

Environment Transformer and Policy Optimization for Model-Based Offline Reinforcement Learning

Pengqin Wang, Meixin Zhu[†], and Shaojie Shen

Abstract—Interacting with the actual environment to acquire data is often costly and time-consuming in robotic tasks. Model-based offline reinforcement learning (RL) provides a feasible solution. On the one hand, it eliminates the requirements of interaction with the actual environment. On the other hand, it learns the transition dynamics and reward function from the offline datasets and generates simulated rollouts to accelerate training. Previous model-based offline RL methods adopt probabilistic ensemble neural networks (NN) to model aleatoric uncertainty and epistemic uncertainty. However, this results in a great increase in training time and computing resource requirements. Furthermore, these methods are easily disturbed by the accumulative errors of the environment dynamics models when simulating long-term rollouts. To solve the above problems, we propose an uncertainty-aware sequence modeling architecture called Environment Transformer. It models the probability distribution of the environment dynamics and reward function to capture aleatoric uncertainty and treats epistemic uncertainty as a learnable noise parameter. Benefiting from the accurate modeling of the transition dynamics and reward function, Environment Transformer can be combined with arbitrary planning, dynamics programming, or policy optimization algorithms for offline RL. In this case, we perform Conservative Q-Learning (CQL) to learn a conservative Q-function. Through simulation experiments, we demonstrate that our method achieves or exceeds state-of-the-art performance in widely studied offline RL benchmarks. Moreover, we show that Environment Transformer’s simulated rollout quality, sample efficiency, and long-term rollout simulation capability are superior to those of previous model-based offline RL methods.

I. INTRODUCTION

Deep reinforcement learning (RL) has obtained significant achievements in a variety of domains by utilizing a great deal of interactions with the environment [1, 2]. However, due to the high expense of online data collection, the trial-and-error method is typically impractical in numerous real-world scenarios such as autonomous driving, robot manipulation, and aerial vehicles [3]–[8]. Offline RL aims to solve the problem of learning a policy completely from a fixed batch of data without interacting with the environment [9]–[11]. This provides an appealing paradigm for a wide range of applications where there exist large and diverse pre-recorded

datasets. Recent studies have shown that model-based RL, which learns the transition dynamics from the batch of data, demonstrates better generalization capability in dealing with offline RL tasks [14]–[18]. Previous model-based offline RL approaches [13]–[18] use probabilistic ensemble NN to capture both aleatoric uncertainty (inherent system stochasticity) and epistemic uncertainty (subjective uncertainty, due to limited data). Nevertheless, this leads to a significant growth in the duration of training and the need for computational resources. Furthermore, the effectiveness of these approaches might be significantly affected by the accumulative errors of the environment dynamics models when simulating long-term rollouts.

To overcome the above issues, we consider the accurate modeling of the transition dynamics and reward function as a sequence-to-sequence task. We propose an uncertainty-aware sequence modeling architecture called Environment Transformer. It models the probability distribution of the environment dynamics and reward function to capture aleatoric uncertainty and treats epistemic uncertainty as a learnable noise parameter. The state-action pairs are sampled from the offline dataset to predict the probability distribution of future state-reward pairs using a causal self-attention mask [31, 32]. Furthermore, Environment Transformer can be combined with arbitrary planning, dynamics programming, or policy learning algorithms for offline RL, thanks to its accurate modeling of the transition dynamics and reward function. In this instance, we conduct Conservative Q-Learning (CQL) [19] to learn a conservative Q-function, as shown in Fig. 1.

We compare the final performance of our proposed method with both model-based and model-free state-of-the-art (SOTA) offline RL methods in widely studied robot continuous control benchmarks [35, 36]. Through simulation experiments, we show that our method achieves the highest score in 9 of 20 tasks and similar performance to SOTA in 8 tasks. In addition, we demonstrate that the simulated rollout quality, sample efficiency, and long-sequence simulation error of Environment Transformer are superior to those of previous model-based offline RL methods.

We summarize the contributions of this paper as follows.

- We propose Environment Transformer, an uncertainty-aware sequence modeling architecture. It models the probability distribution of the environment dynamics and reward function to account for aleatoric uncertainty and considers epistemic uncertainty as a learnable noise parameter. CQL is conducted based on Environment Transformer for offline RL tasks.

Pengqin Wang is with the Division of Emerging Interdisciplinary Areas, the Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail: pwangas@connect.ust.hk).

Meixin Zhu is with the Systems Hub at the Hong Kong University of Science and Technology (Guangzhou), the Civil and Environmental Engineering Department at the Hong Kong University of Science and Technology and Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (e-mail: meixin@ust.hk).

Shaojie Shen is with the Department of Electronic and Computer Engineering, the Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail: eeshaojie@ust.hk).

