

Value-Consistent Representation Learning for Data-Efficient Reinforcement Learning

Yang Yue^{1,2*}, Bingyi Kang^{2 †}, Zhongwen Xu², Gao Huang¹, Shuicheng Yan²

¹Department of Automation, BNRist, Tsinghua University

²Sea AI Lab

yueyang22f@gmail.com, {kangby, xuzw, yansc}@sea.com, gaohuang@tsinghua.edu.cn

Abstract

Deep reinforcement learning (RL) algorithms suffer severe performance degradation when the interaction data are scarce, which limits their real-world application. Recently, visual representation learning has been shown to be effective and promising for boosting sample efficiency in RL. These methods usually rely on contrastive learning and data augmentation to train a transition model for *state* prediction, which is different from how the model is used in RL—performing *value*-based planning. Accordingly, the learned representation by these visual methods may be good for recognition but not optimal for estimating state value and solving the decision problem. To address this issue, we propose a novel method, called *value-consistent representation learning* (VCR), to learn representations that are directly related to decision-making. More specifically, VCR trains a model to predict the future state (also referred to as the “imagined state”) based on the current one and a sequence of actions. Instead of aligning this imagined state with a real state returned by the environment, VCR applies a Q -value head on both states and obtains two distributions of action values. Then a distance is computed and minimized to force the imagined state to produce a similar action value prediction as that by the real state. We develop two implementations of the above idea for the discrete and continuous action spaces respectively. We conduct experiments on Atari 100K and DeepMind Control Suite benchmarks to validate their effectiveness in improving sample efficiency. It has been demonstrated that our methods achieve new state-of-the-art performance for search-free RL algorithms.

1 Introduction

An important research direction in Deep Reinforcement Learning (RL) is to improve data efficiency, which is much demanded by the wide application of deep RL techniques in real-world scenarios. With the state-of-the-art RL algorithms, simple tasks such as video games in Arcade Learning Environment (Bellemare et al. 2013) require billions of frames to achieve human-level performance (Badia et al. 2020). In real-world applications, such as robot controllers and self-driving systems, it is impractical to obtain such a huge amount of

*This work was done when Yang Yue was an intern at Sea AI Lab.

†Corresponding Author.

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

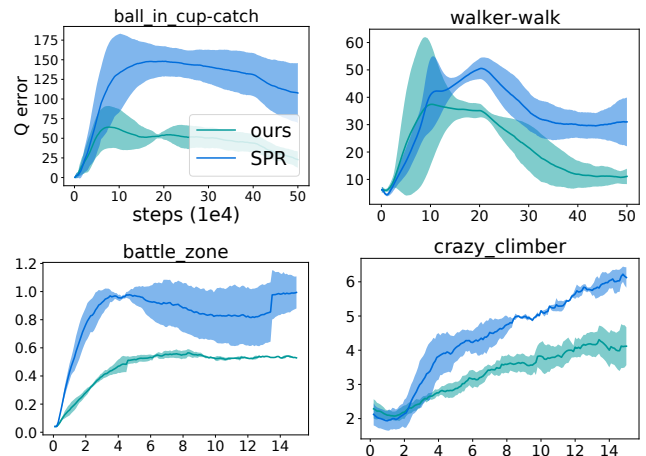


Figure 1: Q error of imagined states. Given a dataset of real trajectories, we have a imagined state for each of the state in the dataset by shooting forward for several steps with the dynamic model. The curves represents the average error between the value estimation from imagined states and the ground-truth values along trajectories. The first row shows Q errors of SPR (blue) and VCR (ours, green) on 2 DeepMind Control environments. The second row shows Q error on a subset of Atari 100K games. The shaded area represents standard deviation. Details are described in Sec. 4.1.

interaction due to the costly data collection process. To enable deep RL to go beyond virtual games and simulators, researchers explore the data efficiency issue from various perspectives, including model-based RL (Hafner et al. 2019, 2020; Kaiser et al. 2020), auxiliary tasks (Jaderberg et al. 2016; Yarats et al. 2021; Laskin, Srinivas, and Abbeel 2020; Schwarzer et al. 2021), data augmentation (Yarats, Kostrikov, and Fergus 2021; Laskin et al. 2020), etc. A majority of the works borrow ideas from the wider deep learning community to create additional training signals that accelerate the agent training process. Most of these techniques are not specifically tailored for RL problems and some heavily rely on extracting high-quality visual features.

Recently, Self-Predictive Representation (SPR; Schwarzer et al. 2021) introduces contrastive learning into transition

model learning. Specifically, SPR aims at learning an embedding space in which the agent can predict future state embeddings. Despite its effectiveness, it solely focuses on learning discriminative state features, ignoring the fact that some information in the raw pixels is unnecessary or even distracting for decision-making. This part of information would be encoded into representation by visual contrastive learning. Moreover, it is possible in practice that two visually similar states result in significantly different returns. In other words, the learned representation in this visual way might be good for recognition but not optimal for estimating state value and solving the decision problem. For example, it’s common in video games that state value highly depends on a very small number of pixels in a state (e.g., a high-reward treasure), while the predictive representation by SPR focus on global visual features. As evidenced by Fig. 1, the predicted values based on imagined states of SPR are consistently deviated from the true values (by Monte-Carlo estimation) across multiple environments.

Similar observations have been made on the Value Equivalence (VE) principle (Grimm et al. 2020) for model-based RL, which advocates that a model should be able to generate the same Bellman updates as the real environment rather than directly modeling state-to-state transitions. The success of some state-of-the-art algorithms such as MuZero (Schrittwieser et al. 2020), Value Prediction Network (Oh, Singh, and Lee 2017), and Predictron (Silver et al. 2017) can be attributed to this principle. However, value equivalence is only studied with state value functions $V(S)$, and is often coupled with search-based algorithms, which makes it non-trivial to apply VE for value-based RL.

