Deep Robotic Learning
using Visual Imagination and Meta-Learning

BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

Project Lead: *Chelsea Finn* **Demo Engineering & Design:** *Annie Xie*, Sudeep Dasari*, Frederik Ebert, Tianhe Yu*

Standard robotics paradigm:

RGB-D image
↓
segment objects
↓
estimate pose & physics of segments
↓
optimize action using
estimated poses & physics

Brittle, hand-engineered pipeline.

Our approach:

How can we enable robots to learn **vision-based manipulation skills** that **generalize** to new objects & goals?

- Learn from **raw pixel observations** (rather than task-specific, engineered representations)
- collect data with a **diverse** range of objects and environments
- **reuse data** from other objects & tasks when learning to perform new task

One-Shot Visual Imitation Learning

Chelsea Finn, Tianhe Yu*, Tianhao Zhang, Pieter Abbeel, Sergey Levine*

Can robots **reuse data** from other tasks to adapt to new objects from **only one visual demonstration**?

Our meta-learning approach: Learn to learn many other tasks using one demo

Meta-training:

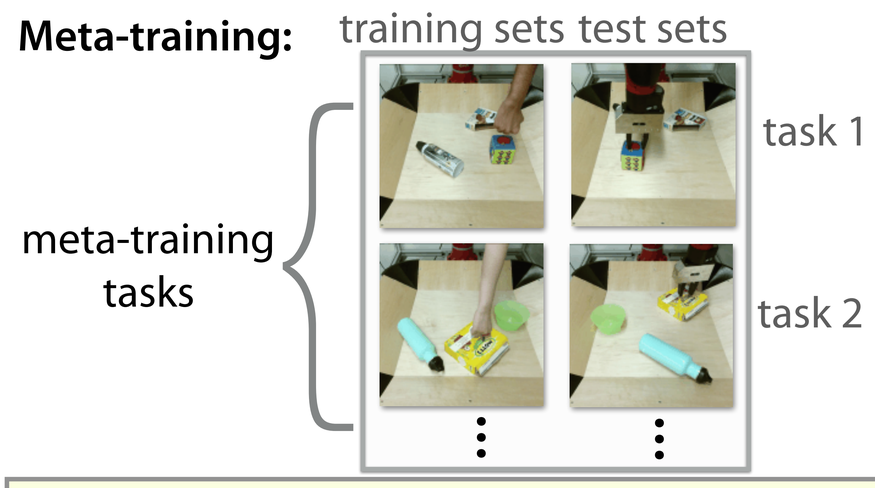

Meta-testing:

Learn new held-out task from 1 demo

Demo: Robot placing, tasks correspond to **different objects**.

Meta-training tasks:

Meta-Imitation Learning using MAML [1,2]

meta-training time: $\min_\theta \sum_{\text{tasks}} \mathcal{L}_{\text{v}}(\theta - \alpha \nabla_\theta \mathcal{L}_{\text{tr}}(\theta))$

meta-training tasks

val demo

training demo

meta-test time: $\theta' \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}(\theta)$

demo of meta-test task with **held-out objects**

shown in demo

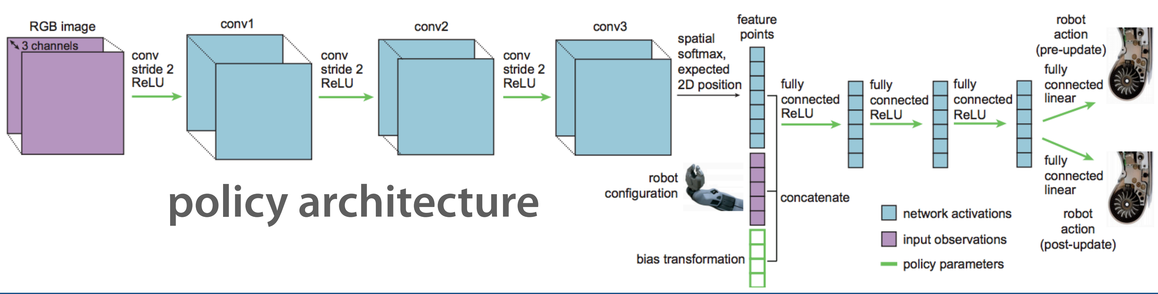

policy architecture

Planning with Visual Foresight

Frederik Ebert, Chelsea Finn, Alex Lee, Sergey Levine

How can robots acquire **general models** and skills using entirely **autonomously-collected data**?

Collect data autonomously

- program initial motions, provide objects
- record camera images and robot actions
- no object supervision, camera calibration, human annotation, etc.

Predict future video for different actions [3,5]

actions:

input image

future predictions

Planning with Visual Foresight [4,5]

User Input/Task Specification

Designated Pixel

Goal Point

Planning Module

Cost Function

Video Prediction Model

Apply action to robot

shown in demo

Sampling-based Planning Optimization

1. Sample many actions sequences
2. Predict the future for each action sequence
3. Rank futures using distance to goal pixel(s)
4. Iteratively refine sampled action sequences

One-Shot Imitation Learning Research

[1] Finn, Abbeel, Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. ICML '17
[2] Finn*, Yu*, Zhang, Abbeel, Levine. One-Shot Visual Imitation Learning via Meta-Learning. CoRL '17

Self-Supervised Visual Foresight Research

[3] Finn, Goodfellow, Levine. Unsupervised Learning for Physical Interaction through Video Prediction. NIPS '16
[4] Finn, Levine. Deep Visual Foresight for Planning Robot Motion. ICRA '17
[5] Ebert, Finn, Lee, Levine. Self-Supervised Visual Planning with Temporal Skip Connections. CoRL '17.