Scaling simulation-to-real transfer by learning composable robot skills

Ryan Julian*, Eric Heiden*, Zhanpeng He, Hejia Zhang,
Stefan Schaal, Joseph Lim, Gaurav Sukhatme, Karol Hausman

USC University of Southern California

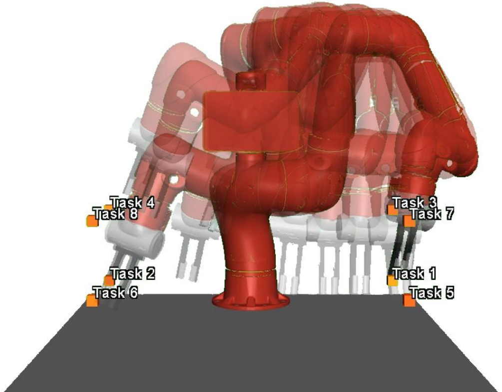

“Deep” RL for Robotics: *Where we are today*

Learning end-to-end

Levine, Sergey, et al. "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection." IJRR 37, 2018.

“Deep” RL for Robotics: *Where we are today*

Learning end-to-end

Levine, Sergey, et al. "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection." *IJRR*, 2018.

sim2real Transfer

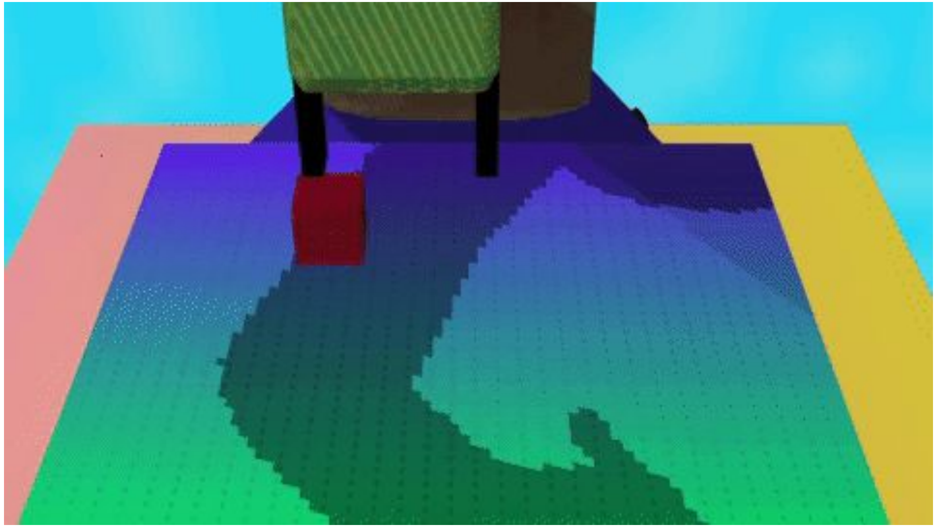

Marcin Andrychowicz, Filip Wolski, Alex Ray, et al. "Hindsight Experience Replay." *NIPS*, 2017.

The Problem

- End-to-end RL is **sample inefficient** per task
- sim2real methods rely on **explicit alignment** or closing the **“reality gap”**
- RL-learned skills are **difficult to reuse** and **compose**