In this paper, we develop a value-consistent metric for action-values (*i.e.*, Q -values) and propose a novel method called *Value-Consistent Representation Learning* (VCR) to boost sample efficiency for action-value-based RL methods. In contrast to existing data-efficient RL ideas, VCR is based on RL semantics other than losses constructed purely upon input states. Specifically, we introduce a dynamics model to predict the next states in a latent representation space induced by a state encoder. When a real trajectory from the environment is provided, our dynamics model can roll out an imagined trajectory from the initial state by taking the same sequence of actions. Then, for each of the real and imagined state pairs, we obtain two Q -value distributions (over all available actions). A value-consistent loss function is applied to align these two distributions. VCR can be viewed as a clean implementation of value equivalence with action-values so that it is highly extendable in a significant family of RL algorithms. The idea is simple yet effective and it can be easily integrated into any value-based RL algorithm. We demonstrate the effectiveness of the idea with two concrete implementations based on Rainbow DQN (Hessel et al. 2018) and Soft Actor-Critic (SAC) (Haarnoja et al. 2018).

We conduct experiments to validate the effectiveness of VCR on two benchmarks: Atari 100K (Kaiser et al. 2020) (discrete action) and DeepMind Control Suite (Tassa et al. 2018) (continuous action). Despite its simplicity, the results show our method can boost sample efficiency significantly and achieve new state-of-the-art in search-free methods.

2 Related Work

2.1 Data-Efficient Reinforcement Learning

Reinforcement learning algorithms suffer severe performance degradation when only a limited number of interactions are available. Various methods are trying to tackle this problem, by either model-based (Kaiser et al. 2020), model-free (van Hasselt, Hessel, and Aslanides 2019; Kielak 2020; Yarats, Kostrikov, and Fergus 2021), or non-parametric learning (Long, Blair, and van Hoof 2022). For example, SimPLe (Kaiser et al. 2020) utilizes a world model learned with collected data to generate imagined trajectories to train the agent. Later, Data-Efficient Rainbow (van Hasselt, Hessel, and Aslanides 2019) and OTRainbow (Kielak 2020) show that Rainbow DQN with hyperparameter tuning can be a strong baseline for the low-data regime by simply increasing the number of steps in multi-step return and allowing more frequent parameter update.

Recently, leveraging computer vision techniques is drawing increasingly more attention from the community to boost representation learning in RL. DrQ (Yarats, Kostrikov, and Fergus 2021) and rQdia (Lerman, Bi, and Xu 2021) makes a successful attempt by introducing image augmentation into RL tasks, while Yarats *et al.* (Yarats et al. 2021) use image reconstruction as an auxiliary loss function. Inspired by the remarkable success of contrastive learning for representation learning (Chen et al. 2020; Grill et al. 2020), the contrastive loss has been integrated into RL as an effective component. For example, CURL (Laskin, Srinivas, and Abbeel 2020) forces different augmentations of the same state to produce similar embeddings and different states to generate dissimilar embeddings, with the contrastive loss jointly optimized with an RL loss.

Recently, SPR (Schwarzer et al. 2021) and KSL (McInroe, Schäfer, and Albrecht 2021) make a sophisticated design by employing contrastive loss in transition model learning. In this way, an agent can learn representations that are predictable when the previous state and action are given. As a follow-up, PlayVirtual (Yu et al. 2021) introduces a backward prediction model that enables the agent to imagine forward and backward to form a cycle. Therefore, arbitrary actions can be taken in imagination to compute cycle-consistent loss. EfficientZero (Ye et al. 2021) introduces SPR loss into MuZero (Schrittwieser et al. 2020), achieving super-human performance on the Atari 100K for the first time.

These successes demonstrate that self-predictive representation learning is indeed a promising way to improve sample efficiency. We also choose to base our method on SPR. However, these methods only focus on making accurate state predictions, ignoring that value prediction is key to decision-making problems. In this work, we propose value-consistent representation learning and show its importance for decision-making.

2.2 Value Equivalence Principle

Some works in model-based RL have proposed a high-level idea of learning the transition model in terms of the value space. Value-Aware Model Learning (VAML) (Farahmand, Barreto, and Nikovski 2017; Farahmand 2018) incorporates

the knowledge of value function into optimization to learn the probabilistic transition model in model-based reinforcement learning. Grimm *et al.* (Grimm et al. 2020) come up with the value equivalence principle which forces the learned transition model to have the same Bellman operator updates with the real environment conditioned on a set of value functions and policies. Empirically successful works like VPN (Oh, Singh, and Lee 2017), VIN (Tamar et al. 2016), and MuZero (Schrittwieser et al. 2020) can be viewed as the instances that follow the value equivalence principle. Self-consistent models and values (Farquhar et al. 2021) boost the performance of Muesli agent (Hessel et al. 2021) by encouraging a value function to satisfy the Bellman equation under a learned transition model and a learned policy.

Although based on similar motivation as the value equivalence principle, our core idea greatly differs from the above works. The above works lie in the category of planning-based RL and leverage value regularization or value equivalence to train a dynamics model for planning. However, our method, a search-free algorithm, harnesses value consistency to learn a transition model for better state representation. Besides, our method employs the reward from the real environment to estimate the target value, while those planning algorithms use predicted rewards with an extra reward predictor.

3 Method

We are especially interested in RL in a low-data regime, where only a limited number of interactions are allowed. In this section, we detail our value-consistent representation (VCR) learning method step by step. First, we prepare the readers with certain preliminaries required by VCR. Then, we introduce the intuition and the overall framework for VCR. In the end, we provide two implementations of VCR for both discrete and continuous action settings.

