

Imagine, Initialize, and Explore: An Effective Exploration Method in Multi-Agent Reinforcement Learning

Zeyang Liu, Lipeng Wan, Xinrui Yang, Zhuoran Chen, Xingyu Chen, Xuguang Lan*

National Key Laboratory of Human-Machine Hybrid Augmented Intelligence

National Engineering Research Center for Visual Information and Application

Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an, China, 710049

{zeyang.liu, wanlipeng, xinrui.yang, zhuoran.chen}@stu.xjtu.edu.cn, xingyuchen1990@gmail.com, xglan@mail.xjtu.edu.cn

Abstract

Effective exploration is crucial to discovering optimal strategies for multi-agent reinforcement learning (MARL) in complex coordination tasks. Existing methods mainly utilize intrinsic rewards to enable committed exploration or use role-based learning for decomposing joint action spaces instead of directly conducting a collective search in the entire action-observation space. However, they often face challenges obtaining specific joint action sequences to reach successful states in long-horizon tasks. To address this limitation, we propose Imagine, Initialize, and Explore (IIE), a novel method that offers a promising solution for efficient multi-agent exploration in complex scenarios. IIE employs a transformer model to imagine how the agents reach a critical state that can influence each other's transition functions. Then, we initialize the environment at this state using a simulator before the exploration phase. We formulate the imagination as a sequence modeling problem, where the states, observations, prompts, actions, and rewards are predicted autoregressively. The prompt consists of timestep-to-go, return-to-go, influence value, and one-shot demonstration, specifying the desired state and trajectory as well as guiding the action generation. By initializing agents at the critical states, IIE significantly increases the likelihood of discovering potentially important under-explored regions. Despite its simplicity, empirical results demonstrate that our method outperforms multi-agent exploration baselines on the StarCraft Multi-Agent Challenge (SMAC) and SMACv2 environments. Particularly, IIE shows improved performance in the sparse-reward SMAC tasks and produces more effective curricula over the initialized states than other generative methods, such as CVAE-GAN and diffusion models.

Introduction

Recent progress in cooperative multi-agent reinforcement learning (MARL) has shown attractive prospects for real-world applications, such as autonomous driving (Zhou et al. 2021) and active voltage control on power distribution networks (Wang et al. 2021). To utilize global information during training and maintain scalability in execution, centralized training with decentralized execution (CTDE) has become a widely adopted paradigm in MARL. Under this

paradigm, value decomposition methods (Rashid et al. 2018; Son et al. 2019; Rashid et al. 2020) factorize the joint Q -value as a function of individual utility functions, ensuring consistency between the centralized policy and the individual policies. Consequently, they have achieved state-of-the-art performance in challenging tasks, such as StarCraft unit micromanagement (Samvelyan et al. 2019).

Despite their success, the simple ϵ -greedy exploration strategy used in these methods has been found ineffective in solving coordination tasks with complex state and reward transitions (Mahajan et al. 2019; Wang et al. 2020c; Zheng et al. 2021). To address this limitation, MAVEN (Mahajan et al. 2019) adopts a latent space for hierarchical control, allowing agents to condition their behavior on shared latent variables and enabling committed exploration. EITI and EDTI (Wang et al. 2020c) quantify and characterize the influence of one agent's behavior on others using mutual information and the difference of expected returns, respectively. By optimizing EITI or EDTI objectives as intrinsic rewards, agents are encouraged to coordinate their exploration and learn policies that optimize team performance. EMC (Zheng et al. 2021), on the other hand, uses prediction errors of individual Q -values to capture the novelty of states and the influence from other agents for coordinated exploration. Unfortunately, these methods suffer from the exponential growth of the action-observation space with the number of agents and become inefficient in long-horizon tasks.

It is possible to decompose the complex task rather than directly conducting collective searches across the entire action-observation space. To this end, RODE (Wang et al. 2020b) decomposes joint action spaces into restricted role action spaces by clustering actions based on their effects on the environment and other agents. Low-level role-based policies explore and learn within these restricted action-observation spaces, reducing execution and training complexity. However, RODE still struggles with long-term tasks as it is still unlikely to obtain a long sequence of specific actions to achieve successful states. The Go-Explore family of algorithms (Ecoffet et al. 2019; Guo et al. 2020; Ecoffet et al. 2021) decomposes the exploration into two phases: returning to the previous state and then starting to explore. These methods store the high-scoring trajectories in an archive and return to sampled states from this archive by running a goal-conditioned policy. However, in the multi-agent field, agents

