



***SUB-PLAY*: Adversarial Policies against Partially Observed Multi-Agent Reinforcement Learning Systems**

Oubo Ma
Zhejiang University
Hangzhou, China
mob@zju.edu.cn

Yuwen Pu*
Zhejiang University
Hangzhou, China
yw.pu@zju.edu.cn

Linkang Du
Xi'an Jiaotong University
Xi'an, China
linkangd@gmail.com

Yang Dai
Laboratory for Big Data and Decision
Changsha, China
daiyang2000@163.com

Ruo Wang
Chinese Aeronautical Establishment
Beijing, China
kurt_ashtray@163.com

Xiaolei Liu
Institute of Computer Application,
China Academy of Engineering
Physics
Mianyang, China
luxaole@gmail.com

Yingcai Wu
Zhejiang University
Hangzhou, China
ycwu@zju.edu.cn

Shouling Ji*
Zhejiang University
Hangzhou, China
sjj@zju.edu.cn

Abstract

Recent advancements in multi-agent reinforcement learning (MARL) have opened up vast application prospects, such as swarm control of drones, collaborative manipulation by robotic arms, and multi-target encirclement. However, potential security threats during the MARL deployment need more attention and thorough investigation. Recent research reveals that attackers can rapidly exploit the victim's vulnerabilities, generating adversarial policies that result in the failure of specific tasks. For instance, reducing the winning rate of a superhuman-level Go AI to around 20%. Existing studies predominantly focus on two-player competitive environments, assuming attackers possess complete global state observation.

In this study, we unveil, for the first time, the capability of attackers to generate adversarial policies even when restricted to partial observations of the victims in multi-agent competitive environments. Specifically, we propose a novel black-box attack (*SUB-PLAY*) that incorporates the concept of constructing multiple subgames to mitigate the impact of partial observability and suggests sharing transitions among subpolicies to improve attackers' exploitative ability. Extensive evaluations demonstrate the effectiveness of *SUB-PLAY* under three typical partial observability limitations. Visualization results indicate that adversarial policies induce significantly different activations of the victims' policy networks. Furthermore, we evaluate three potential defenses aimed at exploring ways to

mitigate security threats posed by adversarial policies, providing constructive recommendations for deploying MARL in competitive environments.

CCS Concepts

• **Security and privacy**; • **Computing methodologies** → **Artificial intelligence**;

Keywords

Adversarial Policy; Multi-Agent Reinforcement Learning; Partially Observable

ACM Reference Format:

Oubo Ma, Yuwen Pu, Linkang Du, Yang Dai, Ruo Wang, Xiaolei Liu, Yingcai Wu, and Shouling Ji. 2024. *SUB-PLAY*: Adversarial Policies against Partially Observed Multi-Agent Reinforcement Learning Systems. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*, October 14–18, 2024, Salt Lake City, UT, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3658644.3670293>

1 Introduction

Multi-agent reinforcement learning (MARL) has succeeded remarkably in diverse domains, from StarCraft II [49] to cyber-physical systems [58], strategic maneuvers [11], and social science [35]. Currently, MARL predominantly emphasizes improving algorithm performance across various tasks, yet there is a noticeable lack of consideration for security aspects.

Recent research [59] has unveiled that even state-of-the-art reinforcement learning (RL) policies exhibit weaknesses and vulnerabilities in competitive environments. Therefore, an attacker can employ adversarial policies to induce the victim's policies to trigger vulnerabilities, resulting in a significant performance decline, possibly even leading to a loss of sequential decision-making capability.

For instance, *Victim-play* [15] is a black-box attack framework designed for adversarial policy generation, where the attacker can

*Yuwen Pu and Shouling Ji are the co-corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '24, October 14–18, 2024, Salt Lake City, UT, USA.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0636-3/24/10

<https://doi.org/10.1145/3658644.3670293>



Figure 1: Three partially observable limitations in multi-agent environments.

interact with the victim without requiring access to the victim’s policy or environmental perturbations. However, *Victim-play* is designed for two-player competitions, where the victim operates as a single-agent system, such as a superhuman-level Go AI [59]. It encounters challenges when striving to sustain a stable attack performance within multi-agent competitions, with the victim functioning as a multi-agent system (MAS), for instance, a drone swarm [71]. The widespread prevalence of partial observability exacerbates this challenge, as attackers are unable to access complete global state information. This may result in adversarial strategies being unable to converge due to fluctuations or getting trapped in poorly performing local optima.

Partial observability is primarily attributed to three limitations (see Figure 1): (1) Uncertainty Limitation: This occurs when partial observability arises due to the constraints imposed by unpredictable environmental events. Examples include obstacle occlusion, noisy measurements, and sensor anomalies. (2) Distance Limitation: This refers to situations where the relative distance between agents exceeds their perceptual range, determined by the sensors deployed by MASs, such as LiDARs and millimeter-wave radars [50]. (3) Region Limitation: Incomplete observations result from privacy concerns, security constraints, or rule restrictions, where specific boundaries define the region and could represent geographical areas or logical ranges. Examples in this regard encompass restricted areas due to permission controls or competitions with incomplete information, such as financial markets or Texas Hold’em poker.

This paper introduces *SUB-PLAY*, a novel black-box attack framework aimed at adversarial policy generation in partially observed multi-agent competitive environments. Our intuition lies in the divide-and-conquer principle, decomposing the attack into multiple subgames. Each subgame is then modeled as a partially observable stochastic game (POSG) [37], and MARL is employed to solve and obtain the corresponding subpolicy. Finally, we integrate all subpolicies in a hard-coded format to generate the ultimate adversarial policy. Our main challenge is the ineffectiveness of attacks caused by data imbalance. Specifically, the attacker records the interactions at each time step in the form of transitions and allocates these transitions to a specific subgame replay buffer based on the observed number of victim agents. However, the number of transitions in each buffer is uneven due to varying probabilities of each subgame occurrence. This imbalance may lead to undertraining of some subpolicies. To mitigate this issue, we propose a transition dissemination mechanism that facilitates the sharing of transition from proximity subgames.

Extensive evaluation results demonstrate that *SUB-PLAY* can effectively address the aforementioned three partial observability limitations and outperform *Victim-play* in Predator-prey and World Communication, two representative multi-agent competitive environments open-sourced by OpenAI [40]. Compared to normal opponents, t-SNE analysis reveals a significant difference in the activations of the victim’s policy network during interactions with adversarial policies. The scalability evaluation indicates that attackers can adjust the granularity of subgame construction to expand the applicability of our method. Moreover, *SUB-PLAY* is algorithm-agnostic, *i.e.*, applicable to both distributed and centralized MARL algorithms.

To explore strategies for mitigating adversarial policies, we evaluate three potential defenses. The results indicate that adversarial retraining is insufficient to counteract adversarial policies, while policy ensemble and fine-tuning could only moderately reduce the effectiveness of attacks. Nevertheless, insights from the evaluation suggest that defenders may mitigate security risks posed by adversarial policies through flexible deployment techniques. For example, periodically updating policies or increasing the diversity of policies in policy ensemble.

In summary, the paper makes the following contributions:

- To the best of our knowledge, *SUB-PLAY*¹ is the first work to investigate the security threats of adversarial policies in multi-agent competitive environments, revealing that attackers can exploit vulnerabilities in the victim’s policy even with partial observations.
- We summarize three partially observable limitations and propose an observable-driven subgame construction method to accommodate these limitations.
- We conduct a systematic evaluation, demonstrating that *SUB-PLAY* outperforms the state-of-the-art attack framework in partially observable multi-agent competitive environments.
- We explore potential defenses, emphasizing that practitioners in MARL should not only focus on improving algorithm performance but also pay attention to deployment details, which is crucial in mitigating security threats posed by adversarial policies.

2 Background

2.1 Multi-Agent RL

MARL refers to scenarios where multiple agents are involved in sequential decision-making, and their policies are updated concurrently. This results in a non-stationary environment where the optimal policy for each agent changes over time, making the Markov property invalid [54].