[†] Corresponding author

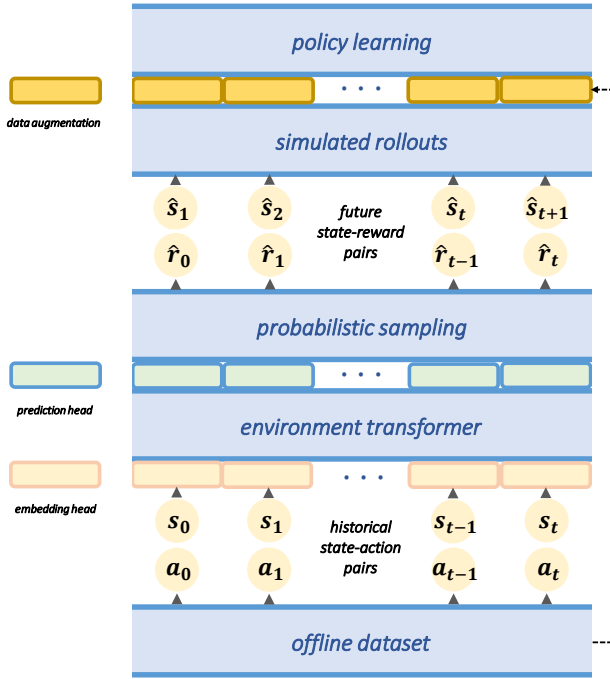


Fig. 1: The framework of the proposed Environment Transformer.

- Simulation experiments are performed on widely studied offline RL benchmarks. The experiment results show that our method achieves the highest score in 9 of 20 tasks and similar performance to SOTA in 8 tasks. Moreover, we demonstrate that Environment Transformer’s simulated rollout quality, sample efficiency, and long-sequence simulation error are superior to those of previous model-based offline RL methods. Therefore, it provides a low-cost, high-efficiency data acquisition method for training real-world robots, with a wide range of application prospects.

II. RELATED WORK

A. Offline RL

The research of Offline RL addresses the challenge of learning a policy from a fixed dataset instead of interacting with the environment. The major issue of Offline RL is extrapolation error [10], which is a generalization error in estimating value functions caused by out-of-distribution actions. One potential approach to solving this problem is to integrate a pessimistic bias for unseen actions into the Q-function. Kumar et al. [19] propose to learn a conservative Q-function such that the expected value of a policy under the Q-function lower-bounds its true value. Kostrikov et al. [24] treat the state value function as a random variable whose randomness is determined by the action, and use the upper expectile to estimate the value of the best actions. An alternative approach is to restrict the policy optimization to remain close to the original data samples. Wang et al. [20] utilize a form of critic-regularized regression to learn policies from data. Nair et al. [21] combine sample efficient dynamic

programming with maximum likelihood policy updates. Wu et al. [22] propose a framework to generalize previous approaches to solve the offline RL problem by regularizing the behavior policy. Zhou et al. [23] propose to learn the policy in the latent action space, which constrains the policy to select actions within the support of the dataset. Xu et al. [25] construct a novel offline-applicable policy learning goal that corresponds to the advantage function value of the behavior policy, multiplied by a state-marginal density ratio. Fujimoto et al. [26] include a behavior cloning component in the policy update of an online RL algorithm and normalize the data to push the policy towards favoring actions in the dataset.

B. Dynamics Modeling

There exists a wealth of prior works to learn the environmental dynamics. Sutton [27] proposes an architecture that maintains a dynamics model of the agent’s transitions. Deisenroth et al. [28] model the environment dynamics as Gaussian processes and incorporate model uncertainty into long-term planning. Levine et al. [29] leverage local linear models to represent the environment dynamics. Chua et al. [13] integrate deep network dynamics models that account for uncertainty with sampling-based uncertainty propagation. Janner et al. [12, 16] use bootstrap ensembles of predictive models to capture aleatoric uncertainty and epistemic uncertainty of the environment dynamics.

Recent breakthroughs in sequence modeling using deep neural networks have resulted in fast improvements in long-horizon prediction accuracy and model efficiency [30]–[32]. Chen et al. [33, 34] consider RL to be a sequence modeling problem, which aims to generate a series of actions to receive high rewards. In our approach, environment modeling is viewed as a sequence-to-sequence problem, which is required to predict the probability distribution for future state-reward pairs based on historical state-action pairs.

III. PRELIMINARIES

A. Transformer

Transformers are proposed by [31] as a framework for effectively modeling sequential data, which consists of stacked encoders and decoders. Both of them are based on attention mechanisms, where the i -th input token x_i is embedded and mapped to key k_i , query q_i and value v_i , and the i -th output token can be represented as:

$$z_i = \sum_{j=1}^n \text{softmax} \left(\left\{ \langle q_i, k_{j'} \rangle \right\}_{j'=1}^n \right)_j \cdot v_j. \quad (1)$$

The generative pre-training transformer (GPT) is proposed by [32], where the encoder-decoder architecture is changed to a causal self-attention mask without any encoders. In our practice, we adopt the GPT framework and develop Environment Transformer to obtain accurate predictions for environment dynamics and reward function.