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

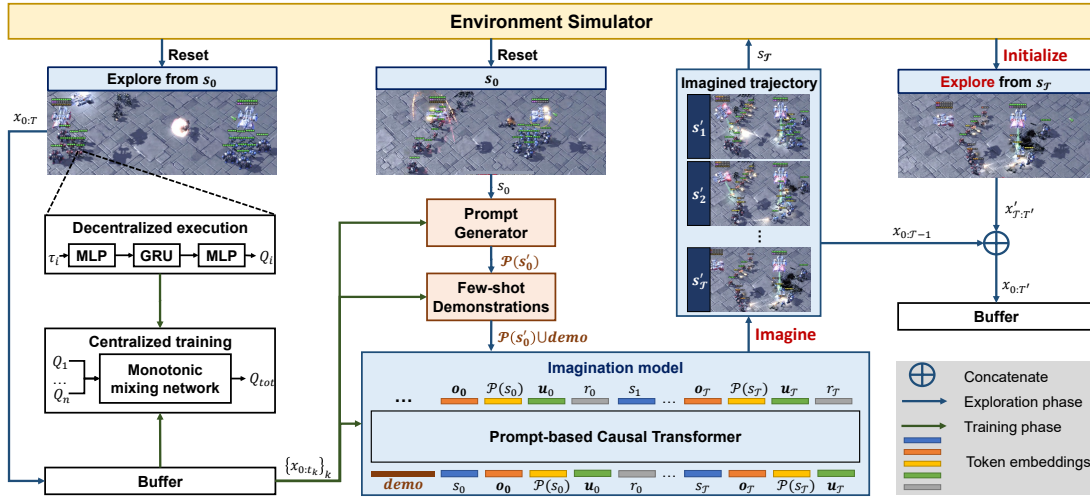


Figure 1: An overview of Imagine, Initialize, and Explore. In the pretraining phase, individual agents collect data from the initial state s_0 provided by the environment simulator. The interaction sequence is divided into several trajectory segments using influence values, which serve as the training dataset for the imagination model and few-shot demonstrations. Given s_0 , the prompt generator is trained to produce critical states, and the imagination model learns to predict how to reach such critical states from s_0 . After pretraining, the imagination model generates a trajectory from the initial state s_0 to a critical state $s_{\mathcal{T}}$ conditioned on $\mathcal{P}(s_0)$ sampled from the prompt generator, and the most related trajectory from the few-shot demonstration dataset. The agents are initialized at $s_{\mathcal{T}}$ by the environment simulator and then interact with the environment using the ϵ -greedy strategy. We concatenate the imagined and the explored trajectory to train the joint policy in the centralized training phase.

should be encouraged to teleport to the states where interactions happen and may lead to critical under-explored regions. These approaches only consider the sparse and deceptive rewards in single-agent settings, making them impractical for MARL with complex reward and transition dependencies among cooperative agents.

Teleporting agents to interaction states that influence each other’s transition function can significantly increase the likelihood of discovering potentially important yet rarely visited states and reduce the exploration space. However, there often exist multiple feasible but inefficient trajectories to reach such interaction states due to the abundance of agents’ tactics and the compositional nature of their functionalities. In light of this, we propose a novel MARL exploration method named Imagine, Initialize, and Explore (IIE). It leverages the GPT architecture (Radford et al. 2018) to imagine trajectories from the initial state to interaction states, acting as a powerful “memorization engine” that can generate diverse agent behaviors. We use target timestep-to-go, return-to-go, influence value, and a one-shot demonstration as prompts to specify the “path” of the imagined trajectory. The influence value is an advantage function that compares an agent’s current action Q -value to a counterfactual baseline that marginalizes this agent. Specifically, we obtain several trajectory segments by dividing the exploration episode, where each segment starts at the initial state from the environment and ends at an interaction state with the highest influence value. Then, the GPT model learns to predict states, observations, prompts, actions, and rewards in an autoregressive manner on these segments.

Fig. 1 presents an overview of IIE architecture. Before the

exploration phase, the agents are given a new initial state and few-shot demonstrations to construct the prompt and generate the imagined trajectory. The agents operating in partially observable environments can benefit from this imagination by utilizing it to initialize the recurrent neural network state. Then, the agents are initialized to the last state of the imagined trajectory by the simulator and interact with the environment to explore. IIE gradually provides high-influence starting points that can be viewed as auto-curricula that help agents collect crucial under-explored regions. To make the best use of the imagination, we also stitch it with the interaction sequence from exploration for policy training.

The main contributions of this paper are threefold: First, it introduces Imagine, Initialize, and Explore, which leverages the GPT architecture to imagine how the agents reach critical states before exploration. This method bridges sequence modeling and transformers with MARL instead of using GPT as a replacement for reinforcement learning algorithms (Chen et al. 2021). Second, empirical results demonstrate significant performance improvements of IIE on partially observable MARL benchmarks, including StarCraft Multi-Agent Challenge (SMAC) with dense and sparse reward settings as well as SMACv2. And third, guided by the target timestep-to-go, return-to-go, influence value, and a one-shot demonstration, the imagination model can produce more effective curricula and outperforms behavior cloning, CVAE-GAN, and classifier-guided diffusion.

Background

Decentralized Partially Observable Markov Decision Process. A fully cooperative multi-agent task in the par-

tially observable setting can be formulated as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) (Oliehoek and Amato 2016), consisting of a tuple $G = \langle A, S, \Omega, O, U, P, r, \gamma \rangle$, where $a \in A \equiv \{1, \dots, n\}$ is a set of agents, S is a set of states, and Ω is a set of joint observations. At each time step, each agent obtains its observation $o \in \Omega$ based on the observation function $O(s, a) : S \times A \rightarrow \Omega$, and an action-observation history $\tau_a \in T \equiv (\Omega \times U)^*$. Each agent a chooses an action $u_a \in U$ by a stochastic policy $\pi_a(u_a | \tau_a) : T \times U \rightarrow [0, 1]$, which forms a joint action $\mathbf{u} \in \mathbf{U}$. It results in a joint reward $r(s, \mathbf{u})$ and a transit to the next state $s' \sim P(\cdot | s, \mathbf{u})$. The formal objective function is to find the joint policy π that maximizes a joint action-value function $Q^\pi(s_t, \mathbf{u}_t) = r(s_t, \mathbf{u}_t) + \gamma \mathbb{E}_{s'} [V^\pi(s')]$, where $V^\pi(s) = \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, \pi]$, and $\gamma \in [0, 1)$ is a discounted factor.