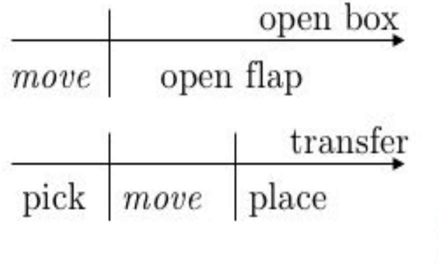

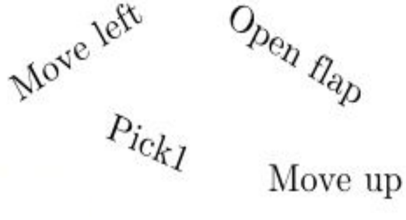

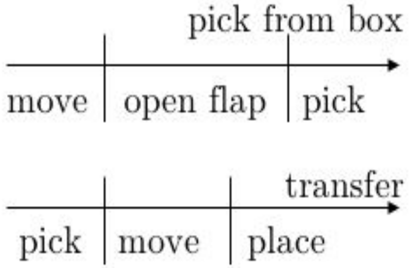

Motivating Example: Pick and Place

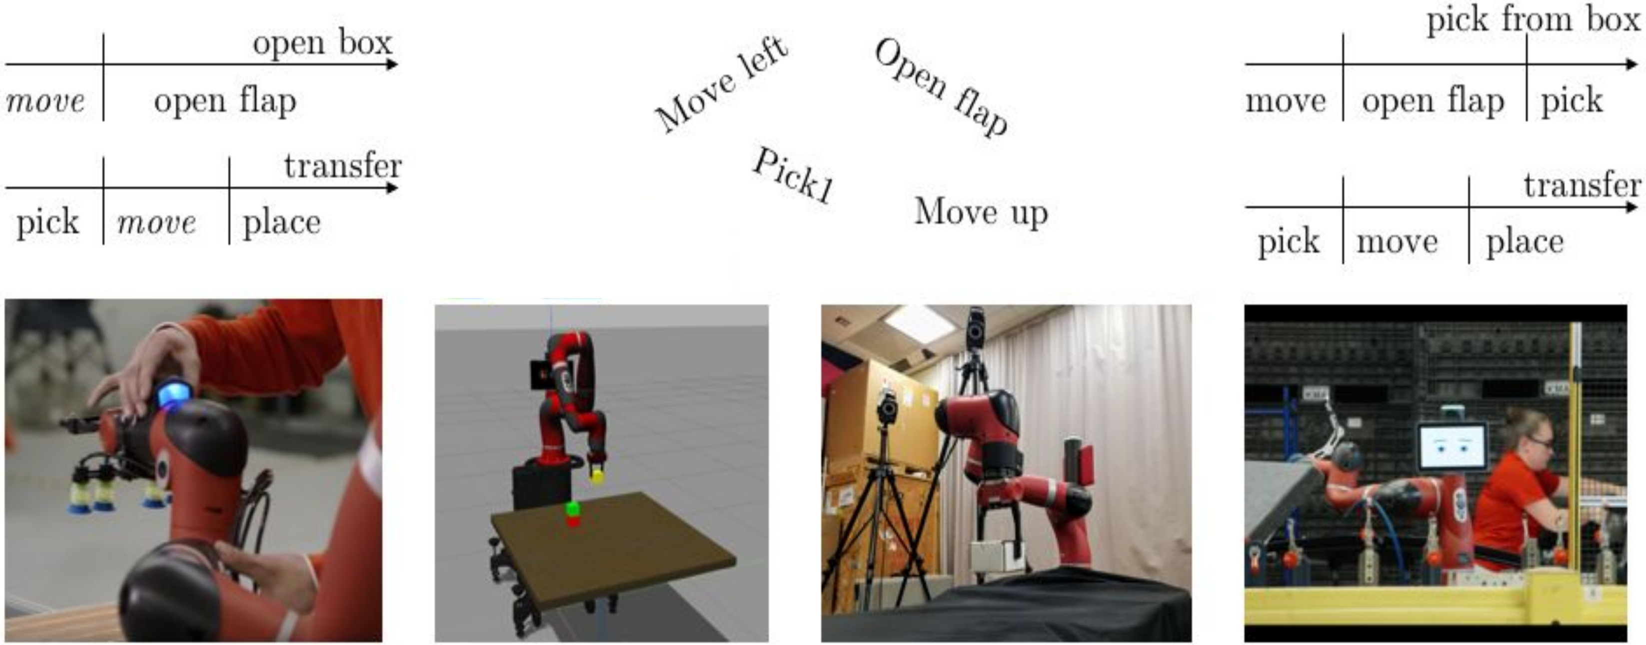

Our Approach

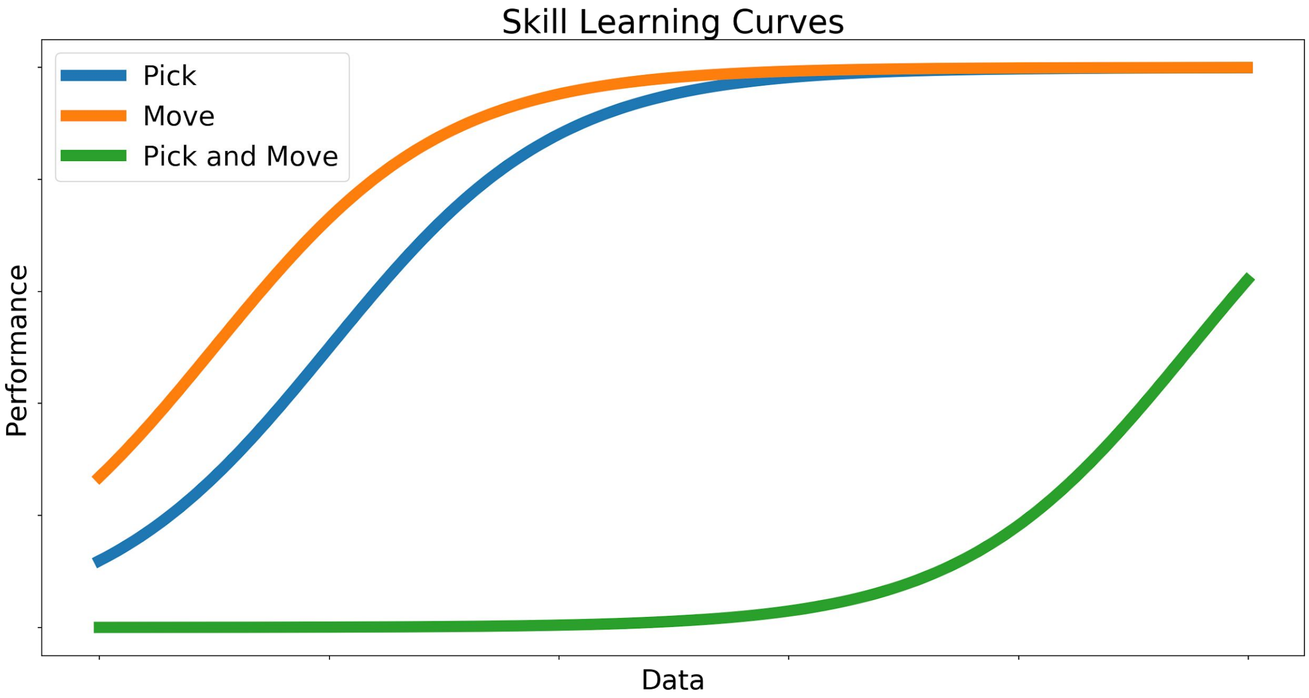

Decomposability → Reusable skills → Simplicity

Skill Embedding Example: Pick, Move, Place