MARL Tasks. Based on the cooperation patterns among agents, MARL tasks are primarily divided into four categories: (1) Fully Cooperative MARL, where agents typically share a common reward function and collaborate to achieve a shared objective. Examples include multi-agent pathfinding and traffic management. (2) Fully Competitive MARL, where agents compete individually to outperform each other and pursue their objectives. These tasks are often

¹Inspired by the previous work [25], the source code of *SUB-PLAY* is responsibly shared with other researchers to avoid potential ethical concerns. Instructions regarding access requests can be found at <https://github.com/maoubo/Repository-Access-SUB-PLAY>.

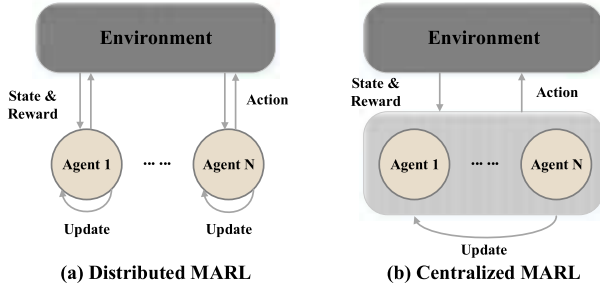


Figure 2: MARL training paradigms.

modeled as two-player zero-sum Markov games, where cooperation between agents is impossible. Examples include Go or arm wrestling. (3) Self-Interested MARL, where agents prioritize their benefits without considering others, as observed in domains like autonomous driving and stock trading. (4) Mixed MARL involves a blend of cooperative and competitive behavior. In most scenarios, two competing MASs exist, but agents within the same MAS collaborate. Examples include military exercises, multi-target encirclement, and team sports. The competitive environment discussed in this paper falls within a typical class of mixed MARL tasks.

Training Paradigm. MARL has two training paradigms based on the presence or absence of a central decision-maker [19]. (1) Distributed MARL: The algorithms assume agents update policies independently, similar to single-agent implementations. *Distributed Training Decentralized Execution (DTDE)*, a typical paradigm (e.g., independent Q-learning), allows efficient training and deployment without communication constraints. However, it may not be suitable for complex environments with many agents, as non-stationarity is neglected. (2) Centralized MARL: *Centralized Training Centralized Execution (CTCE)* is a centralized MARL paradigm with a centralized decision maker in training and deployment. It achieves theoretically optimal performance but needs communication guarantees and suffers from the curse of dimensionality [46]. In contrast, *Centralized Training Decentralized Execution (CTDE)* algorithms, such as QMIX, MADDPG, and MAPPO, guide agents during training but enable independent decision-making during deployment without additional communication, offering state-of-the-art performance. *SUB-PLAY* is not limited by the attacker’s communication capabilities or the number of agents under its control. Therefore, *SUB-PLAY* applies to both distributed and centralized MARL algorithms.

2.2 Adversarial Policy

The adversarial policy is a form of action manipulation attack in which the attacker induces the black-box victim to make suboptimal decisions by controlling the actions of the adversary agents. The training of adversarial policies relies on the competitive relationship between the attacker and the victim (typically zero-sum games), so the attacker only needs to maximize the adversary agents’s reward to autonomously discover and exploit weaknesses and vulnerabilities in the RL policies deployed by the victims.

Existing research [15, 22, 38, 59, 60] predominantly focuses on two-player competitions, assuming an attacker has the privilege

to interact with a victim, can obtain a complete observation of the environment at each time step, and the victim fixedly deploys a well-trained policy. The fundamental reason for the existence of adversarial policies stems from RL’s adoption of *Self-play* for policy training in competitive environments [26]. However, *Self-play* cannot guarantee to reach a Nash equilibrium within finite training. In game theory, non-equilibrium policies are inevitably exploitable. Guided by this intuition, an attacker can manipulate its policy during training, updating it in a direction that maximizes the exploitation of the victim’s vulnerabilities.

For instance, while AlphaGo-style AIs outperform human champions, adversarial policies specifically trained against them still achieve a success rate of over 77% [59]. Remarkably, these adversarial policies fall short when facing ordinary Go enthusiasts. This indicates that adversarial policies are highly targeted, sacrificing generalizability to intensify exploitation against a specific victim. Therefore, an adversarial policy is a complement to RL or a figuration of its weaknesses rather than a substitute.

Finding or approximating a Nash equilibrium in a multi-agent competition is at least as hard as PPAD-complete [4]. This implies that in real deployment scenarios, MARL policies are exploitable. Therefore, MARL must be attentive to the potential risks posed by adversarial policies.

More details about the adversarial policy, including a more nuanced explanation of its existence and the upper limit of attack performance, can be found in Appendix A of [42]².

3 Threat Model and Problem Formulation

3.1 Threat Model

In this paper, we propose an adversarial policy attack for mixed MARL tasks in two-team competitive environments³.

Definition 1. A two-team competitive environment involves two MASs, *Adversary* and *Victim*, which consist of two sets of agents, \mathcal{M} and \mathcal{N} , where $|\mathcal{M}| = M$ and $|\mathcal{N}| = N$. *Adversary* and *Victim* are in full competition, while the agents within each MAS collaborate.

Attacker’s Goal. Maximizing the reduction of *Victim*’s performance on a specific MARL task.

Attacker’s Capabilities. The attacker possesses complete control over the *Adversary* and can update its MARL policy. Additionally, the attacker has interaction privileges with the *Victim*, obtaining partial observations about the environment at each time step. The attacker knows that there are N agents in the *Victim*. Apart from this, *Victim* is regarded as a black box by the attacker. Moreover, the attacker is restricted from manipulating the environment.

We assume that *Victim*’s MARL policy is fixed, i.e., the parameters during the deployment phase are frozen (subsequent evaluations indicate that this assumption can be relaxed). This is common for the deployment of RL on physical entities. For instance, a manufacturer releases an encirclement MAS consisting of multiple-legged robots [13] or drones [31]. These physical entities fixedly deploy policies to avoid unacceptable losses due to exploration, such as

²We provide a long version with appendices on arXiv: <https://arxiv.org/pdf/2402.03741>.

³“Two-team competitive environments” is a subset of “multi-agent competitive environments”. For clarity and readability, we use the first term in sections related to problem formulation and scheme design (Section 3 and Section 4) and the second term in all other sections.

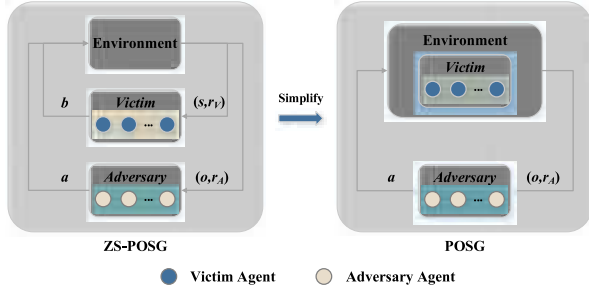


Figure 3: A two-team competitive environment can be simplified from a ZS-POSG to a POSG if the joint policy of *Victim* is fixed.

robot malfunctions or drone crashes. Even if the manufacturer offers periodic policy upgrade services, the parameters of the MAS remain fixed between two consecutive updates, and this interval might span several months or even years. In such scenarios, the manufacturer is a potential victim. Once the attacker generates an adversarial policy, it implies that all users deploying the MAS from that manufacturer are exposed to the threat.

3.2 Problem Formulation

Based on the threat model, we formulate the attack as a zero-sum partially observable stochastic game (ZS-POSG) [37, 69].

Definition 2. A zero-sum partially observable stochastic game is defined as

$$\mathcal{G} = (\mathcal{S}, \{\mathcal{A}^i\}, \{\mathcal{B}^j\}, \{\Omega^i\}, \{O^i\}, \mathcal{P}, \{\mathcal{R}_A^i\}, \{\mathcal{R}_V^j\}, \gamma), \quad (1)$$

where \mathcal{S} denotes the state space, \mathcal{A}^i (resp. \mathcal{B}^j) denotes the action space for agent $i \in \mathcal{M}$ (resp. $j \in \mathcal{N}$). $\mathcal{A} = \prod_{i \in \mathcal{M}} \mathcal{A}^i$ and $\mathcal{B} = \prod_{j \in \mathcal{N}} \mathcal{B}^j$ denote the joint action spaces for the adversary and the victim. Each agent $i \in \mathcal{M}$ receives an observation $o^i \in \Omega^i$, and the observation function $O^i : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \Delta(\Omega^i)$ is a probability distribution over possible subsequent observations given the previous state and the actions of all agents. Each agent $j \in \mathcal{N}$ receives the global state $s \in \mathcal{S}$. $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \Delta(\mathcal{S})$ represents the probability that taking joint action $a \in \mathcal{A}$ and $b \in \mathcal{B}$ in state $s \in \mathcal{S}$ results in a transition to $s' \in \mathcal{S}$. $\mathcal{R}_A^i : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ (resp. $\mathcal{R}_V^j : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$) denotes the reward function for agent $i \in \mathcal{M}$ (resp. $j \in \mathcal{N}$). $\gamma \in [0, 1)$ is the discount factor.

