

Models as Agents: Optimizing Multi-Step Predictions of Interactive Local Models in Model-Based Multi-Agent Reinforcement Learning

Zifan Wu¹, Chao Yu^{1,2*}, Chen Chen³, Jianye Hao³, Hankz Hankui Zhuo¹

¹School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, Guangdong, China
wuzf5@mail2.sysu.edu.cn, {yuchao3,zhuohank}@mail.sysu.edu.cn.

²Pengcheng Laboratory, Shenzhen, Guangdong, China

³Huawei Noah’s Ark Lab, Beijing, China
cclvr@163.com, haojianye@huawei.com

Abstract

Research in model-based reinforcement learning has made significant progress in recent years. Compared to single-agent settings, the exponential dimension growth of the joint state-action space in multi-agent systems dramatically increases the complexity of the environment dynamics, which makes it infeasible to learn an accurate global model and thus necessitates the use of agent-wise local models. However, during multi-step model rollouts, the prediction of one local model can affect the predictions of other local models in the next step. As a result, local prediction errors can be propagated to other localities and eventually give rise to considerably large global errors. Furthermore, since the models are generally used to predict for multiple steps, simply minimizing one-step prediction errors regardless of their long-term effect on other models may further aggravate the propagation of local errors. To this end, we propose Models as AGents (MAG), a multi-agent model optimization framework that reversely treats the local models as multi-step decision making agents and the current policies as the dynamics during the model rollout process. In this way, the local models are able to consider the multi-step mutual affect between each other before making predictions. Theoretically, we show that the objective of MAG is approximately equivalent to maximizing a lower bound of the true environment return. Experiments on the challenging StarCraft II benchmark demonstrate the effectiveness of MAG.

Introduction

Model-Based Reinforcement Learning (MBRL) (Moerland, Broekens, and Jonker 2020; Luo et al. 2022) aims to improve the sample efficiency of model-free methods by learning an approximate world model and then using it to aid policy learning. Despite the success in single-agent settings, there are still limited works concentrating on MBRL in multi-agent systems. In these systems, the exponential dimension growth of the joint state-action space dramatically increases the complexity of the environment dynamics, making it infeasible to learn an accurate global model (Zhang et al. 2021; Wang, Zhang, and Zhang 2022). Thus, a common practice is to make use of local agent-wise models which only require partial information and then predict the most relevant




Figure 1: Intuition for multi-step interactions between the local models and the policies.

information for policy learning of each corresponding agent, so as to alleviate the issue of high dimension and avoid modelling the whole complicated dynamics (Kim, Park, and Sung 2020; Zhang et al. 2021).

One of the most commonly used paradigms in Multi-Agent Reinforcement Learning (MARL) is Centralized Training with Decentralized Execution (CTDE) (Lowe et al. 2017; Foerster et al. 2018; Rashid et al. 2018; Wang et al. 2021; Wu et al. 2021; Fu et al. 2022), which allows the use of global information in the policy training phase, yet retains local observability of each agent during execution. MAMBA (Egorov and Shpilman 2022), a recently proposed multi-agent MBRL method under the CTDE paradigm, achieves state-of-the-art sample efficiency in several challenging benchmarks, especially the StarCraft II challenge (Samvelyan et al. 2019). To take full advantage of the centralized training phase, MAMBA utilizes the Attention mechanism (Vaswani et al. 2017) to extract information for each local model from the global information, i.e., $(f^1, \dots, f^N) = \text{Attention}(l^1, \dots, l^N)$, where N is the number of agents, $l^i := (o^i, a^i)$ denotes the local information of agent i , o^i and a^i are the observation and action of agent i respectively, and f^i denotes the extracted feature for the local model of agent i . However, since the Attention block fuses all local predictions obtained from the local models, the prediction of each local model, i.e., $\hat{P}^i(o^i | f^i)$, can affect the subsequent predictions of other local models in the next rollout step.

Furthermore, while generally trained to simply minimize one-step prediction errors, the local models are usually not able to take into account the aforementioned multi-step er-

*corresponding author

rors induced by the interactions between the local models and policies. As a result, prediction errors of one local model can be propagated to the others and eventually induce large accumulative global errors during multi-step model rollouts, which would hinder the learning of policies.

Figure 1 gives an intuitive illustration of the above discussion: Given local features (f_t^1, f_t^2) w.r.t. agent 1 and 2 at step t , 1) the local models predict $\hat{P}^1(o_{t+1}^1|f_t^1) = 1$ and $\hat{P}^2(o_{t+1}^2|f_t^2) = \hat{P}^2(\tilde{o}_{t+1}^2|f_t^2) = 50\%$; and 2) under the previous joint policy $\pi_{\text{old}}^{1,2}$, both prediction of the next joint observation, i.e., (o_{t+1}^1, o_{t+1}^2) and $(o_{t+1}^1, \tilde{o}_{t+1}^2)$, will lead the trajectory to go into regions with low values predicted by the value function, hence $\pi_{\text{old}}^{1,2}$ is updated to $\pi_{\text{new}}^{1,2}$ to explore regions with potential high values. But under the updated joint policy $\pi_{\text{new}}^{1,2}$, the subsequent rollout trajectory starting from $(o_{t+1}^1, \tilde{o}_{t+1}^2)$ would lead to considerably larger model errors compared to the trajectory starting from (o_{t+1}^1, o_{t+1}^2) . Thus, to reduce accumulative model errors along rollout trajectories, the local models should learn to coordinate with each other while quickly adapting to the current joint policy. Formally, as will be shown in the method section, smaller accumulative model errors could provide stronger performance guarantee.

