

Puzzle it Out: Local-to-Global World Model for Offline Multi-Agent Reinforcement Learning

Sijia Li

The Hong Kong University of Science and Technology
Hong Kong, China
slifg@connect.ust.hk

Shibo Chen*

South China University of Technology
Guangdong, China
shibochen.ustc@gmail.com

Xinran Li

The Hong Kong University of Science and Technology
Hong Kong, China
xinran.li@connect.ust.hk

Jun Zhang*

The Hong Kong University of Science and Technology
Hong Kong, China
eejzhang@ust.hk

ABSTRACT

Offline multi-agent reinforcement learning (MARL) aims to solve cooperative decision-making problems in multi-agent systems using pre-collected datasets. Existing offline MARL methods primarily constrain training within the dataset distribution, resulting in overly conservative policies that struggle to generalize beyond the support of the data. While model-based approaches offer a promising solution by expanding the original dataset with synthetic data generated from a learned world model, the high dimensionality, non-stationarity, and complexity of multi-agent systems make it challenging to accurately estimate the transitions and reward functions in offline MARL. Given the difficulty of directly modeling joint dynamics, we propose a local-to-global (LOGO) world model, a novel framework that leverages local predictions—which are easier to estimate—to infer global state dynamics, thus improving prediction accuracy while implicitly capturing agent-wise dependencies. Using the trained world model, we generate synthetic data to augment the original dataset, expanding the effective state-action space. To ensure reliable policy learning, we further introduce an uncertainty-aware sampling mechanism that adaptively weights synthetic data by prediction uncertainty, reducing approximation error propagation to policies. In contrast to conventional ensemble-based methods, our approach requires only an additional encoder for uncertainty estimation, significantly reducing computational overhead while maintaining accuracy. Extensive experiments across 8 scenarios against 8 baselines demonstrate that our method surpasses state-of-the-art baselines on standard offline MARL benchmarks, establishing a new model-based baseline for generalizable offline multi-agent learning.

KEYWORDS

Offline multi-agent reinforcement learning; Multi-agent model-based reinforcement learning

*Corresponding authors.



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). <https://doi.org/10.65109/IVNK5836>

ACM Reference Format:

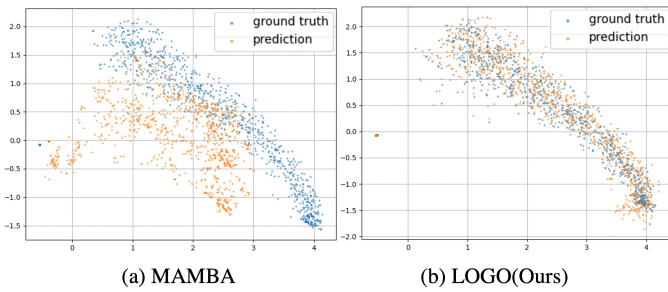
Sijia Li, Xinran Li, Shibo Chen, and Jun Zhang. 2026. Puzzle it Out: Local-to-Global World Model for Offline Multi-Agent Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/IVNK5836>

1 INTRODUCTION

Multi-Agent Reinforcement Learning (MARL) tackles scenarios where multiple agents learn and interact concurrently to solve cooperative tasks with a shared goal. Its ability to model complex dynamics and enable joint optimization has led to widespread applications, ranging from robotics to games and distributed systems [10, 23, 24, 29, 38, 66]. However, online MARL training faces several critical challenges, including sample inefficiency and non-stationarity, arising from its continuous interaction with the environment. This results in high computational costs, poor sample utilization, and safety concerns in online exploration, thereby hindering its deployment in critical real-world applications such as autonomous driving and clinical decision support [18, 27, 68].

In contrast, offline MARL [8] trains agents using pre-collected datasets, eliminating the need for active environment interaction. It improves sample efficiency and stability by learning from fixed dataset, avoiding the costs and risks associated with online exploration. However, offline MARL methods face out-of-distribution (OOD) issues, as the fixed dataset cannot account for unseen states or actions. To address this difficulty, current offline MARL approaches employ conservative value estimation [46, 55] and pessimistic policy training [33, 40] by restricting policies to the dataset support. While effective in mitigating OOD errors, these model-free methods often result in overly conservative behaviors, limiting the optimality of learned policies.

Model-based methods provide a potentially promising solution to alleviate this over-conservatism by expanding the offline dataset with policy rollouts in world models. However, constructing accurate world models [1] for offline MARL introduces significant challenges. First, the inherent interaction complexity and high-dimensionality of multi-agent systems brings excessive computational demands and difficulties in learning accurate dynamics with limited data [17]. Second, these difficulties inevitably induce approximation errors in the learned dynamic model, which can propagate during policy rollouts, degrading policy performance.



MACQL	MACQL+MAMBA	MACQL+LOGO
16.81±1.12	15.63±1.05	19.13±0.98

(c) Performance with different world model

Figure 1: The motivation of LOGO. The results in (a) and (b) demonstrate the next state prediction accuracy of baseline MAMBA and our proposed LOGO method on the SMAC 5m_vs_6m scenario, while (c) presents the performance of MACQL (Multi-Agent Conservative Q-Learning) when integrated with different world models. These results indicate that employing an inaccurate world model can adversely affect overall performance.

As illustrated in Fig. 1, the baseline world model produces inaccurate transition estimations which negatively impact downstream policy learning. These challenges necessitate both more precise multi-agent dynamic modeling and effective management of model-generated data to enhance generalization while ensuring training stability.

To address these limitations, this paper develops a Local-to-Global (LOGO) world model that improves multi-agent dynamic modeling by leveraging local transitions to assist the global transition prediction. Specifically, LOGO first learns local dynamics models from each agent’s observational space, enabling more accurate and efficient learning due to the reduced dimensionality and focusing on individual dynamics. The local predictions are then integrated to infer global state dynamics, piecing them together like a puzzle and capturing inter-agent dependencies and bridging the gap between local and global system behaviors. Furthermore, to account for the inevitable approximation errors in our model, we add an auxiliary state encoder that efficiently estimates prediction uncertainties using the discrepancy between different prediction paths with lower computational cost than conventional ensemble-based methods, enhancing the world model’s robustness.

We summarize our contributions as follows:

- We propose a novel Local-to-Global world model (LOGO) that first captures local observation dynamics then integrating them to deduce the global state, which enables more efficient and accurate modeling of multi-agent transitions and interactions while preserving stability.
- With LOGO, we augment the original offline dataset with synthetic data with uncertainty-aware weighted sampling to avoid model errors propagating to policy learning. Specifically, we propose a lightweight uncertainty estimation mechanism based on prediction discrepancy across different paths.

This approach achieves robust error estimation at significantly lower computational cost than traditional ensemble methods.

- Empirical results on established offline MARL benchmarks demonstrate that our method achieves superior performance compared to existing baselines, highlighting its effectiveness in improving generalization and policy robustness.

2 RELATED WORK

Offline MARL. Offline MARL optimizes policies using pre-collected datasets, removing the need for environment interactions during training and relaxing the requirement for expert demonstrations [60]. Recent advances in offline MARL have leveraged diverse methodological approaches to address the challenges of offline learning. From a multi-agent modeling perspective, value-based methods have been enhanced through counterfactual credit assignment [46], local-to-global value decomposition [55], causal inference frameworks [56], and improved temporal credit assignment [7, 50, 63]. Concurrently, data-centric innovations employ advanced architectures like transformers for local-global policy guidance [52] and diffusion models for trajectory generation [14, 25, 26, 71], while novel data augmentation techniques address dataset limitations [37]. On the policy optimization front, methods range from strict in-distribution behavior constrain [33, 36] to approaches encouraging controlled diversity within dataset constraints [16, 40, 69], with regularization techniques enabling multi-objective optimization [2, 64]. These advances demonstrate progress in tackling offline MARL challenges through complementary value estimation, data utilization, and policy constraint strategies. However, model-free methods often produce overly conservative policies limited to dataset support, motivating model-based approaches that can generate synthetic state-action pairs via pessimistic rollouts while maintaining stability and encouraging generalization.

