

Robust Representation Learning by Clustering with Bisimulation Metrics for Visual Reinforcement Learning with Distractions

Qiyuan Liu¹, Qi Zhou¹, Rui Yang¹, Jie Wang^{1,2*}

¹CAS Key Laboratory of Technology in GIPAS, University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{qiyuanliu,zhouqida,yr0013}@mail.ustc.edu.cn, jiewangx@ustc.edu.cn

Abstract

Recent work has shown that representation learning plays a critical role in sample-efficient reinforcement learning (RL) from pixels. Unfortunately, in real-world scenarios, representation learning is usually fragile to task-irrelevant distractions such as variations in background or viewpoint. To tackle this problem, we propose a novel clustering-based approach, namely **Clustering with Bisimulation Metrics (CBM)**, which learns robust representations by grouping visual observations in the latent space. Specifically, CBM alternates between two steps: (1) grouping observations by measuring their bisimulation distances to the learned prototypes; (2) learning a set of prototypes according to the current cluster assignments. Computing cluster assignments with bisimulation metrics enables CBM to capture task-relevant information, as bisimulation metrics quantify the behavioral similarity between observations. Moreover, CBM encourages the consistency of representations within each group, which facilitates filtering out task-irrelevant information and thus induces robust representations against distractions. An appealing feature is that CBM can achieve sample-efficient representation learning even if multiple distractions exist simultaneously. Experiments demonstrate that CBM significantly improves the sample efficiency of popular visual RL algorithms and achieves state-of-the-art performance on both multiple and single distraction settings. The code is available at <https://github.com/MIRALab-USTC/RL-CBM>.

Introduction

Reinforcement learning (RL) from visual observations has achieved remarkable success in various domains ranging from video games (Mnih et al. 2015; Lample and Chaplot 2017) to robotics manipulation (Levine et al. 2016; Kalashnikov et al. 2018). Recently, representation learning that embeds image inputs into low-dimensional vectors has drawn great attention in visual RL. Popular approaches include learning sequential autoencoders (Yarats et al. 2021c; Lee et al. 2020a), applying data augmentation (Yarats, Kostrikov, and Fergus 2020; Raileanu et al. 2020) or constructing auxiliary tasks (Jaderberg et al. 2016; Srinivas, Laskin, and Abbeel 2020; Stooke et al. 2021). However, recent work has shown that many of these methods can be easily distracted

by task-irrelevant distractions, such as variations in background or viewpoint (Stone et al. 2021; Fu et al. 2021). A fundamental challenge is to learn robust representations that effectively capture task-relevant information and are invariant to the task-irrelevant distractions (Zhang et al. 2021).

In order to learn robust representations, many recent approaches remedy the overfitting of the encoder by applying strong data augmentation, such as color-jitter (Laskin et al. 2020) or random convolution (Lee et al. 2019). However, they require prior knowledge of the distraction to choose semantic-preserving image transformations in RL settings (Fan et al. 2021). Moreover, many of these methods require access to the environments without distractions for sample-efficient policy optimization. (Hansen, Su, and Wang 2021; Hansen and Wang 2020). Another line of work proposes a series of auxiliary tasks. One of the most popular schemes is to introduce contrastive objectives, which can encourage representations to preserve dynamics of the latent space (Okada and Taniguchi 2021; Nguyen et al. 2021) or behavioral similarity between states (Zhang et al. 2021; Agarwal et al. 2021). However, these contrastive-based methods usually rely on carefully constructed positive-negative pairs or a large batch size (Caron et al. 2020).

In this paper, we propose a novel clustering-based approach, namely **clustering with bisimulation metrics (CBM)**, which learns robust representations by discriminating between groups of visual observations with similar behavior. Specifically, at each training step, CBM first computes predicted cluster assignments via the geometric similarity between representations and a set of prototypes. Then, CBM calculates target cluster assignments by measuring the bisimulation distances between visual observations and prototypes. Finally, CBM optimizes the encoder and prototypical representations by enforcing the consistency between the predicted and target assignments. An appealing feature of CBM is that it can improve the robustness against distractions without domain knowledge or carefully constructed pairwise comparison. Moreover, compared with previous clustering-based methods, CBM exploits the properties of RL tasks and achieves task-specific clustering, which enables CBM to extract task-relevant information effectively.

The proposed method CBM is compatible with most visual RL algorithms. In our experiments, we combine CBM with two popular baselines, DrQ (Yarats, Kostrikov, and Fer-

*Corresponding author.

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

gus 2020) and DrQ-v2 (Yarats et al. 2021a). We perform empirical evaluations on Distracting Control Suite (Stone et al. 2021). Experiments demonstrate that our method significantly improves the performance of these two algorithms and achieves state-of-the-art performance both on multiple and single distraction settings. We also conduct quantitative analysis and visualization of the clustering results to show that our method learns representations that hardly contain task-irrelevant information.

To summarize, this paper makes the following contributions: (1) We are the first to incorporate the task-specific properties of RL problems into clustering-based representation learning. (2) We propose a novel approach CBM, which learns robust representations by grouping visual observations with bisimulation metrics. (3) Our experiments demonstrate that the proposed CBM significantly improves performance over popular visual RL baselines and achieves sample-efficient representation learning even if multiple distractions exist simultaneously.

