

# RePreM: Representation Pre-training with Masked Model for Reinforcement Learning

Yuanying Cai<sup>1</sup>, Chuheng Zhang<sup>2</sup>, Wei Shen<sup>3</sup>, Xuyun Zhang<sup>4</sup>, Wenjie Ruan<sup>4,5</sup>, Longbo Huang<sup>1\*</sup>

<sup>1</sup> IIIS, Tsinghua University, Beijing, China

<sup>2</sup> Microsoft Research Asia, Beijing, China

<sup>3</sup> Hulu, Beijing, China

<sup>4</sup> Macquarie University, Sydney, Australia

<sup>5</sup> University of Exeter, Exeter, UK

cai-yy16@mails.tsinghua.edu.cn, chuhengzhang@microsoft.com, longbohuang@tsinghua.edu.cn

## Abstract

Inspired by the recent success of sequence modeling in RL and the use of masked language model for pre-training, we propose a masked model for pre-training in RL, RePreM (**R**epresentation **P**re-training with **M**asked Model), which trains the encoder combined with transformer blocks to predict the masked states or actions in a trajectory. RePreM is simple but effective compared to existing representation pre-training methods in RL. It avoids algorithmic sophistication (such as data augmentation or estimating multiple models) with sequence modeling and generates a representation that captures long-term dynamics well. Empirically, we demonstrate the effectiveness of RePreM in various tasks, including dynamic prediction, transfer learning, and sample-efficient RL with both value-based and actor-critic methods. Moreover, we show that RePreM scales well with dataset size, dataset quality, and the scale of the encoder, which indicates its potential towards big RL models.

## Introduction

Reinforcement learning (RL) has achieved great success in solving various decision-making tasks [Mnih et al. 2015; Silver et al. 2016; Arulkumaran, Cully, and Togelius 2019; Jumper et al. 2021]. Yet, the sample inefficiency issue [Tsividis et al. 2017] and the overfitting problem [Zhang et al. 2018] impede its further application to real-world scenarios. In contrast, humans can typically learn more efficiently and robustly than current RL algorithms. The gap is resulted from the fact that existing RL algorithms usually start learning *tabula rasa* without any prior knowledge [Schwarzer et al. 2021; Hutsebaut-Buysse, Mets, and Latré 2020]. As a result, it requires a large number of samples to learn meaningful representations [Dubey et al. 2018; Lake et al. 2017].

In this paper, we consider learning state representations that incorporate rich knowledge by unsupervised pre-training from offline datasets. While unsupervised pre-training methods have achieved great success in natural language processing (NLP) [Devlin et al. 2018; Brown et al. 2020] and computer vision (CV) [Henaff 2020; He et al.

2020; Grill et al. 2020; Chen et al. 2020], they are relatively underexplored in RL. The main benefit of pre-trained representations is that they can be used for sample-efficient learning for various downstream tasks such as dynamics prediction, transfer learning, and reinforcement learning with different algorithms and different reward functions.

However, existing methods that generate pre-trained representations suffer from the following limitations: 1) They are algorithmically complex. Some of these methods, especially those based on self-supervised learning, rely on hand-crafted data augmentation that needs specific domain knowledge [see e.g., Stooke et al. 2021; Schwarzer et al. 2021], while others are based on the world model and require additional networks to model latent transition or reward functions [see e.g., Seo et al. 2022]. 2) We find that, although several techniques, such as goal-conditioned RL [Schwarzer et al. 2021] and RSSM [Hafner et al. 2019b], are designed to encourage representations to incorporate the information over the multi-step dynamics in the future, previous methods cannot predict long-term dynamics accurately, which limits their performance (see our later experiments).

Recently, sequence modeling (mostly with Transformer) has achieved great success in RL [Chen et al. 2021; Janner, Li, and Levine 2021]. This approach can avoid algorithmic sophistication by following the RvS (RL via supervised learning) paradigm [Emmons et al. 2021] and extract long-term information more effectively by learning the whole sequence. However, most previous work on sequence modeling for RL focuses on the offline RL setting [Chen et al. 2021; Janner, Li, and Levine 2021; Reed et al. 2022; Furuta, Matsuo, and Gu 2021; Akhmetov et al. 2022; Xu et al. 2022], where the agent’s performance is limited by the quality of the offline datasets. There are two exceptions [Zheng, Zhang, and Grover 2022; Lee et al. 2022] that fine-tune the models pre-trained from offline datasets. However, these two methods use the full pre-trained model instead of the representation, which limits the variety of applicable downstream tasks.

To effectively learn presentations that offer flexibility to downstreaming RL tasks, we propose RePreM (**R**epresentation **P**re-training with **M**asked Model) that pre-trains the encoder to predict the masked states or actions in

\*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the offline trajectories with the help of transformer blocks [Vaswani et al. 2017]. Specifically, we treat the trajectory  $(s_1, a_1, s_2, a_2, \dots)$  as a sequence and randomly mask a proportion of states or actions in the sequence. Then, we generate the representations of the states/actions in the sequence with the corresponding encoder and feed these representations to the transformer blocks. We train the encoders and the transformer blocks to predict the embedding of the masked states/actions with a contrastive loss. Predicting the masked states and actions with a uniform encoder-transformer architecture nicely corresponds to the role of the forward and inverse transition models designed in previous pretraining tasks (see, e.g., [Yang and Nachum 2021; Schwarzer et al. 2021]). When a specific downstream task is given, we build the corresponding networks (e.g., the policy or the value network) simply by appending several MLP layers to the fixed state encoder. In this way, we can achieve superior performance on various of tasks sample-efficiently.