Centralized Training with Decentralized Execution. In this paradigm, agents' policies are trained with access to global information in a centralized way and executed only based on local histories in a decentralized way (Kraemer and Banerjee 2016). One of the most significant challenges is to guarantee the consistency between the individual policies and the centralized policy, which is also known as Individual-Global Max (Son et al. 2019):

$$\arg \max_{\mathbf{u}} Q_j(s_t, \mathbf{u}) = \left\{ \begin{array}{c} \arg \max_{u^1} Q^1(s_t, u^1) \\ \dots \\ \arg \max_{u^n} Q^n(s_t, u^n) \end{array} \right\} \quad (1)$$

To address this problem, QMIX (Rashid et al. 2018) applies a state-dependent monotonic mixing network f_s to combine per-agent Q -value functions with the joint Q -value function $Q_j(s, \mathbf{u})$. The restricted space of all $Q_j(s, \mathbf{u})$ that QMIX can be represented as:

$$Q^m := \{Q_j | Q_j = f_s(Q^1(s, u^1), \dots, Q^n(s, u^n))\}, \quad (2)$$

where $Q^a(s, u^a) \in \mathbb{R}$, $\frac{\partial f_s}{\partial Q^a} \geq 0, \forall a \in A$.

Method

This section presents a novel multi-agent exploration method, Imagine, Initialize, and Explore, consisting of three key components: (1) the imagination model, which utilizes a transformer model to generate the trajectory representing how agents reach a target state autoregressively, (2) the prompt generator, which specifies the target state and trajectory for imagination, and (3) the environment simulator, which can teleport the agents to a state instantly.

Imagination Model

It has been found that agents operating in a partially observable environment can benefit from the action-observation history (Hausknecht and Stone 2015; Karkus, Hsu, and Lee 2017; Rashid et al. 2018), e.g., a model that has a recall capability such as gated recurrent unit (GRU). Therefore, it is necessary to imagine the trajectory from the initial state to the selected interaction state. We formulate the problem of trajectory generation as a sequence modeling task, where the sequences have the following form:

$$x = \{\dots, s_t, o_t^1, \dots, o_t^n, \mathcal{P}(s_t), u_t^1, \dots, u_t^n, r_t, s^{t+1}, \dots\}, \quad (3)$$

where t represents the timestep, n is the number of agents, and $\mathcal{P}(s_t)$ is the prompt for action generation at the state s_t , with the definition provided in the next subsection.

We obtain the token embeddings for states, observations, prompts, actions, and rewards through a linear layer followed by layer normalization. Moreover, an embedding for each timestep is learned and added to each token. The transformer model processes the tokens and performs autoregressive modeling to predict future tokens.

The imagination model is trained to focus on interaction states where the agents can influence each other's transition function by prioritizing trajectories with a high-influence last state in the sampled batch. Specifically, we split interaction sequences from the sampled batch into K segments, starting from the initial state by the simulator and ending at the states with top- K high-influence levels. To ensure generalization across the entire explored regions, we also enforce all state-action pairs of the sampled batch to be assigned a minimum probability of $\frac{\lambda}{N}$, where λ is a hyperparameter, and N is the number of state-action pairs in the batch. The influence \mathcal{I} is defined as an advantage function that compares the Q -value for the current action u^a to a counterfactual baseline, marginalizing out u^{a*} at a given state s :

$$\mathcal{I}(s) = \max_{a \in A} \{Q_j(s, \tau, \mathbf{u}) - \mathbb{E}_{u^{a*}} Q_j(s, \tau, (u^{a*}, u^{-a}))\}, \quad (4)$$

where $Q_j(s, \tau, \mathbf{u}) = f_s(Q^1(\tau^1, u^1; \theta), \dots, Q^n(\tau^n, u^n; \theta))$; ϕ represents the joint Q -value function, $-a$ denotes all agents A except agent a , f_s denotes the monotonic mixing network whose non-negative weights are generated by hyper-networks that take the state as input.

The imagination model parameterized by ψ is trained by:

$$L_m = \sum_{t=1}^T \left[\log q^\psi(s_t | x_{<s_t}) + \log q^\psi(r_t | x_{<r_t}) + \sum_{a=1}^n (\log q^\psi(u_t^a | x_{<u_t^a}) + \log q^\psi(o_t^a | s_t)) \right], \quad (5)$$

where T is the episode length. Since the observation is only related to the current state and the vision range of the agents, we filter out the historical memories in $x < o_{t^a}$ and use s_t as the input of the observation model $q^\psi(o_t^a | s_t)$.

Prompt For Imagination

The primary objective of our imagination model is to identify critical states and imagine how to reach them from the initial point. However, multiple feasible trajectories exist to reach these states due to the diverse tactics and compositional nature of agents' functionalities. Despite their feasibility, many of these trajectories are highly inefficient for imagination. For example, agents may wander around their initial points before engaging the enemy, decreasing success rates within a limited episode length. Therefore, a prompt, serving as a condition for action generation, is necessary to specify the target trajectory.

Given the starting state s_i , we propose a prompt generator $\mathcal{P}^\xi(s_i)$ to predict the sequence $\{s_i, \mathcal{I}_i, \mathcal{T}_i, \mathcal{R}_i\}$, where

\mathcal{I}_i is the influence value, \mathcal{T}_i is the timestep-to-go representing how quickly we can achieve the interaction state, and $\mathcal{R}_i = \sum_{t=i}^{i+\mathcal{T}} r_t$ is the return-to-go.

As the training datasets for the imagination model contain a mixture of trajectory segments, directly sampling from the prompt generator is unlikely to produce the critical state consistently. Instead, we aim to control the prompt generator to produce frequently visited but high-influence states and the trajectory leading to them. To this end, we sample the target influence value according to the log-probability:

$$\mathcal{I}_i = \log p(\mathcal{I}|s_i) + \kappa(\mathcal{I} - \mathcal{I}_{low})/(\mathcal{I}_{high} - \mathcal{I}_{low}), \quad (6)$$

where κ is a hyperparameter that controls the preference on high-influence states, \mathcal{I}_{low} and \mathcal{I}_{high} are the lower bound and upper bound of \mathcal{I}_i , respectively. The target timestep-to-go and return-to-go are obtained in the same way.