Related Work

Clustering for representation learning. Learning representations by clustering is one of the most promising approaches for self-supervised learning of neural networks. Caron et al. (2018) cluster deep features by k-means and use the cluster assignments as pseudo-labels to learn convnets. To avoid degenerate solutions, Asano, Rupprecht, and Vedaldi (2020) propose a principled formulation by adding the constraint that the labels must induce equipartition of the data. Later, Caron et al. (2020) learned prototypical representations through contrastive losses. They obtain online assignments and make it possible to scale to any dataset size. Following this, recent approaches (Yarats et al. 2021b; Deng, Jang, and Ahn 2021) combine prototypical representations and dynamics learning in the RL setting. In contrast to these methods, CBM incorporates the task-specific properties of RL problems into the clustering process and thus induces robust representations against distractions.

Representation learning for RL Many recent methods have taken inspiration from the successes of representation learning in computer vision to improve sample efficiency and generalization (Vincent et al. 2008; Chen et al. 2020; Xie et al. 2020). Learning autoencoders by minimizing reconstruction errors has been proven effective in visual RL (Lee et al. 2020a; Yarats et al. 2021c; Hafner et al. 2019). Later, Several researchers Stooke et al. (2021); Zhu et al. (2020) perform contrastive learning that maximizes agreement between augmented versions of the same observation. More recently, adopting data augmentations such as random shifts (Yarats, Kostrikov, and Fergus 2020; Yarats et al. 2021a) has also been demonstrated effective in visual RL. Another line of work (Lee et al. 2020b; Mazouze et al. 2020; Yu et al. 2021) has designed various auxiliary tasks to encourage the model to capture the predictive information. However, when training with distractions, most of these methods suffer from performance degradation due to the interference of task-irrelevant information.

RL with visual distractions. Learning control from pixels with distractions requires robust representation learning.

Recent work has demonstrated the effectiveness of strong augmentation techniques in RL (Laskin et al. 2020). However, the strong augmentation may cause training instability or divergence. Recent approaches resolve this problem by policy distillation (Fan et al. 2021) or Q function stabilization (Hansen, Su, and Wang 2021). Although their results are encouraging, these methods usually require domain knowledge to choose proper augmentation or clean environments without distractions for sample-efficient training. In model-based RL, many papers (Nguyen et al. 2021; Okada and Taniguchi 2021) replace the reconstruction-based objectives with contrastive-based losses, which encourages the encoder to capture controllable and predictable information in the latent space. Another kind of method (Agarwal et al. 2021; Zhang et al. 2021) proposes to learn invariant representations for RL by forcing the latent space to preserve the behavioral similarity. These methods rely on constructing positive pairs or learning accurate dynamics models, which may be difficult in environments with complex distractions.

Preliminaries

In this section, we first present the notation and the definition of bisimulation metrics. Then we briefly introduce the visual RL baselines and previous clustering-based representation learning methods that CBM builds on top of.

Notation

We consider the environment as an infinite horizon Markov decision process defined by a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma \rangle$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{P}(\cdot | s, a)$ is the transitioning probability from state $s \in \mathcal{S}$ to $s' \in \mathcal{S}$, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a reward function, and $\gamma \in [0, 1]$ is the discount factor. Since image observations only provide partial state observability, we define a state s_t by stacking several consecutive image observations $s_t = \{\mathbf{o}_t, \mathbf{o}_{t-1}, \dots, \mathbf{o}_{t-k}\}$, $\mathbf{o} \in \mathcal{O}$, where \mathcal{O} is the high-dimensional observation space (image pixels). The goal is to learn a parameterized policy $\pi_\theta : \mathcal{S} \rightarrow \mathcal{A}$ such that maximizes the cumulative return $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$.

Bisimulation Metrics

Intuitively, two states are behaviorally equivalent if they receive indistinguishable reward sequences given any action. By recursively utilizing the reward signal, bisimulation metrics (Ferns, Panangaden, and Precup 2011) define a distance function $d : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ that measures how behaviorally differently two states are.

Definition 0.1. (Ferns, Panangaden, and Precup 2011) *Given two arbitrary states $s_1, s_2 \in \mathcal{S}$ and $c \in [0, 1]$, the bisimulation distance between s_1, s_2 is defined as*

$$d(s_1, s_2) = \max_{a \in \mathcal{A}} ((1 - c) |r_{s_1}^a - r_{s_2}^a| + c W_1(\mathcal{P}(\cdot | s_1, a), \mathcal{P}(\cdot | s_2, a); d))$$

The definition of bisimulation metrics consists of a distance between rewards and distance between state distributions. The latter is computed by Wasserstein metric W_1 , which denotes the cost of transport mass from one distribution to another (Villani 2021).

Data-Regularized Q for Pixel-RL

Data-Regularized Q (DrQ) (Yarats, Kostrikov, and Fergus 2020) uses an optimality invariant state transformation f , which preserves the Q values to augment the training data when optimizing the Q functions. DrQ generates K samples for each raw image by applying random transformations and estimates the Q function and Q target by averaging over the augmented images. The addition of the augmentation method improves the performance of SAC (Haarnoja et al. 2018) in DeepMind Control Suite (Tassa et al. 2018).