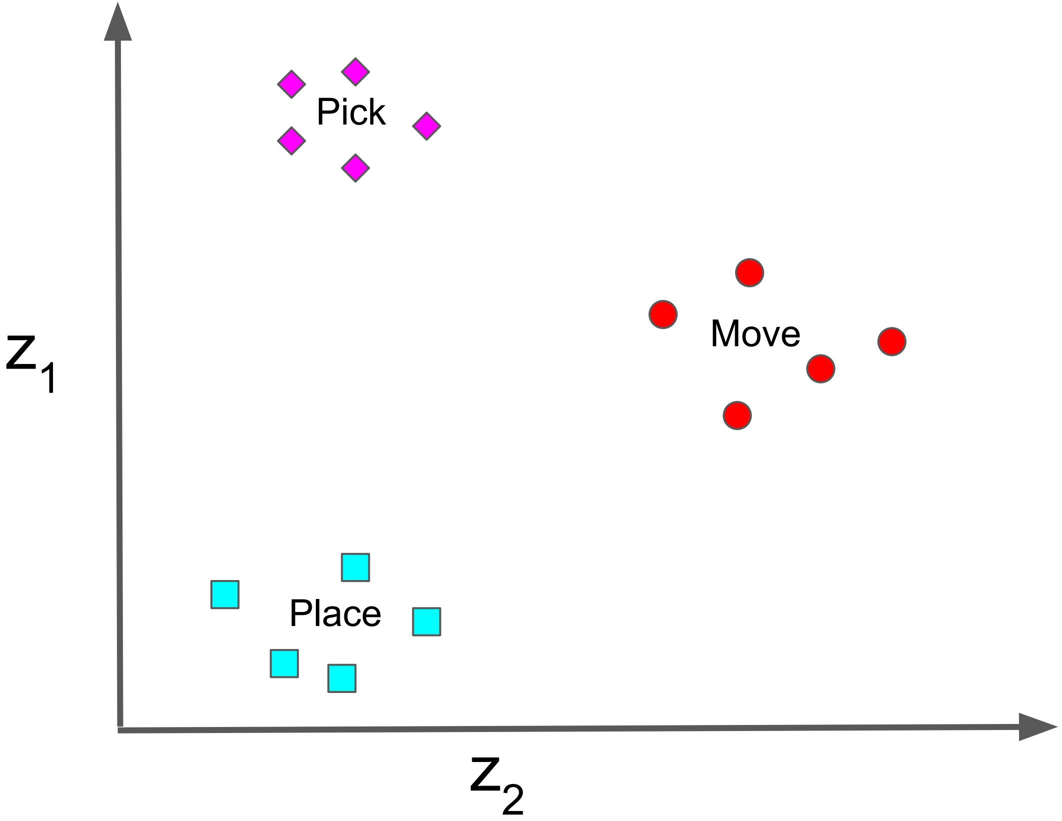

Skill Embedding Example: Pick, Move, Place

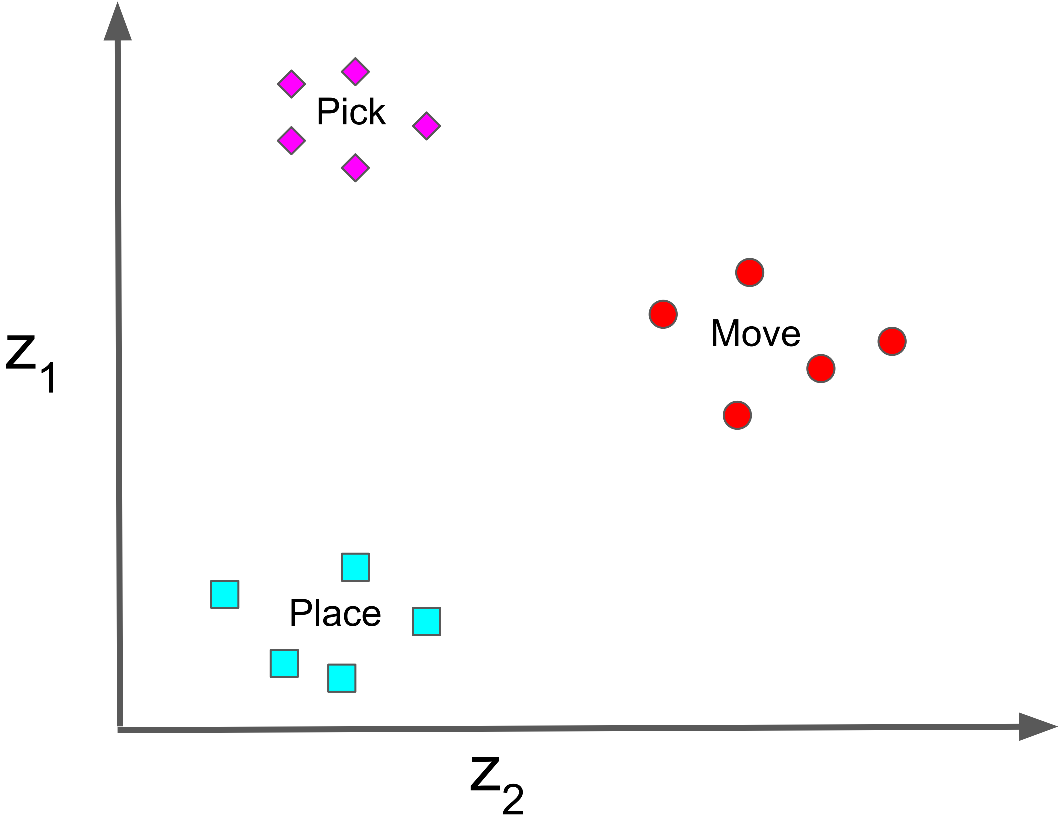

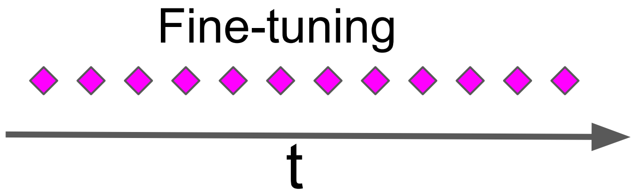

Skill Embedding Example: Pick, Move, Place

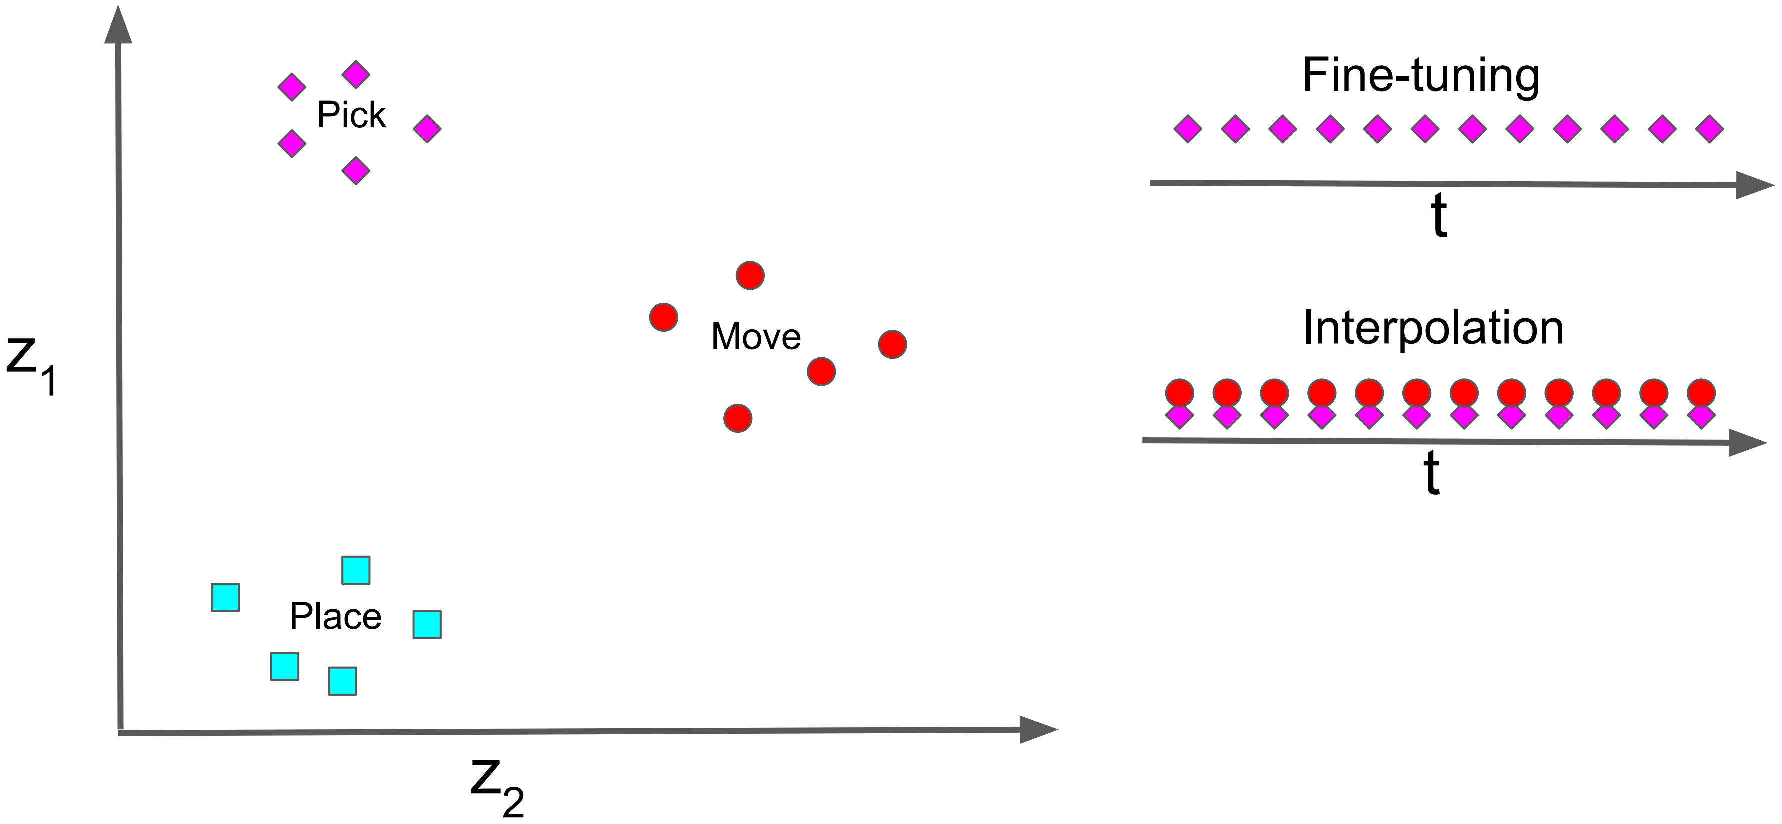

Skill Embedding Example: Pick, Move, Place

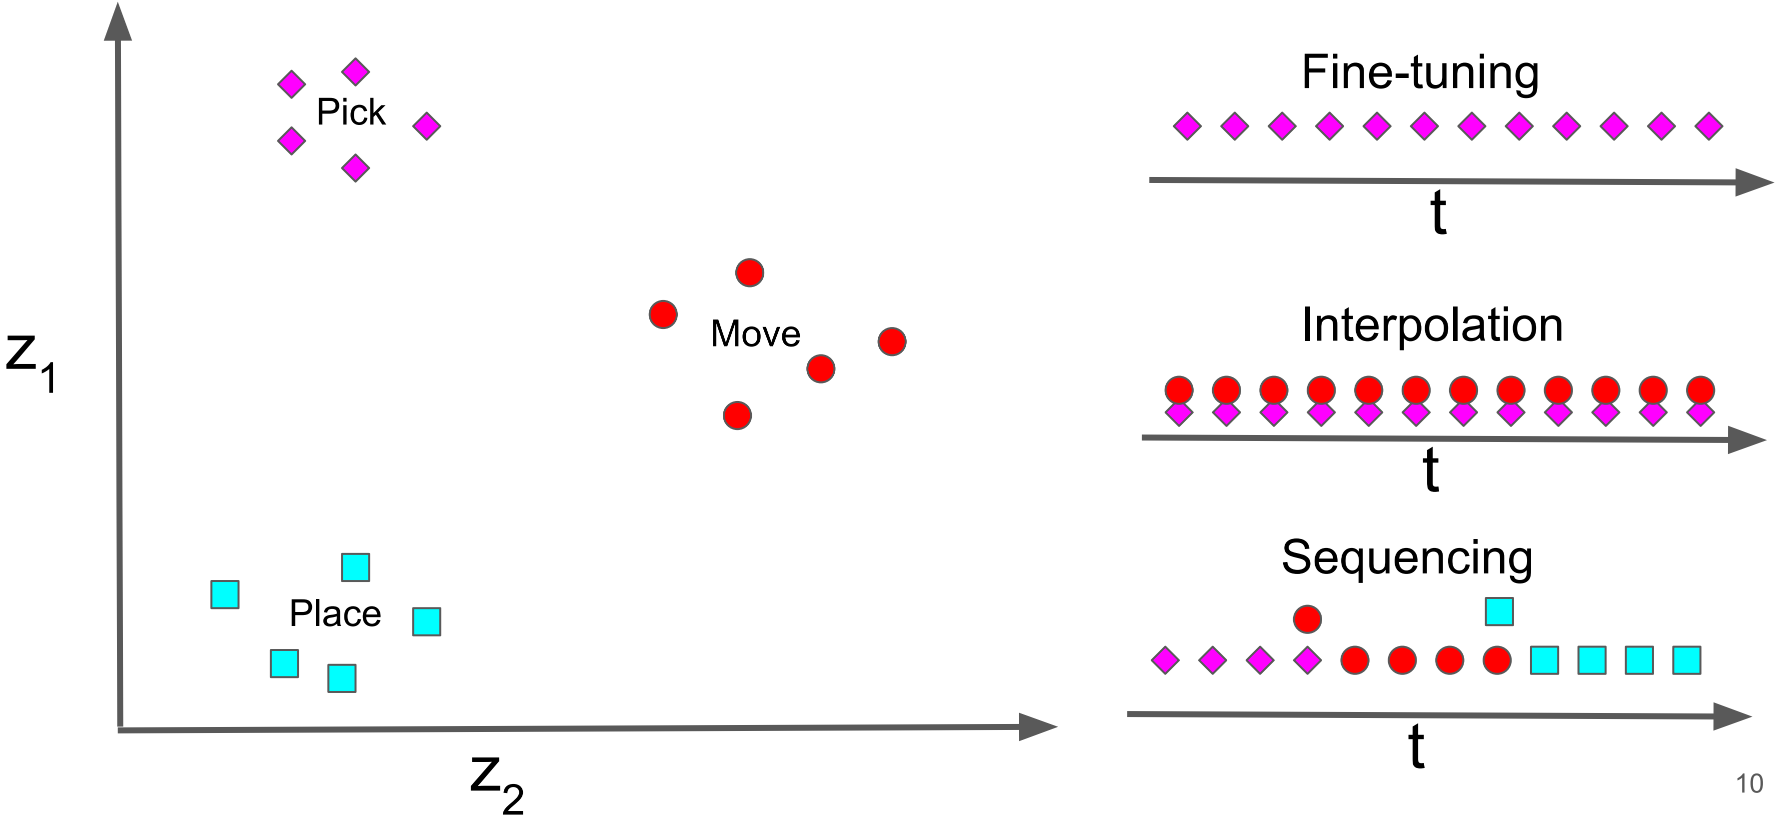

Learned skill space (embedding)

