

# Unifying Model-Free Efficiency and Model-Based Representations via Latent Dynamics

Jashaswimalya Acharjee  
Wadhvani School of Data Science and AI  
Indian Institute of Technology Madras  
Chennai, India  
jashaswimalya.acharjee@dsai.iitm.ac.in

Balaraman Ravindran  
Wadhvani School of Data Science and AI  
Indian Institute of Technology Madras  
Chennai, India  
ravi@dsai.iitm.ac.in

## ABSTRACT

We present Unified Latent Dynamics (ULD), a novel reinforcement learning algorithm that unifies the efficiency of model-free methods with the representational strengths of model-based approaches, without incurring planning overhead. By embedding state-action pairs into a latent space in which the true value function is approximately linear, our method supports a single set of hyperparameters across diverse domains — from continuous control with low-dimensional and pixel inputs to high-dimensional Atari games. We prove that, under mild conditions, the fixed point of our embedding-based temporal-difference updates coincides with that of a corresponding linear model-based value expansion, and we derive explicit error bounds relating embedding fidelity to value approximation quality. In practice, ULD employs synchronized updates of encoder, value, and policy networks, auxiliary losses for short-horizon predictive dynamics, and reward-scale normalization to ensure stable learning under sparse rewards. Evaluated on 80 environments spanning Gym locomotion, DeepMind Control (proprioceptive and visual), and Atari, our approach matches or exceeds the performance of specialized model-free and general model-based baselines — achieving cross-domain competence with minimal tuning and a fraction of the parameter footprint. These results indicate that value-aligned latent representations alone can deliver the adaptability and sample efficiency traditionally attributed to full model-based planning.

## KEYWORDS

Reinforcement Learning; Latent Dynamics; Model-Based Methods; Value Function Approximation; Sample Efficiency; Representation Learning

## ACM Reference Format:

Jashaswimalya Acharjee and Balaraman Ravindran. 2026. Unifying Model-Free Efficiency and Model-Based Representations via Latent Dynamics. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 7 pages. <https://doi.org/10.65109/OSSS8091>



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/OSSS8091>

## 1 INTRODUCTION

Reinforcement learning (RL) is built on a powerful yet simple idea: an agent learns optimal behaviour by interacting with its environment while guided by a reward signal. In theory, this framework is universal — given a clear objective and sufficient interaction data, an RL agent should be able to adapt to any task. However, in practice, most RL algorithms are far from general-purpose. Instead, they are highly specialized architectures, set of hyperparameters, while incorporating certain training strategies depending on the problem domain.

A pertinent example includes algorithms like Rainbow [13] (designed for Atari games [3]) and TD3 [10] (optimized for MuJoCo control tasks [24]) share little in common beyond their RL foundations. Their hyperparameters differ significantly, and neither performs well when applied to the other’s domain. This specialization undermines RL’s promise of a unified learning framework, forcing practitioners to redesign solutions for each new problem.

Some RL methods do claim generality — policy gradient and evolutionary algorithms, for instance, make minimal assumptions about the environment. However, they often suffer from poor sample efficiency and subpar or asymptotic performance compared to specialized domain-specific methods. Worse, they frequently require laborious hyperparameter tuning, making them impractical for real-world use.

Recently, model-based RL techniques like DreamerV3 [11] and TD-MPC2 [12] have demonstrated impressive generalization, achieving strong performance across diverse benchmarks without task-specific tuning. However, these methods come with high computational costs and added complexity, limiting their accessibility.

This paper introduces ULD, a novel model-free algorithm that leverages learned representations inspired by model-based RL while avoiding its computational overhead. The key insight is that the true advantage of model-based methods may not lie in their explicit environment models, but in the rich feature representations they implicitly learn.

ULD maps state-action pairs into a unified embedding space where their relationship to the value predictions is approximately linear. This approach draws from modern representation learning techniques that encode environment dynamics, as well as the theoretical work of Parr et al. showing that model-based and model-free methods converge to similar solutions in linear settings.

By abstracting away environment-specific details, ULD enables a single set of hyperparameters to work across vastly different tasks — from image observation based environments to proprioceptive or vector-based robotic control environments.

We evaluated our algorithm, ULD, on 80 environments spanning across major RL benchmarks. Without any algorithmic or hyperparameter adjustments, it achieves competitive performance against both specialized model-free methods, as well as general model-based approaches. This suggests that model-based representations alone, without full dynamics modeling, can provide the sufficient sample efficiency and adaptability of model-based RL while retaining the simplicity of model-free methods.

By decoupling the benefits of model-based representations from their traditional learning framework, we have developed an algorithm that works out-of-the-box across domains without excessive tuning or computational demands.