In addition, we also provide a one-shot demonstration for the imagination to avoid it being too far away from the target trajectory, also known as the ‘‘hallucination’’ problem in the large language model (McKenna et al. 2023; Manakul, Liusie, and Gales 2023). Specifically, we store the trajectory segments in a few-shot demonstration dataset and use the prompt to characterize the segments. Before imagination, we search for the trajectory whose description has the highest similarity with current prompt $\mathcal{P}(s_i)$ and then prepend it into the original input x . This process only affects the inference procedure of the model – training remains unaffected and can rely on standard next-token prediction frameworks and infrastructure.

Initialize and Explore

Before exploration, we begin by sampling a prompt $\mathcal{P}(s_0) = \{\mathcal{I}, \mathcal{T}, \mathcal{R}\}$ for imagination given the initial state s_0 from the environment simulator. Then, the agents imagine how to reach the interaction state $s_{\mathcal{T}}$ in an autoregressive manner by conditioning on the prompt and the most related trajectory from the few-shot demonstration dataset. After rolling out each imagination step and obtaining the next state, we have the next prompt through decrementing the target timestep-to-go \mathcal{T} and return-to-go \mathcal{R} by 1 and the predicted reward r_t from the imagination model, respectively. We maintain the influence value \mathcal{I} as a constant and repeat this process until the target timestep-to-go is zero. The imagined trajectory $x_{0:\mathcal{T}} = \{s_0, o_0, \mathbf{u}_0, r_0, s_1, \dots, \mathbf{u}_{\mathcal{T}}, r_{\mathcal{T}}\}$ is then used to initialize GRU state of the agent network.

Next, we define a probability α of initializing the agents at state s_0 to emphasize the importance of the target task. With the probability $1 - \alpha$, we initialize the agents at the last state $s_{\mathcal{T}}$ of the imagined trajectory using the simulator. For StarCraft II, we add a distribution over team unit types, start positions, and health points in the reset function of the environment simulator. We anneal α from 1.0 to 0.5 over a fixed number of steps after the pretraining phase of the imagination model. Finally, the agents interact with the environment to collect the online data $x_{\mathcal{T}:T} = \{s_{\mathcal{T}}, o_{\mathcal{T}}, \mathbf{u}_{\mathcal{T}}, r_{\mathcal{T}}, s_{\mathcal{T}+1}, \dots, \mathbf{u}_T, r_T\}$. We stitch the imagined trajectory and the online data $\mathcal{X} = x_{0:\mathcal{T}-1} \oplus x_{\mathcal{T}:T}$ as training datasets for policy training.

The individual Q -value functions $Q^a(o^a, \tau^a, u^a)$ are optimized jointly to minimize the following loss:

$$\min_{\theta, \phi} \sum_{t=0}^{T-1} [Q_j(s_t, \tau_t, \mathbf{u}_t; \theta, \phi) - y(s_t, \tau_t, \mathbf{u}_t)]^2, \quad (7)$$

where $y(s_t, \tau_t, \mathbf{u}_t) = r_t + \gamma \max_{\mathbf{u}} Q'_j(s_{t+1}, \tau_{t+1}, \mathbf{u})$ is the target value function, and Q'_j is the target network whose parameters are periodically copied from Q_j .

Related Work

Multi-agent Exploration. Individual exploration suffers from the inconsistency between local and global information as well as the non-stationary problem in multi-agent settings. To address these limitations, Jaques et al. (2019) introduce intrinsic rewards based on ‘‘social influence’’ to incentivize agents in selecting actions that can influence other agents. Similarly, Wang et al. (2020c) leverage mutual information to capture the interdependencies of rewards and transitions, promoting exploration efficiency and facilitating the policies training. EMC (Zheng et al. 2021) leverages prediction errors from individual Q -values as intrinsic rewards and uses episodic memory to improve coordinated exploration. MAVEN (Mahajan et al. 2019) employs a hierarchical policy for committed and temporally extended exploration, learning multiple state-action value functions for each agent through a shared latent variable. RODE (Wang et al. 2020b) introduces a role selector to enable informed decisions and learn role-based policies in a smaller action space based on the actions’ effects. However, learning in long-horizon coordination tasks remains challenging due to the exponential state-action spaces.

Go-Explore. Go-Explore (Ecoffet et al. 2019) is one of the most famous exploration approaches, particularly well-suited for hard-exploration domains with sparse or deceptive rewards. In the Go-Explore family of algorithms (Ecoffet et al. 2021; Guo et al. 2020), an agent returns to a promising state without exploration by running a goal-conditioned policy or restoring the simulator state and then explores from this state. However, these methods are primarily designed for single-agent scenarios, neglecting the interdependencies between agent policies and the strategies of others in MARL. Moreover, these methods require restoring previously visited simulator states or training a goal-conditioned policy to generate actions for returning. This leads to significant computational costs, especially in coordination scenarios involving complex observations and transitions.

Transformer Model. Several works have explored the integration of transformer models into reinforcement learning (RL) settings. We classify them into two major categories depending on the usage pattern. The first category focuses on representing components in RL algorithms, such as policies and value functions (Parisotto et al. 2020; Parisotto and Salakhutdinov 2021). These methods rely on standard RL algorithms to update policy, where the transformer only provides the large representation capacity and improves feature extraction. Conversely, the second category aims to replace