RePreM is motivated by the repeated success achieved by similar architectures (i.e., transformer blocks combined with the masked model) in NLP [Devlin et al. 2018; Song et al. 2019] and CV [Carion et al. 2020; Dosovitskiy et al. 2020] for sequential information processing where long-range data dependency matters, a successful example of which is BERT [Devlin et al. 2018]. Compared to previous methods that require data augmentation or modeling the transition/reward process explicitly, RePreM is simple since it learns the state encoder with only transformer blocks and a self-supervision loss that predicts the masked elements. Furthermore, with the help of sequence modeling, the pre-trained encoder by RePreM can capture the long-term dynamics effectively.

We conduct extensive experiments on Atari games [Bellefleur et al. 2013] and DeepMind Control Suite (DMControl) [Tassa et al. 2018]. We show that our pre-trained state encoder enables sample-efficient learning on several downstream tasks including dynamic prediction, transfer learning, and sample-efficient RL. For dynamic prediction, our encoder results in a smaller prediction error than the baselines, especially for long-horizon predictions, which demonstrates that RePreM can capture the long-term dynamics effectively. For transfer learning, we pre-train the encoder on the data from a set of 24 Atari games and successfully transfer the representation to unseen games. For sample-efficient RL, we evaluate the pre-trained representation with the 100k benchmark (where the agent is only allowed to interact with the environment for 100k steps) proposed by Kaiser et al. [2019]. The results show that the representation learned by RePreM boosts Rainbow (a value-based method) [Hessel et al. 2018] and SAC (an actor-critic method) [Haarnoja et al. 2018], and achieves superior performance in Atari and DMControl respectively. Moreover, we show that the performance of RePreM scales well with the volume of the dataset, the quality of the dataset, and the size of the encoder. This indicates that this BERT-style pre-training method has great potential in learning representations to extract rich prior knowledge from big data. Finally, we conduct ablation studies on different designs of RePreM.

The contributions of our paper are summarized as follows:

- We propose a simple yet effective pre-training method

called RePreM (Representation Pre-training with Masked Model) that incorporates rich information and captures long-term dynamics for the representations in RL. To the best of our knowledge, we are the first to adopt BERT-style representation pre-training in RL.

- We conduct extensive experiments and show that RePreM can generate representations that enable sample-efficient learning in various downstream tasks. Specifically, we achieve state-of-the-art performance on the Atari-100k benchmark. Moreover, we show that its superior performance comes from its ability to capture long-term dynamics and it scales well with the dataset volume/quality and the encoder size.

## Related Work

**Representation learning in RL.** The studies on representation learning in RL can be broadly divided into two categories: representation learning along with online RL (e.g., [Schwarzer et al. 2020]) and pre-training representations based on offline data (e.g., [Schwarzer et al. 2021]). The objective of representation learning in RL is to generate generalizable representations that can aggregate similar states. However, some of the previous methods ignore the relationship between one state and other states or actions in the trajectory. Therefore, the resultant representations do not contain information about the environmental dynamics. Examples of this type include reconstruction-based methods [Jaderberg et al. 2016; Yarats et al. 2019] and data augmentation-based methods [Bachman, Hjelm, and Buchwalter 2019; Laskin, Srinivas, and Abbeel 2020; Stooke et al. 2021]. Other methods incorporate the dynamics to representations with bisimulation/other abstraction techniques [Castro 2020; Zhang et al. 2020; Liu et al. 2020], prediction-based approaches [Gelada et al. 2019] or Laplacian representations [Mahadevan and Maggioni 2007; Wu, Tucker, and Nachum 2018]. However, these methods only consider the one-step transition dynamics and therefore cannot incorporate the long-term dynamics into the representation. To capture long-term dynamics, recent methods adopt techniques such as predicting multiple steps in the future [Schwarzer et al. 2020], goal-conditioned RL [Schwarzer et al. 2021], recurrent architecture [Guo et al. 2020], or recurrent state space model (RSSM) [Hafner et al. 2019a,b, 2020]. Nevertheless, these methods are complex due to the reliance on multiple models and loss functions. In contrast, our method is simple (only with an additional transformer structure trained by the masked prediction task) and better captures the long-term dynamics.

**Sequence Modeling in RL.** Recently, motivated by the success of sequence modeling in NLP, researchers have begun to adopt it to RL. Decision Transformer [Chen et al. 2021] and Trajectory Transformer [Janner, Li, and Levine 2021] first use the transformer structure [Vaswani et al. 2017] to model the RL agent. However, most of the work focuses on offline RL [Reed et al. 2022; Furuta, Matsuo, and Gu 2021; Akhmetov et al. 2022; Xu et al. 2022] or online RL [Parisotto et al. 2020; Banino et al. 2021]. Two exceptions are Zheng, Zhang, and Grover [2022] and Lee et al. [2022]