Recently, Yarats et al. (2021a) propose DrQ-v2 that combines DrQ and TD3 (Fujimoto, Hoof, and Meger 2018). DrQ-v2 replaces the maximum entropy term (Haarnoja et al. 2018) with a scheduled noise for adjustable exploration and use target policy smoothing to reduce the bias of Q functions. Moreover, DrQ-v2 uses multi-step TD to learn value functions. Experiments show that DrQ-v2 achieves more efficient learning than DrQ in complex visual control tasks.

SwAV and Proto-RL

To learn visual features without supervision, Caron et al. (2020) proposes a clustering-based method, Swapping Assignments between different views (SwAV), which computes cluster assignments in an online fashion and enforces consistency between cluster assignments produced for different augmentations of the same image.

Specifically, given two image features \mathbf{z}_s and \mathbf{z}_t from two different augmentations of a same image, they compute their codes \mathbf{q}_s and \mathbf{q}_t by matching the features to K learnable prototypes $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_K]$. Then they setup a swapped prediction problem with the following loss function:

$$L(\mathbf{z}_t, \mathbf{z}_s) = \ell(\mathbf{z}_t, \mathbf{q}_s) + \ell(\mathbf{z}_s, \mathbf{q}_t).$$

where the ℓ represents the cross entropy loss between the code and the probability obtained by taking a softmax of the dot products of \mathbf{z} and all prototypes.

To ensure equal partitioning of the prototypes across all embeddings, given B feature vectors $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_B]$, the codes $\mathbf{Q} = [\mathbf{q}_1, \dots, \mathbf{q}_B]$ are calculated by

$$\mathbf{Q} = \text{Diag}(\mathbf{u}) \exp\left(\frac{\mathbf{C}^\top \mathbf{Z}}{\varepsilon}\right) \text{Diag}(\mathbf{v}), \quad (1)$$

where $\mathbf{u} \in \mathbb{R}^K$ and $\mathbf{v} \in \mathbb{R}^B$ are renormalization vectors, and ε is a parameter that controls the smoothness of the codes. The renormalization vectors are computed using a small number of matrix multiplications using the iterative Sinkhorn-Knopp algorithm (Cuturi 2013).

Proto-RL (Yarats et al. 2021b) draws inspiration from SwAV and adapts prototypical representation into RL setting. Different from SwAV, they enforce the consistency of cluster assignments from the predicted transition and the ground truth transition, which encourages the representations to preserve one-step dynamical similarity. They also demonstrate that the learned prototypes form the basis of latent space and thus induce efficient downstream exploration.




Figure 1: Measure clustering quality with respect to physical states. The agents are trained with multiple distractions

Methods

In this section, we propose clustering with bisimulation metrics (CBM)—which groups observations in the latent space with bisimulation metrics—to learn robust representations against distractions. Specifically, we first measure the clustering quality of clustering-based methods to study whether they capture task-relevant information and filter out task-irrelevant information. Secondly, we propose to incorporate the task-specific properties in RL problems by introducing bisimulation metrics into the clustering process. Then, we propose an approach to approximate the bisimulation metrics in the clustering process. Finally, we describe the overall algorithm and implementation details of CBM.

Measuring Clustering Quality in RL

In this part, we take Proto-RL (Yarats et al. 2021b) for example to study whether previous clustering-based methods learn robust representations which capture task-relevant information and be invariant to task-irrelevant distractions. To capture task-relevant information, clustering-based methods should discriminate observations that correspond to different physical states. Moreover, to filter out task-irrelevant information, clustering-based methods should assign observations corresponding to similar states to the same group and encode them as neighboring points in the latent space. Therefore, we can evaluate the robustness against distractions by the CH index (Caliński and Harabasz 1974) with respect to the low-dimensional physical states. The CH index is the ratio of between-clusters dispersion and within-cluster dispersion. A higher value of the CH index indicates that the clusters are more dense and well separated. Specifically, we assign each observation to the cluster whose prototype is closest in the latent space. And we obtain the corresponding physical states from the simulator.

The results in Figure 1 show that Proto-RL struggles to learn robust representations while CBM achieves robustness against distractions. We observe that Proto-RL and CBM almost start from equal clustering quality because of the random initialization. However, as the training continues, CBM consistently improves the clustering quality in terms of states while Proto-RL sticks to a low value of CH-index. A potential reason is that the dynamical information used in Proto-RL to cluster observations contains task-irrelevant

