

Learning Embeddings for Sequential Tasks Using Population of Agents

Mridul Mahajan , Georgios Tzannetos , Goran Radanovic and Adish Singla

Max Planck Institute for Software Systems

{mrmahaja, gtzannet, gradanovic, adishs}@mpi-sws.org

Abstract

We present an information-theoretic framework to learn fixed-dimensional embeddings for tasks in reinforcement learning. We leverage the idea that two tasks are similar if observing an agent’s performance on one task reduces our uncertainty about its performance on the other. This intuition is captured by our information-theoretic criterion which uses a diverse agent population as an approximation for the space of agents to measure similarity between tasks in sequential decision-making settings. In addition to qualitative assessment, we empirically demonstrate the effectiveness of our techniques based on task embeddings by quantitative comparisons against strong baselines on two application scenarios: predicting an agent’s performance on a new task by observing its performance on a small quiz of tasks, and selecting tasks with desired characteristics from a given set of options.

1 Introduction

Embeddings are widely used to represent data points as vectors in a space that captures meaningful relations between them [Sun *et al.*, 2014; Sung *et al.*, 2018; Athar *et al.*, 2020; Mikolov *et al.*, 2013; Pennington *et al.*, 2014; Cer *et al.*, 2018; Zhang *et al.*, 2021]. They could also be utilized as representations for tasks, as studied in various areas such as multi-task learning [Zhang *et al.*, 2018], meta-learning [Achille *et al.*, 2019], and domain-adaptation [Peng *et al.*, 2020]. In reinforcement learning (RL), task embeddings could be used to understand the shared structure in sequential decision-making problems if similar tasks are embedded in close proximity. This enables efficient, one-shot computation of task similarity, eliminating the need for time-consuming policy rollouts.

Despite the potential benefits, prior work on learning task embeddings in RL [Qin *et al.*, 2022; Schäfer *et al.*, 2022; Arnekvist *et al.*, 2018; Yoo *et al.*, 2022; Rakelly *et al.*, 2019; Bing *et al.*, 2023; Gupta *et al.*, 2018; Fu *et al.*, 2020; Li *et al.*, 2021; Lan *et al.*, 2019; Walke *et al.*, 2022; Sodhani *et al.*, 2021b; Vuorio *et al.*, 2019] does not explicitly optimize for task similarity. This could primarily be attributed to the lack of a general framework to measure (and reason about) similarities among sequential tasks.

To this end, we introduce an information-theoretic framework to learn fixed-dimensional embeddings for tasks in RL; the inner product in the embedding space captures similarity between tasks, and the norm of the embedding induces an ordering on the tasks based on their difficulties (see Fig. 1). A critical component of the framework is an agent population exhibiting diverse behaviors, which serves as an approximation for the space of agents. It leverages the idea that two sequential tasks are similar to each other if observing the performance of an agent from this population on one task significantly decreases our uncertainty about its performance on the other.

Essentially, there is an underlying notion of skills required to solve sequential tasks, and several tasks require some skills in common. As these skills are unknown, we measure task similarity by gathering task correlation statistics from an agent population that emerges during the multi-task policy learning process. More concretely, we introduce an information-theoretic criterion to measure task similarity (Section 4.1). Through this, we construct a set of ordinal constraints on the embeddings (with each such constraint asserting the relative similarity between a triplet of tasks), and propose a training scheme for an embedding network to learn them (Section 4.2).

Besides assessing the learned embedding spaces through visualizations (Section 5) and comparisons with prior work (Section 6), we ground our framework in two downstream scenarios that are inspired by real-world applications (Section 7). Firstly, we show its utility in predicting an agent’s performance on a new task given its performance on a small quiz of tasks, which is similar to assessing a student’s proficiency in adaptive learning platforms via a compact quiz [He-Yueya and Singla, 2021]. Several methods for curriculum design in RL rely on the agent’s value function [Klink *et al.*, 2020b; Eimer *et al.*, 2021; Yengera *et al.*, 2021]. Thus, one concrete application of task embeddings is in black-box curriculum design, where only the agent’s performance is observable, not the value function. We demonstrate that task embeddings can be used to infer the agent’s value function by observing its performance on a small quiz. Secondly, we demonstrate the application of our framework in selecting tasks with desired characteristics from a given set of options, such as choosing tasks that are slightly harder than a reference task. This is analogous to selecting desired questions from a pool for a personalized learning experience in online education systems [Ghosh *et al.*, 2022]. Through comparisons with strong base-




Figure 1: Schematics of our approach. We learn a task embedding function $f_\phi(\cdot)$ that maps a task s to its fixed-dimensional representation E . In this illustration, we show the properties of the learned embeddings using the MULTIKEYNAV environment in which tasks require the agent (shown as a black circle) to pick-up certain keys (from the gray segments) to unlock the door (the right-most segment) that has certain requirements (shown in color in the form of gates). A possible solution trajectory is depicted using dotted lines. Keys on this trajectory correspond to the ones that the agent possesses at that point in time. For instance, in task s_2 , the agent starts off with the yellow key in possession already. $\langle E_1, E_2 \rangle$ is greater than $\langle E_1, E_3 \rangle$, since tasks s_1 and s_2 have a common requirement of picking the blue key, and thus, are similar. Additionally, $\|E_2\|_2$ is less than both $\|E_1\|_2$ and $\|E_3\|_2$, since task s_2 requires picking a single key, while tasks s_1 and s_3 require picking two keys, which makes them harder than s_2 .

lines on diverse environments, we show the efficacy of our techniques based on task embeddings.

To summarize, our work makes the following contributions:

- I. We introduce an information-theoretic framework to learn task embeddings in RL. As part of the framework, we propose a task similarity criterion which uses a diverse population of agents to measure similarity among sequential tasks (Section 4.1).
- II. We propose a scheme to learn task embeddings by leveraging the ordinal constraints imposed by our similarity criterion (Section 4.2).
- III. To assess our framework, we perform visual assessments of the learned embedding spaces, and introduce two quantitative benchmarks: (a) agent’s performance prediction, and (b) task selection with desired characteristics (Sections 5, 6, and 7).¹

2 Related Work

Task embeddings in RL. Several works in meta-learning and multi-task learning have explored the use of embeddings to model relationships between sequential tasks, where embeddings are either learned explicitly through objectives such as reconstruction [Arnekvist *et al.*, 2018; Yoo *et al.*, 2022; Bing *et al.*, 2023] and trajectory-based contrastive learning [Fu *et al.*, 2020; Li *et al.*, 2021], or implicitly to aid generalization to new tasks [Lan *et al.*, 2019; Walke *et al.*, 2022; Sodhani *et al.*, 2021b; Vuorio *et al.*, 2019]. While these methods integrate embeddings with policies solely to improve performance, we propose a framework to learn general-purpose embeddings that can be used to quantify and analyze task similarities. Furthermore, in our framework, embedding computation is a one-shot operation, unlike prior work that relies on experience data from the policy for the task. These distinctions position our work as complementary to existing methods.

