

MOOSS: Mask-Enhanced Temporal Contrastive Learning for Smooth State Evolution in Visual Reinforcement Learning

Jiarui Sun, M. Ugur Akcal, Girish Chowdhary
 University of Illinois Urbana-Champaign
 Urbana, IL, USA
 {jsun57, makcal2, girishc}@illinois.edu

Wei Zhang
 Visa Research
 Foster City, CA, USA
 wzhan@visa.com

Abstract

In visual Reinforcement Learning (RL), learning from pixel-based observations poses significant challenges on sample efficiency, primarily due to the complexity of extracting informative state representations from high-dimensional data. Previous methods such as contrastive-based approaches have made strides in improving sample efficiency but fall short in modeling the nuanced evolution of states. To address this, we introduce MOOSS, a novel framework that leverages a temporal contrastive objective with the help of graph-based spatial-temporal masking to explicitly model state evolution in visual RL. Specifically, we propose a self-supervised dual-component strategy that integrates (1) a graph construction of pixel-based observations for spatial-temporal masking, coupled with (2) a multi-level contrastive learning mechanism that enriches state representations by emphasizing temporal continuity and change of states. MOOSS advances the understanding of state dynamics by disrupting and learning from spatial-temporal correlations, which facilitates policy learning. Our comprehensive evaluation on multiple continuous and discrete control benchmarks shows that MOOSS outperforms previous state-of-the-art visual RL methods in terms of sample efficiency, demonstrating the effectiveness of our method.

1. Introduction

Visual Reinforcement Learning (RL), *i.e.*, an RL agent learning from visual signals composed of sequences of image-based observations, has long been a significant challenge. Compared to RL that utilizes compact state-based features, Visual RL is notably *sample inefficient*: it requires more environment interactions for a visual RL agent to achieve a comparable performance to its state-based counterparts [51]. This inefficiency primarily stems from the complexity in extracting informative states from high-dimensional visual data (pixels). Despite this, visual RL's




Figure 1. t-SNE [53] visualization of the state representations from a trained visual RL agent on the *reacher-easy* task from DeepMind Control Suite [51]. The state representations are encoded from an observation sequence $\mathbf{o}_{0:19}$ of length 20, guided by random actions. Numbers within the color-coded dots denote the temporal indices. Note that the t-SNE visualization demonstrates a temporal order, suggesting a gradual, smooth evolution of the states.

ability to function without handcrafted features offers broad applicability and a close resemblance to natural learning processes. Therefore, the ability to efficiently learn effective state representations is crucial.

To this end, many approaches improve sample efficiency of visual RL agents through incorporating auxiliary tasks tailored to benefit the learning of informative state representations. These auxiliary tasks often rely on *self-supervision* signals, which are derived from trajectory roll-outs obtained from agent-environment interactions. Examples of these tasks include learning forward [43] or backward [42] predictive features, predicting rewards [46], and applying bisimulation metrics [64]. Among numerous ways to facilitate state representation learning, *contrastive-based* approaches have emerged as a prominent framework, focusing on maximizing agreement between different views of a state. For example, CURL [31] generates positive samples of state through image augmentation techniques; subsequent works such as ATC [47] treat encoded observations

separated by a short temporal difference as positive samples, introducing the temporal concept to the contrastive objective. On the other hand, methods involving masked reconstruction, such as MLR [63], which perform reconstruction from corrupted observations, are less common yet offer unique insights. These auxiliary objectives have shown great improvements in sample efficiency for visual RL.

However, the effectiveness of current methods is limited by their inadequate consideration of *state evolution*. Specifically, if we consider observations or states within adjacent timesteps, as exemplified in Fig. 1, it becomes apparent that they typically exhibit stronger temporal correlations, *i.e.*, more “similar”, due to their inherent causal relationships, as opposed to those further apart. This suggests that state embeddings, encoded from raw observations, are likely to evolve temporally in a gradual and smooth manner, with abrupt changes being less probable. However, existing contrastive methods only consider a *binary distinction* between positive and negative samples, overlooking the gradual evolutionary nature of states. In addition, unlike video models [9] that can process multiple frames simultaneously to capture temporal evolution, RL’s formulation constrains the observation encoder to map *one* observation to *one* state independently. This makes temporal modeling even harder. On the other hand, approaches within the masked reconstruction domain often adopt a uniform masking approach, overlooking the high spatial-temporal correlation of consecutive pixel-based observations. We argue that such reconstruction task does not sufficiently challenge the model to understand the underlying dynamics of the observations, making the learned state representations less informative. These limitations in both contrastive and masked reconstruction methods – the former’s binary view of sample relationships and the latter’s oversight of spatial-temporal nuances – impede a deeper understanding of state dynamics, which is essential for progress in efficiency of visual RL.

To address the above limitations, we propose to explicitly model the state evolution for efficient state representation learning via self-supervision. Our approach, MOOSS, **M**ask-enhanced **tempORal cO**ntrastive learning for **S**mooth **S**tate evolution, explores the potential of combining contrastive learning with spatial-temporal mask modeling. Specifically, as shown in Fig. 2, MOOSS integrates an auxiliary temporal contrastive objective into visual RL agents, which is jointly trained with the main RL objective. This contrastive objective goes beyond the conventional binary distinction by modeling state similarities at *multiple levels*. This allows us to encourage the model to focus on gradual and evolving state changes over various temporal distances. Alongside this, we envision pixel-based observations as a *spatial-temporal graph*, applying a random walk-based masking technique. This presents a complex pre-text task, posing greater challenges than those presented by

standard uniform block-based masking [63], thereby compelling the RL agent to acquire a deeper understanding of observations with deliberately disrupted spatial-temporal connections. By combining these approaches, MOOSS applies the temporal contrastive objective to embeddings from both masked and unmasked observations. This unified strategy enhances the model’s ability to efficiently capture the dynamics of the observations by encouraging the agent to focus on evolving elements, thus facilitating informative state learning and improve policy learning.

Our main contributions are summarized as follows. (1) We propose a novel, auxiliary temporal contrastive objective tailored to visual RL, aimed at emphasizing the temporal continuity and change of states derived from pixel-based observations. (2) We re-cast pixel-based observations as a spatial-temporal graph, employing random walk-based masking to generate contrastive samples with disrupted spatial-temporal correlations. (3) Combining temporal contrastive objective with spatial-temporal masking, we introduce MOOSS. MOOSS is proven effective for improving the sample efficiency of visual RL algorithms across multiple continuous and discrete control benchmarks, including the DeepMind Control Suite [51] and Atari games [4], outperforming previous state of the art. Our detailed ablation studies further validate the efficacy of our method.