Let $\pi_A \in \mathbb{P}_A : \Omega \rightarrow \Delta(\mathcal{A})$ and $\pi_V \in \mathbb{P}_V : \mathcal{S} \rightarrow \Delta(\mathcal{B})$ be the joint policies of *Adversary* and *Victim*, respectively. \mathbb{P}_A and \mathbb{P}_V are their corresponding joint policy spaces.

3.3 Problem Simplification

Inspired by [22], we define the following proposition (the proof is in Appendix B of [42]).

Proposition 1. In a zero-sum partially observable stochastic game, if the victim keeps a fixed joint policy π_V , the state transition of the environment is solely dependent on the adversary's joint policy π_A .

According to Proposition 1, the attacker can treat *Victim* as part of the environment (as shown in Figure 3) and simplify the attack from a ZS-POSG to a POSG:

$$\mathcal{G}_\alpha = (\mathcal{S}, \{\mathcal{A}^i\}, \{\Omega^i\}, \{O_\alpha^i\}, \mathcal{P}_\alpha, \{\mathcal{R}_\alpha^i\}, \gamma), \quad (2)$$

retains the same state space, action space for *Adversary*, observation space, and discount factor as the original \mathcal{G} . However, the observation function, transition function, and reward function for *Adversary* are reconstructed as

$$\begin{aligned} O_\alpha^i(s, a) &= O^i(s, a, b), \\ \mathcal{P}_\alpha(s, a) &= \mathcal{P}(s, a, b), \\ \mathcal{R}_\alpha^i(s, a) &= \mathcal{R}_A^i(s, a, b), \end{aligned} \quad (3)$$

where a and b are the joint actions of *Adversary* and *Victim*, respectively. Eventually, the attacker's objective translates into finding a policy $\pi_A \in \mathbb{P}_A$ that maximizes the accumulated rewards $\sum_i^M \sum_{t=0}^T \gamma^t \mathcal{R}_\alpha^i(s_t, a_t)$ of *Adversary*, where T is the time horizon.

4 Methodology

This section first introduces the framework of *SUB-PLAY* and then describes the design details of each step. The intuition behind *SUB-PLAY* is to adopt a divide-and-conquer strategy, decomposing a complex POSG, as depicted in Equation 2, into multiple relatively simpler POSGs. By tackling these simplified subgames individually, it becomes possible to address the overall complexity of the original POSG more efficiently.

4.1 Attack Framework

SUB-PLAY consists of four main steps: subgame construction, transition dissemination, subpolicy training, and policy combination (see Figure 4).

In the preparation phase, the attacker constructs multiple subgames based on the potentially observed number of victim agents and models each subgame as a POSG. Each subgame initializes its own subpolicy and replay buffers. For specific details of subgame construction, please refer to Section 4.2.

To mitigate the undertraining subpolicies caused by limited interaction transition, the attacker employs a uniform transition dissemination mechanism for all agents in *Adversary*. Then, each agent predefines a transition dissemination table, which is utilized to determine the probability of sharing each transition data among the replay buffers. For specific details of transition dissemination, please refer to Section 4.3.

In the training phase, *Adversary* and *Victim* interact within the environment. The generated transitions are stored in each replay buffer based on the probabilities determined in the transition dissemination table. When a replay buffer accumulates a batch of transitions, the MARL algorithm updates the corresponding subpolicy. The reward of *Adversary* is negatively correlated with the performance of *Victim* in the competition. Therefore, in accordance with the MARL paradigm, each subpolicy tends to minimize the performance of *Victim* to achieve the attack objective. This reward-oriented process does not require any additional knowledge or human intervention. For specific details of subpolicy training, please refer to Section 4.4.

In the deployment phase, the attacker combines subpolicies to form an adversarial policy. Since the attacker has complete control over the adversary and stealthiness is not a concern, we implement

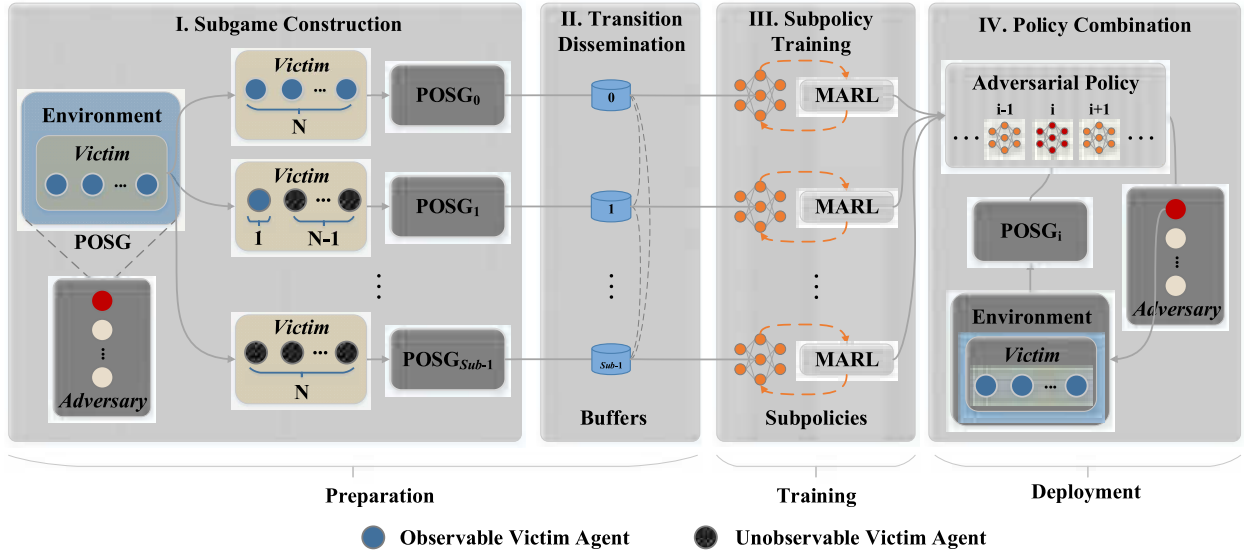


Figure 4: The framework of SUB-PLAY.

the policy combination in a hard-coded manner. When launching an attack, *Adversary* determines which subgame the current competition belongs to and then switches to the corresponding subpolicy to make decisions. More details of policy combination can be referred to Section 4.5.

4.2 Subgame Construction

For a partially observed MAS, the attacker constructs subgames based on the observed agents. We define

$$Sub = N + 1, \quad (4)$$

where *Sub* indicates the number of constructed subgames and *N* is the number of agents belonging to *Victim*. All subgames form a set $\{\mathcal{G}_{\alpha_k}\}_{k \in \mathcal{K}}$. $\mathcal{K} = \{0, 1, \dots, Sub - 1\}$ and each subgame is formulated as a POSG:

$$\mathcal{G}_{\alpha_k} = (\mathcal{S}, \{\mathcal{A}^i\}, \{\Omega_k^i\}, \{O_\alpha^i\}, \mathcal{P}_\alpha, \{\mathcal{R}_\alpha^i\}, \gamma), \quad (5)$$

where the only difference in this equation from Equation 2 is the term $\Omega_k^i \subseteq \Omega^i$. For example, if *Victim* consists of two agents, the attacker constructs three subgames $\{\mathcal{G}_{\alpha_0}, \mathcal{G}_{\alpha_1}, \mathcal{G}_{\alpha_2}\}$, corresponding to the cases where the attacker observes 0, 1, and 2 agents from *Victim*, respectively.

Remark 1. In real-world environments, the partial information about specific agent components may be available to an attacker due to limited perspective. We propose that a conservative strategy can be adopted in high exploration cost environments, treating these agents as unobservable, while a more aggressive strategy can be employed in low exploration cost environments by treating them as observable.