3.1 Preliminaries: Value-Based Reinforcement Learning

Reinforcement Learning (RL) addresses the problem of sequential decision-making, which is usually formulated with a Markov Decision Process (MDP). A typical MDP is represented with a tuple $\langle \mathcal{S}, \mathcal{A}, T, r, \gamma \rangle$. Here \mathcal{S} is a finite set of states; \mathcal{A} is the action space; $T(s, a, s') = P(s'|s, a)$ is the dynamics function describing the probability of transitioning from a state s to s' after taking an action a ; $r(s, a, s')$ and $\gamma \in (0, 1]$ are the reward function and the discount factor respectively. The fundamental goal of RL is to learn an agent maximizing the discounted cumulative reward (*i.e.*, return) $G_t = \sum_{\tau=0}^{\infty} \gamma^\tau r_{t+\tau+1}$ at any time step t . The behavior of an agent is denoted by a policy $\pi(a|s)$ mapping from states to actions, while the action-value function $Q_\pi(s, a) = \mathbb{E}_\pi [G_t | s_t = s, a_t = a]$ predicts the expected return if the agent takes an action a for state s at the time step t , following the policy π .

There are various ways to learn the optimal policy π^* . In this paper, we focus on value-based RL algorithms that are rooted in Q -learning. Q -learning performs approximate dynamic programming by following the Bellman equation (Sutton and Barto 2018). Deep Q Network (DQN) (Mnih et al. 2015) scales Q -learning to large state (*e.g.*, visual inputs)

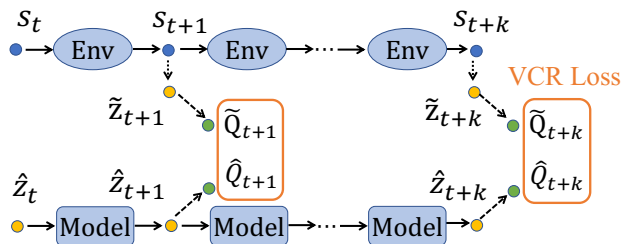


Figure 2: Pipeline of Value-Consistent Representation Learning. The agent interacts with the environment to produce a sequence of states $\{s_{t+k}\}_{k=1:K}$. The sequence is encoded into a latent embedding space, denoted as $\{\tilde{z}_{t+k}\}_{k=1:K}$. Then the parametric transition model predicts the sequence of latent states $\{\hat{z}_{t+k}\}_{k=1:K}$ conditioned on the start point \hat{z}_t and the action sequence, where \hat{z}_t is the latent embedding from s_t . Value-Consistent Representation Learning enforces the value prediction from imagined states to be consistent with the values of a real state from the environment. SPR loss is also utilized to help stabilize the training at the early stage. For simplicity, SPR loss and action inputs to the environment and the transition model are omitted in the figure.

by utilizing neural networks to encode states and generate Q values. Additionally, experience replay and a separated target network are used to stabilize the training of DQN. The overall objective for DQN is given by

$$\mathcal{J}(\theta) = \mathbb{E}_{(s, a, s', r) \sim \mathcal{D}} \left[\left(r + \gamma \max_{a'} Q_{\bar{\theta}}(s', a') - Q_{\theta}(s, a) \right)^2 \right], \quad (1)$$

where Q_{θ} is the Q network parameterized with θ , $Q_{\bar{\theta}}$ is the target Q network, and \mathcal{D} represents the replay buffer which stores the experience tuples. Let $\bar{G}_t^{(n)} = \sum_{\tau=1}^n \gamma^{\tau-1} r_{t+\tau} + \gamma^n \max_a Q_{\bar{\theta}}(s_{t+n}, a)$ be the n -step value target. Note that when $n = 1$, it reduces to the estimator used in Eqn. (1), but $n > 1$ is also commonly used to get a better estimation (Sutton and Barto 2018; Hessel et al. 2018).

As the main goal of our work is to boost the sample efficiency of RL algorithms, without loss of generality, we build our method upon two best-performing algorithm variants, *i.e.*, Rainbow DQN for discrete action domains and Soft Actor-Critic (SAC) for continuous action domains. For more details, please refer to the Appendix or their papers. The appendix can be found in the paper’s arXiv version. Though empirical results are only shown with these two algorithms, our method is general enough to be easily integrated into any other value-based RL algorithms.

3.2 Value-Consistent Representation Learning

Various methods have shown that representation learning can boost the sample efficiency of RL. However, all these methods approach the problem from a computer vision perspective, *i.e.*, encouraging similar states to generate similar embeddings and forcing different states to be discriminative. Despite their effectiveness, visual recognition is not always directly related to decision-making. To alleviate this issue, we base our method on the assumption (Schwarzer et al. 2021;

Yu et al. 2021) that a good representation for RL should be able to predict the resulting state following a sequence of actions. Instead of aligning the predictions in the embedding space, our key idea is the value prediction from an *imagined state* should be consistent with the values of a *real state* from the environment, as illustrated in Fig. 2. Thus, the method is termed Value-Consistent Representation Learning (VCR).

State Prediction with Transition Model. Considering a one-step interaction between an agent and the environment: (s_t, a_t, s_{t+1}) , s_{t+1} can be determined by the transition T in the underlying MDP, given (s_t, a_t) . Similar to SPR and PlayVirtual, we introduce a parametric transition model $h(\cdot, \cdot)$ to mimic the behavior of T in a latent embedding space. More specifically, a (convolutional) neural encoder $f(\cdot)$ is used to encode a pixel-based observation/state s_t into a latent representation $z_t = f(s_t)$. Then $h(\cdot, \cdot)$ operates by $\hat{z}_{t+1} = h(z_t, a_t)$. As shown in Fig. 2, based on the current latent state z_t , following a sequence of future K actions $a_{t:t+K-1}$, a sequence of state predictions $\hat{z}_{t+1:t+K}$ is obtained by applying $h(\cdot, \cdot)$ recursively:

$$\begin{aligned} \hat{z}_t &= z_t = f(s_t) \\ \hat{z}_{t+k+1} &= h(\hat{z}_{t+k}, a_{t+k}), \quad k = 0, 1, \dots, K-1. \end{aligned} \quad (2)$$

