

Contrastive Initial State Buffer for Reinforcement Learning

Nico Messikommer, Yunlong Song, Davide Scaramuzza

Abstract—In Reinforcement Learning, the trade-off between exploration and exploitation poses a complex challenge for achieving efficient learning from limited samples. While recent works have been effective in leveraging past experiences for policy updates, they often overlook the potential of reusing past experiences for data collection. Independent of the underlying RL algorithm, we introduce the concept of a *Contrastive Initial State Buffer*, which strategically selects states from past experiences and uses them to initialize the agent in the environment in order to guide it toward more informative states. We validate our approach on two complex robotic tasks without relying on any prior information about the environment: (i) locomotion of a quadruped robot traversing challenging terrains and (ii) a quadcopter drone racing through a track. The experimental results show that our initial state buffer achieves higher task performance than the nominal baseline while also speeding up training convergence.

Multimedia Material A video is available at <https://youtu.be/RB7mDq2fhho> and code at https://github.com/uzh-rpg/cl_initial_buffer

I. INTRODUCTION

In Reinforcement Learning (RL), achieving efficient learning from limited samples has been a central pursuit [1], [2]. The trade-off between exploration and exploitation lies at the heart of this field, posing a complex problem for researchers and practitioners alike. The paradigm of utilizing past experiences to guide future actions has emerged as an effective strategy, leading to the development of techniques such as experience replay [3], [4], [5]. One common aspect of these techniques is the reliance on a replay buffer, a structure that stores past experiences. However, these methods primarily focus on reusing past experiences for updating the policy, overlooking a potentially significant factor: reusing past experiences for data collection.

Traditionally, agents in benchmark tasks are often initialized randomly based on a predefined state distribution. However, many robotic tasks present a considerable challenge for exploration in reinforcement learning since prohibitive amounts of exploration are required to reach certain states and receive some learning signal.

Let’s consider a task where a humanoid robot needs to learn how to perform a backflip. In this task, directly initializing the robot with initial states drawn from a successful backflip trajectory would significantly reduce the required

The authors are with the Robotics and Perception Group, Department of Informatics, University of Zurich, and Department of Neuroinformatics, University of Zurich and ETH Zurich, Switzerland (<http://rpg.ifi.uzh.ch>, nmessi@ifi.uzh.ch). This work was supported by the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement No. 871479 (AERIAL-CORE) and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

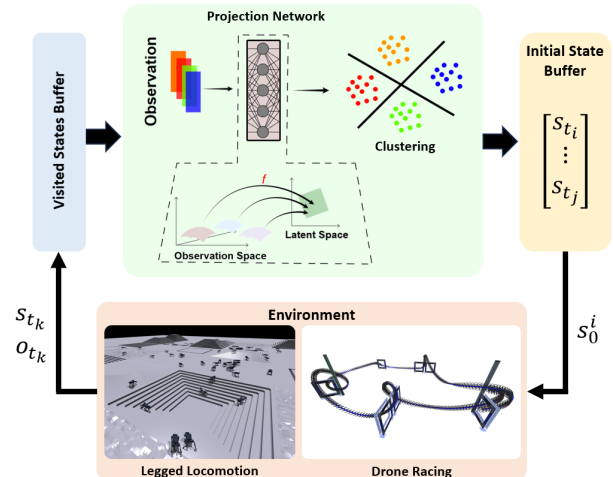


Fig. 1: **Initial State Buffer for Reinforcement Learning.** Our method uses a network to project observations O_{t_k} to an embedding space, in which we apply K-means clustering. In a next step, we add states S_{t_i} close to the cluster center to an Initial State Buffer, which sets the initial states of the robot in the environment s_0^i during the roll-out.

number of environment interactions compared to starting the robot from the ground. This approach leverages prior knowledge about the desired behavior and allows the robot to collect more relevant and informative training data from the beginning. As a result, the movements necessary for successful execution are learned much faster compared to initializing the robot from the ground. However, in most cases, a *successful* trajectory, or even complete information about the environment, is not available beforehand.