In this work, we propose Models as AGents (MAG), a multi-agent model optimization framework which considers the interactions between local models during multi-step model rollout. Based on the MAMBA framework, the whole environment dynamics is decomposed into agent-wise local models, and our key idea lies in reversely considering the local models as multi-step decision makers while fixing the current joint policy to serve as the environment. During model learning, the local models perform multi-step interactions with each other as well as the policies, so as to take the long-term global effect of immediate local predictions into account and generate trajectories with less accumulative errors. Theoretically, we show the necessity of considering the local model interactions and minimizing the multi-step accumulative errors. Empirically, the results on several challenging tasks in the StarCraft II benchmark demonstrate that MAG significantly outperforms MAMBA in low data regime, and the model error analysis further verifies the effectiveness of our model learning mechanism.

Background

In this section, we first introduce the problem setting of MARL and MBRL, and then give a brief description of MAMBA, the aforementioned state-of-the-art model-based MARL method.

MARL In this work, we focus on the fully cooperative multi-agent systems that can be formalized as Dec-POMDPs (Oliehoek and Amato 2016), which are defined by tuple $(N, S, \Omega, O, A, R, P, \gamma)$, where N is the number of agents, S the set of global states, Ω the observation space shared by the agents, $O(s, i)$ the function deriving partial observations for each agent i from a global state $s \in S$, A the action space, $R(s, a^1, \dots, a^N)$ a shared scalar reward function that takes $s \in S$ and $a^i \in A, i \in \{1, \dots, N\}$ as input,

and $\gamma \in [0, 1)$ the discount factor. Each agent has an action-observation history $\tau^i \in T \equiv (\Omega \times A)^*$. We use the bold symbol $\mathbf{o}, \mathbf{a}, \boldsymbol{\pi}$ to denote the joint observation $\{o^1, \dots, o^N\}$, action $\{a^1, \dots, a^N\}$ and policy $\{\pi^1, \dots, \pi^N\}$, respectively. At each timestep, agent i chooses an action $a^i \in A$ according to its policy $\pi^i(a^i|\tau^i)$ (We replace τ^i by o^i in our analysis for brevity). The environment then returns the reward signal $R(s, \mathbf{a})$ and shifts to the next state according to the transition function $P(s'|s, \mathbf{a})$. The expected return of joint policy $\boldsymbol{\pi}$ is defined by $J(\boldsymbol{\pi}) := \mathbb{E}_{\boldsymbol{\pi}}[\sum_{t'=0}^{\infty} \gamma^{t'} R_{t+t'}|s_t, \mathbf{a}_t]$. Some previous works (Wang et al. 2020; Liu et al. 2020) have shown that it is possible to significantly reduce the state space in large environments to only relevant information for the agents' decision making. Hence, in this paper we assume that the joint observation-action, i.e., (\mathbf{o}, \mathbf{a}) , is sufficient to predict the next joint observation \mathbf{o}' and the global reward R . Serving as a special case of Dec-POMDPs, MMDPs (Boutilier 1996) assume global observability of each agent and are adopted to reformulate the model rollout process in our work.

MBRL MBRL methods learn a model \hat{P} that approximates the unknown dynamics P , and then use this model to assist policy learning. While the model can be utilized in various ways (Feinberg et al. 2018; D'Oro and Jaśkowski 2020; Curi, Berkenkamp, and Krause 2020; Song and Sun 2021; Amos et al. 2021), this work focuses on one of the most common usages, i.e., generating pseudo samples to enrich the dataset, so as to accelerate policy learning and reduce interactions with the true environment (Sutton 1991; Chua et al. 2018; Luo et al. 2019; Janner et al. 2019). The expected return of policy $\boldsymbol{\pi}$ predicted by model \hat{P} is denoted as $J^{\hat{P}}(\boldsymbol{\pi}) := \mathbb{E}_{\boldsymbol{\pi}, \hat{P}}[\sum_{t'=0}^{\infty} \gamma^{t'} R_{t+t'}|s_t, \mathbf{a}_t]$. As a state-of-the-art MBRL method in discrete environments, Dreamer V2 (Hafner et al. 2020) makes use of the RSSM model (Hafner et al. 2019) to learn the dynamics of the environment in the latent space by minimizing the evidence lower bound (Kingma and Welling 2013).

MAMBA Building upon Dreamer V2, MAMBA (Egorov and Shpilman 2022) also learns the environment dynamics in the latent space, and makes use of the Attention mechanism (Vaswani et al. 2017) to extract features for each local models from global information. To disentangle the agents' latent space and encourage the local models to be mutually independent when making predictions, MAMBA proposes to maximize the mutual information between the latent state and the previous action of the corresponding agent. In addition, the method allows communicating with the neighbouring agents via discrete messages to sustain world models during the execution phase, thus regarding world models as an instance of communication. To the best of our knowledge, MAMBA is the first model-based MARL method that improves the sample efficiency of model-free methods by an order of magnitude on the challenging StarCraft II benchmark. Nevertheless, compared to the performance of Dreamer V2 in Atari games (Bellemare et al. 2013) and MBPO (Janner et al. 2019) in the MuJoCo (Todorov, Erez, and Tassa 2012) benchmark, the overall improvement

of sample efficiency, as well as the asymptotic performances in some difficult tasks achieved by MAMBA are still relatively limited, which may be due to the high complexity of the dynamics of multi-agent systems.

Method

In this section, we first propose a theoretical result of how the prediction errors of agent-wise local models affect the overall policy performance, based on which we reformulate the model rollout process as a multi-agent sequential decision making problem. In the last subsection, we present the practical implementation of MAG and further detail some important steps in the algorithm.