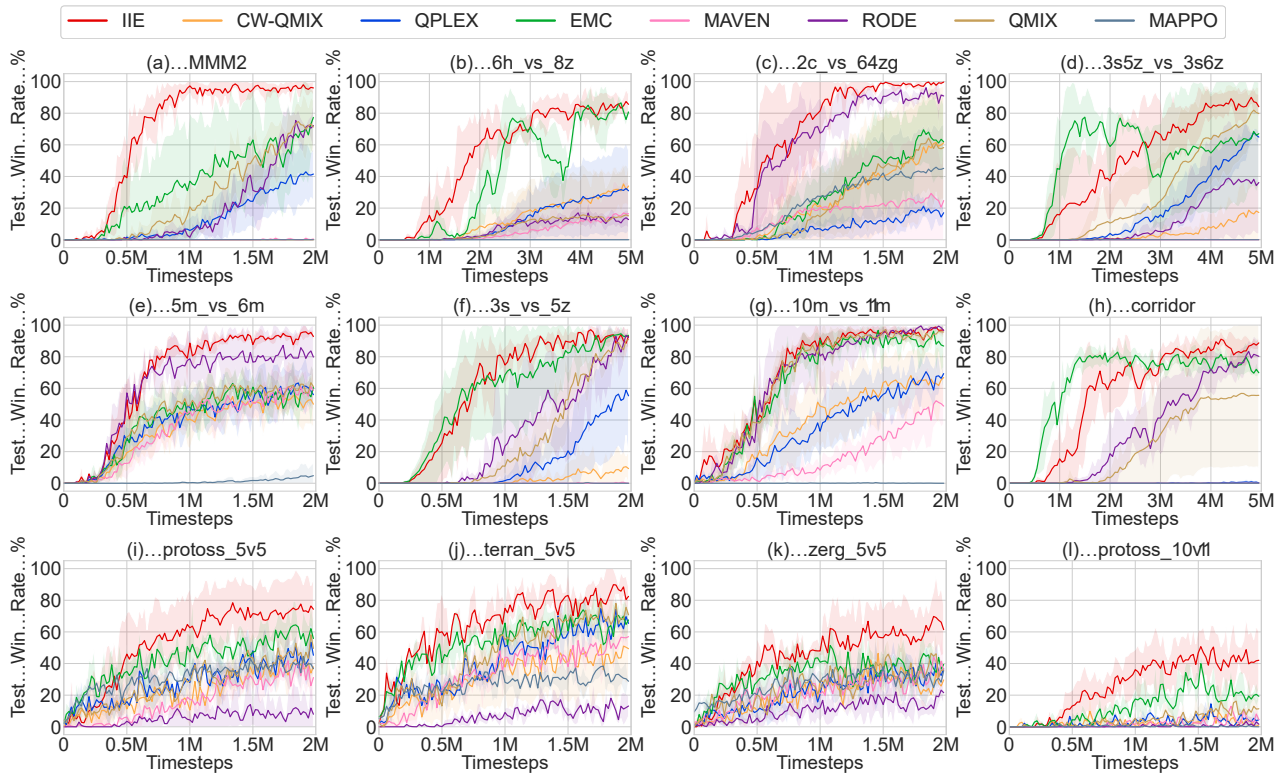


Figure 2: Performance comparisons on the dense-reward SMAC and SMACv2 benchmarks.

the RL pipeline with sequence modeling. They autoregressively generate states, actions, and rewards by conditioning on the desired return-to-go during inference (Chen et al. 2021; Lee et al. 2022; Reed et al. 2022). However, there is a risk of unintended behaviors when offline datasets contain destructive biases. Moreover, it remains an open question of how to extend these supervised learning paradigms to online settings and satisfy the scalability in practice.

We utilize a transformer model to imagine the trajectory to reach the interaction state guided by a prompt, offering increased representational capacity and stability compared to existing Go-Explore methods. In contrast to current multi-agent exploration methods, we initialize agents at the states with high influence values. This form of curriculum learning significantly reduces the exploration space and enhances coordination exploration. We aim to bridge sequence modeling and transformers with MARL rather than replacing conventional RL algorithms.

Results

In this section, we conduct empirical experiments to answer the following questions: (1) Is Imagine, Initialize, and Explore (IIE) better than the existing MARL exploration methods in complex cooperative scenarios or sparse reward settings? (2) Can IIE generate a reasonable curriculum of the last states over timesteps and outperform other returning methods? We also investigate the contribution of each component in the proposed prompt to the imagination model.

We conduct experiments on NVIDIA RTX 3090 GPUs.

Each task needs to train for about 12 to 20 hours, depending on the number of agents and the episode length limit. We evaluate 32 episodes with decentralized greedy action selection every $10k$ timesteps for each algorithm. All figures are plotted using mean and standard deviation with confidence interval 95%. We conduct five independent runs with different random seeds for each learning curve.

Performance Comparison

In our evaluation, we compare the performance of CW-QMIX (Rashid et al. 2020), QPLEX (Wang et al. 2020a), MAVEN (Mahajan et al. 2019), EMC (Zheng et al. 2021), RODE (Wang et al. 2020b), QMIX (Rashid et al. 2018), MAPPO (Yu et al. 2022), and IIE on the StarCraft Multi-Agent Challenge (SMAC) (Samvelyan et al. 2019) and SMACv2 (Ellis et al. 2022) benchmarks.

SMAC is a partially observable MARL benchmark known for its rich environments and high control complexity. It requires learning policies in a large observation-action space, where agents take various actions, such as “move” in cardinal directions, “stop”, and selecting an enemy to attack. The maximum number of actions and the episode length vary across different scenarios, ranging from 7 to 70 actions and 60 to 400 timesteps, respectively. In contrast, SMACv2 presents additional challenges of stochasticity and generalization. It includes procedurally generated scenarios with start positions, unit types, attack range, and sight range. The agents must generalize their learned policies to previously unseen settings during testing.