Population-based techniques. Our framework requires a diverse agent population. This is inline with [Furuta *et al.*, 2021; Tylkin *et al.*, 2021; Vinyals and *et al.*, 2019; Jaderberg and *et al.*, 2019; Parker-Holder *et al.*, 2020], which

¹GitHub repository: <https://github.com/machine-teaching-group/ijcai2024-task-embeddings-rl>.

use agent populations in the RL setting. For instance, Furuta *et al.* [2021] use a randomly generated agent population to empirically estimate *policy information capacity*, an information-theoretic measure of task difficulty in RL.

3 Problem Setup

MDP and Tasks. We use the Markov Decision Process (MDP) framework to define an environment. An MDP \mathcal{M} is defined as a 6-tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \mathcal{S}_{\text{init}}, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $\mathcal{T} : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the transition dynamics, and $\mathcal{S}_{\text{init}} \subseteq \mathcal{S}$ is the set of initial states. Each state $s \in \mathcal{S}_{\text{init}}$ corresponds to a goal-based task (for example, the goal could be to reach a specific destination in a navigation task). Our formalism of tasks as initial states is connected to the contextual MDP setting [Hallak *et al.*, 2015] (see Section 6). The reward is 0 on all transitions but those on which a task gets completed. On task completion, the reward is 1. As an alternative to discounting, at each time step, there is a failure probability of $1 - \gamma$, which incentivises the agent to solve the task quickly. This ensures that the cumulative reward is binary.

Population of agents and task solvability. $p(\theta)$ represents a distribution over the population of agents. Concretely, it is a distribution over the agents’ policy parameters. We use Θ to represent the random variable that takes on the value θ . $\mathcal{O}_{s,\theta} \in \{0, 1\}$ is a Bernoulli random variable that takes on the value 1 if, on a rollout, an agent sampled from $p(\theta)$ could successfully solve the task $s \in \mathcal{S}_{\text{init}}$ (i.e., the cumulative reward is 1), and 0 otherwise. We call $\mathcal{O}_{s,\theta}$ the optimality variable for task s . $\text{POS}(s) := \mathbb{E}[\mathcal{O}_{s,\theta}]$ denotes the probability of success on s , and is the complement of task difficulty.

Task embedding space. Formally, we wish to learn a task embedding function (parameterized by ϕ) $f_\phi : \mathcal{S}_{\text{init}} \rightarrow \mathbb{R}^n$, for an MDP \mathcal{M} and a prior over the population of agents $p(\theta)$, that maps tasks to n -dimensional representations. The range of $f_\phi(\cdot)$ is the task embedding space.

Objective. Our objective is to learn embeddings for sequential tasks with the following properties: (a) the inner product in the embedding space captures task similarity, where the realizations of optimality variables are similar for tasks that are embedded in close proximity, and (b) the norm of the

embedding induces an ordering on the tasks based on their difficulties. We formalize these objectives in Section 4.

4 Learning Framework

In Sections 4.1, we formally define our information-theoretic criterion to measure task similarity in RL. In Section 4.2, we view the problem of learning task embeddings through the lens of ordinal constraint satisfaction.

4.1 Information-Theoretic Task Similarity Measure

Our goal is to measure similarity between sequential tasks. To this end, we propose the mutual information between task optimality variables as a measure of task similarity. This metric captures the intuition that two tasks are similar to each other if observing an agent’s performance on one task reduces our uncertainty about its performance on the other. We begin by formally defining performance uncertainty. Thereafter, we provide a formal definition of our task similarity criterion.

Definition 1 (Performance Uncertainty). *The entropy of the population with prior $p(\theta)$ solving a task s is defined as:*

$$\mathcal{H}(\mathcal{O}_{s,\theta}) = - \sum_{o \in \{0,1\}} P(\mathcal{O}_{s,\theta} = o) \log P(\mathcal{O}_{s,\theta} = o),$$

where $\mathcal{O}_{s,\theta}$ is the optimality variable for s .

Thus, the similarity between tasks $s_i, s_j \in \mathcal{S}_{\text{init}}$ could be measured as the reduction in $\mathcal{H}(\mathcal{O}_{s_i,\theta})$ by observing $\mathcal{O}_{s_j,\theta}$.

Definition 2 (Task Similarity). *Given a prior over the population of agents $p(\theta)$, we measure the similarity between two tasks $s_i, s_j \in \mathcal{S}_{\text{init}}$ as the mutual information $\mathcal{I}(\cdot; \cdot)$ between their optimality variables $\mathcal{O}_{s_i,\theta}, \mathcal{O}_{s_j,\theta}$:*

$$\mathcal{I}(\mathcal{O}_{s_i,\theta}; \mathcal{O}_{s_j,\theta}) = \mathcal{H}(\mathcal{O}_{s_i,\theta}) - \mathcal{H}(\mathcal{O}_{s_i,\theta} | \mathcal{O}_{s_j,\theta}).$$

It quantifies the information obtained about $\mathcal{O}_{s_i,\theta}$ by observing $\mathcal{O}_{s_j,\theta}$.

A comprehensive algorithm detailing the empirical estimation of \mathcal{I} is provided in the longer version of the paper [Mahajan *et al.*, 2023]. Given an MDP \mathcal{M} and a prior distribution of the agent parameters $p(\theta)$, our algorithm uses N samples to estimate $\mathcal{I}(\mathcal{O}_{s_i,\theta}; \mathcal{O}_{s_j,\theta})$. For each sample, it randomly samples $\theta_l \sim p(\theta)$, and performs rollouts of π_{θ_l} from s_i and s_j to obtain estimates of the probability mass functions required for computing \mathcal{I} . The estimation procedure can be invoked with the signature $\text{ESTIMATE}(s_i, s_j, \mathcal{M}, \pi, p(\theta), N)$.