2. Related Work

2.1. Representation Learning for Visual RL

Efficiently learning informative state representations from pixel-based observations is a challenging problem for RL. Unlike the abundance of data in supervised settings, RL relies on experience trajectories collected through costly agent-environment interactions. This makes robust observation encoding from limited samples a complex task. As such, sample efficiency has emerged as a critical focus area for visual RL, with various approaches being developed to address this problem. Some methods involve learning world models [14, 15, 26, 41, 45], where the aim is to construct an internal representation of the environment that aids policy learning. Few other works [17, 23, 29, 30, 37] emphasize enhancing observation diversity through data augmentation techniques. Through enriching training samples, these methods acquire observation encoders that are more robust and generalizable, thereby alleviating the efficiency issue. Facilitated by data augmentation, one major line of work involves leveraging self-supervised auxiliary objectives that are optimized jointly with policy learning objectives. Notable examples include learning forward or backward predictive features [11, 12, 32, 43, 46, 62], and state reconstruction [60, 63, 66]. Within state reconstruction methods, MLR [63] stands out by performing latent reconstruction from corrupted pixels, marking an early exploration of

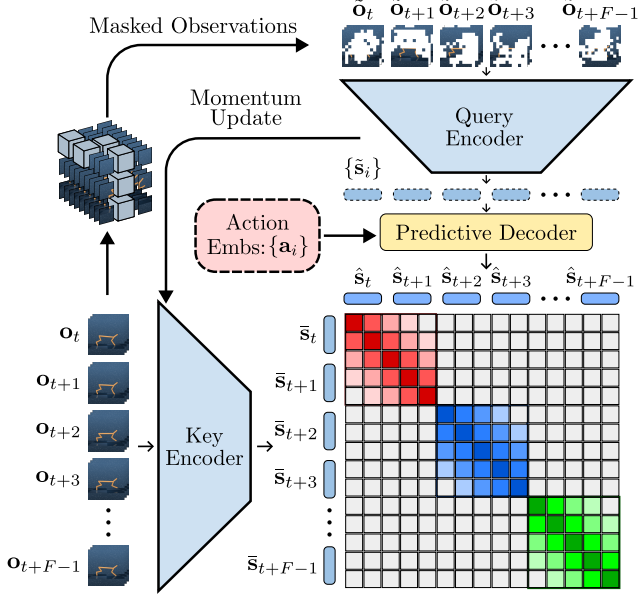


Figure 2. The proposed MOOSS framework. We first perform graph-based spatial-temporal masking on the observation sequence $\mathbf{o}_{t:t+F-1}$. The masked observations are then fed into a query encoder, generating $\tilde{\mathbf{s}}_i$ s. The unmasked observations are processed by a momentum key encoder. The key encoder generates the *key state embeddings* $\bar{\mathbf{s}}_{t:t+F-1}$. A predictive decoder is used to further process the outputs $\tilde{\mathbf{s}}_i$ s of the query encoder, generating the *query state embeddings* $\hat{\mathbf{s}}_{t:t+F-1}$ conditioned on the corresponding action embeddings \mathbf{a}_i s (Embs).

mask-based modeling in visual RL.

Among these auxiliary tasks, contrastive discrimination [2, 31, 34, 38, 40, 47, 65] has emerged as a prominent technique for enhancing state representation learning. The seminal work CURL [31] focuses on maximizing agreement between augmented versions of the same observation. Subsequent works integrate temporal elements into their contrastive objectives. ATC [47] and ST-DIM [2] treat temporally close neighbors as positive samples to emphasize temporal proximity, whereas DRIML [38] and TACO [65] focus on aligning predicted future states with their groundtruth counterparts. In addition to this joint learning scheme, another major direction of research aims to acquire robust, informative state representations from pre-trained encoders before policy learning [35, 36, 44, 57] as a separate stage. Our approach, MOOSS, falls in the auxiliary joint learning framework, explores the potential of combining contrastive learning with mask modeling to explicitly model state evolution.

2.2. Contrastive Learning and Masked Modeling

Contrastive learning, a self-supervised representation learning approach, has gained significant attention and been applied in various fields such as computer vision [7, 19]

and graph learning [59, 61]. The most prominent objective in contrastive learning is the InfoNCE loss [40], designed to maximize the mutual information between positive samples. Formally, given a query q and a key set \mathcal{K} containing its positive key k^+ , the objective \mathcal{L}_q is to ensure that q aligns more closely with k^+ than with other keys in \mathcal{K} :

$$\mathcal{L}_q = -\mathbb{E} \left[\log \frac{\exp(\text{sim}(q, k^+)/\tau)}{\sum_{k \in \mathcal{K}} \exp(\text{sim}(q, k)/\tau)} \right], \quad (1)$$

where $\text{sim}(\cdot)$ measures the similarity of the sample pair, and τ is the temperature parameter. In visual RL, this similarity is typically calculated through a bilinear product [31, 47, 65].

However, despite various principles are used to form the positive pair (q, k^+) , the contrastive objective focuses only one unique positive pair for each query state. This approach, while effective, adheres to a binary distinction, categorizing interactions solely as positives or negatives. Some works from other fields aim to broaden this perspective by allowing multiple positive samples for one query. Approaches such as MIL-NCE [39] and CoCLR [16] incorporate multiple positive keys to one query into their contrastive loss to learn video representations. RINCE [21] further extends the binary distinction by preserving a ranked ordering of positive samples, showing effectiveness in supervised classification task with additional superclass labels and unsupervised video representation learning. Inspired by RINCE, MOOSS is the first visual RL approach using a multi-level temporal contrastive objective to model state evolution.

Masked modeling, with roots dating back to [56], has recently gained prominence in language [8, 49], vision [3, 18], and graph [22, 50] domains. Its effectiveness in training models through self-supervised reconstruction has made it a preferred choice for many studies. While reconstruction has proven to be a powerful pretext task, masking techniques vary significantly among domains. Language models typically perform masking at the token level, obscuring specific words or phrases to encourage the model to predict the missing information based on context. Image models often employ patch masking [10, 18] due to the heavy spatial redundancy of images, while some video models utilize techniques such as tube masking [52, 58] to incorporate the temporal dimension. For graph learners, strategies range from uniform [22] to path-based [33, 48] masking. In our work, we explore the application of graph masking principles to image-based observation sequences in visual RL. Through experiments, we demonstrate that this creates a challenging pretext task, compelling MOOSS to develop a deep understanding of state dynamics and enhancing its ability to interpret complex spatial-temporal patterns of visual data.

3. Preliminaries

The learning process of Visual RL corresponds to a Partially Observable Markov Decision Process (POMDP)




Figure 3. Illustration of our graph-based spatial-temporal masking. The observation sequence η_o with shape $F \times H \times W$ is equally divided into non-overlapping cubes with shape $f \times h \times w$, constructing a spatial-temporal graph \mathcal{G} with adjacent nodes connected. Masking is applied by simulating a random walk on the constructed graph.