To achieve the benefits of sampling from an optimal trajectory but without any prior knowledge of the task or environment, we use the concept of *Initial State Buffer* (ISB), which leverages past experiences for data collection. To optimally select states for the ISB, we propose a novel *Contrastive-Learning Buffer* (CL-Buffer), which maps observations to an embedding space with a learned projection, grouping together states with similar task-relevant experiences, see Fig. 1. The projection is trained in an unsupervised manner, ensuring that states experiencing similar changes in their value function after policy-gradient updates are in close proximity in the embedding space. By using the distance in the embedding space, a diverse set of states featuring a wide range of experiences can be added to the ISB.

We evaluate the effectiveness of our approach on two complex robotics tasks without relying on any prior information about the environment: (i) a quadruped robot learning to walk

from scratch on challenging terrains and (ii) a quadrotor flying through a challenging race track. Our results show that, for the same number of iterations, the introduction of the CL-Buffer leads to overall faster training convergence along with 18.3% higher task performance for the quadruped walking task. Additionally, for the drone racing task, our buffer achieves a success rate of 0.9, in contrast to the baseline without ISB, which has a success rate of 0.2.

II. RELATED WORK

Leveraging past experiences for policy improvement has been studied extensively in the literature. Notably, experience replay [3], [4], [5], [6] via a replay buffer plays a key role in the off-policy setting. However, standard off-policy RL methods can fail in practice due to the overestimation of values induced by the distributional shift between the collected data and the learned policy. Moreover, on-policy algorithms suffer from low sample efficiency since they cannot leverage past experiences for policy improvement.

Researchers are increasingly focusing on enhancing sample efficiency by accumulating more valuable data. These approaches exploit domain knowledge, as demonstrated in studies by [7], [8], [9], [10], or leverage reference demonstrations [11], [12] to initialize the robot with valid initial states. For example, in [7], the initial state of the robot is sampled either from a previous trajectory or a random distribution, with equal probability. Similarly, [9] proposed using an initial state buffer to store valid states in a drone racing task via a heuristic state machine. Compared to our tackled setting, those methods rely on prior knowledge of either the task or the environment, e.g., the race track. For goal-conditioned tasks, [13] proposed using reverse curriculum generation, which teaches an agent to reach the goal from starting positions that are progressively distant from the target. In [14], an archive of visited states is constructed and used to robustify the policy by returning to specified states from the archive. Once the agent reaches the selected archive state, it further explores its environment from there. However, the states are discretized by downscaling images or discretizing robot states, which quickly becomes infeasible for high-dimensional state space common in robotics tasks. In contrast, we propose a contrastive learning embedding combined with k-means to discretize the visited states adaptively to the capabilities of the agent. The sampling of states is also used in other fields, e.g., sampling-based motion planning, where supervised learning from optimal trajectories [15] or Gaussian Mixture Models [16] can guide the optimization.

Finally, exploring the environment can also be induced by directly designing RL agents with intrinsic motivation through adapting the reward formulation [17], [18], [19]. Since our proposed method selects initial states for the roll-out without any assumption on the underlying RL agent, our approach can be applied together with these approaches.

III. PRELIMINARIES

General robot tasks solved by Reinforcement Learning (RL) algorithms can be modeled by a *Markov decision*

process (MDP) which is described by the tuple $M = (S, A, p, p_0, R, \gamma)$ with state space S , action space A , transition probabilities between states p , initial state distribution p_0 and rewards R with discount factor γ .

While the MDP defines the specific task setting, a stochastic control policy $\pi : S \times A \rightarrow \mathbb{R}$ represents a strategy for the robot to navigate in the given setting by proposing actions $a \sim \pi(s)$ that the robot should take in each state $s \in S$. Therefore, starting from an initial state distribution $s_0 \sim p_0$, a policy induces a probabilistic trajectory of states and actions $\tau = \{s_0, a_0, s_1, a_1, \dots, s_T, a_{T-1}\}$. As shorthand, we will write that $\tau \sim \pi$ to denote this relationship. When starting at a specific fixed state s_0 , the possible states and actions reached by a policy will be denoted as $\tau \sim \pi|s_0$. Following the control policy π , the expected future return achieved by the robot at the current state s_i can be expressed with the value function

$$V^\pi(s_i) = \mathbb{E}_{\tau \sim \pi|s_i} \left[\sum_{t=i}^{\infty} \gamma^t R(s_t, a_t) \right]. \quad (1)$$

