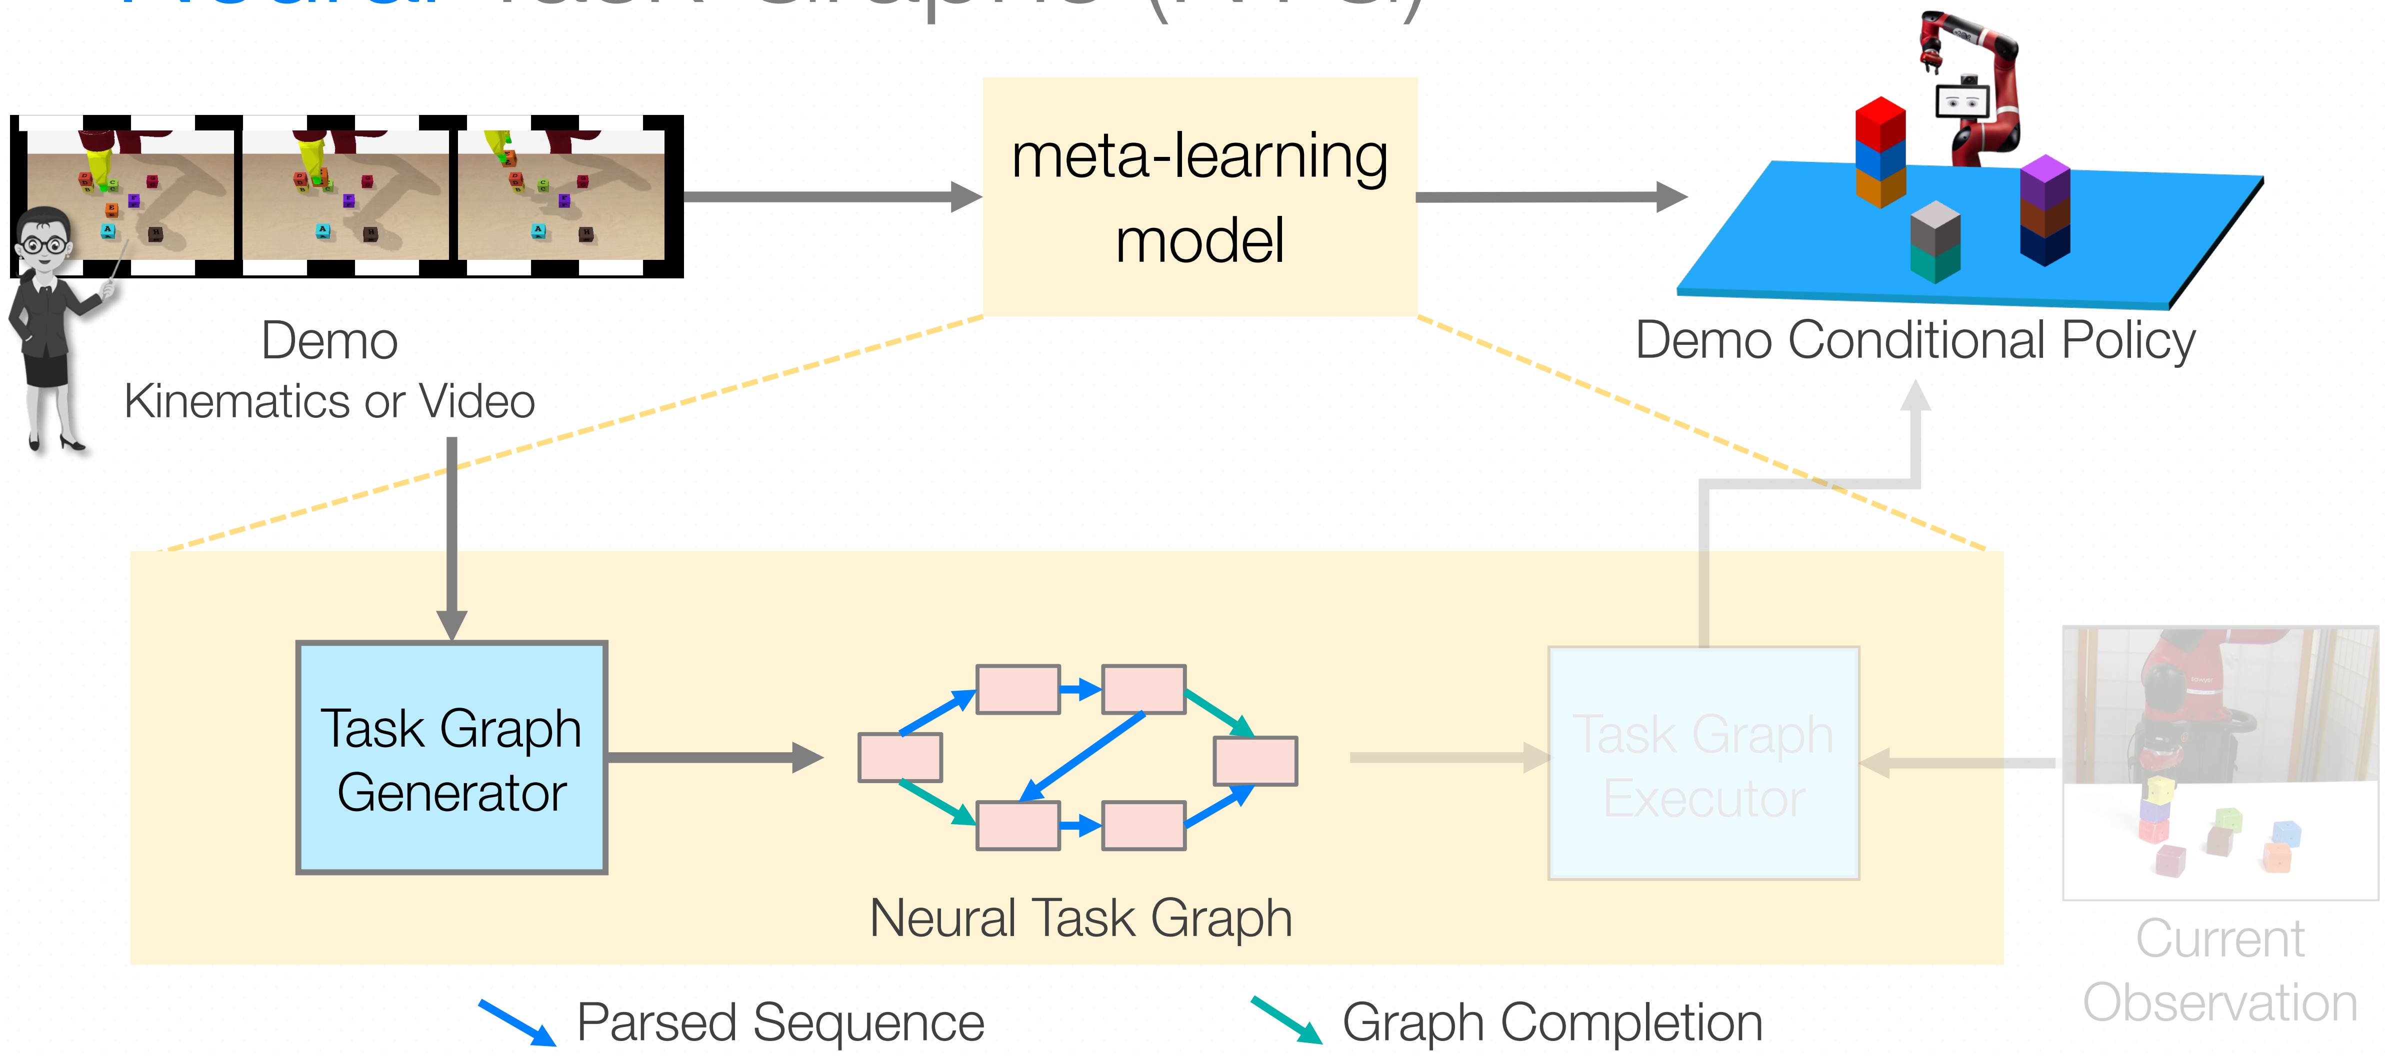

Hierarchical Policy Learning as Graph Induction

[ICRA'18, CVPR '19]

Neural Task Graphs (NTG)

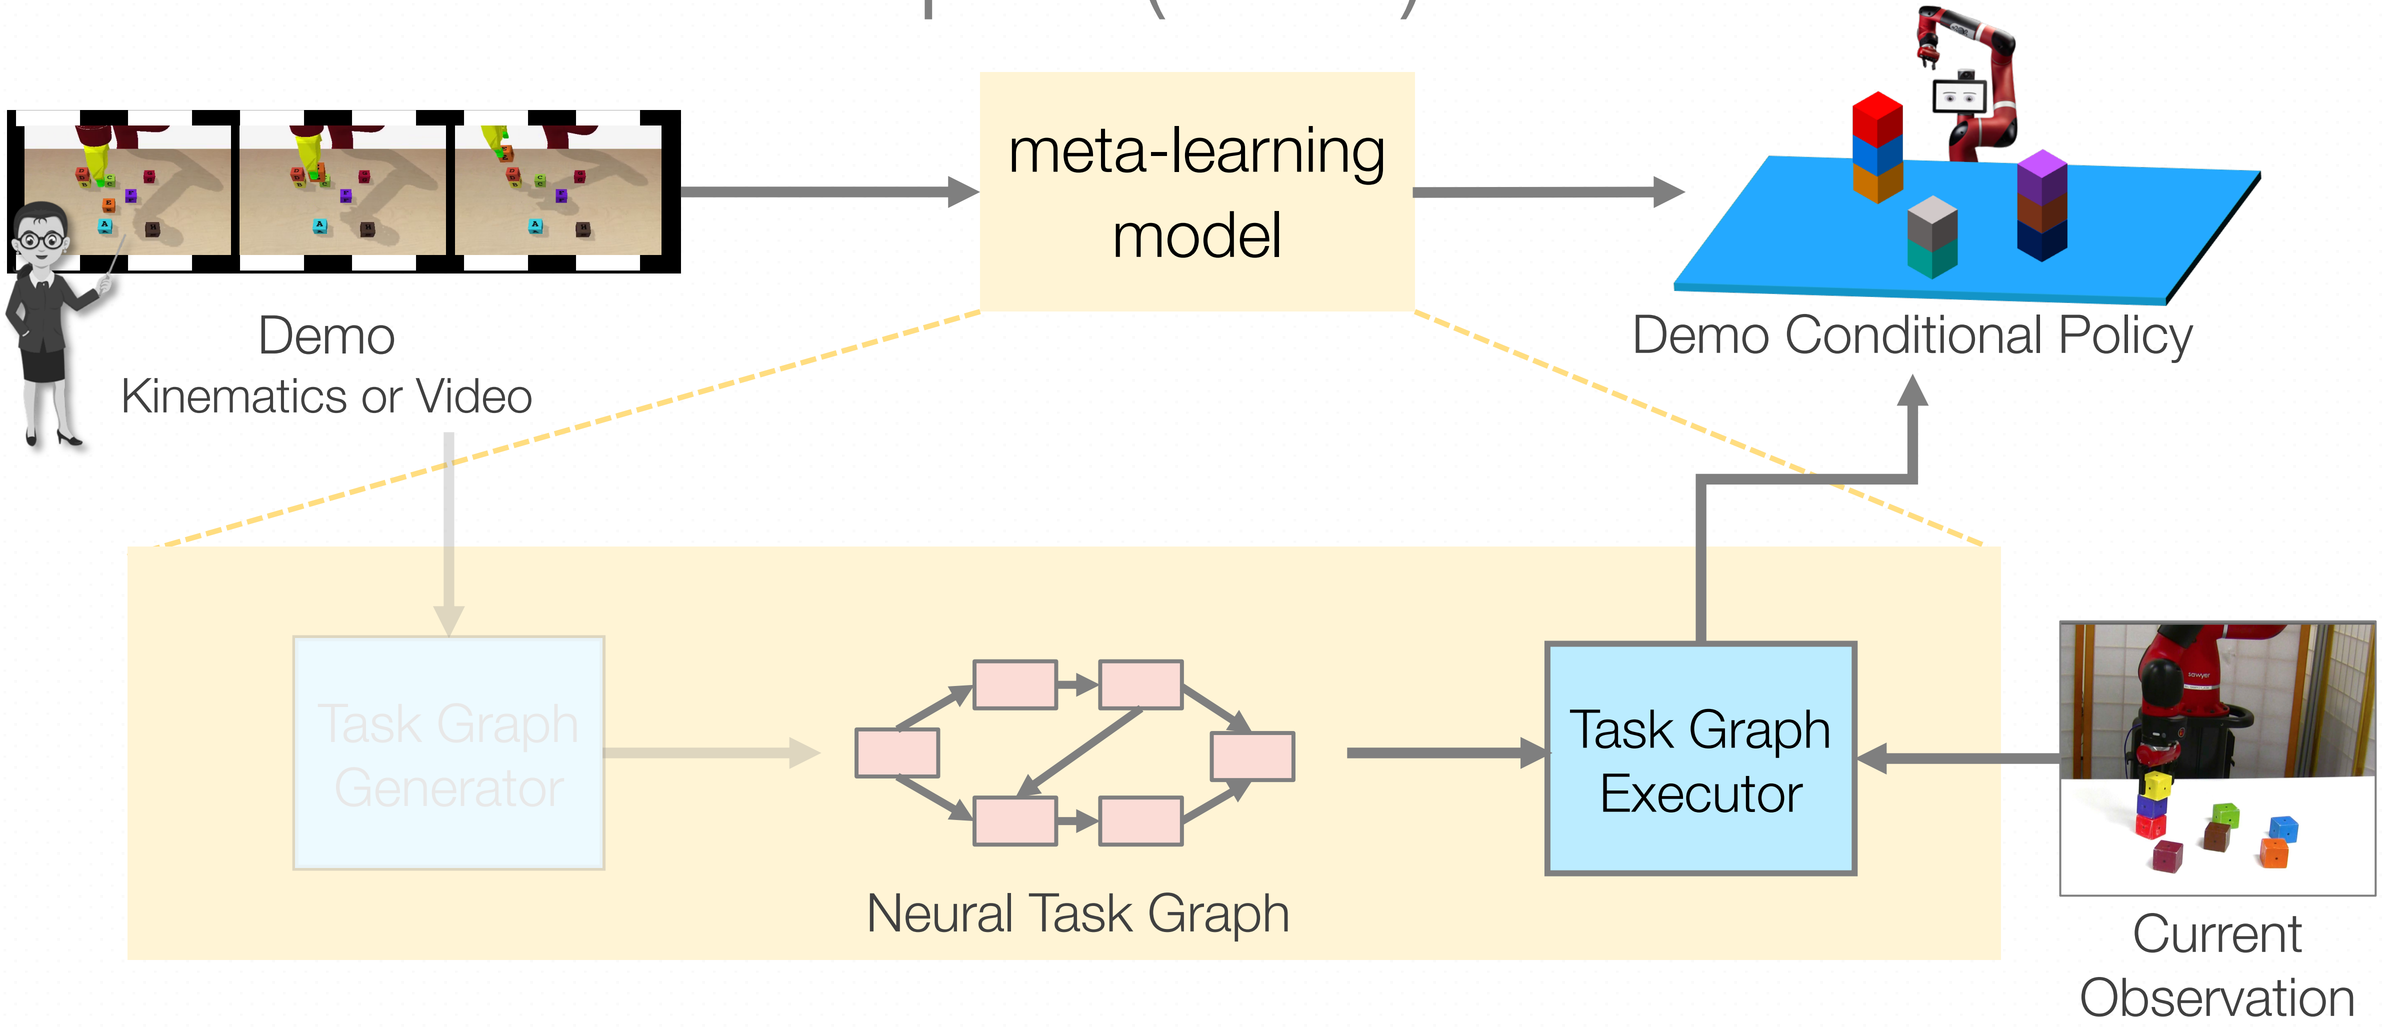

Hierarchical Policy Learning as Graph Induction

[ICRA'18, CVPR '19]

Neural Task Graphs (NTG): Representation

Task Graph

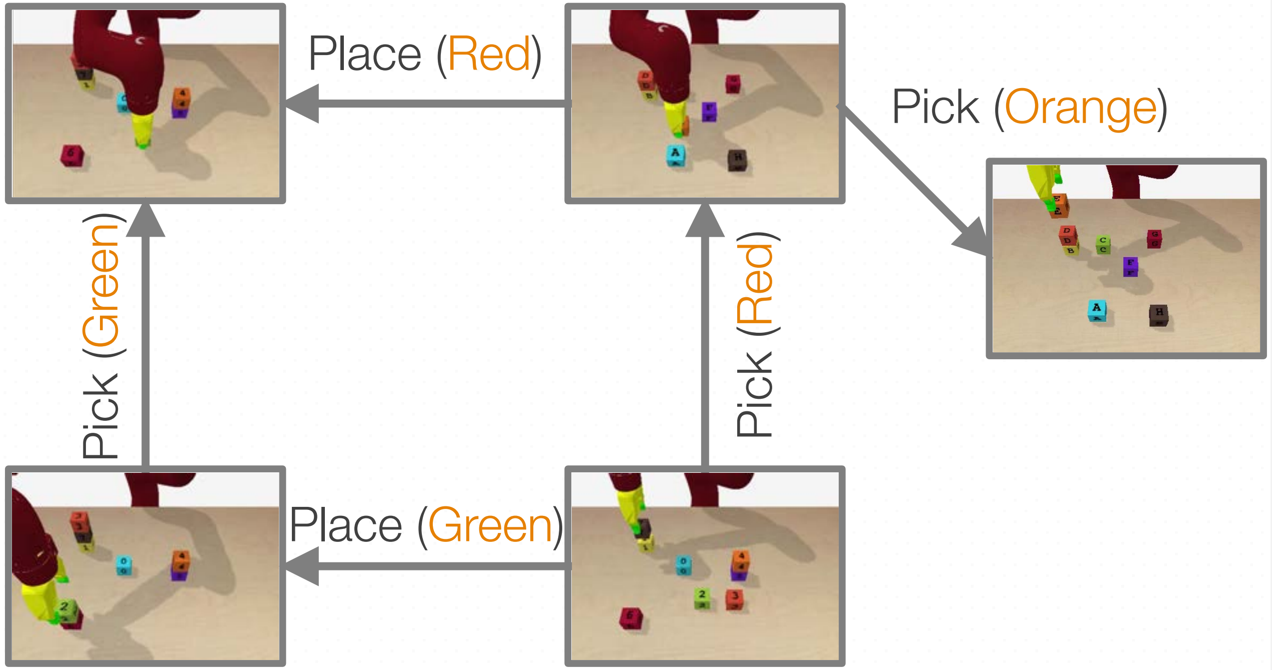

Nodes: States **Combinatorial**

Edges: Action

Conjugate Task Graph

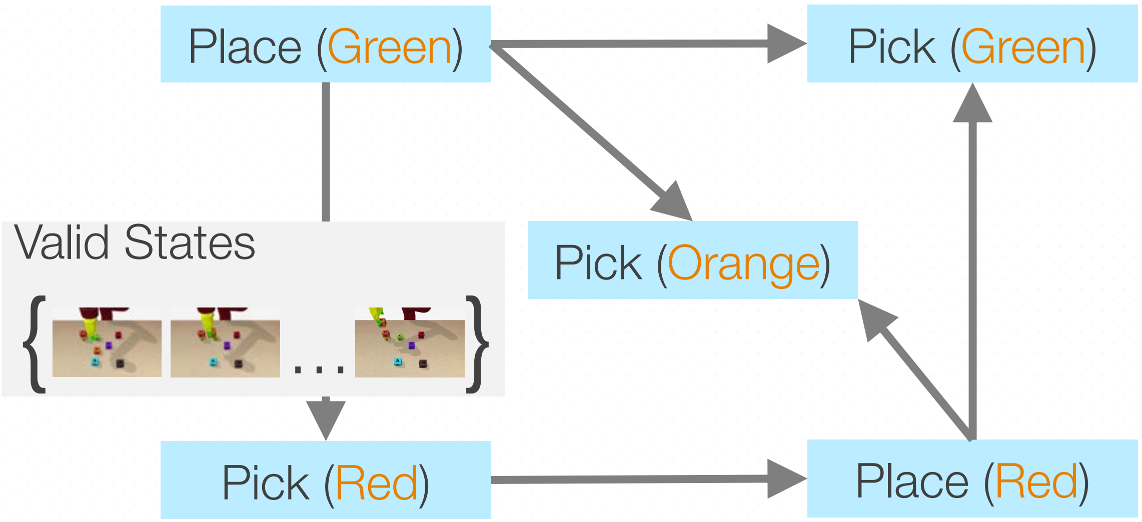

Nodes: Actions **Finite**

Edges: States (Preconditions)