Model-based MARL. Model-based MARL employs learned dynamics models to simulate agent interactions for sample-efficient policy optimization. Current approaches primarily address key multi-agent challenges, including scalability [6], decentralized local observations [58], and latent world models for state-space compression. Specifically, architectures like Dreamer [11, 34, 51] and Transformer-based models [65] employ latent state encoding to autoregressively predict future states, enabling high-fidelity trajectory simulation and efficient policy training. Online MARL methods [15, 41, 45, 59, 67] focus on exploration and sample efficiency, while offline model-based approaches address OOD generalization, requiring careful balance between leveraging learned dynamics for synthetic data and avoiding over-reliance on uncertain extrapolations beyond the dataset support. Several online MARL studies have also adopted a local-global perspective to achieve more accurate prediction, such as MACD [3] and MABL [54]. However, these approaches rely on large amounts of data and continuous environment feedback to refine their world models, which makes them unsuitable for straightforward extension to offline MARL. Our work advances this in offline MARL by developing local-to-global world models with enhanced accuracy and incorporating them into policy optimization through an uncertainty-weighting mechanism.

Offline model-based RL. Model-based offline RL employs a learned dynamics model to augment the dataset and improve generalization, which have demonstrated superior performance compared to model-free approaches in offline settings [20, 49, 62]. The standard pipeline involves jointly training a probabilistic dynamics model and reward function, enabling synthetic trajectory generation through model rollouts. However, due to inevitable approximation errors, careful uncertainty quantification becomes essential - most commonly implemented through ensemble variance [19, 20, 28, 44, 57, 62], while other methods such as count-based methods [21] and Bellman estimations [48] have also demonstrated effectiveness. Alternatively, distributional [47] or adversarial training paradigms [30] offer viable approaches to learn the world model. To ensure policy robustness against model errors, existing methods typically incorporate conservatism at either the policy optimization stage [4, 13, 53, 70] or value estimation level [44, 61]. In this vein, Barde et al. [1] introduces a model-based offline MARL algorithm featuring the inter-agent coordination. While advances in single-agent model-based offline RL such as uncertainty-aware learning [5] and conservative regularization [32] show promise, their ability to extend to offline MARL remains limited by the difficulty of modeling complex interactions among multiple agents and the excessive computational costs of ensembling such models.

3 PRELIMINARIES

Offline MARL. In this work, we adopt the decentralized partially observable Markov decision process (Dec-POMDP) [39] framework for cooperative MARL, formally defined by the tuple $M = \langle S, A, P, R, \Omega, O, n, \gamma \rangle$, where n agents learn cooperative policies under partial observability. Each agent i receives only local observation $o^i \in O$ based on the observation function $Z(s, i) : S \times N \rightarrow O$, and selects actions $a^i \in A$ according to its policy π^i . The environment evolves through joint actions $\mathbf{a} = (a^1, \dots, a^n)$ according to transition dynamics $P(s' | s, \mathbf{a}) : S \times A^n \times S' \rightarrow [0, 1]$, while providing a shared global reward $r = R(s, \mathbf{a}) \in \mathbb{R}$ to guide policy training. The discount factor $\gamma \in [0, 1)$ determines the agent’s temporal preference by exponentially weighting future rewards in the return calculation $\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]$. Offline MARL algorithms aim to derive an optimal joint policy $\pi = (\pi^1, \pi^2, \dots, \pi^n)$ from a static dataset consisting of multi-agent transition tuples, where the data is generated by the unknown behavioral policies.

Offline model-based RL. In offline model-based RL, we learn a dynamics model $\hat{M} = (\hat{T}, \hat{r})$ from a fixed dataset $D = \{(s, a, s', r)\}_{i=1}^N$, where $\hat{T}(s' | s, a)$ is the estimated transition function and $\hat{r}(s, a)$ is the estimated reward function. The model generates synthetic transitions (s, a, s', r) to form a training dataset \mathcal{D}_m for policy optimization, where $s' \sim \hat{T}(s' | s, a)$ and $r \sim \hat{r}(r | s, a)$. To alleviate model errors, a common approach is to employ uncertainty-penalized rewards, which discourage the policy from exploiting regions with high model uncertainty and thereby mitigate the risk of performance degradation induced by model bias or approximation errors. The reward penalty is defined as:

$$\tilde{r}(s, a) = \hat{r}(s, a) - \lambda u(s, a) \quad (1)$$

where $u(s, a)$ estimates state-action uncertainty (typically via model ensembles) and λ controls conservatism.

4 METHOD

As illustrated in Fig. 1, excessive model errors in the learned world model can significantly degrade the performance of offline MARL. Motivated by this observation (Sec. 4.1), we introduce two key innovations in our proposed LOGO: 1) a more accurate multi-agent world model architecture that reconstructs global states through precise local dynamic predictions (Sec. 4.2), and 2) an uncertainty estimation mechanism that quantifies prediction uncertainty in the transition and reward functions. (Sec. 4.3). Leveraging this uncertainty-aware world model, we generate additional high-quality synthetic data to augment policy training through uncertainty weighted sampling, thereby improving the agents’ generalization capability and overall performance.

4.1 Motivation Example

Current model-based MARL methods predominantly rely on direct state prediction to learn world dynamics. However, as illustrated in Fig. 1, model-based MARL methods following this paradigm - including MAMBA - suffer from inherent limitations in state prediction accuracy due to the compounding challenges of multi-agent systems, including partial observability, non-stationary dynamics, and high-dimensional joint action spaces. Such inaccuracies in world modeling further propagate through policy training, ultimately degrading the quality of the learned policy and the performance, as demonstrated in Fig. 1(c).

To address this issue, we propose a novel world model framework that leverages local prediction as a more tractable intermediate representation. Our key insight is that local observations (e.g., agent-centric states or neighborhood information) can be predicted with significantly higher fidelity than global states, as they circumvent the curse of dimensionality inherent in full-state modeling. By deducing the global state with these more reliable local predictions, our approach achieves more accurate global state estimation while maintaining computational efficiency. This hierarchical local-to-global paradigm not only mitigates error accumulation but also enables adaptive focus on critical interaction patterns within the multi-agent system.

4.2 Local-to-global World Model Construction

The key idea of our proposed LOGO world model is to bypass the curse of dimensionality typical in model-based MARL and obtain a more accurate world model by leveraging local dynamic prediction as an intermediate representations to reconstruct the global state, as illustrated in Fig. 2. We model each agent independently, where their observations naturally contain information about nearby agents and the environment. This allows us to train a local predictive model that focuses only on each agent’s immediate dynamics, which are simpler and more stable compared to modeling the full system. Instead of directly predicting complex global state transitions, we first predict each agent’s next local observation and then combine these predictions to reconstruct the complete global state - similar to assembling pieces of a puzzle. This approach improves estimation accuracy through independently modeling agent-level dynamics while leveraging environmental context to implicitly capture inter-agent dependencies. By decoupling local prediction from