[5, 25]: $(\mathcal{O}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{O}, \mathcal{A}, P, R, \gamma$ denote the observation space, the action space, the transition dynamics $\mathcal{O} \times \mathcal{A} \rightarrow \Delta(\mathcal{O})$, the reward function $\mathcal{O} \times \mathcal{A} \rightarrow \mathbb{R}$, and the discount factor, respectively. $\Delta(\mathcal{O})$ is the space of probability distributions over \mathcal{O} , and the reward function at time step t can be written as $r_t = R(\mathbf{o}_t, a_t)$, where a_t is the t^{th} action. For visual RL, each observation $\mathbf{o}_t \in \mathbb{R}^{c \times H \times W}$ consists of c two-dimensional pixel-based feature maps. The objective of the RL agent is to learn a policy $\pi(a_t | \mathbf{o}_t)$ which maximizes the discounted cumulative reward $\mathbb{E}_\pi \sum_{t=0}^{\infty} \gamma^t r_t$, where $\gamma \in [0, 1)$.

4. Methodology

As a method designed for efficient state representation learning in visual RL, MOOSS can be seamlessly integrated with any existing RL algorithms, such as SAC [13] or Rainbow [20]. This integration is achieved by combining policy updates from the chosen RL algorithm with MOOSS’s auxiliary contrastive loss updates. The core idea of MOOSS is to explicitly model state evolution through (1) graph-based spatial-temporal masking on pixel-based observations for contrastive sample generation, and (2) a carefully designed multi-level temporal contrastive objective with the help of the masking approach. In the following subsections, we first present MOOSS’s overall framework, then introduce the proposed masking module with related architectural designs in detail. We then delve into the specifics of the temporal contrastive objective.

4.1. Overall Framework

The MOOSS framework, illustrated in Fig. 2, begins by constructing a spatial-temporal graph from the raw, pixel-based observations. On this graph, a masking operation

is performed. The graph’s masked observations, alongside their unmasked counterparts, are then fed into an observation query encoder and a momentum key encoder, respectively, to produce state embeddings. The masked state embeddings are then passed to a predictive decoder to generate *query* states, while the unmasked observations are used to form *key* states. Finally, the temporal contrastive objective is applied to these query and key state representations, with the aim of modeling the evolution of states over time.

4.2. Graph-based Masking for State Generation

Spatial-Temporal Masking. We perform graph-based spatial-temporal masking to obtain masked observation sequences which are used to generate the query embeddings. The masking process is illustrated in Fig. 3. Let $\eta_o := \{\mathbf{o}_i\}_{i=t}^{t+F-1}$ denote a sequence of observations with F timesteps sampled from the replay buffer. We first stack all observations in η_o as a cuboid of shape $F \times H \times W$.¹ Then, we equally divide the cuboid into non-overlapping cubes with the shape of $f \times h \times w$, where each cube can be thought of as a node on a graph. For two such nodes that are adjacent to each other, *i.e.*, two cubes that are spatial-temporally consecutive, we form an edge in between. As such, we construct a spatial-temporal graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ from the observation sequence. \mathcal{G} contains $\frac{FHW}{fhw}$ nodes by construction.

We then randomly mask a portion of the nodes from \mathcal{G} to obtain a masked observation sequence $\tilde{\eta}_o := \{\tilde{\mathbf{o}}_i\}_{i=t}^{t+F-1}$. Instead of uniformly masking image patches as in previous works [63], we propose to use random walk-based masking on the constructed graph \mathcal{G} . Formally, the set of masked nodes $\mathcal{V}_{\text{mask}}$ with size $|\mathcal{V}| \cdot p_m$ are collected from a sampled random walk $\mathcal{E}_{\text{mask}}$ as:

$$\mathcal{E}_{\text{mask}} \sim \text{RandomWalk}(\mathcal{E}, r), \quad (2)$$

where p_m is the masking ratio, and $r \in \mathcal{V}$ is the root node to start the walk. Then, all cubes corresponding to nodes in $\mathcal{V}_{\text{mask}}$ are masked by setting the corresponding patches to zero to form $\tilde{\eta}_o$. Compared to uniform patch-based masking, our graph-based spatial-temporal masking can more effectively break short-range consecutive information chunks. As the information density of image-based observation sequences is relatively low due to the spatial-temporal redundancy of visual data, our method creates a more challenging pretext task for the subsequent modules to solve.

Observation Encoding. Inspired by works in self-supervised image representation learning [12, 19], two observation encoders are used to generate state embeddings from (1) the masked and (2) the original observations, respectively. The encoders are Convolutional Neural Network

¹Here we omit the feature dimension c for notation simplicity.

(CNN)-based, and their architectural design are taken from previous works [51, 60]. First, one encoder $f_\theta(\cdot)$ is used to process $\tilde{\eta}_o$, which generates a sequence of masked state embeddings $\tilde{\eta}_s := \{\tilde{s}_i\}_{i=t}^{t+F-1}$, $\tilde{s}_i \in \mathbb{R}^d$. The parameters of $f_\theta(\cdot)$ are optimized in an end-to-end manner. At the same time, another momentum observation encoder $f_{\bar{\theta}}(\cdot)$ is used to encode the original observations η_o to produce the *key state embeddings* η_k :

$$\eta_k := \{\bar{s}_i\}_{i=t}^{t+F-1} = f_{\bar{\theta}}(\eta_o). \quad (3)$$

This second encoder $f_{\bar{\theta}}(\cdot)$ shares the same architecture as $f_\theta(\cdot)$, and its parameters $\bar{\theta}$ are updated by an Exponential Moving Average (EMA) of θ with the momentum coefficient $m \in [0, 1)$ as $\bar{\theta} \leftarrow m\bar{\theta} + (1 - m)\theta$.