[ICRA'18, CVPR '19]

Neural Task Graphs (NTG): Execution

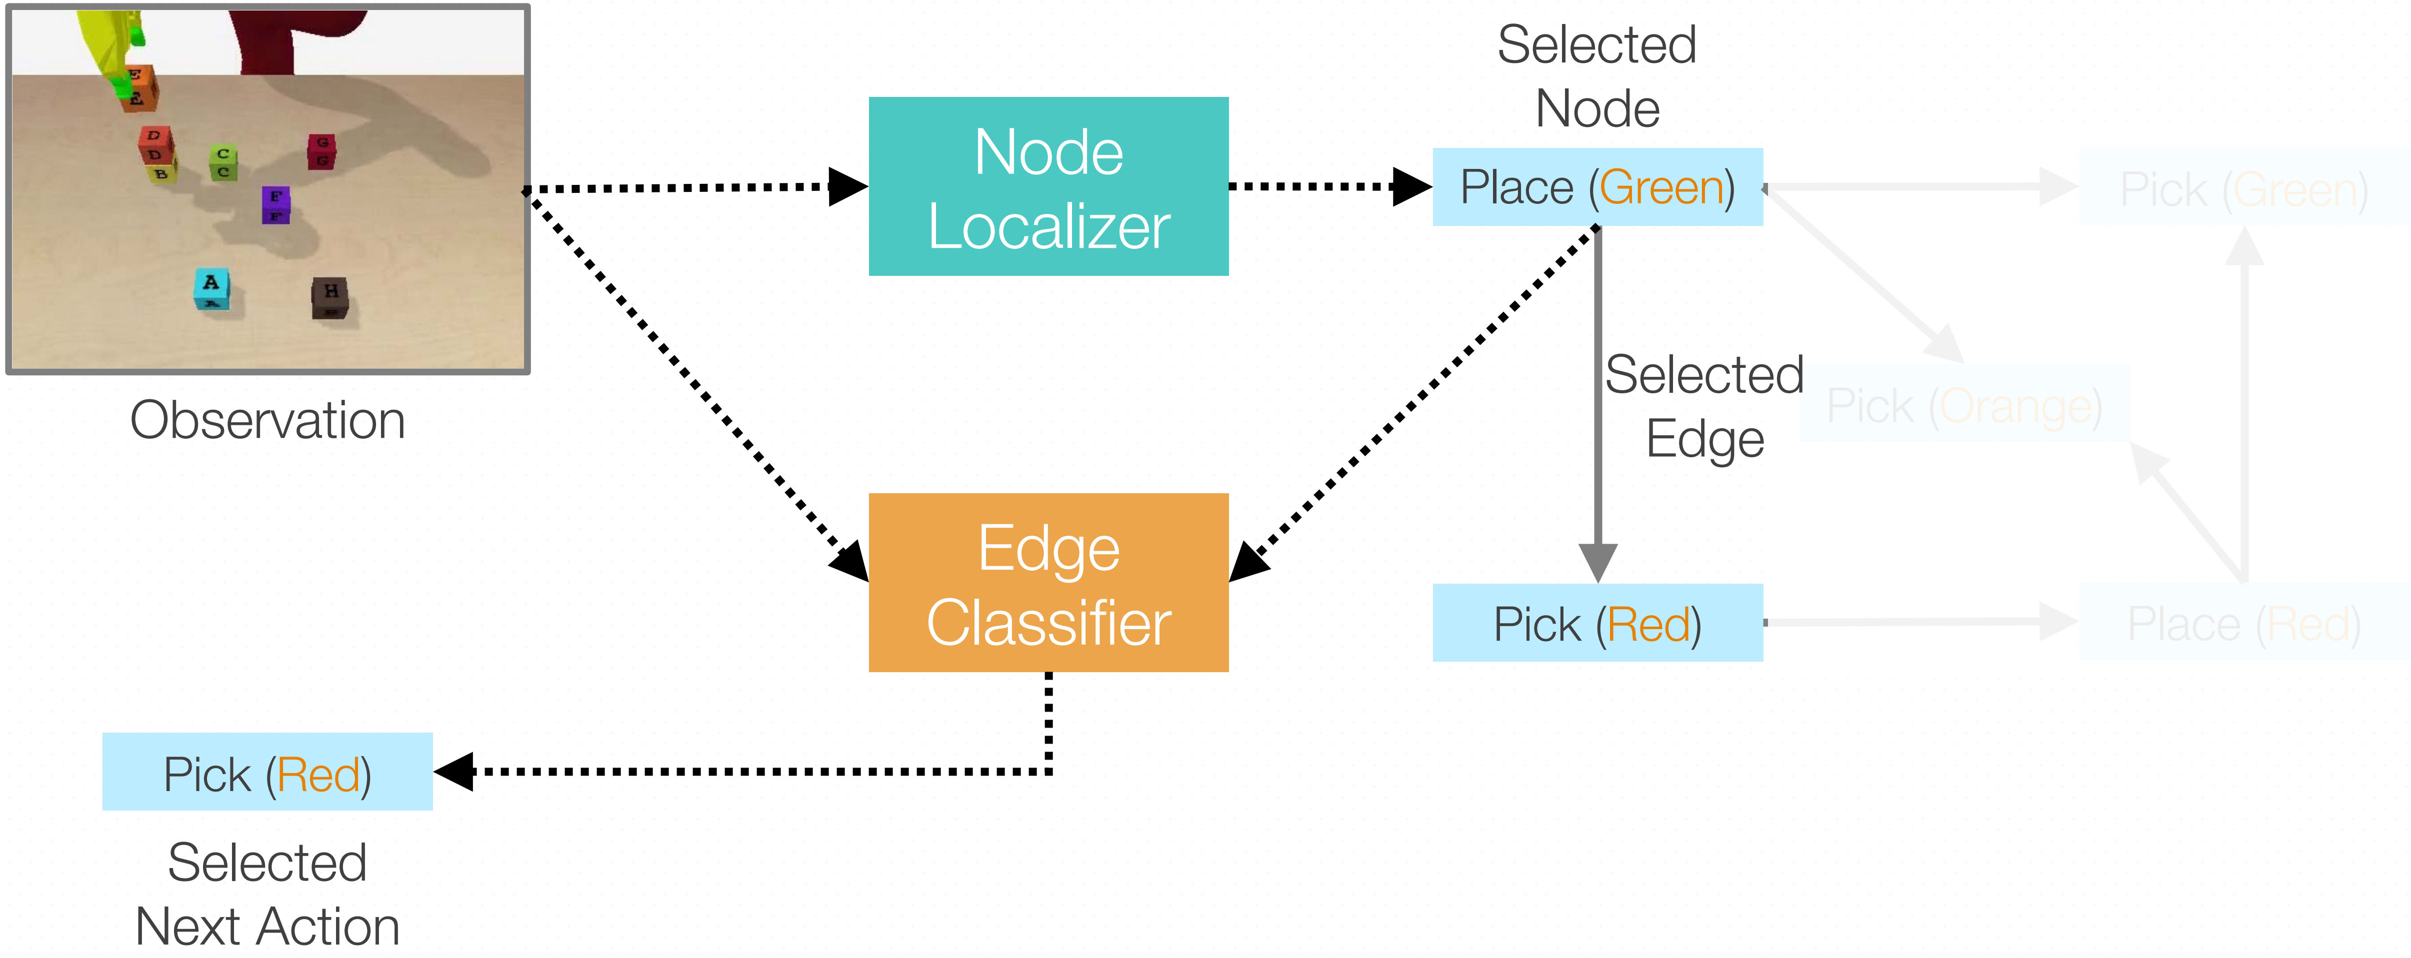

[ICRA'18, CVPR '19]

Neural Task Graphs (NTG)

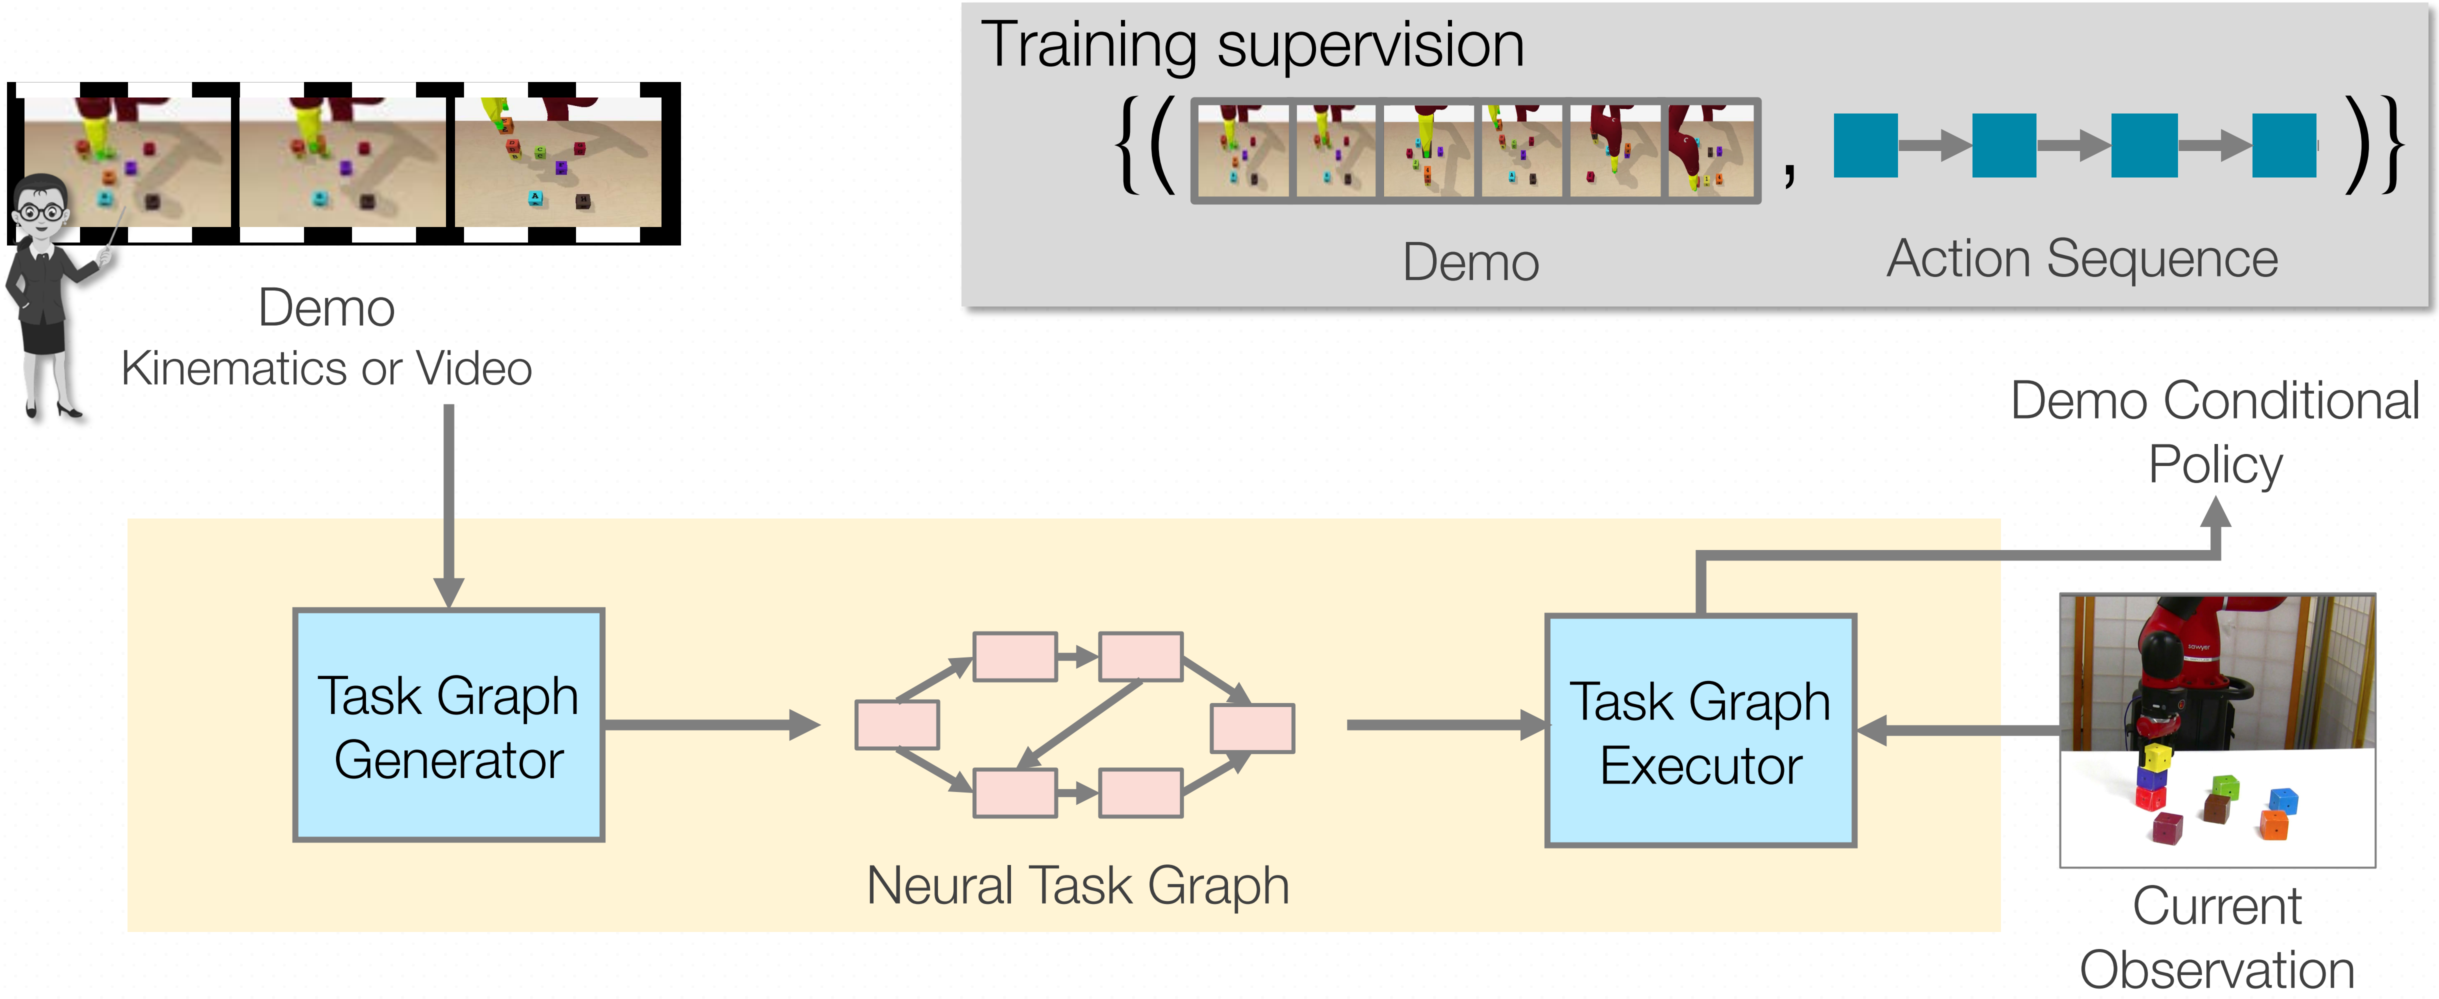

Hierarchical Policy Learning as Graph Induction

[ICRA'18, CVPR '19]