4.2 Learning Task Embeddings

With the task similarity criterion defined, we are interested in learning a task embedding function $f_\phi: \mathcal{S}_{\text{init}} \rightarrow \mathbb{R}^n$ (consequently, an embedding space) that satisfies the desiderata introduced in Section 3. To this end, we pose the problem of learning $f_\phi(\cdot)$ as an ordinal constraint satisfaction problem. Essentially, the similarity criterion \mathcal{I} imposes a set \mathcal{C}_{MI} of triplet ordinal constraints on the task embeddings. $\text{POS}(\cdot)$ imposes another set $\mathcal{C}_{\text{NORM}}$ of pairwise ordinal constraints.

Concretely, \mathcal{C}_{MI} is a collection of ordered triplets of tasks s.t. for each $(s_1, s_2, s_3) \in \mathcal{C}_{\text{MI}}$, $\mathcal{I}(\mathcal{O}_{s_1,\theta}; \mathcal{O}_{s_2,\theta}) > \mathcal{I}(\mathcal{O}_{s_1,\theta}; \mathcal{O}_{s_3,\theta})$. Consequently,

Algorithm 1 Learn the Task Embedding Function (f_ϕ)

```

1: procedure TRAIN(Set of tasks  $\mathcal{S}_{\text{init}}$ , MDP  $\mathcal{M}$ , Policy  $\pi$ , Prior
   distribution of the agent parameters  $p(\theta)$ , Number of samples  $N$ ,
   Hyperparameter  $\lambda$ , Number of iterations  $M$ )
2:   Initialize  $\phi$ .
3:   for  $i \in \{1, \dots, M\}$  do
4:     Sample task  $s_1, s_2, s_3 \sim \mathcal{S}_{\text{init}}$ .
5:      $E_1, E_2, E_3 \leftarrow f_\phi(s_1), f_\phi(s_2), f_\phi(s_3)$ 
6:      $\hat{\mathcal{I}}_{12} \leftarrow \text{ESTIMATE}(s_1, s_2, \mathcal{M}, \pi, p(\theta), N)$ 
7:      $\hat{\mathcal{I}}_{13} \leftarrow \text{ESTIMATE}(s_1, s_3, \mathcal{M}, \pi, p(\theta), N)$ 
8:     if  $\hat{\mathcal{I}}_{12} > \hat{\mathcal{I}}_{13}$  then
9:       loss  $\leftarrow \log(1 + \exp(\langle E_1, E_3 \rangle - \langle E_1, E_2 \rangle))$ 
10:    else
11:      loss  $\leftarrow \log(1 + \exp(\langle E_1, E_2 \rangle - \langle E_1, E_3 \rangle))$ 
12:    Sample task  $s_4, s_5 \sim \mathcal{S}_{\text{init}}$ .
13:     $E_4, E_5 \leftarrow f_\phi(s_4), f_\phi(s_5)$ 
14:    if  $\text{POS}(s_4) > \text{POS}(s_5)$  then
15:      loss  $\leftarrow \text{loss} + \lambda \log(1 + \exp(\|E_4\|_2 - \|E_5\|_2))$ 
16:    else
17:      loss  $\leftarrow \text{loss} + \lambda \log(1 + \exp(\|E_5\|_2 - \|E_4\|_2))$ 
18:    Update  $\phi$  to minimize loss.
19:  return  $\phi$ 

```

we would like to satisfy the constraint $\langle f_\phi(s_1), f_\phi(s_2) \rangle > \langle f_\phi(s_1), f_\phi(s_3) \rangle$. Likewise, $\mathcal{C}_{\text{NORM}}$ is a collection of ordered tuples of tasks s.t. for each $(s_1, s_2) \in \mathcal{C}_{\text{NORM}}$, $\text{POS}(s_1) > \text{POS}(s_2)$. Consequently, we would like to satisfy the constraint $\|f_\phi(s_2)\|_2 > \|f_\phi(s_1)\|_2$ (embeddings for easier tasks have smaller norm).

We learn the task embedding function $f_\phi(\cdot)$, for an MDP \mathcal{M} and a prior over the agent population $p(\theta)$, by optimizing the parameters ϕ to maximize the log-likelihood of the ordinal constraints under the Bradley-Terry-Luce (BTL) model [Luce, 1959]. Concretely, given a task triplet (s_1, s_2, s_3) , we define:

$$P((s_1, s_2, s_3) \in \mathcal{C}_{\text{MI}}) := \frac{\exp(\langle f_\phi(s_1), f_\phi(s_2) \rangle)}{\exp(\langle f_\phi(s_1), f_\phi(s_2) \rangle) + \exp(\langle f_\phi(s_1), f_\phi(s_3) \rangle)}.$$

Similarly, given a tuple of tasks (s_1, s_2) , we define:

$$P((s_1, s_2) \in \mathcal{C}_{\text{NORM}}) := \frac{\exp(\|f_\phi(s_2)\|_2)}{\exp(\|f_\phi(s_1)\|_2) + \exp(\|f_\phi(s_2)\|_2)}.$$

Hence, the task embedding function $f_\phi(\cdot)$ is learned by solving the following optimization problem:

$$\min_{\phi} \left[\mathbb{E}_{(s_1, s_2, s_3) \sim \mathcal{C}_{\text{MI}}} \log \left(1 + \exp(\langle E_1, E_3 \rangle - \langle E_1, E_2 \rangle) \right) + \lambda \mathbb{E}_{(s_4, s_5) \sim \mathcal{C}_{\text{NORM}}} \log \left(1 + \exp(\|E_4\|_2 - \|E_5\|_2) \right) \right],$$

where E_i denotes $f_\phi(s_i)$, and λ is a hyperparameter. The pseudocode for the proposed algorithm to learn the task embedding function $f_\phi(\cdot)$ is given in Algorithm 1.