The goal is to find a policy $\pi^* = \arg \max_{\pi} \mathbb{E}_{s_0 \sim p_0} [V^\pi(s_0)]$, which maximizes the expected cumulative return obtained by starting from initial states $s_0 \sim p_0$. The distribution of initial states is often crucially important in RL since it determines the overall distribution of roll-outs, and may thus affect learning speed and effectiveness. In what follows, we will have a closer look at this distribution and how to leverage it to increase the sample efficiency.

IV. INITIAL STATE BUFFER

One of the main challenges of RL is low sample efficiency, leading to slow training [20]. This problem is amplified for large spaces S , especially when the transition probabilities impede the fast traversal of the state space by the robot. To tackle the sample efficiency problem, we instead propose to sample initial robot states from an *Initial State Buffer (ISB)* which we incrementally build from previously visited states throughout training. We leverage this buffer to (i) reduce the number of initializations of the agent in situations where it must solve already mastered sub-tasks, since this adds training iterations without improving training, and (ii) focus the attention on new or hard situations which require multiple training iterations but are rarely visited.

This greatly diversifies the robot experience leading to enhanced training performance, an observation which was already made in [10], [7], [8], [9], [21], albeit with prior information about the environment and task. By contrast, our ISB can be used in a general setting without domain-specific knowledge. It is thus generally applicable to a diverse set of RL algorithms and problems.

A. Sampling from the ISB

The overview of the ISB for the specific subvariant of the CL-Buffer is visualized in Fig. 1. Since storing all of the already visited states during the roll-out phases is intractable, we use a rolling buffer of states called the *visited states buffer*

$\mathcal{V} = ((s^0, a^0), (s^1, a^1), \dots, (s^k, a^k))^1$, which continuously collects states from the active agents and replaces them in a first-in-first-out fashion. After each roll-out phase, $N = 256$ states from \mathcal{V} are added to the ISB. Every time an episode is terminated, a new episode is started by initializing the robot with a probability of $p = 0.8$ at a state from the ISB and a probability of $1 - p$ at the original start state s_0 . In what follows, we will discuss three variants of ISB, namely the *Random-Buffer*, *Obs-Buffer*, and *Contrastive Learning Buffer* (CL-Buffer).

Random-Buffer The Random-Buffer simply selects a set S of N random states from \mathcal{V}

$$S = \text{RandomSample}_N(\mathcal{V}) \quad (2)$$

While this strategy ensures some diversity, it also tends to oversample states in which the robot spends most of the time during the roll-out episode, e.g. simple transition actions. These states usually also share similar task-relevant experiences. It is thus important to modify the sampling algorithm to increase the diversity of states, which is addressed by the next ISB variant.

Obs-Buffer To ensure efficient learning, it should be avoided to oversample states from \mathcal{V} that are too similar. We thus design an Obs-Buffer, which initiates a K-Means clustering with $K = 64$ in the observation space based on the cosine similarity and then samples the N/K states closest to the obtained cluster centers.

$$\{C_k\}_{k=1}^K = \text{KMeansCluster}(\mathcal{V}) \quad (3)$$

$$S = \bigcup_{k=1}^K \text{NearestNeighbor}_{N/K}(C_k), \quad (4)$$

By doing so, frequently visited states sharing the same observations are all assigned to the same cluster and, thus, are not oversampled. While this modification alleviates the problem of oversampling, clustering in the observation space tends to group states together that, actually, lead to completely different task outcomes. This is because even small changes in the overall state, especially in unstable configurations, e.g., when a robot is about to crash, can have drastically different expected rewards. It is thus beneficial to reenvision how “closeness” is defined when performing clustering, which is discussed next in the *Contrastive Learning Buffer*.