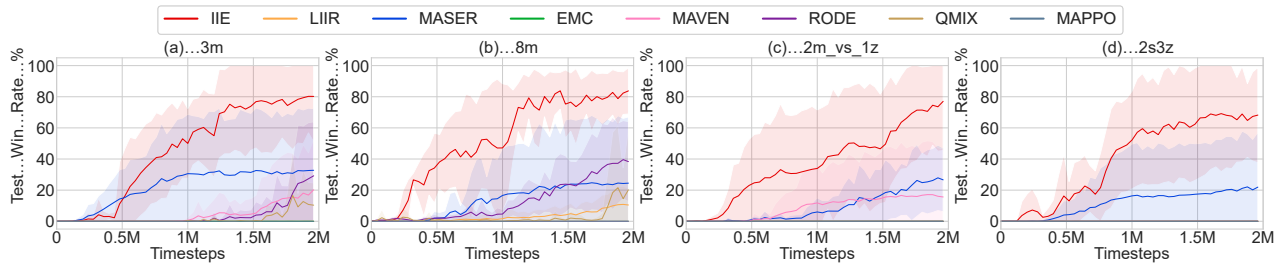


Figure 3: Performance comparisons on the sparse-reward SMAC benchmark.

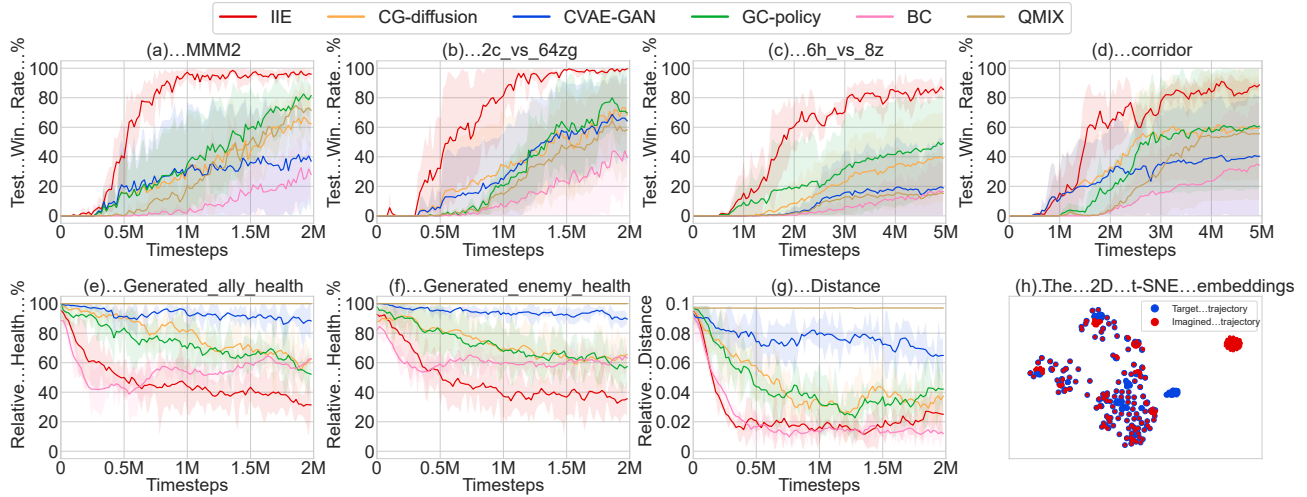


Figure 4: (a-d) Performance comparisons with different returning methods on the SMAC benchmark. (e-g) The mean health of allies and enemies, as well as the relative distance between two groups at the last state in the MMM2 scenario. (h) The 2D t-SNE embeddings of the trajectory returned from IIE in the MMM2 scenario after pretraining.

Fig. 2 shows that IIE considerably outperforms the state-of-the-art MARL methods in both SMAC and SMACv2 maps with the dense reward setting. This result demonstrates that IIE significantly enhances learning speed and coordination performance in complex tasks. In specific scenarios like *3s5z_vs_3s6z* and *corridor*, EMC shows faster learning in the beginning, which can be attributed to the fact that IIE requires more time to pre-train the imagination model in more complex tasks. However, as the training progresses, IIE excels in providing a more targeted and efficient exploration of complex coordination scenarios, leading to the best final performance across all scenarios.

Sparse-reward Benchmark

We investigate the performance of IIE, EMC, MAVEN, RODE, QMIX, and MAPPO, in addition to two multi-agent exploration methods designed for sparse rewards, including LIIR (Du et al. 2019) and MASER (Jeon et al. 2022), on the SMAC benchmark with the sparse-reward setting. In this setting, global rewards are sparsely given only when one or all enemies are defeated, with no additional reward for state information such as enemy and ally health. This problem is difficult for credit assignments because credit must be propagated from the end of the episode.

The imagination model in IIE can improve robustness in these settings because it makes minimal assumptions on the density of the reward. As shown in Fig. 3, sparse and delayed rewards minimally affect IIE. We hypothesize that the transformer architecture in the imagination model can be effective critics and enable more accurate value prediction. In contrast, QMIX and MAPPO fail to solve these tasks since they heavily rely on densely populated rewards for calculating temporal difference targets or generalized advantage estimates. LIIR, MASER, RODE, and MAVEN exhibit slow learning and instability, particularly on the heterogeneous map *2s3z*, suggesting the difficulty of learning such intrinsic rewards, subgoals, roles, or noise-based hierarchical policies in hard-exploration scenarios.