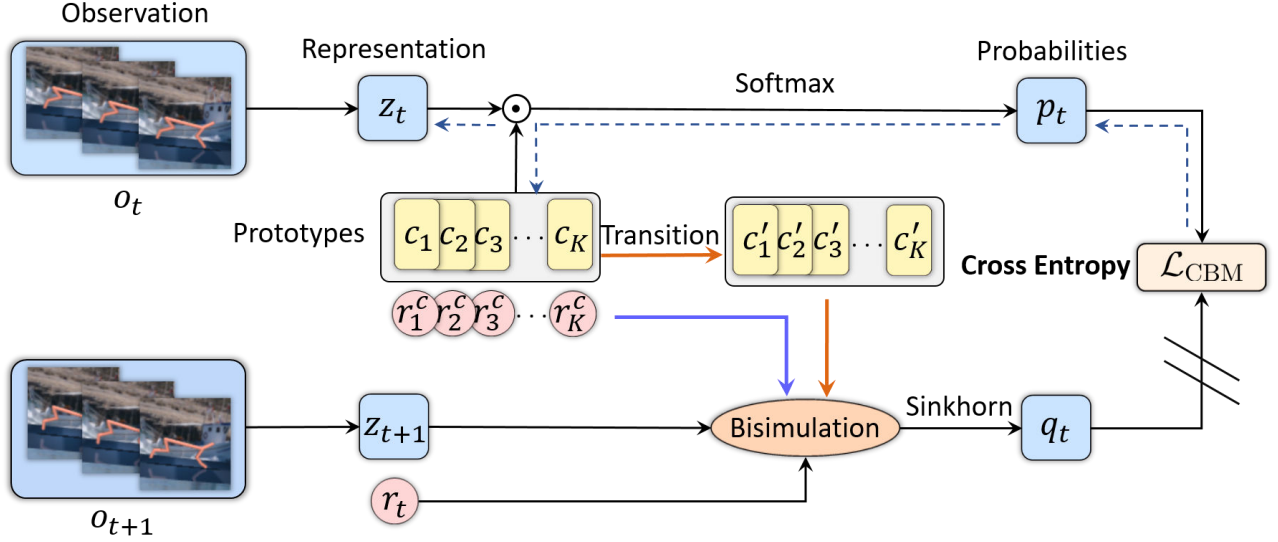


Figure 2: The framework of CBM. The dashed line represents the backpropagation of gradients. CBM enforces consistency between predicted cluster assignments and target cluster assignments. We update prototypes’ rewards and transitions by taking the average of batch rewards and using the learned transition model, respectively.

information about distractions and thus is not task-specific.

Clustering with Bisimulation Metrics

In RL, reward signals provide essential task-specific information. As stated in Section , bisimulation metrics recursively utilize the reward signal and quantify how behaviourally different two observations are. Therefore, we propose to calculate the cluster assignments according to the bisimulation metrics to the prototypes. The following proposition shows that a cluster formed in this manner will contain observations with close expected returns.

Proposition 0.2. *Let V^* be the optimal value function for a given discount factor γ . If $c \geq \gamma$, given the bisimulation metric d and a prototype’s underlying state s_c , then for any two states s_1, s_2 such that $d(s_1, s_c) < \epsilon, d(s_2, s_c) < \epsilon$, we have*

$$|V^*(s_1) - V^*(s_2)| < \frac{2\epsilon}{1-c}.$$

Proof in Appendix. The proposition extends Theorem 5.1 in Ferns, Panangaden, and Precup (2004) to the clustering setting. In CBM, we calculate a bisimulation distance matrix $\mathbf{D} = (d_{ij})_{K \times B}$ from the batch of observation encodings $\{z_i\}_{i=1}^B$ and a set of prototypes $\{c_j\}_{j=1}^K$, where $d_{ij} = d(z_i, c_j)$. Then we employ the Sinkhorn-Knopp clustering procedure to ensure equal partitioning of the prototypes across all embeddings and produce the target codes $\mathbf{Q} = [q_1, \dots, q_B]$.

$$\mathbf{Q} = \text{Diag}(\mathbf{u}) \exp\left(\frac{-\mathbf{D}}{\epsilon}\right) \text{Diag}(\mathbf{v}). \quad (2)$$

Consequently, we calculate soft cluster assignments to the set of prototypes, where we assign a high probability to the prototype that is behaviourally similar to the observation.

Approximating Bisimulation Metrics in CBM

In this part, we introduce the approximation to the bisimulation metrics between observations and prototypes. The bisimulation metrics consist of differences in rewards and transition. We can store the rewards and transitions of the observations in the replay buffer. However, the prototypes are representations in the latent space. Therefore, we cannot directly obtain their corresponding rewards and transitions from the environment.

To tackle this problem, we first define the rewards of the prototypes. We approximate the prototypes’ rewards by taking the weighted average over the rewards in a batch $\mathbf{r} = [r_1, \dots, r_B]^\top$. We estimate the prototype reward by

$$\hat{r}_k^c = \frac{K}{B} \mathbf{w}_k^\top \mathbf{r} \quad (3)$$

where the weight $\mathbf{w}_k = [q_{k1}, \dots, q_{kB}]^\top$ is taken from the target codes. Since the target codes should satisfy $\sum_{j=1}^B q_{kj} = \frac{B}{K}$ after the Sinkhorn-Knopp clustering procedure, we multiply the weighted average by a factor K/B for normalization. At the beginning of the training, we initialize the prototype reward by randomly sampling from the replay buffer. We update the prototype reward at each training step by exponential moving average.

$$r_k^c \leftarrow \beta \hat{r}_k^c + (1 - \beta) r_k^c \quad (4)$$

Secondly, to obtain the transitions of the prototypes, we train a deterministic latent dynamics model \mathcal{P} concurrently. We obtain the prototype transition c'_k from prediction of the model \mathcal{P} given c_k . Since prior work has shown that minimizing a quadratic loss is prone to collapse (Schwarzer et al. 2020; Gelada et al. 2019), we train the latent dynamics model by one-step contrastive predictive coding (CPC) (Oord, Li, and Vinyals 2018).