B. Contrastive Learning Buffer

The Contrastive Learning Buffer (CL-Buffer) differs from the Obs-Buffer, in that it trains a neural network f , here a three-layer MLP with Tanh activations, to map the observation to an embedding space, before clustering is performed. However, how this feature space is constructed is non-trivial since it should encode task knowledge to work effectively. Intuitively, we want the network to map states to close-by embeddings if they provide similar experiences to the RL agent since these may require the same “skills”, i.e. subroutine of actions, to solve. We use the term experience here

¹Note here states may come from multiple agents, and thus we use superscripts to avoid confusion with states within one rollout.

to describe, in general, the distribution of tuples containing the observation, action, and reward information acquired at specific states. For example, let us consider the skill of a quadruped to climb stairs facing the stairs upward. In this case, states at the lower and upper part of the stairs provide similar experiences to train the robot to climb stairs. In contrast, letting a quadruped walk backward on a flat surface will provide different experiences, which are less relevant for climbing stairs. Thus, the states on the stairs should be clustered together, whereas the state on a flat surface should be further away in the embedding space.

To quantify such a “closeness”, we argue that states exhibiting a comparable change in the value function based on a policy gradient update contain shared learning experiences. We formalize this similarity by analyzing the value function increase at those states after one policy gradient step, which updates a policy π_0 to π_1

$$\begin{aligned} \Delta V^{\pi_1}(s) &\doteq V^{\pi_1}(s) - V^{\pi_0}(s) \\ &= \mathbb{E}_{(s_0, a_0, \dots) \sim \pi_0} \left[\sum_{t=0}^{\infty} \gamma^t A_{\pi_1}(s_t, a_t) \right] \quad (5) \\ &\approx \sum_{t=0}^{\infty} \gamma^t A_{GAE}^{\pi_1}(s_t, R_t, s_{t+1}, R_{t+1}, \dots). \quad (6) \end{aligned}$$

Here we approximate the change of the value function following [22] by using the *advantage function* $A_{\pi_1}(s_t, a_t)$.

Training the Embedding Network At the beginning of each policy update phase, we select a fixed set of states s^0, s^1, \dots, s^i from the visited state buffer \mathcal{V} , see Fig. 2. For each of those states, we can define a sub-MDP, which shares all of the same properties as the original MDP, i.e., $M = (S, A, p, R, \gamma)$ except for a different start distribution p_0 . At those selected states, we can compute the value function increase $\Delta V^{\pi_1}(s)$ after each policy update. Since we only have access to the roll-out trajectories obtained with π_0 , we leverage the fact that Eq. 5 uses the expectation over the trajectories $\tau \sim \pi_0$, which we approximate by considering only one sample τ obtained during the roll-out phase. Additionally, we approximate the advantage function in Eq. 5 with the *General Advantage Estimator* introduced in [23]. Eq. 6 enables us to quantify for each policy gradient step the improvement over the original policy, for which we have collected the roll-out samples. This difference in the value function at different states can then be used to cluster states with similar task-relevant experiences together.

For each policy update step, the predefined states can be ranked according to the approximated value function increase $\Delta V^{\pi_1}(s)$. The top-k states with the highest increase in the expected return are then used as a positive set P , whereas the lowest-k states are used as a negative set N . Additionally, we select the embedding $\mathbf{x}_i = f(s_i)$ of a state from the positive sets P as an anchor point. Similarly to [24], we use the anchor point, the positive and negative set to formulate the soft-nearest neighbor loss [25], which is a version of the




Fig. 2: **sub-MDP**. A standard MDP problem can be divided into multiple sub-MDPs.

InfoNCE loss [26], as follows

$$L_{\text{contra}} = -\log \frac{\sum_{j \in P} \exp(-g(\mathbf{x}_i, \mathbf{x}_j)/\tau)}{\sum_{k \in N \cup P} \exp(-g(\mathbf{x}_i, \mathbf{x}_k)/\tau)} \quad (7)$$

The function g computes, in our case, the cosine similarity between two different embedding vectors x_i, x_j while τ is a temperature parameter. By backpropagating through the loss, we can update the parameters of the projection network f . Using this network, we map the states in the visited state buffer \mathcal{V} to the embedding space and then perform K-Means clustering. The sampling procedure thus becomes

$$\mathcal{X} = \{x_1, \dots, x_N\} = f(\mathcal{V}) \quad (8)$$

$$\{C_k\}_{k=1}^K = \text{KMeansCluster}(\mathcal{X}) \quad (9)$$

$$S = \bigcup_{k=1}^K \text{NearestNeighbor}_{N/K}(C_k), \quad (10)$$

where we compute the nearest neighbors in the embedding space. Note that while states are regularly mapped to the embedding space to perform sampling, the network is continually trained throughout RL training, which means that it can learn to shift the focus from states as they no longer change in terms of the value function. This type of clustering thus adds a dynamic behavior to the ISB.