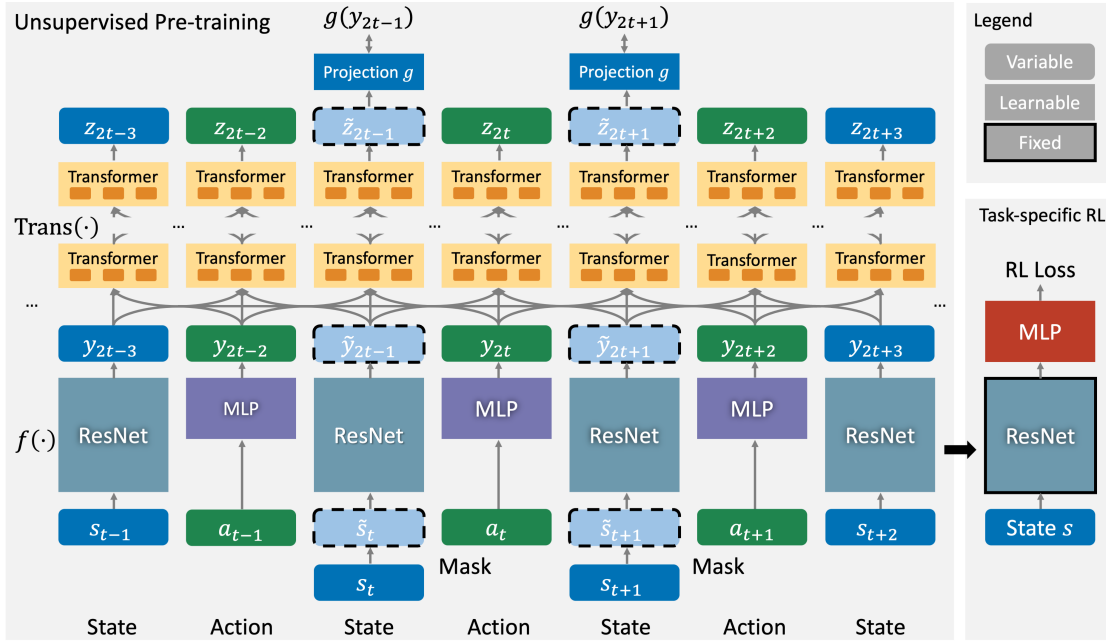


Figure 1: Unsupervised pre-training phase and the task-specific reinforcement learning phase of our method.

that consider the pre-training setting and fine-tune the pre-trained model as the agent/policy later. Different from these methods that use the pre-trained model, we aim to learn a pre-trained encoder, and this approach has two benefits: First, pre-trained representation is more portable for various downstream tasks. Second, since we only take the encoder from the pre-training phase, we can adopt simpler and more efficient bi-directional models instead of autoregressive ones.

## Methods

In this section, we will introduce how RePreM (Representation Pre-training with Masked Model) trains the state encoder in the pre-training phase and how we use the pre-trained state encoder in downstream RL tasks.

### Pre-training in RL

As in many previous papers [e.g., Eysenbach et al. 2019; Schwarzer et al. 2021], we consider pre-training from a reward-free dataset of trajectories that include only states and actions. Compared with pre-training with rewards [e.g., Gelada et al. 2019], this setting has two benefits: 1) reward-free datasets are more feasible in practice; and 2) the knowledge extracted from reward-free datasets can be used for different tasks in similar environments. The objective of this setting is to learn generalizable representations that are independent of the behavior policy (i.e., the policy used to collect the dataset) and the reward function. In this way, the representation can be trained with the trajectories collected by arbitrary behavior policies and used in tasks with different reward functions. We denote the offline dataset with  $N$  trajectories as  $\mathcal{D} = \{(s_1^i, a_1^i, s_2^i, a_2^i, \dots, s_T^i)\}_{i=1}^N$ , where  $T$

is the length of the trajectory.

Previously, various self-supervised learning tasks have been proposed to extract information from provided agent experiences. There are two categories of frequently used tasks: the forward prediction task in which the agent is trained to predict the immediate states [Guo et al. 2018; Gelada et al. 2019; Stooke et al. 2021; Schwarzer et al. 2020], and the backward prediction task where the agent is trained to predict the actions that can generate the specified transitions [Pathak et al. 2017; Shelhamer et al. 2016; Du et al. 2019]. Recent success in representation pre-training relies on the combination of different self-supervision tasks [Schwarzer et al. 2021]. Based on the previous self-supervision tasks, our objective is to design a simple and uniform self-supervision task that can realize both of these two types of tasks. Moreover, we hope that our self-supervision task can incentivize the representation to capture the long-term dynamics instead of only predicting the one-step dynamics.

Accordingly, we propose a self-supervision task that predicts the masked states or actions in the trajectory. We present our pre-training model in Figure 1 and introduce the pre-training model as follows:

**Mask.** Given a trajectory  $(s_1^i, a_1^i, s_2^i, a_2^i, \dots, s_T^i)$  in the offline dataset, we first randomly select an anchor item (i.e., state or action) in the sequence and mask several consecutive items of the same kind starting from this anchor item. Each item can be selected as the anchor with probability  $p$  and the number of consecutive masked items is sampled from  $\text{Unif}(n)$  (a discrete uniform distribution over  $1 \sim n$ ), where  $p \in \mathbb{R}$  and  $n \in \mathbb{Z}$  are hyperparameters. Taking the case in Figure 1 as an example, we select the state  $s_t$  and

mask 2 consecutive states starting from this state (i.e.,  $s_t$  and  $s_{t+1}$ ). Similar to BERT [Devlin et al. 2018], we replace the masked elements with 1) a fixed token 80% of the time; 2) an element of the same kind randomly sampled in the dataset 10% of the time; and 3) the unchanged element 10% of the time. Each trajectory typically contains thousands of steps, which can lead to a large amount of computational cost for the self-attention module. To reduce computational cost, we randomly cut the trajectory into several segments, each of which contains up to  $K$  states.

**Architecture.** We train the encoder  $f(\cdot)$  and the transformer blocks  $\text{Trans}(\cdot)$  to predict the masked elements. We maintain two different encoders for states and actions respectively. The states in our later experiments are image-based, and we use a ResNet [He et al. 2016] architecture as a state encoder. Notice that, compared with the encoder architecture used for image-based states in popular RL models (such as [Liu and Abbeel 2020; Stooke et al. 2021]) with only a few convolution layers, we use a larger encoder architecture. This is motivated by the recent observation that a larger network architecture can incorporate richer information and is more suitable for the pre-training setting than the online setting in RL [Schwarzer et al. 2021]. Next, we use a two-layer MLP as the action encoder to generate the action representation based on the raw action. For discrete actions, the raw actions are represented as one-hot vectors. We use an improved version of transformer blocks called GTrXL [Parisotto et al. 2020] that reorders the identity mapping and additionally adopts a gating layer based on the original transformer structure [Vaswani et al. 2017]. GTrXL improves the stability and learning speed of the original transformer structure and is shown to be suitable for RL. We use  $L$  GTrXL blocks in our model.