Predictive Decoding. RL naturally operates sequentially: an agent’s current state is determined by its past states and actions. Thus, the actions stored in the trajectory roll-outs provide crucial guidance in state evolution. Considering this, we utilize both states and actions as the inputs to a causal Transformer-based predictive decoder for query state generation, reducing possible ambiguities to facilitate the subsequently described temporal contrastive objective. Formally, the decoder $g_\phi(\cdot)$ takes as inputs of the masked state embeddings $\tilde{\eta}_s$ and the actions $\{a_i\}_{i=t}^{t+F-1}$, both of which can be taken from the replay buffer. The actions are firstly embedded as d -dimensional tokens $\{\mathbf{a}_i\}_{i=t}^{t+F-1}$ with linear layers. Then, state and action embeddings are summed with positional encodings [55] to obtain positional information, and ordered alternatively to form a state-action sequence:

$$\tilde{\eta}_{s,a} := \text{Flat}(\{\tilde{s}_i, \mathbf{a}_i\}_{i=t}^{t+F-1}) + \text{Flat}(\{\mathbf{p}_i, \mathbf{p}_i\}_{i=t}^{t+F-1}), \quad (4)$$

where $\tilde{\eta}_{s,a} \in \mathbb{R}^{2S \times d}$ is the input to the Transformer layers, $\mathbf{p}_i \in \mathbb{R}^d$ is the i^{th} positional encoding, and Flat. denotes the flatten operation. Then, we gather outputs at the state indices from the Transformer layers, and use a Multi-Layer Perceptron (MLP)-based projection head to obtain the learned representations. The causality is enforced through masked self-attention within each Transformer layer. Let η_q denote the *query state embeddings*. We have:

$$\eta_q := \{\hat{s}_i\}_{i=t}^{t+F-1} = g_\phi(\tilde{\eta}_{s,a}). \quad (5)$$

4.3. Temporal Contrastive Learning

The guiding principle of MOOSS is to learn state representations that evolve temporally in a gradual, smooth fashion, similar to the slowness and variability principles firstly proposed in [24]. Recall that $\eta_q = g_\phi(\tilde{\eta}_{s,a})$, $\eta_k = f_{\bar{\theta}}(\eta_o)$ are the query and key trajectories encoded from η_o , respectively. In addition, let $\{\eta'_k\} = f_{\bar{\theta}}(\{\eta'_o\})$ be the set of key trajectories encoded from other observation sequences of the

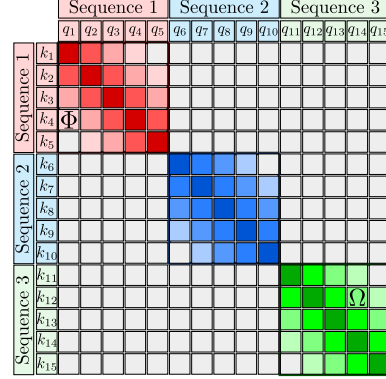


Figure 4. Illustration of the temporal contrastive objective. This mock setup contains 3 sampled sequences with 15 query-key pairs in total (observation length is $F = 5$; batch size is 3), and models four similarity levels with $L = 3$. If embeddings are learned from the same sequence, they share the same color scheme. The temporal contrastive objective aims to capture a ranked order of state similarities, indicated by the diminishing color intensity from the main diagonal to the off-diagonal cells. In this example, $\Phi = \text{sim}(\mathbf{q}_1, \mathbf{k}_4) = \text{sim}(\mathbf{q}, \mathbf{k}_{\Delta=3})$, and $\Omega = \text{sim}(\mathbf{q}_{14}, \mathbf{k}_{12}) = \text{sim}(\mathbf{q}, \mathbf{k}_{\Delta=2})$. The gray cells denote learned similar scores between \mathbf{q} and \mathbf{k}' , *i.e.*, query-key pairs either belonging to different sampled sequences, or have temporal distance larger than 3. These pairs belong to the lowest similarity level.

same batch, *i.e.*, $\eta_k \notin \{\eta'_k\}$. Then, for any query $\mathbf{q} \in \eta_q$, we can form its corresponding sets of ranked keys $\{\mathcal{K}_{\Delta=l}\}_{l=0}^L$, to encourage \mathbf{q} is more similar to its temporally adjacent neighbors than those further apart. That is:

$$\text{sim}(\mathbf{q}, \mathbf{k}_{\Delta=0}) > \text{sim}(\mathbf{q}, \mathbf{k}_{\Delta=1}) > \dots > \text{sim}(\mathbf{q}, \mathbf{k}_{\Delta=L}) > \text{sim}(\mathbf{q}, \mathbf{k}'), \forall \mathbf{k}_{\Delta=l} \in \mathcal{K}_{\Delta=l}, \mathbf{k}' \in \{\eta'_k\} \cup \mathcal{K}_{\Delta>l}, \quad (6)$$

where $\mathbf{k}_{\Delta=l} \in \eta_k$ denotes key states that are l units temporally away from \mathbf{q} , $\mathbf{k}' \in \{\eta'_k\}$ are key states that do not come from η_k , and L is the temporal window size on which the contrastive objective focuses. Figure 4 illustrates this pattern.

To model such decaying query-key similarities at multiple levels, inspired by [21], we use the InfoNCE loss shown in Eq. (1) in a recursive manner from $l = 0$ to $l = L$. Specifically, at the l^{th} temporal distance level, the corresponding loss treats $\mathbf{k}_{\Delta=l}$ as positive keys, while the negatives consist of (1) keys from the same trajectory that are temporally further away and (2) keys from other trajectories in the batch. Formally, let $\mathcal{L}_{\text{MOOSS}} = \sum_{l=0}^L \mathcal{L}_{\mathbf{q}}^l$ denote MOOSS’s objective for query \mathbf{q} , where $\mathcal{L}_{\mathbf{q}}^l$ be the l^{th} -level temporal contrastive loss. We have:

$$\mathcal{L}_{\mathbf{q}}^l = -\log \frac{\sum_{\mathbf{k}_{\Delta=l}} \exp(\text{sim}(\mathbf{q}, \mathbf{k})/\tau_l)}{\sum_{\mathbf{k}_{\Delta \geq l} \cup \mathbf{k}'} \exp(\text{sim}(\mathbf{q}, \mathbf{k})/\tau_l)}, \quad (7)$$

where $\mathbf{k}_{\Delta \geq l} \in \eta_k$ denotes key states that are more than or equal to l -temporally away from \mathbf{q} , and $\tau_l < \tau_{l+1}$. MOOSS’s