## 2 RELATED WORK

Learning latent embeddings from environment dynamics has emerged as a powerful paradigm for improving RL robustness. Model-free approaches predict future states using learned features (Schwarzer et al., 2023; Guo et al., 2022), while model-based methods like PlaNet (Hafner et al., 2019) and TD-MPC (Hansen et al., 2024) use dynamics models for planning. These efforts share a core insight: predictive representations simplify value estimation and policy optimization. Our work builds on these principles but introduces critical innovations. Unlike TD7 (Fujimoto et al., 2024) – which processes raw inputs – ULD operates exclusively in a learned embedding space, discarding environment-specific input structures. We further incorporate couple of auxiliary losses and integrate multi-step returns and categorical distributional learning component for cross-domain stability.

### 2.1 Theoretical Foundations for State Abstraction

Linear MDP frameworks [2, 14] and linear spectral representation methods [20] formalize low-dimensional spaces where value functions behave linearly. They aim to learn a low-rank decomposition of the transition dynamics of the MDP and to establish a linear correspondence between an embedding and its corresponding value function.

Similarly, bisimulation metrics [5, 7, 8] and MDP homomorphisms [19, 21, 27, 28] cluster states by behavioral equivalence, enabling efficient abstraction. Two-stage RL approaches [6, 15] decouple representation learning from the RL policy, aligning closely with our methodology.

ULD extends these ideas by learning a state-action embedding that approximates linear value dynamics while accommodating both discrete and continuous tasks. Unlike model-based generalists (e.g., DreamerV3), we avoid simulated rollouts; unlike prior model-free representation learners (e.g., TD7), we fully abstract environment-specific inputs.

## 3 BACKGROUND

The general notion of Reinforcement Learning problems are described by a Markov Decision Process (MDP), which is typically defined by a tuple consisting of  $(S, A, p, R, \gamma)$ , where the state space is denoted by  $S$ , action space by  $A$ , the transition dynamics function as  $p$ , reward function by  $R$  and finally the discount factor as  $\gamma$ . In addition to the above, the value-based RL methods learn a value

function,  $Q^\pi(s, a) := \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, a_0 = a]$  which models the expected discounted sum of rewards  $r_t \sim R(s_t, a_t)$ , following a policy  $\pi$ .

The true value function  $Q^\pi$  in our work is estimated by an approximate value function  $Q_\theta$  – subscripts are used to indicate network parameters  $\theta$ . The target networks have parameters denoted by with parameters of the form  $Q_{\theta'}$ , which are intended to introduce stationarity in the predictions. The parameters are periodically synced with current network parameters ( $\theta' \leftarrow \theta$ ).

## 4 MODEL-BASED REPRESENTATIONS

### 4.1 Linear Value Decomposition

We consider a linear decomposition of the action-value function  $Q(s, a)$  using state-action embeddings  $\mathbf{z}_{sa}$  and linear weights  $\mathbf{w}$ :

$$Q(s, a) = \mathbf{z}_{sa}^\top \mathbf{w} \quad (1)$$

Our primary objective is to learn embeddings  $\mathbf{z}_{sa}$  that maintain an approximately linear relationship with the true value function  $Q^\pi$ . Since this relationship may not be exact, we ultimately employ these features as inputs to a non-linear function  $\hat{Q}(\mathbf{z}_{sa})$  rather than relying solely on linear approximation.

### 4.2 Learning Approaches

Given a dataset  $\mathcal{D}$  of transitions  $(s, a, r, s', a')$ , we analyze two learning paradigms:

**Model-Free Approach:** Semi-gradient temporal difference learning:

$$\mathbf{w} \leftarrow \mathbf{w} - \alpha \mathbb{E}_{\mathcal{D}} \left[ \nabla_{\mathbf{w}} \left( \mathbf{z}_{sa}^\top \mathbf{w} - [r + \gamma \mathbf{z}_{s'a'}^\top \mathbf{w}]_{\text{sg}} \right)^2 \right] \quad (2)$$

where  $[\cdot]_{\text{sg}}$  denotes the stop-gradient operator.