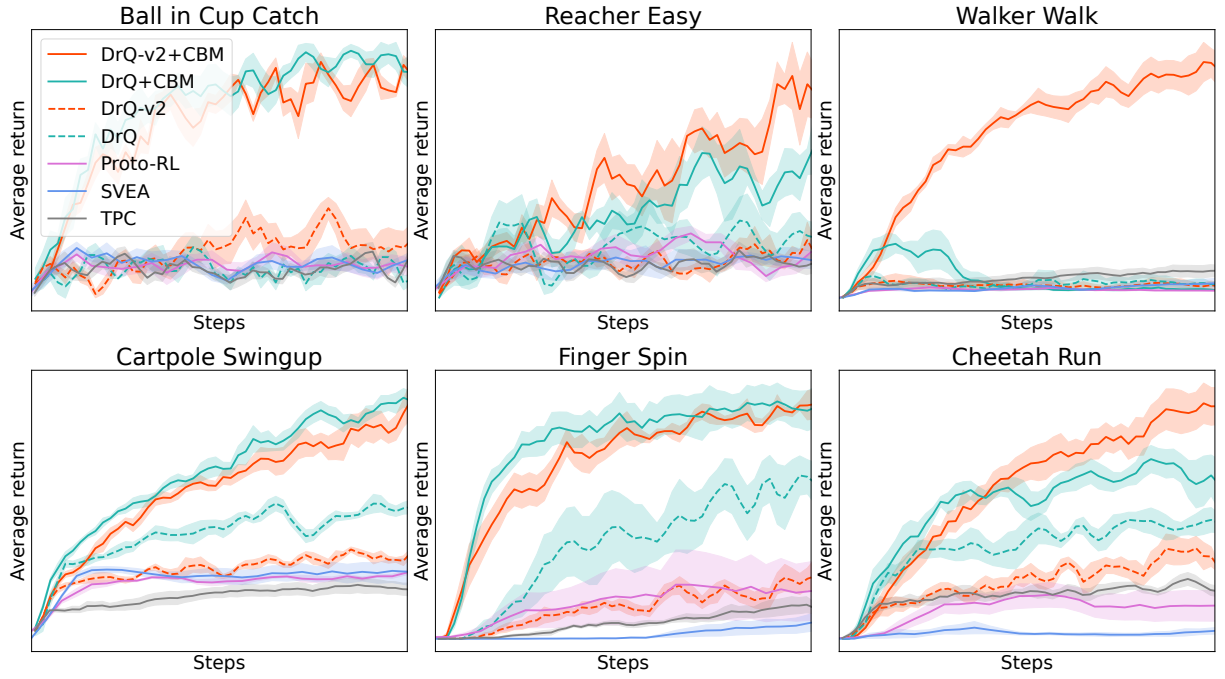


Figure 3: Performance under multiple distractions. The solid curves correspond to the mean and the shaded region to the standard deviation. CBM significantly improves the sample efficiency for both DrQ and DrQ-v2.

Finally, we approximate the bisimulation distance $d(\mathbf{z}_i, \mathbf{c}_j)$ as $\|r_i - r_j^c\| + \|\mathbf{z}_i' - \mathbf{c}_j'\|$ (Castro 2020) using the prototypes’ rewards and transitions, and thus we can compute the distance matrix \mathbf{D} in Equation 2.

Algorithm

In this part, we describe the training process of CBM. Our algorithm repeatedly performs the following steps. Firstly, we sample a batch of data $\{\mathbf{o}_i, \mathbf{a}_i, r_i, \mathbf{o}_i', \}_{i=1}^B$ and map visual observations $\mathbf{o}_i, \mathbf{o}_i'$ into latent states $\mathbf{z}_i, \mathbf{z}_i'$ using the encoder. Then, we predict the cluster assignment \mathbf{p}_i by projecting \mathbf{z}_i into a set of prototypes $\{\mathbf{c}_j\}_{j=1}^K$ and taking the softmax

$$\mathbf{p}_i = \frac{\exp(\frac{1}{\tau} \mathbf{z}_i^\top \mathbf{c}_k)}{\sum_{j=1}^K \exp(\frac{1}{\tau} \mathbf{z}_i^\top \mathbf{c}_j)},$$

where τ is a temperature parameter. Next, we compute the target cluster assignment \mathbf{q}_i by Equation 2. To approximate the bisimulation metrics, we obtain prototypes’ rewards by Equation 3, 4 and transitions by using the latent dynamics model, respectively. Finally, we update the components in CBM. Specifically, we update prototypical representations $\{\mathbf{c}_j\}_{j=1}^K$ by optimizing the cross entropy loss

$$\mathcal{L}_{\text{CBM}} = - \sum_{k=1}^B \mathbf{q}_k \log \mathbf{p}_k.$$

We update the dynamics model by one-step contrastive predicting coding. We allow the gradients from \mathcal{L}_{CBM} and dynamics loss to propagate to the encoder. Note that CBM is compatible with most visual RL algorithms. We illustrate

the framework of CBM in Figure 2, and we provide the pseudocodes of CBM in the Appendix.