100k Step Scores	Dreamer	SAC+AE	CURL	DrQ	PlayVirtual	MLR	Base	MOOSS
Finger, spin	341 ± 70	740 ± 64	767 ± 56	901 ± 104	915 ± 49	907 ± 58	853 ± 112	822 ± 6
Cartpole, swingup	326 ± 27	311 ± 11	582 ± 146	759 ± 92	816 ± 36	806 ± 48	784 ± 63	873 ± 1
Reacher, easy	314 ± 155	274 ± 14	538 ± 233	601 ± 213	785 ± 142	866 ± 103	593 ± 118	969 ± 7
Cheetah, run	235 ± 137	267 ± 24	299 ± 48	344 ± 67	474 ± 50	482 ± 38	399 ± 80	506 ± 15
Walker, walk	277 ± 12	394 ± 22	403 ± 24	612 ± 164	460 ± 173	643 ± 114	424 ± 281	798 ± 42
Ball in cup, catch	246 ± 174	391 ± 82	769 ± 43	913 ± 53	926 ± 31	933 ± 16	648 ± 287	944 ± 30
Mean	289.8	396.2	559.7	688.3	729.3	772.8	616.8	818.6
Median	295.5	351.0	560.0	685.5	800.5	836.0	620.5	847.5
500k Step Scores								
Finger, spin	796 ± 183	884 ± 128	926 ± 45	938 ± 103	963 ± 40	973 ± 31	944 ± 97	977 ± 8
Cartpole, swingup	762 ± 27	735 ± 63	841 ± 45	868 ± 10	865 ± 11	872 ± 5	871 ± 4	878 ± 0
Reacher, easy	793 ± 164	627 ± 58	929 ± 44	942 ± 71	942 ± 66	957 ± 41	943 ± 52	977 ± 12
Cheetah, run	570 ± 253	550 ± 34	518 ± 28	660 ± 96	719 ± 51	674 ± 37	602 ± 67	712 ± 7
Walker, walk	897 ± 49	847 ± 48	902 ± 43	921 ± 45	928 ± 30	939 ± 10	818 ± 263	957 ± 22
Ball in cup, catch	879 ± 87	794 ± 58	959 ± 27	963 ± 9	967 ± 5	964 ± 14	960 ± 10	974 ± 15
Mean	782.8	739.5	845.8	882.0	897.3	896.5	856.3	912.5
Median	794.5	764.5	914.0	929.5	935.0	948.0	907.0	965.5

Table 1. Quantitative results for DMC-100k and DMC-500k, as reported in their respective works. **Bold** values indicate best performance.

similarity score is measured by bilinear product as in previous works [31, 47] through $\text{sim}(\mathbf{q}, \mathbf{k}) = \mathbf{q}^T \mathbf{W} \mathbf{k}$, where \mathbf{W} is a learnable weight matrix.

4.4. Overall Objective

The temporal contrastive objective $\mathcal{L}_{\text{MOOSS}}$ serves as an auxiliary loss for RL algorithms. Let \mathcal{L}_{rl} denote the loss for the base RL algorithm. The overall learning objective for the visual RL agent with MOOSS is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rl}} + \lambda \mathcal{L}_{\text{MOOSS}}, \quad (8)$$

where λ is a hyper-parameter trading off the main RL loss and MOOSS’s temporal contrastive loss. We note that the proposed predictive decoder $g_{\phi}(\cdot)$ is only used during training. During evaluation, only the observation encoder $f_{\theta}(\cdot)$ is kept to encode raw, unmasked observations to states.

5. Experiments

5.1. Benchmark Environments

Sample efficiency of MOOSS is studied on both the continuous control benchmark DeepMind Control Suite (DMC) [51] and the discrete control benchmark Atari [4]. For continuous control, 6 tasks from DMC are used following prior works [62, 63], including *Finger-spin*, *Cartpole-swingup*, *Reacher-easy*, *Cheetah-run*, *Walker-walk* and *Ball in cup-catch*. Algorithms are evaluated at 100k and 500k environment steps, referred as DMC-100k and DMC-500k. For discrete control, the Atari-100k benchmark is used [31, 63]. It contains 26 Atari games, and performance is evaluated at 100k interaction steps (*i.e.*, 400k environment steps with action repeat of 4) between the game and RL agents.

5.2. Baselines and Metrics

For DMC, MOOSS is compared with sample-efficient RL methods tailored to continuous control, including Dreamer [14], SAC+AE [60], CURL [31], DrQ [29], PlayVirtual [62] and MLR [63]. Following previous works, per-task mean (with standard deviation) over 10 episodic runs with different seeds are reported. We also report the overall mean and median scores to reflect the general performance. For Atari experiments, MOOSS is compared with DER [54], OTR [27], CURL [31], DrQ [29], SPR [43], PlayVirtual [62] and MLR [63]. Each Atari game is evaluated through 100 episodic runs across 3 random seeds following [63]. We leverage the Interquartile Mean (IQM) and the Optimality Gap (OG) metrics with percentile Confidence Intervals (CIs) proposed in Rliable [1] to study MOOSS’s sample efficiency on Atari. As Atari games are highly non-deterministic with high variances across different games and runs, these aggregate metrics can provide a more rigorous and robust evaluation on algorithmic performance than raw scores. We report these aggregate metrics alongside individual game scores on Atari-100k with 95% CIs.

5.3. Implementation

SAC [13] and Rainbow [20] are used as continuous and discrete RL algorithms on DMC and Atari environments, respectively. Following previous works [63], data augmentation including random crop and random intensity are employed as they are proved helpful [29, 30] in improving sample efficiency of RL algorithms. Based on these, *Base* models [63] are firstly devised, which only rely on \mathcal{L}_{rl} for policy updates by setting $\lambda = 0$. Then, we integrate MOOSS into the *Base* models. For all DMC and Atari experiments, we set $\lambda = 0.1$ to balance \mathcal{L}_{rl} and $\mathcal{L}_{\text{MOOSS}}$. By default, we set the temporal window size $L = 6$ and the mask ratio

