Generalization in Planning 2023

Towards General-Purpose In-Context Learning Agents

Louis Kirsch, James Harrison, C. Daniel Freeman,
Jascha Sohl-Dickstein, Jürgen Schmidhuber

@LouisKirschAI

Main focus of Meta-Learning

- Excels in adaptation to unseen but similar tasks

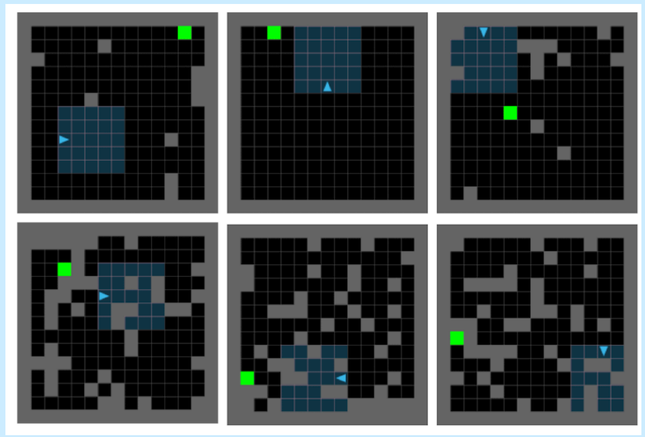

Env Distribution A

Meta Learner A
→ Learning Algorithm A

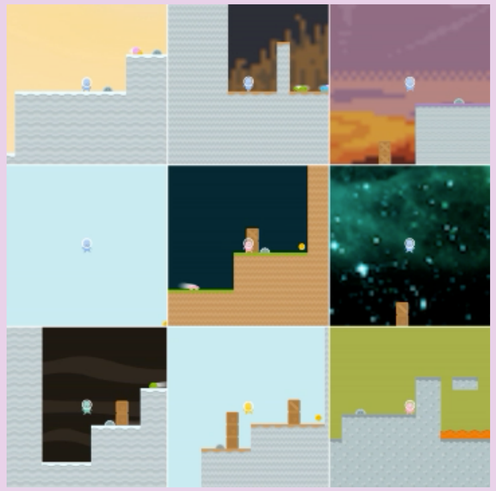

Env Distribution B

Meta Learner B
→ Learning Algorithm B

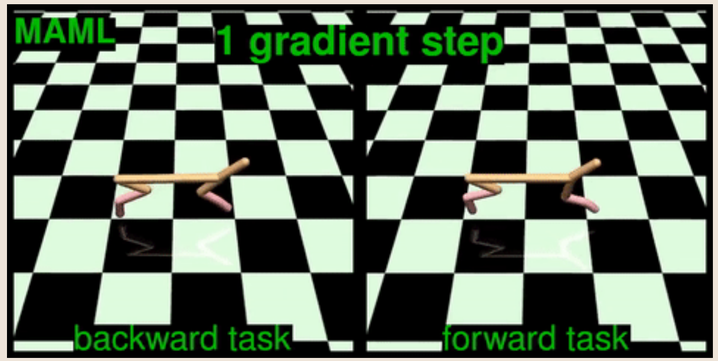

Env Distribution C

Meta Learner C
→ Learning Algorithm C

General-Purpose Meta-Learning

Can we train an agent that can efficiently in-context **learn and act in any environment**?

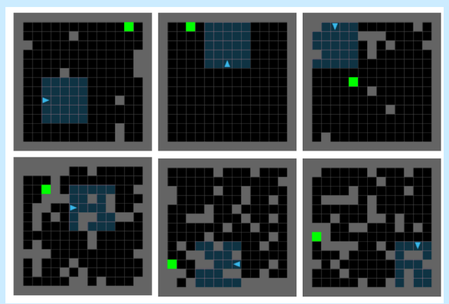

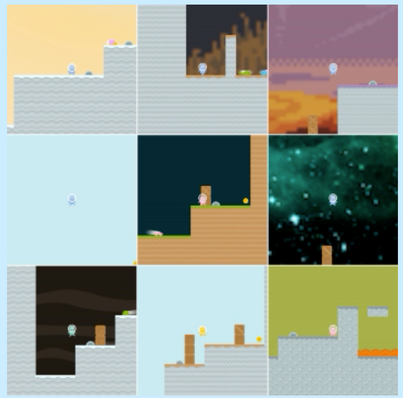

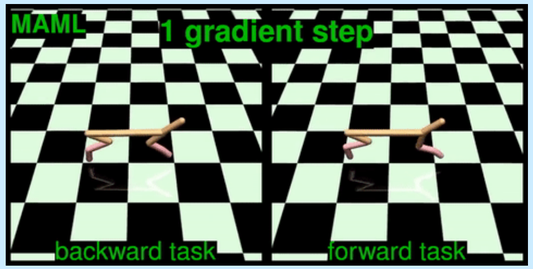

Env Distribution

Meta Learner
→ General-Purpose Learning Algorithm

Generalize

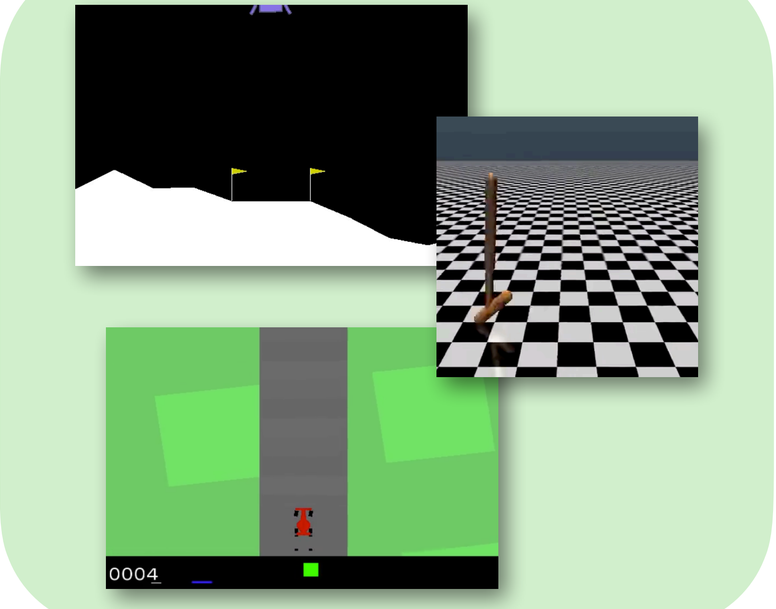

From gradient-based learning ...