5 Experiments: Visualizing Embedding Spaces

In this section, we visualize the embedding spaces to gather qualitative insights, addressing the following research questions: (i) Can distinct clusters of tasks be identified by visualizing the embedding space? (ii) How does regularization

Environment	Task Variability	Action	State	Number of Tasks
MULTIKEYNAV	Reward Function	7	$\mathbb{R} \times \{0, 1\}^6$	Infinite
CARTPOLEVAR	Dynamics	2	$\mathbb{R}^5 \times \{0, 1\} \times [200]$	Infinite
POINTMASS	Dynamics	\mathbb{R}^2	\mathbb{R}^7	Infinite
KAREL	Reward Function + Dynamics	52	$\{0, 1\}^{51840}$	73688
BASICKAREL	Reward Function + Dynamics	6	$\{0, 1\}^{88}$	24000




Figure 2: We evaluate our framework on a diverse set of environments. (a) compares the characteristics of these environments. (b) illustrates these environments for a better understanding of the tasks.

through $\mathcal{C}_{\text{NORM}}$ affect the embedding space? (iii) What influence do agent population and environment specification have on the embedding space? We begin by discussing the rationale for environment selection, describing these environments. Subsequently, we provide an overview of the embedding networks' training process, followed by the qualitative results.

5.1 Environments

We evaluate our framework on environments with diverse characteristics to demonstrate its generality and scalability to different sequential decision-making problems (see Fig. 2). As the running example, we use MULTIKEYNAV (based on [Devizze *et al.*, 2021]) because of its compositional nature in which the agent needs to compose different actions for picking keys (with four distinct key types, each requiring a specific action to be picked) in a task-specific manner to unlock the door. This also makes it suitable for ablation experiments. Task variability comes from the agent's initial position, the keys that it possesses initially, and the door type (with each type requiring a unique set of keys).

Given that task variability in MULTIKEYNAV comes from the reward function, we use CARTPOLEVAR to highlight our framework's applicability to environments where it comes from the dynamics instead. This environment is a variation of the classic control task from OpenAI gym [Brockman *et al.*, 2016], and also takes inspiration from [Sodhani *et al.*, 2021a] in which the forces could be negative as well. Tasks require keeping a pole attached by an unactuated joint to a cart upright for 200 timesteps by applying forces to the left (action 0) or to the right (action 1) of the cart. Task variability comes from the force F applied on the cart by each action, and the TaskType $\in \{0, 1\}$. Tasks of Type 0 involve "Pulling" with action 0 pulling the cart from the left and action 1 pulling the cart from the right, while tasks of Type 1 involve "Pushing".

We select POINTMASS (introduced in [Klink *et al.*, 2020a]) to test if our framework can handle continuous action spaces. Here, the agent applies forces to control a point mass inside a walled square. Tasks require reaching a fixed goal position through a gate, with variability arising from the gate width and

position, along with the space's kinetic friction coefficient.

Finally, we investigate our framework's scalability using the real-world environment KAREL from [Bunel *et al.*, 2018], which is a challenging environment (with high-dimensional states) that has applications in programming education. Tasks require the agent to synthesize a program, potentially containing control flow constructs such as loops and conditionals, satisfying a given specification comprising input-output examples. This program serves as a controller for an avatar navigating a grid, where each cell could contain marker(s), correspond to a wall, or be empty. The avatar can traverse the grid and manipulate it by picking or placing markers. Thus, an example in the specification comprises the Pre-Grid and the corresponding Post-Grid. In addition, we evaluate it on BASICKAREL [Tzannetos *et al.*, 2023], which is a simpler variant of KAREL that excludes control flow constructs.

5.2 Training Process

To learn the embedding function, we first obtain the agent population by taking snapshots while training a neural network policy using either behavioral cloning [Bain and Sammut, 1995] or policy gradient methods [Sutton *et al.*, 1999]. Concretely, a snapshot is recorded if the average performance on a validation set of tasks (denoted as $\mathcal{S}_{\text{snap}}$) improves by δ_{snap} compared to the previously recorded snapshot. A snapshot of the untrained policy is recorded by default. Different subpopulations, obtained by either masking actions or by using biased task distributions during training, are combined to form the final population. Here, masking an action corresponds to setting its logit to a large negative number. Using biased task distribution during training is another way to inject diversity into the population. In MULTIKEYNAV, for instance, this could correspond to assigning low probability mass to tasks with certain types of doors in the initial state distribution during training. Finally, we parameterize the embedding function $f_{\phi}(\cdot)$ with a neural network, optimizing its parameters as described in Algorithm 1. We provide additional details in the longer version of the paper [Mahajan *et al.*, 2023].

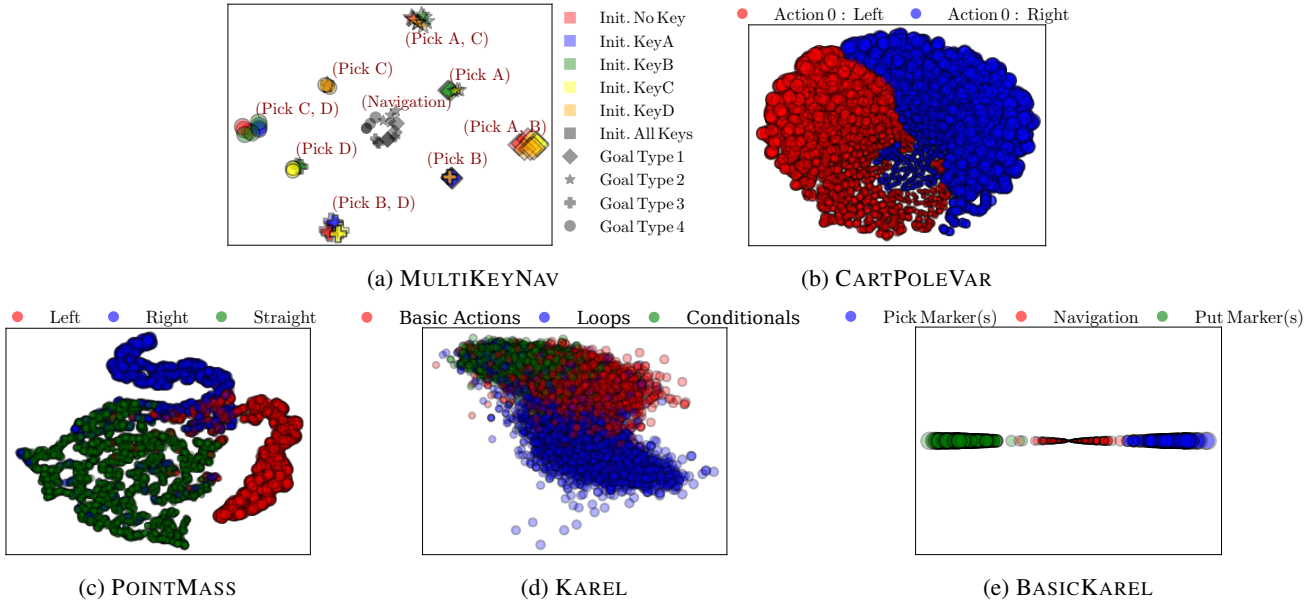


Figure 3: Visualization of the task embedding spaces learned through our framework. Each point represents a task, and the size of the points is proportional to the norm of the embeddings.