Different Returning Methods

In this section, we seek insight into whether the imagination model in IIE can be thought of as performing efficient curriculum learning and is better than other returning methods. To investigate this, we compare it with behavior cloning (BC), a goal-conditioned policy method (GC-policy) (Ecoffet et al. 2021), and generative models, including CVAE-GAN (Bao et al. 2017) and classifier-guided diffusion (CG-diffusion) (Ajay et al. 2023). Before exploration, we initial-

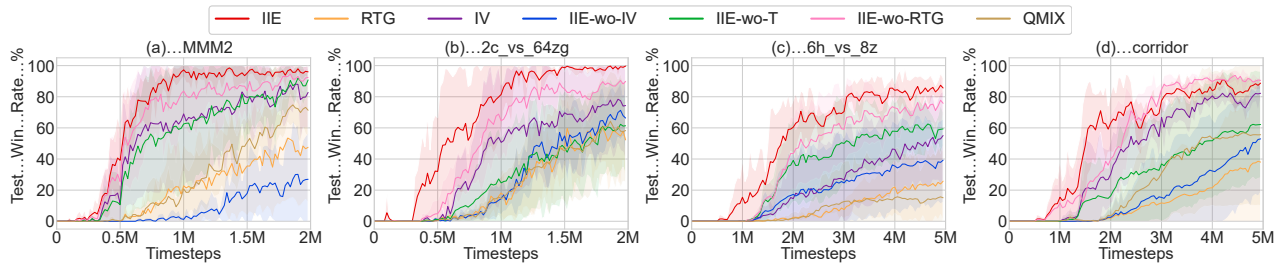


Figure 5: IIE with different prompts on the SMAC benchmark.

ize the agents to the last state of the imagination trajectory produced by BC, CVAE-GAN, and CG-diffusion using the environment simulator or run a goal-conditioned policy from GC-policy to return agents to the target state.

We show the performance comparison and visualization results in Fig. 4. IIE provides agents with interaction points for exploration with the shortest relative distance between agents and enemies and decreasing but comparable health levels as training progresses. This contributes to early and more frequent multi-agent interactions against the enemies quickly, reducing the complexity of the multi-agent exploration space. As a result, IIE outperforms BC, GC-policy, CVAE-GAN, and CG-diffusion across all tasks, indicating that the prompt-based imagination can be more effective and has better generalization than simply performing imitation learning or other generative models. We also illustrate the data distribution of imagined trajectories in a 2D space with dimensional reduction via T-SNE. The results show that the imagined trajectory will be close to the target trajectory, which verifies that IIE can capture and learn the dynamics of the multi-agent environment based on the prompt.

CG-diffusion and GC-policy have plateaued in performance, showing a similar trend of emergent complexity with IIE but taking far longer to learn the joint policy. On the one hand, the continuous diffusion models have been extremely successful in vision and audio domains, but they do not perform well in text because of the inherently discrete nature of text (Li et al. 2022). The imagination of the trajectory is more related to text than image generation because it has low redundancy, and the transition between two connected states is essential. On the other hand, the applicability of the current goal-conditioned methods with complex observations is limited, as their flexibility is often constrained to task spaces using low-dimensional parameters and requires well-defined task similarity. CVAE-GAN does not show positive results and keeps generating similar health levels. We hypothesize that the mode collapse problem is a bottleneck for CVAE-GAN. The generator can find data that can easily fool the discriminator because of the unbalanced training data from the exploration. BC shows the worst results because it imitates all past interaction sequences and lacks the generalization ability to avoid sub-optimal solutions. From Fig. 4e-g, we can see that BC prioritizes states with short distances but does not provide reasonable health levels for enemies - the mean health level of enemies far exceeds that of allies, making it difficult or even impossible to achieve any success.

Different Prompts for Imagination

In this section, we conduct ablation studies to analyze the contributions of each component in the prompts, including the timesteps, the accumulated return-to-go (RTG), the constant influence value (IV), and the most related trajectory (T). We integrate the desired timesteps into the following prompts: (1) RTG, (2) IV, (3) RTG with the influence value, denoted as IIE-wo-T, (4) RTG with the most related trajectory, denoted as IIE-wo-IV, (5) IV with the most related trajectory, denoted as IIE-wo-RTG. We compare them with IIE and QMIX on the SMAC benchmark.

Fig. 5 shows that RTG and IIE-wo-IV achieve poor performance in `MMM2`, `6h_vs_8z`, and `corridor` due to their ignorance of the importance of interactions in multi-agent exploration and limited positive samples for value decomposition. IV and IIE-wo-T perform worse than IIE-wo-IV and IIE with a considerable gap, as they do not exploit one-shot demonstration to guide action generation, leading to a potential risk that the imagined trajectories may deviate significantly from the target trajectory, especially in long-horizon tasks with complex transition functions. IIE outperforms IIE-wo-IV across `MMM2`, `6h_vs_8z`, and `2c_vs_64zg` maps because RTG can further specify a trajectory, improving the efficiency and robustness of the imagination learning. However, IIE-wo-IV performs better than IIE in `corridor`, implying that the imagination model may have some degree of generalization without RTG and provide more diverse samples for policy training.

Conclusion

We proposed Imagine, Initialize, and Explore, which enables us to break down a difficult multi-agent exploration problem into a curriculum of subtasks created by initializing agents at the interaction state - the last state in the imagined trajectory. We empirically evaluated our algorithm and found it outperforms current multi-agent exploration methods and generative models on various benchmarks. We also show that the prompt-based imagination model performs efficient conditional sequence generation and has one-shot generalization. We hope this work will inspire more investigation of sequence-prediction models' applications in multi-agent reinforcement learning (MARL) rather than using them to replace conventional MARL. In future work, we consider learning continuous prompts to specify the imagination, as opposed to the simple influence value used in this paper.

Acknowledgments

This work was supported in part by National Key R&D Program of China under grant No. 2021ZD0112700, NSFC under grant No. 62125305, No. 62088102, No. 61973246, No. 62203348, and No. U23A20339, the Fundamental Research Funds for the Central Universities under Grant xtr072022001.