MetaGenRL

Parameterized Objective / Loss function

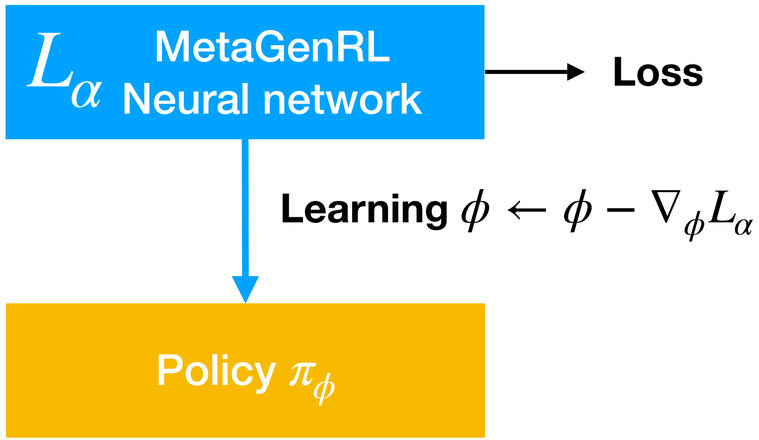

Kirsch et al [ICLR 2020]

Learned Policy Gradient

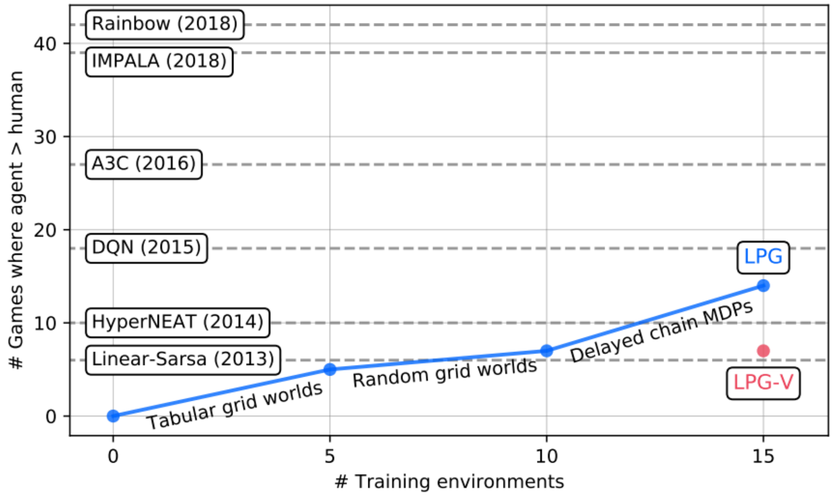

Oh et al [NeurIPS 2020]

Evolving RL Algorithms

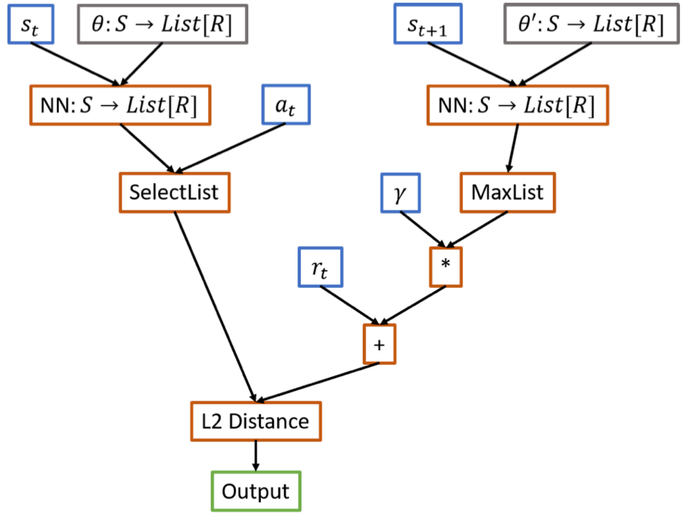

Co-Reyes et al [ICLR 2021]

… to in-context learning

MetaGenRL, LPG, etc

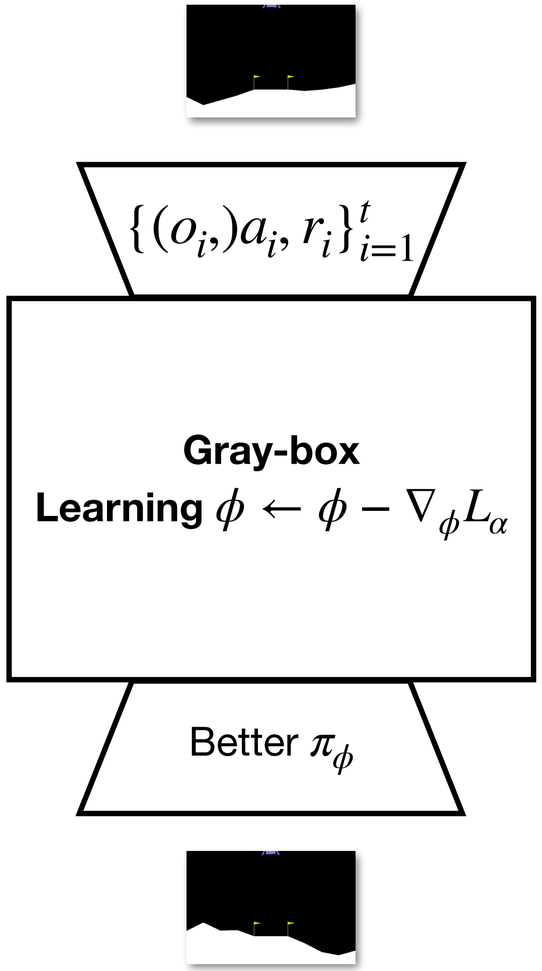

Good generalization

In-context meta-RL

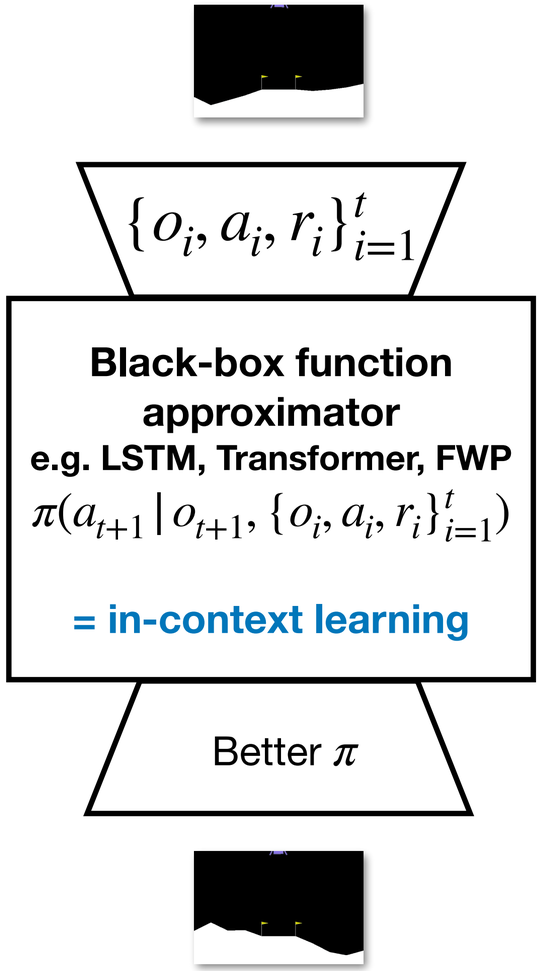

[Schmidhuber since 1992, RL² Duan et al 2016, Wang et al 2016, SymLA Kirsch et al 2022, Adaptive Agents 2023, etc]

Difficult generalization

This paper:
How to fix this?