The transition model $h(\cdot, \cdot)$ is usually optimized to minimize the prediction error between $\hat{z}_{t+1:t+K}$ and $\tilde{z}_{t+1:t+K}$, a sequence of latent features extracted directly from the raw observations / states with $\tilde{z}_{t+k} = f_T(s_{t+k})$. For example, the self-predictive representations (SPR) (Schwarzer et al. 2021) method utilizes a cosine similarity for prediction error, leading to the following objective:

$$\mathcal{L}_{\text{SPR}} = - \sum_{k=1}^K \left(\frac{\hat{z}_{t+k}}{\|\hat{z}_{t+k}\|_2} \right)^\top \left(\frac{\tilde{z}_{t+k}}{\|\tilde{z}_{t+k}\|_2} \right), \quad (3)$$

where \tilde{z} is often referred to as the target embedding and generated from a target encoder $f_T(\cdot)$, which is a stop-gradient version of the online encoder $f(\cdot)$.

Value Consistency. Intuitively, for a good encoder and a reasonable transition model in RL algorithms, the predictive representations should contain abundant information such that precise value estimation can be made by feeding the latent embedding into a value head. With a slight abuse of notation, we denote the Q -value predictions by a value head for the embedding-action pair (z_t, a_t) as $Q(z_t, a_t)$. For simplicity, we use $Q(z_t, \cdot)$ to denote the action values from the embedding z_t for all possible actions. In this way, the value predictions for imagined states $\hat{z}_{t+1:t+K}$ and target states $\tilde{z}_{t+1:t+K}$ can be written as $\{Q(\hat{z}_{t+k}, \cdot)\}_{k=1:K}$ and $\{Q_T(\tilde{z}_{t+k}, \cdot)\}_{k=1:K}$ respectively. Then our value-consistent representation learning loss is obtained by minimizing the distance between these two value predictions, described by:

$$\mathcal{L}_{\text{VCR}} = \sum_{k=1}^K d_{\text{VCR}}(Q(\hat{z}_{t+k}, \cdot), Q_T(\tilde{z}_{t+k}, \cdot)), \quad (4)$$

where d_{VCR} is a distance metric for action-values to be detailed below, and $Q_T(\tilde{z}_{t+k}, \cdot)$ is the target value prediction generated with a target head Q_T .

Overall Objective. At the early stage of training, the value estimation is not accurate and therefore sole VCR may not provide good supervision signals for training the dynamics model. SPR loss is introduced to help stabilize the training of the dynamics model. Now, we are ready to present our overall training objective as below:

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{J}(\theta) + \lambda_{\text{SPR}} \mathcal{L}_{\text{SPR}}(\theta) + \lambda_{\text{VCR}} \mathcal{L}_{\text{VCR}}(\theta), \quad (5)$$

where θ denotes all the model parameters used for computing the above loss function, and λ_{SPR} , λ_{VCR} are the hyperparameters to weight different losses. Specific values for hyperparameters are at the apped.

Value-Consistent Distance Metric. Here we develop our distance metric $d_{\text{VCR}}(\cdot, \cdot)$ and provide two different implementations for both discrete action and continuous settings. It seems easy to come up with an idea for $d_{\text{VCR}}(\cdot, \cdot)$. For example, one can simply apply mean-squared loss to the imagined q-values and the target q-values over all possible actions: $d_{\text{MSE}} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} [Q(\hat{z}_t, a) - Q_T(\tilde{z}_t, a)]^2$. Actually, this can not work through because the target action-value $Q(\tilde{z}_t, a_t)$ for the real action a_t keeps evolving as the training proceeds. Based on $\mathcal{J}(\theta)$ in Eqn. (1), at each iteration, the Q value $Q(\tilde{z}_t, a_t)$ (where $\tilde{z}_t = f(s_t)$) for a real state-action pair (s_t, a_t) will be updated towards a n -step target estimation $\bar{G}_t^{(n)}$. If we still use d_{MSE} , it means that we are optimizing $Q(\tilde{z}_t, a_t)$ towards a sub-optimal point. Based on the above observation, we propose the following distance function:

$$\begin{aligned} d_{\text{VCR}} &= \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} [Q(\hat{z}_t, a) - \bar{Q}(\tilde{z}_t, a)]^2, \\ \text{where } \bar{Q}(\tilde{z}_t, a) &= \begin{cases} \bar{G}_t^{(n)} & \text{if } a = a_t, \\ Q_T(\tilde{z}_t, a) & \text{if } a \neq a_t. \end{cases} \end{aligned} \quad (6)$$

With above equations, we let the value prediction of a real state-action pair (s_t, a_t) align with a n -step target estimation $\bar{G}_t^{(n)}$, while other pairs align with the corresponding target action-value $Q_T(\tilde{z}_t, a)$. To avoid the trivial solution problem(Grill et al. 2020), we stop the gradient of the target value $Q_T(\tilde{z}_t, a)$. Additionally, the varying rewards introduced by $\bar{G}_t^{(n)}$ also disqualify the constant output as a solution.

Discrete Action Implementation. For discrete actions, the Q network or head is implemented to directly generate $|\mathcal{A}|$ outputs representing the Q values for the corresponding actions. We can simply enumerate all of them to calculate the above distance in Eqn. (6). The above $\bar{G}_t^{(n)}$ is given by $\bar{G}_t^{(n)} = \sum_{\tau=1}^n \gamma^{\tau-1} r_{t+\tau} + \gamma^n \max_a Q_{\bar{\theta}}(s_{t+n}, a)$, which adopt the same form with n -step estimation of Q -learning to mitigate possible gradient conflict in multi-task learning(Sener and Koltun 2018; Yu et al. 2020; Jean, Firat, and Johnson 2019). As our method is based on Rainbow, each value is divided into bins to build a distribution.