**Model-Based Approach:** Dynamics-based rollout approximation:

$$\mathbf{w}_{\text{mb}} := \sum_{t=0}^{\infty} \gamma^t \mathbf{W}_p^t \mathbf{w}_r \quad (3)$$

with reward and transition weights:

$$\mathbf{w}_r := \arg \min_{\mathbf{w}} \mathbb{E}_{\mathcal{D}} \left[ \left( \mathbf{z}_{sa}^\top \mathbf{w} - r \right)^2 \right] \quad (4)$$

$$\mathbf{W}_p := \arg \min_{\mathbf{W}} \mathbb{E}_{\mathcal{D}} \left[ \left\| \mathbf{z}_{sa}^\top \mathbf{W} - \mathbf{z}_{s'a'}^\top \right\|^2 \right] \quad (5)$$

### 4.3 Equivalence of Approaches

Following foundational work in linear RL, we establish the equivalence between these paradigms:

**THEOREM 4.1 (SOLUTION EQUIVALENCE).** *The fixed point of the model-free update (2) and the model-based solution (3) are identical.*

**PROOF.** See Supplementary Material [1, A.1] for the complete derivation.  $\square$

**4.3.1 Value Error Bound.** From Theorem 4.1, we establish a bound on the approximation error of the value function, which may be defined in the following manner:

$$\text{VE}(s, a) := Q(s, a) - Q^\pi(s, a) \quad (6)$$

**THEOREM 4.2 (ERROR BOUND).** *The value error for the solution in Theorem 4.1 is bounded by:*

$$|VE(s, a)| \leq \frac{1}{1-\gamma} \left( \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} |z_{sa}^\top \mathbf{w}_r - \mathbb{E}[r|s, a]| + \max_i |w_i| \sum |z_{sa}^\top \mathbf{W}_p - \mathbb{E}_{s',a'}[z_{s'a'}]| \right) \quad (7)$$

In other words it is bounded by the accuracy of the estimated dynamics and reward.

**PROOF.** The complete proof appears in Supplementary Material [1, A.2].  $\square$

#### 4.4 Practical Limitations and Adjustments

While theoretically sound, the joint optimization objective:

$$\mathcal{L}(z_{sa}, \mathbf{w}_r, \mathbf{W}_p) = \mathbb{E}_{\mathcal{D}} \left[ (z_{sa}^\top \mathbf{w}_r - r)^2 \right] + \lambda \mathbb{E}_{\mathcal{D}} \left[ \|z_{sa}^\top \mathbf{W}_p - z_{s'a'}\|^2 \right] \quad (8)$$

encounters two significant practical limitations:

- (1) *Policy Dependency:* The target  $z_{s'a'}$  depends on actions  $a' \sim \pi$ , creating coupling between policy updates and representation learning.
- (2) *Optimization Challenges:* Joint optimization can lead to degenerate solutions similar to Bellman residual minimization, particularly with sparse rewards or incomplete state coverage.

To address these issues, we propose a modified objective:

$$\mathcal{L}(z_{sa}, \mathbf{w}_r, \mathbf{W}_p) = \mathbb{E}_{\mathcal{D}} \left[ (z_{sa}^\top \mathbf{w}_r - r)^2 \right] + \lambda \mathbb{E}_{\mathcal{D}} \left[ \|z_{sa}^\top \mathbf{W}_p - \bar{z}_{s'}\|^2 \right] \quad (9)$$

with two key modifications:

- (1) Use state-only embeddings  $z_{s'}$  as dynamics targets instead of  $z_{s'a'}$
- (2) Employ a target network  $\phi'_s(s')$  with slowly updated parameters  $\theta'$  to generate  $\bar{z}_{s'}$ .

#### 4.5 Non-Linear Value Representation

Although our adjustments break the theoretical linear relationship between the embedding  $z_{sa}$  and the value function, we establish conditions for non-linear representability:

**THEOREM 4.3 (NON-LINEAR REPRESENTATION).** *Given state encoder  $\phi_s(s) = z_s$  and action embedding function  $\phi_{sa}(z_s, a) = z_{sa}$ , if there exist functions  $\hat{p}$  and  $\hat{R}$  satisfying:*

$$\mathbb{E}_{\hat{R}}[\hat{R}(z_{sa})] = \mathbb{E}_R[R(s, a)] \quad (10)$$

$$\hat{p}(z_{s'}|z_{sa}) = \sum_{\hat{s}: \phi(\hat{s})=z_{s'}} p(\hat{s}|s, a) \quad (11)$$

then for any policy  $\pi$  with corresponding  $\hat{\pi}(a|z_s) = \pi(a|s)$ , there exists  $\hat{Q}$  such that:

$$\hat{Q}(z_{sa}) = Q^\pi(s, a) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (12)$$

Moreover, equation 12 guarantees that there exists an optimal policy  $\hat{\pi}^*(a|z_s) = \pi^*(a|s)$ .

**PROOF.** See Supplementary Material [1, A.3] for the constructive proof.  $\square$

This theoretical foundation motivates our practical algorithm to learn embeddings using the adjusted objective (9), which capture approximately the linear value relationships, then employ a non-linear  $\hat{Q}$  function to compensate for approximation errors.

## 5 ALGORITHM

We now present ULD, our practical algorithm that implements the aforementioned theoretical insights. The core principle is to learn state-action embeddings  $z_{sa}$  that maintain an approximately linear relationship with the true value function  $Q^\pi$ , while using these embeddings as inputs to non-linear function approximators to handle approximation errors.