Automated task / environment generation

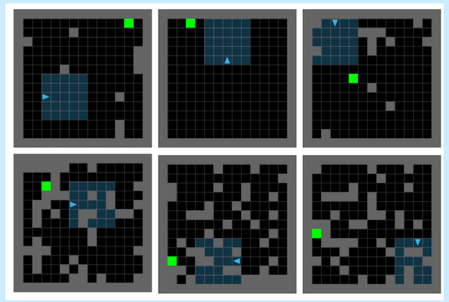

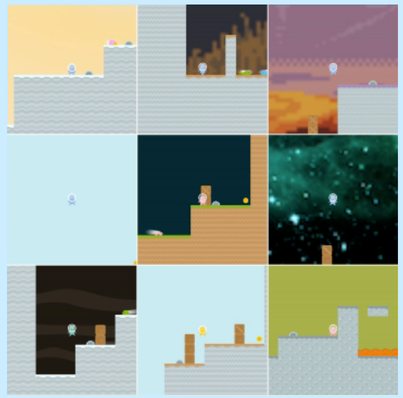

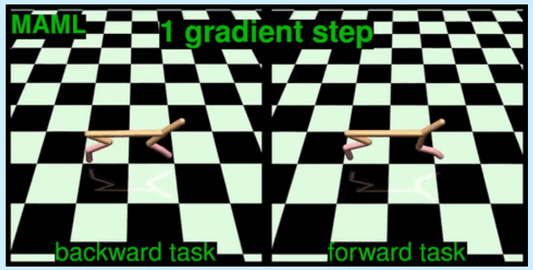

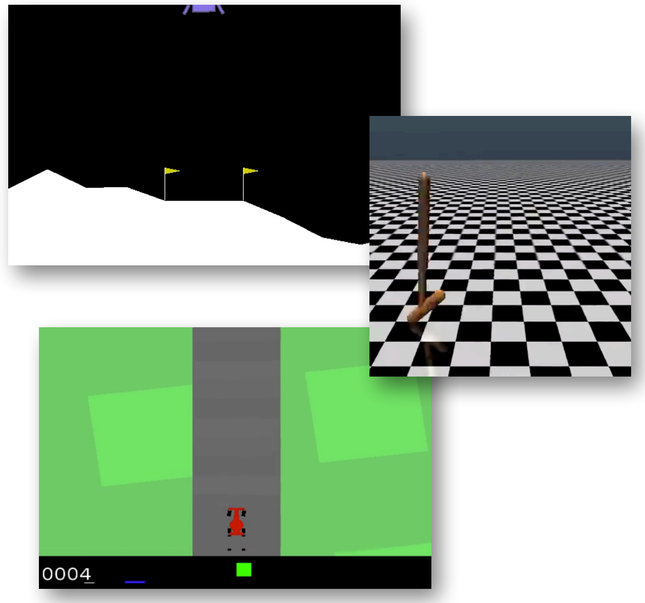

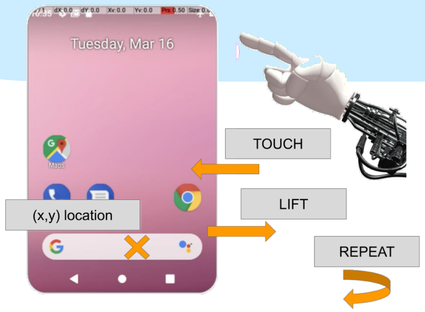

Env Distribution

GPICL: From memorization to general learn-to-learn

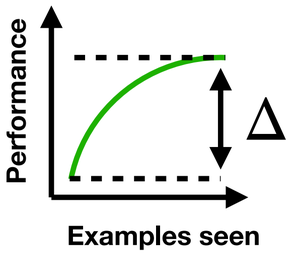

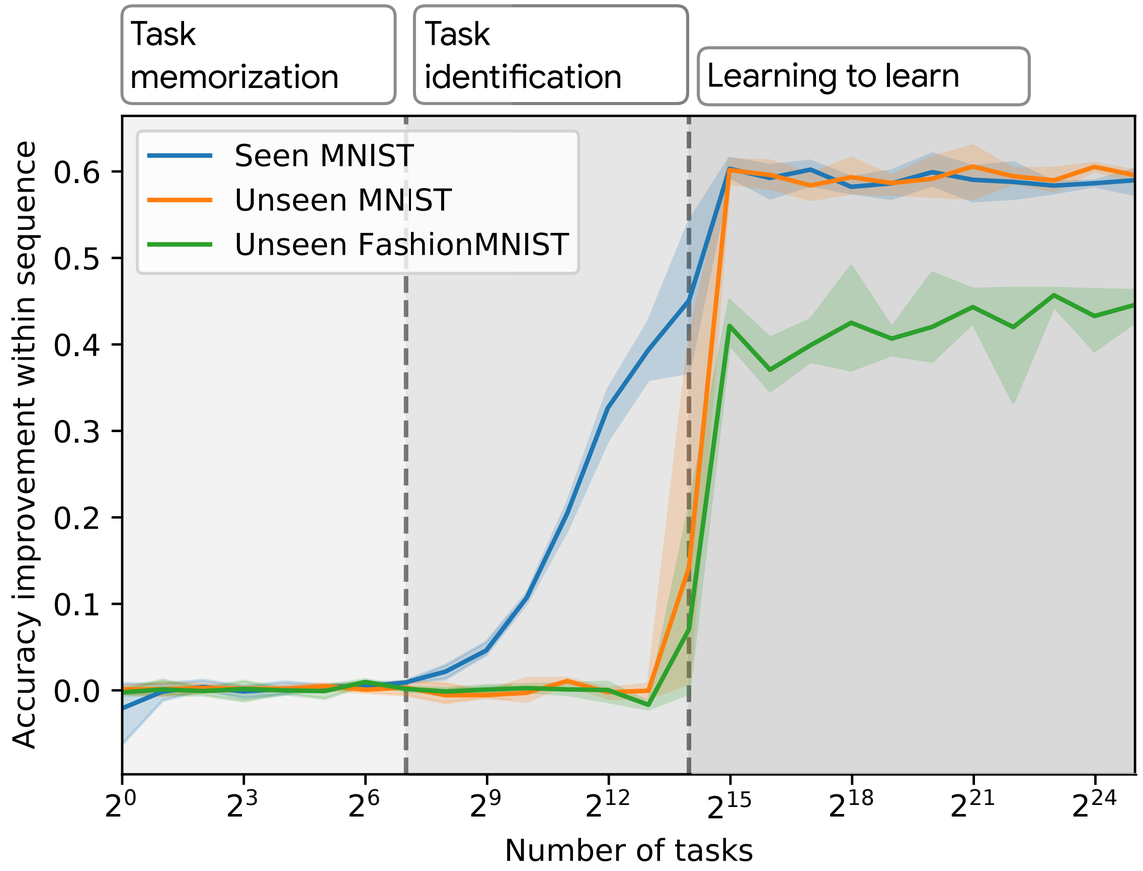

Learning	Generalization	Algorithm Description
× No	× No	Task memorization
✓ Yes	× No	Task identification
✓ Yes	✓ Yes	General-purpose learning algorithm

Transformers exhibit three different phases in terms of meta-learned behavior.