Remark 2. Subgame construction is scalable. In scenarios involving more agents, the attacker can choose a coarser granularity to determine the scope and number of subgames. For instance, when the victim has eight agents, attackers can construct three subgames $\{\mathcal{G}_{\alpha_{0-2}}, \mathcal{G}_{\alpha_{3-5}}, \mathcal{G}_{\alpha_{6-8}}\}$ instead of nine. Furthermore, the attacker

could construct subgames based on regions or apply our method to scenarios where the victim is a single-agent system. In the latter case, subgames could be constructed based on the observability of different components of the single agent.

4.3 Transition Dissemination

The attacker records interactions with *Victim* in the form of transitions. Each transition is represented as a tuple (o_t, a_t, r_t, o_{t+1}) , where o_t is the observation at time step t , a_t is the joint action of *Adversary*, r_t is the instantaneous reward, o_{t+1} is the observation at the next time step, and $t \in [0, T - 1]$, where T is the time horizon. Each agent in *Adversary* maintains a set of replay buffers $\{\mathcal{E}_k^i\}$ that stores the transitions for each subgame, where $i \in \mathcal{M}$ denotes the agent's identifier and $k \in \mathcal{K}$ denotes the subgame's identifier.

Empirical Study. We introduce a concept of *Occupancy Rate (OR)*, which quantifies the occurrence frequency of a subgame. The occupancy rates are related to the number of transitions and serve as the foundation for constructing the transition dissemination table. We conduct an empirical study to explore the relationship between occupancy rates and three partially observable limitations (please refer to Section 5.1 for setup details). The results in Figure 5 reveal two key observations.

Observation 1. (Heterogeneity Property) *The occupancy rate of subgames exhibit variations, which is further affected by the limitations of partial observability.*

Using the first column of Figure 5(c) as an example, the occupancy rates for four subgames $\mathcal{G}_{\alpha_0}, \mathcal{G}_{\alpha_1}, \mathcal{G}_{\alpha_2}$ and \mathcal{G}_{α_3} are 0.04, 0.08, 0.24, and 0.65, respectively. This highlights that some subgames, such as \mathcal{G}_{α_2} and \mathcal{G}_{α_3} , occur frequently, leading to higher occupancy rates (0.24 and 0.65). These subgames provide a sufficient number of transitions, contributing to a more comprehensive dataset for learning and analysis. On the other hand, subgames like \mathcal{G}_{α_0} and \mathcal{G}_{α_1} occur infrequently, resulting in lower occupancy

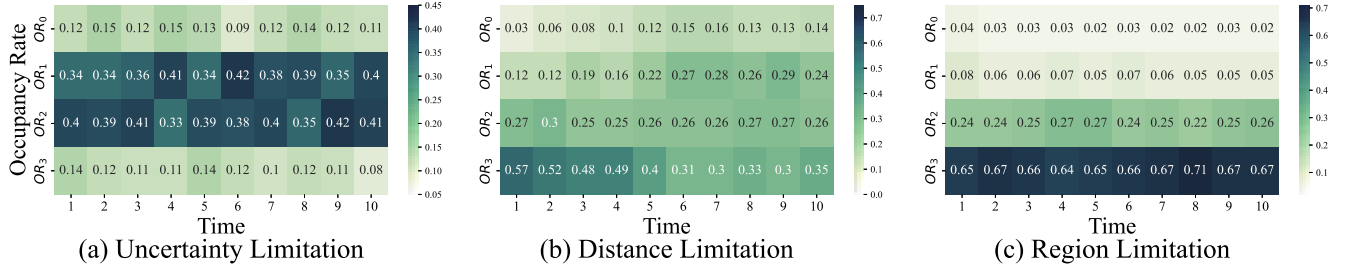


Figure 5: The occupancy rate of subgames under three different limitations. The environment is Predator-prey, with *Victim* consisting of three agents. The vertical coordinate contains the occupancy rate of four subgames $\{\mathcal{G}_{\alpha_0}, \mathcal{G}_{\alpha_1}, \mathcal{G}_{\alpha_2}, \mathcal{G}_{\alpha_3}\}$. As time progresses, the attacker’s policy will be updated.

rates (0.04 and 0.08). The infrequent occurrence of these subgames leads to a scarcity of transitions, which may present challenges for effective learning and decision-making within specific contexts.

Observation 2. (Dynamics Property) *Under distance limitations, the occupancy rate of subgames is influenced by variations in the attacker’s policy.*

The occupancy rates remain stable under uncertainty and region limitations, but they exhibit significant shifts under distance limitations. For example, OR_0 increases from 0.03 to 0.14, while OR_3 decreases from 0.57 to 0.35. The distance limitation is associated with an observable range represented by a circle centered around each agent in *Adversary*. Therefore, the observation by an agent is determined by its behavior pattern or policy.

The heterogeneity property underscores the non-uniform nature of occupancy rates across different subgames. We propose three methods aim to address this non-uniformity and determine occupancy rate values in such scenarios.

Static Estimation. Under uncertainty limitations, if we assume that the uncertainty in the observation of any *Victim*’s agent is the same, OR_k is considered as the probability of repeatedly observing N agents in *Victim* with exactly k successes. The probability of successful observation is μ , which can be introduced as priori knowledge or obtained from historical observation of *Victim*. Thus, the occupancy rate obeys a binomial distribution with parameters N and μ , i.e.,

$$OR_k = \binom{N}{k} \cdot \mu^k \cdot (1 - \mu)^{N-k}, \quad (6)$$

where $\binom{N}{k}$ is the binomial coefficient, which represents the number of ways to choose k successes from N trials.

Static Observation. Alternatively, the attacker calculates the occupancy rate for each subgame \mathcal{G}_{α_k} by counting the number of related transitions in all replay buffers.

$$OR_k = \frac{|\mathcal{E}_k|}{|\mathcal{E}_0| + |\mathcal{E}_1| + \dots + |\mathcal{E}_{Sub-1}|}, \quad (7)$$

where $|\mathcal{E}_k| = \sum_{i=1}^M |\mathcal{E}_k^i|$ and $|\mathcal{E}_k^i|$ indicates the number of transitions stored in \mathcal{E}_k^i . These transitions can either be direct interactions with *Victim* or observations of *Victim*’s interactions with other MASs. Static observation does not rely on additional assumptions or prior knowledge and is applicable to all three types of limitations discussed in this paper.

Dynamic Observation. The dynamics property reveals that occupancy rates may exhibit significant fluctuations and migration in a competitive environment. To account for this, the attacker can utilize *Exponentially Weighted Averages* to accommodate the dynamics of occupancy rates:

$$OR_{k_t} = \beta \cdot OR_{k_{t-1}} + (1 - \beta) \cdot OR'_{k_t}, \quad (8)$$

where $OR_{k_{t-1}}$ denotes the current weighted average of \mathcal{G}_{α_k} ’s occupancy rate, β denotes the weight of historical transition (e.g., $\beta = 0.9$ means that the occupancy rate is approximately to the average of 10 episodes), and OR'_{k_t} denotes the occupancy rate obtained from the current episode statistics. Dynamic observation is applicable to all the three limitations and does not require additional assumptions or prior knowledge. Section 5.2 elucidates the strengths and weaknesses of these three methods through experiments and analysis, guiding on their selection.

Transition Dissemination Table. The attacker determines the *Dissemination Rate* (DR) between each pair of replay buffers based on the occupancy rate. As shown in Figure 6, for each agent $i \in \mathcal{M}$, we define $DR_{\hat{k} \rightarrow k}^i \in [0, 1]$ as the probability of transition transmission from buffer $\mathcal{E}_{\hat{k}}^i$ to buffer \mathcal{E}_k^i .

