

TEA: Trajectory Encoding Augmentation for Robust and Transferable Policies in Offline Reinforcement Learning

Batıkan Bora Ormancı^{1,2}, Phillip Swazinna¹, Steffen Udluft¹, Thomas A. Runkler^{1,2} *

1- Siemens AG, Munich, Germany

2- Technical University of Munich, Germany

Abstract. In this paper, we investigate offline reinforcement learning (RL) with the goal of training a single robust policy that generalizes effectively across environments with unseen dynamics. We propose a novel approach, Trajectory Encoding Augmentation (TEA), which extends the state space by integrating latent representations of environmental dynamics obtained from sequence encoders, such as autoencoders. Our findings show that incorporating these encodings with TEA improves the transferability of a single policy to novel environments with new dynamics, surpassing methods that rely solely on unmodified states. These results indicate that TEA captures critical, environment-specific characteristics, enabling RL agents to generalize effectively across dynamic conditions.

1 Introduction & Related Work

Reinforcement Learning (RL) has seen notable success in addressing complex decision-making tasks across different fields. One ongoing challenge is to enable RL agents to generalize effectively when exposed to environments with varying dynamics, especially when the training needs to happen offline. This issue limits RL’s practical applications in real-world situations, where conditions may change unexpectedly. While transfer learning in RL usually seeks to tackle this issue by leveraging online data in both source and target environments, an open challenge remains: developing policies to new, unseen dynamics through adjusted offline training, purely with data from the source environments.

Some methods involve training agents across environments with varying dynamics to foster robustness and adaptability. Domain randomization [1] introduces environmental variations during training to improve robustness, facilitating transfer from simulation to reality. Information bottlenecks [2] can also be used for dynamics generalization by enabling agents to capture only the task-relevant information while ignoring irrelevant environmental variations.

Meta-reinforcement learning (Meta-RL) algorithms, such as Probabilistic Embeddings for Actor-Critic Reinforcement Learning (PEARL) [3] and Variational Bayes-Adaptive Deep RL (VariBAD) [4], learn latent task representations to facilitate rapid adaptation to new tasks through online interactions. A Gaussian Process-based approach has been proposed to transfer knowledge across

*The project this report is based on was supported with funds from the German Federal Ministry of Education and Research under project number 01IS24087A. The sole responsibility for the report’s contents lies with the authors.

systems with slight variations by modeling correlations between members, focusing on multi-task settings [5]. In contrast, our method employs sequence encoders to capture latent dynamics, ensuring generalization in offline reinforcement learning without requiring task-specific identifiers.

Sequence encoders, like autoencoders and recurrent neural networks, have been used to compress sequences into compact representations in RL. For example, the world models framework [6] learns latent representations of the environment to facilitate planning.

In offline RL, where learning is confined to static datasets, algorithms like Batch-Constrained Q-learning (BCQ) [7, 8] and Learning in Interactive Offline eNvironments (LION) [9] address challenges such as value function overestimation and out-of-distribution actions. However, these techniques do not directly tackle the transferability of policies to new environments with unobserved dynamics.

In this study, we investigate an offline method in RL to enhance robustness and generalization, concentrating on the classic cart-pole environment with variations in pole lengths and cart masses. We introduce a sequence encoder to capture latent representations of environment dynamics from state-action sequences. By adding these latent encodings to the state space, we aim to improve the agent’s adaptability to different environments. Our work provides three main contributions:

1. We introduce the use of sequence encoders, particularly autoencoders, to derive meaningful representations of environment dynamics from state-action sequences.
2. We demonstrate that adding these encodings to the state space enhances the transferability of policies to environments with new dynamics.
3. We compare our approach to a baseline method and demonstrate that our method enables a single policy to maintain strong performance across new environments after being trained solely on the source environments.

2 Experimental Setup




Fig. 1: Illustration of cart-poles with varying pole lengths and cart masses.

Environment Design: We use the classic *CartPole-v1* environment [10] as our testbed, introducing variability in dynamics by modifying its transition-defining parameters. Five source environments were created by altering pole length and

cart mass, with values drawn uniformly from the ranges $[0.1, 2.0]$ and $[0.5, 2.0]$, respectively. These parameters are set for each environment as follows:

$$\text{env_properties} = \left\{ \begin{array}{l} \text{pole length : } l_i \sim U(0.1, 2.0), \\ \text{cart mass : } m_i \sim U(0.5, 2.0) \end{array} \right\}. \quad (1)$$

The goal is to train a single policy using data from the source environments that can generalize to new environments without additional training. Our method is designed to achieve robust generalization through offline training on the source environments. Figure 2 displays the variability in pole length and cart mass across source environments and highlights the individual new environments N1 to N10 used for evaluation.




Fig. 2: Scatter plot showing pole length and cart mass for source environments (blue) and new environments (red).

Sequence Encoder: To capture the dynamics of each environment, we introduce the notion of a sequence encoder. For this, we used autoencoders (AE), however other architectures, like LSTMs or even MLPs are conceivable. The AE is trained to reconstruct sequences of state-action pairs, compressing the sequence into a latent representation. The latent encoding is derived from the bottleneck layer of the AE.

Our experiments indicated that the AE with a four-dimensional latent space produced the best performance, so we focus our analysis on this configuration.

Data Collection: We initially train a Deep Q-Network (DQN) [11] on the standard *CartPole-v1* environment, using an implementation from the authors of the BCQ algorithm [7]. This trained agent is then deployed on a number of source environments with randomly sampled dynamics, as visualized in Figure 2.