### 5.1 Architecture Overview

ULD consists of three main components trained end-to-end:

- (1) **State Encoder**  $\phi_s(s) \rightarrow z_s$ : Maps observations to state embeddings
- (2) **State-Action Encoder**  $\phi_{sa}(z_s, a) \rightarrow z_{sa}$ : Combines state embeddings with actions
- (3) **Environment Model  $\mathbf{m}$** : Linear predictor for next state, reward, and termination

The state encoder  $\phi_s(s)$  is modular and can be adapted to different observation modalities by changing its architecture while keeping the remaining components unchanged. Since  $z_s$  is a fixed-size vector, all downstream networks use standard feedforward architectures regardless of the input observation space.

### 5.2 Learning Objective

Given a transition tuple  $(s, a, r, d, s')$  from the replay buffer, where  $d$  indicates episode termination, our learning objective consists of three components:

*5.2.1 Representation Learning.* Following the adjusted loss from Equation 9, we learn representations by unrolling the learned model over a finite horizon  $H_{enc}$ . Starting from an initial state  $s^{(0)}$ , we recursively apply:

$$\tilde{z}_{s_t}, \tilde{r}_t, \tilde{d}_t = \phi_{sa}(\tilde{z}_{s_{t-1}}, a_{t-1})^T \mathbf{m} \quad (13)$$

where  $\tilde{z}_{s_0} = \phi_s(s_0)$  and  $\mathbf{m}$  is the linear environment model.

The representation loss combines three terms:

$$\mathcal{L}_{rep}(\phi_s, \phi_{sa}, \mathbf{m}) = \sum_{t=1}^{H_{enc}} [\lambda_r \mathcal{L}_{reward}(\tilde{r}_t, r_t) \quad (14)$$

$$+ \lambda_d \mathcal{L}_{dynamics}(\tilde{z}_{s_t}, \tilde{z}_{s_t}) \quad (15)$$

$$+ \lambda_t \mathcal{L}_{terminal}(\tilde{d}_t, d_t)] \quad (16)$$

**Reward Loss:** To handle sparse rewards and varying magnitudes robustly, we use a categorical representation with cross-entropy loss:

$$\mathcal{L}_{reward}(\tilde{r}, r) = \text{CrossEntropy}(\tilde{r}, \text{TwoHot}(r)) \quad (17)$$

The two-hot encoding uses non-uniform spacing based on  $\text{symexp}(x) = \text{sign}(x)(\exp(|x|) - 1)$  to handle diverse reward scales.

**Dynamics Loss:** We minimize the squared error between predicted and target state embeddings:

$$\mathcal{L}_{\text{dynamics}}(\bar{z}_s, \bar{z}_s) = \|\bar{z}_s - \bar{z}_s\|_2^2 \quad (18)$$

where  $\bar{z}_s = \phi'_s(s)$  comes from the target encoder with slowly-updated parameters.

**Terminal Loss:** We predict episode termination using mean squared error:

$$\mathcal{L}_{\text{terminal}}(\tilde{d}, d) = (\tilde{d} - d)^2 \quad (19)$$

The terminal loss coefficient  $\lambda_t$  is set to zero until the first terminal transition is observed, following standard practice in model-based RL.

**5.2.2 Value Function Learning.** We employ a modified TD3-style approach with several enhancements for cross-domain stability. We train two critic networks  $Q_{\theta_1}$  and  $Q_{\theta_2}$  that operate on state-action embeddings:

$$\mathcal{L}_{\text{value}}(\tilde{Q}) = \text{Huber}\left(\tilde{Q}, \frac{1}{\bar{r}} \left[ \sum_{k=0}^{H_Q-1} \gamma^k r^{(t+k)} + \gamma^{H_Q} \tilde{Q}'_j \right]\right) \quad (20)$$

where  $\tilde{Q}'_j = \bar{r}' \min_{j=1,2} Q_{\theta'_j}(z_{sa_\pi})$  represents the target value scaled by the target reward normalization factor  $\bar{r}'$ .

Key modifications include:

- **Multi-step Returns:** We use  $H_Q$ -step returns for improved sample efficiency.
- **Huber Loss:** Reduces bias from prioritized experience replay.
- **Reward Normalization:** Targets are scaled by  $\bar{r}$ , the running average of absolute rewards, for consistent loss magnitudes across domains

The target action  $a^{(\pi)}$  is computed using the target policy  $\pi'$  with added noise:

$$a^{(\pi)} = \begin{cases} \text{clip}(\pi'(z_{s'}) + \epsilon, -1, 1) & \text{continuous actions} \\ \pi'(z_{s'}) & \text{discrete actions} \end{cases} \quad (21)$$

where  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  is clipped Gaussian noise.