DR s are determined by four factors: (1) $DR_{\hat{k} \rightarrow k}$ is negatively correlated with the occupancy rate of the destination buffer \mathcal{E}_k^i since the buffer with sufficient transitions does not require additional transitions. (2) $DR_{\hat{k} \rightarrow k}$ is negatively correlated with the distance between $\mathcal{G}_{\alpha_{\hat{k}}}$ and \mathcal{G}_{α_k} since the transitions with higher similarity are provided between buffers close to each other. (3) $DR_{\hat{k} \rightarrow k}$ is positively correlated with the number of constructed subgames since the mean value of occupancy rates decreases as *Sub* increases. (4) $DR_{\hat{k} \rightarrow k}$ is negatively correlated with the dispersion of occupancy rates since high dispersion is prone to multiple occupancy rates with lower values. In summary, we define $DR_{\hat{k} \rightarrow k}$ as

$$DR_{\hat{k} \rightarrow k} = \begin{cases} \text{clip}((\lambda - OR_k + \sigma)^{\frac{|k-k|}{\sqrt{Sub}}}, 0, 1), & \text{if } OR_k \leq \lambda \\ \frac{|k-k|}{\sigma \sqrt{Sub}}, & \text{if } OR_k > \lambda \end{cases} \quad (9)$$

where λ is an adjustment parameter positively correlated with the complexity of the two-team competitive environment. σ is the standard deviation of all occupancy rates and is used to measure their dispersion. $|k - k|$ indicates the distance between two subgames.

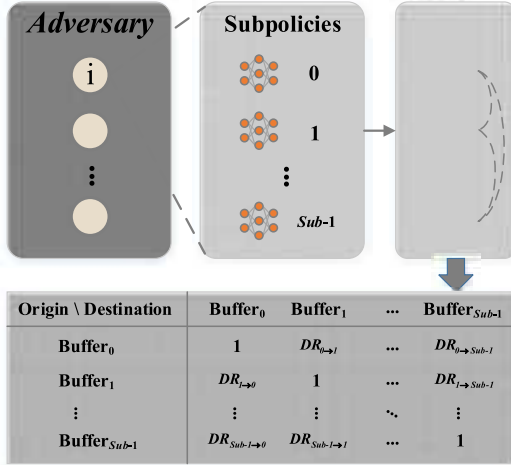


Figure 6: Each agent in *Adversary* maintains a transition dissemination table.

\sqrt{Sub} indicates that as the number of constructed subgames increases, the transition dissemination between buffers will become more frequent, meaning that the value of $DR_{\hat{k} \rightarrow k}$ will increase. We use the clip function to limit the value of $DR_{\hat{k} \rightarrow k}$ to the range $[0,1]$.

All DR s collectively form a transition dissemination table. When a new transition is generated, the agent allocates it to replay buffers based on the probabilities recorded in this table. If static estimation or static observation is employed to determine occupancy rates, this table is static. In contrast, if dynamic observation is used to determine occupancy rates, this table changes dynamically. The impact of transition dissemination is presented in Figure 20 of [42].

4.4 Subpolicy Training

Each agent $i \in \mathcal{M}$ maintains a set $\{\pi_{\alpha_k}^i\}_{k \in \mathcal{K}}$ consisting of all subpolicies. We perform *Policy Meritocracy (PM)* to preserve top-performing subpolicies based on the harmonic mean of their test performance across L metrics, mitigating performance fluctuations caused by non-stationarity.

$$PM_k^i = \frac{L}{\sum_{l=1}^L 1/\eta_{k_l}^i}, \quad (10)$$

where $\eta_{k_l}^i$ indicates the test performance of the subpolicy $\pi_{\alpha_k}^i$ with respect to metric l . In the policy pool, only one subpolicy is retained for each subgame to minimize storage overhead. Replacements occur when a subpolicy with superior test performance emerges. The detailed update process for the subpolicies is outlined in Algorithm 1 of [42].

4.5 Policy Combination

The attacker combines all subpolicies $\{\pi_{\alpha_k}^i\}_{i \in \mathcal{M}, k \in \mathcal{K}}$ to generate the final adversarial policy. Specifically, at each time step of the deployment phase, *Adversary* obtains a joint observation (o_1, o_2, \dots, o_M) and then outputs a joint action (a_1, a_2, \dots, a_M) . The subgame construction ensures that each observation exclusively belongs to a

single subgame, without any overlap or intersection between the subgames, *i.e.*, $\Omega_{\hat{k}}^i \cap \Omega_k^i = \emptyset$, where $\hat{k} \neq k$ ($\hat{k}, k \in \mathcal{K}$), $i \in \mathcal{M}$. Furthermore, unlike backdoor attacks [17, 57], the attacker can modify the program structure without concerns about stealthiness. Hence, we propose that the attacker can implement the policy combination in a hard-coded manner.

Taking the distributed MARL as an example, each agent $i \in \mathcal{M}$ determines which subgame a received observation o_i belongs to and subsequently selects the corresponding subpolicy to guide its action decision-making:

$$\pi_{\alpha}^i = \begin{cases} \pi_{\alpha_0}^i, & \text{if } o_i \in \Omega_0^i \\ \dots & \\ \pi_{\alpha_{Sub-1}}^i, & \text{if } o_i \in \Omega_{Sub-1}^i \end{cases} \quad (11)$$

This approach has two advantages: it provides a straightforward logic and greater flexibility, allowing each subgame to employ a separate MARL algorithm. The implementation of policy combination can be found in Algorithm 2 of [42].

5 Evaluation

This section first introduces the evaluation setup and then evaluates *SUB-PLAY* through six perspectives: attack performance, ablation study, transferability, scalability, overhead, and potential defenses.

5.1 Setup

Environment. We adopt two-dimensional environments, Predator-prey and World Communication, within the Multi Particle Environments (MPE) [40] framework developed by OpenAI. MPE is a gym-based benchmark designed for cooperative, competitive, and mixed MARL tasks.

- **Predator-prey.** There are N slower predators controlled by *Victim* and M preys controlled by *Adversary*. The predators cooperate to collide with preys, while the preys cooperate to avoid collisions. The environment is initialized randomly at the beginning of each episode, including the positions of agents and obstacles.
- **World Communication.** This environment shares similarities with the Predator-prey setup but includes additional features. (1) There are M foods that preys are rewarded for being close to. (2) The environment randomly initializes a forest, making all agents invisible within it initially. (3) A leader predator exists, having full visibility of all agents and the ability to communicate with other predators to enhance cooperation.

The state and action spaces of both environments are continuous. The evaluation of the attack performance is conducted in five different scenarios (MvN): 1v3, 2v3, 3v3, 2v2, and 4v2.

Partial Observability Implementation. Observation is a multi-dimensional vector. We additionally introduce a Mask vector consisting of 0s and 1s, which has the same dimensions as Observation. The Mask is determined by specific rules for partial observability. We multiply each element of Observation with the corresponding element of the Mask, and the result is partial observation. For specific implementation, please consult Algorithm 3 of [42].

In Predator-prey, we conduct evaluations under both uncertainty and distance limitations. The uncertainty rate (*i.e.* μ in Equation 6)

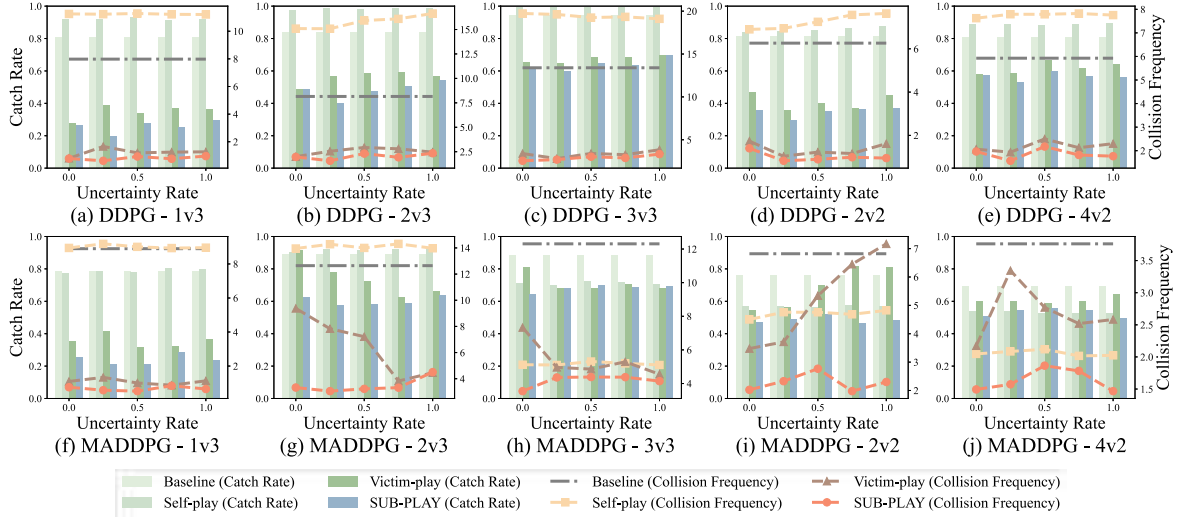


Figure 7: The attack performance of *SUB-PLAY* under uncertainty limitations in Predator-prey.

ranges from $\{0.00, 0.25, 0.50, 0.75, 1.00\}$. For example, $\mu = 0$ indicates complete observability, and the mask array consists of all 1s. The observable distance ranges from $\{0.5, 1.0, 1.5, 2.0\}$.