B. Conservative Q-Learning

We consider a Markov decision process (MDP), defined by the tuple $(\mathcal{S}, \mathcal{A}, T, r, \gamma)$. \mathcal{S} and \mathcal{A} represent state and action space. We consider state space and action space to be both continuous. The transition dynamics is denoted as $T(s'|s, a)$, and the reward function is written as $r(s, a)$, and $\gamma \in (0, 1)$ represents the discount factor.

The goal of standard RL algorithms is to obtain the optimal policy π^* such that the cumulative reward is maximized. The optimal policy π^* can be written as:

$$\pi^* = \arg \max_{\pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]. \quad (2)$$

Q-learning methods maintain a Q-function $Q(s, a)$ that measures the discounted return based on the state s and action a , under current policy π . Given current policy $\pi(a|s)$, the Bellman backup for obtaining the corresponding Q function gives:

$$\mathcal{B}^{\pi} Q(s, a) := r(s, a) + \gamma \mathbb{E}_{s' \sim T(s'|s, a), a' \sim \pi(a'|s')} [Q(s', a')]. \quad (3)$$

The optimal policy's Q-function satisfies the following Bellman optimal operator:

$$\mathcal{B}^* Q(s, a) := r(s, a) + \gamma \mathbb{E}_{s' \sim T(s'|s, a)} \left[\max_{a' \in \mathcal{A}} Q(s', a') \right]. \quad (4)$$

Online interaction with the environment is impractical in offline RL settings. Only previously collected datasets $\mathcal{D} = \{s_t, a_t, r_t, s_{t+1}, d_t\}_{t=1}^N$ are accessible, where d is the terminal flag showing whether the episode is ended. CQL approach [19] learns a conservative Q-function such that the expected value of a policy under this Q-function lower-bounds its true value, which gives the following iterative update for training:

$$\begin{aligned} \min_Q \max_{\mu} \alpha & \left(\mathbb{E}_{s \sim \mathcal{D}, a \sim \mu(a|s)} [Q(s, a)] - \right. \\ & \mathbb{E}_{s \sim \mathcal{D}, a \sim \hat{\pi}_{\beta}(a|s)} [Q(s, a)] + \\ & \left. \frac{1}{2} \mathbb{E}_{s, a, s' \sim \mathcal{D}} \left[\left(Q(s, a) - \hat{\mathcal{B}}^{\pi_k} \hat{Q}^k(s, a) \right)^2 \right] + \right. \\ & \left. \mathcal{R}(\mu), \right) \end{aligned} \quad (5)$$

where $\hat{\pi}_{\beta}(a|s)$ represents the data distribution and π_k is the policy derived from the Q-function. \mathcal{D} is the collected offline dataset augmented by the simulated rollouts and \mathcal{R} is a regularizer.

IV. ENVIRONMENT TRANSFORMER

A. Environment Modeling

We consider the environment dynamics as a Gaussian distribution with diagonal covariance to capture aleatoric uncertainty and treats epistemic uncertainty as a learnable noise parameter:

TABLE I: Hyperparameters of Environment Transformer

Hyperparameter	Value
Number of layers	8
Number of attention heads	16
Embedding dimension	1024
Nonlinearity function	ReLU
Batch size	16
Sequence length	100
Dropout	0.1
Learning rate	10^{-4}
Weight decay	10^{-4}
Optimizer	HybridAdam
HybridAdam eps	10^{-4}

$$Pr(s_{t+1}, r_t | s_t, a_t) = \mathcal{N}(\mu_{au}(s_t, a_t), \Sigma_{au}(s_t, a_t)) + \epsilon_t, \quad (6)$$

where (s_t, a_t) denotes the state-action pairs at time step t . $\mu_{au}(s_t, a_t)$ and $\Sigma_{au}(s_t, a_t)$ represent the mean and covariance of the Gaussian distribution considering aleatoric uncertainty at time step t , respectively. ϵ_t is a learnable parameter for epistemic uncertainty at time step t .

B. Training

The training inputs for Environment Transformer are state-action pairs from the offline datasets containing thousands of trajectories. The i -th trajectory training input τ^i can be represented as:

$$\tau^i = \{(s_t^i, a_t^i)\}_{t=0}^{T-1}, \quad (7)$$

where t denotes the timestep of state-action pairs, i is training trajectory index and T is the sequence length.