**5.2.3 Policy Learning.** The policy operates on state embeddings  $z_s$  and is trained using the deterministic policy gradient algorithm:

$$\mathcal{L}_{\text{policy}}(a_\pi) = -\frac{1}{2} \sum_{i=1,2} \tilde{Q}_i(z_{sa_\pi}) + \lambda_{\text{pre}} \mathbf{u}_\pi^2 \quad (22)$$

where  $a_\pi = \text{activation}(\mathbf{u}_\pi)$  and  $\mathbf{u}_\pi$  represents the pre-activation values of the state.

To handle both continuous and discrete action spaces uniformly, we use:

- **Continuous actions:** Tanh activation with Gaussian exploration noise
- **Discrete actions:** Gumbel-Softmax activation with noise added to one-hot encodings

The pre-activation regularization term  $\lambda_{\text{pre}} \mathbf{u}_\pi^2$  helps avoid local minima in sparse reward environments.

## 5.3 Training Protocol

ULD follows a synchronized update schedule to reduce non-stationarity:

---

### Algorithm 1 Agent Evaluation Protocol

---

**Require:** Policy  $\pi$ , Evaluation environment  $\mathcal{E}_{\text{eval}}$ , Episodes  $N_{\text{eval}}$

```

1: Initialize:  $\mathbf{R} \leftarrow \text{zeros}(N_{\text{eval}})$ 
2: for ep = 0 to  $N_{\text{eval}} - 1$  do
3:    $s_0, d_0 \leftarrow \mathcal{E}_{\text{eval}}.\text{reset}(), \text{False}$ 
4:    $t \leftarrow 0$ 
5:   while NOT  $d_t$  do
6:      $a_t \leftarrow \pi(z_{s_t})$  {Deterministic policy}
7:      $s_{\tau+1}, r_\tau, d_t \leftarrow \mathcal{E}_{\text{eval}}.\text{step}(a_t)$ 
8:      $\tau \leftarrow \tau + 1$ 
9:   end while
10:   $\mathbf{R}[\text{ep}] \leftarrow \sum_{\tau=0}^{T_{\text{ep}}} r^{(\tau)}$ 
11: end for
12:  $\mathcal{J}_{\text{eval}} \leftarrow \frac{1}{N_{\text{eval}}} \sum_{i=0}^{N_{\text{eval}}-1} \mathbf{R}[i]$ 
13: Store evaluation score:  $\mathcal{J}_{\text{eval}}$ 

```

---



---

### Algorithm 2 ULD Training Loop

---

```

1: Initialize:  $t \leftarrow 0$ , start_time  $\leftarrow$  current_time()
2: Initialize networks:  $\phi_s, \phi_{sa}, \mathbf{m}, Q_{\theta_1}, Q_{\theta_2}, \pi$ 
3: Initialize target networks:  $\phi'_s \leftarrow \phi_s, \phi'_{sa} \leftarrow \phi_{sa}, Q'_{\theta_1} \leftarrow Q_{\theta_1},$ 
    $Q'_{\theta_2} \leftarrow Q_{\theta_2}, \pi' \leftarrow \pi$ 
4: Initialize replay buffer  $\mathcal{D}$  and reward scaling  $\bar{r} \leftarrow 0$ 
5:  $s \leftarrow \mathcal{E}_{\text{train}}.\text{reset}()$ 
6: while  $t \leq T_{\text{total}}$  do
7:   if  $t \bmod f_{\text{eval}} = 0$  then
8:     call EVALUATE( $\pi', \mathcal{E}_{\text{eval}}, N_{\text{eval}}$ )
9:   end if
10:   $a \leftarrow \pi(z_s^{(t)}) + \epsilon^{(t)}$ 
11:   $s', r, done \leftarrow \mathcal{E}_{\text{train}}.\text{step}(a)$ 
12:  replay_buffer.add( $s, a, s', r, done$ ) to  $\mathcal{D}$ 
13:   $s \leftarrow s'$ 
14:  Sample batch  $\mathcal{B}$  from  $\mathcal{D}$ 
15:  Update  $\bar{r} \leftarrow$  running average of  $|r|$  over  $\mathcal{B}$ 
16:  Compute:  $\mathcal{L}_{\text{rep}}, \mathcal{L}_{\text{value}}$  and  $\mathcal{L}_{\text{policy}}$ 
17:  if  $t \bmod T_{\text{target}} = 0$  then
18:    Update target networks:  $\phi'_s \leftarrow \phi_s, \phi'_{sa} \leftarrow \phi_{sa}, Q'_{\theta_1} \leftarrow Q_{\theta_1},$ 
    $Q'_{\theta_2} \leftarrow Q_{\theta_2}, \pi' \leftarrow \pi$ 
19:    Update target reward scaling:  $\bar{r}' \leftarrow \bar{r}$ 
20:  end if
21:   $t \leftarrow t + 1$ 
22: end while

```

---

We use prioritized experience replay with priorities based on TD errors, and update target networks every  $T_{\text{target}}$  steps to maintain training stability. The reward scaling factor  $\bar{r}$  is computed as the running average of absolute rewards and helps normalize loss magnitudes across different environments.