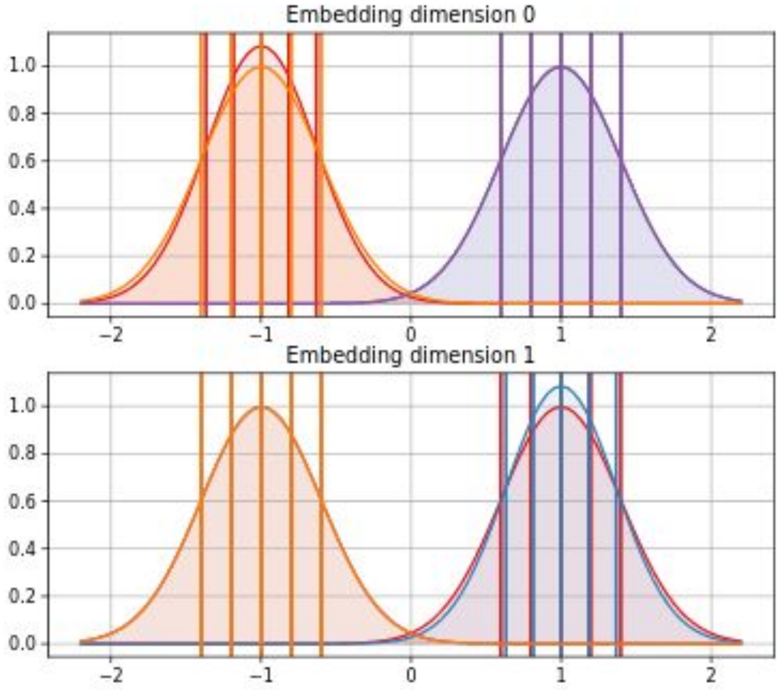

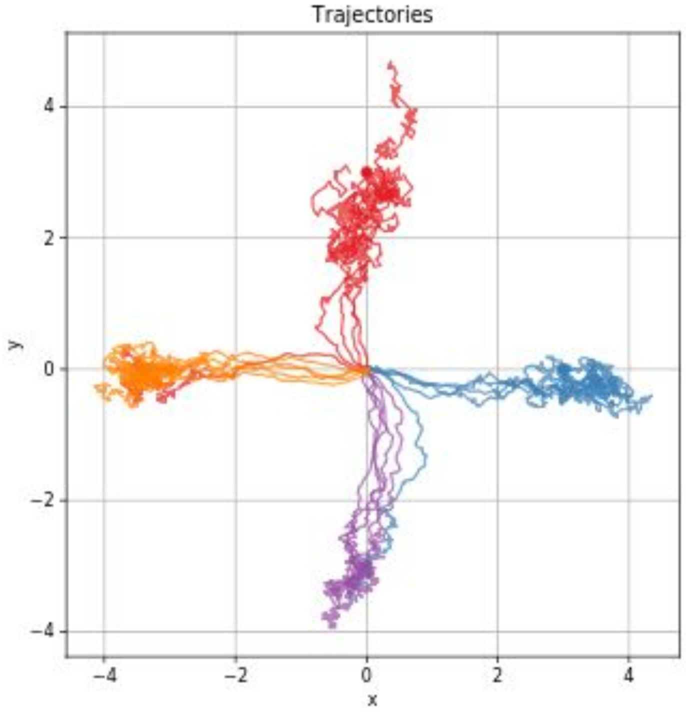

Training evolution of skill embeddings

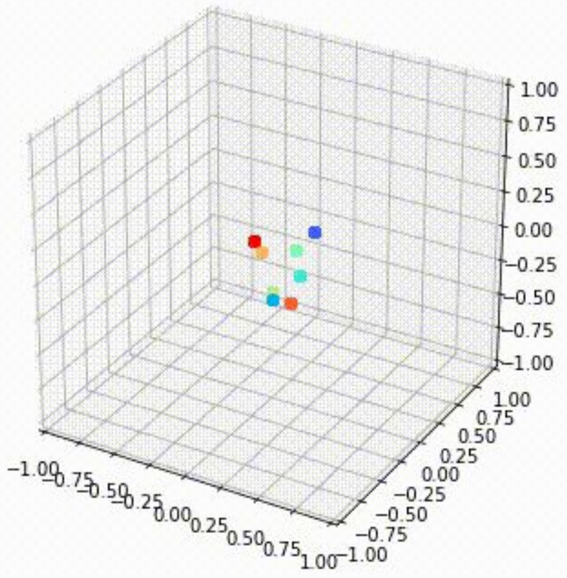

8-goal Reacher

Training evolution of skill embeddings

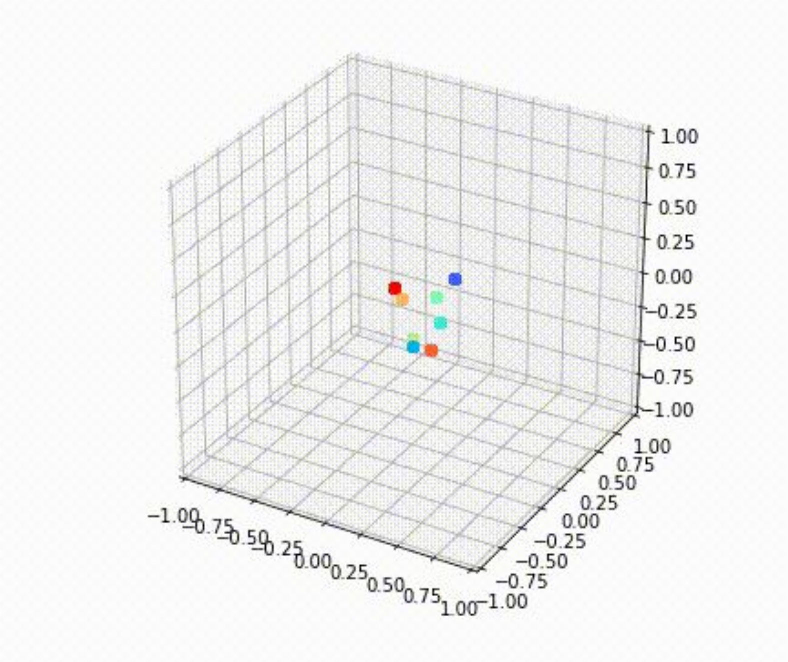

8-Goal Reacher

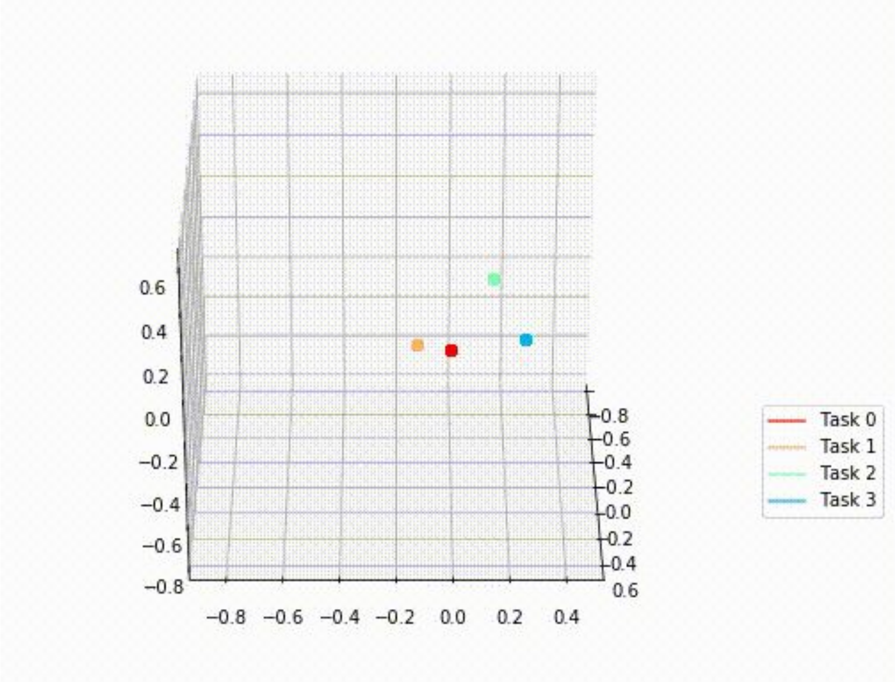

4-Goal Pusher

Method

Assumptions

- Useful library of “low-level” skills is known before learning
  - Diversity is important
  - Simpler skills → easier sim2real transfer
- All skills can be trained at once

Method

Decompose and Simulate to Scale

1. Jointly **learn** diverse low-level “skill” policies in simulation, parameterized by a (learned) latent space
2. Directly **transfer** policies to the robot
3. Quickly **search** in latent space for effective policies (or sequences thereof) for real-world tasks

Building Our Method

Single-Task Reinforcement Learning

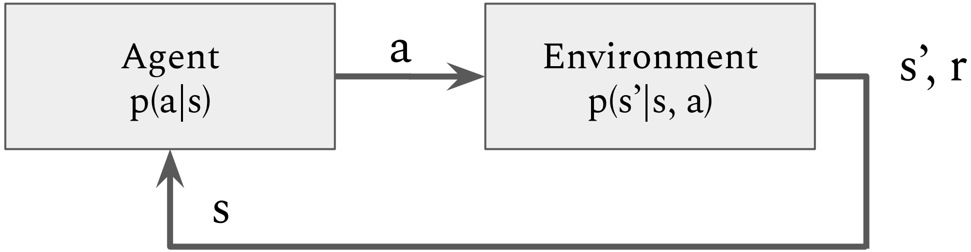

Sutton and Barto. "Reinforcement Learning: An Introduction." MIT Press, 1998.