In practice, we consider the epistemic uncertainty parameter ϵ_t to be a sample from a Gaussian distribution with zero mean and learnable covariance conditioned on input state-action pairs. After the prediction and probabilistic sampling, the output of Environment Transformer ϕ^i can be written as:

$$f_t^i \sim \mathcal{N}(\mu_{au}(s_t^i, a_t^i), \Sigma_{au}(s_t^i, a_t^i)), \quad (8)$$

$$\epsilon_t^i \sim \mathcal{N}(0, \Sigma_{eu}(s_t^i, a_t^i)), \quad (9)$$

$$\phi^i = \{(f_t^i + \epsilon_t^i)\}_{t=0}^{T-1}, \quad (10)$$

where $\Sigma_{eu}(s_t^i, a_t^i)$ represents the learnable covariance of the Gaussian distribution considering epistemic uncertainty conditioned on the input state-action pair (s_t^i, a_t^i) . $\mu_{au}(s_t^i, a_t^i)$ and $\Sigma_{au}(s_t^i, a_t^i)$ represent the mean and covariance of the Gaussian distribution considering aleatoric uncertainty at time step t of the i -th trajectory.

After obtaining the predicted future state-reward pairs for the i -th trajectory, we can calculate the mean square error (mse) loss between the predictions and ground-truth:

$$loss = mse(\phi^i, \{(s_{t+1}^i, r_t^i)\}_{t=0}^{T-1}). \quad (11)$$




Fig. 2: The robot continuous control offline RL benchmarks including ant, halfcheetah, hopper and walker2d.

Environment	Dataset	Ours	MOPO	RAMBO	COMBO	CQL	IQL
HalfCheetah	Full-Replay	80.8±0.4	86.4	89.0	-	-	-
	Medium-Expert	96.4±0.6	63.3	79.3	90.0	62.4	86.7
	Medium-Replay	45.1±0.9	53.1	67.0	55.1	46.2	44.2
	Medium	49.5±0.1	42.3	71.0	54.2	44.4	47.4
	Random	22.7±2.1	35.4	33.5	38.8	35.4	-
Hopper	Full-Replay	106.9±0.9	108.1	107.6	-	-	-
	Medium-Expert	106.9±2.7	23.7	89.5	111.1	111.0	91.5
	Medium-Replay	94.9±6.3	67.5	97.6	73.1	48.6	94.7
	Medium	67.3±6.5	28.0	91.2	94.9	86.6	66.3
	Random	10.1±0.7	11.7	15.5	17.9	10.8	-
Walker2d	Full-Replay	97.5±1.2	56.9	52.6	-	-	-
	Medium-Expert	84.7±2.4	44.6	63.1	96.1	98.7	109.6
	Medium-Replay	88.9±4.6	39.0	88.5	56.0	32.6	73.9
	Medium	76.6±3.1	17.8	89.1	75.5	74.5	78.3
	Random	14.2±1.7	13.6	0.2	7.0	7.0	-
Ant	Full-Replay	138.4±0.8	27.8	119.3	-	-	-
	Medium-Expert	141.5±2.3	26.8	95.7	-	-	-
	Medium-Replay	105.7±1.9	28.4	49.7	-	-	-
	Medium	116.1±0.5	20.4	67.0	-	-	-
	Random	32.6±8.2	15.8	29.8	-	-	-
Average		78.8	40.5	69.8	64.1	54.9	77.0

TABLE II: Comparisons on Offline RL Benchmarks

We train Environment Transformer on 4 Nvidia RTX 3080 10GB and Intel(R) Xeon(R) Platinum 8255C CPU. We use Gemini, the heterogeneous memory space manager of Colossal-AI [40], to train Environment Transformer. The hyperparameters are listed in Table I.

V. EXPERIMENTS

In this section, we design simulation experiments on widely studied offline RL benchmarks [37]–[39] to evaluate the performance with both model-based and model-free SOTA offline RL algorithms. Then we compare the simulated rollout quality and sample efficiency with previous model-based offline RL methods based on probabilistic ensemble NN. Finally, we evaluate the model’s capacity for long-term forecasting using both offline datasets and online environments. The benchmarks include four environments (Ant, HalfCheetah, Hopper, and Walker2d) and five dataset

types (full-replay, medium-expert, medium-replay, medium, and random), as shown in Fig. 2. The Ant is a 3D robot consisting of one torso with four legs, whose goal is to coordinate the four legs to move forward. The HalfCheetah is a two-dimensional robot. The objective is to apply torque to its joints so that the robot can run forward as quickly as possible. The Hopper is a one-legged, two-dimensional figure. Similarly, the goal is to make forward-moving jumps. The Walker2d is a two-dimensional figure with two legs, whose goal is to control the forward movement.