V. EXPERIMENTS

Setup We evaluate the different ISB strategies on two robotic platforms: a quadruped and a quadrotor.

The first task involves command following for a quadruped platform, based on the implementation of [10]. The input to the RL agent is the heading direction, which is translated to an angular rotation, and linear velocities commands for the x - and y -axis expressed in its body frame. Based on the commands, the quadruped needs to traverse across different terrain, including upwards- and downwards stairs, rough terrain, smooth terrain, discrete terrain with different elevation planes, and flat surfaces. These terrains are created as square areas of $8\text{m} \times 8\text{m}$ in a simulated IsaacGym [27] environment. During training, the robot is initialized in the middle of the whole environment and needs to explore the other terrain types. However, during validation, we initialize the robot in the center of all of the different terrain types, which enables

us to directly assess how well the robot learned to walk on the different terrains. These validation runs are conducted throughout the training process in time intervals of 100 roll-out phases. We perform 2000 training iteration steps with 1024 environments, which delivers the best trade-off between result expressibility and run-time cost.

In the second task, we use the different ISB strategies to train an agent for drone racing, which is a challenging task that requires pushing the drone to its physical limit [21], [28]. For this task, we train our agent in the Flightmare simulator [29]. The agent receives as input the relative gate corner of the next gate and the angular and linear velocity of the quadrotor and outputs the collective thrust and the body rates. Like the quadruped task, we evaluate each method after 100 roll-out phases. As in [21], we use 2000 training iterations with 100 environments.

Implementation We evaluate the different ISB variants: Random-Buffer, Obs-Buffer, and CL-Buffer, as well as the nominal *Vanilla* baseline without an ISB. To avoid introducing bias to the evaluation of the ISBs, we adopt the same PPO algorithm and reward function design used in [10] for the quadruped platform and the framework from [21] for the quadrotor. Furthermore, the same environments and training frameworks were shown multiple times to transfer to the real world for the quadruped [10], [30], [31] and for the quadrotor [21], [32]. This confirms the validity of the simulation for real-world robotic applications.

We design a simple pre-selection of states to be added to the visited state buffer \mathcal{V} . For the quadrupedal locomotion, we exclude states less than 15 timesteps away from their episode start and states with an accumulated reward below zero since this indicates that the command following is not performed. The ablation experiments in Fig. 3 (d) conducted with our CL-Buffer show that combining both filtering steps leads to the best training performance throughout the training. For drone racing, we only add states to the buffer, which are visited in trajectories initialized at the original start state to avoid diverging from the race track. Generally, our method is not significantly impacted by parameters like cluster numbers and buffer size for visited states.

VI. RESULTS

A. Quadrupedal Locomotion

Method	Reward \uparrow		
	Overall	Easy	Hard
Vanilla	9.21 \pm 0.43	10.19 \pm 0.46	7.62 \pm 0.44
Random-Buffer	9.58 \pm 0.36	10.45 \pm 0.35	8.14 \pm 0.43
Obs-Buffer	10.40 \pm 0.33	11.13 \pm 0.30	9.19 \pm 0.38
CL-Buffer	10.90\pm0.15	11.81\pm0.08	9.46\pm0.36

TABLE I: **Quadrupedal Locomotion**. The validation performance averaged over five different training seeds for different terrain difficulties according to the parametrization of [10].