Continuous Action Implementation. For continuous actions, the Q network has a different implementation, which usually takes both the state s and the action a and then outputs a scalar as the Q value $Q(s, a)$. Note that SAC is chosen to be our baseline algorithm, in which the target Q value is estimated by one-step return. For the case that

action a is real action a_t , the above $\bar{G}_t^{(n)}$ is calculated by $\bar{G}_t^{(n)} = r_t + \gamma Q_{\phi_{\text{target}}}(s_{t+1}, a_{\theta}(s_{t+1}))$. We randomly sample a few actions as actions that are not equal to a_t , which together with a_t constitute the set \mathcal{A} . Considering soft Q values are used in SAC, we also employ soft q values when calculating the value-consistent loss.

4 Experiments

In this section, we first investigate the quality of dynamic models given by SPR in terms of value estimation, and show that VCR is able to improve it substantially. Further, we empirically evaluate the performance of our VCR on boosting data efficiency in Atari 100K (Bellemare et al. 2013; Kaiser et al. 2020) and DeepMind Control Suite (Tassa et al. 2018). We then conduct ablation studies to analyze the important components in our method.

4.1 Value of Imagined States

Intuitively, it’s a good property for a reasonable encoder and transition model that the predicted representations contain abundant information such that precise value estimation can be made. We measure the absolute difference between the values of imagined state \hat{z}_{t+k} and true values, where \hat{z}_{t+k} is given by Eqn. 2. Here we use the Monte-Carlo return as the true value. For a complete evaluation trajectory of length T , the average error is as:

$$\frac{1}{TK} \sum_{t=1}^T \sum_{k=1}^K |Q(\hat{z}_{t+k}, a_{t+k}) - G_{t+k}| \cdot \mathbb{1}_{(t+k \leq T)}. \quad (7)$$

The timesteps that go beyond the terminal timestep are masked. Out of implementation efficiency, we collect evaluation trajectories of a total length of 1000 and test the error every 1000 train steps. In finite-horizon Atari games, G_{t+k} is calculated as the discounted return. DeepMind Control Suite environments are infinite control problems with time limit $L = 1000$, so G_{t+k} is derived by $L - t - k$ step bootstrap of $Q(s_L, a_{\theta}(s_L))$. For a fair comparison, we set prediction step K for SPR and VCR to be the same ($K = 5$ for Atari 100K, and $K = 3$ for DeepMind Control Suite).

Fig. 1 compares the average and shows the results in a subset of Atari 100K and DeepMind Control Suite. For most curves of SPR, as the policy immediately achieves high returns at the beginning steps, the Q error of SPR grows rapidly and then keeps high until the end. Instead, VCR consistently reduces Q -error in all 6 tasks of DeepMind Control (only show part of results in the figure). VCR achieves an average of 36% reduction in 16 out of Atari games, and comparable Q -error in other games. Note that we are using a high-variance Monte-Carlo return estimator. The variance may dominate the result. If we replace it with a low-variance N -step return, Q -error is reduced in all Atari games as shown in the appendix. This clearly illustrates the effectiveness of VCR.

4.2 Setup for Empirical Evaluation

Environments. We benchmark VCR in environments where the number of interactions is limited. Specifically, we

choose Atari 100K for discrete control and DeepMind Control Suite for continuous control. For DeepMind Control Suite, following Hafner et al. (2019) and Yarats, Kostrikov, and Fergus (2021), we use six environments (*i.e.*, ball-in-cup, finger-spin, reacher-easy, cheetah-run, walker-walk, and cartpole-swingup) for benchmarking with 100K and 500K environment steps.

Baselines. For Atari 100K, we take SPR as a *strong baseline*. Also, SimPLe (Chen et al. 2020), DER (van Hasselt, Hessel, and Aslanides 2019), OTR (Kielak 2020), CURL (Laskin, Srinivas, and Abbeel 2020), DrQ (Yarats, Kostrikov, and Fergus 2021) are chosen as baselines because all of them were state-of-the-arts in Atari 100K at their publications. PlayVirtual (Yu et al. 2021) is chosen as another state-of-the-art baseline, which is also a representation learning method based on SPR requiring a little more computation and memory. For DeepMind Control Suite, we choose Dreamer (Hafner et al. 2020), SAC+AE (Yarats et al. 2021), SLAC (Lee et al. 2020), CURL, DrQ, SPR and PlayVirtual as our baselines. Since SPR is designed for discrete tasks, we adopt a modified SPR based on SAC for continuous tasks¹. EfficientZero (Ye et al. 2021), a method based on Monte-Carlo Tree Search, has achieved excellent performance on both Atari 100K and DeepMind Control 100K. However, considering the success comes at the cost of one order of magnitude more GPU and CPU computation complexity, we do not compare with EfficientZero here.

Implementation Details. For discrete action tasks, we base our implementation of VCR on the official code² of SPR. For Q head, noisy parameters (Fortunato et al. 2018) are reserved because we verify that the noisy q -value output does not have any negative influence on Value-Consistent representation learning (see the appendix). Different from SPR which has an asymmetric prediction head at the end of the online encoding branch, we validate that the removal of the prediction head would not impair the performance (see the appendix). Thus we directly build q head following the encoding branch. The prediction step is set $K = 5$. Q -learning loss and Value-Consistent loss are optimized jointly by an Adam Optimizer (Kingma and Ba 2015), where the batch size is 32. For continuous control tasks, the modified SPR for continuous control is chosen as our codebase. The prediction step is set $K = 3$. Actor loss, critic loss, and Value-Consistent loss are optimized separately by three Adam optimizers, where the batch size for the actor-critic update is 512 and the batch size for VCR update is 128. For more details, please refer to the appendix. Code will be open-sourced upon acceptance.

Evaluation Metrics. According to Agarwal et al. (2021), we choose *interquartile mean (IQM)* and *optimality gap* as *main evaluation indicator* considering their good properties. IQM computes the mean Human Normalized Score (HNS) of the middle 50% runs over all games and seeds. The optimality gap denotes the gap between algorithms and the target performance. Higher IQM and lower optimality gap are better. We also present performance profile curves. On Atari-100K

¹Link: <https://github.com/microsoft/Playvirtual>, MIT License.