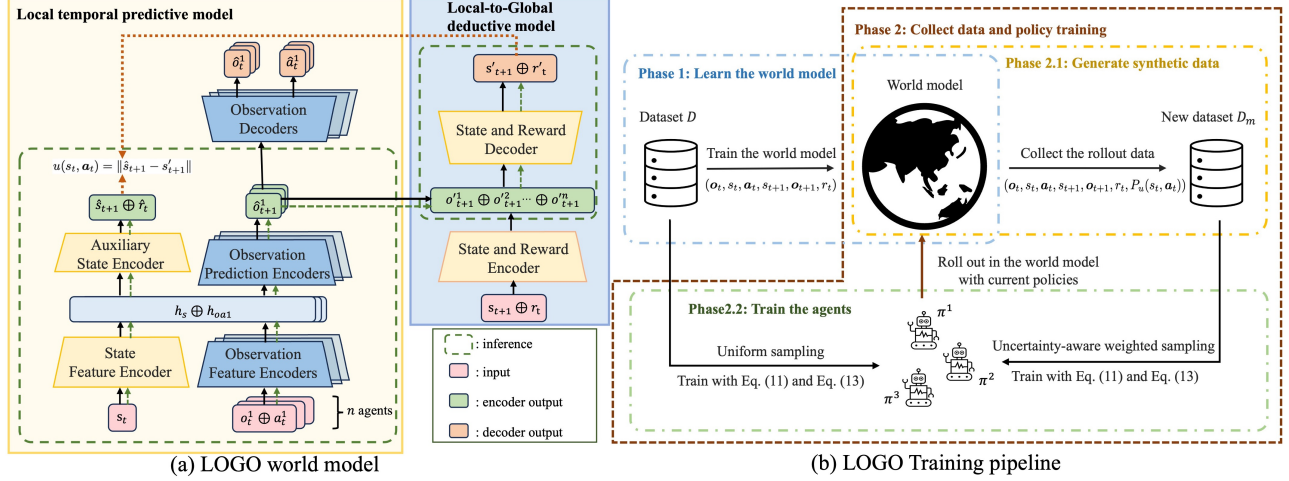


Figure 2: Overall framework of LOGO. (a) We first train the LOGO world model with both prediction loss and reconstruction loss. Subsequently, during the dynamic data generation phase, the predictive model (yellow part) predicts \hat{o}_{t+1} with (o_t, a_t, s_t) . These predictions are then propagated through the deductive model’s (blue part) decoder (indicated by the green dashed pathway) to compute the subsequent state s'_{t+1} and reward r'_t . Simultaneously, the state outputs from the uncertainty layer \hat{s}_{t+1} along with the deduced state s'_{t+1} are utilized to estimate the prediction uncertainty (indicated by the orange dashed pathway). (b) For policy training, we sample from the generated data using uncertainty-aware weighting to prioritize more reliable transitions.

the unstable and high-dimensional global prediction, the method reduces error propagation and enhance computational efficiency. The local-to-global abstraction aligns with modular MARL architectures, empirically improving training stability and generalization.

As shown in Fig. 2 (a), our proposed LOGO world model consists of a local world model (predictive model for each agent) illustrated in yellow background, and a deductive model illustrated in blue background. The training path is illustrated by solid lines, whereas the inference path is represented by dashed lines.

During the training phase, the predictive model first takes (o_t^i, a_t^i) as input for each agent i and use state information s_t as auxiliary information to predict the \hat{o}_{t+1}^i . The objective is to enable the encoder part of the local predictive model to accurately predict the next local observation. The deductive model will take the next state s_{t+1} and the reward r_t as input to construct the next joint observations $(\hat{o}_{t+1}^1, \hat{o}_{t+1}^2, \dots, \hat{o}_{t+1}^n)$. It then uses these observations to deduce a revised next state s'_{t+1} and the reward r'_t , thereby enabling the decoder to reconstruct the global state from local observations. For implementation, we initiate both local predictive models and deductive model as auto-encoder structures as they could capture critical dynamic features in a self-supervised manner for more robust representation. This local-to-global approach ensures that the model leverages both local agent interactions and global state information for more accurate and robust prediction. The predictive model’s output is denoted by the hat symbol $\hat{\cdot}$, while the output of the deductive model is indicated by the prime notation $'$.

During inference, the local predictive model first encodes the current state s_t via a state feature encoder to get state representation h_s , while agent-specific observations o_t^i and action a_t^i are processed through local observation feature encoders to obtain representation h_{oai} . The observation predictive encoders then generate per-agent next-observation estimates \hat{o}_{t+1} . These predictions are

subsequently concatenated as $(\hat{o}_{t+1}^1, \hat{o}_{t+1}^2, \dots, \hat{o}_{t+1}^n)$ and are processed within the decoder module of the deductive model to infer the next global state s'_{t+1} and estimated reward r' .

Local predictive model. During the world model training stage, we minimize the prediction error and the reconstruction error for the predictive model with the loss given as:

$$L_p = \underbrace{\mathbb{E}_{(o^i, s_t, a_t, o_{t+1}^i) \sim \mathcal{D}} \left[-\log \widehat{T}(o_{t+1}^i | s_t, o_t^i, a_t^i) \right]}_{\text{encoder loss}} + \underbrace{\mathbb{E}_{\hat{o}_{t+1}^i \sim \mathcal{E}_p(o_t^i, a_t^i, s_t)} \left[\left\| [o_t^i \oplus a_t^i] - \mathcal{R}_p(\hat{o}_{t+1}^i) \right\|^2 \right]}_{\text{decoder loss}}, \quad (2)$$

where \mathcal{E}_p and \mathcal{R}_p are the encoder and decoder of the predictive model and \widehat{T} is dynamics model trained with maximum likelihood estimation. The encoder predict \hat{o}_{t+1} and the decoder output \hat{o}_t and \hat{a}_t . Here \oplus denotes the vector concatenation operation. The decoder incorporates an auxiliary reconstruction loss, which is the second term in Eq. (2), to stabilize world model training and boost representational capacity.

Local-to-Global deductive model. After obtaining accurate observation predictions \hat{o}_{t+1}^i from the encoder of each local predictive model, we construct the global state by aggregating these predictions into a joint observation $(\hat{o}_{t+1}^1, \hat{o}_{t+1}^2, \dots, \hat{o}_{t+1}^n)$. This phase, referred to as the deduction stage, ensures a coherent integration of local predictions into a global context.

To further enhance the robustness modeling state-observation relationships and capture the underlying information, we also adopt an auto-encoder-based framework in which the encoder takes the next state and reward (s_{t+1}, r_t) as input and learns to reconstruct

them as training objectives. Meanwhile, the aggregated observations derived from the predictive model serve as the decoder’s input during the training phase. The loss functions of the deductive model are given as:

$$L_d = \underbrace{\mathbb{E}_{(s_{t+1}, r_t, o_{t+1}^i) \sim \mathcal{D}} \left[-\log \widehat{P}(o_{t+1}^i | s_{t+1}, r_t) \right]}_{\text{encoder loss}} + \underbrace{\mathbb{E}_{o_{t+1}^i \sim \mathcal{E}_d(s_{t+1}, r_t)} \left[\left\| [s_{t+1} \oplus r_t] - \mathcal{R}_d(o_{t+1}^i) \right\|^2 \right]}_{\text{decoder loss}}, \quad (3)$$

where \mathcal{E}_d and \mathcal{R}_d are the encoder and decoder of the deductive model and \widehat{P} is a probability function with maximum likelihood estimation.

To further enhance the model’s deductive capability, we introduce an additional loss term that regularizes the deductive model to improve its inference accuracy on predicted output:

$$L_{\mathcal{R}_d} = \mathbb{E}_{\delta_{t+1}^i \sim \mathcal{E}_p(o_{t+1}^i, s_t)} \left[\left\| [s_{t+1} \oplus r_t] - \mathcal{R}_d(\delta_{t+1}^i) \right\|^2 \right], \quad (4)$$

where \mathcal{E}_p represents the predictive model’s encoder. The decoder of the deductive model is therefore trained with both Eq. (4) and Eq. (3).