Performance of CL-Buffer Fig. 3 (a) shows the validation performance of the tested methods at the different training iteration steps. Since we initialize the robot at multiple

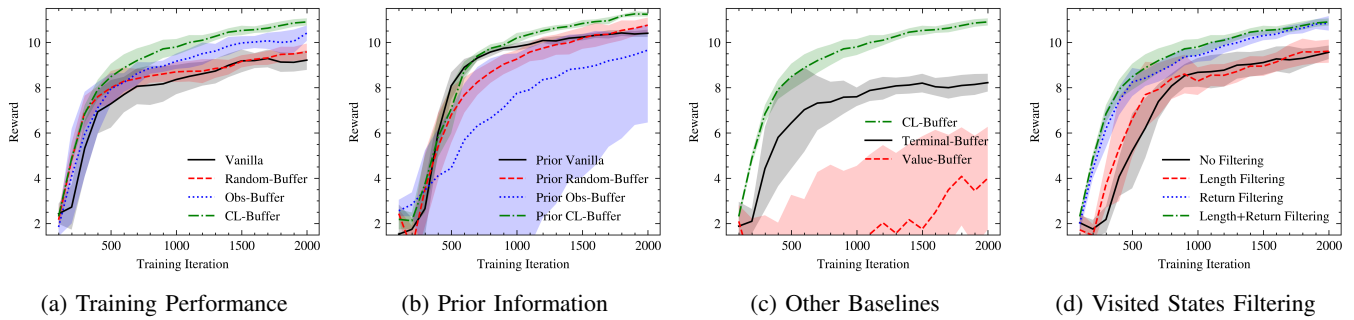


Fig. 3: **Quadrupedal Locomotion.** The mean validation performance at different iteration steps obtained with five different training seeds for all of the tested methods.

terrains during validation, we can directly assess the learned capabilities for different terrain difficulties. It can be observed that initializing the states using an ISB speeds up the learning at the start of the learning, as confirmed by the higher validation performance achieved by all of the ISB methods until the training iteration 1000. However, in the later stages of the training, the Random-Buffer converges to the same performance as the Vanilla approach. This confirms the importance of strategically adding states to the ISB.

In terms of final validation performance, our CL-Buffer outperforms Obs-Buffer, the closest ISB method, by 4.8% and improves the performance compared to the Vanilla baseline by 18.3%, as also reported in Tab. I. This confirms the effectiveness of our proposed state selection based on a learned projection network. It needs to be emphasized that the performance increase is achieved without changing the underlying RL algorithm, which shows the potential of our proposed ISB for general RL tasks.

In Fig. 4 (a), we visualize the state distributions for the tested methods during a roll-out phase performed in the middle of the training. The quadruped states are visualized in a top-down view of the environment using different colors for different methods. The figure shows that the Vanilla approach has a limited exploration radius since it always initializes the trajectories in the middle of the environment. As can be further observed, the Random-Buffer increases the state distribution evenly, which leads to the sampling of redundant experiences. To select a more diverse set of initial states with more relevant experiences, the Obs-Buffer uses observations as distance metrics in the K-Means clustering. However, the observation space can cluster uncorrelated dimensions, making it an unsuitable metric for selecting diverse states. This can lead to states being selected for the ISB, which share similar experiences and are close in the environment, as visualized in Fig. 4 (a). Instead, our CL-Buffer uses an embedding space that clusters states with similar experiences. The result is a diverse state distribution, leading to an improved performance confirmed by the results.

Furthermore, we evaluate the methods if prior information about the environment is available by initializing the agents at the center of all terrains during the training. As can be observed in Fig. 3 (b), the Prior CL-Buffer leads to the best training performance and, compared to the Prior

Vanilla, also has a steeper reward curve at the end of the training, indicating that the performance gain could even further increase. In comparison, the Prior Random-Buffer struggles in selecting useful states if the training data already contains a widespread of situations present in the tested environment. For the Prior Obs-Buffer, the uniform terrain initialization can result in initial state clusters of excessively high difficulty, causing some training runs to collapse, which explains the high variance of the reward curve. If we compare the Prior Vanilla with our CL-Buffer in Fig. 3 (a), which does not have the significant advantage of having access to environment information, the Prior Vanilla only performs similarly and even slightly worse than our proposed CL-Buffer. Furthermore, we evaluate two other simple ISB strategies in Fig. 3 (c): a Terminal-Buffer, which selects states close to the terminal states of previously conducted roll-outs as initial states, and a Value-Buffer, which selects states from the visited state buffer with the highest value predicted by the estimated value function. Both strategies lead to an underperforming training performance.