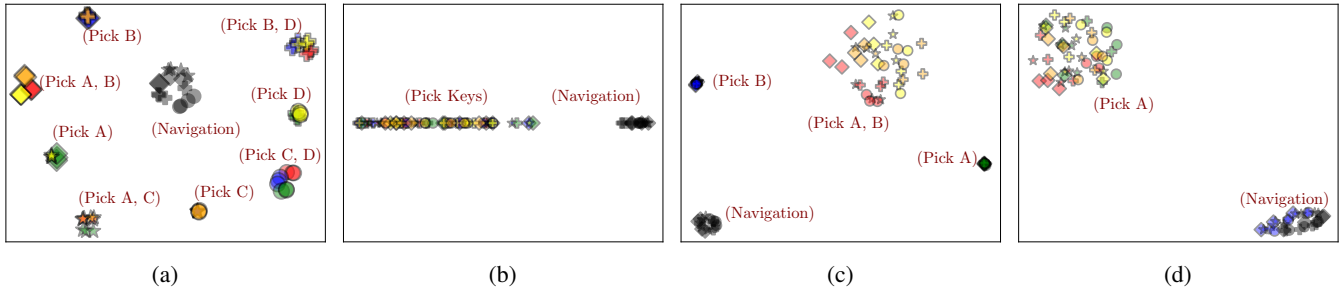


Figure 4: Task embedding spaces for the MULTIKEYNAV environment: (a) without C_{NORM} , (b) `pickKey` actions masked, (c) all doors require KeyA, KeyB, and (d) all doors require KeyA.

5.3 Visualizations and Qualitative Results

We visualize the embedding spaces on a 2-dimensional map using t-SNE [van der Maaten and Hinton, 2008] to identify distinct clusters of tasks. Although t-SNE preserves the local structure, it does not necessarily preserve the embeddings’ norm. For this reason, we scale the points in proportion to the norm of the embeddings. Additionally, we provide PCA plots in the longer version of the paper [Mahajan *et al.*, 2023].

Visualizations. For MULTIKEYNAV (Fig. 3a), our framework discovers distinct clusters of tasks, with each cluster corresponding to a unique set of keys that need to be picked. The norm of the embeddings is in accordance with the number of keys that need to be picked (with tasks requiring navigation only having the smallest norm). Additionally, tasks in clusters adjacent to each other share a common key requirement. For CARPOLEVAR (Fig. 3b), it discovers that each task exhibits one of two types of underlying dynamics. In one (+ve F and Type 0, or -ve F and Type 1), action 0 moves the cart to the left, while in the other (-ve F and Type 0, or +ve F and Type 1), it moves the cart to the right. For POINTMASS (Fig. 3c),

it discovers three clusters of tasks based on the behavior that the agent needs to exhibit near the gate. The first cluster includes tasks in which the agent need not steer to cross the gate, while the second and third clusters contain tasks in which it must steer left or right to cross the gate, respectively. For KAREL and BASICKAREL (Fig. 3d and 3e), it discovers different clusters based on whether the solution code requires loops or conditionals, and whether the agent needs to pick or put markers in the grid, respectively. In the longer version of the paper [Mahajan *et al.*, 2023], we show that the learned embedding network can generalize to new tasks.

Ablation w.r.t. C_{NORM} . Fig. 4a shows the task embedding space learned without the norm ordinal constraints C_{NORM} (i.e., λ is set to 0). As expected, the norm of the embeddings is not proportional to the number of keys that need to be picked. Instead, the points are nearly uniform in size.

Ablation w.r.t. population specification. To understand the effect of population on the embedding space, we learn the embedding function $f_\phi(\cdot)$ for MULTIKEYNAV using an agent population in which `pickKey` actions are masked (Fig. 4b).

In this case, we obtain two distinct clusters of tasks – one of the clusters contains tasks that cannot be solved (these require picking key(s)), and the other contains tasks that require navigation only. These results emphasize the importance of the population’s quality in learning a good task embedding space.

Ablation w.r.t. environment specification. We check the impact of the environment specification on the embedding space through the following `MULTIKEYNAV` variants: (a) each door requires KeyA and KeyB (Fig. 4c), i.e., they have identical key requirements, and (b) each door requires KeyA only (Fig. 4d). Modifying the environment specification changes the task semantics, thereby impacting the embedding space. Thus, these results are inline with our intuition.

6 Experiments: Comparison with Prior Work

In this section, we compare our framework with prior work using silhouette scores, which assess clustering quality in the embedding spaces based on intuitively identified task clusters. To enable this comparison, we introduce the *PredModel* baseline, inspired by prior work (e.g., PEARL [Rakelly et al., 2019]). We also compare our method against embedding networks with random weights (*RandomModel*).

Most existing methods use variational inference to learn latent context from task-specific experience data, where the inference network could be trained to reconstruct the MDP for the task through predictive models of reward and dynamics. To adapt this approach to our setting, we connect our formalism of tasks as initial states to the contextual MDP setting [Hallak et al., 2015], where each context (e.g., `MULTIKEYNAV`’s context: agent’s initial position, possessed keys initially, door type) corresponds to a distinct task represented by a separate MDP with context-dependent transitions and rewards. By including context variables as part of the state, this set of MDPs can be converted into an equivalent MDP, where each initial state represents a task.

The modifications needed for *PredModel* are as follows: Firstly, since context is observable in our setup, we condition the approximate posterior over the embeddings on the initial state, eliminating the need for experience data. Secondly, we train the predictive models on states with context variables removed, ensuring the utilization of the embedding that the model is conditioned on. We provide additional details in the longer version of the paper [Mahajan et al., 2023].