Experiments

In this section, we discuss the empirical results of CBM. We combine CBM with two popular algorithms, DrQ (Yarats, Kostrikov, and Fergus 2020) and DrQ-v2 (Yarats et al. 2021a), which improve SAC and TD3 for visual control tasks. To evaluate robustness against distractions, we test all methods on multiple and single distraction settings. Then, we analyze the impact of each component in CBM. Finally, we visualize the embedding space of CBM to demonstrate that CBM achieves task-specific clustering. All experiments report results over five seeds.

Control Tasks We evaluate all agents on a challenging benchmark for vision-based control, Distracting Control Suite (DCS) (Stone et al. 2021). DCS extends DeepMind Control (DMC) (Tassa et al. 2018) with three kinds of visual distractions (background, color and camera pose). In multiple distractions settings, we use DCS ”easy” setting, where the background images are sampled from 4 videos, and the scale for camera and color distraction is 0.1.

Baseline Methods Besides DrQ and DrQ-v2, we compare CBM with three state-of-the-art methods including Proto-RL (Yarats et al. 2021b), which learn representations by grouping observations with dynamics information, SVEA (Hansen, Su, and Wang 2021), which regularizes the representations by utilizing strong augmentation like random convolution, TPC (Nguyen et al. 2021), which contrastively learns a latent model for planning.

Architectures When combining CBM with DrQ and DrQ-




Figure 4: Evaluation results after training with single distraction for 5e5 environment steps. CBM achieves state-of-the-art performance under the camera and background distraction.

v2, we use a shared convolutional encoder followed with a linear projector with 50 outputs. CBM learns $K=128$ prototypes, each parameterized as a 50-dimensional vector.

Learning Control with Multiple Distractions

In this part, we evaluate methods under multiple distractions in terms of sample efficiency and asymptotic performance. Stone et al. (2021) find that multiple distractors have a compounding effect: combined distractors degrade performance more than individually. The results in Figure 3 show that CBM significantly improves the sample efficiency and asymptotic performance of DrQ and DrQ-v2. Moreover, CBM outperforms all baselines in six environments. In some tasks like finger spin or cheetah run, DrQ and DrQ-v2 are able to learn reasonable behaviour, but they require more samples to train the agent. The poor performance of Proto-RL, SVEA, and TPC indicate clustering with dynamics information, utilizing strong augmentation and learning latent dynamics contrastively struggle to overcome the challenging compounding effect of the multiple distractions.

Learning Control with Single Distractions

In this part, we train agents with each single distraction in different magnitude and evaluate them at the end of training. We present the results in Figure 4. CBM significantly improves the performance of both DrQ and DrQ-v2, especially under strong distractions. We observe that Proto-RL performs fairly well when the distraction is mild, while its performance declines dramatically as the distraction scale increases. This result indicates that clustering with dynamics information can help to handle a mild distraction but performs poorly under stronger distraction. TPC only achieves about 600 average returns in the standard setting without distractions. We argue that its poor performance partly comes from the low sample efficiency of model training. SVEA is quite robust to color distraction. However, it suffers from




Figure 5: CBM works better than training a latent model alone or contrastively utilizing the bisimulation metrics.

severe degradation under the camera and background distraction. A potential reason is that the particular augmentation choice in SVEA (random convolution) can simulate the color variations well but can not approximate the camera and background variations. Therefore, we conclude that the data augmentation method can perform well with proper augmentation choice according to the domain knowledge but perform much worse than our method CBM when such domain knowledge is missing.

Ablation Study

This part analyzes the impact of clustering-based objectives, the transition model, and the number of prototypes. We conduct the experiments on cartpole swingup and walker walk in the multiple distraction setting.

Comparison with contrastive objectives We compare the clustering-based objective in CBM with a contrastive-based objective. Concretely, for every observation encoding, we select the positive sample as the closest observations from the sampled batch with respect to bisimulation metrics. We keep all hyperparameters fixed and approximate the bisimulation metrics in the same way as in CBM. We optimize the