**Self-supervision task.** Different from the masked prediction task in BERT [Devlin et al. 2018] where the targets are discrete, our model should predict image-based states. Therefore, we use the transformer blocks to predict the embeddings of the image-based states or actions instead of the original ones. We design a contrastive loss based on SimCLR [Chen et al. 2020] for the self-supervision task as follows: Given a segment of  $K$  steps (corresponding to a sequence of length  $2K - 1$ ), we denote the index set of masked items as  $\mathcal{T}$ . Consider a masked item  $\tilde{x}_t$  with index  $t \in \mathcal{T}$  that is generated from the original item  $x_t$ . The item  $x$  can be the state  $s$  or the action  $a$ . We denote the embedding of the masked item (i.e., the output of the encoder) as  $\tilde{y}_t := f(\tilde{x}_t)$  for some  $t \in \mathcal{T}$  and the embedding of the original item as  $y_t := f(x_t)$  for some  $t \in [2T - 1]$ . We feed these embeddings to the GTrXL blocks, i.e.,  $z_1, \dots, \tilde{z}_t, \dots, z_K := T(y_1, \dots, \tilde{y}_t, \dots, y_K)$ , and obtain the predicted embedding for the masked item  $\tilde{z}_t$  for some  $t \in \mathcal{T}$ . Our aim is to maximize the similarity between the predicted embedding  $\tilde{z}_t$  and the target embedding  $y_t$  that is the representation of the original item  $x_t$ . To better calculate the similarity of two high-dimensional vectors, we use a projection function  $g(\cdot)$  to map the embeddings to lower-dimensional vectors. The loss

of our self-supervision task is defined as:

$$\mathcal{L} := -\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left[ \log \frac{\exp(\text{sim}(g(\tilde{z}_t), g(y_t)))}{\sum_{\tau \in [2K-1]} \mathbb{I}[\tau \neq t] \exp(\text{sim}(g(\tilde{z}_t), g(y_\tau)))} + \log \frac{\exp(\text{sim}(g(y_t), g(\tilde{z}_t)))}{\sum_{\tau \in [2K-1]} \mathbb{I}[\tau \neq t] \exp(\text{sim}(g(y_t), g(z_\tau)))} \right], \quad (1)$$

where the cosine similarity is defined as  $\text{sim}(z_1, z_2) := z_1^T z_2 / (\|z_1\| \|z_2\|)$ .

**Discussion.** This self-supervision task consists of two types of sub-tasks depending on the type of the masked items: the state prediction task and the action prediction task. When the model tries to predict the first several states in a sequence of masked states, it essentially predicts the future states based on the past states and actions. This is similar to the forward prediction task in previous papers. When the model tries to predict the masked action, this corresponds to the previous backward prediction task. Moreover, our self-supervision task emphasizes more on predicting based on the sequence instead of single states or actions, which is more suitable for the partially observable setting where the transition cannot be modeled accurately using the one-step transition.

## Downstream RL Tasks

In specific RL tasks, we only use the state encoder learned in the pre-training phase. Notice that the action encoder also incorporates useful prior knowledge, and it should be beneficial to utilize it in specific RL tasks, but we leave it as a future research direction. Broadly speaking, there are two approaches to utilizing the pre-trained encoder: fine-tuning and using as fixed features. In our experiments, we find that using fixed features works better for RL. This has also been validated in previous papers such as [Seo et al. 2022] in which the authors explain that the fine-tuning scheme can quickly erase the useful knowledge in the encoder. Therefore, we fix the state encoder and append several additional MLP layers as the value network, the policy network, the network for dynamic prediction, etc., according to different downstream tasks.

## Experiments

We design our experiments to answer the following questions:

- (Q1, Atari) Can RePreM improve the performance of value-based RL algorithms on the Atari-100k benchmark compared with previous algorithms?
- (Q2, Atari) How does the performance of RePreM change with the size of the dataset, the quality of the dataset, and the size of the state encoder?
- (Q3, Atari) Can the representations generated by RePreM be used for dynamic prediction and capture the long-term dynamics better than the previous algorithms?
- (Q4, Atari) Can the representations generated by RePreM be transferred to accelerate the learning for unseen tasks?

Methods	Dataset	Med.	Mean	>H	>0
<i>a1</i> Rainbow		-0.090	0.000	0	11
<i>a2</i> SimPLe		0.144	0.443	2	26
<i>a3</i> DER	–	0.161	0.285	2	26
<i>a4</i> DrQ		0.268	0.357	2	24
<i>a5</i> SPR		0.415	0.704	7	26
<i>b1</i> ATC	W-3M	0.219	0.587	4	26
<i>b2</i> ATC	M-3M	0.204	0.780	5	26
<i>b3</i> BC w/o FT	M-3M	0.139	0.227	0	23
<i>b4</i> BC	M-3M	0.548	0.858	8	26
<i>b5</i> SGI-M	W-5M	0.589	1.144	8	26
<i>b6</i> SGI-S	M-3M	0.423	0.914	8	26
<i>b7</i> SGI-M	M-3M	0.679	1.149	9	26
<i>b8</i> SGI-L	M-6M	0.753	1.598	9	26
<i>c1</i> RePreM-M	R-10M	0.190	0.289	4	26
<i>c2</i> RePreM-M	W-10M	0.774	1.335	9	26
<i>c3</i> RePreM-M	M-3M	0.717	1.298	9	26
<i>c4</i> RePreM-M	M-5M	0.788	1.436	9	26
<i>c5</i> RePreM-S	M-10M	0.835	1.384	10	26
<i>c6</i> RePreM-M	M-10M	0.953	1.970	13	26
<i>c7</i> RePreM-L	M-10M	<b>1.085</b>	<b>2.038</b>	15	26

Table 1: The performance of different algorithms on Atari-100K. The performance in group *a* and *b* are quoted from previous papers or calculated based on the scores provided in previous papers. The performance in group *c* are the mean across 10 random seeds and 100 evaluation trajectories.