Neural Task Graph
Results

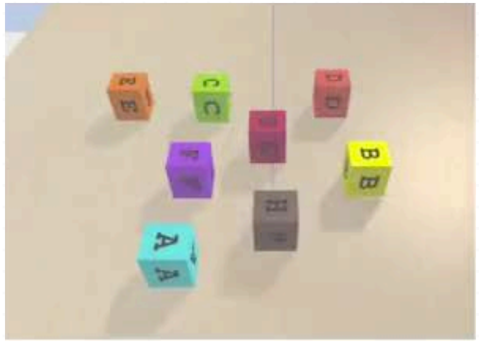

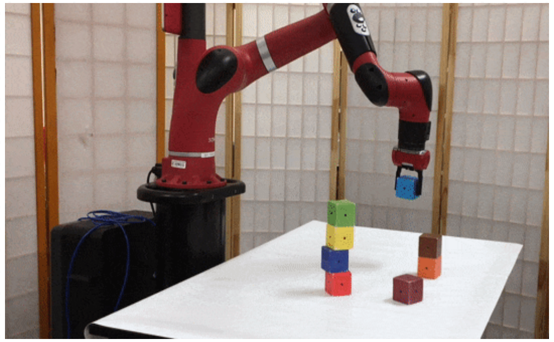

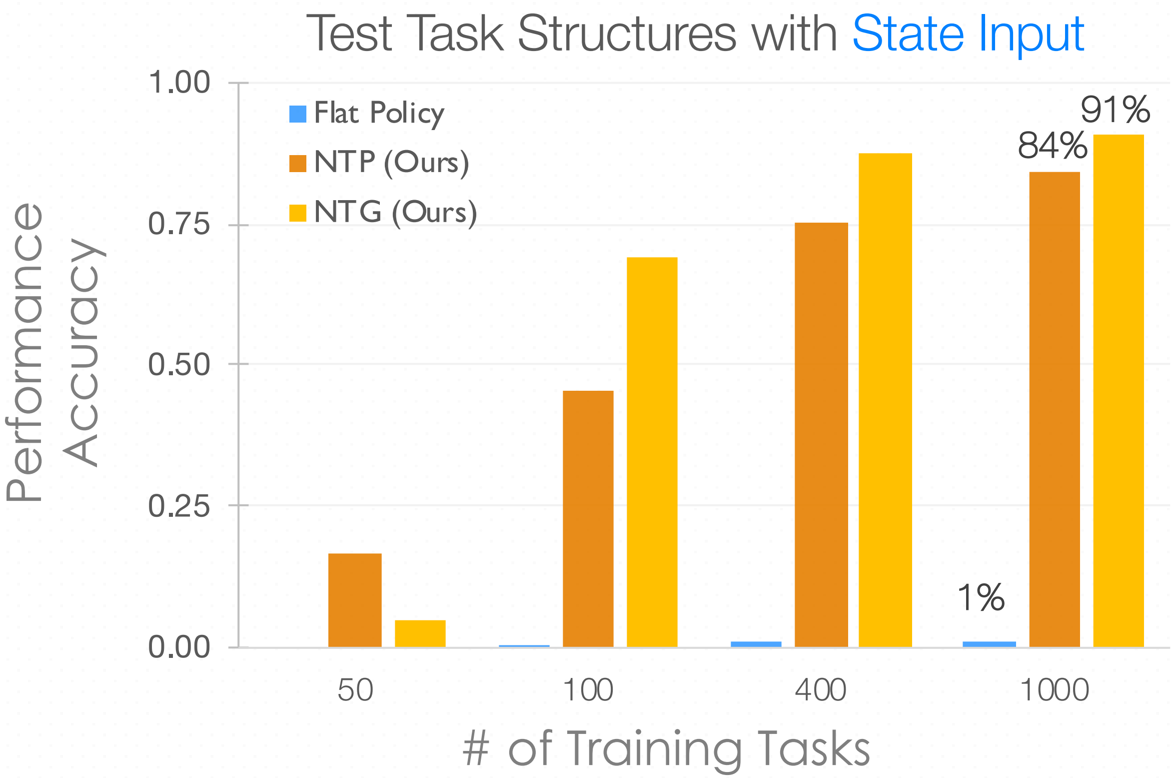

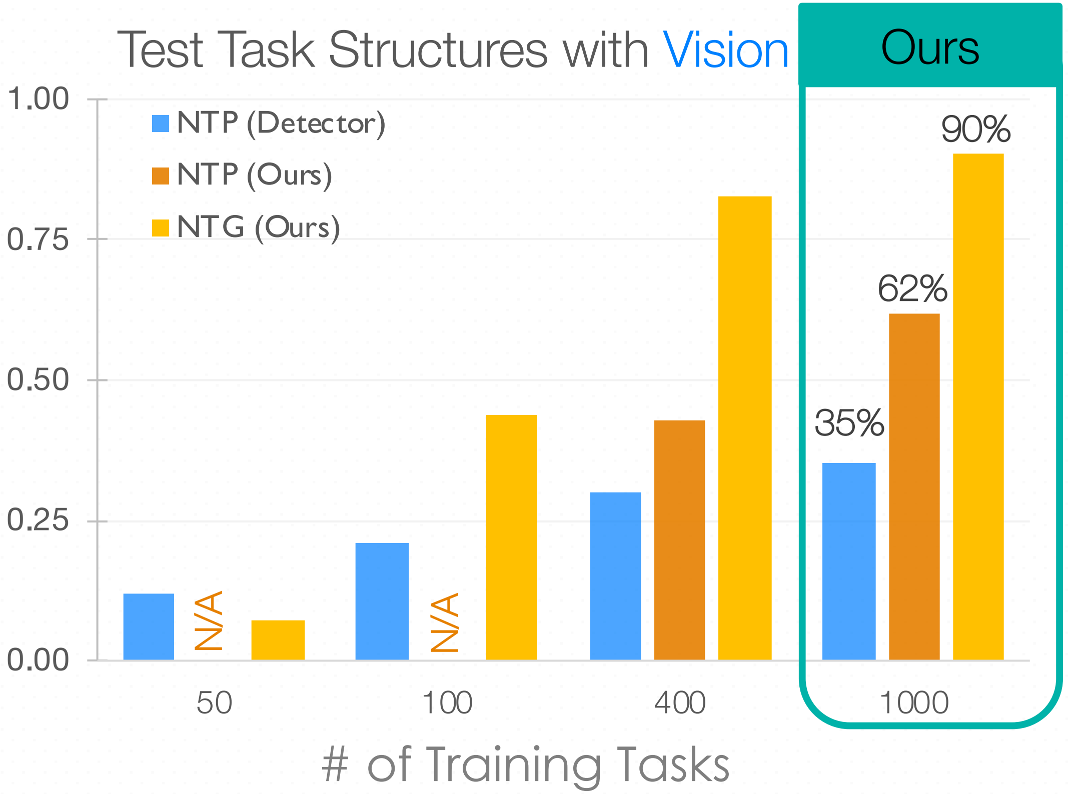

Weaker Supervision and Better Generalization

[ICRA'18, CVPR '19]

Compositional Planning: NTP and NTG

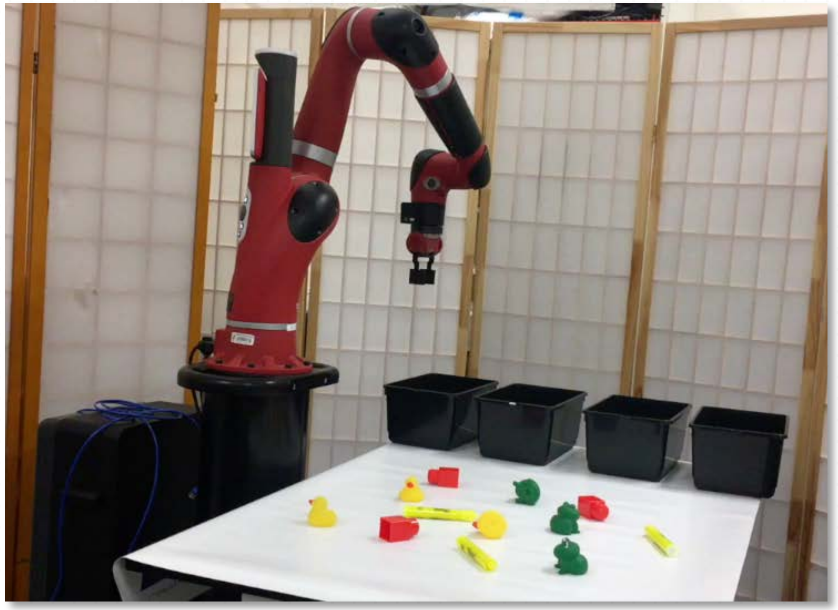

Object Sorting
(NTP)

Table Clean Up
(NTP)

Sequential Search and Prediction
AI2 Thor with NTG

[ICRA'18, CVPR '19]

Task Structure Learning

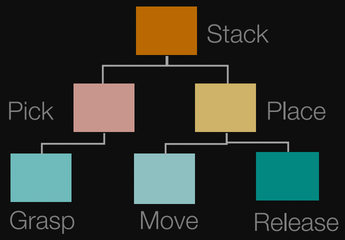

Object Arrangement

ICRA 2018

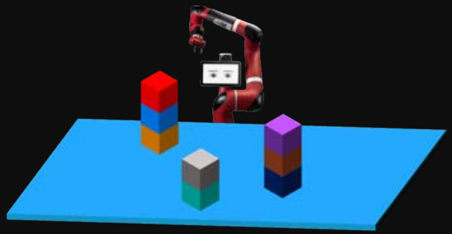

CVPR 2019 (oral)

Compositional priors with modular structure enable generalizable learning in hierarchical domains