²Link: <https://github.com/mila-ijqa/spr>, MIT License.

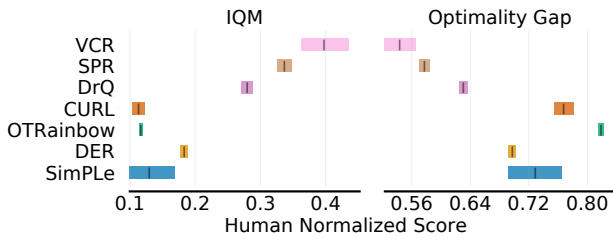


Figure 3: Aggregate IQM and optimality gap of methods. Higher IQM and lower optimality gap are better. The shaded bar shows 95% stratified bootstrap confidence intervals (Agarwal et al. 2021). VCR runs 10 seeds over 26 games.

and DeepMind Control, we use 10 seeds for each game to evaluate VCR. More explanation about the metrics can be seen at the appendix.

4.3 Results of Empirical Evaluation

Atari-100K. As shown in Fig. 3, VCR achieves the best performance on both IQM HNS and optimality gap. It can be seen from the nonoverlapping confidence intervals that our improvement over SPR is statistically significant. Fig. 4 reveals that VCR is nearly above baselines along the whole axis, demonstrating consistent improvement over baselines. The improvement is particularly noteworthy when focusing on the HNS interval between 0.2 and 1.0. The numerical performance of methods is presented in Table 1. By adding value consistency constraint, our method gets a boost over the baseline SPR by **6.0% IQM HNS**, which is significant when considering the 5.7% IQM HNS improvement of SPR over DrQ and 3.7% of PlayVirtual over SPR. Compared to PlayVirtual, the other regularization method based on SPR architecture by producing virtual cycle trajectories, VCR achieves better performance (higher IQM HNS by 2.3%) with less computation and memory consumption (See the appendix). All scores of individual games are shown in the appendix.

DeepMind Control Suite. Within very limited interactions 100K, VCR achieves the best performance on 3 out of 6 tasks, as shown in Table 2. In addition, VCR achieves the best IQM HNS and optimality gap. Compared with the baseline SPR, VCR has a relatively 6.4% higher IQM and 15.5% lower optimality gap. When 500K interactions are allowed, VCR is close to the perfect score in 4 environments, and achieves a comparable median score with DrQ, SPR and PlayVirtual.

4.4 Analysis

In this section, we conduct ablation studies with Atari 100K to analyze the important components in VCR and the effectiveness of VCR compared with other auxiliary losses.

Value-Consistent Distance Metric. We find that the Value-Consistent distance metric has a great influence on the performance. We use a mixed target of n -step target estimation $\bar{G}_t^{(n)}$ and action-value $Q_T(\tilde{z}_t, a)$, i.e., d_{VCR} (see Eqn. 6). Here, we test two variants to validate the proposed Value-Consistent distance metric. The first is to simply apply MSE

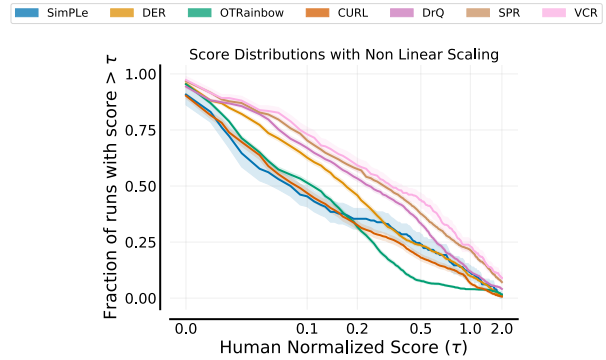


Figure 4: Performance profiles based on score distributions with linear and non-linear scaling on Atari 100K.

loss to the imagined q-values and the target q-values over real state-action pairs: $d_{MSE} = [Q(\hat{z}_t, a_t) - Q_T(\tilde{z}_t, a_t)]^2$. Further, we can enforce value consistency over all possible actions: $d_{MSE-A} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} [Q(\hat{z}_t, a) - Q_T(\tilde{z}_t, a)]^2$. Note that MSE loss is replaced with a cross-entropy for distributional RL.

We evaluate these two variants on Atari 100K. d_{MSE} achieves 30.9% IQM HNS and 60% optimality gap, while d_{MSE-A} achieves 34.0% IQM HNS and 57.9% optimality gap. Compared to d_{VCR} (39.7% IQM HNS and 54.4% optimality gap), two variants d_{MSE} and d_{MSE-A} have inferior performance. Especially for d_{MSE} , it is lower than the baseline (33.7% IQM HNS and 57.7% optimality gap). That may be because, in these two variants, VCR loss has conflicting gradients with policy learning loss (i.e., DQN loss), a common issue in multi-task learning (Sener and Koltun 2018; Yu et al. 2020; Jean, Firat, and Johnson 2019). Q -learning loss pushes the value function towards the value distribution under the optimal policy while d_{MSE} and d_{MSE-A} aim to align with the current value approximation. Intuitively, d_{MSE} has bigger conflict than d_{MSE-A} because only real state-action pairs lead to an explicit conflict, which are only $1/|\mathcal{A}|$ of state-action pairs used in d_{MSE-A} , where $|\mathcal{A}|$ is the size of the action set. It may explain why d_{MSE} has a larger drop in performance.