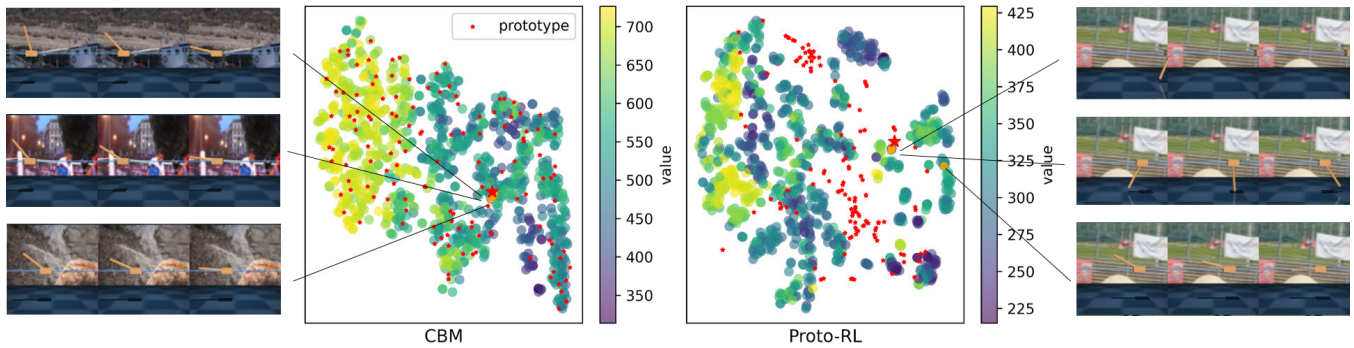


Figure 6: t-SNE visualization of representations learned by CBM and Proto-RL. The color represents the state value, the red points represent the prototypes. We show an example prototype and three encodings that have biggest assignment probability to it. The cluster formed by CBM is more behaviorally similar and the prototypes disperse more uniformly in the latent space.




Figure 7: Performance of CBM with various number of prototypes. Zero prototypes represent the implementation that only train the transition model.

encoder by the InfoNCE loss (Oord, Li, and Vinyals 2018) and the dynamics loss. As shown in Figure 5, we observe that the contrastive-based method performs poorly on the multiple distraction setting, while CBM exploits the bisimulation metrics better and is much more sample-efficient. A potential reason is that the contrastive-based method is sensitive to negative samples and requires a large batch size.

Impact of transition model To approximate the bisimulation metrics, we concurrently train a latent model by one-step CPC. Previous work has shown that learning latent dynamics can help discriminate task-relevant information (Nguyen et al. 2021). To study the separate effect of training a latent model in our method, we update the encoder by training the transition model alone. From Figure 5, we investigate that training the latent model improves the performance of DrQ-v2. Adding the clustering loss \mathcal{L}_{CBM} can further improve the sample efficiency and asymptotic performance, which proves the effectiveness of the clustering.

Impact of number of prototypes We test CBM with a different number of prototypes ranging from 0 to 512. From

Figure 7, we find that even a small number of prototypes, such as eight, can significantly improve performance. We also observe that further increasing K to 128 can make a slight improvement, while a larger K may not always bring further improvement. These results suggest that the number of prototypes has little impact as long as it is "enough". Since using more prototypes increases the computational cost, we recommend training CBM with 128 prototypes.

Visualization Analysis

In this part, we visualize the embeddings trained with background distraction on cartpole swingup. We apply t-SNE algorithms on randomly selected observation encodings and the set of prototypical representations. Figure 6 shows that CBM maps observations with similar values to neighboring regions while Proto-RL does not. This indicates that CBM effectively captures task-relevant information.

To visualize the clustering results, we select a prototype and present the three nearest observations in terms of cosine similarity. We observe that observations obtained from the embedding of CBM exhibit similar behavior, while those obtained from Proto-RL behave quite differently. This observation demonstrates that CBM achieves task-specific clustering. We put more clustering results of CBM from different environments in the Appendix.

Conclusion

Learning robust representations is critical for sample-efficient reinforcement learning from images. In this paper, we propose CBM, a novel clustering-based method that learns robust representations by grouping visual observations with bisimulation metrics. By incorporating the properties of RL tasks, CBM effectively captures task-relevant information and filters out task-irrelevant information. Experiments demonstrate that CBM significantly improves the sample efficiency of popular visual RL algorithms and achieves state-of-the-art performance on both multiple and single distraction settings.

Acknowledgments

We would like to thank all the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundations of China grants U19B2026, U19B2044, 61836006, and 62021001, and the Fundamental Research Funds for the Central Universities grant WK3490000004.