- (Q5, DMControl) Can RePreM also improve the performance of actor-critic algorithms on the DMControl-100k benchmark compared with previous algorithms?
- (Q6, Ablation) What is the contribution of each proposed design in RePreM?

**Dataset.** We conduct the following experiments on Atari [Bellemare et al. 2013] and DMControl [Tassa et al. 2018]. For Atari, as in many previous papers (e.g., [Kaiser et al. 2019]), we select a set of 26 games. For pre-training on Atari games, we collect the following three types of datasets with different qualities for each game: The **Random** dataset is collected by executing uniformly randomly sampled actions at a number of consecutive steps sampled from a Geometric distribution with  $p = \frac{1}{3}$ . The **Weak** dataset is collected from the first 1M transitions generated by DQN. The **Mixed** dataset is obtained by concatenating multiple checkpoints evenly throughout the training of DQN. The quality of the dataset increases from **Random** to **Mixed**. We also evaluate the algorithms on datasets of different sizes. Larger datasets can be obtained by running the above procedure multiple times with different random seeds. For DMControl, we collect the offline dataset in a similar way to the procedure to collect the Mixed dataset using SAC [Haarnoja et al. 2018].

**Baselines algorithms.** In our experiments, we compare our algorithm with a wide range of previous algorithms including both sample-efficient RL algorithms and pre-training algorithms for RL.

For Atari games, we incorporate the representation pre-trained by RePreM with Rainbow [Hessel et al. 2018] in downstream tasks (except for dynamic prediction). The

baseline algorithms include not only sample-efficient on-line RL algorithms (such as Rainbow, SimPLe [Kaiser et al. 2019], data-efficient Rainbow/DER [van Hasselt, Hessel, and Aslanides 2019], DrQ [Yarats, Kostrikov, and Fergus 2020], and SPR [Schwarzer et al. 2020]) but also pre-training RL methods (such as ATC [Stooke et al. 2021], SGI [Schwarzer et al. 2021], and behavior cloning/BC which is shown to be a strong baseline in [Schwarzer et al. 2021]).

For DMControl, we incorporate the representation pre-trained by RePreM with SAC [Haarnoja et al. 2018] into downstream tasks. We compare our algorithm with SAC without pre-training and APV [Seo et al. 2022] which is a state-of-the-art pre-training method. This method beats many strong sample-efficient RL algorithms, including DrQ [Yarats, Kostrikov, and Fergus 2020] and Dreamer-v2 [Hafner et al. 2020].

## Q1: Performance of RePreM on Atari-100k

**Evaluation metrics.** Atari-100k is a popular benchmark for sample-efficient RL where the agent is allowed to interact with the environment for only 100k steps, equivalent to approximately two hours of human experience. We report the evaluation metrics following [Schwarzer et al. 2021]. The metrics include the mean and median human normalized score (HNS) averaging across 10 random seeds and over 100 evaluation trajectories at the end of the training process, the number of games in which the agent achieves super-human performance (>H) and greater-than-random performance (>0). HNS is calculated as  $\frac{\text{AgentScore} - \text{RandomScore}}{\text{HumanScore} - \text{RandomScore}}$  and we use the agent/human scores listed in [Badia et al. 2020].

**Experiment results.** To answer the first question, we run RePreM on the *Mixed* dataset with 3M samples, which is comparable to the datasets in previous pre-training methods. We present the result in Table 1. We can observe that RePreM (*c3* in Table 1) outperforms not only the previous learning-from-scratch sample efficient algorithms (*a1-a5* in Table 1) but also recently proposed pre-training methods (*b2-b4*, *b6-b7* in Table 1). We also note that the superior performance of RePreM is achieved using Rainbow (which is a popular model-free algorithm) as the base RL algorithm, but the performance of this base algorithm alone is the worst among all the listed algorithms (cf. *a1* in Table 1). This indicates that the representation pre-trained by RePreM can turn Rainbow into a strong algorithm. The results of this experiment indicate that the representation pre-trained by RePreM enables sample-efficient RL effectively.

## Q2: Scalability of RePreM

**Variants of RePreM.** To evaluate the scalability of RePreM in terms of the dataset size, the dataset quality, and the scale of the encoder, we obtain several variants of the representations pre-trained by RePreM. First, we use the datasets of 3M, 5M, and 10M samples, respectively, to evaluate the scalability w.r.t. the dataset size. Second, we use the **Random**, **Weak**, and **Mixed** datasets to evaluate the scalability w.r.t. the quality of the dataset. Third, we vary the size of the ResNet used in the encoder to evaluate the scalability w.r.t. the scale of the encoder, resulting in **Small**, **Medium**,

and Large variants.

**Experiment results.** To evaluate how performance changes with the quality of the dataset, we can compare the performance of RePreM on Random-10M, Weak-10M and Mixed-10M (see *c1*, *c2*, and *c6* in Table 1 respectively). On the basis of these results, we conclude that the performance of RePreM scales with the quality of the datasets. Moreover, although the performance worsens when the data quality degenerates, RePreM from Weak (see *c2*) still outperforms the previous pre-training methods that learn from Weak datasets (see *b1* and *b5*).

To evaluate how performance varies with the size of the dataset, we can compare the performance of RePreM learning from the Mixed dataset with 3M, 5M, and 10M samples (see *c3*, *c4*, and *c6* in Table 1 respectively). We can observe that the performance also scales with the dataset size, which indicates that our method can benefit from a larger dataset.