Kirsch et al [2022]

From classification to supervised meta-RL

IDSIA

Adaptation to RL

Instead of training on a supervised dataset $\left(\{x_i, y_i\}_{i=1}^{N_D}, x' \right) \mapsto y'$

Train on RL data $\left(\{s_i, a_i, r_i, d_i\}_{i=1}^{N_D}, s \right) \mapsto a$

Our recipe: **Offline meta-training + augmented RL data → Generalization**

Training is offline, but agent can learn online!

How it works - PPO Data collection

Data Collection

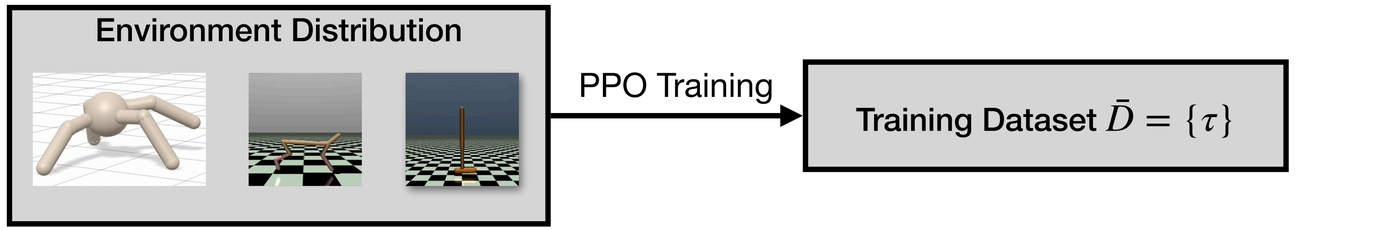