**Table 1: TD3-Normalised Benchmark comparison summary**

Method	Gym	DMC-P	DMC-V	Atari
TD7	<b>1.00</b>	0.87	-	-
TD-MPC2	0.92	<b>0.95</b>	0.85	0.78
DreamerV3	0.76	0.82	0.76	<b>1.00</b>
PPO	0.68	0.71	0.62	0.82
ULD	0.97	<b>0.96</b>	<b>0.92</b>	0.94

## 5.4 Computational Complexity

ULD maintains the computational efficiency of model-free methods while benefiting from model-based representations. The additional overhead consists of:

- Forward passes through the environment model during representation learning
- Unrolling predictions over horizon  $H_{enc}$  (typically 3-5 steps)
- Computing auxiliary losses for reward, dynamics, and terminal prediction

Unlike full model-based methods, we do not perform explicit planning or long-horizon rollouts, keeping the computational cost manageable while achieving the representational benefits of dynamics modeling.

## 6 EXPERIMENTS

We assess ULD’s effectiveness across four established reinforcement learning benchmarks comprising 80 distinct environments. Our comparative analysis includes: (1) Domain-specialized methods that represent current state-of-the-art in each benchmark, (2) General model-based approaches (DreamerV3 [11], TD-MPC2 [12]), and (3) A widely adopted model-free algorithm (PPO [22]). Rather than pursuing benchmark-specific state-of-the-art results, our primary objective is to demonstrate consistent competence across diverse domains using identical hyperparameters. All baselines employ author-recommended configurations without domain-specific tuning.

### 6.1 Benchmark Performance

Figure 1 presents aggregated learning curves, with comprehensive results detailed in Supplementary Material [1, B].

**Continuous Control - Gym Locomotion:** We evaluate on 5 MuJoCo locomotion tasks [25] with continuous actions and low-dimensional states within the Gym framework [4, 26]. Agents train for 1M environment steps without specialized preprocessing. Comparative baselines include: TD7 [9] (near state-of-the-art), TD-MPC2, DreamerV3, and PPO. Performance is normalized against TD3 [10].

**Robotics Control - Proprioceptive DMC:** Using the DeepMind Control suite [23], we examine 23 robotics tasks with vector observations. Episode rewards are capped at 1000 for standardized comparison. Agents train for 500k steps (equivalent to 1M environment frames due to action repetition). Baselines include TD7, TD-MPC2 (near state-of-the-art), DreamerV3, and PPO.

**Vision-Based Control - Visual DMC:** Employing identical tasks to the proprioceptive benchmark but with pixel observations,

agents train for 500k steps. Comparative methods include: DrQ-v2 [29] (model-free state-of-the-art), TD-MPC2, DreamerV3, and PPO.

**Discrete Decision-Making - Atari:** Using the Arcade Learning Environment [3], we evaluate on 43 games with pixel observations and discrete actions. Standard preprocessing includes sticky actions [16]. Agents train for 2.5M steps (10M frames). Baselines include: DreamerV3, DQN [17], Rainbow [13], and PPO. Scores are normalized relative to human performance.

### 6.2 Analysis

Our results reveal the characteristic “no free lunch” tradeoff – top-performing methods in one benchmark typically underperform in others. Despite this, ULD demonstrates exceptional versatility:

- Achieves superior performance on both proprioceptive and visual DMC benchmarks, demonstrating effective cross-modal representation learning.
- Maintains competitive performance on Gym locomotion tasks, trailing only the specialized TD7 approach.
- Delivers strongest overall performance across continuous control domains.
- Outperforms well-known model-free baselines (PPO, DreamerV3, Rainbow) on discrete-action Atari games.

Notably, while DreamerV3 excels on Atari, this comes with substantial computational overhead and diminished performance on other benchmarks. ULD achieves its cross-domain competence without such tradeoffs, validating our design philosophy.

Table 1 summarizes normalized performance across domains, illustrating ULD’s balanced competence. The algorithm’s consistent performance across observation spaces (vector vs. pixel) and action types (continuous vs. discrete) confirms its general-purpose utility.