Comparison with Reward Loss. One may simply attribute the improvement of VCR to the introduction of reward to representation learning in $\bar{G}_t^{(n)}$. Here, we construct a baseline by adding reward prediction based on SPR, where the dynamics model outputs predicted reward and next state conditioned on the current state and action. The predicted reward is supervised by a real reward. We denote this baseline as *SPR with reward loss*, which achieves 35.4% IQM HNS and 56.8% optimality gap. We see that although *SPR with reward loss* achieves a little higher IQM HNS than baseline, it is still far behind VCR. This implies that the effectiveness of our method should be attributed to consistent value prediction rather than involving additional reward prediction.

Comparison with Policy Learning Loss. Value-Consistent loss and policy learning loss are similar in terms of the formula form, both of which leverage Q -values to update the encoder and Q -value head. For example, in

Game	SimPLe	DER	OTR	CURL	DrQ	SPR	VCR	PlayVirtual†
IQM HNS (%)	13.0	18.3	11.7	11.3	28.0	33.7	39.7	37.4
Optimality Gap (%)	72.9	69.8	81.9	76.8	63.1	57.7	54.4	55.8

Table 1: Aggregated scores achieved by different methods on Atari-100k. † denotes using virtual trajectories.

100k Step Scores	Dreamer	SAC+AE	SLAC	CURL	DrQ	SPR	VCR	PlayVirtual†
Finger, spin	341 ± 70	740 ± 64	693 ± 141	767 ± 56	901 ± 104	840 ± 143	795 ± 157	683 ± 189
Cartpole, swingup	326 ± 27	311 ± 11	-	582 ± 146	759 ± 92	815 ± 48	815 ± 47	812 ± 66
Reacher, easy	314 ± 155	274 ± 14	-	538 ± 233	601 ± 213	684 ± 186	763 ± 112	663 ± 214
Cheetah, run	235 ± 137	267 ± 24	319 ± 56	299 ± 48	361 ± 67	452 ± 117	422 ± 54	510 ± 38
Walker, walk	277 ± 12	394 ± 22	361 ± 73	403 ± 24	634 ± 160	397 ± 220	650 ± 143	499 ± 161
Ball in cup, catch	246 ± 174	391 ± 82	512 ± 110	769 ± 43	914 ± 51	807 ± 165	858 ± 85	939 ± 20
IQM HNS	-	-	-	-	731	700	745	690
Optimality Gap	710	603	-	440	305	335	283	316
500k Step Scores								
Finger, spin	796 ± 183	884 ± 128	673 ± 92	926 ± 45	938 ± 103	924 ± 132	972 ± 25	963 ± 40
Cartpole, swingup	762 ± 27	735 ± 63	-	841 ± 45	868 ± 10	870 ± 12	854 ± 26	865 ± 11
Reacher, easy	793 ± 164	627 ± 58	-	929 ± 44	942 ± 71	925 ± 79	938 ± 37	942 ± 66
Cheetah, run	570 ± 253	550 ± 34	640 ± 19	518 ± 28	660 ± 96	716 ± 47	661 ± 32	719 ± 51
Walker, walk	897 ± 49	847 ± 48	842 ± 51	902 ± 43	921 ± 45	916 ± 75	930 ± 18	928 ± 30
Ball in cup, catch	879 ± 87	794 ± 58	852 ± 71	959 ± 27	963 ± 9	963 ± 8	958 ± 4	967 ± 5
Median Score	794.5	764.5	757.5	914.0	929.5	920.0	934.0	935.0

Table 2: Scores (mean and standard deviation) achieved by different methods on the DeepMind Control. We run VCR with 10 seeds. On 500K steps, single run scores of SPR and PlayVirtual are missing to calculate IQM, so we follow Yu et al. (2021) to report median scores to profile the overall performance. † denotes using virtual trajectories.

the discrete setting, both \mathcal{L}_{VCR} and \mathcal{L}_{DQN} align predicted Q -values to n -step estimation, while \mathcal{L}_{VCR} is computed over K times of state-action pairs, where K is prediction steps. Thus a possible concern is that the boost of our method comes from more state-action pairs to update the Q -value head. Here we replace \mathcal{L}_{VCR} on imagined state-action pairs with \mathcal{L}_{DQN} , equivalent to increasing the mini-batch size of \mathcal{L}_{DQN} . *SPR-L* achieves 27.2% IQM HNS and 61.5% optimality gap, while *SPR-XL* achieves 17.6% IQM HNS and 70.1% optimality gap. Two variants display performance drop, which implies VCR helps train the Q -value head but the main gain comes from its regularization on representation learning of the encoder and transition model.

Influence of Prediction Steps K . We increase the number of prediction steps from 5 to 9 to test if more improvement can be obtained. $K = 9$ achieves 39.7% IQM HNS and 54.3% optimality gap, which is roughly comparable to $K = 5$. That means increasing prediction steps in a range would not change the overall performance although the performance on a subset of games increases much (see the appendix), at the cost of more computation and memory.

5 Conclusion and Limitation

To boost the sample efficiency of value-based reinforcement learning algorithms, we propose a novel *Value-Consistent Representation Learning* (VCR) method. The intuition behind VCR is that an agent should be capable of making an

imagination of future states from its behaviors and obtaining correct value predictions based on the imagined states. The property becomes more demanding when the environment is stochastic and learning a precise transition model is impossible. Some previous works have validated the effectiveness of this idea with search-based methods. We develop a value-consistent metric for Q values and introduce it into value-based RL algorithms for the first time. We further show that the method is compatible with any value-based methods by providing two implementations dealing with both discrete and continuous actions. We evaluate our method on two benchmarks including Atari 100K for discrete control and DeepMind Control 100K for continuous control. The results clearly show that VCR can improve sample efficiency significantly and achieve new state-of-the-arts on both tasks. However, there are still some limitations to our method. In RL, except jointly optimizing network with RL loss, representation learning can also be used to pre-train the encoder (Lee et al. 2020). Considering VCR relies on value estimation, VCR may need a proxy value network or work with value-based offline RL methods to enable pre-training. Besides, we derive the value-consistent distance metric by simply employing MSE, which might not be robust enough. We leave these investigations as future work.