Building Our Method

“Vanilla” Multi-Task RL

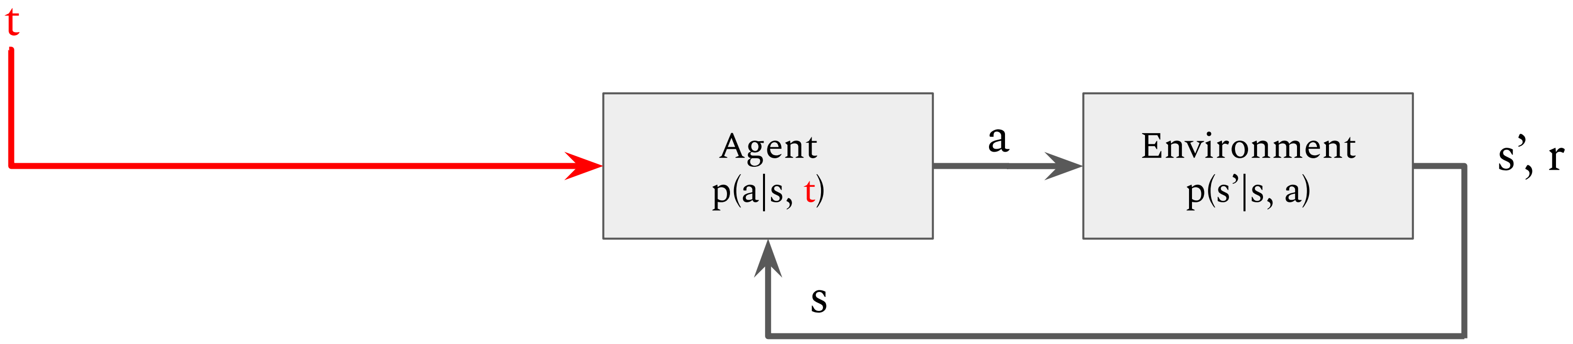

Building Our Method

Multi-Task RL with Embeddings (our method)