To evaluate how performance varies with the size of the state encoder, we can observe the performance of RePreM (-S/-M/-L) learned from Mixed-10M (see *c5-c7* in Table 1). In contrast to the online RL setting, where a large network structure does not lead to performance improvement, RePreM can benefit from a larger network structure.

Finally, combined with a large moderate-quality dataset (Mixed-10M) and a large state encoder, RePreM can achieve a performance that significantly outperforms the other algorithms and variants. The scaling trends of RePreM indicate that the performance of RePreM should further increase with the scale of parameters and datasets.

### Q3: RePreM for Dynamic Prediction

We evaluate the representations learned by different pre-training algorithms with the dynamic prediction task that predicts the future states based on the learned representation on Atari games. Specifically, we fix the state encoder pre-trained with different methods and append two MLP layers to the encoder. We train this network to predict the future RAM state  $\bar{s}_{t+k}$  (divided by 255) based on the current state  $s_t$ . We collect the training and testing samples using the same DQN policy and evaluate the prediction error using the mean-squared error between the predicted RAM state and the ground truth. We present the results obtained on Alien in Figure 2 and the results are similar across different games. We can observe that the representations generated by RePreM and SGI result in a more accurate prediction for the states in the near future than those of BC and ATC. This may be due to the fact that RePreM and SGI use forward prediction as one of the self-supervision tasks. Moreover, we find that, compared with SGI, the model based on RePreM can predict the states in the far future with higher accuracy. This indicates that RePreM generates representations that can better capture long-term information. This may explain the reason why RePreM outperforms the other representation pre-training algorithms in the Atari-100k benchmark.

### Q4: Transfer to Unseen Tasks

To evaluate how the representations pre-trained by RePreM can transfer the knowledge to unseen tasks, we select Alien




Figure 2: The prediction error versus different prediction horizons on Alien. We fix the pre-trained representations from different algorithms and append a trainable prediction network (i.e., a two-layer MLP) to predict the future RAM state. We train the prediction network with 10 random seeds and show the mean.




Figure 3: Comparison on the performance between the RePreM representation pre-trained on the same task and that pre-trained on the remaining 24 tasks. The shaded area indicates the standard deviation over 10 random seeds.

and Freeway as the target tasks and pre-train the representation using RePreM on a dataset collected on the 24 remaining Atari games with 12M samples (with 0.5M samples from the Mixed dataset for each game). The pre-trained representation is used in the same way as Atari-100k. We compare it with learning curves from the representation pre-trained on the homogeneous task (with 10M Mixed samples) and learning-from-scratch using the same network structure. We present the results in Figure 3. The result indicates that the representation pre-trained by RePreM can extract the knowledge transferable to other tasks.

### Q5: Performance of RePreM on DMC-100k

We first collect the offline dataset by SAC on the Walker and Quadruped environments from DMControl. Then, we pre-train the representation using RePreM and run SAC based on the pre-trained representation. The pre-trained representation is used as the encoder of both the critic and the actor in SAC. We compare RePreM with SAC from scratch and APF. For fair comparison, instead of using the original APF model pre-trained on different domains, we pre-train their model using the same dataset as ours and this improves the

Task	SAC	APV	RePreM
Walker-Walk	42±12	444±36	<b>633±52</b>
Walker-Run	67±27	496±74	<b>530±43</b>
Walker-Stand	142±21	742±81	<b>899±38</b>
Quadruped-Walk	127±184	420±87	<b>799±132</b>
Quadruped-Run	93±36	<b>698±96</b>	579±85

Table 2: The performance of different algorithms on Walker from DMControl-100K.

performance over the version provided in their paper.

We show the performance in Table 2. First, we can see that SAC based on the RePreM representation performs much better than learning from scratch, which indicates the effectiveness of the pre-trained representation. Second, we note that, different from RePreM that only uses the pre-trained representation, APF uses not only the pre-trained representation but also the latent transition model and the image decoder. Nevertheless, with the representation pre-trained by RePreM, SAC achieves comparable performance with this state-of-the-art pre-training method. In addition, note that we use the same RePreM representation for these tasks, and RePreM+SAC achieves good performance. This indicates that the RePreM representation can work for downstream tasks with different reward functions.

## Q6: Ablation Studies

To evaluate the contribution of each proposed design, we implement the following ablated versions on top of the RePreM version in *c3* of Table 1. The results are presented in Table 3.

- *RePreM+FT* that fine-tunes the network in the downstream tasks instead of fixing the pre-trained state encoder. We can see that fine-tuning in downstream tasks degenerates the performance. This result is consistent with [Seo et al. 2022] who finds that fine-tuning may quickly lose the pre-trained representation.
- *RePreM+short* that models shorter sequences with  $T = 10$  (instead of  $T = 50$  in RePreM). The experiment result shows that using a shorter sequence length significantly degenerates the performance. This may indicate that capturing the long-term dynamics with a large  $T$  is important to the success of RePreM.
- *RePreM-BERT* that replaces all masked elements with a fixed token instead of following the masking scheme in BERT. The result indicates that following the masking scheme in BERT also contributes to the success of RePreM.
- *RePreM-GTrXL* that uses the original transformer block instead of GTrXL. We can see that using GTrXL can improve performance, possibly because of its stability.
- *RePreM+ReLIC* that uses the ReLIC [Mitrovic et al. 2020] loss as the contrastive self-supervision loss instead of the SimCLR loss for the masked prediction (as in CoBERL [Banino et al. 2021]). We can see that the performance of the simple SimCLR loss is similar to that of the ReLIC loss, but ReLIC is more complex. For simplicity, we use SimCLR in our proposed method.