Theoretical Result

Since general MBRL methods optimize the policy by maximizing the expected return predicted by the model, one of the most crucial theoretical problems for MBRL is to bound the gap between the model predicted return and the true environment return. Our major theoretical result is the following theorem that bounds the performance gap:

Theorem 1. Denoting the set of local models by $\hat{P} := \{\hat{P}^i\}_{i=1}^N$ and the data-collecting policy obtained in the last iteration by π_D , the gap between the expected return of the model and the environment can be bounded as ¹:

$$\left| J(\pi) - J^{\hat{P}}(\pi) \right| \leq \frac{R_{max}}{(1-\gamma)^2} \left(2\epsilon_\pi + (1-\gamma) \sum_{t=1}^{\infty} \gamma^t \epsilon_{m_t} \right), \quad (1)$$

where $\epsilon_\pi := \max_{\mathcal{O}} D_{TV}(\pi_D(\cdot|\mathcal{O})\|\pi(\cdot|\mathcal{O}))$ denotes the distribution shift of the joint policy between two consecutive iterations, $\epsilon_{m_t} :=$

$$\mathbb{E}_{\mathcal{O} \sim \hat{P}_{t-1}(\cdot|\pi)} \left[\max_{\mathbf{a}} \sqrt{2 \sum_{i=1}^N \mathbb{E}_{\mathcal{O}' \sim \hat{P}^i(\cdot|\mathcal{O}, \mathbf{a})} \left[\log \frac{\hat{P}^i(\mathcal{O}'|\mathcal{O}, \mathbf{a})}{N \sqrt{P(\mathcal{O}'|\mathcal{O}, \mathbf{a})}} \right]} \right]$$

denotes the upper bound of the i -th model's error at timestep t of the model rollout trajectory, $\hat{P}_{t-1}(\mathcal{O}; \pi)$ denotes the distribution of joint observation at $t-1$ under \hat{P} and π , and $R_{max} := \max_{s, \mathbf{a}} R(s, \mathbf{a})$.

Proof. Please refer to Appendix A. \square

It is worth noting that Theorem 1 is not simply a multi-agent version of the results that have been derived in the single-agent setting (Luo et al. 2019; Janner et al. 2019). The key difference is that Theorem 1 does not scale up the step-wise model prediction errors (i.e., ϵ_{m_t}) to their maximum over timesteps, which not only leads to a tighter bound that provides stronger guarantee for policy improvement (see Appendix A for proof), but also indicates how the interactions between local models affect the overall performance error bound: Note that by definition the model error

¹In our theoretical analysis, the reward function is assumed to be known. Note that this is a commonly adopted assumption since the sample complexity of learning the reward function with supervised learning is a lower order term compared to the one of learning the transition model (Gheshlaghi Azar, Munos, and Kappen 2013).

at step t , i.e., ϵ_{m_t} , depends on the distribution of the joint observation at the last timestep, i.e., $\hat{P}_{t-1}(\mathcal{O}; \pi)$, and except for the first step of rollout trajectories, this distribution further depends on the current policies π and the prediction of other local models at the last timestep, i.e., $\hat{P}_t(\mathcal{O}; \pi) = \mathbb{E}_{\mathcal{O} \sim \hat{P}_{t-1}(\cdot|\pi), \mathbf{a} \sim \pi(\cdot|\mathcal{O})} \left[\prod_{i=1}^N \hat{P}^i(\mathcal{O}^i|\mathcal{O}, \mathbf{a}) \right]$. Thus, the errors of the local models can affect each other during multi-step model rollout, and this mutual affect can largely determine the tightness of the overall error bound.

Based on this result, a performance lower bound with regard to the policy shift and the model error can be written as: $J(\pi) \geq J^{\hat{P}}(\pi) - 2C(\hat{P}, \pi)$, where $C(\hat{P}, \pi)$ denotes the right hand side of Eq. (1). Then, in an ideal manner, applying the following update rule repeatedly can **guarantee the monotonic improvement** of the joint policy:

$$\hat{P}, \pi \leftarrow \arg \max_{\hat{P}, \pi} J^{\hat{P}}(\pi) - C(\hat{P}, \pi). \quad (2)$$

The update rule in Eq. (2) is often impractical since it involves an exhaustive search in the joint state-action space to compute C , and requires full-horizon rollouts in the model for estimating the accumulative model errors. Thus, similar to how algorithms like TRPO (Schulman et al. 2015) approximate their theoretically monotonic version, this update rule can be approximated by maximizing the expected model return (i.e., $J^{\hat{P}}(\pi)$) while keeping the accumulative model error (i.e., $\sum_t \gamma^t \epsilon_{m_t}$) small. As for the policy shift term ϵ_π , though the bound suggests that this term should also be constrained, we found empirically that it is sufficient to only control the model error. This may be explained by the relatively small scale of policy shift w.r.t. the model error, as observed in (Janner et al. 2019).

By treating $-\epsilon_{m_t}$ as the ‘‘reward’’ shared by the local models at timestep t , the learning of the local models can be regarded as an optimization process of multi-step predictions of the local models, where the objective is to minimize the global prediction errors accumulated along the model rollout trajectories. Note that the definition of ϵ_{m_t} involves the expectation under the current joint policy π , thus during model learning, the joint policy can be fixed to serve as a background environment, while the local models reversely play the role of decision-makers that should learn to maximize the ‘‘expected return’’ (i.e., $-\mathbb{E}_{\pi, \hat{P}}[\sum_t \gamma^t \epsilon_{m_t}]$) under the current joint policy. Building on the above theoretical intuition, we now propose the MAG framework in the next subsection.