The overall loss function of the world model is defined as the sum of the losses of each component:

$$L_{world} = L_p + L_d + L_{\mathcal{R}_d}.$$

4.3 Weighted Sampling with Uncertainty Estimation

As demonstrated in Fig. 1, inaccurate world model predictions adversely affect policy learning performance through synthetic data generated by the world model that contains significant approximation errors. A common method to address this problem is reward penalty, as mentioned in Sec. 3. However, if the reward function itself is underestimated, applying a penalty could further degrade performance. To mitigate this issue, we propose an uncertainty-based reweighted sampling strategy. This approach prioritizes synthetic trajectories generated by the world model that exhibit high prediction confidence, while down-weighting those associated with high uncertainty.

Uncertainty estimation. To achieve this uncertainty-aware reweighting, we implement an uncertainty estimation module that quantifies prediction confidence by measuring the discrepancy between state-reward s_{t+1}, r_t predictions from different computational paths (the predictive model and deductive model as shown in Fig. 2). This structural divergence between models naturally captures prediction variability, yielding a principled measure of epistemic uncertainty. The uncertainty is defined as:

$$u(s_t, \mathbf{a}_t) = \|\hat{s}_{t+1} - s'_{t+1}\|, \quad (5)$$

where \hat{s}_{t+1} is the state output of the auxiliary state encoder in the predictive model and s'_{t+1} is the state output in the deductive model. The auxiliary state encoder that assist the uncertainty estimation above (detailed in Fig. 2) is trained with the auxiliary loss given as:

$$L_{\mathcal{E}_{ps}} = \mathbb{E}_{(o^i, s_t, \mathbf{a}_t, s_{t+1}, r_t) \sim \mathcal{D}} \left[-\log \widehat{T}(s_{t+1}, r_t | s_t, o_t^i, \mathbf{a}_t^i) \right], \quad (6)$$

where \mathcal{E}_{ps} is the uncertainty layer in the predictive model and \widehat{T} is the dynamics model trained with maximum likelihood estimation.

Weighted sampling. This constrained uncertainty metric is then utilized as a sampling weight during data selection from the world model’s generated dataset D_m . Specifically, when a predicted experience (s_{t+1}, r_t) exhibits higher uncertainty, we assign it a correspondingly lower sampling probability in D_m . This weighting scheme prioritizes high-confidence predictions during policy training, mitigating the impact of uncertain or inaccurate model-generated transitions to prevent training degradation.

In our framework, each synthetic experience is stored as

$$(\mathbf{o}_t, s_t, \mathbf{a}_t, s_{t+1}, \mathbf{o}_{t+1}, r_t, P_u(s_t, \mathbf{a}_t)),$$

where $P_u(s_t, \mathbf{a}_t) = \text{clip}(C - u(s_t, \mathbf{a}_t), [0, C])$ is defined as the clipped version with constant C , ensuring valid bounds while prioritizing high-confidence (low-uncertainty) samples during training. While we uniformly sample from the original dataset D to preserve its distribution, we apply uncertainty-weighted sampling exclusively to world model-generated dataset D_m . The sampling weight for D_m is inversely proportional to its associated uncertainty is defined as:

$$w(s_t, \mathbf{a}_t) = \frac{\exp(P_u(s_t, \mathbf{a}_t))}{\sum_{s_i, \mathbf{a}_i \in D_m} \exp(P_u(s_i, \mathbf{a}_i))}. \quad (7)$$

This weighting scheme reduces the influence of high-uncertainty generated experience during policy optimization, improving training robustness and mitigating policy degradation due to prediction error.

Under our proposed framework, the overall probability of sampling data from the generated dataset D_m is maintained equal to the overall probability of sampling from the true dataset D . Formally, if we denote the optimal Q-value function estimated under uniform sampling as

$$Q^*(s_t, \mathbf{a}_t) = \mathbb{E} \left[r(s_t, \mathbf{a}_t) + \gamma \max_{\mathbf{a}_{t+1}} Q^*(s_{t+1}, \mathbf{a}_{t+1}) \right], \quad (8)$$

the optimal Q-value under uncertainty-aware weighted sampling estimation can be expressed as:

$$Q^*(s_0, \mathbf{a}_0) = \mathbb{E} \left[w(s_t, \mathbf{a}_t) [r(s_t, \mathbf{a}_t) + \gamma \max_{\mathbf{a}_{t+1}} Q^*(s_{t+1}, \mathbf{a}_{t+1})] \right], \quad (9)$$

where $\mathbb{E}[w(s_t, \mathbf{a}_t)] = 1$. We establish the estimation error bound of the optimal Q-function under the generated dataset D_m in Theorem 1 (see Appendix ?? for proof).

Theorem 1. *With the assumption that Q-function and reward function are Lipschitz continuous functions with Lipschitz constants L_Q and L_r , and the assumption that the estimated errors of r, s and Q are bounded by $\epsilon_s, \epsilon_r, \epsilon_Q$ respectively, we have the estimation error between generated optimal Q-value $\widehat{Q}^*(s, a)$ and true optimal Q-value $Q^*(s, a)$ bounded as:*

$$\|\widehat{Q}^*(s_t, \mathbf{a}_t) - Q^*(s_t, \mathbf{a}_t)\| \leq (L_r + \gamma L_Q) \epsilon_s + \epsilon_r + \gamma \epsilon_Q.$$

This mechanism enhances sample efficiency and mitigates the potential negative impact of erroneous model predictions on policy optimization.

4.4 Overall Training Framework and Objectives

As illustrated in Fig 2.(b), we first train the LOGO world model with losses Eq. (2), Eq. (3), Eq. (4) and Eq. (6). After pre-training, we iteratively generate additional data with the trained world model and use these synthetic data D_m alongside the original dataset D to update the policy with the uncertainty-aware priority $P_u(s_t, \mathbf{a}_t)$.

For policy training, our method can be integrated with any existing offline MARL algorithm. In our implementation, we adopt MACQL [22] as the backbone. The Q-value loss function is:

$$\min_Q \underbrace{\alpha \cdot \mathbb{E}_{s \sim \mathcal{D} \cup \mathcal{D}_m} [\mathbb{E}_{s \sim \mathcal{D}, a \sim \pi(a)} [Q_\theta(s, a)] - \mathbb{E}_{(s,a) \sim \mathcal{D}} [Q_\theta(s, a)]]}_{\text{Regularization}} + \underbrace{\frac{1}{2} \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} [(Q_\theta(s, a) - \mathcal{T}^\pi Q_\theta(s, a))^2]}_{\text{Bellman error}}, \quad (10)$$

where in the multi-agents setting the Bellman operator \mathcal{T} is defined as:

$$\mathcal{T}^\pi Q_\theta(s, \mathbf{a}) = r(s, \mathbf{a}) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, \mathbf{a})} \left[\max_{\mathbf{a}'} Q_\theta(s', \mathbf{a}') \right]. \quad (11)$$

For the policy network optimization, we employ a composite objective function that combines policy gradient maximization with behavior cloning (BC) regularization, formulated as:

$$\mathcal{L}_{\text{policy}} = \underbrace{\mathbb{E}_{(s,a) \sim \mathcal{D} \cup \mathcal{D}_m} [Q_\theta(s, \pi_\phi(s))]}_{\text{Policy Gradient}} - \lambda \cdot \underbrace{\mathbb{E}_{(o_i, a_i) \sim \mathcal{D} \cup \mathcal{D}_m} [\|\pi_{\phi_i}(o^i) - a^i\|_2^2]}_{\text{BC Regularization}}, \quad (12)$$

where λ is a hyperparameter that controls the trade-off between policy optimization and behavior cloning regularization. We include the psseudo code for the overall training pipeline in Appendix ??.