References

- Ajay, A.; Du, Y.; Gupta, A.; Tenenbaum, J. B.; Jaakkola, T. S.; and Agrawal, P. 2023. Is Conditional Generative Modeling all you need for Decision Making? In *The Eleventh International Conference on Learning Representations*.
- Bao, J.; Chen, D.; Wen, F.; Li, H.; and Hua, G. 2017. CVAE-GAN: fine-grained image generation through asymmetric training. In *Proceedings of the IEEE international conference on computer vision*, 2745–2754.
- Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.
- Du, Y.; Han, L.; Fang, M.; Liu, J.; Dai, T.; and Tao, D. 2019. Liir: Learning individual intrinsic reward in multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 32.
- Ecoffet, A.; Huizinga, J.; Lehman, J.; Stanley, K. O.; and Clune, J. 2019. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*.
- Ecoffet, A.; Huizinga, J.; Lehman, J.; Stanley, K. O.; and Clune, J. 2021. First return, then explore. *Nature*, 590(7847): 580–586.
- Ellis, B.; Moalla, S.; Samvelyan, M.; Sun, M.; Mahajan, A.; Foerster, J. N.; and Whiteson, S. 2022. SMACv2: An Improved Benchmark for Cooperative Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2212.07489*.
- Guo, Y.; Choi, J.; Moczulski, M.; Feng, S.; Bengio, S.; Norouzi, M.; and Lee, H. 2020. Memory based trajectory-conditioned policies for learning from sparse rewards. *Advances in Neural Information Processing Systems*, 33: 4333–4345.
- Hausknecht, M.; and Stone, P. 2015. Deep recurrent q-learning for partially observable mdps. In *2015 aaai fall symposium series*.
- Jaques, N.; Lazaridou, A.; Hughes, E.; Gulcehre, C.; Ortega, P.; Strouse, D.; Leibo, J. Z.; and De Freitas, N. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International conference on machine learning*, 3040–3049. PMLR.
- Jeon, J.; Kim, W.; Jung, W.; and Sung, Y. 2022. MASER: Multi-Agent Reinforcement Learning with Subgoals Generated from Experience Replay Buffer. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 10041–10052. PMLR.
- Karkus, P.; Hsu, D.; and Lee, W. S. 2017. Qmdp-net: Deep learning for planning under partial observability. *Advances in neural information processing systems*, 30.
- Kraemer, L.; and Banerjee, B. 2016. Multi-agent reinforcement learning as a rehearsal for decentralized planning. *Neurocomputing*, 190: 82–94.
- Lee, K.-H.; Nachum, O.; Yang, M. S.; Lee, L.; Freeman, D.; Guadarrama, S.; Fischer, I.; Xu, W.; Jang, E.; Michalewski, H.; et al. 2022. Multi-game decision transformers. *Advances in Neural Information Processing Systems*, 35: 27921–27936.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35: 4328–4343.
- Mahajan, A.; Rashid, T.; Samvelyan, M.; and Whiteson, S. 2019. Maven: Multi-agent variational exploration. *Advances in Neural Information Processing Systems*, 32.
- Manakul, P.; Liusie, A.; and Gales, M. J. 2023. Self-checkgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- McKenna, N.; Li, T.; Cheng, L.; Hosseini, M. J.; Johnson, M.; and Steedman, M. 2023. Sources of Hallucination by Large Language Models on Inference Tasks. *arXiv preprint arXiv:2305.14552*.
- Oliehoek, F. A.; and Amato, C. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- Parisotto, E.; and Salakhutdinov, R. 2021. Efficient Transformers in Reinforcement Learning using Actor-Learner Distillation. In *International Conference on Learning Representations*.
- Parisotto, E.; Song, F.; Rae, J.; Pascanu, R.; Gulcehre, C.; Jayakumar, S.; Jaderberg, M.; Kaufman, R. L.; Clark, A.; Noury, S.; et al. 2020. Stabilizing transformers for reinforcement learning. In *International conference on machine learning*, 7487–7498. PMLR.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training. In *OpenAI*. OpenAI.
- Rashid, T.; Farquhar, G.; Peng, B.; and Whiteson, S. 2020. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 33: 10199–10210.
- 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.
- Reed, S.; Zolna, K.; Parisotto, E.; Colmenarejo, S. G.; Novikov, A.; Barth-Maron, G.; Gimenez, M.; Sulsky, Y.; Kay, J.; Springenberg, J. T.; et al. 2022. A generalist agent. *arXiv preprint arXiv:2205.06175*.

Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*.

Son, K.; Kim, D.; Kang, W. J.; Hostallero, D. E.; and Yi, Y. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, 5887–5896. PMLR.

Wang, J.; Ren, Z.; Liu, T.; Yu, Y.; and Zhang, C. 2020a. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *International Conference on Learning Representations*.

Wang, J.; Xu, W.; Gu, Y.; Song, W.; and Green, T. C. 2021. Multi-agent reinforcement learning for active voltage control on power distribution networks. *Advances in Neural Information Processing Systems*, 34: 3271–3284.

Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2020b. Rode: Learning roles to decompose multi-agent tasks. *arXiv preprint arXiv:2010.01523*.

Wang, T.; Wang, J.; Wu, Y.; and Zhang, C. 2020c. Influence-Based Multi-Agent Exploration. In *International Conference on Learning Representations*.

Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.

Zheng, L.; Chen, J.; Wang, J.; He, J.; Hu, Y.; Chen, Y.; Fan, C.; Gao, Y.; and Zhang, C. 2021. Episodic Multi-agent Reinforcement Learning with Curiosity-driven Exploration. *Advances in Neural Information Processing Systems*, 34.

Zhou, M.; Luo, J.; Vilella, J.; Yang, Y.; Rusu, D.; Miao, J.; Zhang, W.; Alban, M.; FADAKAR, I.; Chen, Z.; et al. 2021. SMARTS: An Open-Source Scalable Multi-Agent RL Training School for Autonomous Driving. In *Conference on Robot Learning*, 264–285. PMLR.