References

- Agarwal, R.; Schwarzter, M.; Castro, P. S.; Courville, A. C.; and Bellemare, M. 2021. Deep reinforcement learning at the edge of the statistical precipice. *NeurIPS*.
- Badia, A. P.; Piot, B.; Kapturowski, S.; Sprechmann, P.; Vitvitskyi, A.; Guo, Z. D.; and Blundell, C. 2020. Agent57: Outperforming the Atari human benchmark. In *ICML*.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICML*.
- Farahmand, A.-m. 2018. Iterative value-aware model learning. *NeurIPS*.
- Farahmand, A.-m.; Barreto, A.; and Nikovski, D. 2017. Value-aware loss function for model-based reinforcement learning. In *Artificial Intelligence and Statistics*.
- Farquhar, G.; Baumli, K.; Marinho, Z.; Filos, A.; Hessel, M.; van Hasselt, H. P.; and Silver, D. 2021. Self-Consistent Models and Values. *NeurIPS*.
- Fortunato, M.; Azar, M. G.; Piot, B.; Menick, J.; Osband, I.; Graves, A.; Mnih, V.; Munos, R.; Hassabis, D.; Pietquin, O.; et al. 2018. Noisy networks for exploration. In *ICLR*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent-a new approach to self-supervised learning. *NeurIPS*.
- Grimm, C.; Barreto, A.; Singh, S.; and Silver, D. 2020. The value equivalence principle for model-based reinforcement learning. *NeurIPS*.
- Haara, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Hafner, D.; Lillicrap, T.; Ba, J.; and Norouzi, M. 2020. Dream to Control: Learning Behaviors by Latent Imagination. In *ICLR*.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019. Learning latent dynamics for planning from pixels. In *ICML*.
- Hessel, M.; Danihelka, I.; Viola, F.; Guez, A.; Schmitt, S.; Sifre, L.; Weber, T.; Silver, D.; and Van Hasselt, H. 2021. Muesli: Combining improvements in policy optimization. In *ICML*.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI*.
- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2016. Reinforcement learning with unsupervised auxiliary tasks. In *ICLR*.
- Jean, S.; Firat, O.; and Johnson, M. 2019. Adaptive scheduling for multi-task learning. *arXiv preprint arXiv:1909.06434*.
- Kaiser, L.; Babaeizadeh, M.; Miłos, P.; Osipiński, B.; Campbell, R. H.; Czechowski, K.; Erhan, D.; Finn, C.; Kozakowski, P.; Levine, S.; Mohiuddin, A.; Sepassi, R.; Tucker, G.; and Michalewski, H. 2020. Model Based Reinforcement Learning for Atari. In *ICLR*.
- Kielak, K. P. 2020. Do recent advancements in model-based deep reinforcement learning really improve data efficiency?
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Laskin, M.; Lee, K.; Stooke, A.; Pinto, L.; Abbeel, P.; and Srinivas, A. 2020. Reinforcement learning with augmented data. *NeurIPS*.
- Laskin, M.; Srinivas, A.; and Abbeel, P. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. In *ICML*.
- Lee, A. X.; Nagabandi, A.; Abbeel, P.; and Levine, S. 2020. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *NeurIPS*.
- Lerman, S.; Bi, J.; and Xu, C. 2021. rQdia: Regularizing Q-Value Distributions With Image Augmentation. *OpenReview*.
- Long, A.; Blair, A.; and van Hoof, H. 2022. Fast and Data Efficient Reinforcement Learning from Pixels via Non-Parametric Value Approximation. *AAAI*.
- McInroe, T.; Schäfer, L.; and Albrecht, S. V. 2021. Learning Temporally-Consistent Representations for Data-Efficient Reinforcement Learning. *arXiv preprint arXiv:2110.04935*.
- 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.
- Oh, J.; Singh, S.; and Lee, H. 2017. Value prediction network. *NeurIPS*.
- Schrittwieser, J.; Antonoglou, I.; Hubert, T.; Simonyan, K.; Sifre, L.; Schmitt, S.; Guez, A.; Lockhart, E.; Hassabis, D.; Graepel, T.; et al. 2020. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839): 604–609.
- Schwarzter, M.; Anand, A.; Goel, R.; Hjelm, R. D.; Courville, A.; and Bachman, P. 2021. Data-Efficient Reinforcement Learning with Self-Predictive Representations. In *ICLR*.
- Sener, O.; and Koltun, V. 2018. Multi-task learning as multi-objective optimization. *NeurIPS*.
- Silver, D.; Hasselt, H.; Hessel, M.; Schaul, T.; Guez, A.; Harley, T.; Dulac-Arnold, G.; Reichert, D.; Rabinowitz, N.; Barreto, A.; et al. 2017. The Predictron: End-to-end learning and planning. In *ICML*.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tamar, A.; Wu, Y.; Thomas, G.; Levine, S.; and Abbeel, P. 2016. Value iteration networks. *NeurIPS*.
- 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*.

- van Hasselt, H. P.; Hessel, M.; and Aslanides, J. 2019. When to use parametric models in reinforcement learning? *NeurIPS*.
- Yarats, D.; Kostrikov, I.; and Fergus, R. 2021. Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels. In *ICLR*.
- Yarats, D.; Zhang, A.; Kostrikov, I.; Amos, B.; Pineau, J.; and Fergus, R. 2021. Improving sample efficiency in model-free reinforcement learning from images. *AAAI*.
- Ye, W.; Liu, S.; Kurutach, T.; Abbeel, P.; and Gao, Y. 2021. Mastering atari games with limited data. *NeurIPS*.
- Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020. Gradient surgery for multi-task learning. *NeurIPS*.
- Yu, T.; Lan, C.; Zeng, W.; Feng, M.; Zhang, Z.; and Chen, Z. 2021. Playvirtual: Augmenting cycle-consistent virtual trajectories for reinforcement learning. *NeurIPS*.