Problem Reformulation

To formalize the intuition of reversing the roles of the models and the agents during model learning, we first define the *model MMDP* to reformulate the model rollout process and then outline the overall model optimization of MAG as a generic solution to the reformulated problem.

Definition 1. The model MMDP is defined by tuple $(N, \gamma, S_m, A_m, P_m, R_m)$, where N is the number of local models, γ is the discount factor, S_m, A_m, P_m and R_m are

the model-state space, the model-action space, the model-transition function and the scalar model-reward function, respectively.

At each timestep t , each local model \hat{P}^i receives model-state $s_{m_t} := (\mathbf{o}_t, \mathbf{a}_t) \in S_{m_t}$, then takes a model-action $a_{m_t}^i := o_{t+1}^i$ according to its ‘‘policy’’ $\hat{P}^i(a_{m_t}^i | s_{m_t})$. After that, the model-transition function returns the next model-state by $P_m(s_{m_{t+1}} | s_{m_t}, \mathbf{a}_{m_t}) := \pi(\mathbf{a}_{t+1} | \mathbf{o}_{t+1}) \prod_{i=1}^N \hat{P}(o_{t+1}^i | \mathbf{o}_t, \mathbf{a}_t)$, while the model-reward function returns a scalar reward by $R_m(s_{m_t}, \mathbf{a}_{m_t}) := \sum_{i=1}^N \log \frac{\hat{P}^i(o_{t+1}^i | \mathbf{o}_t, \mathbf{a}_t)}{N \sqrt{P(\mathbf{o}_{t+1} | \mathbf{o}_t, \mathbf{a}_t)}}$.

Using the model MMDP formulation, the model learning phase can be viewed as a multi-agent learning problem, where the current joint policy is fixed to serve as the environment dynamics and the local models, now as the decision makers, interact with each other and learn to minimize the accumulative prediction error under the current joint policy. From this perspective, the local models trained by minimizing one-step prediction errors for each individual can be intuitively interpreted as greedy independent learners, which are often considered shortsighted and may struggle to learn cooperative behaviors. To minimize the accumulative global errors, the local models must instead consider the long-term global effect of immediate local predictions.

Note that in the competitive or mixed cooperative-competitive scenarios, the goal of each local model is generally to assist policy learning of only one individual agent, thus in those scenarios the local models would aim at minimizing the individual accumulative errors instead of the global summation of model errors. Consequently, in those scenarios, the model rollout process can be defined as a *Markov Game* (Shapley 1953), where the reward function can be defined respectively for each local model. Since the major focus of this work is the fully cooperative scenarios, we leave the above discussion as a possible motivation for future work.

Similar to the optimization of the policy, the objective of model learning can be written as $\arg \max_{\hat{P}} J^\pi(\hat{P})$ where $J^\pi(\hat{P}) := \mathbb{E}_{\pi, \hat{P}}[\sum_t \gamma^t R_m(s_{m_t}, \mathbf{a}_{m_t})]$. Due to this duality between the learning of the policies and the models, we call this overall model-based MARL method by Models as AGents (MAG). Specifically, during model learning, the local models first generate samples by actively interacting with the current joint policy (now viewed as the background environment), and then optimize the expected return $J^\pi(\hat{P})$ accordingly.

Practical Implementation

To give a practical solution to the model MMDP, we describe the implementation of MAG in this subsection.

The Overall Algorithm Algorithm 1 gives the overall algorithm design of MAG. In each outer loop, the current joint policy is applied in the real environment to collect an episode of real-world data, which is then added to the environment dataset \mathcal{D}_e (Line 3). Then, the local models are pre-trained by traditional one-step prediction loss $\sum_{i=1}^N \|\hat{o}^{i'} -$

Algorithm 1: MAG

- 1: Initialize joint policy π , predictive local models $\hat{P}^i, i = 1, 2, \dots, N$, model-reward predictor \hat{R}_m , environment dataset \mathcal{D}_e and model dataset \mathcal{D}_m .
 - 2: **for** N episodes **do**
 - 3: Collect an episode of real-environment data using π and then add the data to \mathcal{D}_e ;
 - 4: Train models $\{\hat{P}^i\}_{i=1}^N$ on \mathcal{D}_e via one-step prediction loss;
 - 5: Train \hat{R}_m on \mathcal{D}_e via supervised learning on \mathcal{D}_e ;
 - 6: **for** M model rollouts **do**
 - 7: Sample joint observations \mathbf{o} uniformly from \mathcal{D}_e and use them as the initial observations for our rollout trajectories;
 - 8: **for** k rollout steps **do**
 - 9: Each agent takes action a^i according to $\pi^i(\cdot | \mathbf{o}^i)$;
 - 10: Initialize $s_{m_0} = (\mathbf{o}, \mathbf{a})$ and perform L parallelized rollouts for H steps by actively interacting the local models with π : $\{s_{m_{0,j}}, \mathbf{a}_{m_{0,j}}, R_{m_{0,j}}, s_{m_{1,j}}, \mathbf{a}_{m_{1,j}}, \dots, s_{m_{H,j}}\}_{j=1}^L$, where $s_{m_{0,j}} = s_{m_0}, \forall j \in \{1, 2, \dots, N\}$;
 - 11: Compute $r_{m_j} = \sum_{t=0}^{H-1} \hat{R}_m(s_{m_{t,j}}, \mathbf{a}_{m_{t,j}})$ for each rollout trajectory j ;
 - 12: Take $(\mathbf{o}', \mathbf{R}) = \mathbf{a}_{m_{0, \arg \max_j r_{m_j}}}$;
 - 13: Store $(\mathbf{o}, \mathbf{a}, \mathbf{R}, \mathbf{o}')$ to \mathcal{D}_m and then let $\mathbf{o} = \mathbf{o}'$;
 - 14: **end for**
 - 15: **end for**
 - 16: **for** G gradient updates **do**
 - 17: Update π using data sampled from \mathcal{D}_m ;
 - 18: **end for**
 - 19: **end for**
-