## 7 DISCUSSION AND CONCLUSION

This work presented ULD, a general-purpose model-free deep reinforcement learning algorithm that achieves consistently strong performance across diverse continuous control and visual benchmarks. Inspired by insights from model-based representation learning, ULD demonstrates that model-free deep RL—when equipped with powerful representation objectives and well-grounded design choices—can match or even surpass model-based approaches, while remaining computationally efficient and conceptually simple.

*Bridging model-based and model-free paradigms.* Although ULD is entirely model-free during execution, it draws upon principles traditionally associated with model-based methods, particularly in its representation learning stage. This hybrid perspective shows that the benefits often attributed to model-based RL may, in fact, arise primarily from effective latent representations rather than explicit planning or trajectory simulation. Unlike algorithms such as DreamerV3 and TD-MPC2, which rely on short-horizon rollouts and planning, ULD achieves competitive or superior performance without these computationally expensive components. This observation suggests that, for many standard benchmarks, learning robust state-action embeddings may be sufficient to close the performance gap between model-based and model-free RL.

*Benchmark generalization.* Our experimental results reveal a pronounced lack of positive transfer between common RL benchmarks. Despite similar underlying dynamics (e.g., Gym and DMC both

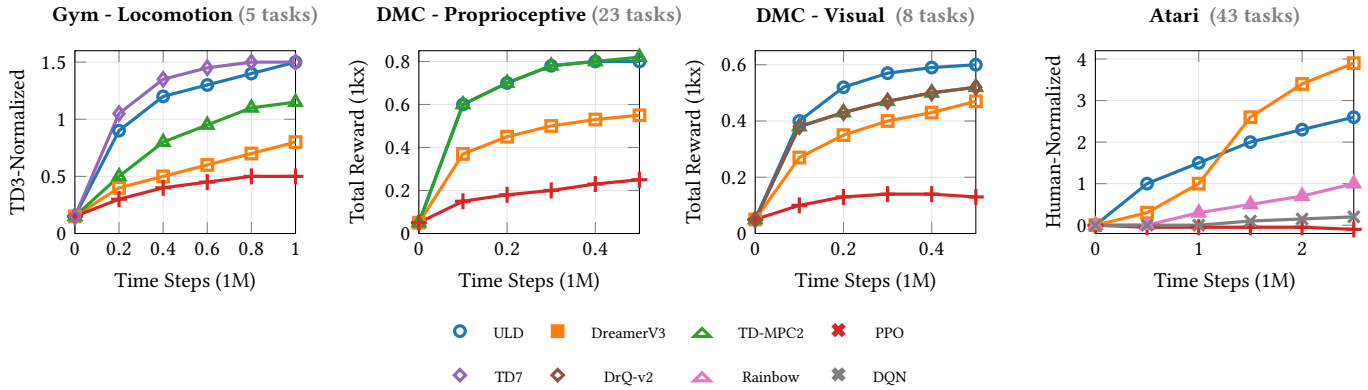


Figure 1: Aggregate learning curves. Average performance over each benchmark. Results are over 10 seeds. Due to action repeat, 500k time steps in DMC correspond to 1M frames in the original environment and 2.5M time steps in Atari corresponds to 10M frames in the original environment.




Figure 2: Aggregate metrics comparison: (Left) Gym Locomotion, (Right) DMC Proprioceptive.

using the MuJoCo simulator), methods that dominate one domain often fail to generalize to the other. In particular, while DreamerV3 performs strongly on Atari, it underperforms in continuous control tasks, struggling in complex locomotion environments such as DMC-Dog and DMC-Humanoid. Conversely, algorithms tuned for proprioceptive control, including TD-MPC2, fail to replicate their advantage in visual or high-dimensional observation spaces. These findings reinforce that progress on a single benchmark does not imply broader generalization, highlighting the need for more comprehensive, cross-domain evaluation protocols.

*Limitations and future directions.* ULD represents one of the early step towards scalable and general model-free deep RL. Nevertheless, several challenges remain. The current formulation does not explicitly address tasks requiring long-horizon reasoning, hard exploration, or non-Markovian dependencies. Moreover, our evaluation is limited to widely adopted RL benchmarks for controlled comparison. Extending these methods to real-world domains — such as multi-agent coordination, robotics, or language-conditioned control — will be necessary to demonstrate true generality. Established

baselines like PPO have already proven adaptable in such domains, from complex team games to drone racing and large language models. Achieving comparable versatility remains an open frontier for future model-free algorithms.

*Conclusion.* In summary, ULD challenges the conventional divide between model-based and model-free RL, showing that strong representation learning and architectural regularization can yield general-purpose, high-performance agents. The results across Gym, DMC-Proprioceptive, DMC-Visual, and Atari benchmarks demonstrate the promise of pursuing simplicity without sacrificing generality, paving the way for a new generation of unified RL algorithms.