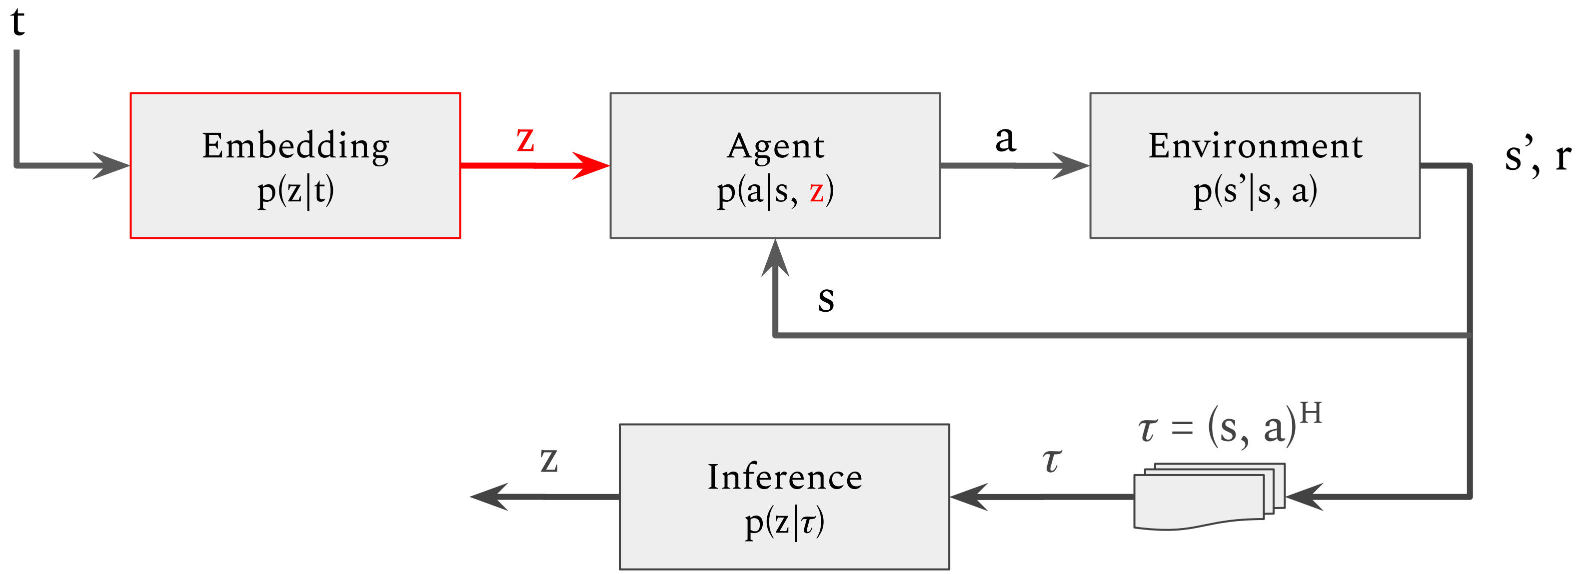

Hausman, et al. “An Embedding Space for Transferrable Robot Skills”. ICLR 2018

Experiments

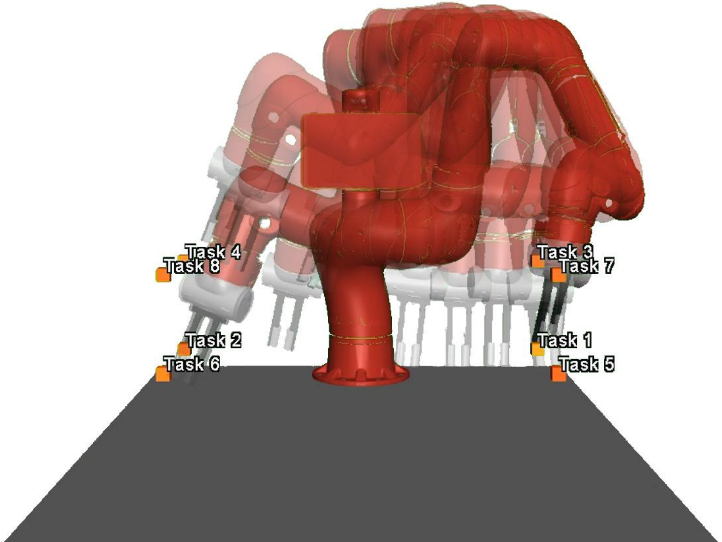

Sawyer Reacher

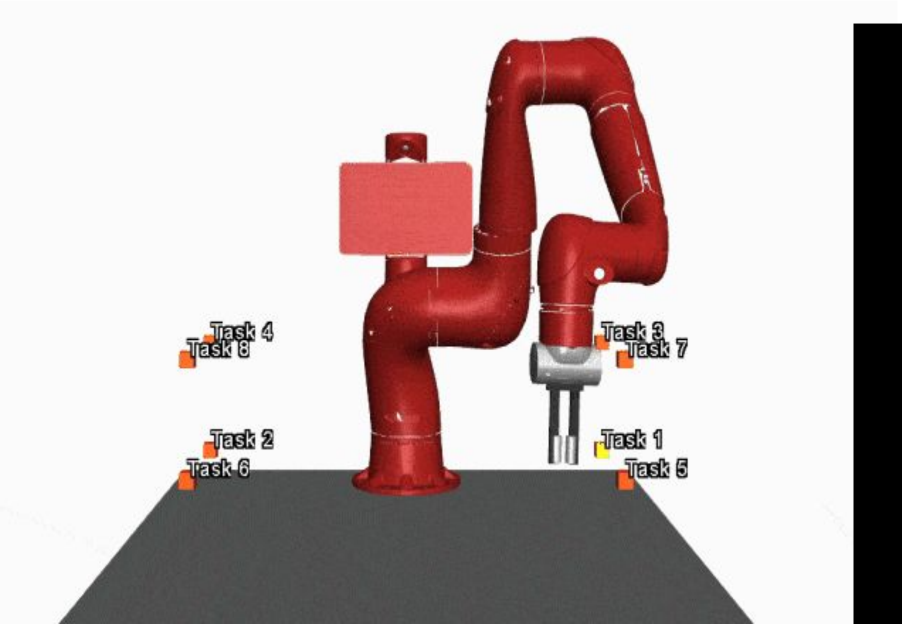

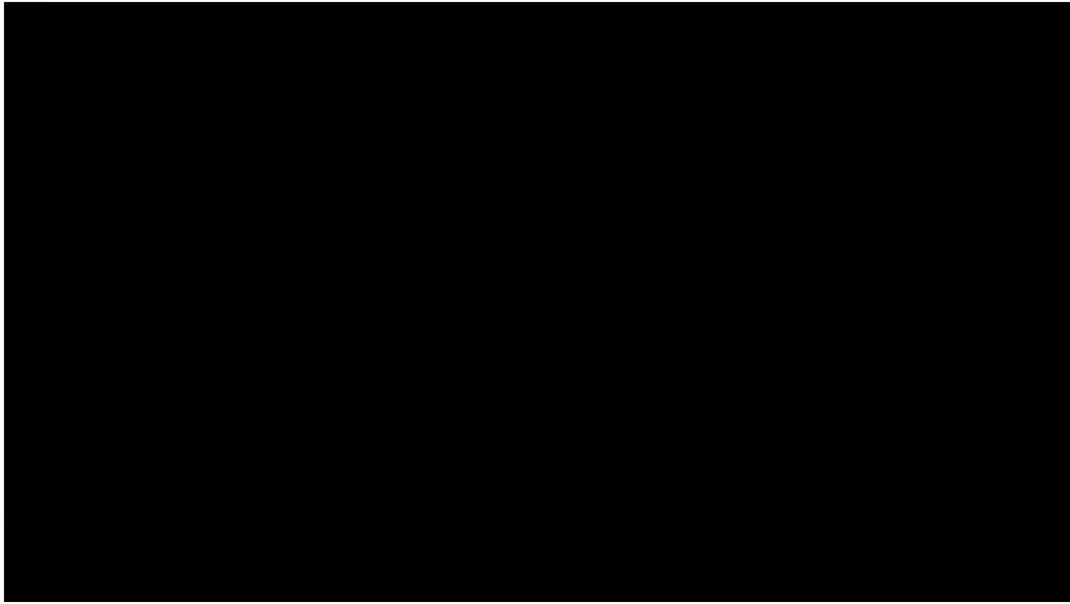

Sawyer Reacher - Sim vs. Real

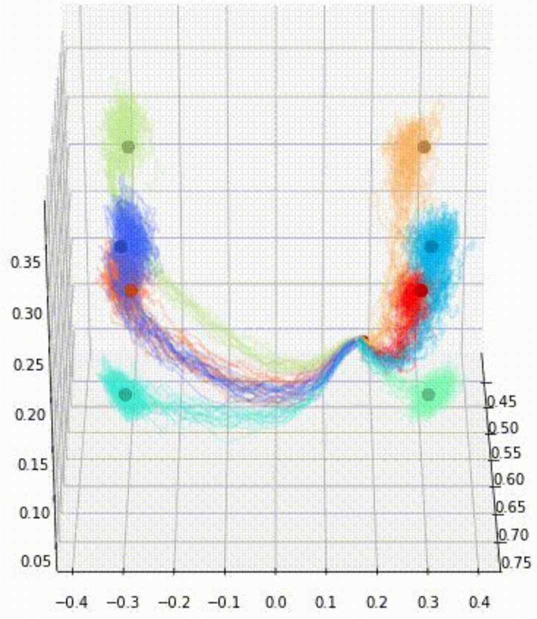

Simulation

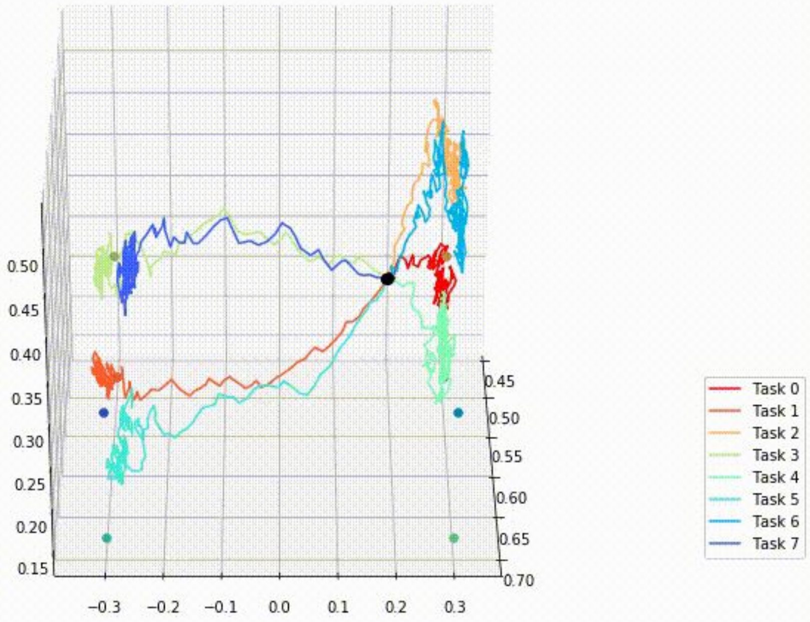

Reality

Sawyer Reacher - Composition (UCS)