We collect interaction data from this agent, storing transitions in a regular replay buffer. For the autoencoder training, we use sequences of state-action pairs of length 16 (comprising 16 states and 15 actions in between), generated by the DQN agent trained on the regular cart-pole environment.

Sequence Encoder Training: The AE is trained to reconstruct input sequences, compressing each sequence into a four-dimensional latent space. The training objective is to minimize the mean squared error between input and reconstructed sequences.

State Space Augmentation: After training the encoder, the state space of each source environment is augmented. For each source environment in our system that we used for training, five trajectories of length 16 are taken from the buffer and encoded. The average latent encoding is then computed, creating a four-dimensional vector that reflects the dynamics of each environment.

This encoding is appended to the original states, expanding the state space. To avoid additional data collection from the source environments, we retroactively extend states in the existing dataset by appending the relevant encodings.

Offline RL Training with BCQ: We employ the BCQ algorithm in its discrete form to train policies exclusively using pre-collected datasets from the source environments, without any additional interactions with the environments during training. We conduct the following experiments:

1. **Baseline:** BCQ trained on the data from the source environments without augmented states.
2. **Trajectory Encoding Augmentation (TEA):** BCQ trained on the data from the source environments with augmented states using learned encodings of trajectories.

Evaluation Protocol: To assess policy transferability, we measure each policy’s performance on new environment(s) which were randomly sampled from the same distribution as the source environments and are visualized in Figure 2, reporting average returns as well as standard error over 100 random seeds after 20,000 timesteps of training. Note that five sequences of length 16 were needed to be seen from these new environments to facilitate the creation of a first encoding.

3 Discussion

Results: The performance of each experiment on the new environment(s) after 20,000 timesteps of training on the source environments is shown in Table 1. In Figure 3 we see the performance of trajectory encoding augmentation (TEA) divided by the baseline performance for each new environment.

Environment	Baseline	TEA
N1	36.2 ± 2.2	72.1 ± 6.6
N2	38.4 ± 0.2	45.3 ± 2.2
N3	43.3 ± 1.1	50.1 ± 2.9
N4	49.8 ± 0.7	55.7 ± 2.3
N5	36.3 ± 2.2	72.5 ± 6.7
N6	35.1 ± 0.6	35.2 ± 1.5
N7	47.3 ± 0.8	47.7 ± 1.7
N8	101.5 ± 1.4	145.5 ± 7.4
N9	58.9 ± 1.9	60.7 ± 2.5
N10	40.3 ± 0.9	52.6 ± 13.7
Average	48.7 ± 0.4	63.7 ± 1.9

Table 1: Performance Across N1 to N10 (mean \pm SEM over 100 seeds)




Fig. 3: Plot illustrating the performance improvement achieved through TEA across various new environments, compared to the baseline performance. Notably, all performance ratios exceed 1, demonstrating consistent gains.

Analysis: The results in Table 1 demonstrate a statistically significant improvement in performance for the TEA method over the baseline condition, as indicated by the uncertainties accompanying the average values. These findings show that augmenting the state space with latent encodings, derived from a sequence encoder, enables the trained policy to generalize more effectively across environments with new dynamics. The observed results confirm that the approach based on latent encodings leads to better outcomes under changing dynamics, suggesting that this form of augmentation enhances policy transferability compared to the baseline.

4 Conclusion

In this paper, we addressed the problem of enabling quick adaptation to new environments with different dynamics using offline pretraining from source environments. Our proposed Trajectory Encoding Augmentation (TEA) method demonstrates that using sequence encoders to augment the state space in offline RL can notably improve policy transferability to respective new environments.

By utilizing latent representations of environment-trajectories, agents perform well in new environments.

This technique offers a promising path toward developing RL agents more resilient to changing environmental dynamics.

Future Work: Our method can be expanded to explore the *offline-to-online* learning paradigm. Specifically, we plan to initiate training using BCQ and subsequently continue with regular Q-learning. This extension would enable us to examine the agent’s adaptability during the shift from offline to online learning, yielding deeper insights into the effectiveness of our approach and offering a more comprehensive perspective on its application in dynamic environments. In addition, exploring environments with higher-dimensional state-action spaces and experimenting with alternative encoder architectures could further validate and enhance the versatility of our approach.

References

- [1] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30. IEEE, 2017.
- [2] Xingyu Lu, Kimin Lee, Pieter Abbeel, and Stas Tiomkin. Dynamics generalization via information bottleneck in deep reinforcement learning. *arXiv preprint arXiv:2008.00614*, 2020.
- [3] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International Conference on Machine Learning*, pages 5331–5340. PMLR, 2019.
- [4] Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. VariBAD: a very good method for Bayes-adaptive deep RL via meta-learning. In *International Conference on Learning Representations*, 2020.
- [5] Timothy Verstraeten, Pieter Libin, and Ann Nowe. Fleet control using coregionalized gaussian process policy iteration. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 1571–1578. IOS Press, August 2020.
- [6] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems*, 31, 2018.
- [7] Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pages 2052–2062, 2019.
- [8] Scott Fujimoto, Edoardo Conti, Mohammad Ghavamzadeh, and Joelle Pineau. Benchmarking batch deep reinforcement learning algorithms. *arXiv preprint arXiv:1910.01708*, 2019.
- [9] Phillip Swazinna, Steffen Udfluft, and Thomas Runkler. User-interactive offline reinforcement learning. In *The 11th International Conference on Learning Representations*, 2023.
- [10] Andrew G Barto, Richard S Sutton, and Charles W Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, (5):834–846, 1983.
- [11] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.