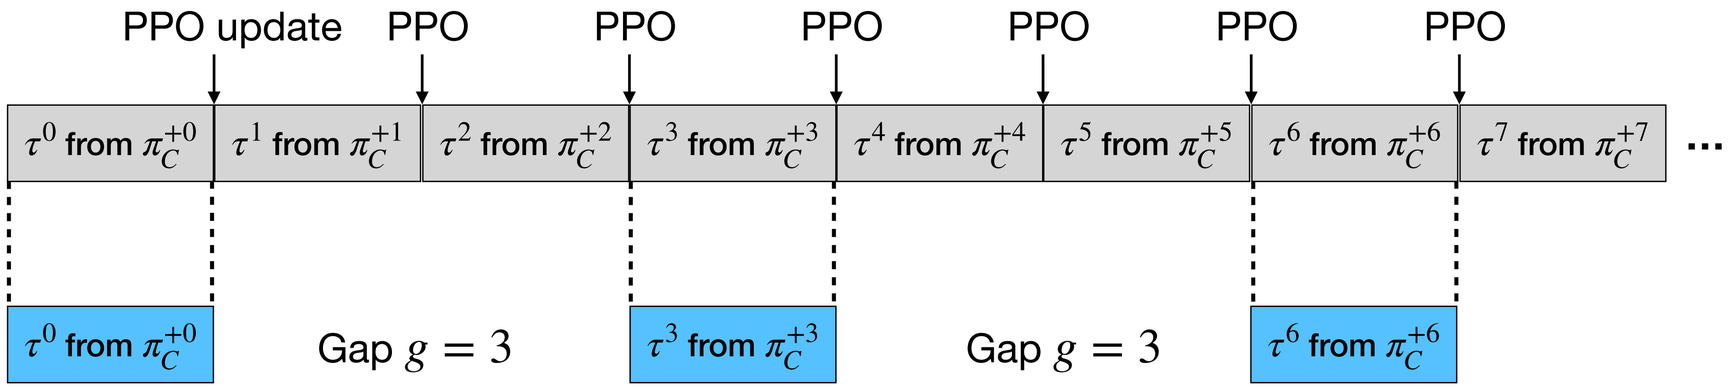

How it works - Model learning algorithm

Data Collection

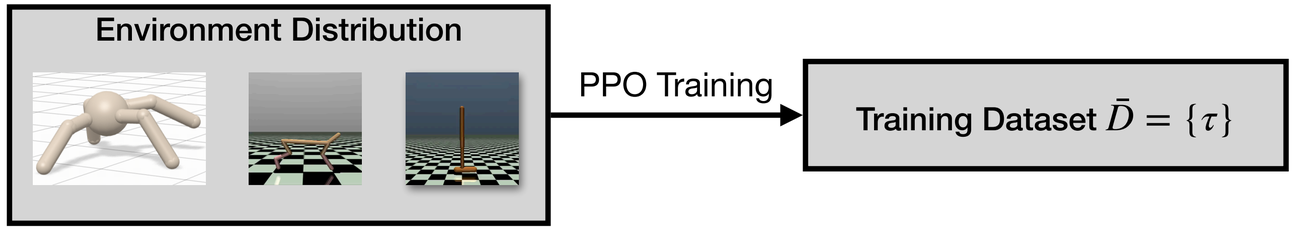

Take a subset of the collected data

Model an in-context learning algorithm

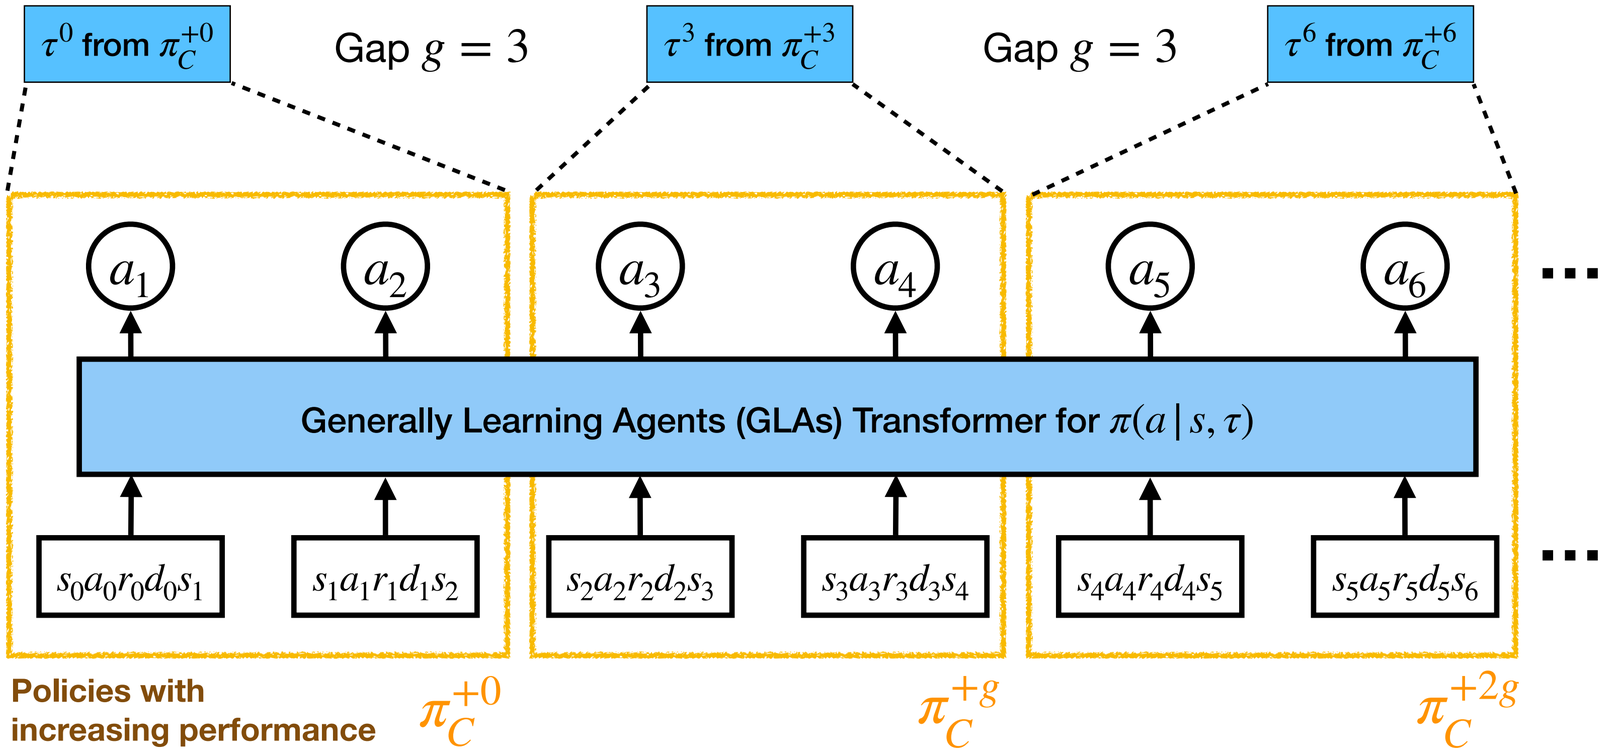

How it works - Augmentations

Data Collection

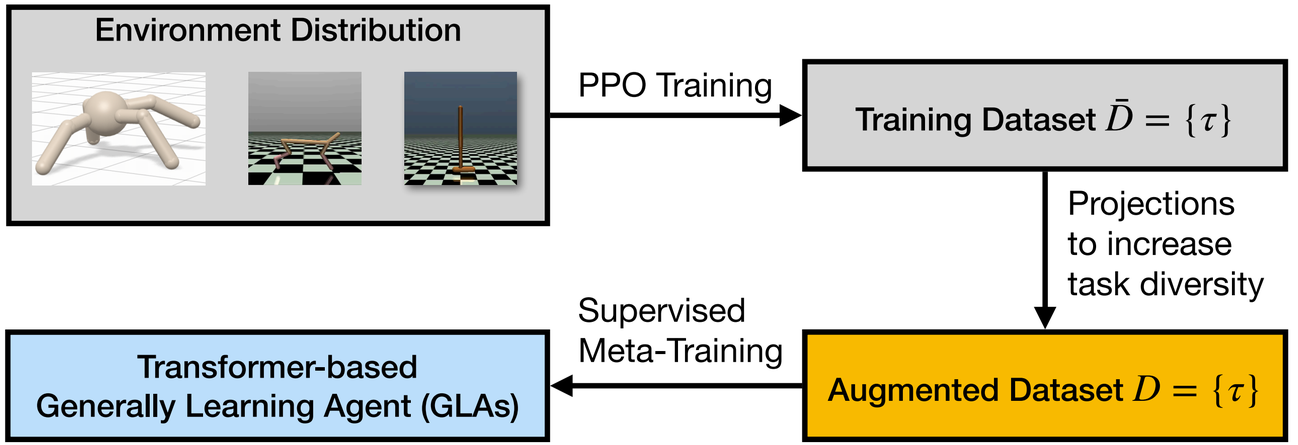

Meta-Training

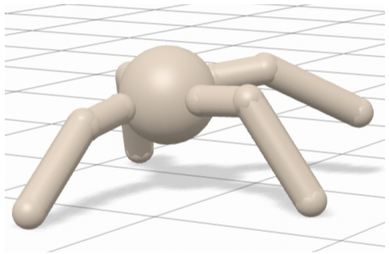

Sensor observations & actions