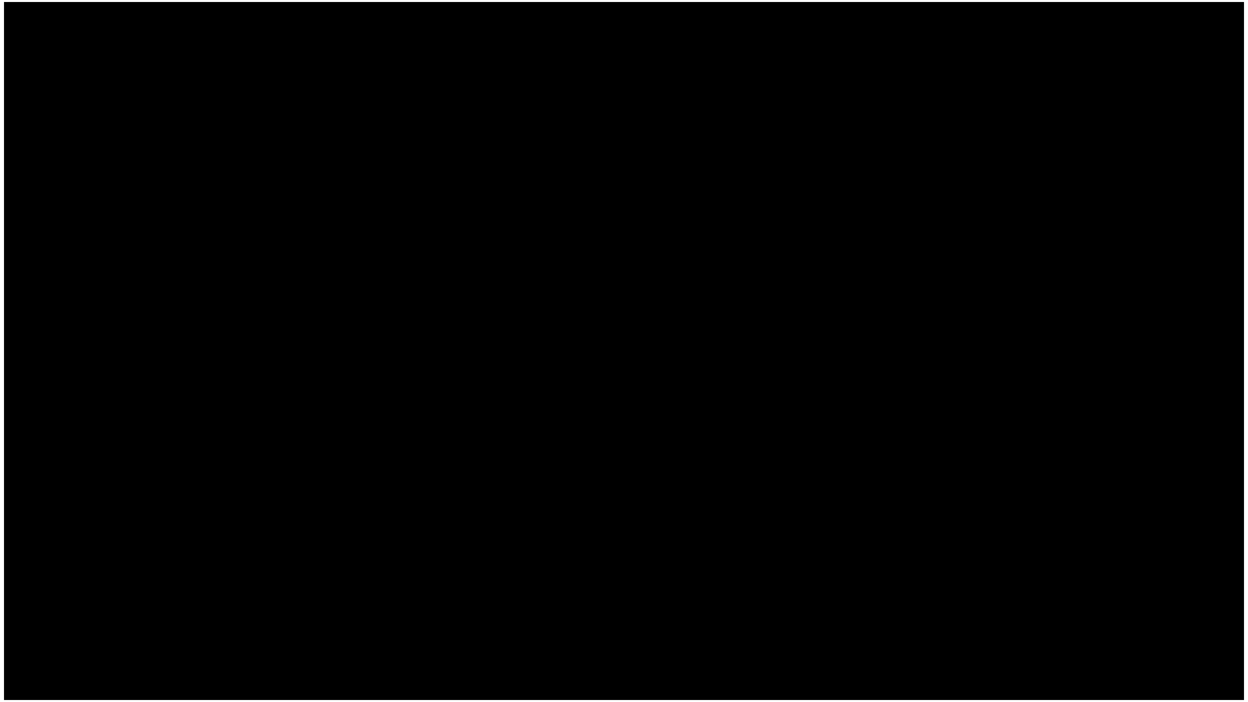

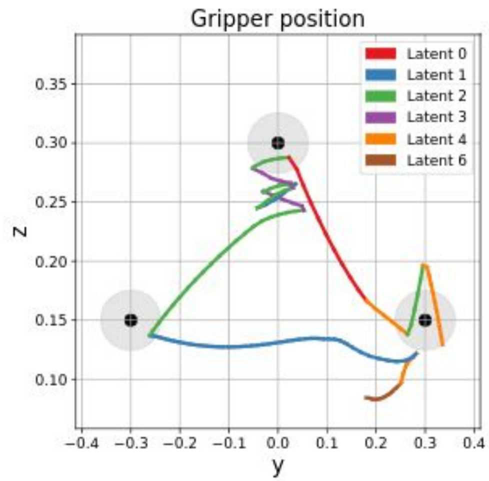

Search-based sequencing of task latents

Sawyer Pusher

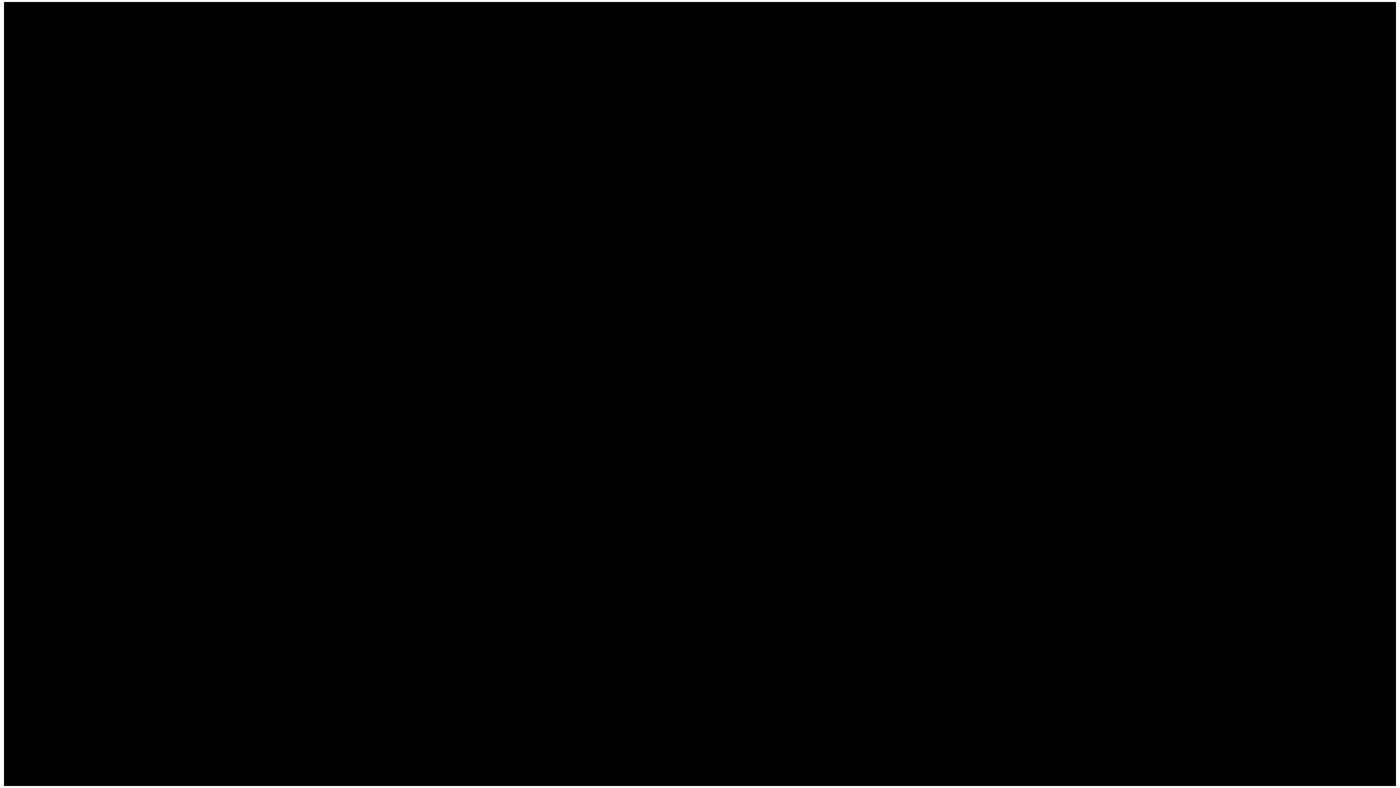

Sawyer Pusher - Single tasks

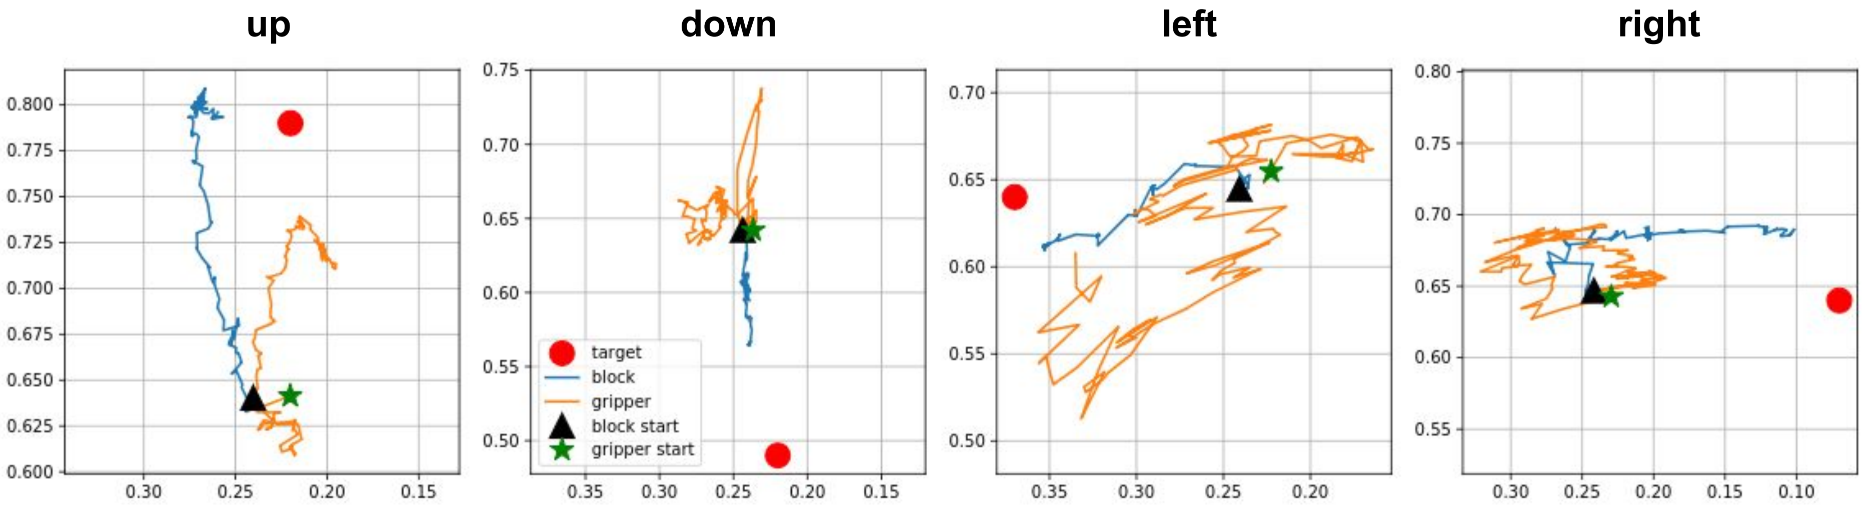

Sawyer Pusher - Composition (interpolation)

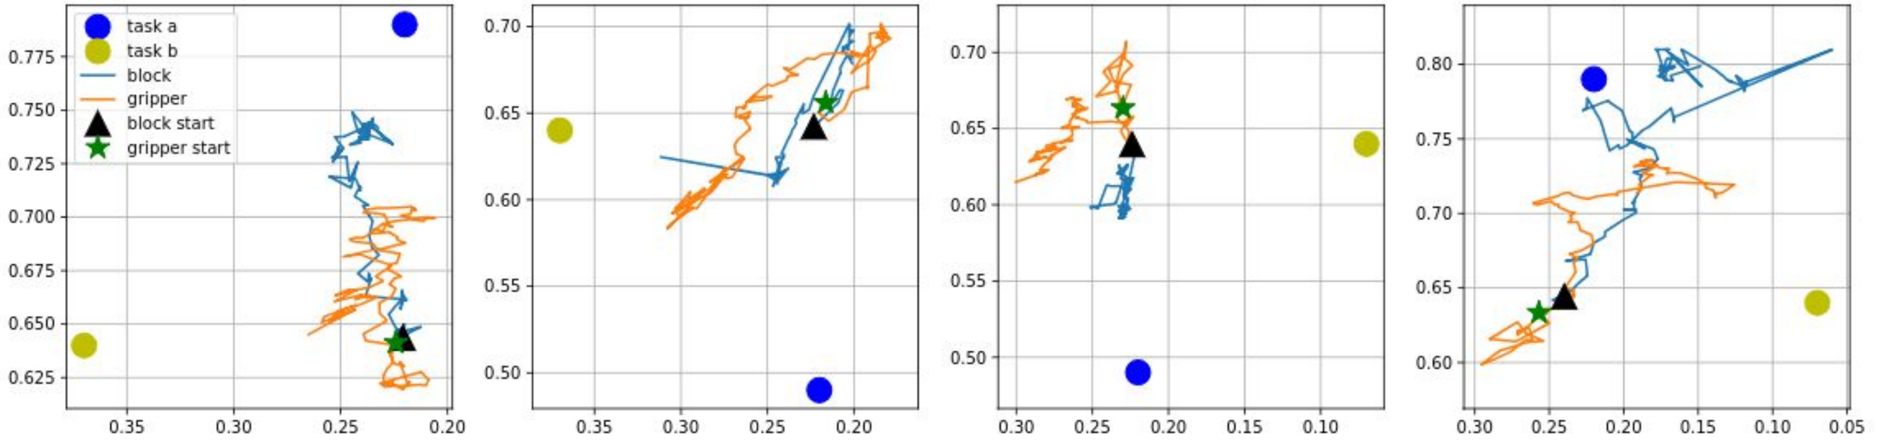

Sawyer Pusher - Composition (UCS)