$\hat{o}^{i'} + \|\hat{R}^i - R\|$, where $\hat{o}^{i'}, \hat{R}^i \sim \hat{P}^i(\cdot, \cdot | \mathbf{o}, \mathbf{a})$ and each transition $(\mathbf{o}, \mathbf{a}, R, \mathbf{o}')$ is sampled from the environment dataset (Line 4). Since the reward function R is generally not available in practice, each local model is also trained to predict the global reward respectively given (\mathbf{o}, \mathbf{a}) . Besides, it deserves to be noted that we do not directly use the above pre-trained local models to obtain the predictions during model rollout, but instead optimize the multi-step predictions of local models via a planning process. In Line 5, MAG trains the \hat{R}_m network to approximate R_m , since by definition the model-reward R_m involves the true environment dynamics and thus cannot be directly computed. The approximation of R_m will be detailed later on. Lines 6-15 give the model rollout process where M parallelized trajectories of length k are generated based on different initial observations sampled from \mathcal{D}_e . For each rollout step, before predicting the next observation, the local models first treat the current joint policy as the ‘‘dynamics’’ and then perform H -step ($H \leq k$) planning to obtain the best predictions for the current step (Lines 10-12). This is the core of MAG and will be detailed later. Finally, the pseudo samples generated by the model are added to the model dataset, which is then used for policy learning. Specifically, we adopt PPO (Schulman et al. 2017) as the underlying policy optimization method and use global

information that has been processed by the Attention block as the input of the critic.

Approximating R_m We approximate R_m by training a neural network \hat{R}_m which takes transitions $(\mathbf{o}, \mathbf{a}, R, \mathbf{o}')$ sampled from the environment dataset as inputs, and the model prediction errors on these transitions as labels. The prediction error of an environment transition is computed via $\sum_{i=1}^N \|\hat{o}^{i'} - o^{i'}\| + \|\hat{R}^i - R\|$, where $\hat{o}^{i'}$ and \hat{R}^i are sampled from $\hat{P}^i(\cdot, \cdot | \mathbf{o}, \mathbf{a})$. Intuitively, \hat{R}_m can be seen as an indicator that informs the models where their “weaknesses” lie in. Additionally, since Dreamer V2 utilizes VAE (Kingma and Welling 2013) and learns the dynamics in a latent space, the actual loss of the dynamics consists of a reconstruction loss of the auto-encoder and a KL divergence loss that aim to minimize the distance between the prior and the posterior of the latent state. Consequently, computing the model errors in Dreamer V2 can be interpreted as computing the prediction errors in the latent space, and thus is equivalent to computing $\sum_{i=1}^N \|\hat{o}^{i'} - o^{i'}\| + \|\hat{R}^i - R\|$ in principle.

Planning to Predict Since in our problem formulation the “dynamics” of the model rollout process (i.e., the current joint policy π) is accessible, one of the simplest yet effective approaches to learn the models can be the Model Predictive Control (MPC) (Camacho and Alba 2013), which utilizes the dynamics to plan and optimize for a sequence of actions. Given the state s_{m_t} at step t , the MPC controller first optimizes the sequence of actions $\mathbf{a}_{m_{t:t+H}}$ over a finite horizon H , and then employs the first action of the optimal action sequence, i.e., $\mathbf{a}_{m_{H,t}} := \arg \max_{\mathbf{a}_{m_{t:t+H}}} \mathbb{E}_{\hat{P}, \pi} \sum_{t'=t}^{t+H-1} R_m(s_{m_{t'}}, \mathbf{a}_{m_{t'}})$. Computing the exact $\arg \max_{\mathbf{a}_{m_{t:t+H}}}$ requires a complete search in a space of dimension $|A_m|^{N \cdot H}$, which is impractical in most scenarios. Thus, as specified from Lines 10-12 in Algorithm 1, we adopt the random-sampling shooting method (Rao 2009) which generates L random action sequences, executes them respectively, and chooses the one with the highest return predicted by the dynamics. Essentially, this planning process is a simulation of the interactions between the local models and the current joint policy, according to which each local model chooses the best prediction that approximately minimizes the global model error in concert with the other local predictions, thus achieving the coordination between local models.

Experiments

In this section, we present an empirical study of MAG on the challenging StarCraft II benchmark (SMAC) (Samvelyan et al. 2019). In the first subsection, we provide the overall comparison between MAG and several baselines. Then, we provide a quantitative analysis on the multi-step prediction loss to verify the effectiveness of our algorithm design in model learning. In the last subsection, we conduct ablation studies to show how the choices of the planning horizon (i.e., H in Algorithm 1) and the number of random shooting trajectories (i.e., L in Algorithm 1) affect the overall performance.