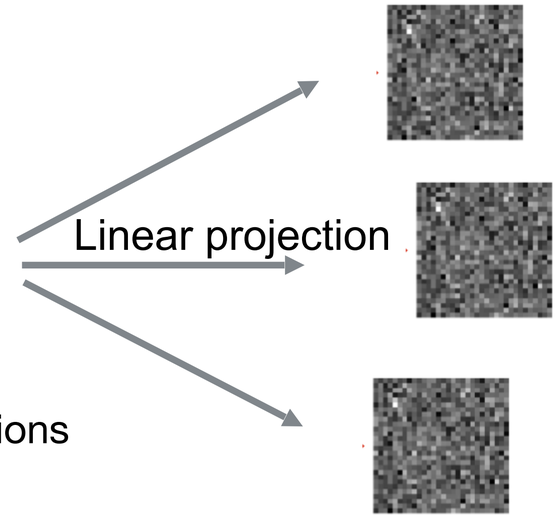

Transition dataset

$$\bar{D} = \{s_i, a_i, r_i, d_i\}$$

Create n tasks $\{\mathcal{S}_j, \mathcal{A}_j\}_{j=1}^{n}$

$$D_j = \{\mathcal{S}_j s_i, \mathcal{A}_j a_i, r_i, d_i\}$$

Linear observation projection with $\mathcal{S}_j$

Linear action projection with $\mathcal{A}_j$

How it works - Meta-Testing

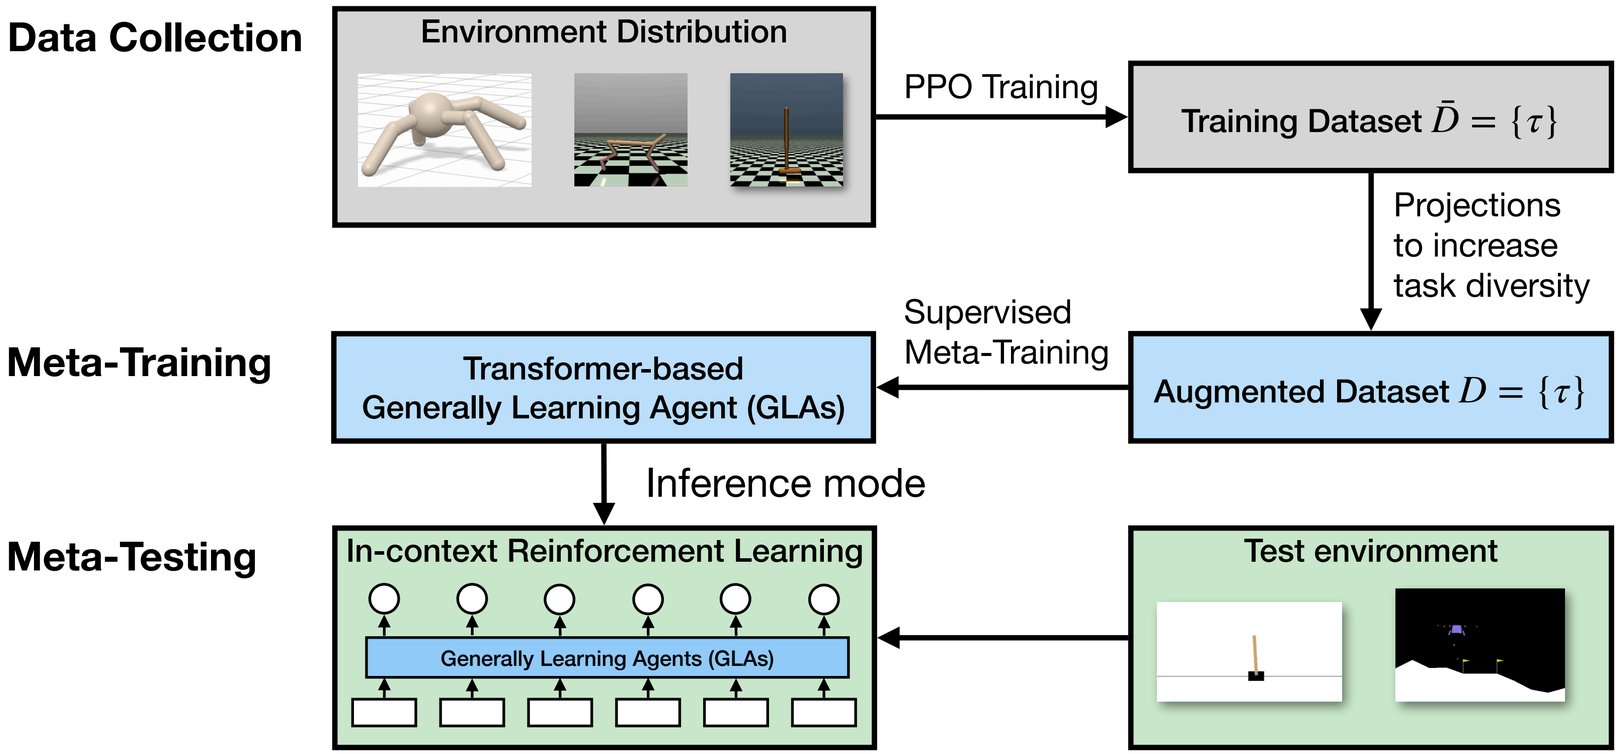

Evaluation: Single Environment

In-context learning policies that encode a learning algorithm on a specific task