Generalizable Autonomy in Robot Manipulation

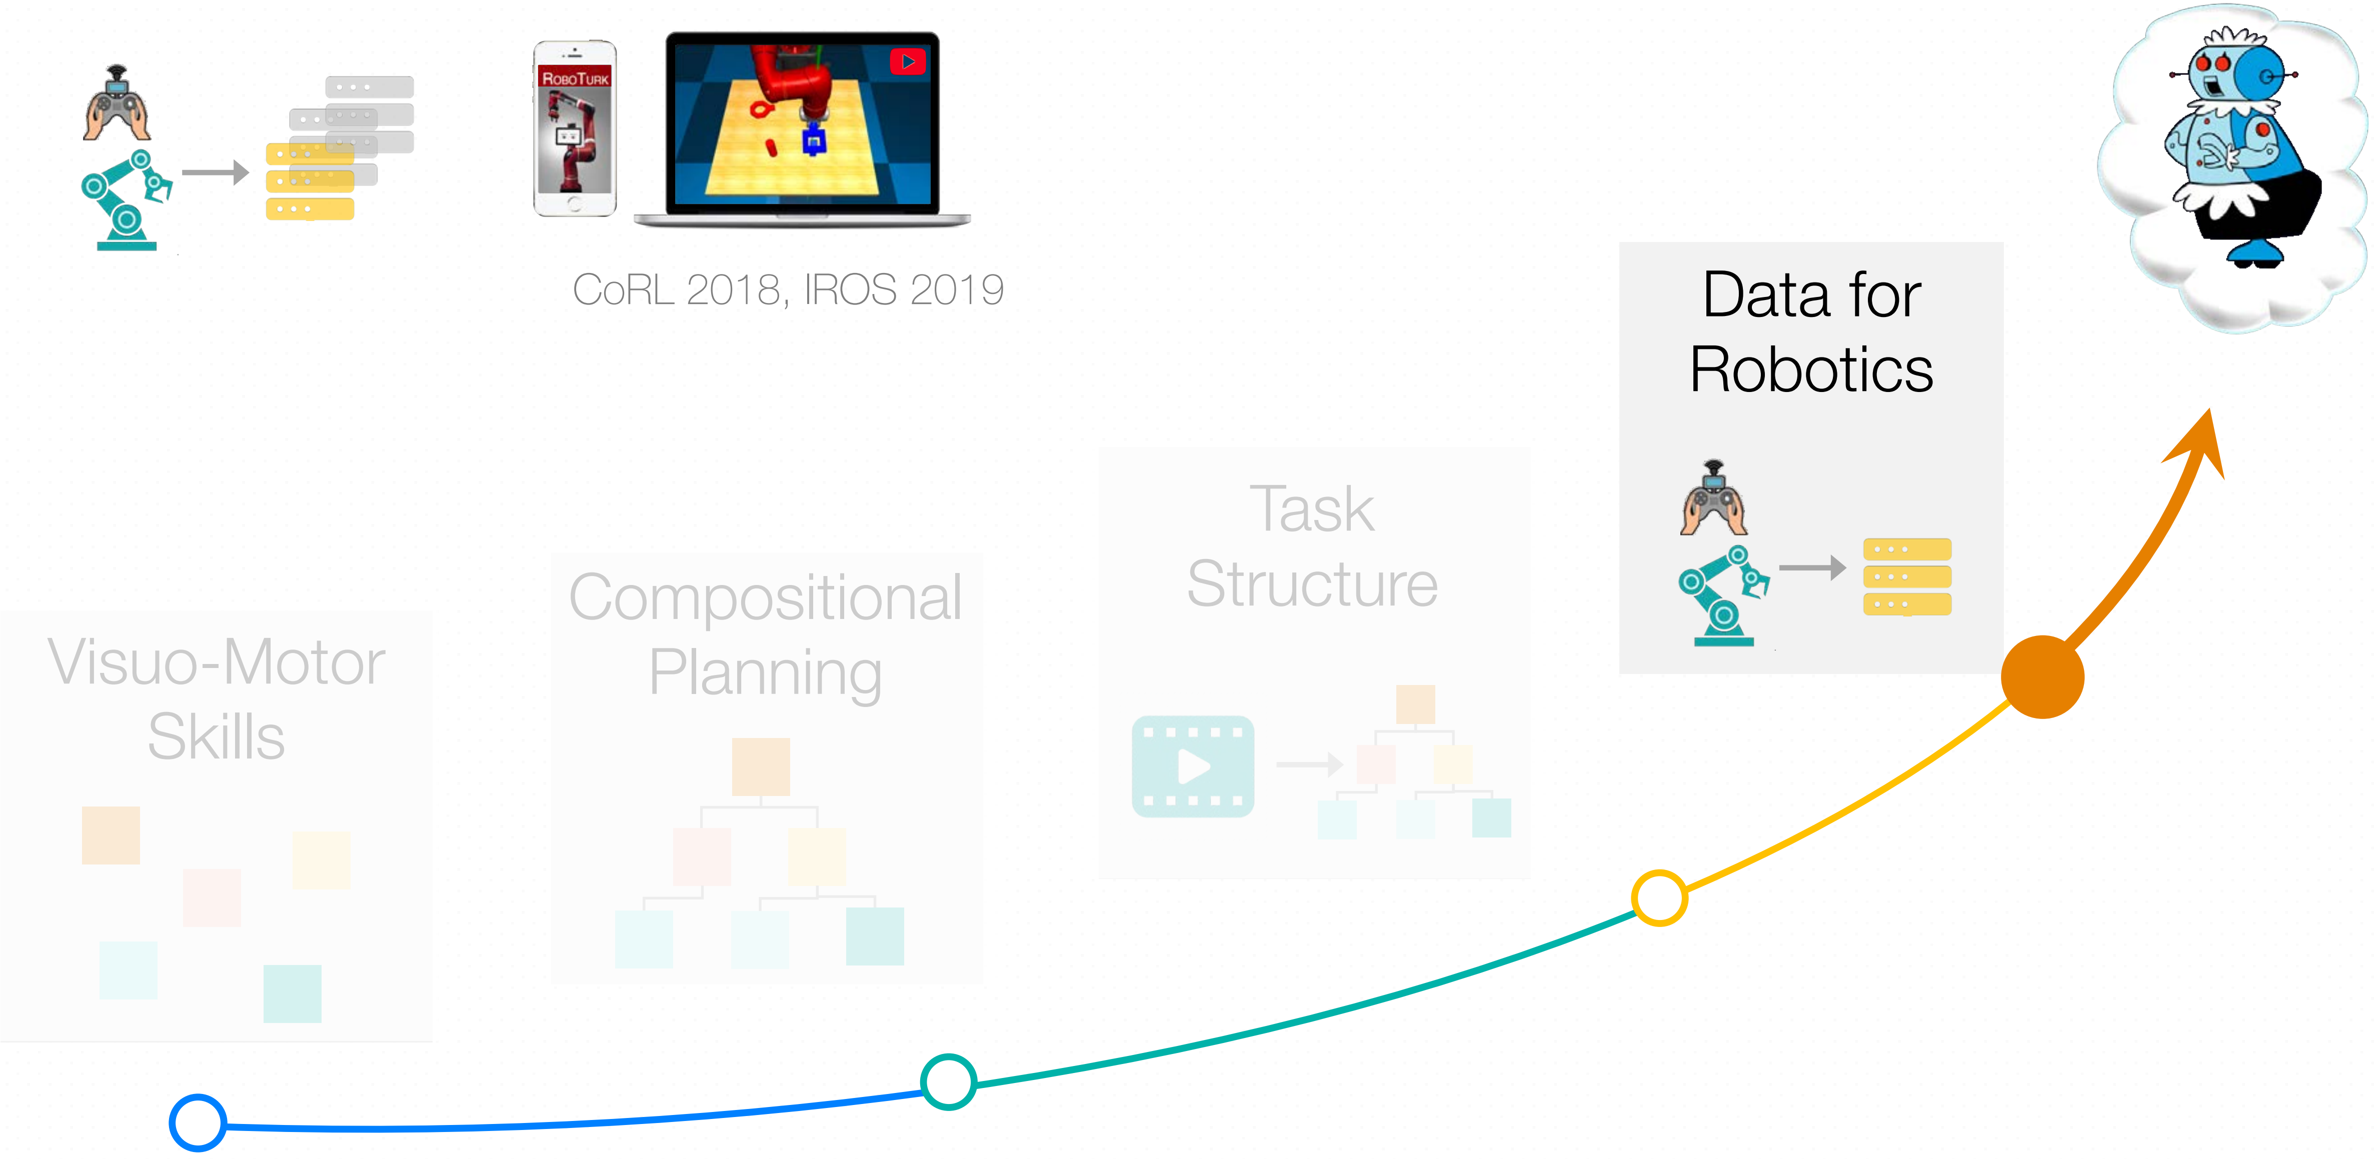

Data for Robotics

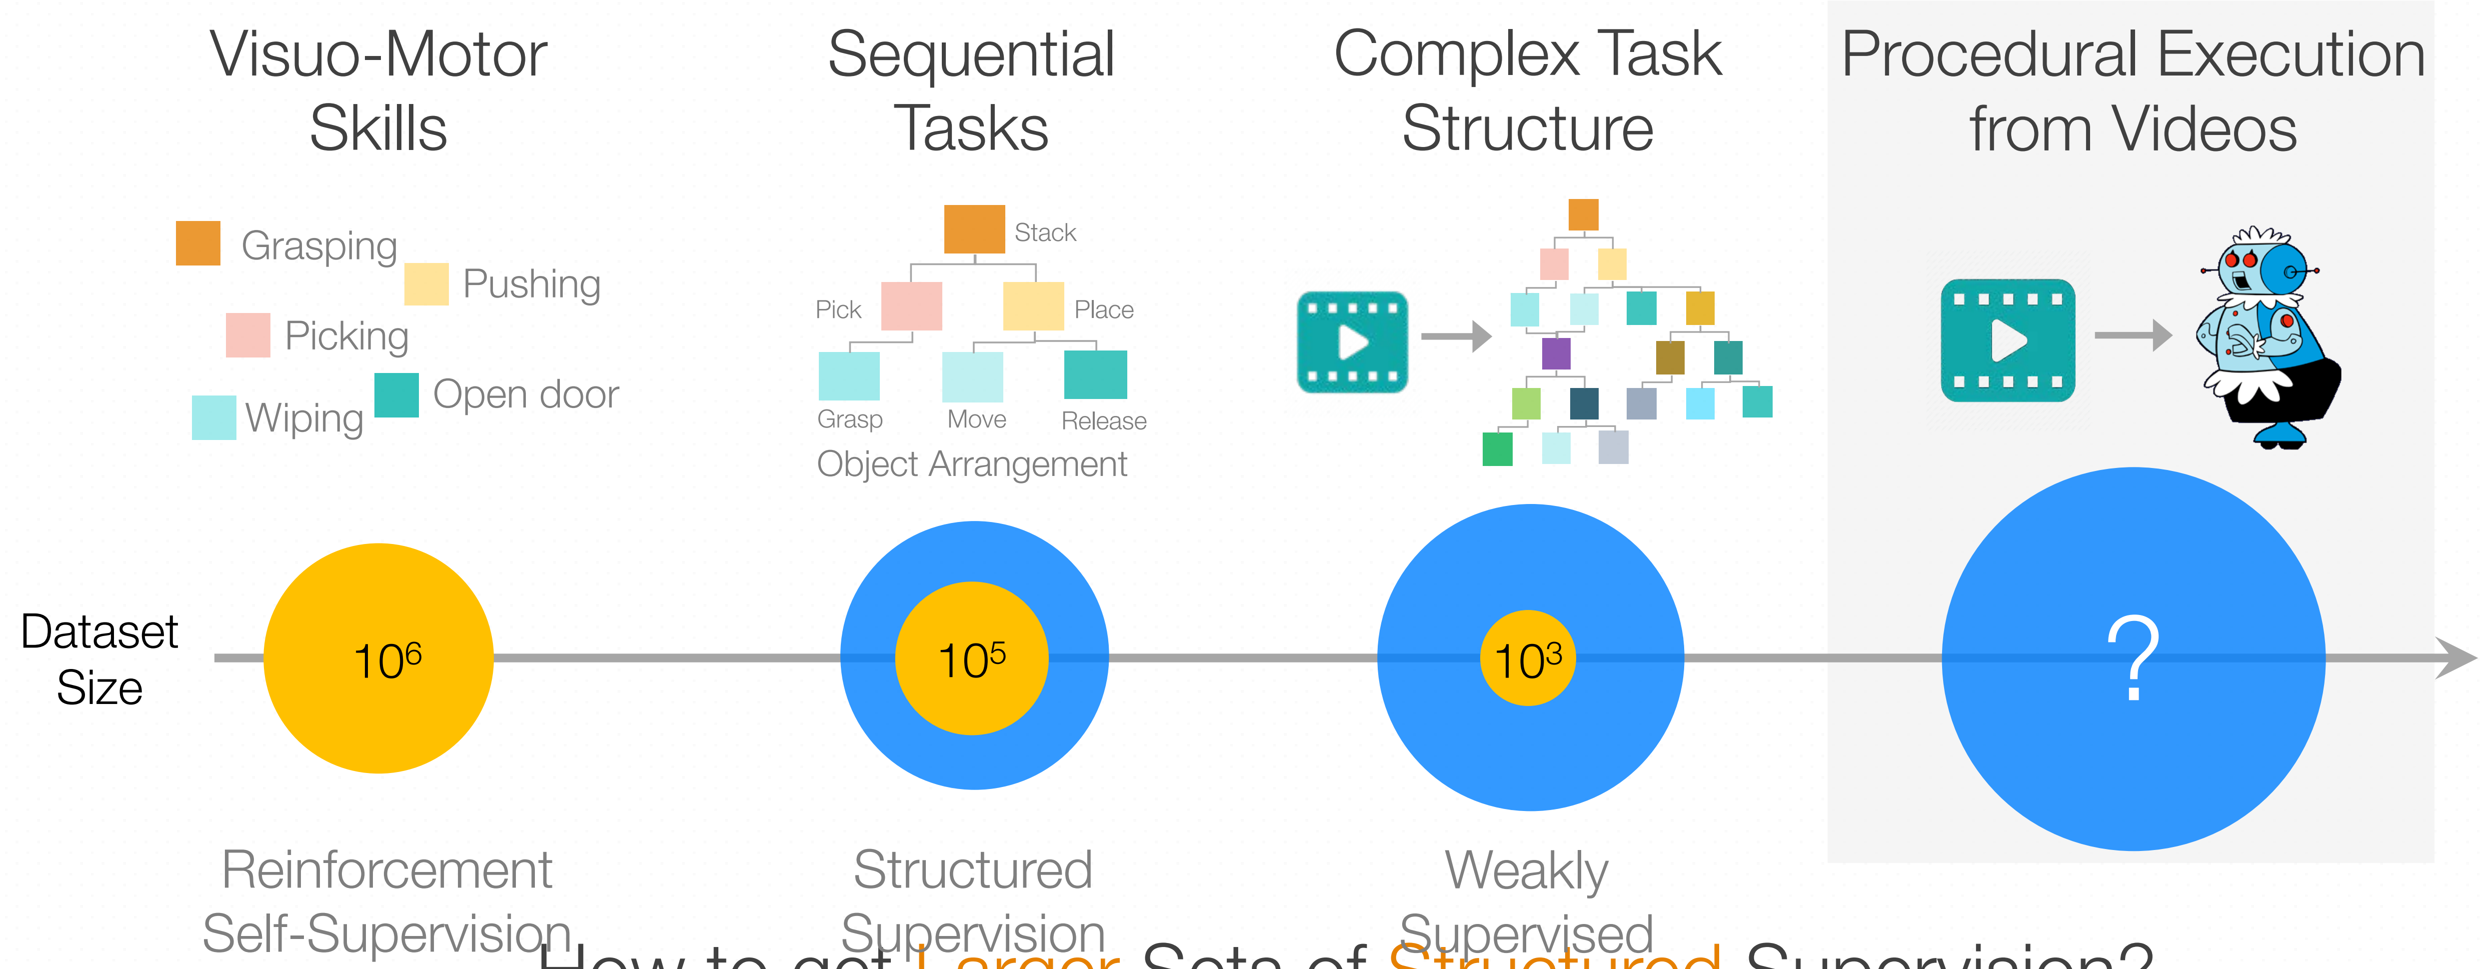

How to get Larger Sets of Structured Supervision?

Data for Robotics: Imitation + RL

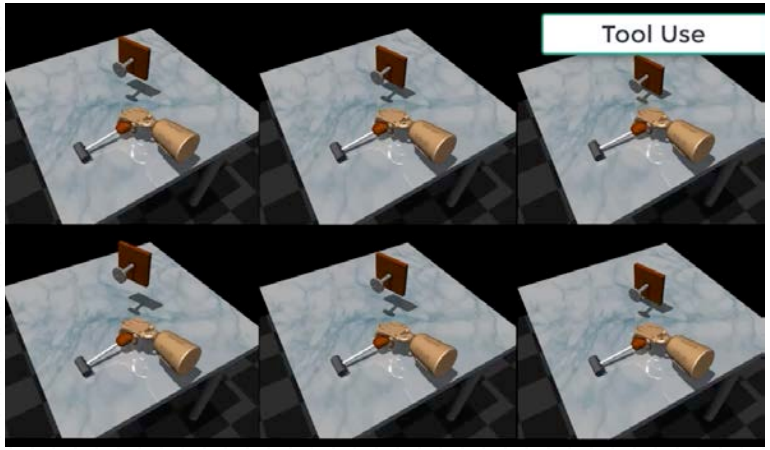

Rajeswaran et al. (2018)

25 demonstrations

~ 10 Minutes

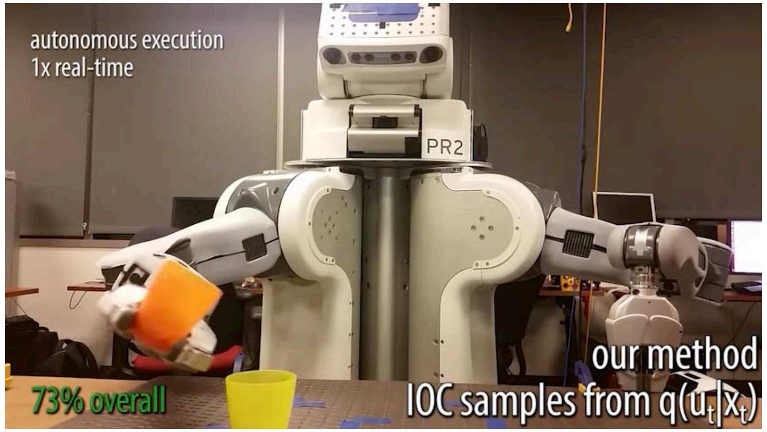

Finn et al. (2017)

30 demonstrations

~ 10 Minutes

Vecerik et al. (2017)