Game	Human	Random	DER	OTR	CURL	DrQ	SPR	PlayVirtual	MLR	Base	MOOSS
Alien	7127.7	227.8	802.3	570.8	711.0	734.1	841.9	947.8	990.1	678.5	951.1
Amidar	1719.5	5.8	125.9	77.7	113.7	94.2	179.7	165.3	227.7	132.8	207.5
Assault	742.0	222.4	561.5	330.9	500.9	479.5	565.6	702.3	643.7	493.3	667.0
Asterix	8503.3	210.0	535.4	334.7	567.2	535.6	962.5	933.3	883.7	1021.3	1140.0
Bank Heist	753.1	14.2	185.5	55.0	65.3	153.4	345.4	245.9	180.3	288.2	288.0
Battle Zone	37187.5	2360.0	8977.0	5139.4	8997.8	10563.6	14834.1	13260.0	16080.0	13076.7	11363.3
Boxing	12.1	0.1	-0.3	1.6	0.9	6.6	35.7	38.3	26.4	14.3	22.4
Breakout	30.5	1.7	9.2	8.1	2.6	15.4	19.6	20.6	16.8	16.7	16.8
Chopper Cmd	7387.8	811.0	925.9	813.3	783.5	792.4	946.3	922.4	910.7	878.7	1477.0
Crazy Climber	35829.4	10780.5	34508.6	14999.3	9154.4	21991.6	36700.5	23176.7	24633.3	28235.7	21093.3
Demon Attack	1971.0	152.1	627.6	681.6	646.5	1142.4	517.6	1131.7	854.6	310.5	904.0
Freeway	29.6	0.0	20.9	11.5	28.3	17.8	19.3	16.1	30.2	30.9	20.3
Frostbite	4334.7	65.2	871.0	224.9	1226.5	508.1	1170.7	1984.7	2381.1	994.3	2898.5
Gopher	2412.5	257.6	467.0	539.4	400.9	618.0	660.6	684.3	822.3	650.9	731.4
Hero	30826.4	1027.0	6226.0	5956.5	4987.7	3722.6	5858.6	8597.5	7919.3	4661.2	9531.2
Jamesbond	302.8	29.0	275.7	88.0	331.0	251.8	366.5	394.7	423.2	270.0	326.3
Kangaroo	3035.0	52.0	581.7	348.5	740.2	974.5	3617.4	2384.7	8516.0	5036.0	6122.7
Krull	2665.5	1598.0	3256.9	3655.9	3049.2	4131.4	3681.6	3880.7	3923.1	3571.3	4195.9
Kung Fu Master	22736.3	258.5	6580.1	6659.6	8155.6	7154.5	14783.2	14259.0	10652.0	10517.3	19402.3
Ms Pacman	6951.6	307.3	1187.4	908.0	1064.0	1002.9	1318.4	1335.4	1481.3	1320.9	1362.2
Pong	14.6	-20.7	-9.7	-2.5	-18.5	-14.3	-5.4	-3.0	4.9	-3.1	-4.14
Private Eye	69571.3	24.9	72.8	59.6	81.9	24.8	86.0	93.9	100.0	93.3	100.0
Qbert	13455.0	163.9	1773.5	552.5	727.0	934.2	866.3	3620.1	3410.4	553.8	3398.0
Road Runner	7845.0	11.5	11843.4	2606.4	5006.1	8724.7	12213.1	13429.4	12049.7	12337.0	19077.0
Seaquest	42054.7	68.4	304.6	272.9	315.2	310.5	558.1	532.9	628.3	471.9	455.5
Up N Down	11693.2	533.4	3075.0	2331.7	2646.4	3619.1	10859.2	10225.2	6675.7	4112.8	6963.9
Interquartile Mean	1.000	0.000	0.183	0.117	0.113	0.224	0.337	0.374	0.432	0.292	0.433
Optimality Gap	0.000	1.000	0.698	0.819	0.768	0.692	0.577	0.558	0.522	0.614	0.524

Table 2. Quantitative results for Atari-100k. The best results are highlighted in bold.

$p_m = 50\%$, and these key hyper-parameters are further studied in the supplementary material. More implementation details are also provided in the supplementary material.

5.4. Comparison with *Base* and State of the Art

DMC. We first compare MOOSS with state-of-the-art visual RL methods and its *Base* model on DMC-100k and DMC-500k. The evaluation results are summarized in Tab. 1. From the table, we first observe that MOOSS consistently improves the performance of its corresponding *Base* model on all tasks by large margins on both benchmarks. In particular, MOOSS achieves relative improvements of **33%** in mean scores and **37%** in median scores on DMC-100k, and **7%** in mean scores and **6%** in median scores on DMC-500k, respectively. These improvements clearly demonstrate MOOSS’s ability in improving sample efficiency of visual RL algorithms on continuous control tasks. Second, MOOSS-equipped RL agents outperform previous state-of-the-art methods. For both DMC-100k and DMC-500k, MOOSS secures the top performance in five out of six tasks, and obtain the best mean and median scores. These results indicate that MOOSS is effective in both sample efficiency and asymptotic performance.

Atari. In Tab. 2, we summarize MOOSS’s quantitative results on Atari-100k. From the table, we again observe that

MOOSS significantly improves the performance of its corresponding *Base* model, having a **48%** relative improvement on IQM and a **15%** relative improvement on OG, respectively. This indicates MOOSS can greatly improve sample efficiency of visual RL algorithms on discrete control tasks. In addition, MOOSS also performs competitively with the current state-of-the-art method MLR, achieving the best IQM score and the second best OG score. These results demonstrate that MOOSS has the highest sample efficiency and performs close to human-level performance.

5.5. Ablation Study

In this section, we conduct an ablation analysis on DMC-100k to investigate how different design choices of MOOSS affect its efficacy in improving sample efficiency. All ablation results are obtained through 10 evaluation runs across different seeds. Additional ablations are provided in the supplementary material.

General Framework Components. MOOSS enhances RL algorithms through its (1) temporal contrastive objective facilitated by the (2) random walk-based spatial-temporal masking. We first evaluate the individual contributions of these components to MOOSS’s performance. Specifically, in addition to MOOSS, we test four variants of our framework: (1) First, as previously mentioned, the *Base* model does not rely on $\mathcal{L}_{\text{MOOSS}}$ updates. (2) We then introduce

Model Variants \ Task	Finger	Cartpole	Reacher	Cheetah	Walker	Ball	Mean	Median
<i>Base</i> , $\lambda = 0$	853 ± 112	784 ± 63	593 ± 118	399 ± 80	424 ± 281	648 ± 287	616.8	620.5
$L = 0, p_m = 0$	829 ± 9	795 ± 1	702 ± 409	401 ± 49	68 ± 41	766 ± 190	593.3	734.0
$L = 6, p_m = 0$	840 ± 20	870 ± 1	873 ± 291	491 ± 11	52 ± 24	931 ± 35	800.9	871.5
$L = 6, p_m = 50\%$ as [63]	656 ± 5	862 ± 9	676 ± 435	454 ± 53	547 ± 91	930 ± 35	687.4	666.0
MOOSS	822 ± 6	873 ± 1	969 ± 7	506 ± 15	798 ± 42	944 ± 30	818.6	847.5

Table 3. Ablation on MOOSS’s general framework components.

Task	<i>Base</i>	MOOSS-NoTrans	MOOSS-S	MOOSS-SAR	MOOSS
Finger	853 ± 112	975 ± 6	938 ± 10	827 ± 16	822 ± 6
Cartpole	784 ± 63	837 ± 2	527 ± 19	790 ± 9	873 ± 1
Reacher	593 ± 118	778 ± 387	872 ± 286	683 ± 441	969 ± 7
Cheetah	399 ± 80	427 ± 5	543 ± 19	559 ± 7	506 ± 15
Walker	424 ± 281	670 ± 120	284 ± 107	701 ± 63	798 ± 42
Ball	648 ± 287	956 ± 17	888 ± 58	899 ± 74	944 ± 30
Mean	616.8	773.7	675.4	743.2	818.6
Median	620.5	807.5	707.5	745.5	847.5

Table 4. Ablation on MOOSS’s predictive decoder $g_\phi(\cdot)$.

the contrastive objective into the *Base* model without masking ($p_m = 0$). At the same time, we set $L = 0$ such that the model does not consider temporally adjacent states thus does not model state evolution. (3) Next, we improve upon the second model by leveraging the temporal contrastive objective ($L = 6$), while keeping the masking ratio to 0. (4) In the fourth variant, we additionally leverage masking with $p_m = 50\%$. However, instead of doing random walk-based spatial-temporal masking, we apply cube masking [63], which masks the observation cubes uniformly.