A. Comparisons on Offline RL Benchmarks

We evaluate the final performance after $1e6$ steps policy learning, in comparison with model-based approaches MOPO, COMBO and RAMBO [16]–[18], and model-free methods CQL and IQL [19, 24]. We report the mean and variance over three random seeds. The results are shown in Table II. Our approach is the strongest by a significant margin on all datasets in Ant environment with 111-dimensional

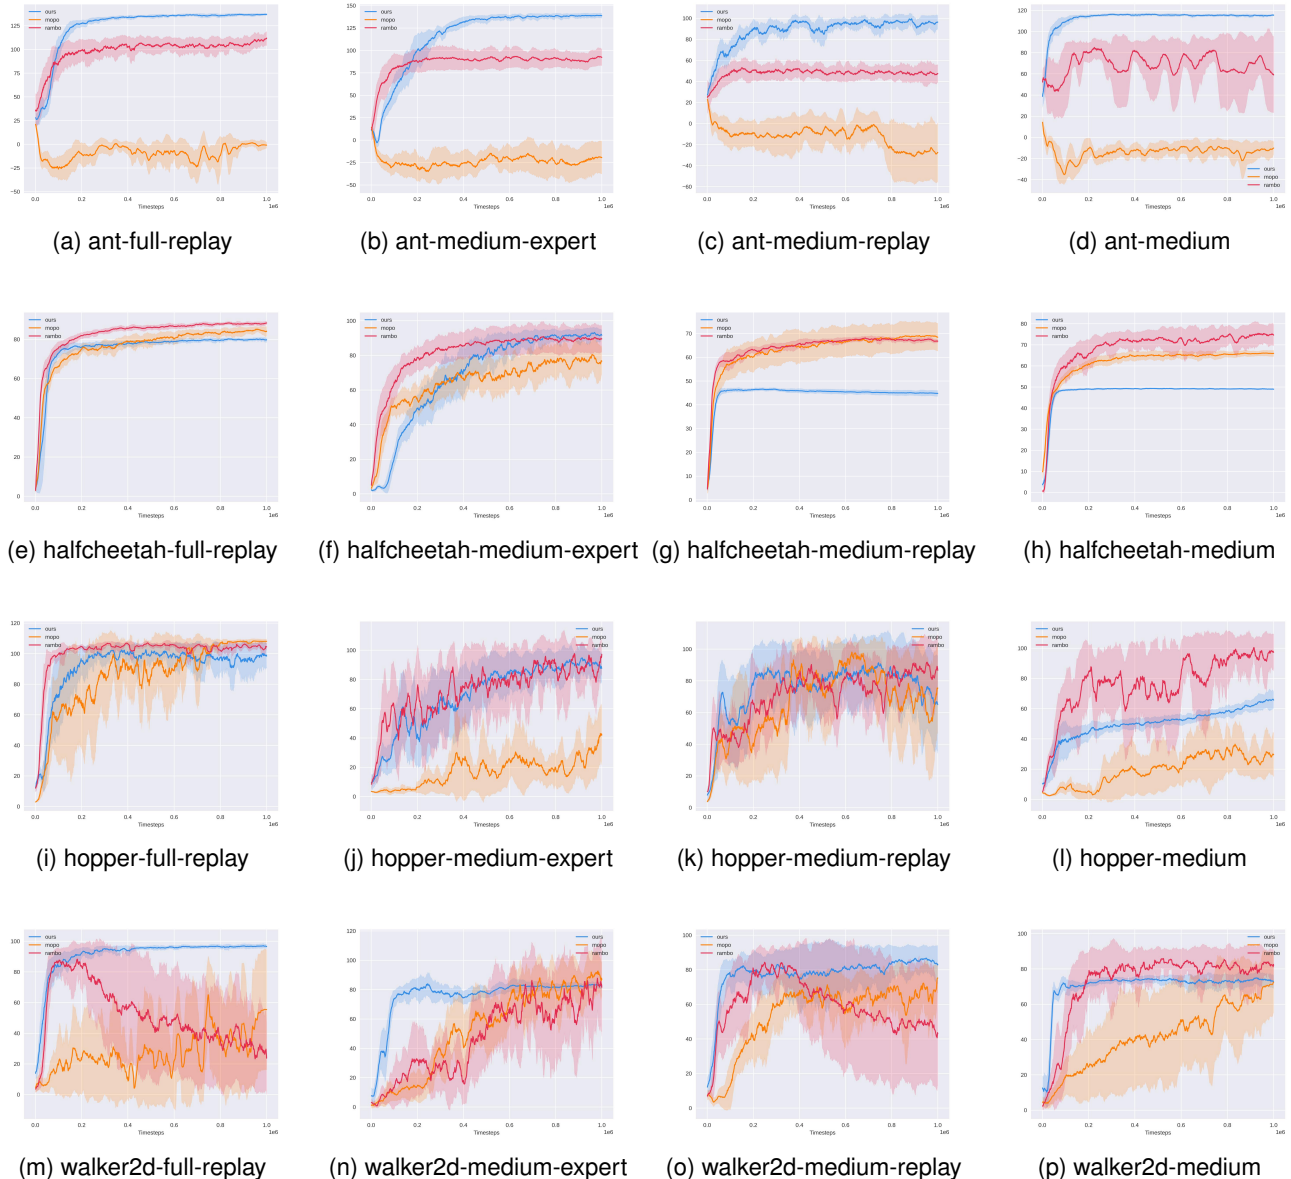
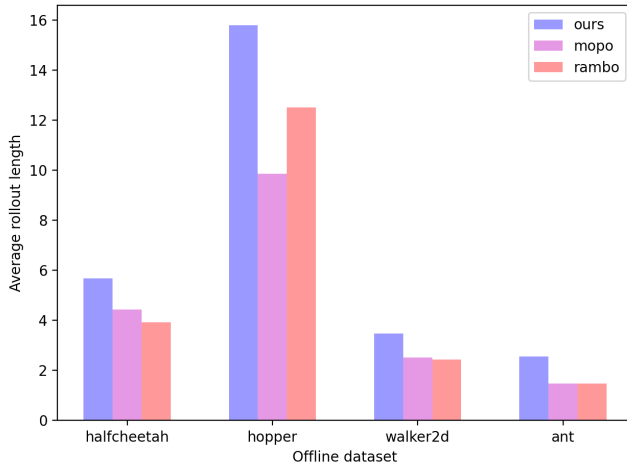