100 demonstrations

~ 30 Minutes

Data for Robotics

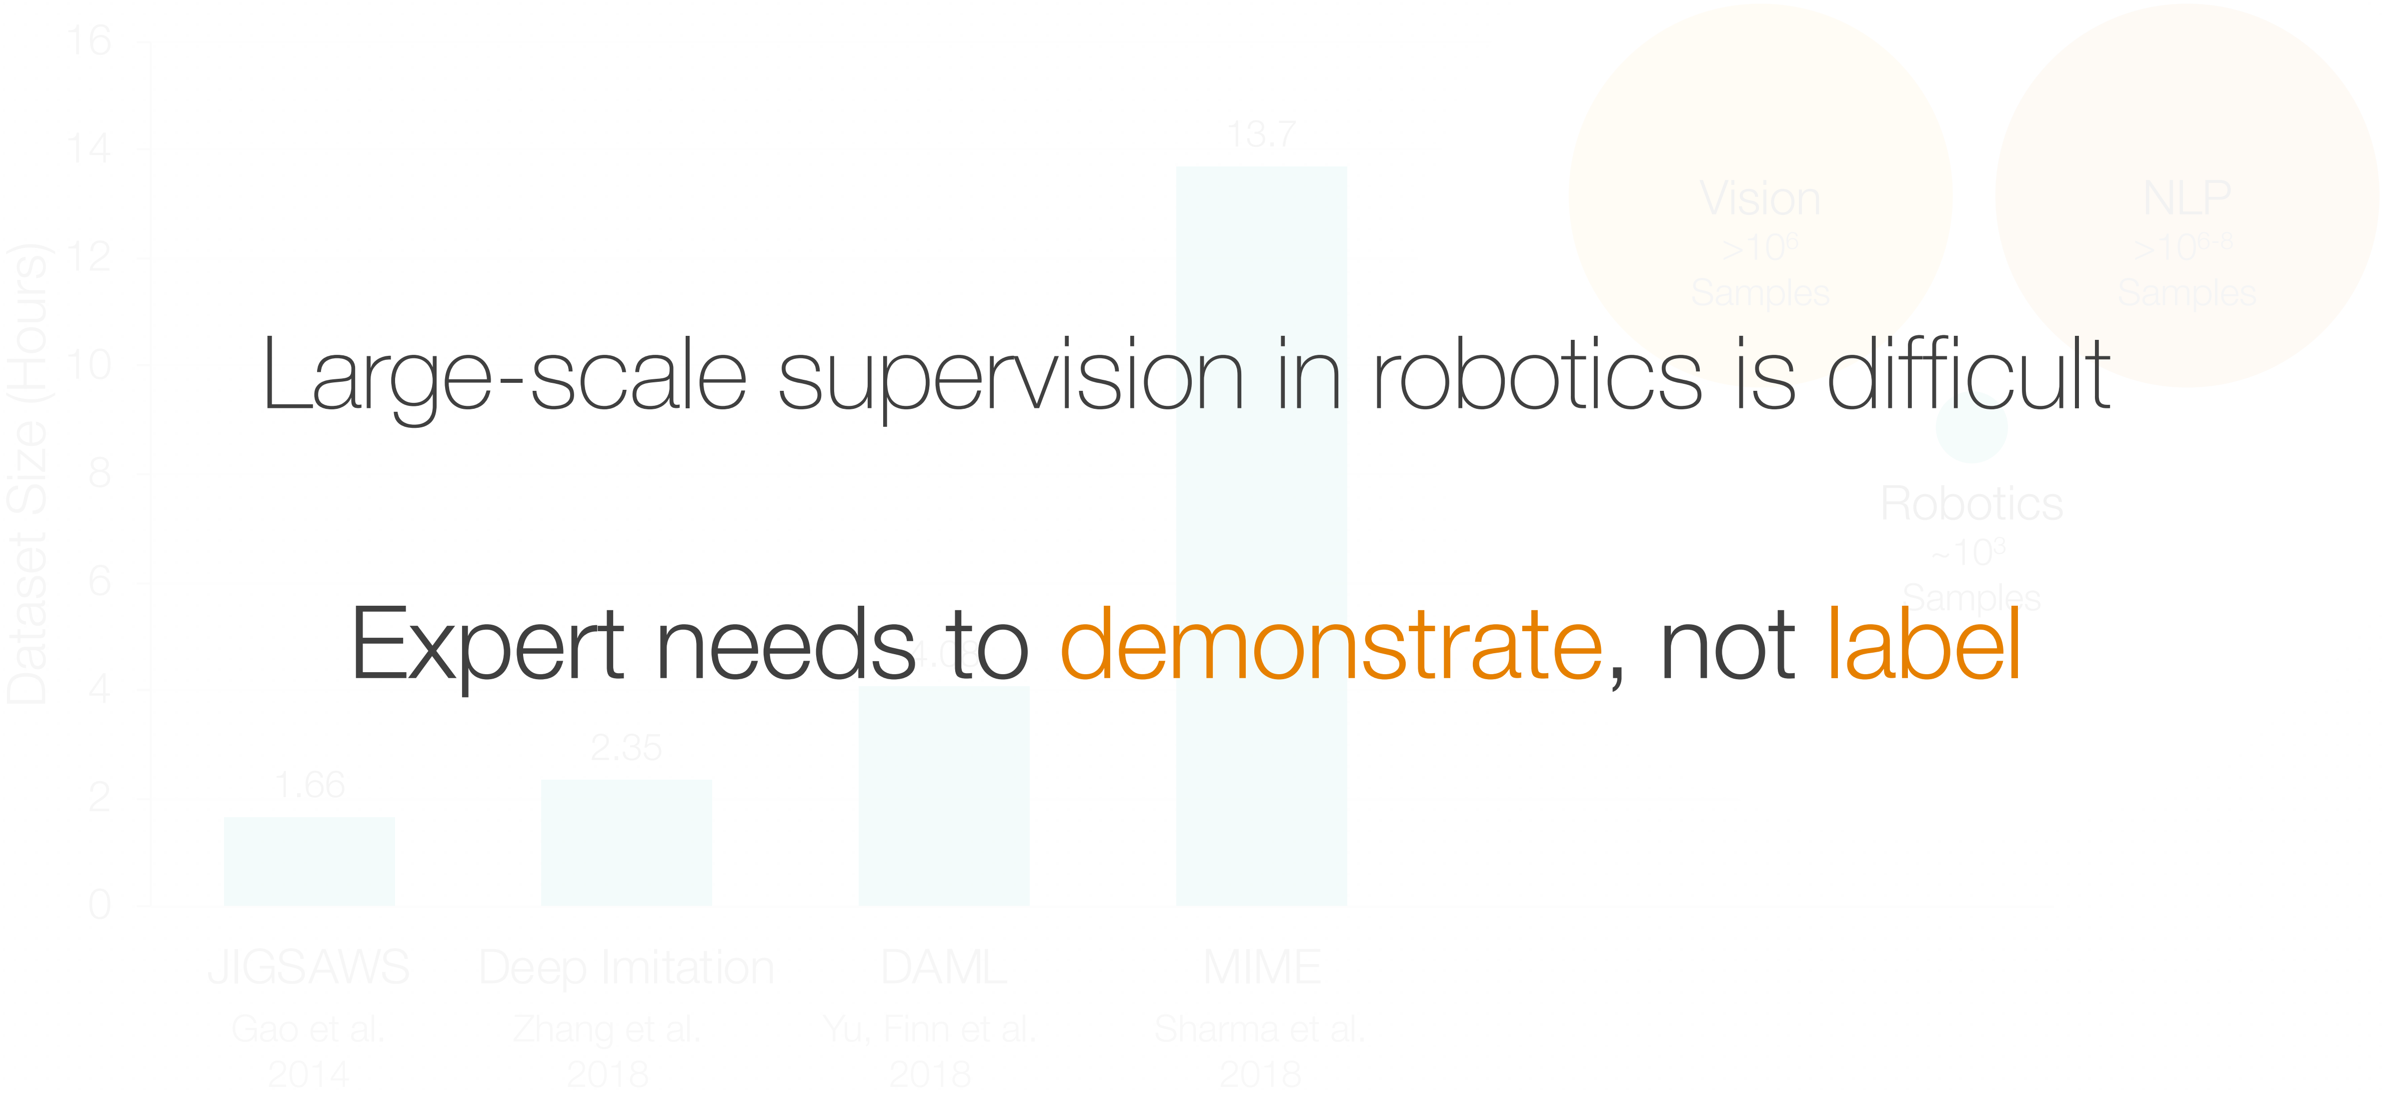

Large-scale supervision in robotics is difficult

Expert needs to demonstrate, not label

Data for Robotics: RoboTurk

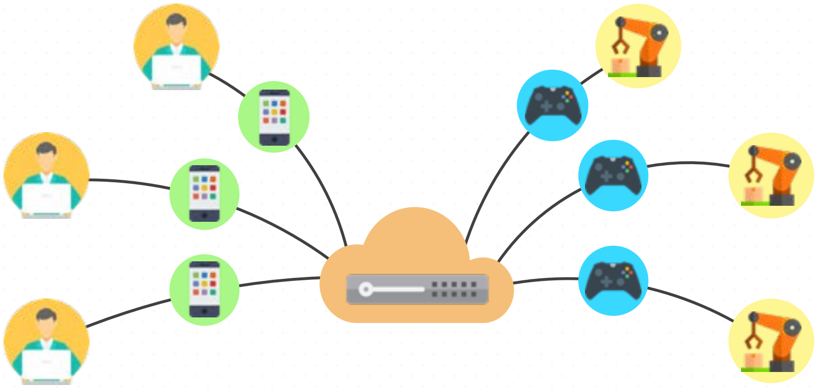

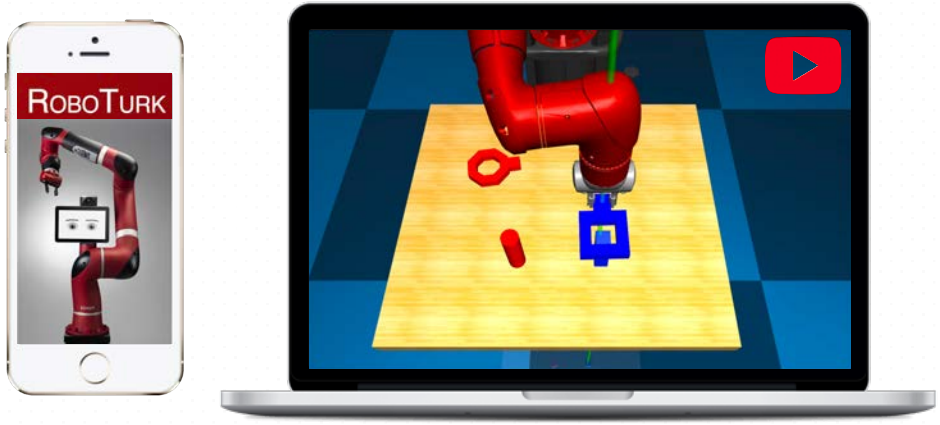

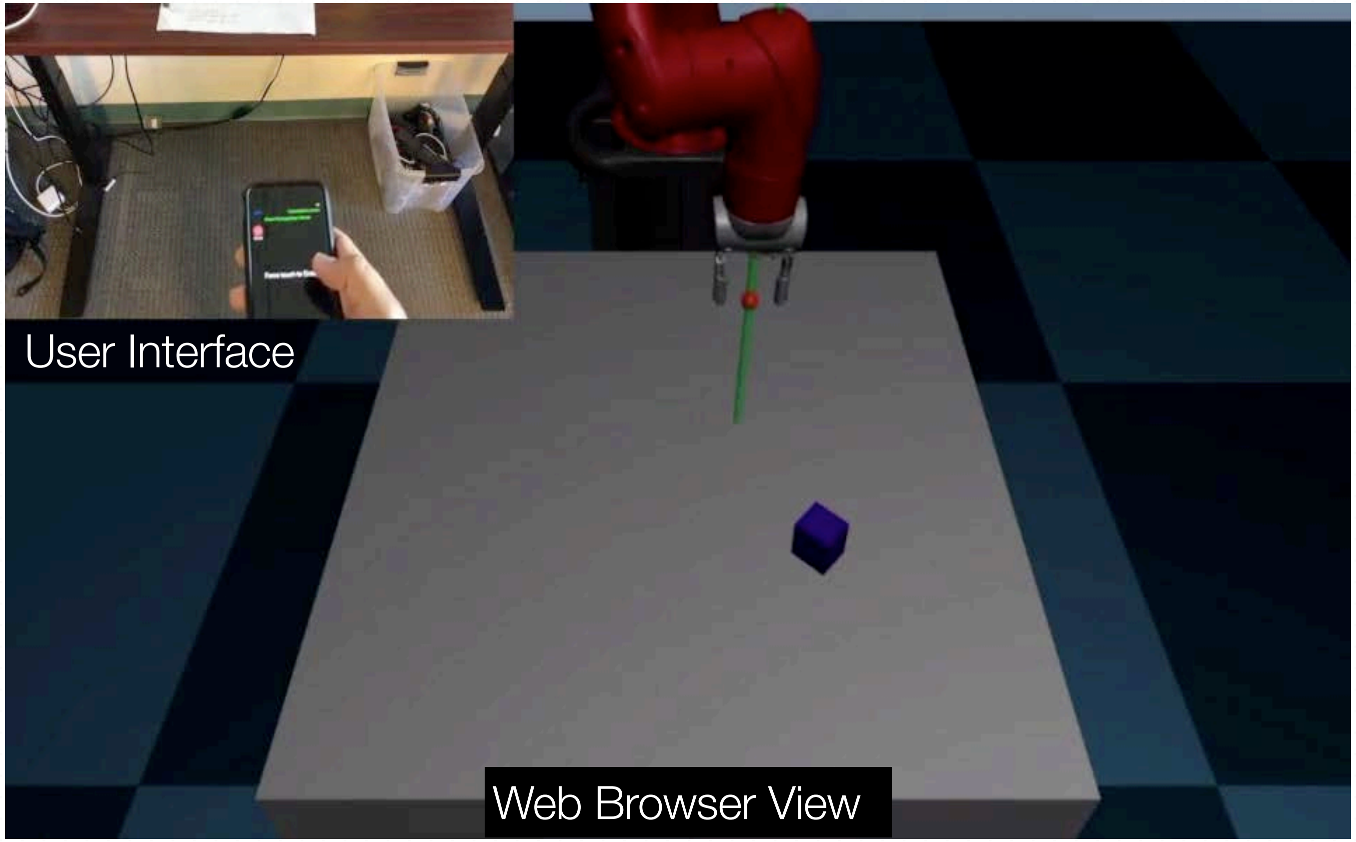

+ Scales easily with commodity hardware
+ Natural 6-DoF Free Space Control

[CoRL '18, IROS 2019]