In World Communication, we conduct evaluations under region limitations. This environment has implemented partial observability, where agents located within a specific region are not visible (a forest with a fixed size but randomly initialized positions).

Implementation Details. Our evaluations are conducted on four servers with Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz, 32GB RAM. Python and PyTorch are used for code implementation.

Selected MARL algorithms include DDPG and MADDPG, serving as representatives for DTDE and CTDE architectures, respectively. DDPG and MADDPG are actor-critic algorithms commonly used in reinforcement learning. These algorithms consist of an actor network and a critic network. We utilize a two-layer ReLU MLP with 128 units in each layer to parameterize all policies. The actor network’s output layer incorporates a Tanh activation function. For weight initialization, we use Xavier normal with a gain of 1.0 for all layers in both victim and adversarial policy training. Biases are initialized with zeros. The chosen optimizer is Adam, with a learning rate of 0.001 and ϵ set to 10^{-8} .

To ensure smooth policy updates, we employ *Exponential Moving Average (EMA)* with a decay rate of 0.95. Random noise sampled from a normal distribution with a standard deviation of 0.01 is added to the output actions to promote exploration. The discount factor γ for RL is set to 0.95. The adjustment parameter λ in Equation 9 is set to 0.5. All reported results are averaged over 1,000 test runs. For additional implementation details, parameter settings, and training results of *Victims*, please refer to Appendix C of [42].

Comparison Methods. (1) *Baseline*: *Victim*’s normal performance on a specific task. Specifically, *Adversary* deploys a heuristic policy where the agent’s movement is characterized by a fixed speed and direction, with random updates after collisions. (2) *Self-play* [1, 3]: Similar to the process of victim training, this setup grants the attacker complete access to the environment. The attacker randomly

initializes both *Adversary* and *Victim*, allowing them to compete with each other and undergo updates. Ultimately, *Adversary* is retained to carry out the attack. (3) *Victim-play* [15, 22, 38, 59, 60]: Apart from substituting the algorithm with MARL, retain the other fundamental settings of this framework (refer to Section 2.2), specifically fixing *Victim* and updating *Adversary*.

Metrics. We adopt *Catch Rate (CR)* and *Collision Frequency (CF)* as evaluation metrics.

$$CR = \frac{Num_c}{Num_e}, \quad (12)$$

where Num_c indicates the number of episodes where the *Adversary* is caught; Num_e indicates the total number of episodes.

$$CF = \frac{1}{Num_e} \sum_{e=1}^{Num_e} \sum_{j=1}^N Num_{ej}, \quad (13)$$

where Num_{ej} denotes the number of collisions between the j th predator and the prey in a specific episode e .

The attacker aims to evade pursuit and minimize *Victim*’s *CR* and *CF*. Policy meritocracy is based on the harmonic mean of these two metrics,

$$PM = 2 \cdot \frac{CR \cdot CF}{CR + CF}. \quad (14)$$

We calculate the average improvement in attack performance of *SUB-PLAY* compared to *Victim-play* by

$$\frac{(PM_B - PM_S) - (PM_B - PM_V)}{PM_B - PM_V}, \quad (15)$$

where PM_B , PM_V , and PM_S represent the victim’s average performance when the attacker executes *Baseline*, *Victim-play*, and *SUB-PLAY*, respectively.

5.2 Attack Performance

We evaluate the attack performance of *SUB-PLAY* across different environments and limitations. The training process of *SUB-PLAY* can be found in Figure 17 of [42].

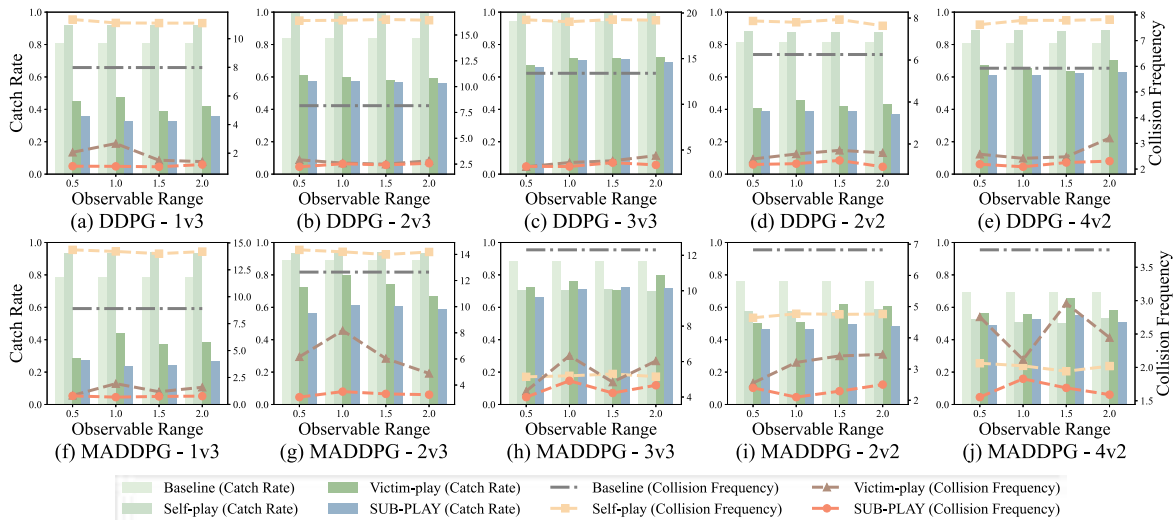


Figure 8: The attack performance of *SUB-PLAY* under distance limitations in Predator-prey.

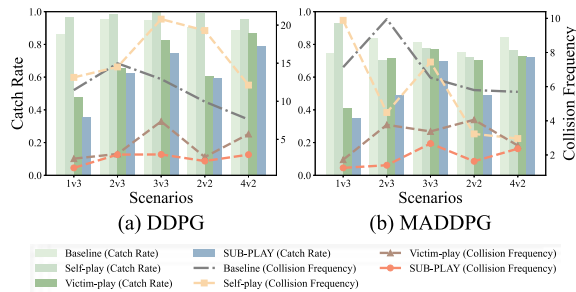


Figure 9: The attack performance of *SUB-PLAY* under region limitations in World Communication.

Uncertainty Limitation. Under uncertainty limitations, *SUB-PLAY*, on average, reduces the victim’s performance to 51.98% of the baseline (see Figure 7). Moreover, *SUB-PLAY* outperforms the other two methods and minimizes the victim’s catch and collision rates in 94.0% (47/50) and 98.0% (49/50) scenarios. Compared to *Victim-play*, *SUB-PLAY* demonstrates an average improvement of 32.22% in attack performance (refer to Equation 15).

The results also show that when the MARL algorithm is MADDPG (resulting in increased input dimensionality due to information sharing among agents), *SUB-PLAY* demonstrates more stable attack performance, showcasing its potential to handle more complex environments. In addition, an unexpected outcome is that the attack remains effective even when the uncertainty is set to 1.00 (*i.e.*, the attacker has no observations of the victim). According to Silver *et al.* [52], maximizing rewards is sufficient to drive intelligent behavior. In our scenarios, this suggests that the attacker can find adversarial policies by focusing on maximizing the rewards obtained, even without direct observation of the victim.

Distance Limitation. Under distance limitations, *SUB-PLAY*, on average, reduces the victim’s performance to 55.71% of the baseline

(see Figure 8). Moreover, *SUB-PLAY* outperforms the other two methods and minimizes the victim’s catch and collision rates in 97.5% (39/40) scenarios. Compared to other methods, the impact of the observation range on *SUB-PLAY* is relatively minor, indicating that *SUB-PLAY* demonstrates better adaptability to dynamic environments. Compared to *Victim-play*, *SUB-PLAY* demonstrates an average improvement of 27.16% in attack performance.