Ablated Versions	Median	Mean	>H	>0
RePreM+FT	0.512	0.839	7	26
RePreM+short	0.322	0.806	5	26
RePreM-BERT	0.588	1.166	8	26
RePreM-GTrXL	0.709	1.220	9	26
RePreM+ReLIC	0.715	1.201	9	26
RePreM+decoder	0.107	0.210	0	24
RePreM ( <i>c3</i> in Table 1)	0.717	1.298	9	26

Table 3: The performance of different ablated variants on Atari-100K pre-trained using the Mixed dataset with 3M samples. These ablated versions use the medium-sized ResNet as the state encoder.

- *RePreM+decoder* that predicts the raw states/actions with state/action decoders instead of predicting their embeddings using the SimCLR loss (as in [Jaderberg et al. 2016; Seo et al. 2022]). We can see that using the encoder to predict the raw states/actions underperforms the proposed RePreM that predicts the embeddings significantly. The reason may be that the raw states are high-dimensional, and reconstructing the raw states may mislead the encoder and the transformer to learn spurious details.

## Conclusion

In this paper, we propose a representation pre-training method called RePreM (Representation Pre-training with Masked Model) for RL. RePreM is algorithmically simple but effective, through learning the representation with only a transformer-based sequence model and masked prediction. Although RePreM adopts sequence modeling that is different from the previous representation pre-training baselines for RL, it performs on a par with or better than previous RL pre-training methods. Empirically, we find that the success of RePreM can be attributed to its strong ability to capture long-term dynamics in the trajectories. This is a benefit of modeling the whole sequence instead of one-step dynamics as in previous methods. Through extensive ablation studies, we show that the representation pre-trained using RePreM can be used for various downstream tasks. We also find that RePreM not only scales well with the dataset size, the dataset quality, and the size of the encoder, but can also generalize across different tasks. This indicates its potential to extract knowledge from big data and use the knowledge for efficient and robust learning. Currently, RePreM only uses the state encoder learned in the pre-training phase but discards the transformer blocks and action encoder that also contain useful knowledge. Thus, how to utilize the knowledge in the transformer blocks and action encoder will be an interesting future research direction.

## Acknowledgements

The work of Yuanying Cai and Longbo Huang is supported by the Technology and Innovation Major Project of the Ministry of Science and Technology of China under Grant 2020AAA0108400 and 2020AAA0108403, the Tsinghua

University Initiative Scientific Research Program, and Tsinghua Precision Medicine Foundation 10001020109. Dr. Xuyun Zhang is the recipient of an ARC DECRA (project No. DE210101458) funded by the Australian Government.

## References

- Akhmetov, Y.; Nurmanova, V.; Bagheri, M.; Zollanvari, A.; and Gharehpetian, G. B. 2022. A Bootstrapping Solution for Effective Interpretation of Transformer Winding Frequency Response. *IEEE Transactions on Instrumentation and Measurement*, 71: 1–11.
- Arulkumaran, K.; Cully, A.; and Togelius, J. 2019. AlphaStar: An evolutionary computation perspective. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 314–315.
- Bachman, P.; Hjelm, R. D.; and Buchwalter, W. 2019. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32.
- Badia, A. P.; Piot, B.; Kapturowski, S.; Sprechmann, P.; Vitvitskiy, A.; Guo, Z. D.; and Blundell, C. 2020. Agent57: Outperforming the atari human benchmark. In *International Conference on Machine Learning*, 507–517. PMLR.
- Banino, A.; Badia, A. P.; Walker, J.; Scholtes, T.; Mitrovic, J.; and Blundell, C. 2021. Coberl: Contrastive bert for reinforcement learning. *arXiv preprint arXiv:2107.05431*.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. In *Proceedings of Advances in Neural Information Processing Systems*, volume 33, 1877–1901.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Castro, P. S. 2020. Scalable methods for computing state similarity in deterministic markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10069–10076.
- Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Deghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, S.; Krishnamurthy, A.; Jiang, N.; Agarwal, A.; Dudik, M.; and Langford, J. 2019. Provably efficient RL with rich observations via latent state decoding. In *Proceedings of the 36th International Conference on Machine Learning*, 1665–1674. PMLR.
- Dubey, R.; Agrawal, P.; Pathak, D.; Griffiths, T. L.; and Efros, A. A. 2018. Investigating human priors for playing video games. *arXiv preprint arXiv:1802.10217*.
- Emmons, S.; Eysenbach, B.; Kostrikov, I.; and Levine, S. 2021. RvS: What is Essential for Offline RL via Supervised Learning? *arXiv preprint arXiv:2112.10751*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2019. Diversity is all you need: Learning skills without a reward function. In *Proceedings of the 7th International Conference on Learning Representation*.
- Furuta, H.; Matsuo, Y.; and Gu, S. S. 2021. Generalized decision transformer for offline hindsight information matching. *arXiv preprint arXiv:2111.10364*.
- Gelada, C.; Kumar, S.; Buckman, J.; Nachum, O.; and Bellemare, M. G. 2019. Deepmdp: Learning continuous latent space models for representation learning. In *International Conference on Machine Learning*, 2170–2179. PMLR.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33: 21271–21284.
- Guo, Z. D.; Azar, M. G.; Piot, B.; Pires, B. A.; and Munos, R. 2018. Neural predictive belief representations. *arXiv preprint arXiv:1811.06407*.
- Guo, Z. D.; Pires, B. A.; Piot, B.; Grill, J.-B.; Altché, F.; Munos, R.; and Azar, M. G. 2020. Bootstrap latent-predictive representations for multitask reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*, 3875–3886. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning*, 1861–1870. PMLR.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2019a. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019b. Learning latent dynamics for planning from pixels. In *Proceedings of the 36th International Conference on Machine Learning*, 2555–2565. PMLR.
- Hafner, D.; Lillicrap, T.; Norouzi, M.; and Ba, J. 2020. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Henaff, O. 2020. Data-efficient image recognition with contrastive predictive coding. In *International Conference on Machine Learning*, 4182–4192. PMLR.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 3215–3222.
- Hutsebaut-Buysse, M.; Mets, K.; and Latré, S. 2020. Pre-trained word embeddings for goal-conditional transfer learning in reinforcement learning. *arXiv preprint arXiv:2007.05196*.

- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2016. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*.
- Janner, M.; Li, Q.; and Levine, S. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286.
- Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873): 583–589.
- Kaiser, L.; Babaeizadeh, M.; Milos, P.; Osinski, B.; Campbell, R. H.; Czechowski, K.; Erhan, D.; Finn, C.; Kozakowski, P.; Levine, S.; et al. 2019. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and brain sciences*, 40.
- Laskin, M.; Srinivas, A.; and Abbeel, P. 2020. CURL: Contrastive unsupervised representations for reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*, 5639–5650. PMLR.
- Lee, K.-H.; Nachum, O.; Yang, M.; Lee, L.; Freeman, D.; Xu, W.; Guadarrama, S.; Fischer, I.; Jang, E.; Michalewski, H.; et al. 2022. Multi-Game Decision Transformers. *arXiv preprint arXiv:2205.15241*.
- Liu, G.; Zhang, C.; Zhao, L.; Qin, T.; Zhu, J.; Jian, L.; Yu, N.; and Liu, T.-Y. 2020. Return-Based Contrastive Representation Learning for Reinforcement Learning. In *International Conference on Learning Representations*.
- Liu, H.; and Abbeel, P. 2020. Unsupervised Active Pre-Training for Reinforcement Learning. In *International conference on Learning Representations*.
- Mahadevan, S.; and Maggioni, M. 2007. Proto-value Functions: A Laplacian Framework for Learning Representation and Control in Markov Decision Processes. *Journal of Machine Learning Research*, 8(10).
- Mitrovic, J.; McWilliams, B.; Walker, J.; Buesing, L.; and Blundell, C. 2020. Representation learning via invariant causal mechanisms. *arXiv preprint arXiv:2010.07922*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemaire, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529.
- Parisotto, E.; Song, F.; Rae, J.; Pascanu, R.; Gulcehre, C.; Jayakumar, S.; Jaderberg, M.; Kaufman, R. L.; Clark, A.; Noury, S.; et al. 2020. Stabilizing transformers for reinforcement learning. In *International Conference on Machine Learning*, 7487–7498. PMLR.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 16–17.
- Reed, S.; Zolna, K.; Parisotto, E.; Colmenarejo, S. G.; Novikov, A.; Barth-Maron, G.; Gimenez, M.; Sulsky, Y.; Kay, J.; Sprinzenberg, J. T.; et al. 2022. A generalist agent. *arXiv preprint arXiv:2205.06175*.
- Schwarzer, M.; Anand, A.; Goel, R.; Hjelm, R. D.; Courville, A.; and Bachman, P. 2020. Data-efficient reinforcement learning with self-predictive representations. *arXiv preprint arXiv:2007.05929*.
- Schwarzer, M.; Rajkumar, N.; Noukhovitch, M.; Anand, A.; Charlin, L.; Hjelm, R. D.; Bachman, P.; and Courville, A. C. 2021. Pre-training representations for data-efficient reinforcement learning. *Advances in Neural Information Processing Systems*, 34.
- Seo, Y.; Lee, K.; James, S.; and Abbeel, P. 2022. Reinforcement Learning with Action-Free Pre-Training from Videos. *arXiv preprint arXiv:2203.13880*.
- Shelhamer, E.; Mahmoudieh, P.; Argus, M.; and Darrell, T. 2016. Loss is its own reward: Self-supervision for reinforcement learning. *arXiv preprint arXiv:1612.07307*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019. MASS: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- Stooke, A.; Lee, K.; Abbeel, P.; and Laskin, M. 2021. Decoupling representation learning from reinforcement learning. In *Proceedings of the 38th International Conference on Machine Learning*, 9870–9879. PMLR.
- Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; Casas, D. d. L.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; et al. 2018. DeepMind control suite. *arXiv preprint arXiv:1801.00690*.
- Tsividis, P. A.; Pouncey, T.; Xu, J. L.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Human learning in Atari. In *2017 AAAI spring symposium series*.
- van Hasselt, H. P.; Hessel, M.; and Aslanides, J. 2019. When to use parametric models in reinforcement learning? *Advances in Neural Information Processing Systems*, 32.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proceedings of the International Conference on Neural Information Processing Systems*, 5998–6008.
- Wu, Y.; Tucker, G.; and Nachum, O. 2018. The laplacian in rl: Learning representations with efficient approximations. *arXiv preprint arXiv:1810.04586*.
- Xu, M.; Shen, Y.; Zhang, S.; Lu, Y.; Zhao, D.; Tenenbaum, J.; and Gan, C. 2022. Prompting Decision Transformer for Few-Shot Policy Generalization. In *International Conference on Machine Learning*, 24631–24645. PMLR.
- Yang, M.; and Nachum, O. 2021. Representation matters: offline pretraining for sequential decision making. In *International Conference on Machine Learning*, 11784–11794. PMLR.
- Yarats, D.; Kostrikov, I.; and Fergus, R. 2020. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*.
- Yarats, D.; Zhang, A.; Kostrikov, I.; Amos, B.; Pineau, J.; and Fergus, R. 2019. Improving sample efficiency in model-free reinforcement learning from images. *arXiv preprint arXiv:1910.01741*.
- Zhang, A.; McAllister, R.; Calandra, R.; Gal, Y.; and Levine, S. 2020. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*.
- Zhang, C.; Vinyals, O.; Munos, R.; and Bengio, S. 2018. A study on overfitting in deep reinforcement learning. *arXiv preprint arXiv:1804.06893*.
- Zheng, Q.; Zhang, A.; and Grover, A. 2022. Online decision transformer. *arXiv preprint arXiv:2202.05607*.