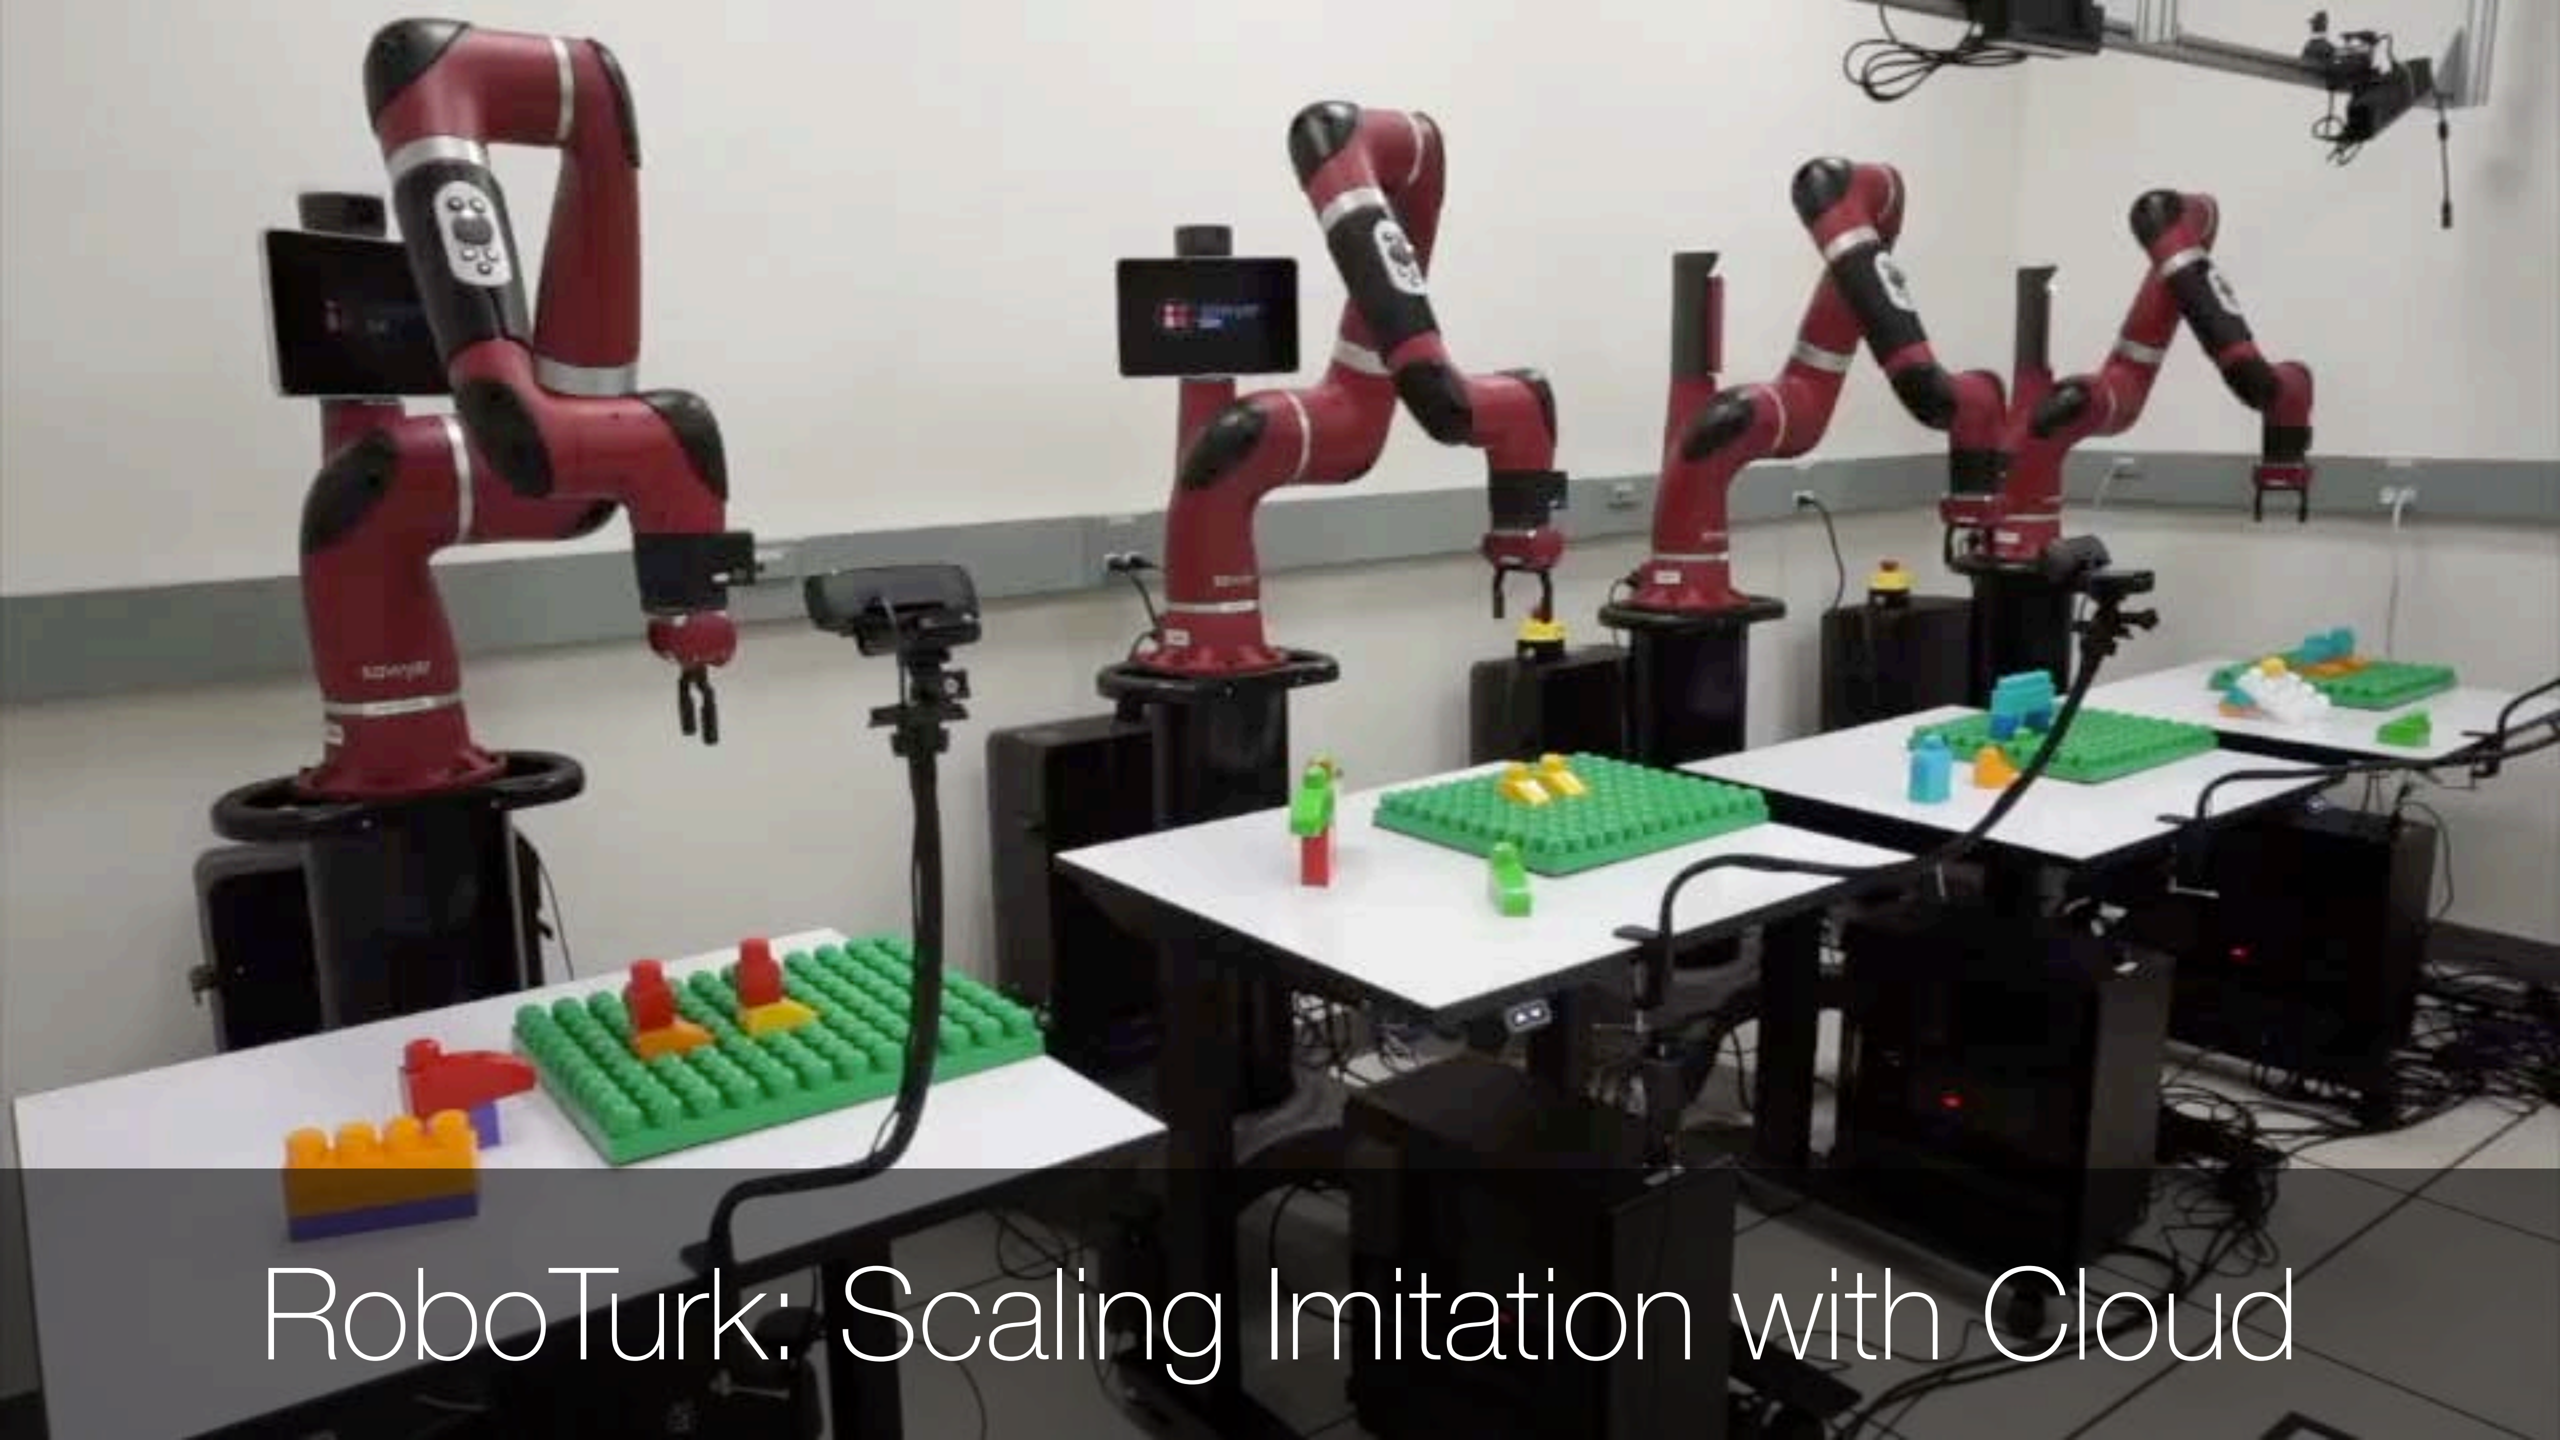
RoboTurk: Scaling Imitation with Cloud

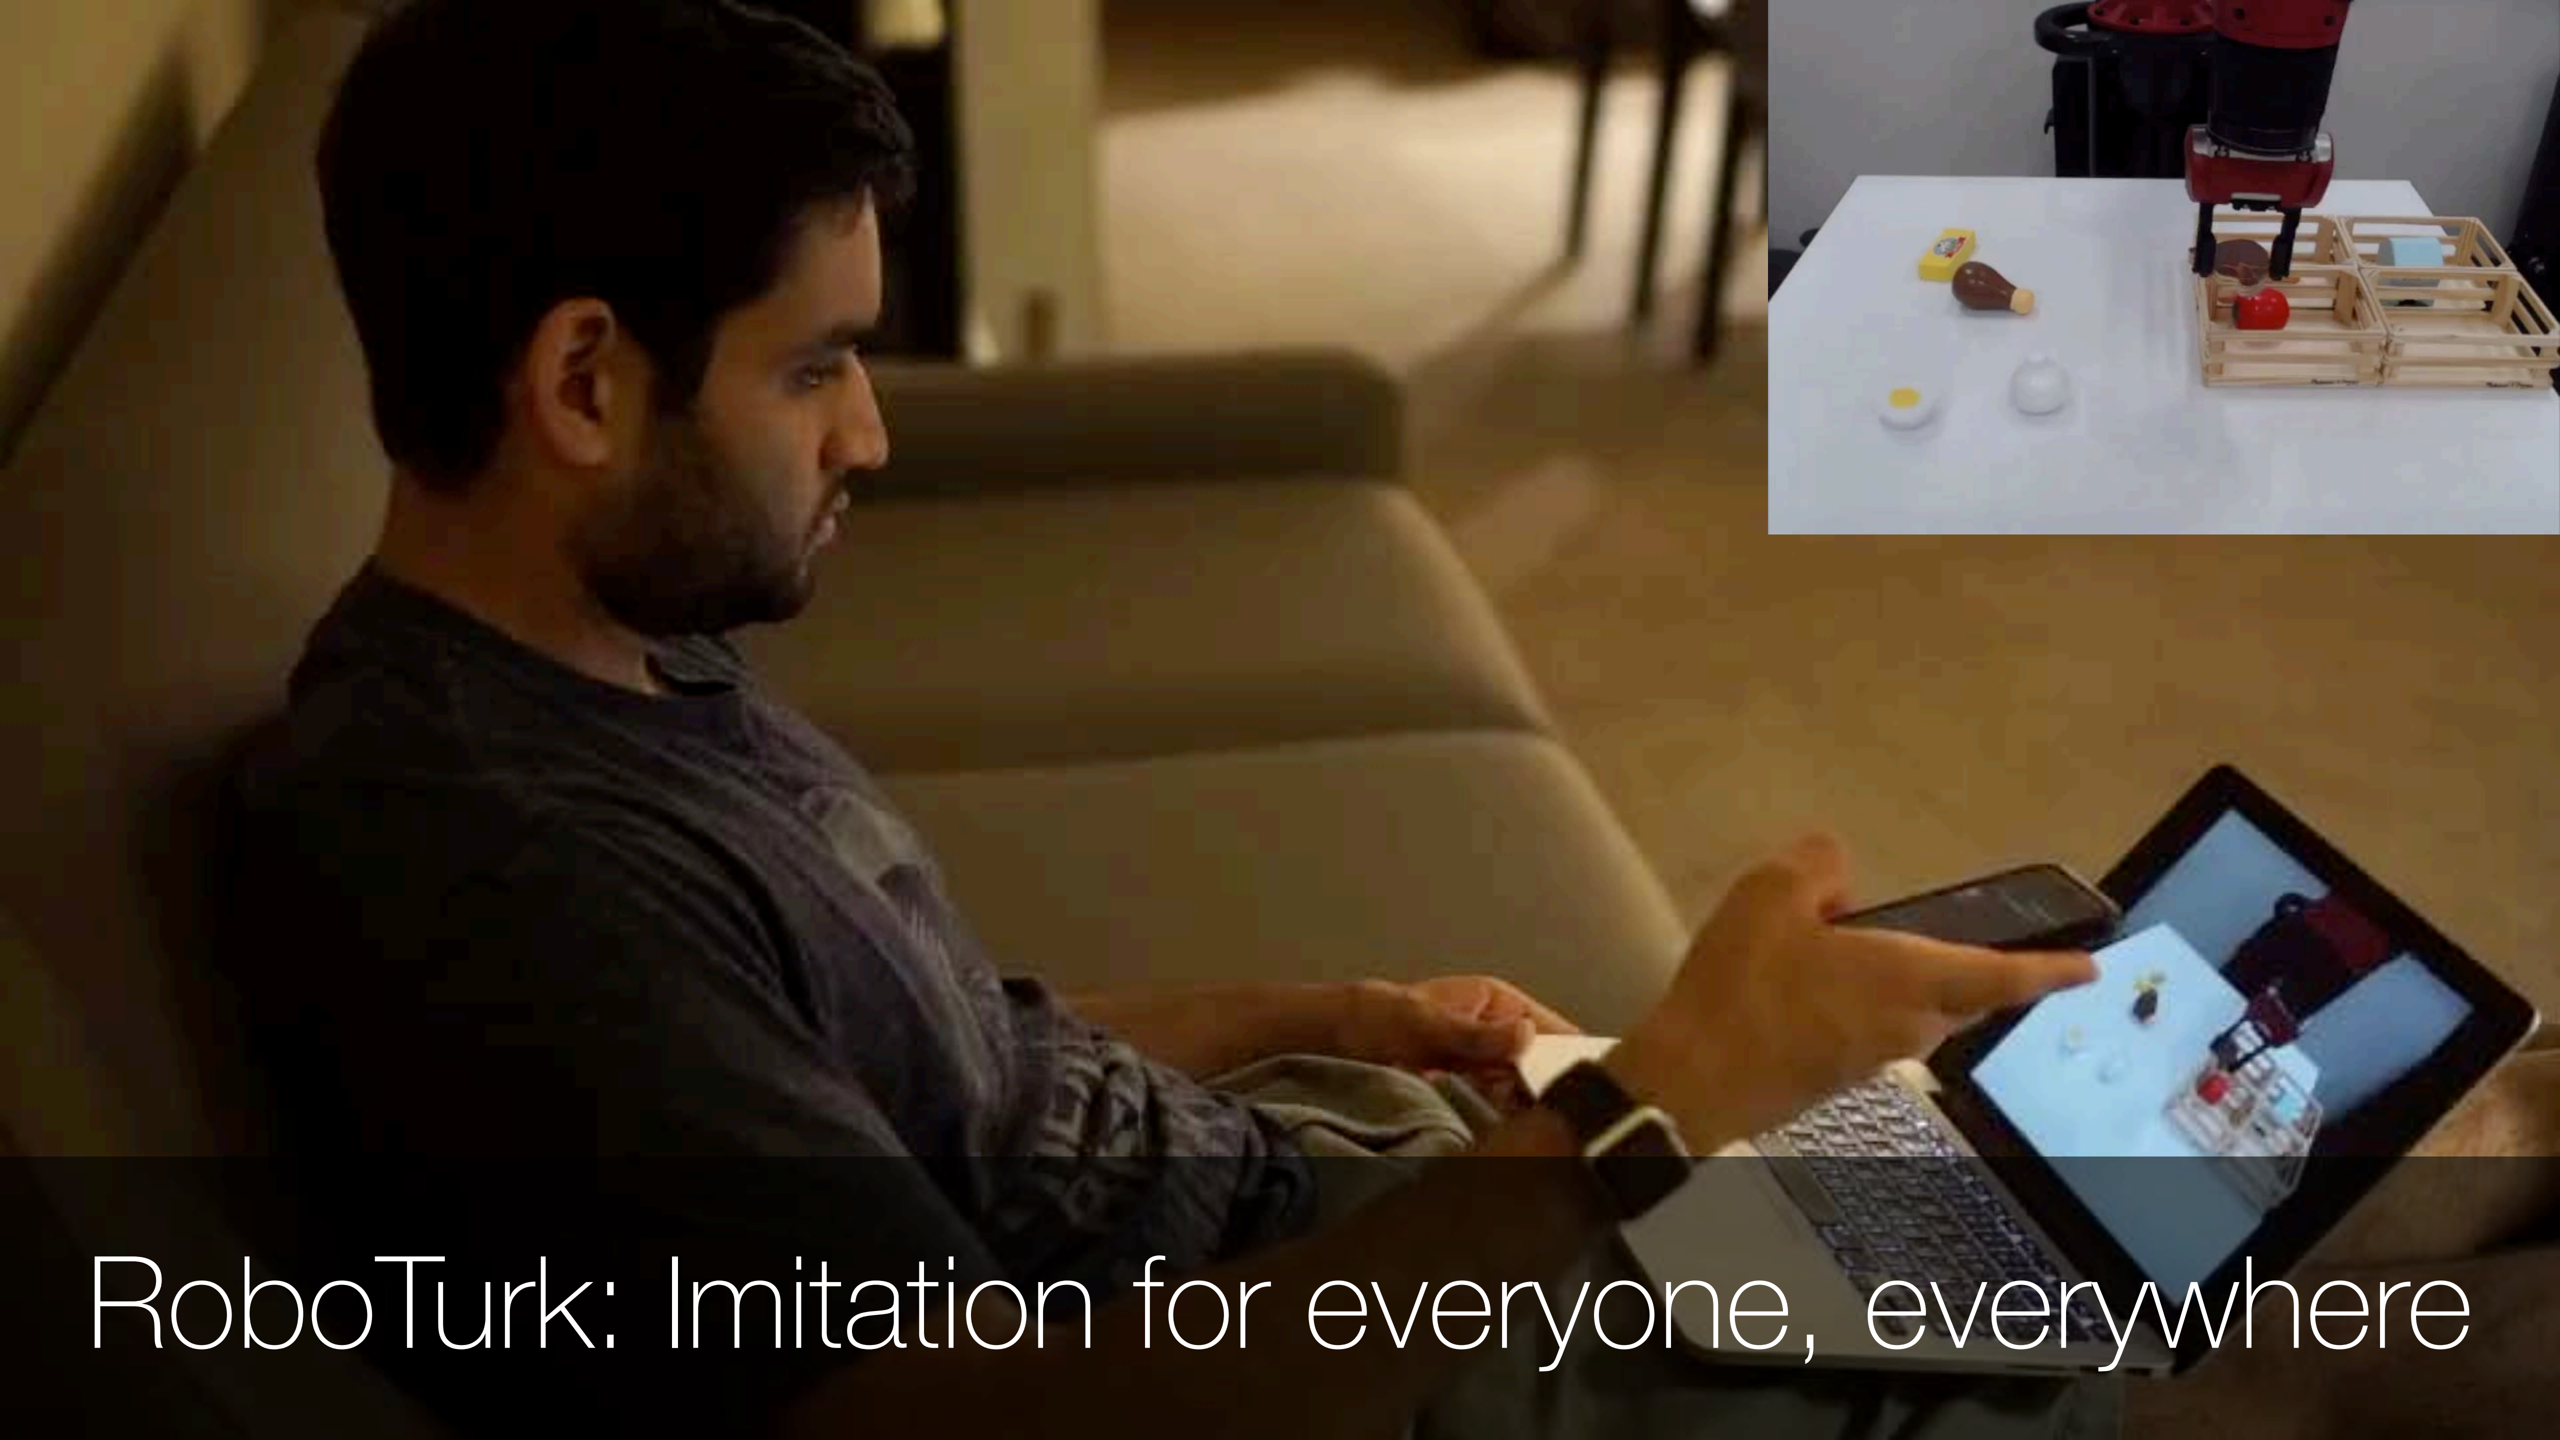

RoboTurk: Imitation for everyone, everywhere

RoboTurk Pilot Datasets

Simulated Data	Real Robot Data
137.5 hours of demonstrations	111 hours of robot demos
22 hours of total platform usage	1 week of data collection
3 dexterous manipulation tasks	3 dexterous manipulation tasks
3224 total attempted demos	2144 total demonstrations
15 novice, remote users	54 non-expert users

Data for Robotics: RoboTurk

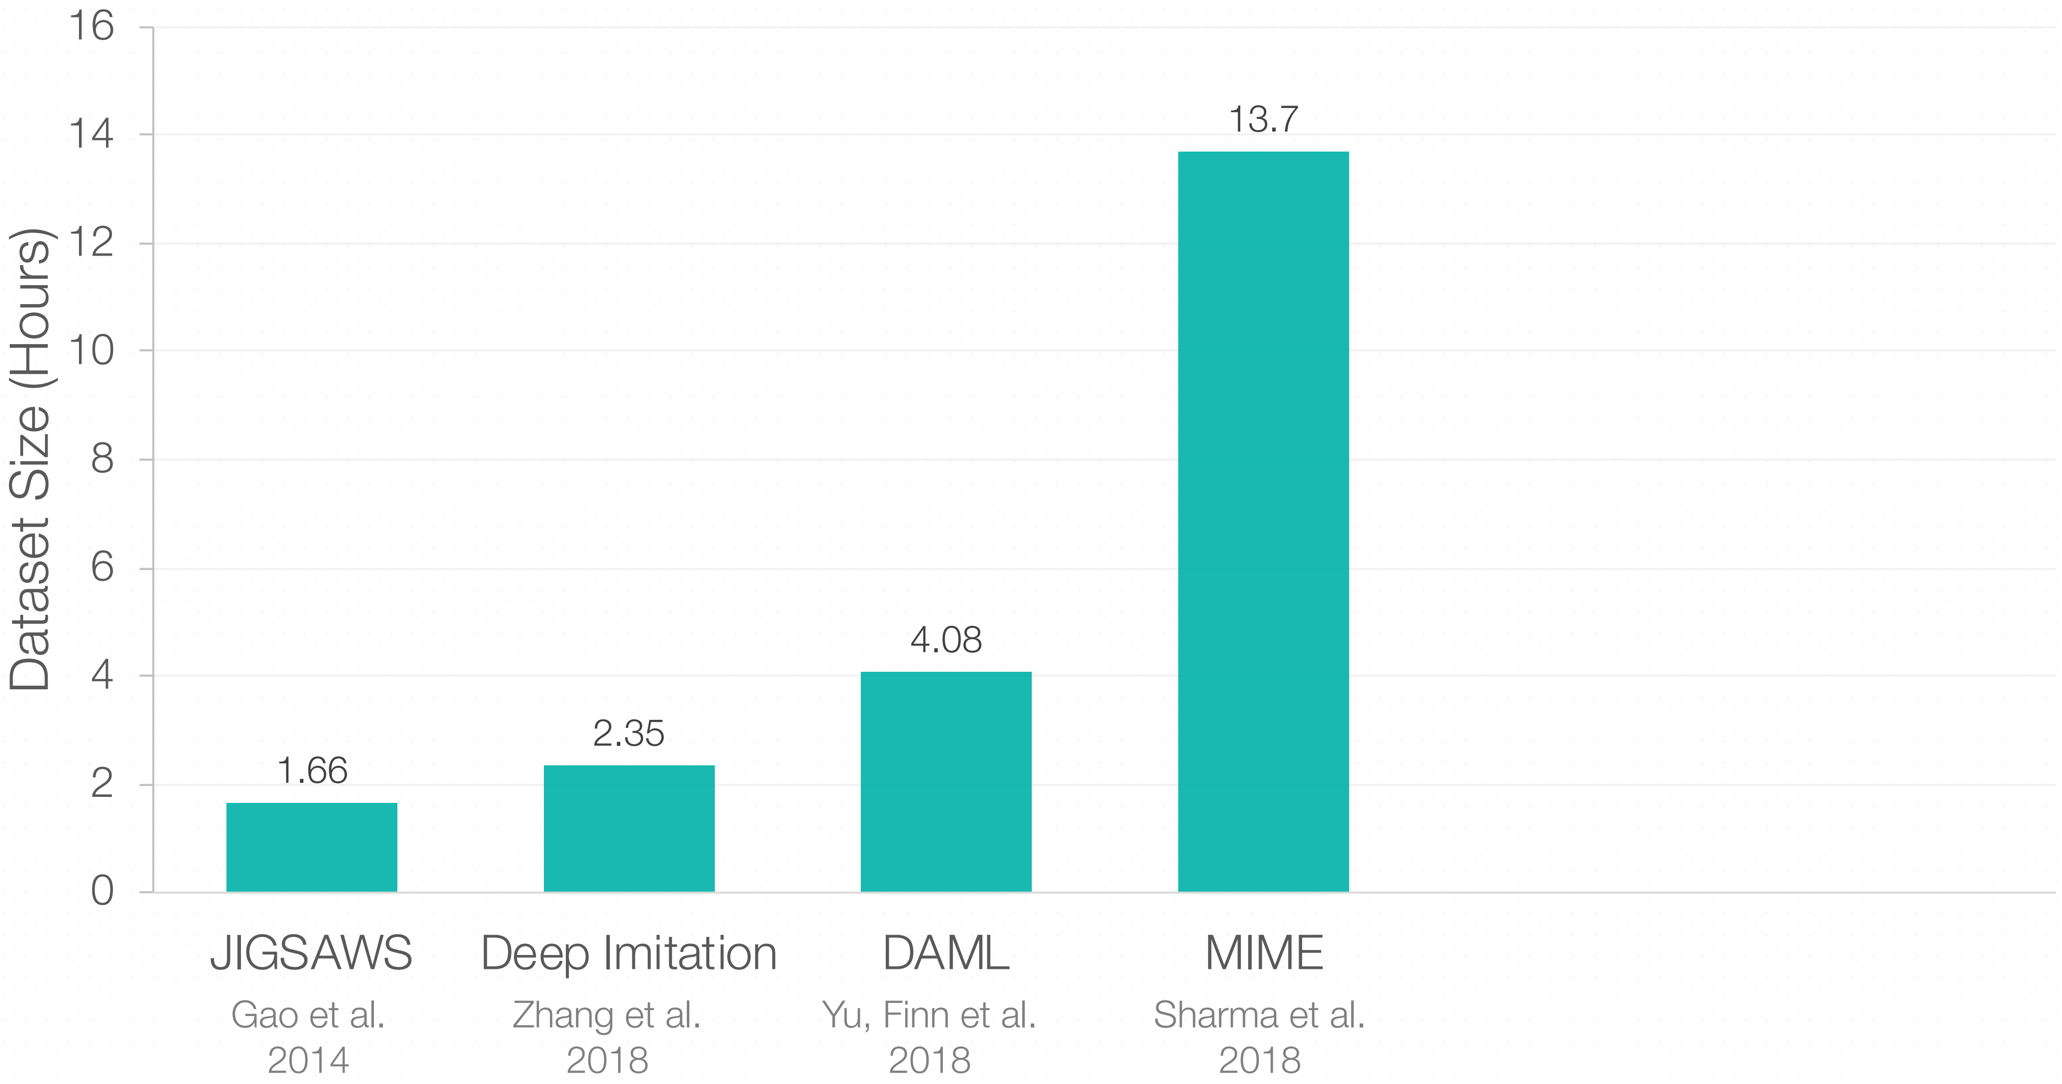

[CoRL '18, IROS 2019]