5 EXPERIMENTS

In this section, we conduct experiments across established offline MARL benchmarks, including the Multi-Agent MuJoCo Environment (MaMuJoCo) [42], and the StarCraft Multi-Agent Environment (SMAC) [35]. Our experimental dataset follows the same setup as prior works [46, 55], with different data qualities. The baselines are summarized in Table 1.

5.1 Main Results

In the experiments, we set the rollout horizon to 15 steps. As demonstrated in Table 7, LOGO achieves superior long-term prediction ability compared with existing model-based approaches in the offline setting. All results are presented as mean values with 95% confidence intervals (mean \pm margin of error). The results are shown in Table 2 and Table 3, which demonstrate that LOGO outperforms both model-based and model-free approaches, particularly with medium-quality datasets. This empirically demonstrates that LOGO achieves superior generalization with superior performance than prior model-free offline methods, while delivering more accurate predictions and more stable policy training compared to model-based offline approaches.

5.2 Generalization Capability with Model Predictive Control (MPC)

To validate the generalization capability of our world model, we further implement the LOGO-enabled Model Predictive Control (MPC) [12] and evaluate its effectiveness. We train the MPC-enhanced MARL algorithm using the loss function:

$$\mathcal{L}_{\text{MPC}} = \mathcal{L}_{\text{policy}} + \mathbb{E}_{(o_i, a_i) \sim \mathcal{D}} [\|\pi_{\phi_i}(o^i) - a_{max}^i\|_2^2]$$

, where a_{max}^i represents the optimal action selected through model rollout from three candidate actions proposed by the policy network. Results from the 6h_vs_8z map in SMAC are presented in Table 4, showing that LOGO-enabled MPC enjoys significant performance improvements and enhanced stability. This experiments show the adaptivity of the proposed the multi-agent world model LOGO across MBRL algorithms and MPC applications.

5.3 Efficiency of LOGO

To evaluate computational efficiency, we conducted extensive benchmarking against ensemble-based world model. Wall-clock time was measured over 500 independent trajectories, each consisting of 10 rollout steps under identical environmental conditions. The results show that our method achieves substantially higher efficiency, with up to a 3 \times reduction in inference runtime while maintaining comparable accuracy (Table 5). These runtime comparisons will be included in the revised manuscript to better highlight the practical advantages of our approach.

5.4 Ablation Study and Visualization Analysis

Local-to-global Prediction Visualization. In this experiment, we demonstrate that compared to directly predicting the global state, our method achieves higher accuracy, which models the local dynamics and infers the subsequent state from the predicted observations. We employ the same architecture as the local predictive model to perform direct state prediction. As demonstrated in Fig. 3 (d) and (e), the proposed approach, which predicts local observations and subsequently infers the system state from these local predictions, yields significantly improved accuracy compared to conventional methods.

Visualization of Prediction Accuracy Comparison. We visualize the discrepancy between predicted and ground-truth values with data out of the dataset. The results demonstrate that, compared to other model-based approaches, our framework yields more accurate predictions. The visualization results presented in Fig. 3 and the one-step reward prediction error reported in the last row of Table 7 demonstrate that our method achieves more accurate predictions compared to the baseline model-based approaches.

Uncertainty estimation. In offline model-based RL, conventional uncertainty estimation often penalizes rewards to conservatively estimate values and mitigate model errors. Through comparative experiments (Table 6) in the medium quality dataset, we demonstrate that our uncertainty-aware sampling outperforms this reward-penalty approach. We identify reward prediction inaccuracies in complex multi-agent settings as a key factor, where reward estimation errors propagate to value approximations - especially with reward shaping. Our uncertainty-weighted sampling addresses this

Table 1: Comparison with Baseline Methods. Off/On:Offline/Online. MB/MF:Model-Based/Model-Free. Gl/Lo:Global/Local. UE: Uncertainty estimation.

Method	Off/On	MB/MF	Gl/Lo	UE	Other Notes
OMAR [40]	Off	MF	Lo	×	A policy-based method under the decentralized framework
MACQL [22]	Off	MF	Gl	×	An adaption of CQL in the multi-agent domain
OMIGA [55]	Off	MF	Gl	×	A global-to-local value-based method
CFCQL [46]	Off	MF	Gl	×	A value-based method with counterfactual credit assignment mechanisms
Morel [20]	Off	MB	Lo	✓	An ensemble-based approach with pessimistic value regularization
SUMO [43]	Off	MB	Lo	✓	An ensemble-based approach that estimates uncertainty with search-based method
MAMBA [6]	On	MB	Gl	×	A DreamerV2-based method in MARL
MAZero [31]	On	MB	Gl	×	A online method with Monte Carlo Tree Search (MCTS) for policy search
LOGO(Ours)	Off	MB	Lo&Gl	✓	A local-to-global world model which explicitly estimates the uncertainty with different prediction paths

Table 2: Performance comparison of different algorithms in SMAC. The best result in each setting are highlighted in blue. MR: Medium-Replay.

Map	Data	MACQL	OMAR	CFCQL	Morel	SUMO	MAMBA	MAZero	LOGO(ours)
2s3z	MR	14.18±1.45	15.15±1.45	16.10±1.38	16.44±1.34	16.53±1.40	16.37±1.37	15.76±1.40	17.06±1.33
	Medium	13.32±1.37	12.68±1.41	13.29±1.52	13.08±1.37	13.17±1.38	13.27±1.29	13.13±1.34	13.82±1.40
	Expert	19.77±0.26	19.61±0.29	19.93±0.04	19.84±0.04	19.86±0.04	19.77±0.42	19.69±0.33	19.94±0.20
	Mixed	12.94±1.24	13.20±1.14	13.21±1.16	13.44±1.18	13.64±1.13	13.54±1.50	13.26±1.44	14.14±1.35
3s_vs_5z	MR	16.27±1.34	14.72±0.91	19.07±0.73	16.64±1.35	17.27±1.29	16.82±1.11	16.89±1.35	19.38±0.98
	Medium	18.86±1.20	20.09±1.03	20.88±1.26	18.91±1.39	19.51±1.34	17.54±1.01	2.01±0.66	21.48±1.45
	Expert	21.03±0.18	21.40±0.29	21.49±0.15	21.12±0.27	21.32±0.31	21.54±0.52	20.65±1.00	21.63±0.30
	Mixed	17.56±1.22	20.03±1.16	21.13±1.23	19.24±0.96	19.40±1.06	20.67±1.47	19.88±1.12	21.43±1.20
5m_vs_6m	MR	10.39±1.27	10.43±1.37	11.35±1.52	11.17±1.59	11.01±1.61	11.82±1.58	11.12±1.55	11.67±1.49
	Medium	9.44±0.96	11.84±1.60	12.29±1.85	11.63±1.32	11.82±1.56	12.33±1.56	10.94±1.53	12.63±1.73
	Expert	16.35±1.61	17.56±1.28	17.56±1.51	8.77±0.96	9.82±0.72	16.35±1.61	17.01±1.15	18.51±1.33
	Mixed	7.10±0.58	17.56±1.45	17.52±1.57	13.68±1.88	14.23±1.71	15.46±1.54	15.66±1.51	18.21±1.43
6h_vs_8z	MR	15.36±1.08	12.57±0.74	16.13±1.17	15.04±1.23	15.54±1.16	15.65±0.95	16.35±0.91	17.14±1.51
	Medium	11.88±1.04	17.24±1.19	17.88±1.04	17.02±1.12	17.34±1.21	16.89±1.08	17.72±1.30	18.70±0.99
	Expert	16.81±1.12	17.51±0.72	18.76±0.85	16.78±1.22	17.71±1.25	15.63±1.05	15.90±1.66	19.13±0.98
	Mixed	12.15±0.84	17.36±0.83	17.72±0.97	17.58±1.13	17.81±1.23	17.13±1.18	17.61±1.05	18.68±1.00
3s5z_vs_3s6z	MR	8.37±0.54	9.64±0.52	10.08±0.61	10.01±0.60	10.03±0.56	10.13±0.42	9.83±0.67	10.37±0.70
	Medium	10.23±0.70	11.54±0.67	11.80±0.61	10.49±0.53	10.97±0.65	11.41±0.73	11.35±0.74	12.17±0.66
	Expert	10.71±0.76	11.81±0.74	12.05±0.74	11.92±0.62	11.91±0.74	10.81±0.40	11.48±0.77	12.31±0.76
	Mixed	9.87±0.67	11.22±0.57	11.53±0.56	10.49±0.63	10.65±0.61	9.95±0.62	10.18±0.51	11.12±0.72

by adaptively prioritizing high-confident samples to reduce estimation bias.