Results. Table 1 reports the silhouette scores, averaged across 3 random seeds, with 1000 tasks per seed (5000 for `BASICKAREL`). The scores for the models learned through our framework are consistently better. While *PredModel* clusters similar tasks together in the embedding space for `CARTPOLEVAR`, it fails to do so in rest of the environments. In contrast to `CARTPOLEVAR`, where task variability comes from dense differences in the dynamics, task variability in other environments comes from sparse differences in the reward function and/or dynamics. Thus, we hypothesize that *PredModel* fails on environments with sparse variability across tasks.

7 Experiments: Application Scenarios

In this section, we evaluate our framework on two application scenarios: performance prediction, and task selection. We con-

Environment	RandomModel	PredModel	Ours
<code>MULTIKEYNAV</code>	+0.036 ± 0.048	-0.037 ± 0.003	+0.753 ± 0.001
<code>CARTPOLEVAR</code>	+0.015 ± 0.016	+0.242 ± 0.007	+0.325 ± 0.009
<code>POINTMASS</code>	+0.104 ± 0.026	-0.010 ± 0.004	+0.380 ± 0.019
<code>BASICKAREL</code>	-0.058 ± 0.007	-0.002 ± 0.003	+0.811 ± 0.019

Table 1: Comparison of silhouette scores (higher is better) based on intuitively identified clusters of tasks in the learned embedding spaces. The scores for our models are consistently better.

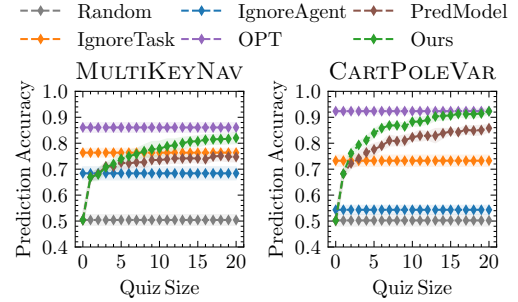


Figure 5: Results for performance prediction using task embeddings. Our method (listed as *Ours*) is competitive with the *OPT* baseline, which is the best one could do on this benchmark.

duct this evaluation on `MULTIKEYNAV` and `CARTPOLEVAR`, as they cover two distinct sources of task variability, namely reward function and dynamics.

7.1 Performance Prediction

First, we assess the learned task embeddings by using them to predict an agent’s performance on a task $s_{\text{test}} \in \mathcal{S}_{\text{init}}$ after observing its performance on a quiz $\mathcal{S}_{\text{quiz}} \subseteq \mathcal{S}_{\text{init}}$. This requires estimating its probability of success on a task, which is precisely its value function when the reward is binary and sparse. Specifically, we seek to answer the following research question: Would an agent show similar performance on tasks that are in close proximity in the learned embedding space? We begin by creating a benchmark for this application scenario, and then compare our method against various baselines.

Benchmark. Given the realizations of the optimality variables of a set of tasks for an agent θ , we are interested in predicting the most probable realization of the optimality variable of a new task for this agent without observing θ . We create benchmarks for this scenario by generating datasets for quiz sizes ranging from 1 to 20, with 5000 examples for both training and testing. Each example is generated by randomly sampling a quiz $\mathcal{S}_{\text{quiz}}$ of desired size, along with a task s_{test} from $\mathcal{S}_{\text{init}}$, and then recording the performance of a randomly sampled agent θ on these tasks. Performance prediction techniques are evaluated by partitioning each dataset into 10 folds, and measuring the mean prediction accuracy across the folds, along with the standard error.

Our approach. We perform soft-nearest neighbor matching of s_{test} with $\mathcal{S}_{\text{quiz}}$ in the task embedding space to predict performance on s_{test} . Concretely, given the embedding function $f_{\phi}(\cdot)$, the prediction is $\mathbf{1}_{c>0.5}$, where c equals $\frac{\sum_{s \in \mathcal{S}_{\text{quiz}}} o_s \exp(-\beta \|f_{\phi}(s) - f_{\phi}(s_{\text{test}})\|_2^2)}{\sum_{s \in \mathcal{S}_{\text{quiz}}} \exp(-\beta \|f_{\phi}(s) - f_{\phi}(s_{\text{test}})\|_2^2)}$, o_s is the realization of




Figure 6: Task selection results using embeddings (dark bars represent Top-3 accuracy and light bars represent Top-1). Our method (listed as *Ours*) is competitive with \widehat{OPT}_{50} . Further, it outperforms *Ours_{woNorm}* on Type-2 queries, highlighting $\mathcal{C}_{\text{NORM}}$'s significance in our framework.

the optimality variable for task s , and β is a hyperparameter.

Baselines. Besides *PredModel*, we compare against different levels of oracle knowledge: (i) *Random*: Randomly predicts the agent’s performance. (ii) *IgnoreTask*: Predicts the agent to succeed on s_{test} iff the probability that it succeeds on a random task exceeds 0.5. (iii) *IgnoreAgent*: Predicts success on s_{test} iff the probability that a random agent succeeds on it exceeds 0.5. (iv) *OPT*: Predicts the agent to succeed on s_{test} iff the probability that it succeeds on s_{test} exceeds 0.5.

Results. Fig. 5 shows the prediction accuracies of various methods. Our method is competitive with *OPT*, which is an upper-bound on the prediction accuracy but relies on the unrealistic assumption of full observability of both the agent and task. Additionally, our proposed task similarity criterion, which is used to learn the embeddings, captures performance correlations that remain valid for new agents (see the longer version of the paper [Mahajan *et al.*, 2023]).

7.2 Task Selection

Next, we assess the learned embeddings by using them to select tasks with desired characteristics. Specifically, we seek to answer the following research questions: (i) Does the inner product in the learned embedding space capture task similarity according to our information-theoretic criterion? (ii) Does the norm of the embedding learned by our framework induce an ordering on the tasks based on their difficulties? We begin by creating a benchmark for this scenario, and then compare our method based on task embeddings against various baselines.