Region Limitation. Under region limitations, *SUB-PLAY* reduces the victim’s performance to an average of 59.07% of the baseline (see Figure 9). Moreover, *SUB-PLAY* outperforms the other two methods and minimizes the victim’s catch and collision rates in 100.0% (10/10) scenarios. Compared to *Victim-play*, *SUB-PLAY* demonstrates an average improvement of 50.22% in attack performance. These results indicate that *SUB-PLAY* exhibits more significant attack potential in complex environments.

Furthermore, we compare *SUB-PLAY* with two additional variants of *Victim-play* [22, 59], further validating *SUB-PLAY*’s attack performance in partially observable environments. Details regarding the settings of these two variants and the final results are provided in Appendix D of [42].

Visualization Results. The visualization results (Figure 18 of [42]) show that in Predator-prey, the preys tend to flee to the edge of the map at the maximum speed from different directions and then stay while trying to bypass predators and obstacles. They quickly get rid of collisions if they occur.

Similarly, in World Communication, the preys do not hide in the forest or approach foods for additional rewards. These observations demonstrate that adversarial policies effectively utilize the speed advantage of the preys to evade predators, mitigating the disadvantages of partial observability. t-SNE (see Figure 10) shows that the activations of the victim’s policy network are significantly different when facing a normal opponent compared to an adversarial policy. More results can be found in Figure 19 of [42].

Table 1: The impact of three occupancy rate determination methods on attack performance under three partial observability limitations. The performance is measured by two metrics (CR/CF). Acronyms: Static Estimation (SE), Static Observation (SO), Dynamic Observation (DO).

Limitations		DDPG			MADDPG		
		SE	SO	DO	SE	SO	DO
Uncertainty	0.00	0.489 / 1.969	0.527 / 2.115	0.571 / 2.527	0.626 / 3.327	0.622 / 4.007	0.606 / 3.846
	0.25	0.401 / 1.573	0.538 / 2.205	0.575 / 2.532	0.579 / 3.053	0.607 / 3.506	0.725 / 6.460
	0.50	0.477 / 2.309	0.530 / 2.430	0.619 / 3.362	0.583 / 3.228	0.605 / 3.389	0.697 / 6.209
	0.75	0.502 / 1.927	0.532 / 2.395	0.580 / 2.475	0.591 / 3.318	0.614 / 3.909	0.639 / 3.904
	1.00	0.543 / 2.345	0.566 / 2.563	0.592 / 2.919	0.638 / 4.506	0.647 / 4.019	0.687 / 4.880
Distance	0.5	-	0.598 / 2.554	0.574 / 2.240	-	0.609 / 3.335	0.563 / 3.075
	1.0	-	0.657 / 3.059	0.571 / 2.529	-	0.647 / 4.148	0.614 / 3.500
	1.5	-	0.609 / 3.025	0.564 / 2.421	-	0.667 / 4.430	0.606 / 3.324
	2.0	-	0.620 / 2.786	0.561 / 2.579	-	0.654 / 4.464	0.589 / 3.264
Region	1	-	0.664 / 2.917	0.626 / 2.998	-	0.594 / 1.855	0.489 / 1.397

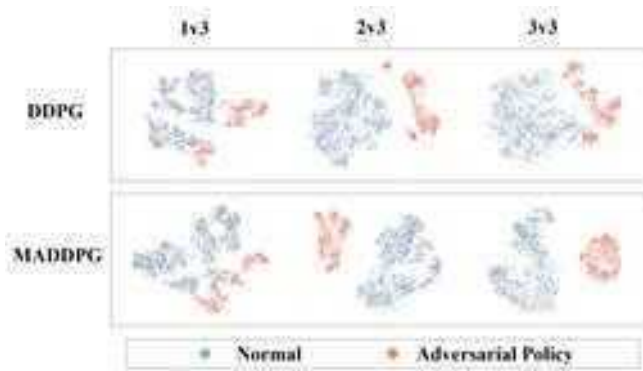


Figure 10: t-SNE activations of the victim when playing against different opponents.

Proactive Masking. Intuitively, partial observability poses a challenge for adversarial policy generation. However, evaluations indicate that proactive masking of environment observations may enhance the attack performance in some scenarios. For example, in Figure 7(a), the attack performance of *SUB-PLAY* is superior when the uncertainty rate is 0.25 compared to the fully observable scenario. The reason is that the minor partial observability corresponds to simplifying the input to the attacker. This suggests that proactive masking may facilitate early training of adversarial policies in complex environments. Recent work has introduced curriculum learning to address sparse rewards in adversarial policy generation [59], but it requires the access to different versions of victims. In contrast, proactive masking does not rely on a similar assumption.

Impact of Occupancy Rates. To compare the impact of the three methods proposed in Section 4.3 for determining occupancy rates, we evaluate them under three partial observability limitations (static estimation, static observation, and dynamic observation). In the static observation, the attacker pre-observes the victim for 1000 episodes, while in the dynamic observation, the parameter β is set to 0.9.

Table 1 shows that *SUB-PLAY* achieves optimal attack performance with static estimation under uncertainty limitations. This is because static estimation provides more accurate occupancy rate estimations. Furthermore, *SUB-PLAY* performs superior attacks using dynamic observation under distance and region limitations. The

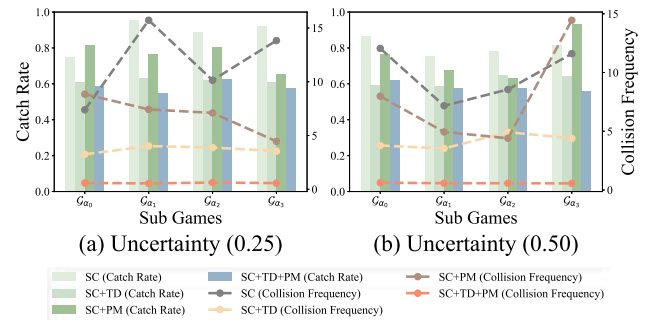


Figure 11: The attack performance of subpolicies under uncertainty limitations.

former is due to the change in occupancy rates as the adversarial policy updates, while the latter is attributed to the additional dynamics introduced by the randomly initialized position of the forest in World Communication.

Therefore, in all evaluations, we set that the attacker adopts static evaluation under uncertainty limitations and dynamic observation under distance and region limitations.

5.3 Ablation Study

Unless additional specifications exist, the subsequent evaluations use the 2v3 scenario, and the MARL algorithm is set to MADDPG.

Component Evaluation. We perform an ablation study to assess the contribution of each component (subgame construction, transition dissemination, and policy meritocracy in subpolicy training) in *SUB-PLAY*. Table 2 demonstrates that when the subgame construction is applied in isolation, the attack performance is inferior. However, when transition dissemination is also performed, the attack performance is significantly improved. This highlights the crucial role of transition dissemination in enhancing performance in partially observable environments. Policy meritocracy also contributes to improved attack performance and is compatible with subgame construction and transition dissemination.

Subgame Evaluation. We continue to explore the reasons behind the performance improvements achieved by *SUB-PLAY*. Figure 11 presents the attack performance of subpolicies in their respective

Table 2: The ablation results of components in SUB-PLAY measured by two metrics (CR/CF). Acronyms: Subgame Construction (SC), Transition Dissemination (TD), Policy Meritocracy (PM).