Rollout horizon selection. To assess the effectiveness of different world models using varying trajectory horizons, we conduct comparative experiments using rollout horizons of 5, 10, 15, and 20 steps. As demonstrated in Table 7, our proposed LOGO consistently outperforms baseline methods across multiple rollout horizons. Specifically, some baseline methods suffer from performance degradation when using a long rollout horizon (e.g., 20 steps), while ours achieves optimal performance even at 15 steps (compared to their

10-step limit), due to its uncertainty reweighted sampling against error accumulation over extended sequences.

6 CONCLUSION

In this paper, we integrate model-based approaches into offline MARL by introducing a novel Local-to-Global World Model (LOGO) to effectively capture the underlying multi-agent transition dynamics and reward functions. We propose a local-to-global framework capable of capturing precise dynamics from observational

Table 3: Performance comparison of different algorithms in MaMuJoCo. The best result in each row are highlighted in blue. Ho: Hopper-v2. An: Ant-v2. Ha: HalfCheetah-v2. MR: Medium-Reply. ME: Medium. EX: Expert. MI: Mixed.

En	Data	MACQL	OMIGA	CFCQL	Morel	SUMO	MAMBA	MAZero	LOGO(ours)
Ho	MR	330.23±6.23	774.18±197.38	854.23±135.43	1412.36±182.96	1218.48±124.63	350.57±9.33	906.63±102.94	1659.10±130.30
	ME	489.23±102.34	1148.62±169.02	1231.34±176.34	976.89±92.26	932.46±106.25	1087.28±140.63	1172.12±146.9	1302.7±189.24
	EX	403.26±126.34	872.39±204.27	899.34±186.24	683.77±211.72	823.53±367.35	903.12±375.8	684.72±307.22	1086.44±221.65
	MI	376.23±63.23	718.01±285.79	845.23±253.23	1038.11±182.96	1089.32±162.76	303.33±104.67	876.69±198.86	1245.65±244.26
An	MR	804.43±83.22	1029.13±21.27	1038.23±43.23	1052.67±22.34	1077.67±25.82	1019.55±26.25	1012.11±16.03	1253.01±18.69
	ME	1337.74±56.66	1417.37±4.11	1457.23±7.23	1398.05±5.12	1408.34±5.02	1380.37±4.75	1382.64±4.65	1441.08±6.18
	EX	1634.76±125.38	2053.23±3.34	2042.75±4.26	2047.27±3.39	2044.23±2.43	2042.84±2.38	2035.96±4.48	2063.73±2.89
	MI	1474.34±86.43	1717.23±37.23	1687.34±58.94	1701.48±29.11	1745.23±32.16	1522.10±48.91	1711.13±41.30	1853.01±32.63
Ha	MR	2042.69±196.34	2406.70±91.13	2334.87±98.71	2393.57±69.75	2437.57±49.85	2434.29±67.19	2593.38±95.02	2566.74±67.63
	ME	1594.34±263.84	2646.26±39.51	2780.06±43.22	2669.34±35.24	2713.48±46.38	2436.24±50.57	2652.67±49.2	2805.25±56.36
	EX	2347.13±298.33	3138.54±227.20	3345.80±164.23	3259.01±204.09	3305.89±183.26	2883.30±201.34	3079.40±233.58	3534.30±191.75
	MI	2174.25±216.52	2825.58±312.33	2645.23±232.53	2853.10±241.48	2902.35±233.15	2801.10±224.26	2947.46±208.75	3060.46±299.17




Figure 3: The prediction comparison of model-based methods. We employ Principal Component Analysis (PCA) [9] to visualize and compare the ground truth next state with the predicted next state generated by world models across various model-based approaches in SMAC 6h_vs_8z map.

Table 4: The performance of MPC.

Method	Med-Rep	Medium	Expert	Mixed
MACQL	15.36±1.08	11.18±1.04	12.81±1.12	12.15±0.84
MACQL+MPC	16.36±0.74	17.36±0.63	17.98±0.55	17.27±0.84

Table 5: The Comparison of Efficiency.

Map	Ensemble	LOGO (ours)
2s3z	0.227	0.057
5m_vs_6m	0.225	0.042
6h_vs_8z	0.247	0.063

space—which is more tractable to learn—and subsequently inferring the global state in a puzzle-like manner to reconstruct the global dynamics with greater accuracy while implicitly modeling inter-agent relationships. LOGO expands the dataset by generating synthetic data based on the learned world model, effectively extending data coverage and improving policy generalization. To further enhance policy reliability, we propose an uncertainty-based weighted sampling mechanism, which adaptively weights synthetic

Table 6: Ablation study of uncertainty. WS: Weight sampling. RP: Reward penalty.

Method	Hopper-v2	Ant-v2	HalfCheetah-v2
WS	1302.34±189.24	1436.08±6.18	2805.25±56.36
RP	1133.46±165.64	1245.23±7.45	2539.43±74.23

Table 7: Performance of different rollout horizon and one-step reward prediction error. The best result in each setting are highlighted in blue.

	Morel	SUMO	MAMBA	MAZero	LOGO(ours)
horizon = 5	17.56±0.47	17.31±0.42	17.46±0.54	17.14±0.72	18.14±0.46
horizon = 10	18.02±0.57	18.05±0.40	18.06±0.57	17.14±0.72	18.46±0.56
horizon = 15	16.78±0.61	17.34±0.58	16.89±0.54	17.72±0.65	18.70±0.49
horizon = 20	17.23±0.72	17.58±0.59	16.03±0.48	16.03±0.80	18.57±0.58
prediction error	0.1316	0.0928	0.0673	0.0257	0.0051

data by prediction uncertainty. This mechanism mitigates the impact of approximation errors on policy learning while maintaining computational efficiency by avoiding the overhead of traditional ensemble-based uncertainty estimation methods.