Comparative Evaluation

Baselines We compare MAG with a model-based baseline and several model-free baselines. The model-based baseline is MAMBA (Egorov and Shpilman 2022), a recently proposed multi-agent MBRL method that achieves state-of-the-art sample efficiency in several SMAC tasks. The model-free baselines include 1) Attention-PPO, the model-free counterpart of both MAG and MAMBA which equips PPO (Schulman et al. 2017) with centralized attention-critics and communication during execution; 2) G2A (Liu et al. 2020), which adopts a two-stage attention architecture to realize communication between agents; and 3) CommNet (Sukhbaatar, Fergus et al. 2016), which applies LSTM (Hochreiter and Schmidhuber 1997) to learn continuous communication protocols for partially observable environments. In addition, it deserves to note that MAG is essentially a flexible plug-in component which can be employed by most model-based methods to improve the learning of the model. In our comparisons, we plug the model learning process of MAG into MAMBA.

Environments The methods are evaluated on 8 maps of SMAC, ranging from *Easy* maps (2s_vs_1sc, 2s3z, 3s_vs_3z), *Hard* maps (3s_vs_4z, 3s_vs_5z, 2c_vs_64zg) and *Super Hard* maps (corridor, 3s5z_vs_3s6z).

Implementation Details The implementation of MAG is overall built on MAMBA.² For more details of the hyperparameter settings, please refer to Appendix B.

Results The overall results shown in Figure 2 demonstrate that MAG consistently outperforms all the baselines in low data regime. The comparison between MAG and MAMBA verifies the effectiveness of optimizing multi-step prediction errors that are induced by the interactions between local models. Besides, note that except for Attention-PPO in 2c_vs_64zg, all model-free baselines fail to even achieve a non-zero win rate in such low data regimes, showing the significant improvement of sample efficiency resulted from using a world model.

Model Error Analysis

Based on the theoretical result presented in the method section, the core idea of MAG is to reverse the roles played by the local models and the current joint policy, thus treating the models as decision-makers interacting with each other and aiming at minimizing the global accumulative model error. To validate the effectiveness of this algorithmic design, we empirically study the accumulative prediction error on the 2c_vs_64zg map. While the real dynamics is unavailable during training, the error is approximated by a neural network trained on the environment dataset, i.e., \hat{R}_m .

The result in Figure 3 demonstrates that as the model rollout trajectories go longer, the accumulative model errors of MAMBA become significantly larger than that of MAG, which not only validates the effectiveness of MAG in reducing the accumulative model errors, but also provides a solid support for our theoretical result derived in the

²Code available at <https://github.com/ZifanWu/MAG>.

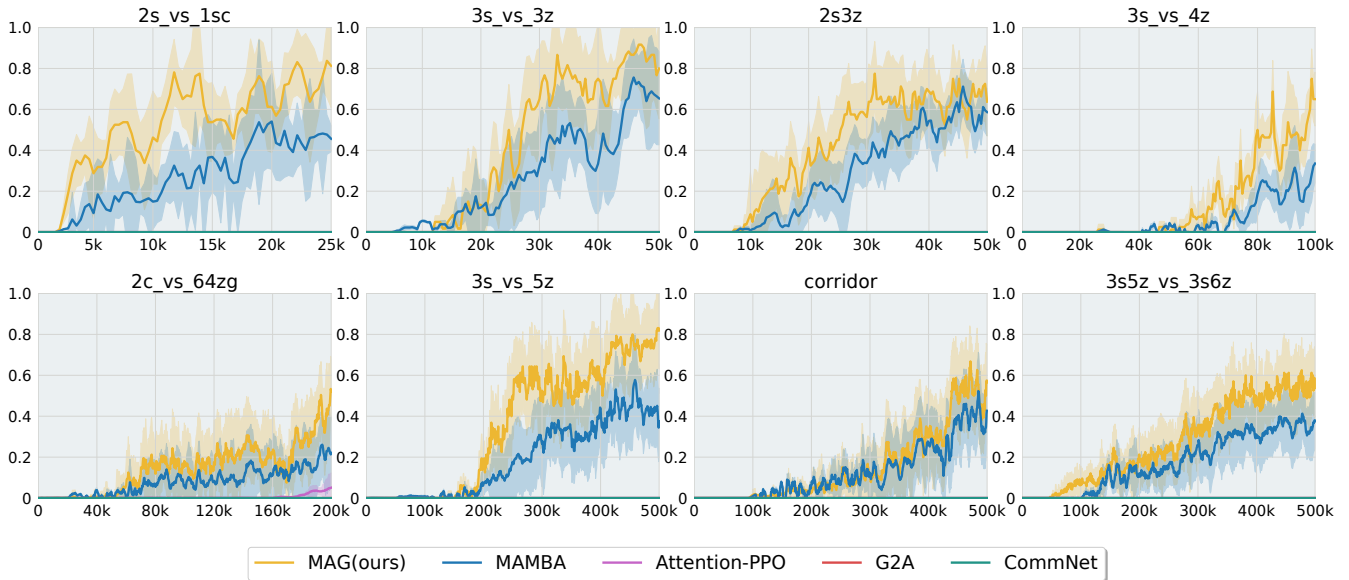


Figure 2: Comparisons against baselines on SMAC. Solid curves represent the mean of runs over 5 different random seeds, and shaded regions correspond to standard deviation among these runs. X axis denotes the number of steps taken in the real environment and Y axis denotes the win rate.




Figure 3: The difference of the accumulative model errors between MAMBA and MAG (the accumulative errors of MAMBA minus the accumulative errors of MAG), on the 2c_vs_64zg map.

method section, i.e., the method inducing less accumulative error is likely to achieve better performance. Besides, we can also observe that in the first two steps the model errors induced by MAG are slightly larger than the errors of the baseline. This further agrees with the intuition mentioned in the method section by showing that MAG is able to trade the one-step greedy model error for the accumulative error by considering the long-term effect of the immediate prediction.