Benchmark. Amongst several options of tasks $\mathcal{S}_{\text{options}}$, we are interested in choosing the task that best matches the desired characteristics, which we categorize into two query types: *Type-1*: Select the task that is the most similar to a given reference task s_{ref} . The ground-truth answer to this query is $\arg \max_{s \in \mathcal{S}_{\text{options}}} \mathcal{I}(\mathcal{O}_{s_{\text{ref}}}, \Theta; \mathcal{O}_s, \Theta)$. *Type-2*: Select the task that is the most similar to (but harder than) a given reference task s_{ref} . Out of all tasks in $\mathcal{S}_{\text{options}}$ that are harder than s_{ref} , the ground-truth answer to this query is the task most similar to it. To create benchmarks for this scenario, we generate a dataset of 50 examples. Each example consists of a randomly sampled s_{ref} and 10 tasks that form $\mathcal{S}_{\text{options}}$. The benchmarks also include 5 easy tasks for reference (determined by ranking a randomly sampled pool of 500 tasks). We evaluate task selection methods by reporting mean selection accuracy across 4 randomly sampled datasets, along with the standard error.

Our approach. We use task embeddings to rank the options according to similarity and/or difficulty, based on which the

selection is made. We additionally compare our method based on embeddings learned without $\mathcal{C}_{\text{NORM}}$ (listed as *Ours_{woNorm}*).

Baselines. Besides *PredModel*, we compare against the following baselines: (i) *Random*: Randomly selects answers from $\mathcal{S}_{\text{options}}$. (ii) *StateSim*: Measures task similarity based on state representation distances. For queries of type 2, it considers a task s_1 to be harder than s_2 iff the similarity between s_1 and the task most similar to it in the set of easy reference tasks, is less than that for s_2 . (iii) *TrajectorySim*: Measures task similarity using the edit distance between expert trajectories. (iv) *OPT*: Estimates task similarity and difficulty using the entire agent population. Given the variance in the estimation process, this is the best one could do on this benchmark. (v) \widehat{OPT}_{50} : Estimates task similarity and difficulty using a randomly sampled 50% of the population.

Results. Fig. 6 compares different methods on the task selection benchmark. Our method outperforms *Random*, *StateSim*, *TrajectorySim*, and *PredModel*, and is competitive with \widehat{OPT}_{50} . This suggests that the inner product in the learned task embedding space successfully captures similarity between tasks. Notably, our technique significantly outperforms *Ours_{woNorm}* on Type-2 queries, indicating that the norm of the embedding effectively orders tasks by difficulty.

8 Concluding Discussion

In this work, we introduced an information-theoretic framework for learning task embeddings in sequential decision-making settings. Through experiments on diverse environments, we empirically demonstrated that the inner product in the embedding space captures task similarity, and the norm of the embedding induces an ordering on the tasks based on their difficulties. A limitation of our current framework is the requirement for tasks to be goal-based, which we plan to address in future work. This could involve using the difference between the cumulative reward obtained during the rollout and the maximum achievable cumulative reward for the given task to parameterize the optimality variable. Additionally, it would be interesting to explore more principled methods for constructing the agent population that explicitly optimize for diversity. Further, empirically estimating the proposed similarity criterion by directly estimating the underlying mass functions could be sample-inefficient for some environments. Thus, a promising direction is to construct sample-efficient estimators for it. Moreover, evaluation in multi-agent settings, where the task embedding could encode the behavior of non-ego agents, is another interesting direction.

Acknowledgments

Funded/Co-funded by the European Union (ERC, TOPS, 101039090). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [Achille *et al.*, 2019] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhansu Maji, Charless C. Fowlkes, Stefano Soatto, and Pietro Perona. Task2Vec: Task Embedding for Meta-Learning. In *ICCV*, 2019.
- [Arnekvist *et al.*, 2018] Isac Arnekvist, Danica Kragic, and Johannes Andreas Stork. VPE: Variational Policy Embedding for Transfer Reinforcement Learning. In *ICRA*, 2018.
- [Athar *et al.*, 2020] Ali Athar, Sabarinath Mahadevan, Aljoša Ošep, Laura Leal-Taixé, and Bastian Leibe. STEM-Seg: Spatio-temporal Embeddings for Instance Segmentation in Videos. In *ECCV*, 2020.
- [Bain and Sammut, 1995] Michael Bain and Claude Sammut. A Framework for Behavioural Cloning. In *Machine Intelligence, Intelligent Agents*, 1995.
- [Bing *et al.*, 2023] Z. Bing, D. Lerch, K. Huang, and A. Knoll. Meta-Reinforcement Learning in Non-Stationary and Dynamic Environments. In *IEEE TPAMI*, 2023.
- [Brockman *et al.*, 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.
- [Bunel *et al.*, 2018] Rudy Bunel, Matthew J. Hausknecht, Jacob Devlin, Rishabh Singh, and Pushmeet Kohli. Leveraging Grammar and Reinforcement Learning for Neural Program Synthesis. In *ICLR*, 2018.
- [Cer *et al.*, 2018] Daniel Matthew Cer, Yinfei Yang, Shengyi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. Universal Sentence Encoder for English. In *EMNLP*, 2018.
- [Devidze *et al.*, 2021] Rati Devidze, Goran Radanovic, Parameswaran Kamalaruban, and Adish Singla. Explorable Reward Design for Reinforcement Learning Agents. In *NeurIPS*, 2021.
- [Eimer *et al.*, 2021] Theresa Eimer, André Biedenkapp, Frank Hutter, and Marius Lindauer. Self-Paced Context Evaluation for Contextual Reinforcement Learning. In *ICML*, 2021.
- [Fu *et al.*, 2020] Haotian Fu, Hongyao Tang, Jianye Hao, Chen Chen, Xidong Feng, Dong Li, and Wulong Liu. Towards Effective Context for Meta-Reinforcement Learning: an Approach based on Contrastive Learning. In *AAAI*, 2020.
- [Furuta *et al.*, 2021] Hiroki Furuta, Tatsuya Matsushima, Tadashi Kozuno, Yutaka Matsuo, Sergey Levine, Ofir Nachum, and Shixiang Shane Gu. Policy Information Capacity: Information-Theoretic Measure for Task Complexity in Deep Reinforcement Learning. In *ICML*, 2021.
- [Ghosh *et al.*, 2022] Ahana Ghosh, Sebastian Tschiatschek, Sam Devlin, and Adish Singla. Adaptive Scaffolding in Block-based programming via Synthesizing New Tasks as Pop quizzes. In *AIED*, 2022.
- [Gupta *et al.*, 2018] Abhishek Gupta, Russell Mendonca, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Meta-Reinforcement Learning of Structured Exploration Strategies. In *NIPS*, 2018.
- [Hallak *et al.*, 2015] Assaf Hallak, Dotan Di Castro, and Shie Mannor. Contextual Markov Decision Processes. *CoRR*, abs/1502.02259, 2015.
- [He-Yueya and Singla, 2021] Joy He-Yueya and Adish Kumar Singla. Quizzing Policy Using Reinforcement Learning for Inferring the Student Knowledge State. In *EDM*, 2021.
- [Jaderberg and et al., 2019] Max Jaderberg and et al. Human-level Performance in 3D Multiplayer Games with Population-based Reinforcement Learning. In *Science*, 2019.
- [Klink *et al.*, 2020a] P. Klink, C. D’Eramo, J. Peters, and J. Pajarinen. Self-Paced Deep Reinforcement Learning. In *NeurIPS*, 2020.
- [Klink *et al.*, 2020b] Pascal Klink, Carlo D’Eramo, Jan R Peters, and Joni Pajarinen. Self-Paced Deep Reinforcement Learning. In *NeurIPS*, 2020.
- [Lan *et al.*, 2019] Lin Lan, Zhenguo Li, Xiaohong Guan, and Pinghui Wang. Meta Reinforcement Learning with Task Embedding and Shared Policy. In *IJCAI*, 2019.
- [Li *et al.*, 2021] Lanqing Li, Rui Yang, and Dijun Luo. FOCAL: Efficient Fully-Offline Meta-Reinforcement Learning via Distance Metric Learning and Behavior Regularization. In *ICLR*, 2021.
- [Luce, 1959] R. Duncan Luce. *Individual Choice Behavior: A Theoretical analysis*. Wiley, 1959.
- [Mahajan *et al.*, 2023] Mridul Mahajan, Georgios Tzannetos, Goran Radanovic, and Adish Singla. Learning Embeddings for Sequential Tasks Using Population of Agents. *CoRR*, abs/2306.03311, 2023.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, 2013.
- [Parker-Holder *et al.*, 2020] Jack Parker-Holder, Aldo Pacchiano, Krzysztof M Choromanski, and Stephen J Roberts. Effective Diversity in Population Based Reinforcement Learning. In *NeurIPS*, 2020.
- [Peng *et al.*, 2020] Xingchao Peng, Yichen Li, and Kate Saenko. Domain2Vec: Domain Embedding for Unsupervised Domain Adaptation. In *ECCV*, 2020.