Data for Robotics: RoboTurk

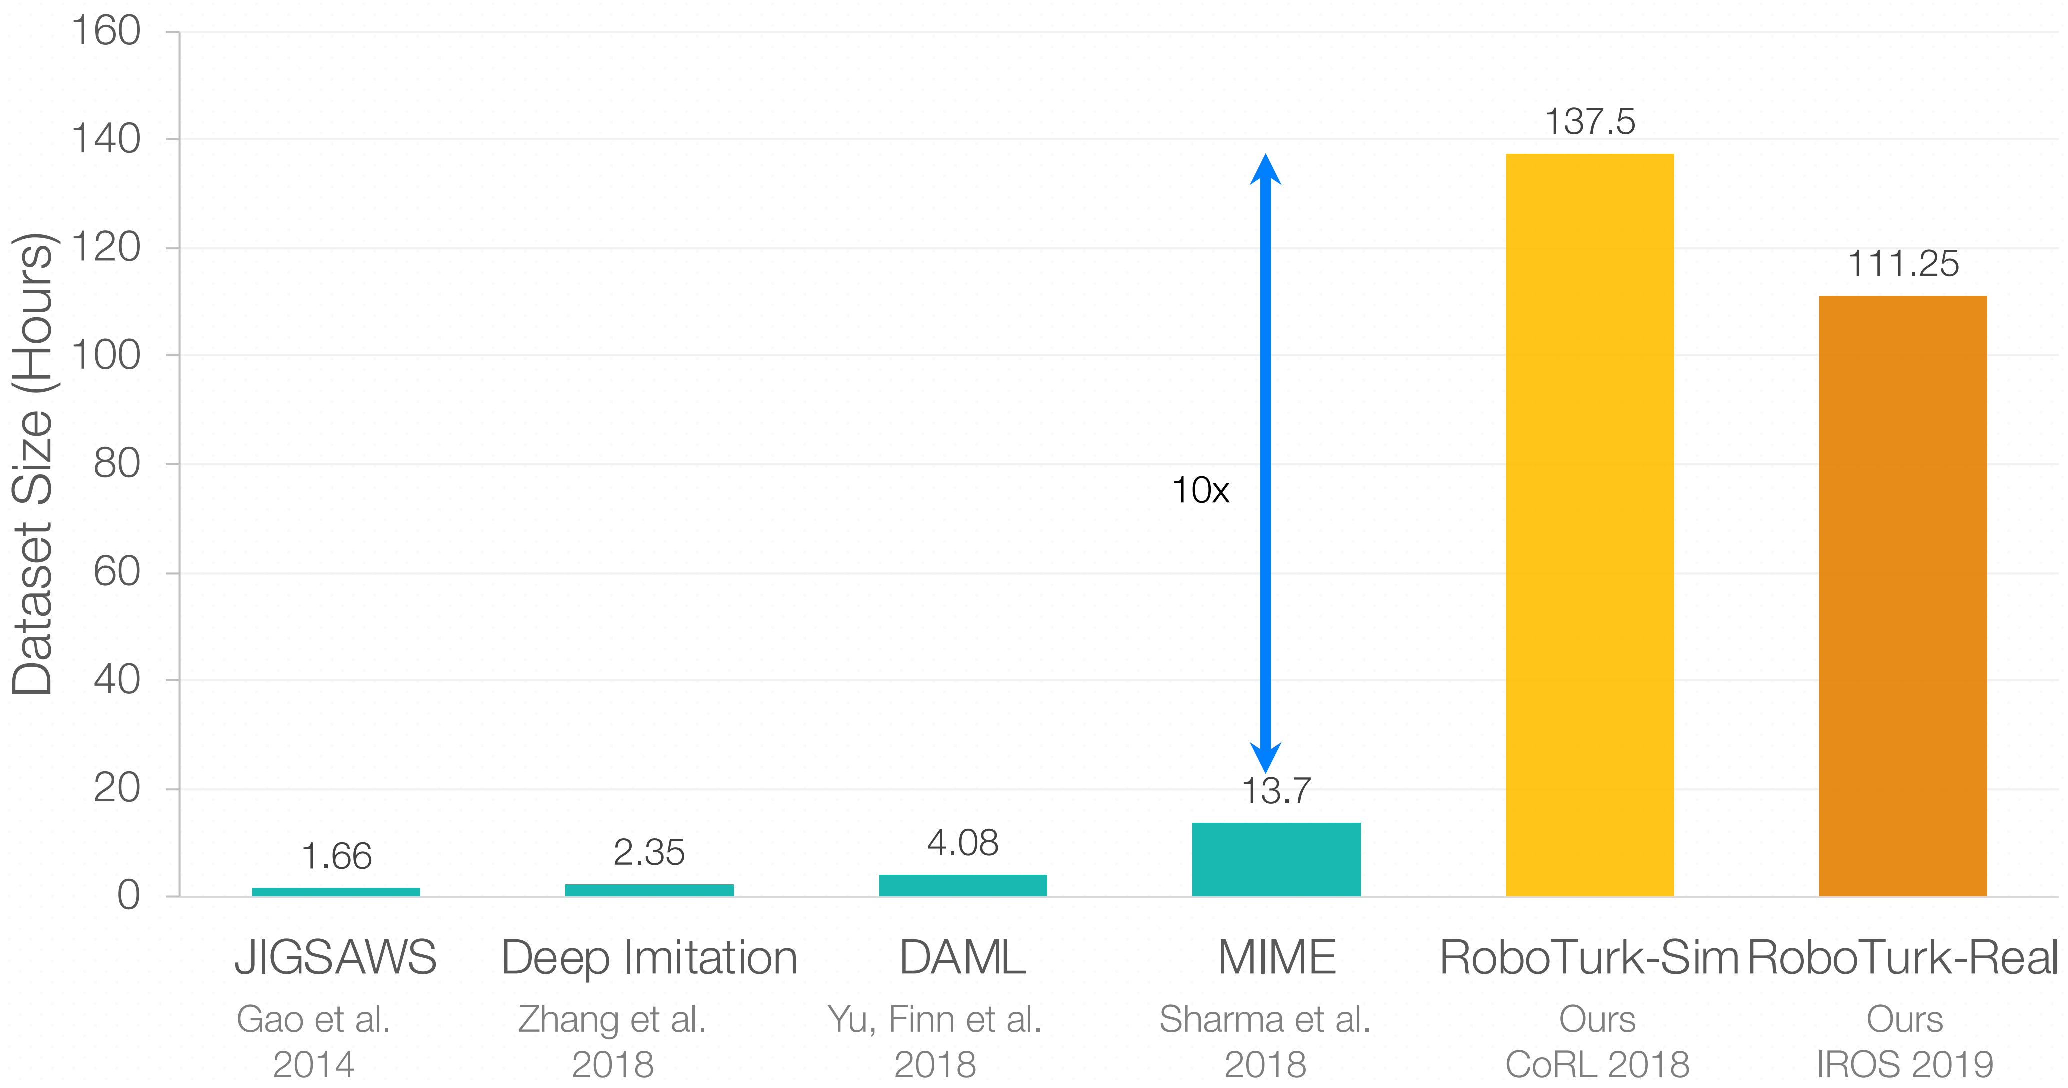

[CoRL '18, IROS 2019]

Data for Robotics: RoboTurk

Imitation + RL

Task Performance vs. Number of Demonstrations

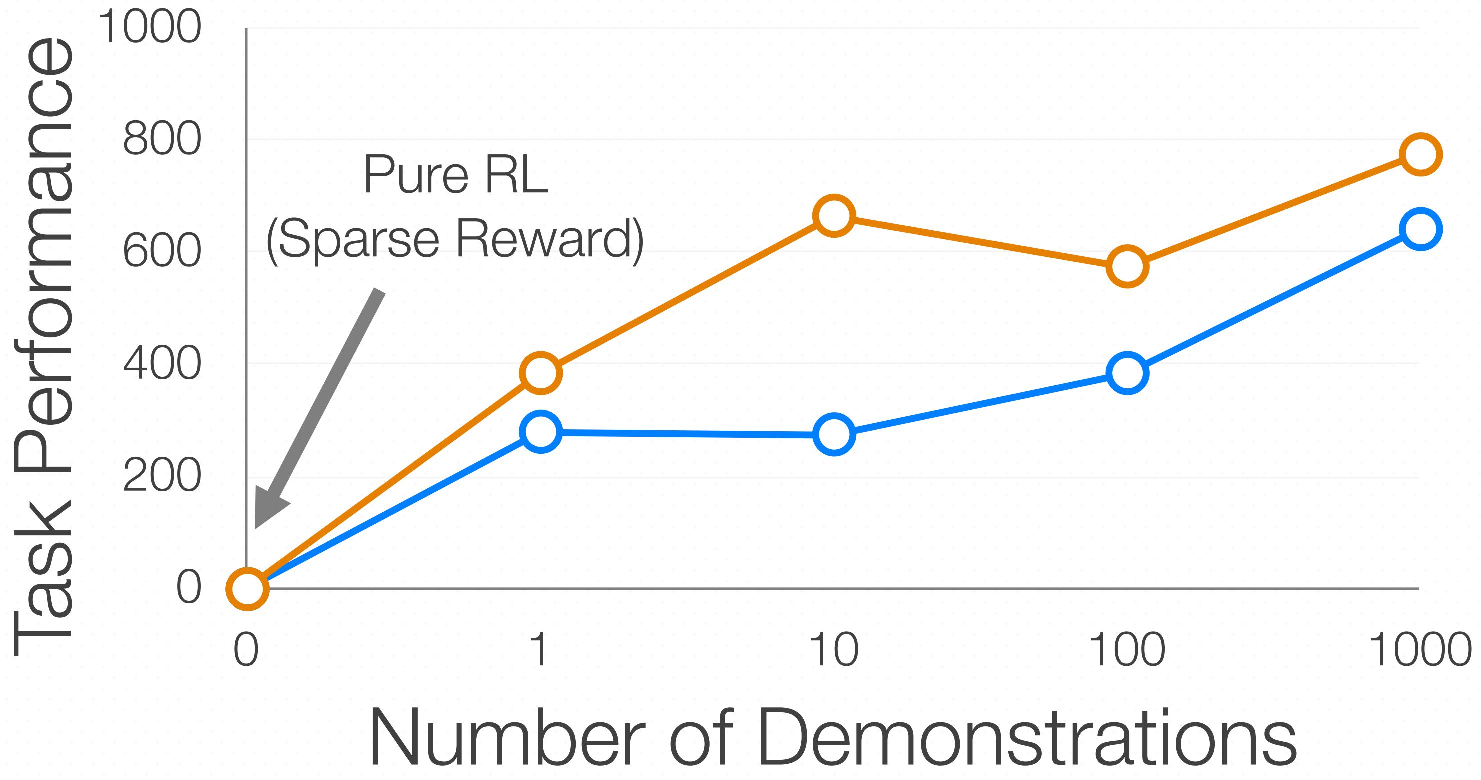

Trained Policy Rollout

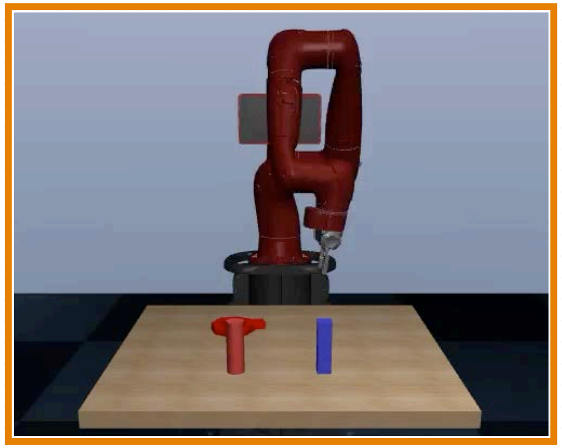

Nut Assembly

Bin Picking

Distributed PPO with curriculum: 32 workers, 24/48 hours

[CoRL '18, IROS 2019]

Data for Robotics: RoboTurk

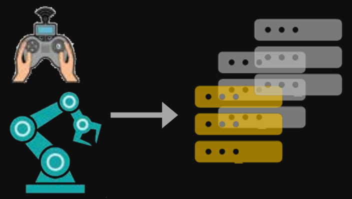

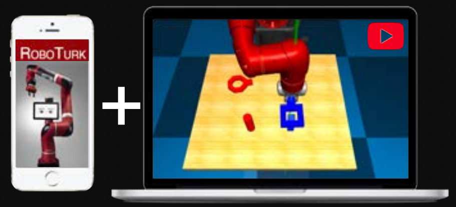

CoRL 2018, IROS 2019

Structured supervision for Robotics through scalable crowdsourcing can empower robot learning in complex tasks.

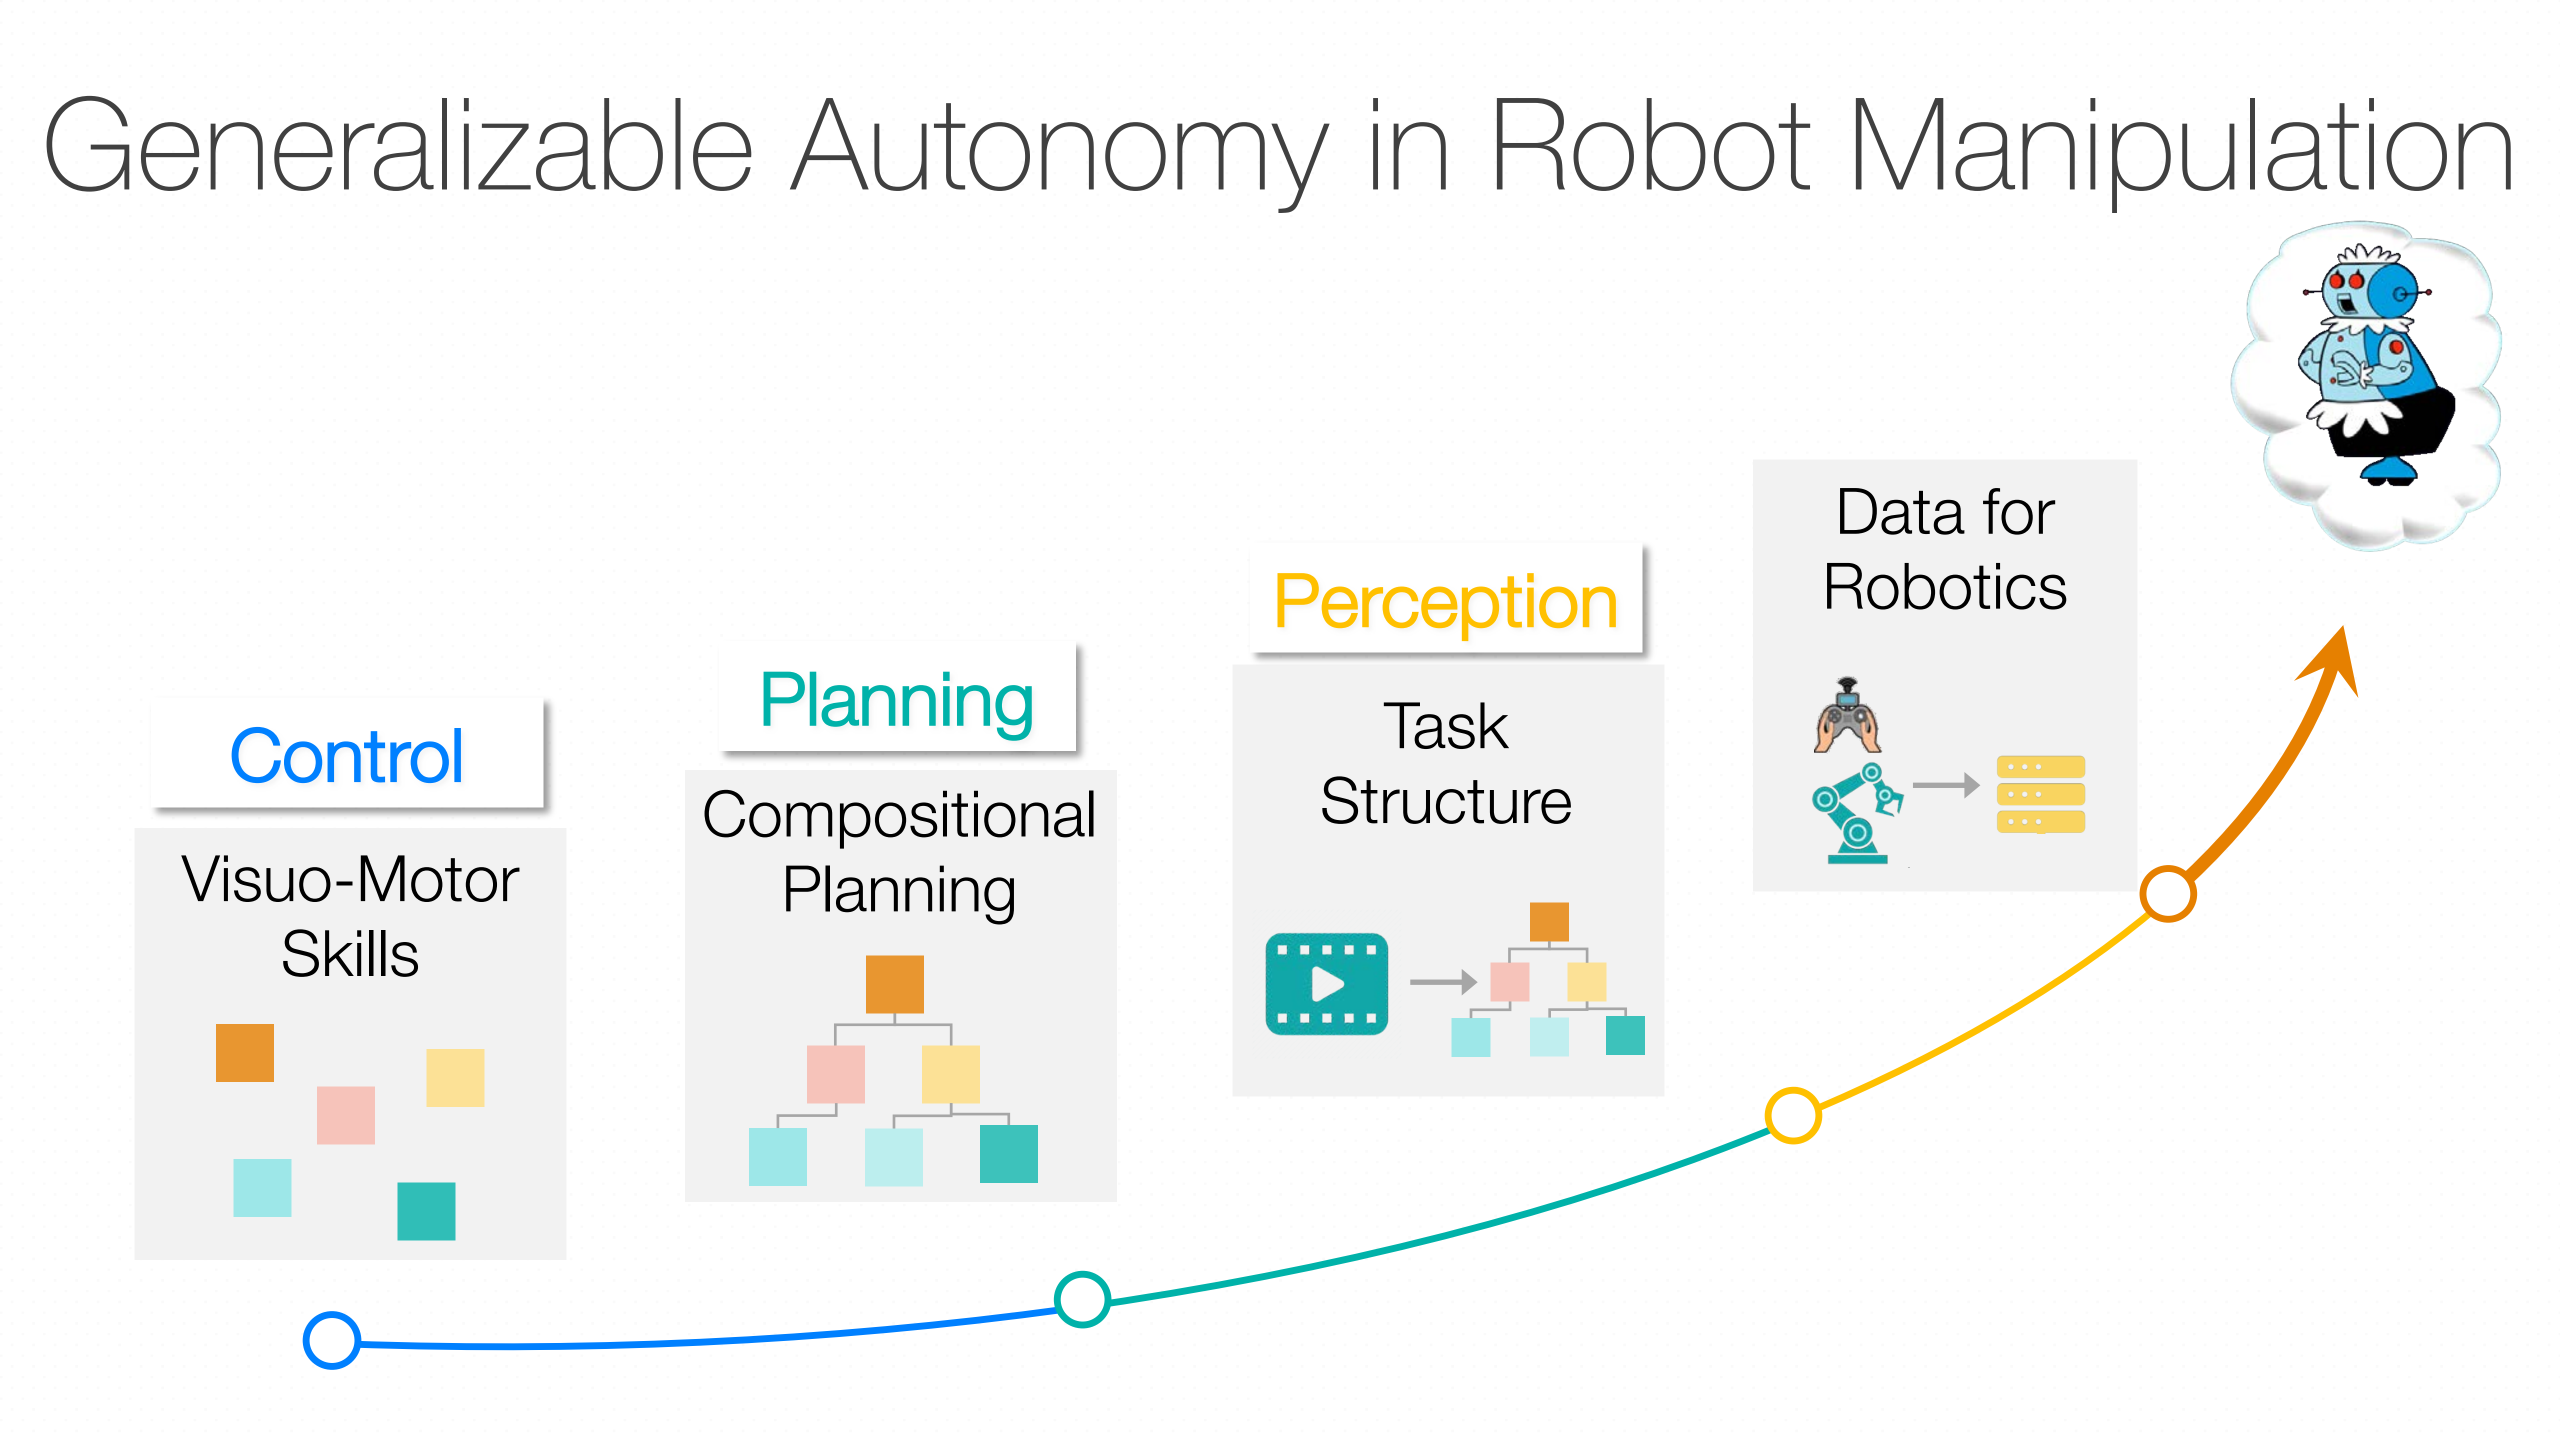
Generalizable Autonomy in Robot Manipulation
Control
Visuo-Motor Skills
Planning
Compositional Planning
Perception
Task Structure
Data for Robotics