Ablation Studies

According to the descriptions in Algorithm 1, apart from easy implementation, another advantage of utilizing MPC in

optimizing the accumulative model-reward is that this only introduces a small number of extra hyperparameters. Specifically, there are mainly two extra hyperparameters that need to be tuned, i.e., the planning horizon H and the number of random shooting trajectories L . The ablation results of these two hyperparameters are shown in Figure 4 and Figure 5 respectively, which indicate that longer planning horizons and more random trajectories always induce better performance. Since increasing H and L leads to a rapid growth in terms of the computational complexity and the memory cost, the finally adopted settings of the two hyperparameters, which are detailed in Appendix B, can be regarded as a compromise between this practical limit and the performance.

Related Works

The research of MBRL can be roughly divided into two lines: the model usage and the model learning. This work focuses on model learning and adopts the most common model usage, that is, generating pseudo samples to enrich the data buffer, so as to reduce the interaction with the environment and accelerate policy learning (Sutton 1990, 1991; Deisenroth et al. 2013; Kalweit and Boedecker 2017; Luo et al. 2019; Janner et al. 2019; Pan et al. 2020). Most of previous works in MBRL train the model simply by minimizing each one-step prediction error for transitions available in the environment dataset (Kurutach et al. 2018; Chua et al. 2018; Janner et al. 2019). However, in the multi-agent setting, the dimension of the joint observation-action space grows rapidly w.r.t. the number of agents, making it impractical to learn a global model for such complex environments (Zhang et al. 2021; Wang, Zhang, and Zhang 2022). Thus, a common approach is to train a local model for each agent which




Figure 4: Ablation study of the planning horizon H .




Figure 5: Ablation study of the the number of random shooting trajectories L .

takes partial observations as input and predicts relevant information for the agent’s policy (Kim, Park, and Sung 2020; Zhang et al. 2021). Therefore, MAMBA (Egorov and Shpilman 2022) proposes to extract relevant information for each local model from the global information via the Attention mechanism (Vaswani et al. 2017), so as to avoid modelling the whole complicated dynamics and accelerate the model learning. Although the Attention mechanism is effective in extracting information for different local models, the fuse of local information during multi-step model rollout may lead to the propagation of prediction errors between different local models, as discussed in the first section. To address this issue, we reformulate the model rollout process as the *model MMDP* where the current joint policy is fixed to serve as the background environment and the local models are reversely regarded as decision-makers aiming at minimizing the global accumulative model error.

In the single-agent setting, some works have attempted to learn the model by treating the model rollout process as a sequential decision-making problem. Shang et al. (2019) propose an environment reconstruction method which models the influence of the hidden confounder on the environment by treating the platform, the user and the confounder as three agents interacting with each other. They focus on the offline setting (i.e., RL-based recommendation) and simultaneously train the model and the policy using a multi-agent imitation learning method. Xu, Li, and Yu (2020) treat the model as a dual agent and analyze the error bounds of the model. They propose to train the model using imitation learning methods. Chen et al. (2022) also consider multi-step model errors, yet they mainly focus on handling counterfactual data queried by adversarial policies. Note that both (Xu, Li, and Yu 2020) and (Chen et al. 2022) focus solely on model learning in the single-agent setting and do not combine with the policy learning phase.

There are also works considering multi-step prediction loss in the single-agent setting (Nagabandi et al. 2020; Luo et al. 2019). The essential difference between their multi-step loss and ours is that their loss is computed over the trajectories sampled from the environment dataset (collected by previous policies), while MAG minimizes the multi-step loss on the trajectories generated by active interactions between the local models as well as the current joint policy. From the theoretical perspective, the model error term in Theorem 1 is defined by the expectation over the current joint policy and the current local models, thus computing the multi-step loss on the trajectories generated by these current policy and current models can better approximate the lower bound, which guarantees better policy improvement.

Conclusion and Future Work

In this work, we first study how the prediction errors of agent-wise local models affect the performance lower bound, which necessitates the considerations of the interactions between models during multi-step model rollout. Based on this theoretical result, we reformulate the model rollout process as the *model MMDP* by treating the local models as multi-step decision-makers and the current policies as the background environment. We then propose a multi-agent model learning framework, i.e., MAG, to maximize the accumulative global “model-reward” defined in the *model MMDP* by considering the interactions between local models. We provide a practical implementation of MAG to optimize the above objective using the model predictive control. Empirical results on the challenging StarCraft II benchmark validates the effectiveness of our algorithmic design by showing that MAG outperforms the compared baselines. In the future, we plan to study the problem of learning local models in the competitive or mixed cooperative-competitive scenarios, which can be seen as learning in a *Markov Game*.

Acknowledgments

We gratefully acknowledge the support from the National Natural Science Foundation of China (No.62076259), the Fundamental and Application Research Funds of Guangdong province (No.2023A1515012946), and the Fundamental Research Funds for the Central Universities-Sun Yat-sen University.