Analysis of CL-Buffer Clusters To provide more insights into the state clustering of the learned embedding space, three samples closest to three cluster centers are visualized in Fig. 5, which shows the state of the quadruped and the given command visualized as a red line. The visualized clusters can be characterized into different skills, e.g., "Walking Forward," "Walking Sideways Up," and "Failure State." Notably, these skill-based clusters evolved during training and were not directly enforced. This shows that our CL-Buffer selects a diverse set of states featuring different experiences for different skills and thus increases the training convergence and final performance, as confirmed by the validation performance. Furthermore, the capability of our projection network to cluster states in an unsupervised fashion into skills relevant to the given task and adaptively to the current policy throughout the training opens up the door for various applications of the proposed clustering. For example, the projection network can be used to design training environments in a curriculum-learning fashion by interpolating between visited states, similarly done as in previous works in Deep Learning [33] and RL [34], [35]. We believe that our projection technique will provide the tool for many future applications.




Fig. 4: **State Distribution.** (a) The quadruped states during a roll-out phase in the middle of the training in a top-down view. (b) For flying through a race track, if the agent can not fly yet through the complete racetrack, (c) the initial state clusters from the CL-Buffer are located around the struggling gate. (d) Once the agent can finish the racetrack, the CL-Buffer clusters are more spread out while still focusing on the difficult parts around the gates and less on the straight lines.




Fig. 5: **Cluster State Visualization.** Our proposed projection network can be used to cluster the embeddings of different states according to the corresponding experience. It can be observed that each cluster represents a specific skill, i.e., *Walking Forward*, *Walking Sideways Up*, and *Failure State*.

B. Drone Racing

Performance of CL-Buffer To show that the ISB framework works on other environments and platforms without any major changes, we evaluate the methods for drone racing. In drone racing, the objective is to train a policy to pass through a sequence of gates in minimum time, as visualized in a top-down view in Fig. 4 (b). Using RL to address this problem has one inherent challenge: delayed gratification. Progress in early gates is critical to reaching and exploring later gates. The agent may need to pass through gates 1 and 2 before seeing gate 3. Thus, a wrong action in gate 1 might not manifest its negative consequence until gate 3. This long causal chain makes the credit assignment difficult. As shown in Fig. 6, the experimental results clearly show the benefit of using an ISB. Our proposed CL-Buffer leads to the fastest training convergence and the highest success rate, indicating how many of the agents trained with different seeds can complete one round of the race track. The achieved lap times of successful passes are comparable for the different ISB while the CL-Buffer leads to more successful passes.

Analysis of CL-Buffer Clusters The states found by the

CL-Buffer are visualized in Fig. 4 (c) and (d) as dots with different colors, while initial states corresponding to the same cluster center have the same color and are close in proximity. The top-down view in Fig. 4 (c) shows that if the agent has not mastered flying through the entire racetrack, most of the cluster samples are located on the struggling gate, which is, in this case, the last gate. Once the agent can finish the racetrack, the cluster centers are more spread out while focusing on the difficult parts around the gates, see Fig. 4 (d). This shows the evolution ability of our proposed CL-Buffer to adapt to the capabilities of the current trained agent.




Fig. 6: **Drone Racing.** The mean validation reward (left), and the success rate (right) obtained with ten different seeds.

VII. CONCLUSION

This work evaluated the concept of Initial State Buffers (ISB), which strategically select states from past experiences to initialize the robot in the environment in order to guide it toward more informative states. To ensure that the ISB contains states with diverse task-relevant experiences, we proposed a novel Contrastive Learning Buffer (CL-Buffer), which relies on a learned projection to an embedding space in which close-by states share similar experiences. Our results show that the CL-Buffer leads to overall faster training convergence along with 18.3% higher task performance for the quadruped walking task. Additionally, for the drone racing task, our buffer achieves a success rate of 0.9, in contrast to the baseline without ISB, which has a success rate of 0.2. By leveraging task and environment information, future work can extend our CL-Buffer to consider a prioritized sampling from the initial state buffer.