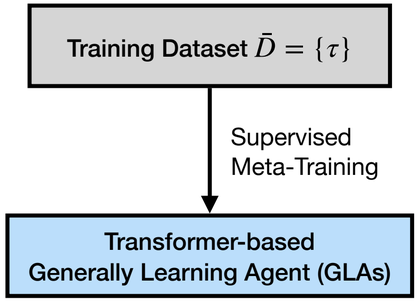

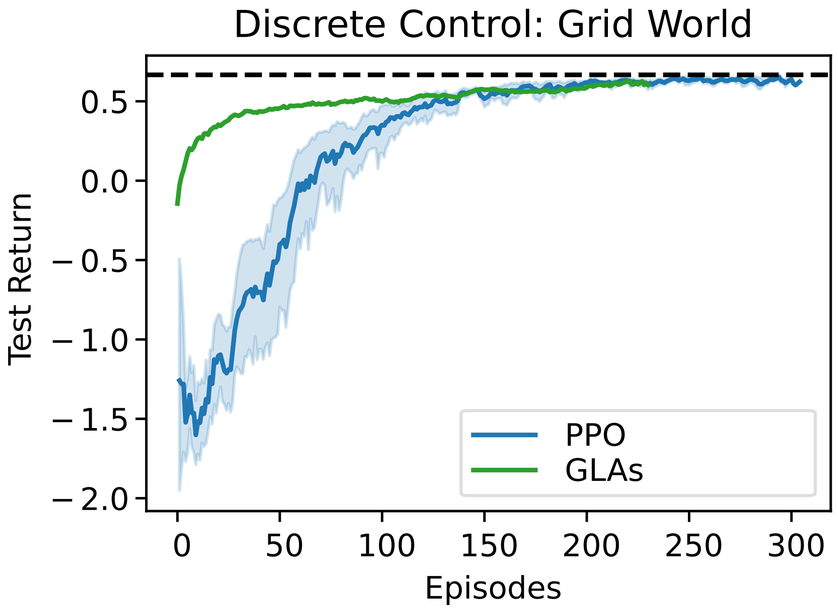

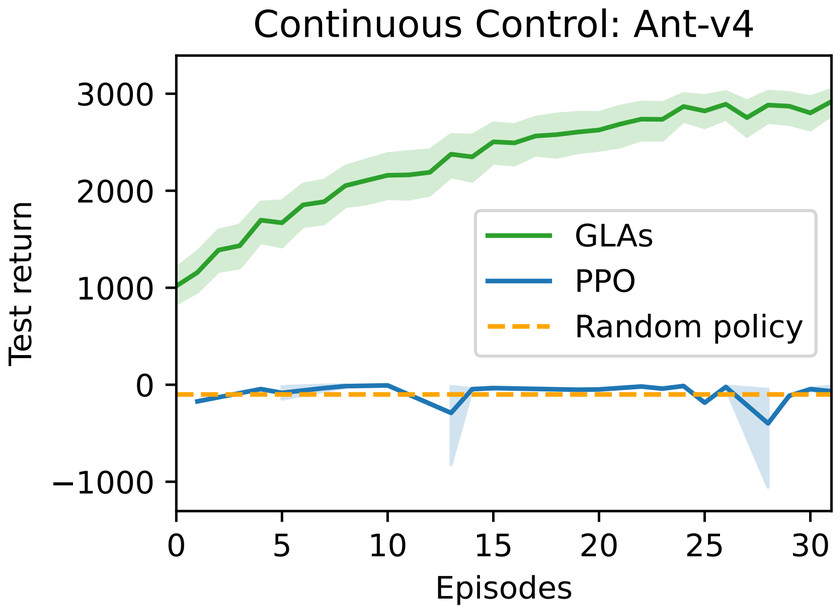

Effect of the gap

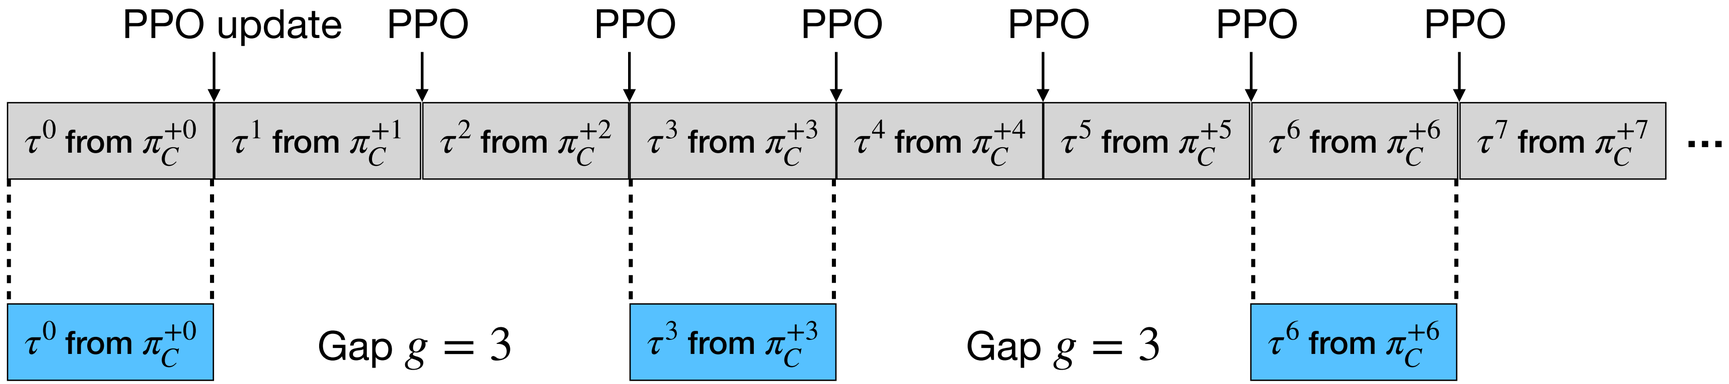

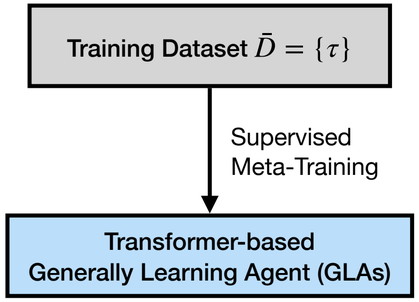

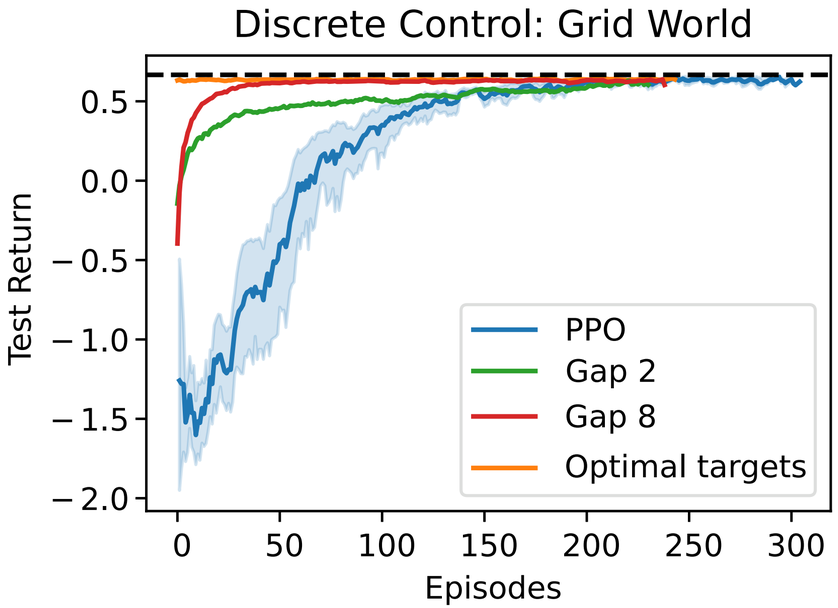

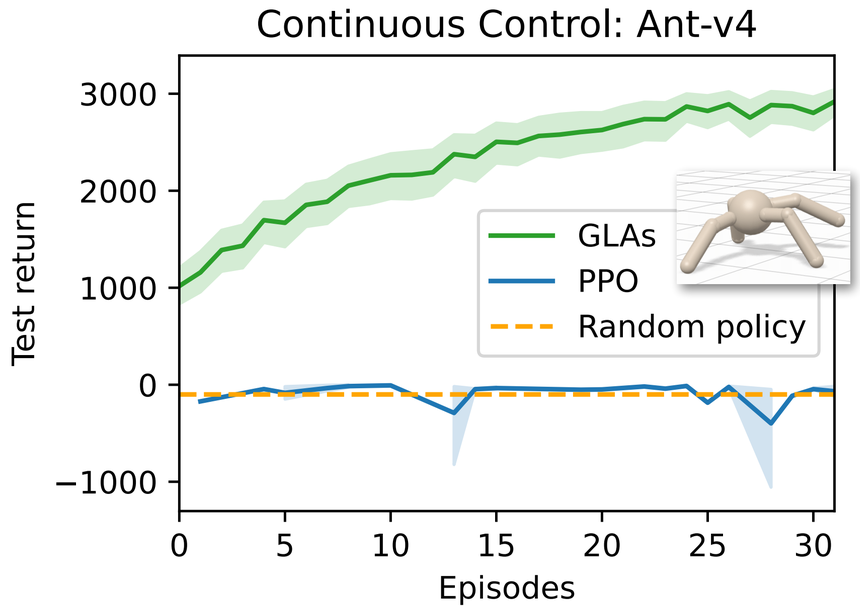

Generalize across environments?

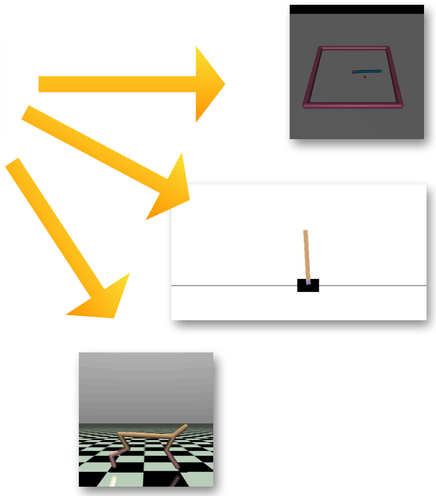

Introducing random projections

Generalize across environments?