References

- Amos, B.; Stanton, S.; Yarats, D.; and Wilson, A. G. 2021. On the model-based stochastic value gradient for continuous reinforcement learning. In *Learning for Dynamics and Control*, 6–20. PMLR.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Boutillier, C. 1996. Planning, learning and coordination in multiagent decision processes. In *Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge*, 195–210.
- Camacho, E. F.; and Alba, C. B. 2013. *Model predictive control*. Springer science & business media.
- Chen, X.-H.; Yu, Y.; Zhu, Z.-M.; Yu, Z.; Chen, Z.; Wang, C.; Wu, Y.; Wu, H.; Qin, R.-J.; Ding, R.; et al. 2022. Adversarial Counterfactual Environment Model Learning. *arXiv preprint arXiv:2206.04890*.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in neural information processing systems*, 31.
- Curi, S.; Berkenkamp, F.; and Krause, A. 2020. Efficient model-based reinforcement learning through optimistic policy search and planning. *Advances in Neural Information Processing Systems*, 33: 14156–14170.
- Deisenroth, M. P.; Neumann, G.; Peters, J.; et al. 2013. A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2): 388–403.
- D’Oro, P.; and Jaśkowski, W. 2020. How to learn a useful critic? Model-based action-gradient-estimator policy optimization. *Advances in Neural Information Processing Systems*, 33: 313–324.
- Egorov, V.; and Shpilman, A. 2022. Scalable Multi-Agent Model-Based Reinforcement Learning. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 381–390.
- Feinberg, V.; Wan, A.; Stoica, I.; Jordan, M. I.; Gonzalez, J. E.; and Levine, S. 2018. Model-based value estimation for efficient model-free reinforcement learning. *arXiv preprint arXiv:1803.00101*.
- Foerster, J.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Fu, W.; Yu, C.; Xu, Z.; Yang, J.; and Wu, Y. 2022. Revisiting Some Common Practices in Cooperative Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*, 6863–6877. PMLR.
- Gheshlaghi Azar, M.; Munos, R.; and Kappen, H. J. 2013. Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3): 325–349.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2019. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, 2555–2565. PMLR.
- Hafner, D.; Lillicrap, T. P.; Norouzi, M.; and Ba, J. 2020. Mastering Atari with Discrete World Models. In *International Conference on Learning Representations*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Janner, M.; Fu, J.; Zhang, M.; and Levine, S. 2019. When to trust your model: Model-based policy optimization. *Advances in Neural Information Processing Systems*, 32.
- Kalweit, G.; and Boedecker, J. 2017. Uncertainty-driven imagination for continuous deep reinforcement learning. In *Conference on Robot Learning*, 195–206. PMLR.
- Kim, W.; Park, J.; and Sung, Y. 2020. Communication in multi-agent reinforcement learning: Intention sharing. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*.
- Liu, Y.; Wang, W.; Hu, Y.; Hao, J.; Chen, X.; and Gao, Y. 2020. Multi-agent game abstraction via graph attention neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7211–7218.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Luo, F.-M.; Xu, T.; Lai, H.; Chen, X.-H.; Zhang, W.; and Yu, Y. 2022. A Survey on Model-based Reinforcement Learning. *arXiv preprint arXiv:2206.09328*.
- Luo, Y.; Xu, H.; Li, Y.; Tian, Y.; Darrell, T.; and Ma, T. 2019. Algorithmic Framework for Model-based Deep Reinforcement Learning with Theoretical Guarantees. In *ICLR (Poster)*.
- Moerland, T. M.; Broekens, J.; and Jonker, C. M. 2020. Model-based reinforcement learning: A survey. *arXiv preprint arXiv:2006.16712*.
- Nagabandi, A.; Konolige, K.; Levine, S.; and Kumar, V. 2020. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, 1101–1112. PMLR.
- Oliehoek, F. A.; and Amato, C. 2016. *A concise introduction to decentralized POMDPs*. Springer.

- Pan, F.; He, J.; Tu, D.; and He, Q. 2020. Trust the model when it is confident: Masked model-based actor-critic. *Advances in neural information processing systems*, 33: 10537–10546.
- Rao, A. V. 2009. A survey of numerical methods for optimal control. *Advances in the Astronautical Sciences*, 135(1): 497–528.
- Rashid, T.; Samvelyan, M.; Schroeder, C.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2018. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International conference on machine learning*, 4295–4304. PMLR.
- Samvelyan, M.; Rashid, T.; de Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J. N.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. In *AAMAS*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shang, W.; Yu, Y.; Li, Q.; Qin, Z.; Meng, Y.; and Ye, J. 2019. Environment reconstruction with hidden confounders for reinforcement learning based recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 566–576.
- Shapley, L. S. 1953. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100.
- Song, Y.; and Sun, W. 2021. Pc-mlp: Model-based reinforcement learning with policy cover guided exploration. In *International Conference on Machine Learning*, 9801–9811. PMLR.
- Sukhbaatar, S.; Fergus, R.; et al. 2016. Learning multiagent communication with backpropagation. *Advances in neural information processing systems*, 29.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Machine learning proceedings 1990*, 216–224. Elsevier.
- Sutton, R. S. 1991. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4): 160–163.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2021. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *International Conference on Learning Representations*.
- Wang, W.; Yang, T.; Liu, Y.; Hao, J.; Hao, X.; Hu, Y.; Chen, Y.; Fan, C.; and Gao, Y. 2020. From few to more: Large-scale dynamic multiagent curriculum learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7293–7300.
- Wang, X.; Zhang, Z.; and Zhang, W. 2022. Model-based Multi-agent Reinforcement Learning: Recent Progress and Prospects. *arXiv preprint arXiv:2203.10603*.
- Wu, Z.; Yu, C.; Ye, D.; Zhang, J.; Zhuo, H. H.; et al. 2021. Coordinated proximal policy optimization. *Advances in Neural Information Processing Systems*, 34: 26437–26448.
- Xu, T.; Li, Z.; and Yu, Y. 2020. Error bounds of imitating policies and environments. *Advances in Neural Information Processing Systems*, 33: 15737–15749.
- Zhang, W.; Wang, X.; Shen, J.; and Zhou, M. 2021. Model-based multi-agent policy optimization with adaptive opponent-wise rollouts. *arXiv preprint arXiv:2105.03363*.