REFERENCES

- [1] Paul Barde, Jakob Foerster, Derek Nowrouzezahrai, and Amy Zhang. 2023. A model-based solution to the offline multi-agent reinforcement learning coordination problem. *arXiv preprint arXiv:2305.17198* (2023).
- [2] The Viet Bui, Thanh Hong Nguyen, and Tien Mai. 2024. ComaDICE: Offline Cooperative Multi-Agent Reinforcement Learning with Stationary Distribution Shift Regularization. *arXiv preprint arXiv:2410.01954* (2024).
- [3] Jiajun Chai, Yuqian Fu, Dongbin Zhao, and Yuanheng Zhu. 2024. Aligning credit for multi-agent cooperation via model-based counterfactual imagination. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 281–289.
- [4] Yassine Chemingui, Aryan Deshwal, Trong Nghia Hoang, and Janardhan Rao Doppa. 2024. Offline model-based optimization via policy-guided gradient search. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 10, 11230–11239.
- [5] Christopher Diehl, Timo Sievernich, Martin Krüger, Frank Hoffmann, and Torsten Bertram. 2021. Umbrella: Uncertainty-aware model-based offline reinforcement learning leveraging planning. *arXiv preprint arXiv:2111.11097* (2021).
- [6] Vladimir Egorov and Aleksei Shpilman. 2022. Scalable multi-agent model-based reinforcement learning. *arXiv preprint arXiv:2205.15023* (2022).
- [7] Eslam Eldeeb, Houssein Sifaou, Osvaldo Simeone, Mohammad Shehab, and Hirley Alves. 2024. Conservative and Risk-Aware Offline Multi-Agent Reinforcement Learning. *IEEE Transactions on Cognitive Communications and Networking* (2024).
- [8] Juan Formanek, Callum R Tilbury, Louise Beyers, Jonathan Shock, and Arnun Pretorius. 2024. Dispelling the mirage of progress in offline marl through standardised baselines and evaluation. *Advances in Neural Information Processing Systems* 37 (2024), 139650–139672.
- [9] Detlef Groth, Stefanie Hartmann, Sebastian Klie, and Joachim Selbig. 2013. Principal components analysis. *Computational Toxicology: Volume II* (2013), 527–547.
- [10] Shangding Gu, Jakub Grudzien Kuba, Yuanpei Chen, Yali Du, Long Yang, Alois Knoll, and Yaodong Yang. 2023. Safe multi-agent reinforcement learning for multi-robot control. *Artificial Intelligence* 319 (2023), 103905.
- [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, Xiaolong Wang, and Hao Su. 2022. Temporal difference learning for model predictive control. *arXiv preprint arXiv:2203.04955* (2022).
- [13] Toru Hishinuma and Kei Senda. 2021. Weighted model estimation for offline model-based reinforcement learning. *Advances in neural information processing systems* 34 (2021), 17789–17800.
- [14] Jianuo Huang. 2024. Diffusion Models for Offline Multi-agent Reinforcement Learning with Safety Constraints. *arXiv preprint arXiv:2407.00741* (2024).
- [15] Jiawei Huang, Niao He, and Andreas Krause. 2024. Model-based rl for mean-field games is not statistically harder than single-agent rl. *arXiv preprint arXiv:2402.05724* (2024).
- [16] Songjun Huang, Chuanneng Sun, Ruo-Qian Wang, and Dario Pompili. 2024. Multi-behavior multi-agent reinforcement learning for informed search via offline training. *2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*, 19–26.
- [17] Zhao Huang. 2024. *Modeling-Based Optimization for Robotic Manipulation*. University of California, San Diego.
- [18] Tao Jiang, Lei Yuan, Lihe Li, Cong Guan, Zongzhang Zhang, and Yang Yu. 2024. Multi-Agent Domain Calibration with a Handful of Offline Data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- [19] Rahul Kidambi, Jonathan Chang, and Wen Sun. 2021. Mobile: Model-based imitation learning from observation alone. *Advances in Neural Information Processing Systems* 34 (2021), 28598–28611.
- [20] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. 2020. Morel: Model-based offline reinforcement learning. *Advances in neural information processing systems* 33 (2020), 21810–21823.
- [21] Byeongchan Kim and Min-hwan Oh. 2023. Model-based offline reinforcement learning with count-based conservatism. *International Conference on Machine Learning*, 16728–16746.
- [22] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems* 33 (2020), 1179–1191.
- [23] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems* 30 (2017).
- [24] Hyeoksoo Lee, Jiwoo Hong, and Jongpil Jeong. 2022. MARL-based dual reward model on segmented actions for multiple mobile robots in automated warehouse environment. *Applied Sciences* 12, 9 (2022), 4703.
- [25] Chao Li, Ziwei Deng, Chenxing Lin, Wenqi Chen, Yongquan Fu, Weiquan Liu, Chenglu Wen, Cheng Wang, and Siqi Shen. [n.d.]. DoF: A Diffusion Factorization Framework for Offline Multi-Agent Reinforcement Learning. In *The Thirteenth International Conference on Learning Representations*.
- [26] Zhuoran Li, Ling Pan, and Longbo Huang. 2023. Beyond conservatism: Diffusion policies in offline multi-agent reinforcement learning. *arXiv preprint arXiv:2307.01472* (2023).
- [27] Sidian Lin, Soroush Saghafian, Jessica Lipschitz, and Katherine Burdick. 2024. A Multi-Agent Reinforcement Learning Algorithm for Personalized Recommendations in Bipolar Disorder. Available at SSRN 5073570 (2024).
- [28] Sen Lin, Jialin Wan, Tengyu Xu, Yingbin Liang, and Junshan Zhang. 2022. Model-based offline meta-reinforcement learning with regularization. *arXiv preprint arXiv:2202.02929* (2022).
- [29] Yiheng Lin, Guannan Qu, Longbo Huang, and Adam Wierman. 2020. Distributed reinforcement learning in multi-agent networked systems. *arXiv preprint arXiv:2006.06555* (2020).
- [30] Zichuan Lin, Garrett Thomas, Guangwen Yang, and Tengyu Ma. 2020. Model-based adversarial meta-reinforcement learning. *Advances in Neural Information Processing Systems* 33 (2020), 10161–10173.
- [31] Qihan Liu, Jianing Ye, Xiaoteng Ma, Jun Yang, Bin Liang, and Chongjie Zhang. 2024. Efficient multi-agent reinforcement learning by planning. *arXiv preprint arXiv:2405.11778* (2024).
- [32] Xiao-Yin Liu, Xiao-Hu Zhou, Mei-Jiang Gui, Xiao-Liang Xie, Shi-Qi Liu, Shuang-Yi Wang, Hao Li, Tian-Yu Xiang, De-Xing Huang, and Zeng-Guang Hou. 2023. Domain: Mildly conservative model-based offline reinforcement learning. *arXiv preprint arXiv:2309.08925* (2023).
- [33] Zongkai Liu, Qian Lin, Chao Yu, Xiawei Wu, Yile Liang, Donghui Li, and Xuetao Ding. 2025. Offline Multi-Agent Reinforcement Learning via In-Sample Sequential Policy Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39, 19068–19076.
- [34] Kenzo Lobos-Tsunekawa, Akshay Srinivasan, and Michael Spranger. 2022. MA-Dreamer: Coordination and communication through shared imagination. *arXiv preprint arXiv:2204.04687* (2022).
- [35] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [36] Daiki E Matsunaga, Jongmin Lee, Jaeseok Yoon, Stefanos Leonardos, Pieter Abbeel, and Kee-Eung Kim. 2023. Alberdice: addressing out-of-distribution joint actions in offline multi-agent rl via alternating stationary distribution correction estimation. *Advances in Neural Information Processing Systems* 36 (2023), 72648–72678.
- [37] Linghui Meng, Xi Zhang, Dengpeng Xing, and Bo Xu. 2024. A New Pre-Training Paradigm for Offline Multi-Agent Reinforcement Learning with Suboptimal Data. *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7520–7524.
- [38] Miquel Noguera i Alonso and Abdel Mfougouon Njupoun. 2024. Game Theory and Multi-Agent Reinforcement Learning: A Mathematical Overview. (2024).
- [39] Frans A Oliehoek, Christopher Amato, et al. 2016. *A concise introduction to decentralized POMDPs*, Vol. 1. Springer.
- [40] Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu. 2022. Plan better amid conservatism: Offline multi-agent reinforcement learning with actor rectification. In *International conference on machine learning*. PMLR, 17221–17237.
- [41] Barna Pasztor, Ilija Bogunovic, and Andreas Krause. 2021. Efficient model-based multi-agent mean-field reinforcement learning. *arXiv preprint arXiv:2107.04050* (2021).
- [42] Bei Peng, Tabish Rashid, Christian Schroeder de Witt, Pierre-Alexandre Kamieny, Philip Torr, Wendelin Böhm, and Shimon Whiteson. 2021. Facmac: Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems* 34 (2021), 12208–12221.
- [43] Zhongjian Qiao, Jiafei Lyu, Kechen Jiao, Qi Liu, and Xiu Li. 2025. Sumo: Search-based uncertainty estimation for model-based offline reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 19, 20033–20041.
- [44] Marc Rigter, Bruno Lacerda, and Nick Hawes. 2022. Rambo-rl: Robust adversarial model-based offline reinforcement learning. *Advances in neural information processing systems* 35 (2022), 16082–16097.
- [45] Pier Giuseppe Sessa, Maryam Kamgarpour, and Andreas Krause. 2022. Efficient model-based multi-agent reinforcement learning via optimistic equilibrium computation. *International Conference on Machine Learning*, 19580–19597.
- [46] Jianzhun Shao, Yun Qu, Chen Chen, Hongchang Zhang, and Xiangyang Ji. 2024. Counterfactual conservative Q learning for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 36 (2024).
- [47] Laixi Shi and Yuejie Chi. 2024. Distributionally robust model-based offline reinforcement learning with near-optimal sample complexity. *Journal of Machine Learning Research* 25, 200 (2024), 1–91.
- [48] Yihao Sun, Jiayi Zhang, Chengxing Jia, Haoxin Lin, Junyin Ye, and Yang Yu. 2023. Model-Bellman inconsistency for model-based offline reinforcement learning. In *International Conference on Machine Learning*. PMLR, 33177–33194.
- [49] Phillip Swazinna, Steffen Udluft, Daniel Hein, and Thomas Runkler. 2022. Comparing model-free and model-based algorithms for offline reinforcement learning. *IFAC-PapersOnLine* 55, 15 (2022), 19–26.