References

- Agarwal, R.; Machado, M. C.; Castro, P. S.; and Bellemare, M. G. 2021. Contrastive Behavioral Similarity Embeddings for Generalization in Reinforcement Learning. In *International Conference on Learning Representations*.
- Asano, Y. M.; Rupprecht, C.; and Vedaldi, A. 2020. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations (ICLR)*.
- Caliński, T.; and Harabasz, J. 1974. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1): 1–27.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep Clustering for Unsupervised Learning of Visual Features. *European conference on computer vision*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33: 9912–9924.
- 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, 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.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Deng, F.; Jang, I.; and Ahn, S. 2021. DreamerPro: Reconstruction-Free Model-Based Reinforcement Learning with Prototypical Representations. *arXiv preprint arXiv:2110.14565*.
- Fan, L.; Wang, G.; Huang, D.-A.; Yu, Z.; Fei-Fei, L.; Zhu, Y.; and Anandkumar, A. 2021. SECANT: Self-Expert Cloning for Zero-Shot Generalization of Visual Policies. *International Conference on Machine Learning*.
- Ferns, N.; Panangaden, P.; and Precup, D. 2004. Metrics for Finite Markov Decision Processes. In *UAI*, volume 4, 162–169.
- Ferns, N.; Panangaden, P.; and Precup, D. 2011. Bisimulation metrics for continuous Markov decision processes. *SIAM Journal on Computing*, 40(6): 1662–1714.
- Fu, X.; Yang, G.; Agrawal, P.; and Jaakkola, T. 2021. Learning task informed abstractions. In *International Conference on Machine Learning*.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*.
- 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*.
- 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 *International conference on machine learning*.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2019. Dream to control: Learning behaviors by latent imagination. *arXiv preprint*.
- Hansen, N.; Su, H.; and Wang, X. 2021. Stabilizing Deep Q-Learning with ConvNets and Vision Transformers under Data Augmentation. In *Conference on Neural Information Processing Systems*.
- Hansen, N.; and Wang, X. 2020. Generalization in reinforcement learning by soft data augmentation. *arXiv preprint arXiv:2011.13389*.
- 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*.
- Kalashnikov, D.; Irpan, A.; Pastor, P.; Ibarz, J.; Herzog, A.; Jang, E.; Quillen, D.; Holly, E.; Kalakrishnan, M.; Vanhoucke, V.; et al. 2018. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, 651–673. PMLR.
- Lample, G.; and Chaplot, D. S. 2017. Playing FPS games with deep reinforcement learning. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Laskin, M.; Lee, K.; Stooke, A.; Pinto, L.; Abbeel, P.; and Srinivas, A. 2020. Reinforcement learning with augmented data. *Advances in Neural Information Processing Systems*.
- Lee, A. X.; Nagabandi, A.; Abbeel, P.; and Levine, S. 2020a. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *Advances in Neural Information Processing Systems*, 33: 741–752.
- Lee, K.; Lee, K.; Shin, J.; and Lee, H. 2019. Network Randomization: A Simple Technique for Generalization in Deep Reinforcement Learning. In *International Conference on Learning Representations*.
- Lee, K.-H.; Fischer, I.; Liu, A.; Guo, Y.; Lee, H.; Canny, J.; and Guadarrama, S. 2020b. Predictive information accelerates learning in rl. *Advances in Neural Information Processing Systems*, 33: 11890–11901.
- Levine, S.; Finn, C.; Darrell, T.; and Abbeel, P. 2016. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1): 1334–1373.
- Mazouze, B.; Tachet des Combes, R.; Doan, T. L.; Bachman, P.; and Hjelm, R. D. 2020. Deep reinforcement and infomax learning. *Advances in Neural Information Processing Systems*, 33: 3686–3698.

- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Nguyen, T. D.; Shu, R.; Pham, T.; Bui, H.; and Ermon, S. 2021. Temporal Predictive Coding For Model-Based Planning In Latent Space. In *International Conference on Machine Learning*.
- Okada, M.; and Taniguchi, T. 2021. Dreaming: Model-based reinforcement learning by latent imagination without reconstruction. In *International Conference on Robotics and Automation*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint*.
- Raileanu, R.; Goldstein, M.; Yarats, D.; Kostrikov, I.; and Fergus, R. 2020. Automatic data augmentation for generalization in deep reinforcement learning. *arXiv preprint*.
- Schwarzer, M.; Anand, A.; Goel, R.; Hjelm, R. D.; Courville, A.; and Bachman, P. 2020. Data-Efficient Reinforcement Learning with Self-Predictive Representations. In *International Conference on Learning Representations*.
- Srinivas, A.; Laskin, M.; and Abbeel, P. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv preprint arXiv:2004.04136*.
- Stone, A.; Ramirez, O.; Konolige, K.; and Jonschkowski, R. 2021. The Distracting Control Suite—A Challenging Benchmark for Reinforcement Learning from Pixels. *arXiv preprint*.
- Stooke, A.; Lee, K.; Abbeel, P.; and Laskin, M. 2021. Decoupling representation learning from reinforcement learning. In *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*.
- Villani, C. 2021. *Topics in optimal transportation*, volume 58. American Mathematical Soc.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, 1096–1103.
- Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33: 6256–6268.
- Yarats, D.; Fergus, R.; Lazaric, A.; and Pinto, L. 2021a. Mastering Visual Continuous Control: Improved Data-Augmented Reinforcement Learning. In *International Conference on Learning Representations*.
- Yarats, D.; Fergus, R.; Lazaric, A.; and Pinto, L. 2021b. Reinforcement learning with prototypical representations. In *International Conference on Machine Learning*, 11920–11931. 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. 2021c. Improving Sample Efficiency in Model-Free Reinforcement Learning from Images. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Yu, T.; Lan, C.; Zeng, W.; Feng, M.; and Chen, Z. 2021. PlayVirtual: Augmenting Cycle-Consistent Virtual Trajectories for Reinforcement Learning. *Advances in Neural Information Processing Systems*.
- Zhang, A.; McAllister, R.; Calandra, R.; Gal, Y.; and Levine, S. 2021. Learning invariant representations for reinforcement learning without reconstruction. *International Conference on Learning Representations*.
- Zhu, J.; Xia, Y.; Wu, L.; Deng, J.; Zhou, W.; Qin, T.; and Li, H. 2020. Masked contrastive representation learning for reinforcement learning. *arXiv preprint*.