Fig. 3: The episode reward curves during policy learning with three random seeds.

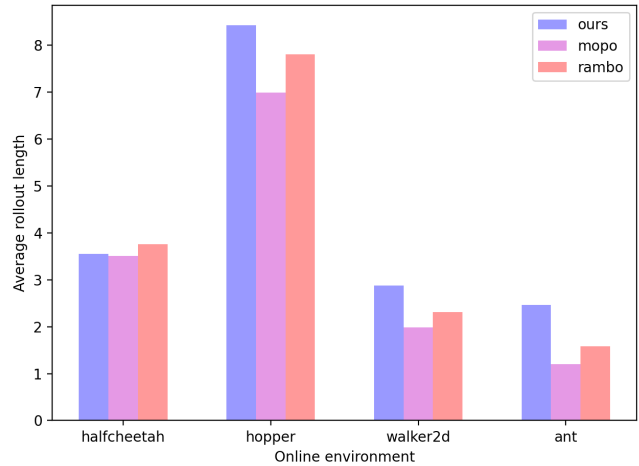
observation space, which demonstrates that Environment Transformer can accurately model the aleatoric and epistemic uncertainty for complex environments. Our method achieves the strongest performance of full-replay, medium-replay, and random datasets in Walker2d environment because it is the environment with the second-highest complexity. However, our approach performs less well in Hopper environment whose observation space is only 11 dimensional. We analyze the reason that for simpler cases the data utilization efficiency of the Environment Transformer is not as good as smaller, lower-capacity models such as ensemble dynamics. On the whole, our method achieves the highest score in 9 of 20 tasks and similar performance to SOTA in 8 tasks. In conclusion, our approach performs the best among both model-based and model-free SOTA offline RL algorithms.

B. Simulated Rollout Quality and Sample Efficiency

We visualize the episode reward curves during $1e6$ policy learning, compared with SOTA model-based offline RL algorithms MOPO and RAMBO, as shown in Fig. 3. Due to the complexity of the Ant environment, MOPO and RAMBO can not accurately model the environment, which leads to the failure to achieve stable policy improvement. However, the strategies trained by our approach all meet or far exceed expert performance, which demonstrates that Environment Transformer’s simulated rollout quality is much greater than the ensemble dynamics models of MOPO and RAMBO. Our approach shows the highest or similar performance with the lowest variance and fastest converging speed in most tasks, which proves our sample efficiency. Our method performs less well on medium datasets. The main reason is that conservative Q-Learning suffers from the lack of action



(a) Long-term evaluation in offline datasets



(b) Long-term evaluation in online environments

Fig. 4: The long-term rollout evaluation using both offline datasets and online environments.

diversity in medium datasets. As a result, learning a policy that generalizes well becomes more difficult.

C. Long-term Rollout Evaluation

We design offline and online experiments to evaluate the long-term rollout generation capability for Environment Transformer and dynamics models of MOPO and RAMBO. In the offline setting, we random sample a long-duration (usually one thousand steps) trajectory from the datasets and then perform rollout simulation for each state-action pair on the trajectory. We record the rollout length when the mse between the normalized predicted state-reward pairs and ground-truth is greater than a threshold. We repeat the above steps hundreds of times and calculate the average for all datasets within the same environment. In the online setting, we run random, medium, and expert policies respectively to collect transitions. Similarly, we perform rollout simulation for each state-action pair, update the online environment and calculate the mse between the normalized predictions and ground-truth. We repeat the rollout simulation procedure until the mse is greater than a threshold and record the episode length. At last, we calculate the average for different policies under the same environment. The threshold is 0.01 for both offline and online settings. The results of the long-term rollout evaluation are shown in Fig. 4. The average rollout length of our method exceeds MOPO and RAMBO in 7 of 8 environments, which demonstrates that the long-term rollout simulation ability of Environment Transformer is superior to previous model-based offline RL methods.