- [50] Qi Tian, Kun Kuang, Furui Liu, and Baoxiang Wang. 2023. Learning from good trajectories in offline multi-agent reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 10, 11672–11680.
- [51] Edan Toledo and Amanda Prorok. 2024. Codreamer: Communication-based decentralised world models. *arXiv preprint arXiv:2406.13600* (2024).
- [52] Wei-Cheng Tseng, Tsun-Hsuan Johnson Wang, Yen-Chen Lin, and Phillip Isola. 2022. Offline multi-agent reinforcement learning with knowledge distillation. *Advances in Neural Information Processing Systems* 35 (2022), 226–237.
- [53] Masatoshi Uehara and Wen Sun. 2021. Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226* (2021).
- [54] Aravind Venugopal, Stephanie Milani, Fei Fang, and Balaraman Ravindran. 2023. Mabl: Bi-level latent-variable world model for sample-efficient multi-agent reinforcement learning. *arXiv preprint arXiv:2304.06011* (2023).
- [55] Xiangsen Wang, Haoran Xu, Yinan Zheng, and Xianyuan Zhan. 2024. Offline multi-agent reinforcement learning with implicit global-to-local value regularization. *Advances in Neural Information Processing Systems* 36 (2024).
- [56] Ziyang Wang, Yali Du, Yudi Zhang, Meng Fang, and Biwei Huang. 2023. Macca: Offline multi-agent reinforcement learning with causal credit assignment. *arXiv preprint arXiv:2312.03644* (2023).
- [57] Fan Wu, Rui Zhang, Qi Yi, Yunkai Gao, Jiaming Guo, Shaohui Peng, Siming Lan, Husheng Han, Yansong Pan, Kaizhao Yuan, et al. 2024. OCEAN-MBRL: offline conservative exploration for model-based offline reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 14, 15897–15905.
- [58] Zifan Wu, Chao Yu, Chen Chen, Jianye Hao, and Hankz Hankui Zhuo. 2023. Models as agents: Optimizing multi-step predictions of interactive local models in model-based multi-agent reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 9, 10435–10443.
- [59] Zhiwei Xu, Bin Zhang, Yuan Zhan, Yunpeng Baiia, Guoliang Fan, et al. 2022. Mingling foresight with imagination: Model-based cooperative multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 11327–11340.
- [60] Yiqin Yang, Xiaoteng Ma, Chenghao Li, Zewu Zheng, Qiyuan Zhang, Gao Huang, Jun Yang, and Qianchuan Zhao. 2021. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 34 (2021), 10299–10312.
- [61] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. 2021. Combo: Conservative offline model-based policy optimization. *Advances in neural information processing systems* 34 (2021), 28954–28967.
- [62] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. 2020. Mopo: Model-based offline policy optimization. *Advances in Neural Information Processing Systems* 33 (2020), 14129–14142.
- [63] Lei Yuan, Yuqi Bian, Lihe Li, Ziqian Zhang, Cong Guan, and Yang Yu. 2025. Efficient multi-agent offline coordination via diffusion-based trajectory stitching. *The Thirteenth International Conference on Learning Representations*.
- [64] Wenhao Zhan, Scott Fujimoto, Zheqing Zhu, Jason D Lee, Daniel R Jiang, and Yonathan Efroni. 2024. Exploiting Structure in Offline Multi-Agent RL: The Benefits of Low Interaction Rank. *arXiv preprint arXiv:2410.01101* (2024).
- [65] Fengzhuo Zhang, Boyi Liu, Kaixin Wang, Vincent Tan, Zhuoran Yang, and Zhuoran Wang. 2022. Relational reasoning via set transformers: Provable efficiency and applications to MARL. *Advances in Neural Information Processing Systems* 35 (2022), 35825–35838.
- [66] Jiayi Zhang, Ziheng Liu, Yiyang Zhu, Enyu Shi, Bokai Xu, Chau Yuen, Dusit Niyato, Mérouane Debbah, Shi Jin, Bo Ai, et al. 2025. Multi-Agent Reinforcement Learning in Wireless Distributed Networks for 6G. *arXiv preprint arXiv:2502.05812* (2025).
- [67] Kaiqing Zhang, Sham M Kakade, Tamer Basar, and Lin F Yang. 2023. Model-based multi-agent rl in zero-sum markov games with near-optimal sample complexity. *Journal of Machine Learning Research* 24, 175 (2023), 1–53.
- [68] Ruiqi Zhang, Jing Hou, Florian Walter, Shangding Gu, Jiayi Guan, Florian Röhrbein, Yali Du, Panpan Cai, Guang Chen, and Alois Knoll. 2024. Multi-agent reinforcement learning for autonomous driving: A survey. *arXiv preprint arXiv:2408.09675* (2024).
- [69] Yihe Zhou, Yuxuan Zheng, Yue Hu, Kaixuan Chen, Tongya Zheng, Jie Song, Mingli Song, and Shunyu Liu. 2025. Cooperative Policy Agreement: Learning Diverse Policy for Offline MARL. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 21, 23018–23026.
- [70] Jin Zhu, Chunhui Du, and Geir E Dullerud. 2024. Model-Based Offline Reinforcement Learning with Uncertainty Estimation and Policy Constraint. *IEEE Transactions on Artificial Intelligence* (2024).
- [71] Zhengbang Zhu, Minghuan Liu, Liyuan Mao, Bingyi Kang, Minkai Xu, Yong Yu, Stefano Ermon, and Weinan Zhang. 2024. Madiff: Offline multi-agent learning with diffusion models. *Advances in Neural Information Processing Systems* 37 (2024), 4177–4206.