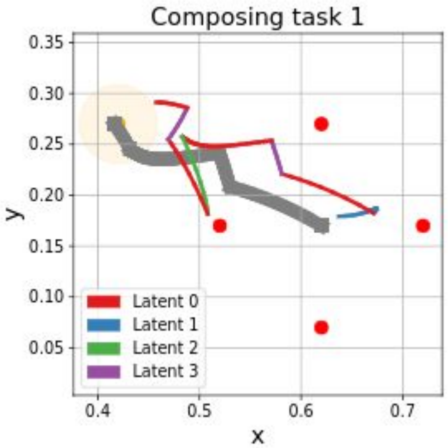

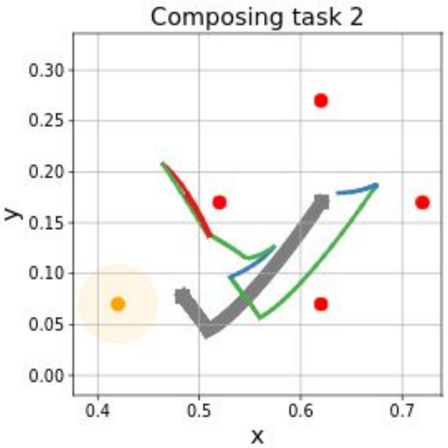

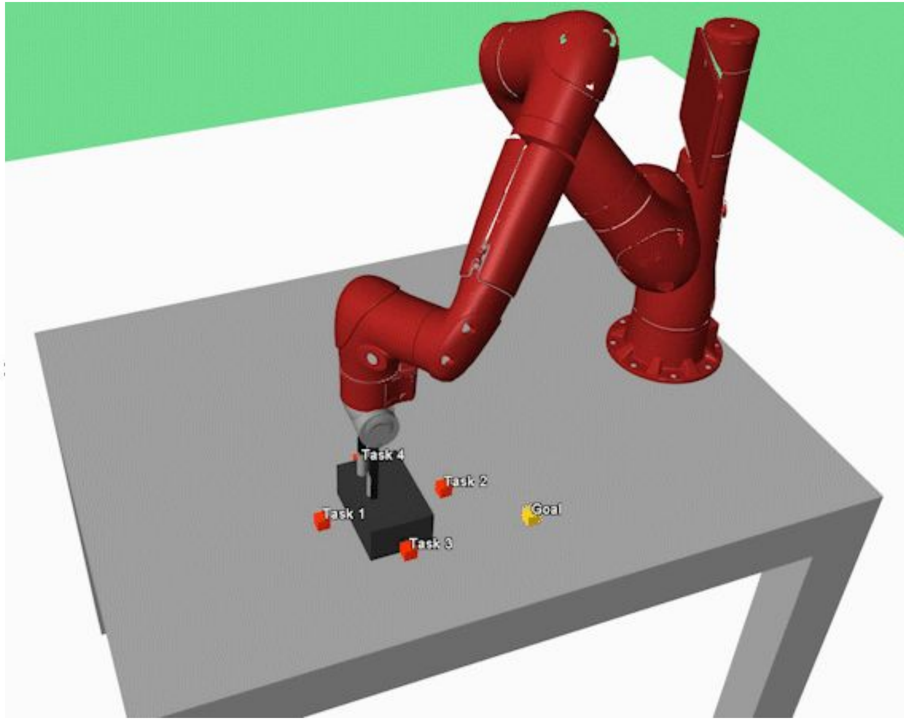

Conclusions

Learning Embedding Space of Composable Robot Skills

Alternative approach to achieving robust sim2real transfer

- *Faster transfer and fine-tuning*
- *Share training time among many tasks*

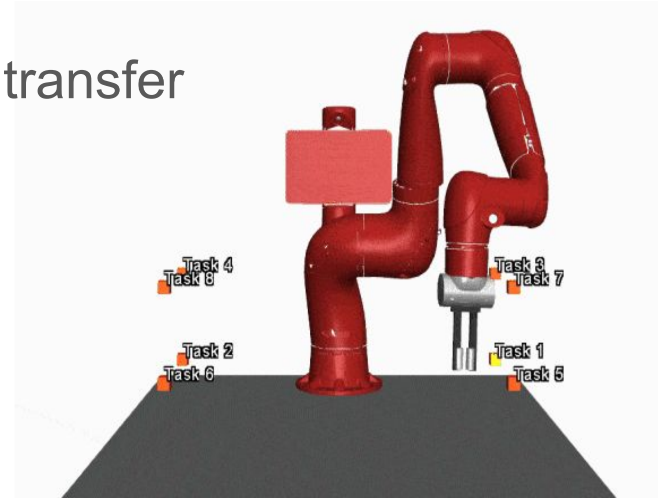

Combine proven robotics methods
(e.g. search) with data-driven learning

- *See our ICRA submission "Zero-Shot Skill Composition and Simulation-to-Real Transfer by Learning Task Representations"* arxiv.org/abs/1810.02422

Use our code!

RL research with real robots requires sophisticated infrastructure and experience

Please use our code!

- github.com/rlworkgroup/garage ← framework
- github.com/ryanjulian/embed2learn ← this paper

Happy to talk offline about tips and tricks for getting started

USC University of Southern California

Our approach

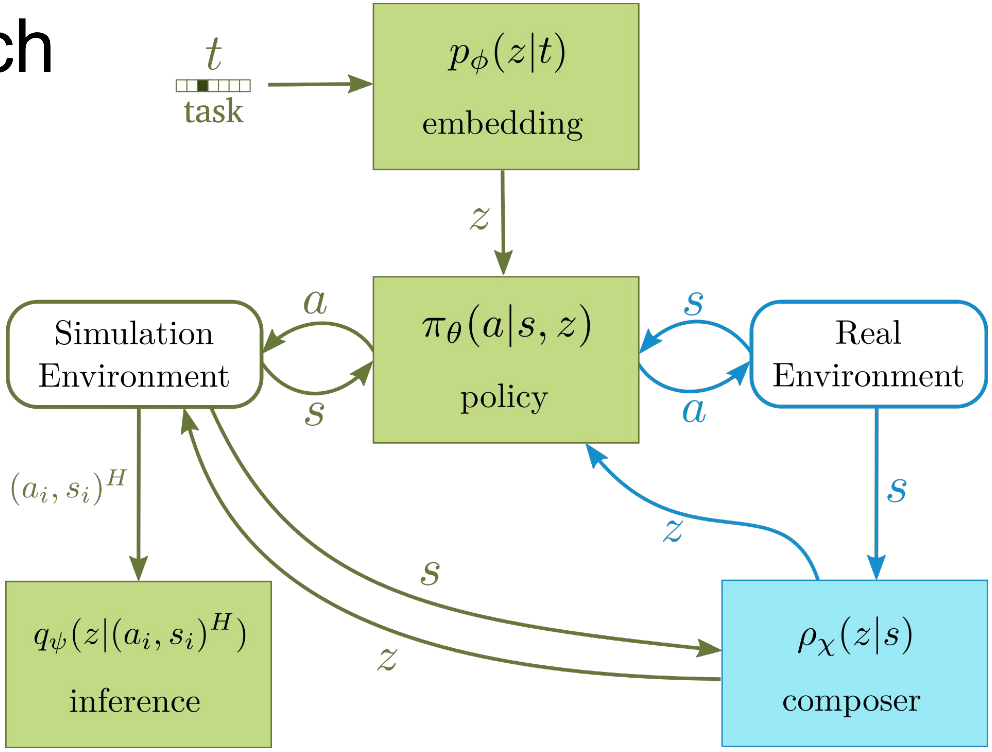

Learning an Embedding Space for Transferable Robot Skills.
International Conference on Learning Representations (ICLR), 2018.
K. Hausman, J.T. Springenberg, Z. Wang, N. Heess, M. Riedmiller

Evidence for Learned Representation