VI. CONCLUSIONS

In this paper, we study model-based offline RL approaches. Existing model-based offline RL works utilize probabilistic ensemble NN to capture aleatoric and epistemic uncertainty, which leads to an great increase in training time and computing resource requirements. Furthermore, these methods suffer from the accumulative errors when simulating

long-duration rollouts. In order to solve the above issues, we propose Environment Transformer, an uncertainty-aware sequence modeling architecture. It models the probability distribution of the transition dynamics and reward function to capture aleatoric uncertainty and considers epistemic uncertainty as a learnable noise parameter. Thanks to the accurate modeling of environment dynamics and reward function, it can be combined with arbitrary planning, dynamic programming, or policy learning methods. In this case, We conduct Conservative Q-Learning (CQL) to learn a conservative Q-function based on Environment Transformer for offline RL tasks.

We design simulation experiments on widely studied offline RL benchmarks to evaluate the performance. The experimental results show that our method achieves or exceeds SOTA performance of both model-based and model-free offline RL algorithms. Moreover, we demonstrate that Environment Transformer’s simulated rollout quality, sample efficiency, and long-term rollout simulation capability are superior to those of previous model-based offline RL methods.

Environment Transformer provides a new paradigm for merging offline RL and robot tasks, as it offers a low-cost and high-efficiency data acquisition approach for training real-world robots. we believe that it has broad applicability potential. The limitation of our method is insufficient real-world tests. In the future, we intend to optimize this work further and implement more difficult real-world evaluations.

ACKNOWLEDGMENT

The authors would like to thank Prof. Barto, Michael Janner, and Costa Huang for insightful discussions.

This study is supported by the National Natural Science Foundation of China under Grant 52302379, Guangzhou Basic and Applied Basic Research Projects under Grants 2023A03J0106 and 2024A04J4290, Guangdong Province General Universities Youth Innovative Talents Project under Grant 2023KQNCX100, and Guangzhou Municipal Science and Technology Project 2023A03J0011.