Through analysing the results presented in Tab. 3, we have the following observations: (1) Both the temporal contrastive objective and the spatial-temporal masking technique improve the sample efficiency of RL algorithms. All variants equipping $\mathcal{L}_{\text{MOOSS}}$ perform better than the *Base* model in terms of mean and median scores. (2) The temporal contrastive objective is essential to MOOSS, as it brings a mean score improvement of 35% and a median score improvement of 19% when masking is not applied. (3) Masking is important to the performance of MOOSS on certain tasks. We observe that if masking is not used, the *Walker* task shows inferior performance even compared with the *Base* model. (4) MOOSS achieves superior performance compared to the *Base* model and its variants on most tasks, having the best mean score performance and the second best median score performance. This indicates the integration of temporal contrastive objective and the spatial-temporal masking technique can enhance RL agent’s understanding of the environment.

Decoder Setups. During training, MOOSS utilizes an additional predictive decoder $g_\phi(\cdot)$ to generate query states. We investigate different design choices of $g_\phi(\cdot)$: (1) MOOSS-NoTrans indicates no Transformer layers are used in the decoder. The masked state embeddings $\tilde{\eta}_s$ are only

decoded via an MLP head. (2) For the MOOSS-S case, only state embeddings are used as inputs to the Transformer-based decoder. (3) MOOSS-SAR indicates states, actions and rewards are all used as inputs to the decoder for query generation. From the results summarized in Tab. 4, we confirm that using states and actions as the inputs to MOOSS’s predictive decoder provides the best overall mean and median performance scores. This indicates the meaningful guidance provided by action signals in modeling state evolution across time. We also observe that MOOSS stays competitive on most tasks even without the predictive decoder. This suggests that the core principle of MOOSS – to capture the essential dynamics of states by modeling their evolution across time – is robust and effective.

6. Conclusion

In this work we present MOOSS, a novel framework with a self-supervised auxiliary objective to improve sample efficiency of visual RL algorithms. Facilitated by a graph-based spatial-temporal masking approach, MOOSS’s temporal contrastive objective goes beyond the binary distinction between positive and negative samples, modeling multiple levels of state similarities across the temporal dimension. In this way, we encourage the observation encoder to focus on the smoothly evolving nature of state changes over various temporal distances. The results obtained from extensive experiments and analyses confirm that MOOSS achieves significant sample efficiency gains over the base method and state-of-the-art works on both DMControl and Atari benchmarks, demonstrating the efficacy of our method.

Acknowledgements: This work is supported in part by Navy N00014-19-1-2373, the joint NSF-USDA CPS Frontier project CNS #1954556, USDA-NIFA #2021-67021-34418, and Agriculture and Food Research Initiative (AFRI) grant no. 2020-67021-32799/project accession no.1024178 from the USDA National Institute of Food and Agriculture: NSF/USDA National AI Institute: AIFARMS. Work is supported in part by NSF MRI grant #1725729 [28]. Work also used Delta GPU at NCSA Delta through allocation CIS230331 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program [6], which is supported by NSF grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

- [1] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021. 6
- [2] Ankesh Anand, Evan Racah, Sherjil Ozair, Yoshua Bengio, Marc-Alexandre Côté, and R Devon Hjelm. Unsupervised state representation learning in atari. *Advances in neural information processing systems*, 32, 2019. 3
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 3
- [4] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013. 2, 6
- [5] Richard Bellman. A markovian decision process. *Journal of mathematics and mechanics*, pages 679–684, 1957. 4
- [6] Timothy J Boerner, Stephen Deems, Thomas R Furlani, Shelley L Knuth, and John Towns. Access: Advancing innovation: Nsf’s advanced cyberinfrastructure coordination ecosystem: Services & support. In *Practice and Experience in Advanced Research Computing*, pages 173–176. 2023. 8
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [9] Blattmann et.al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 2
- [10] Peng Gao, Teli Ma, Hongsheng Li, Ziyi Lin, Jifeng Dai, and Yu Qiao. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*, 2022. 3
- [11] Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G Bellemare. Deepmdp: Learning continuous latent space models for representation learning. In *International conference on machine learning*, pages 2170–2179. PMLR, 2019. 2
- [12] Zhaohan Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-Bastien Grill, Florent Althé, Rémi Munos, and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. In *International Conference on Machine Learning*, pages 3875–3886. PMLR, 2020. 2, 4
- [13] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018. 4, 6
- [14] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019. 2, 6
- [15] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR, 2019. 2
- [16] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *Advances in Neural Information Processing Systems*, 33:5679–5690, 2020. 3
- [17] Nicklas Hansen, Hao Su, and Xiaolong Wang. Stabilizing deep q-learning with convnets and vision transformers under data augmentation. *Advances in neural information processing systems*, 34:3680–3693, 2021. 2
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 3
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3, 4
- [20] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 4, 6
- [21] David T Hoffmann, Nadine Behrmann, Juergen Gall, Thomas Brox, and Mehdi Noroozi. Ranking info noise contrastive estimation: Boosting contrastive learning via ranked positives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 897–905, 2022. 3, 5
- [22] Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, Chunjie Wang, and Jie Tang. Graphmae: Self-supervised masked graph autoencoders. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 594–604, 2022. 3
- [23] Yangru Huang, Peixi Peng, Yifan Zhao, Guangyao Chen, and Yonghong Tian. Spectrum random masking for generalization in image-based reinforcement learning. *Advances in Neural Information Processing Systems*, 35:20393–20406, 2022. 2
- [24] Rico Jonschkowski, Roland Hafner, Jonathan Scholz, and Martin Riedmiller. Pves: Position-velocity encoders for unsupervised learning of structured state representations. *arXiv preprint arXiv:1705.09805*, 2017. 5
- [25] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998. 4
- [26] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey

- Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019. [2](#)
- [27] Kacper Piotr Kielak. Do recent advancements in model-based deep reinforcement learning really improve data efficiency? *arXiv preprint arXiv:2003.10181v1*, 2019. [6](#)
- [28] Volodymyr Kindratenko, Dawei Mu, Yan Zhan, John Maloney, Sayed Hadi Hashemi, Benjamin Rabe, Ke Xu, Roy Campbell, Jian Peng, and William Gropp. Hal: Computer system for scalable deep learning. In *Practice and Experience in Advanced Research Computing*, PEARC '20, page 41–48, New York, NY, USA, 2020. Association for Computing Machinery. [8](#)
- [29] Ilya Kostrikov, Denis Yarats, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. *arXiv preprint arXiv:2004.13649*, 2020. [2](#), [6](#)
- [30] Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *Advances in neural information processing systems*, 33:19884–19895, 2020. [2](#), [6](#)
- [31] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, pages 5639–5650. PMLR, 2020. [1](#), [3](#), [6](#)
- [32] Alex X Lee, Anusha Nagabandi, Pieter Abbeel, and Sergey Levine. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *Advances in Neural Information Processing Systems*, 33:741–752, 2020. [2](#)
- [33] Jintang Li, Ruofan Wu, Wangbin Sun, Liang Chen, Sheng Tian, Liang Zhu, Changhua Meng, Zibin Zheng, and Weiqiang Wang. What’s behind the mask: Understanding masked graph modeling for graph autoencoders. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1268–1279, 2023. [3](#)
- [34] Guoqing Liu, Chuheng Zhang, Li Zhao, Tao Qin, Jinhua Zhu, Jian Li, Nenghai Yu, and Tie-Yan Liu. Return-based contrastive representation learning for reinforcement learning. *arXiv preprint arXiv:2102.10960*, 2021. [3](#)
- [35] Hao Liu and Pieter Abbeel. Aps: Active pretraining with successor features. In *International Conference on Machine Learning*, pages 6736–6747. PMLR, 2021. [3](#)
- [36] Hao Liu and Pieter Abbeel. Behavior from the void: Unsupervised active pre-training. *Advances in Neural Information Processing Systems*, 34:18459–18473, 2021. [3](#)
- [37] Guozheng Ma, Linrui Zhang, Haoyu Wang, Lu Li, Zilin Wang, Zhen Wang, Li Shen, Xueqian Wang, and Dacheng Tao. Learning better with less: Effective augmentation for sample-efficient visual reinforcement learning. *arXiv preprint arXiv:2305.16379*, 2023. [2](#)
- [38] Bogdan Mazouze, Remi Tachet des Combes, Thang Long Doan, Philip Bachman, and R Devon Hjelm. Deep reinforcement and infomax learning. *Advances in Neural Information Processing Systems*, 33:3686–3698, 2020. [3](#)
- [39] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020. [3](#)
- [40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [3](#)
- [41] Minting Pan, Xiangming Zhu, Yunbo Wang, and Xiaokang Yang. Iso-dream: Isolating and leveraging noncontrollable visual dynamics in world models. *Advances in neural information processing systems*, 35:23178–23191, 2022. [2](#)
- [42] Keiran Paster, Sheila A McIlraith, and Jimmy Ba. Planning from pixels using inverse dynamics models. *arXiv preprint arXiv:2012.02419*, 2020. [1](#)
- [43] Max Schwarzer, Ankesh Anand, Rishab Goel, R Devon Hjelm, Aaron Courville, and Philip Bachman. Data-efficient reinforcement learning with self-predictive representations. *arXiv preprint arXiv:2007.05929*, 2020. [1](#), [2](#), [6](#)
- [44] Max Schwarzer, Nitarshan Rajkumar, Michael Noukhovitch, Ankesh Anand, Laurent Charlin, R Devon Hjelm, Philip Bachman, and Aaron C Courville. Pretraining representations for data-efficient reinforcement learning. *Advances in Neural Information Processing Systems*, 34:12686–12699, 2021. [3](#)
- [45] Younggyo Seo, Danijar Hafner, Hao Liu, Fangchen Liu, Stephen James, Kimin Lee, and Pieter Abbeel. Masked world models for visual control. In *Conference on Robot Learning*, pages 1332–1344. PMLR, 2023. [2](#)
- [46] Evan Shelhamer, Parsa Mahmoudieh, Max Argus, and Trevor Darrell. Loss is its own reward: Self-supervision for reinforcement learning. *arXiv preprint arXiv:1612.07307*, 2016. [1](#), [2](#)
- [47] Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, pages 9870–9879. PMLR, 2021. [1](#), [3](#), [6](#)
- [48] Jiarui Sun, Yujie Fan, Chin-Chia Michael Yeh, Wei Zhang, and Girish Chowdhary. Revealing the power of masked autoencoders in traffic forecasting. *arXiv preprint arXiv:2309.15169*, 2024. [3](#)
- [49] Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*, 2019. [3](#)
- [50] Qiaoyu Tan, Ninghao Liu, Xiao Huang, Soo-Hyun Choi, Li Li, Rui Chen, and Xia Hu. S2gae: Self-supervised graph autoencoders are generalizable learners with graph masking. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 787–795, 2023. [3](#)
- [51] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018. [1](#), [2](#), [5](#), [6](#)
- [52] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. [3](#)

- [53] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [1](#)
- [54] Hado P Van Hasselt, Matteo Hessel, and John Aslanides. When to use parametric models in reinforcement learning? *Advances in Neural Information Processing Systems*, 32, 2019. [6](#)
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [5](#)
- [56] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008. [3](#)
- [57] Che Wang, Xufang Luo, Keith Ross, and Dongsheng Li. Vrl3: A data-driven framework for visual deep reinforcement learning. *Advances in Neural Information Processing Systems*, 35:32974–32988, 2022. [3](#)
- [58] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14733–14743, 2022. [3](#)
- [59] Dongkuan Xu, Wei Cheng, Dongsheng Luo, Haifeng Chen, and Xiang Zhang. Infogcl: Information-aware graph contrastive learning. *Advances in Neural Information Processing Systems*, 34:30414–30425, 2021. [3](#)
- [60] Denis Yarats, Amy Zhang, Ilya Kostrikov, Brandon Amos, Joelle Pineau, and Rob Fergus. Improving sample efficiency in model-free reinforcement learning from images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10674–10681, 2021. [2](#), [5](#), [6](#)
- [61] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33:5812–5823, 2020. [3](#)
- [62] Tao Yu, Cuiling Lan, Wenjun Zeng, Mingxiao Feng, Zhizheng Zhang, and Zhibo Chen. Playvirtual: Augmenting cycle-consistent virtual trajectories for reinforcement learning. *Advances in Neural Information Processing Systems*, 34:5276–5289, 2021. [2](#), [6](#)
- [63] Tao Yu, Zhizheng Zhang, Cuiling Lan, Yan Lu, and Zhibo Chen. Mask-based latent reconstruction for reinforcement learning. *Advances in Neural Information Processing Systems*, 35:25117–25131, 2022. [2](#), [4](#), [6](#), [8](#)
- [64] Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*, 2020. [1](#)
- [65] Ruijie Zheng, Xiyao Wang, Yanchao Sun, Shuang Ma, Jieyu Zhao, Huazhe Xu, Hal Daumé III, and Furong Huang. Taco: Temporal latent action-driven contrastive loss for visual reinforcement learning. *arXiv preprint arXiv:2306.13229*, 2023. [3](#)
- [66] Jinhua Zhu, Yingce Xia, Lijun Wu, Jiajun Deng, Wengang Zhou, Tao Qin, Tie-Yan Liu, and Houqiang Li. Masked contrastive representation learning for reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3421–3433, 2022. [2](#)