Early results!
* **Longer contexts**
* **Training dynamics**

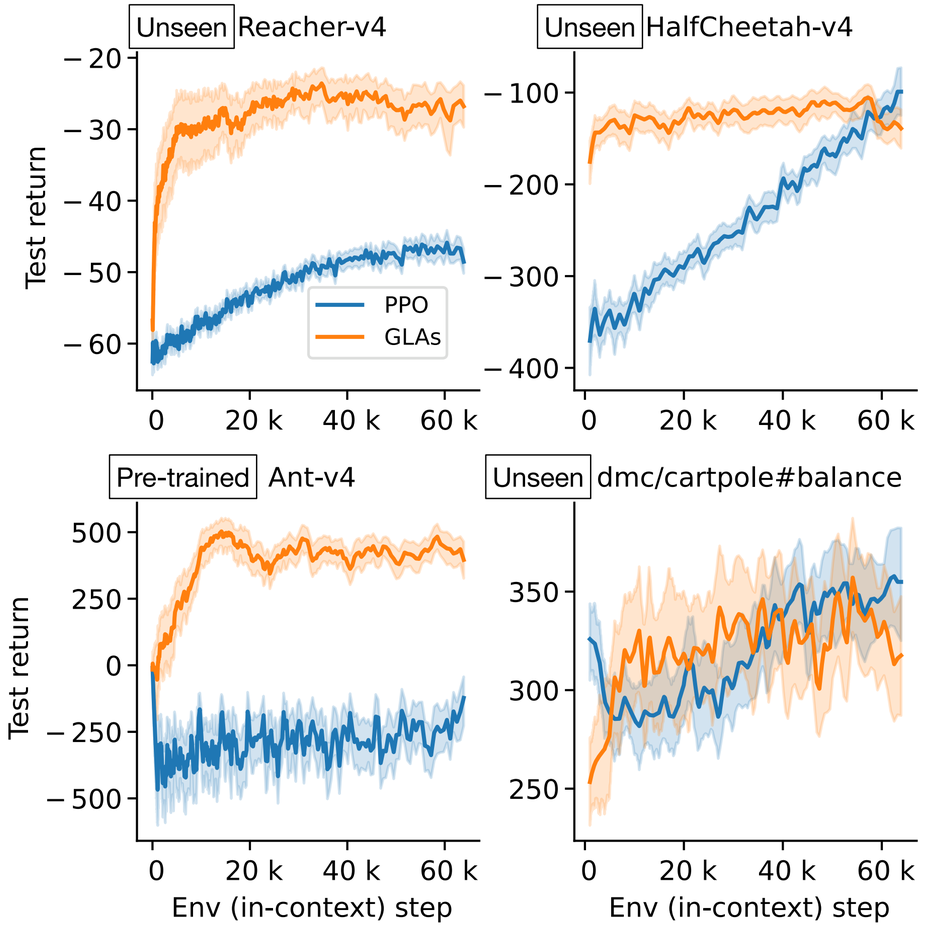

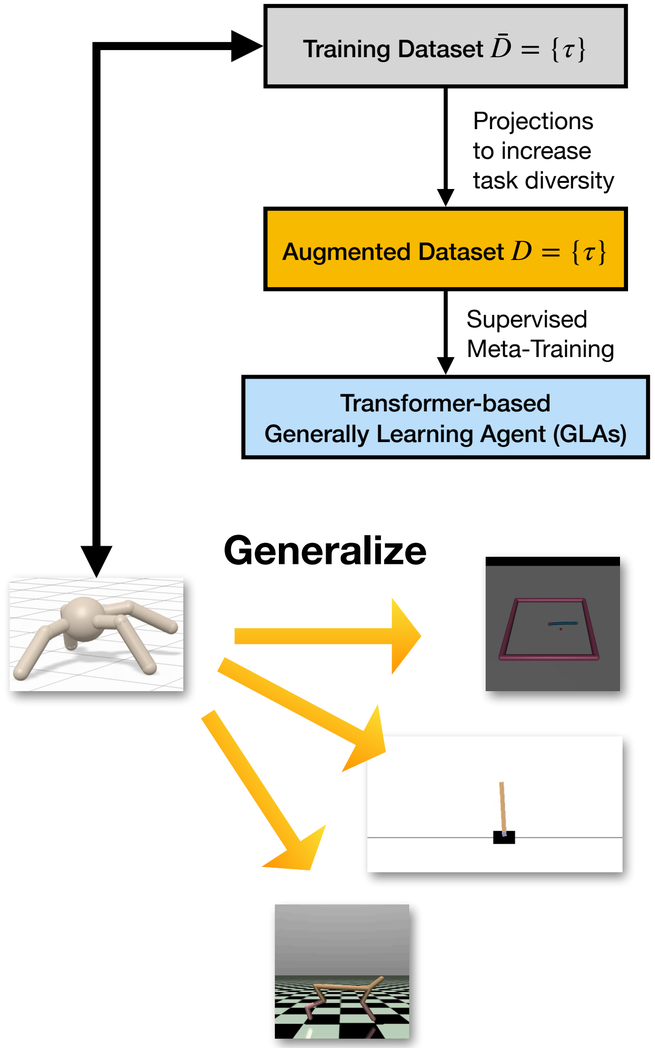

Automated task / environment generation

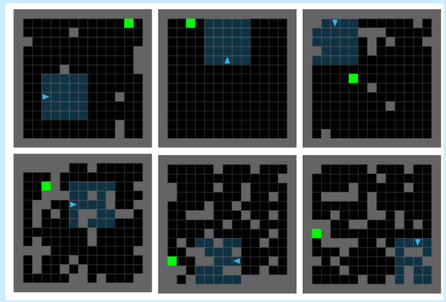

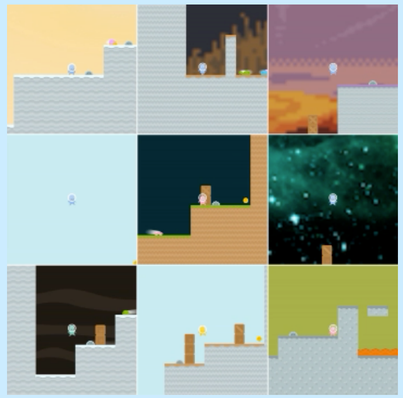

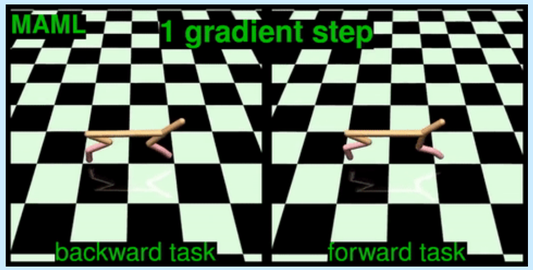

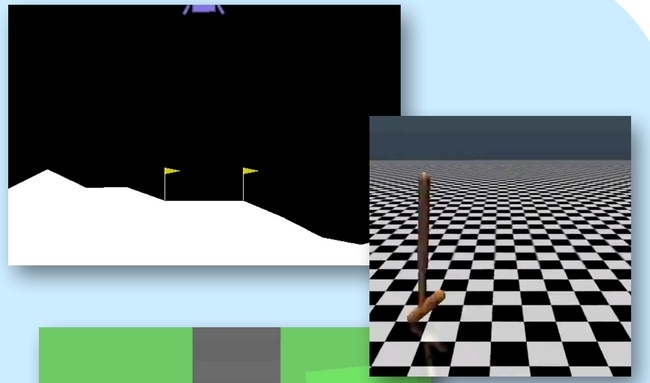

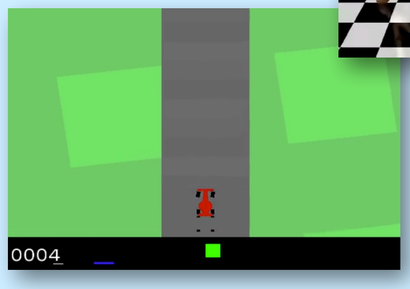

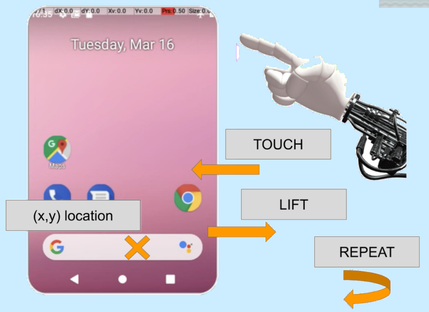

Env Distribution

→ General-Purpose Learning Algorithm

Where to find me

louiskirsch.com

@LouisKirschAI

Summary:

- **We distill an accelerated PPO into a Transformer**
- **Strong data augmentation helps the Transformer implement a general-purpose online RL algorithm generalizing across domains**