REFERENCES

- [1] D. Silver et al., “Mastering the game of Go with deep neural networks and tree search,” *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [2] O. Vinyals et al., “Grandmaster level in StarCraft II using multi-agent reinforcement learning,” *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [3] A. X. Lee et al., “How to Spend Your Robot Time: Bridging Kick-starting and Offline Reinforcement Learning for Vision-based Robotic Manipulation,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [4] J. Li, C. Tang, M. Tomizuka, and W. Zhan, “Hierarchical planning through goal-conditioned offline reinforcement learning,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10216–10223, 2022.
- [5] J. Jin, D. Graves, C. Haigh, J. Luo and M. Jagersand, “Offline learning of counterfactual predictions for real-world robotic reinforcement learning,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 3616–3623.
- [6] T. Z. Zhao et al., “Offline meta-reinforcement learning for industrial insertion,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 6386–6393.
- [7] F. Gao, L. Wang, B. Zhou, X. Zhou, J. Pan, and S. Shen, “Teach-repeat-replan: A complete and robust system for aggressive flight in complex environments,” *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1526–1545, 2020.
- [8] B. Zhou, H. Xu, and S. Shen, “Racer: Rapid collaborative exploration with a decentralized multi-uav system,” *IEEE Transactions on Robotics*, 2023.
- [9] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International conference on machine learning*, 2018, pp. 1861–1870.
- [10] S. Fujimoto, D. Meger, and D. Precup, “Off-Policy Deep Reinforcement Learning without Exploration,” in *International Conference on Machine Learning*, 2019, pp. 2052–2062.
- [11] S. Lange, T. Gabel, and M. Riedmiller, “Batch Reinforcement Learning,” in *Reinforcement Learning*, 2012, pp. 45–73.
- [12] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to Trust Your Model: Model-Based Policy Optimization,” in *Advances in Neural Information Processing Systems*, 2019.
- [13] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models,” in *Advances in Neural Information Processing Systems*, 2018, vol. 31.
- [14] J. Wang, W. Li, H. Jiang, G. Zhu, S. Li, and C. Zhang, “Offline reinforcement learning with reverse model-based imagination,” in *Advances in Neural Information Processing Systems*, 2021, vol. 34, pp. 29420–29432.
- [15] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims, “Morel: Model-based offline reinforcement learning,” in *Advances in neural information processing systems*, 2020, vol. 33, pp. 21810–21823.
- [16] T. Yu et al., “MOPO: Model-based Offline Policy Optimization,” *arXiv preprint arXiv:2005.13239*, 2020.
- [17] T. Yu, A. Kumar, R. Rafailov, A. Rajeswaran, S. Levine, and C. Finn, “COMBO: Conservative Offline Model-Based Policy Optimization,” in *Advances in Neural Information Processing Systems*, 2021, vol. 34, pp. 28954–28967.
- [18] M. Rigter, B. Lacerda, and N. Hawes, “Rambo-rl: Robust adversarial model-based offline reinforcement learning,” in *Advances in neural information processing systems*, vol. 35, pp. 16082–16097, 2022.
- [19] A. Kumar, A. Zhou, G. Tucker, and S. Levine, “Conservative Q-Learning for Offline Reinforcement Learning,” *arXiv preprint arXiv:2006.04779*, 2020.
- [20] Z. Wang et al., “Critic Regularized Regression,” in *Advances in Neural Information Processing Systems*, 2020, vol. 33, pp. 7768–7778.
- [21] A. Nair, A. Gupta, M. Dalal, and S. Levine, “AWAC: Accelerating Online Reinforcement Learning with Offline Datasets,” *arXiv preprint arXiv:2006.09359*, 2020.
- [22] Y. Wu, G. Tucker, and O. Nachum, “Behavior regularized offline reinforcement learning,” *arXiv preprint arXiv:1911.11361*, 2019.
- [23] W. Zhou, S. Bajracharya, and D. Held, “PLAS: Latent Action Space for Offline Reinforcement Learning,” *arXiv preprint arXiv:2011.07213*, 2020.
- [24] I. Kostrikov, A. Nair, and S. Levine, “Offline Reinforcement Learning with Implicit Q-Learning,” *arXiv preprint arXiv:2110.06169*, 2021.
- [25] H. Xu, X. Zhan, J. Li, and H. Yin, “Offline Reinforcement Learning with Soft Behavior Regularization,” *arXiv preprint arXiv:2110.07395*, 2021.
- [26] S. Fujimoto and S. (shane) Gu, “A Minimalist Approach to Offline Reinforcement Learning,” in *Advances in Neural Information Processing Systems*, 2021, vol. 34, pp. 20132–20145.
- [27] R. S. Sutton, Dyna, “an Integrated Architecture for Learning, Planning, and Reacting,” *SIGART Bull.*, vol. 2, no. 4, pp. 160–163, Jul. 1991.
- [28] M. P. Deisenroth and C. E. Rasmussen, “PILCO: A Model-Based and Data-Efficient Approach to Policy Search,” in *Proceedings of the 28th International Conference on Machine Learning*, 2011, pp. 465–472.
- [29] S. Levine and V. Koltun, “Guided Policy Search,” in *Proceedings of the 30th International Conference on Machine Learning*, 2013, vol. 28, pp. 1–9.
- [30] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to Sequence Learning with Neural Networks,” in *Advances in Neural Information Processing Systems*, 2014, vol. 27.
- [31] A. Vaswani et al., “Attention is All you Need,” in *Advances in Neural Information Processing Systems*, 2017, vol. 30.
- [32] A. Radford, K. Narasimhan, T. Salimans and I. Sutskever, “Improving language understanding by generative pre-training,” 2018.
- [33] L. Chen et al., “Decision Transformer: Reinforcement Learning via Sequence Modeling,” *arXiv preprint arXiv:2106.01345*, 2021.
- [34] M. Janner, Q. Li, and S. Levine, “Offline Reinforcement Learning as One Big Sequence Modeling Problem,” in *Advances in Neural Information Processing Systems*, 2021, vol. 34, pp. 1273–1286.
- [35] J. Fu, A. Kumar, O. Nachum, G. Tucker, and S. Levine, “D4rl: Datasets for deep data-driven reinforcement learning,” *arXiv preprint arXiv:2004.07219*, 2020.
- [36] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ international conference on intelligent robots and systems*, 2012, pp. 5026–5033.
- [37] P. Wawrzyński, “A cat-like robot real-time learning to run,” in *Adaptive and Natural Computing Algorithms: 9th International Conference*, 2009, pp. 380–390.
- [38] T. Erez, Y. Tassa, and E. Todorov, “Infinite-horizon model predictive control for periodic tasks with contacts,” *Robotics: Science and systems VII*, pp. 73–80, 2012.
- [39] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, “High-dimensional continuous control using generalized advantage estimation,” *arXiv preprint arXiv:1506.02438*, 2015.
- [40] S. Li et al., “Colossal-AI: A unified deep learning system for large-scale parallel training,” *arXiv preprint arXiv:2110.14883*, 2021.