### ACKNOWLEDGMENTS

The authors gratefully acknowledge the support received from the Prime Minister’s Research Fellowship (PMRF), Ministry of Education, Government of India, for providing necessary funding and resources to carry out this research.

## REFERENCES

- [1] Jashaswimalya Acharjee and Balaraman Ravindran. 2026. Unifying Model-Free Efficiency and Model-Based Representations via Latent Dynamics. arXiv:2602.12643 [cs.LG] <https://arxiv.org/abs/2602.12643>
- [2] Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. 2020. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in neural information processing systems* 33 (2020), 20095–20107.
- [3] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of artificial intelligence research* 47 (2013), 253–279.
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. *arXiv preprint arXiv:1606.01540* (2016).
- [5] Pablo Samuel Castro. 2020. Scalable methods for computing state similarity in deterministic Markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 10069–10076.
- [6] Wesley Chung, Somjit Nath, Ajin Joseph, and Martha White. 2019. Two-timescale networks for nonlinear value function approximation. In *International Conference on Learning Representations (ICLR)*.
- [7] Norm Ferns, Prakash Panangaden, and Doina Precup. 2004. Metrics for finite Markov decision processes. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI)*. 162–169.
- [8] Norm Ferns, Prakash Panangaden, and Doina Precup. 2011. Bisimulation metrics for continuous Markov decision processes. *SIAM J. Comput.* 40, 6 (2011), 1662–1714.
- [9] Scott Fujimoto, David Guo, and David Meger. 2024. TD7: Transformer-based Representation for Continuous Control. *arXiv preprint arXiv:2406.08672* (2024).
- [10] Scott Fujimoto, Herke Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*. PMLR, 1587–1596.
- [11] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2023. Mastering Diverse Domains through World Models. *arXiv preprint arXiv:2301.04104* (2023).
- [12] Nicklas Hansen, Hao Su, and Xiaolong Wang. 2024. TD-MPC2: Scalable, Robust World Models for Continuous Control. *arXiv preprint arXiv:2405.17577* (2024).
- [13] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [14] Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. 2023. Provably efficient reinforcement learning with linear function approximation. *Mathematics of Operations Research* 48, 3 (2023), 1496–1521.
- [15] Nir Levine, Tom Zahavy, Daniel J Mankowitz, Aviv Tamar, and Shie Mannor. 2017. Shallow updates for deep reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 30.
- [16] Marlos C Machado, Marc G Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. 2018. Revisiting the Arcade Learning Environment: Evaluation protocols and open problems for general agents. *JAIR* 61 (2018).
- [17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015).
- [18] Ronald Parr, Lihong Li, Gavin Taylor, Christopher Painter-Wakefield, and Michael L Littman. 2008. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *Proceedings of the 25th International Conference on Machine Learning (Helsinki, Finland) (ICML '08)*. Association for Computing Machinery, New York, NY, USA, 752–759. <https://doi.org/10.1145/1390156.1390251>
- [19] Balaraman Ravindran. 2004. *An algebraic approach to abstraction in reinforcement learning*. PhD thesis. University of Massachusetts Amherst.
- [20] Tongzheng Ren, Tianjun Zhang, Csaba Szepesvári, and Bo Dai. 2022. A free lunch from the noise: Provable and practical exploration for representation learning. In *Uncertainty in Artificial Intelligence*. PMLR, 1686–1696.
- [21] Sahand Rezaei-Shoshtari, Rosie Zhao, Prakash Panangaden, David Meger, and Doina Precup. 2022. Continuous MDP homomorphisms and homomorphic policy gradient. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [22] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. In *Proceedings of the 34th ICML*.
- [23] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andy Lefrancq, et al. 2018. DeepMind Control Suite. *arXiv preprint arXiv:1801.00690* (2018).
- [24] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 5026–5033. <https://doi.org/10.1109/IROS.2012.6386109>
- [25] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ IROS*.
- [26] Mark Towers, Jordan K Terry, Ariel Kwiatkowski, John Balis, Gianluca de Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Arjun KG, Markus Krimmel, et al. 2024. Gymnasium. *arXiv preprint arXiv:2401.12961* (2024).
- [27] Elise van der Pol, Thomas Kipf, Frans A Oliehoek, and Max Welling. 2020. Plannable approximations to MDP homomorphisms: Equivariance under actions. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1431–1439.
- [28] Elise van der Pol, Daniel Worrall, Herke van Hoof, Frans Oliehoek, and Max Welling. 2020. MDP homomorphic networks: Group symmetries in reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 33. 4199–4210.
- [29] Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. 2022. DrQ-v2: Improved Data-Augmented Deep Reinforcement Learning. *arXiv preprint arXiv:2205.12730* (2022).