Opportunity: Personal Robotics

Instructional Youtube Video

How to make Meatball Pasta?

Where / How should Rosie start?

What is the recipe?

How to execute the plan?

How to plan?

Reasoning for Physical Interaction

Understanding Purpose

Ideal Tool
During Training

Task-Based Tool Adaptation
During Execution

Grounding: So many ways to "make eggs"

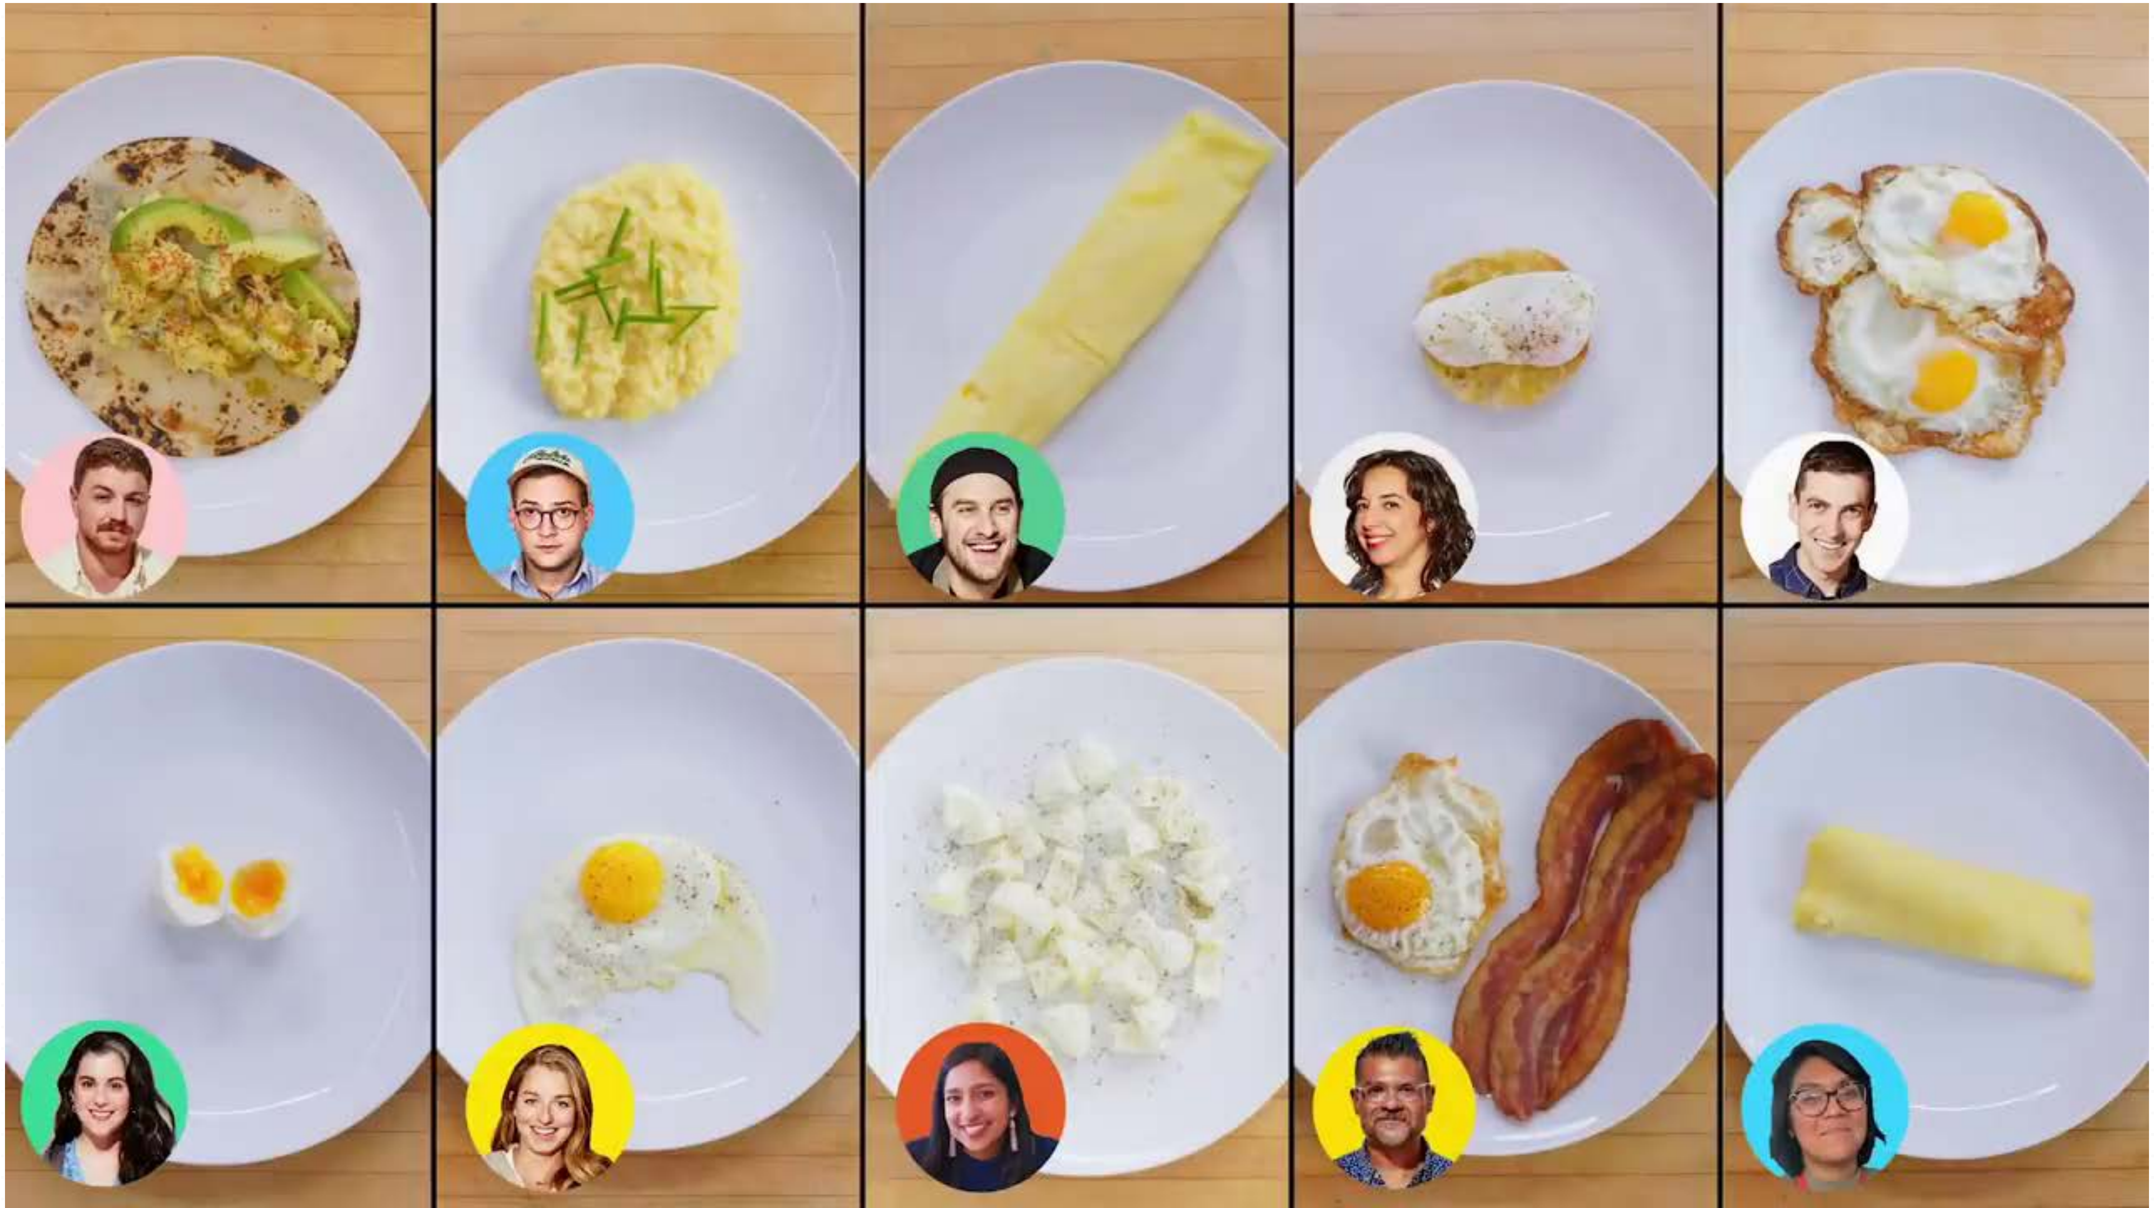

Generalizable Autonomy in Robot Manipulation

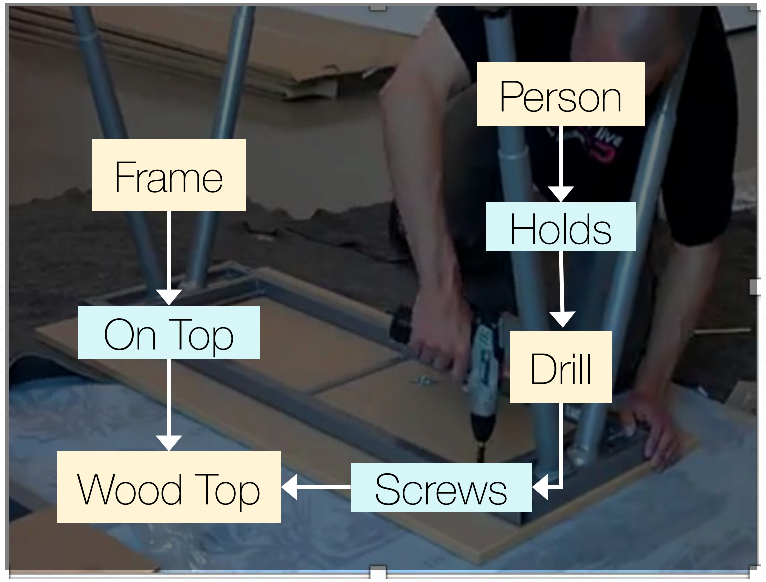

Higher-Order Semantics

What makes an object a hammer?

Failure

Success

State Change: Breaking Eggs

- Perception for Physical Interaction
- Reasoning through Learned Dynamics
- Transfer Learning with Formal Guarantees
- Continual Skill Adaptation & Accumulation

Generalizable Autonomy in Robot Manipulation

Learning with Structured Inductive Bias and Priors

- Efficiency and Generalization

- Combine: Domain Expertise + Data-Driven Methods

Generalizable Autonomy in Robot Manipulation

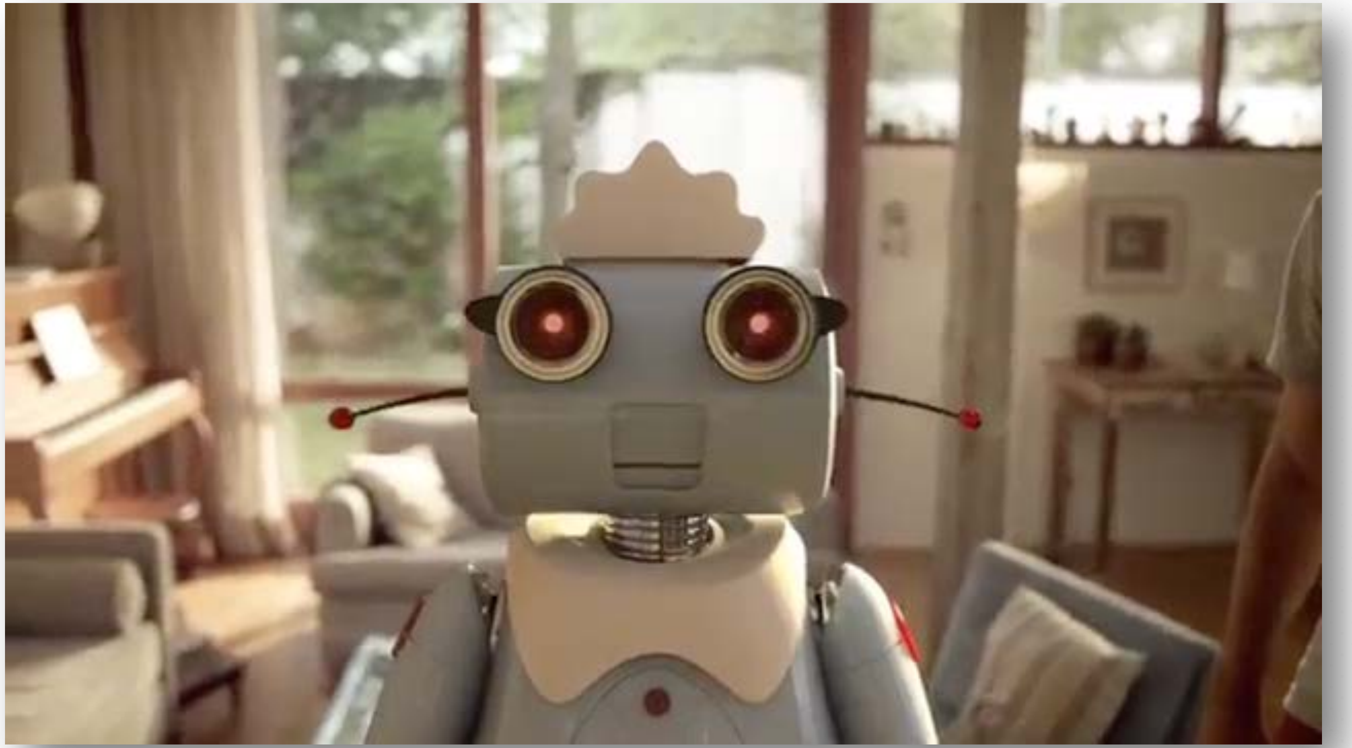

garg@cs.toronto.edu

@Animesh_Garg

Animesh Garg | UNIVERSITY OF TORONTO VECTOR INSTITUTE NVIDIA