REFERENCES

- [1] K. Arulkumar, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017.
- [2] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in neural information processing systems*, vol. 31, 2018.
- [3] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [4] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine learning*, vol. 8, pp. 293–321, 1992.
- [5] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [6] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. Pieter Abbeel, and W. Zaremba, "Hindsight experience replay," *Advances in neural information processing systems*, vol. 30, 2017.
- [7] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [8] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 143:1–143:14, July 2018.
- [9] Y. Song, M. Steinweg, E. Kaufmann, and D. Scaramuzza, "Autonomous drone racing with deep reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1205–1212.
- [10] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: <https://openreview.net/forum?id=wK2fDDJ5VcF>
- [11] T. Salimans and R. Chen, "Learning montezuma's revenge from a single demonstration," 2018.
- [12] X. B. Peng, A. Kanazawa, J. Malik, P. Abbeel, and S. Levine, "SfV: Reinforcement learning of physical skills from videos," *ACM Trans. Graph.*, vol. 37, no. 6, Nov. 2018.
- [13] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel, "Reverse curriculum generation for reinforcement learning," in *Conference on robot learning*. PMLR, 2017, pp. 482–495.
- [14] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, "First return, then explore," *Nature*, vol. 590, pp. 580–586, 2021.
- [15] B. Ichter, J. Harrison, and M. Pavone, "Learning sampling distributions for robot motion planning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7087–7094.
- [16] J. Huh, B. Lee, and D. D. Lee, "Adaptive motion planning with high-dimensional mixture models," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 3740–3747.
- [17] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," *Advances in neural information processing systems*, vol. 29, 2016.
- [18] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, "Exploration by random network distillation," in *International Conference on Learning Representations*, 2018.
- [19] A. P. Badia, P. Sprechmann, A. Vitvitskiy, D. Guo, B. Piot, S. Kapturowski, O. Tieleman, M. Arjovsky, A. Pritzel, A. Bolt, et al., "Never give up: Learning directed exploration strategies," in *International Conference on Learning Representations*, 2019.
- [20] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 698–721, 2021.
- [21] Y. Song, A. Romero, M. Mueller, V. Koltun, and D. Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, 2023.
- [22] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 1889–1897. [Online]. Available: <https://proceedings.mlr.press/v37/schulman15.html>
- [23] J. Schulman, P. Moritz, S. Levine, M. I. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2016. [Online]. Available: <http://arxiv.org/abs/1506.02438>
- [24] J. Xing, L. Bauersfeld, Y. Song, C. Xing, and D. Scaramuzza, "Contrasting learning for enhancing robust scene transfer in vision-based agile flight," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.
- [25] N. Frosst, N. Papernot, and G. Hinton, "Analyzing and improving representations with the soft nearest neighbor loss," in *International Conference on Machine Learning*. PMLR, 2019, p. 2012–2020.
- [26] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [27] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, "Isaac gym: High performance gpu-based physics simulation for robot learning," 2021.
- [28] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone-racing with deep reinforcement learning," *Nature*, 2023.
- [29] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A flexible quadrotor simulator," in *Conference on Robot Learning*, 2020.
- [30] J. Frey, D. Hoeller, S. Khattak, and M. Hutter, "Locomotion policy guided traversability learning using volumetric representations of complex environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 5722–5729.
- [31] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abk2822>
- [32] A. Romero, Y. Song, and D. Scaramuzza, "Actor-critic model predictive control," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.
- [33] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018.
- [34] R. Sander, W. Schwarting, T. Seyde, I. Gilitschenski, S. Karaman, and D. Rus, "Neighborhood mixup experience replay: Local convex interpolation for improved sample efficiency in continuous control tasks," in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 954–967.
- [35] J. Lin, Z. Huang, K. Wang, X. Liang, W. Chen, and L. Lin, "Continuous transition: Improving sample efficiency for continuous control problems via mixup," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 9490–9497.