- [Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global Vectors for Word Representation. In *EMNLP*, 2014.
- [Qin *et al.*, 2022] Rongjun Qin, F. Chen, Tonghan Wang, Lei Yuan, Xiaoran Wu, Zongzhang Zhang, Chongjie Zhang, and Yang Yu. Multi-Agent Policy Transfer via Task Relationship Modeling. *CoRR*, abs/2203.04482, 2022.
- [Rakelly *et al.*, 2019] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient Off-Policy Meta-Reinforcement Learning via Probabilistic Context Variables. In *ICML*, 2019.
- [Schäfer *et al.*, 2022] Lukas Schäfer, Filippos Christianos, Amos J. Storkey, and Stefano V. Albrecht. Learning Task Embeddings for Teamwork Adaptation in Multi-Agent Reinforcement Learning. *CoRR*, abs/2207.02249, 2022.
- [Sodhani *et al.*, 2021a] Shagun Sodhani, Ludovic Denoyer, Pierre-Alexandre Kamienny, and Olivier Delalleau. MTEnv - Environment interface for multi-task reinforcement learning. GitHub, 2021.
- [Sodhani *et al.*, 2021b] Shagun Sodhani, Amy Zhang, and Joelle Pineau. Multi-Task Reinforcement Learning with Context-based Representations. In *ICML*, 2021.
- [Sun *et al.*, 2014] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep Learning Face Representation from Predicting 10,000 Classes. In *CVPR*, 2014.
- [Sung *et al.*, 2018] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. Learning to Compare: Relation Network for Few-Shot Learning. In *CVPR*, 2018.
- [Sutton *et al.*, 1999] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *NIPS*, 1999.
- [Tylkin *et al.*, 2021] Paul Tylkin, Goran Radanovic, and David C. Parkes. Learning Robust Helpful Behaviors in Two-Player Cooperative Atari Environments. In *AAMAS*, 2021.
- [Tzannetos *et al.*, 2023] Georgios Tzannetos, Bárbara Gomes Ribeiro, Parameswaran Kamalaruban, and Adish Singla. Proximal Curriculum for Reinforcement Learning Agents. In *TMLR*, 2023.
- [van der Maaten and Hinton, 2008] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. In *JMLR*, 2008.
- [Vinyals and et al., 2019] Oriol Vinyals and et al. Grandmaster Level in StarCraft II Using Multi-agent Reinforcement Learning. In *Nature*, 2019.
- [Vuorio *et al.*, 2019] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal Model-Agnostic Meta-Learning via Task-Aware Modulation. In *NeurIPS*, 2019.
- [Walke *et al.*, 2022] Homer Walke, Jonathan Yang, Albert Yu, Aviral Kumar, Jędrzej Orbik, Avi Singh, and Sergey Levine. Don't Start From Scratch: Leveraging Prior Data to Automate Robotic Reinforcement Learning. In *CoRL*, 2022.
- [Yengera *et al.*, 2021] Gaurav Yengera, Rati Devidze, Parameswaran Kamalaruban, and Adish Singla. Curriculum Design for Teaching via Demonstrations: Theory and Applications. In *NeurIPS*, 2021.
- [Yoo *et al.*, 2022] Minjong Yoo, Sangwoo Cho, and Honguk Woo. Skills Regularized Task Decomposition for Multi-task Offline Reinforcement Learning. In *NeurIPS*, 2022.
- [Zhang *et al.*, 2018] Yu Zhang, Ying Wei, and Qiang Yang. Learning to Multitask. In *NeurIPS*, 2018.
- [Zhang *et al.*, 2021] Amy Zhang, Rowan Thomas McAllister, Roberto Calandra, Yarın Gal, and Sergey Levine. Learning Invariant Representations for Reinforcement Learning without Reconstruction. In *ICLR*, 2021.