Methods	Limitations				
	Uncertainty (0.25)	Uncertainty (0.50)	Distance (0.5)	Distance (2.0)	Region (1)
<i>Self-play</i>	0.920 / 14.280	0.916 / 13.998	0.936 / 14.349	0.935 / 14.187	0.704 / 4.486
<i>Victim-play</i>	0.782 / 7.823	0.727 / 7.215	0.728 / 6.163	0.670 / 4.891	0.718 / 3.763
SUB-PLAY (SC)	0.830 / 8.402	0.759 / 7.604	0.765 / 6.296	0.708 / 5.982	0.835 / 6.563
SUB-PLAY (SC+TD)	0.617 / 3.740	0.627 / 4.438	0.700 / 6.552	0.672 / 4.675	0.688 / 3.309
SUB-PLAY (SC+PM)	0.731 / 6.059	0.708 / 6.318	0.735 / 6.113	0.677 / 4.576	0.561 / 1.634
SUB-PLAY (SC+TD+PM)	0.579 / 3.053	0.583 / 3.228	0.563 / 3.075	0.589 / 3.264	0.489 / 1.397

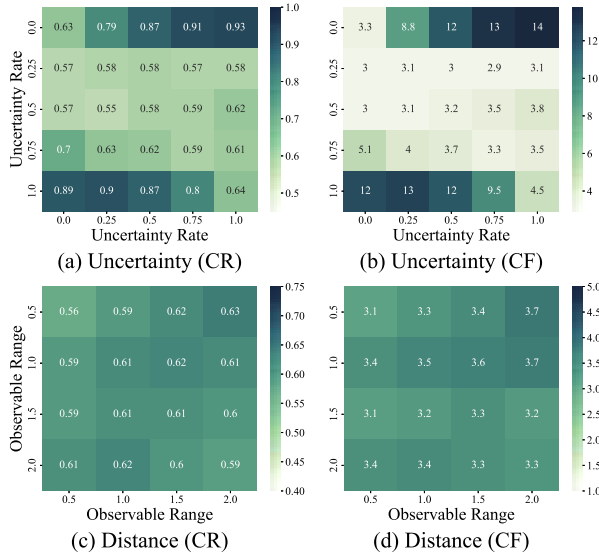


Figure 12: The transferability of adversarial policies across different scenarios. The vertical coordinate represents the limitation set in the training phase and the horizontal coordinate represents the limitation set in the testing phase.

corresponding subgames under uncertainty limitations. The results reveal that when the attacker only applies subgame construction and policy meritocracy, the resulting adversarial policies exhibit substantial performance differences across different subgames. However, with the inclusion of transition dissemination, all subpolicies show notable improvements in their attack performance. Furthermore, the performance disparities among the subpolicies are significantly reduced.

5.4 Transferability

In real-world scenarios, the limitations on partial observability are not static; for instance, environmental factors like weather can impact sensor performance, leading to variations in the observable distance for the adversary agents. Therefore, we evaluate the transferability of adversarial policies across different parameter settings within the same limitation.

Figure 12(a)-(b) illustrate that adversarial policies can be transferred in specific scenarios (*i.e.*, when the uncertainty rate is 0.25,

0.50, and 0.75). However, transferability diminishes when the uncertainty rate is 0.00 or 1.00. This is due to the significant differences in the modeling between fully unobservable and fully observable scenarios compared to partially observable scenarios. Additionally, subgame construction and transition dissemination are ineffective in these two cases. Figure 12(c)-(d) demonstrate that adversarial policies exhibit transferability under distance limitations. In conclusion, adversarial policies derived from SUB-PLAY demonstrate transferability in similar partially observable environments.

However, the transferability of adversarial policies encounters limitations when applied to victims with substantial differences. Table 3 illustrates that the adversarial policy proves ineffective in reducing the catch rate of heuristic victims, although it succeeds in diminishing their collision frequency. The reduction in collision frequency is attributed to the adversarial policy’s capacity to swiftly escape captures, whereas heuristic attackers frequently engage in repeated collisions. The inability to decrease the victims’ catch rate results from adversaries consistently moving at the maximum speed in the adversarial policy, thereby increasing the probability of collision – especially when preys are slow, potentially concluding the episode before colliding with predators. Moreover, preys lose the ability to evade nearby predators.

5.5 Scalability

SUB-PLAY is scalable, meaning the attacker can adjust the granularity of subgame construction (Section 4.2) based on the scenario and requirements. Equation 4 demonstrates the most potent attack by constructing $Sub = N + 1$ sub-games against N victim agents. However, in resource-constrained environments, the attacker can adopt a coarser granularity, *i.e.*, $Sub < N + 1$.

We validate the scalability of SUB-PLAY in Predator-prey (algorithm = MADDPG, distance = 0.5, 1.0, 1.5, 2.0, scenarios = 2v3), constructing 1, 2, 3, and 4 subgames. Figure 13 illustrates that under four different granularity settings, SUB-PLAY can decrease the victim’s catch rate to 0.736, 0.641, 0.611, and 0.593, and the collision frequency decreases to 6.315, 3.738, 3.504, and 3.291, respectively. The results indicate that finer granularity leads to better attack effectiveness, while the rate of performance improvement diminishes as the number of subgames increases, gradually approaching saturation. Additionally, the training time of SUB-PLAY shows a positive correlation with the number of subgames (1171.97s, 1495.39s, 2031.03s, and 2556.17s, respectively), indicating that attackers can manipulate the granularity of subgame construction to scale SUB-PLAY to more complex environments.

Table 3: The attack performance of adversarial policies with heuristic victims. The performance is measured by two metrics (CR↓/CF↓).

Attacker	Heuristic		SUB-PLAY
Victim	Heuristic		
Uncertainty	0.25	0.261 / 0.699	0.307 / 0.599
	0.50	0.259 / 0.694	0.286 / 0.515
Distance	0.5	0.225 / 0.613	0.275 / 0.508
	2.0	0.265 / 0.692	0.313 / 0.600
Region	1	0.304 / 0.802	0.439 / 0.807
Victim	MADDPG		
Uncertainty	0.25	0.900 / 12.643	0.579 / 3.053
	0.50	0.918 / 13.258	0.583 / 3.228
Distance	0.5	0.942 / 13.540	0.563 / 3.075
	2.0	0.920 / 13.224	0.589 / 3.264
Region	1	0.837 / 9.954	0.489 / 1.397

5.6 Overhead

The main overhead of adversarial policy-based attacks arises from three aspects: interaction costs, training costs, and decision delays. SUB-PLAY avoids extra interaction costs (Section 4.3) and decision delays (Section 4.5) but incurs additional training costs due to subgame construction (Section 4.2) and separate training for each subgame (Section 4.4), which scales linearly with the number of subgames.

We explore the training time of SUB-PLAY in Predator-prey (algorithm = DDPG, MADDPG, distance = 0.5, scenarios = 1v3, 2v3, 3v3, 2v2, 4v2). The results in Table 4 indicate that under the condition of maximizing the number of subgames ($Sub = N+1$), the average training times of SUB-PLAY in distributed and centralized MARL algorithms are 1529s and 2781s, respectively (1.73 and 2.02 times that of Victim-play). Nevertheless, the training time of SUB-PLAY remains significantly lower than that of the well-trained victim (only 2.39%). Additionally, as demonstrated in Section 5.5, the attacker can adjust the granularity of subgame construction to reduce training costs.

5.7 Potential Defenses

Adversarial policies pose significant challenges to the real-world deployment of MARL. In response, we explore viable defense approaches to mitigate these security threats.

Adversarial Retraining. An intuitive approach involves the victim adopting adversarial retraining during the training phase [23]. This allows the victim to identify and address vulnerabilities in its policy continuously. However, as shown in Figure 14, this does not render the adversarial policy ineffective. This is because there is no theoretical evidence indicates that the victim’s MARL policy, after adversarial retraining, can approach a Nash equilibrium more closely. Thus, the exploitable space persists, albeit potentially shifting with adversarial retraining. Moreover, as previously mentioned, finding or approximating a Nash equilibrium in a multi-agent competition is at least as tricky as PPAD-complete.

Policy Ensemble. While it is challenging to eliminate the threat posed by adversarial policies through training alone, victims could mitigate this threat by adjusting MARL’s deployment strategy. One such strategy involves adopting a policy ensemble approach, where the victim prepares a set of policies and consistently selects policies

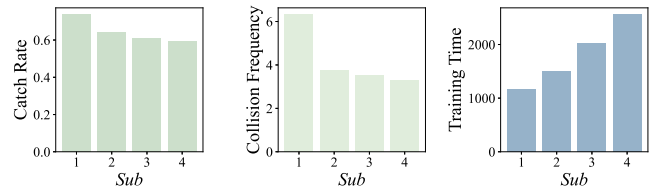


Figure 13: Scalability evaluation.

Table 4: The difference in training time between adversarial policies and victims (s).

Methods	Scenarios					Average
	1v3	2v3	3v3	2v2	4v2	
DDPG						
Victim	53240	68101	80800	53314	70364	65164
Victim-play	625	766	982	738	1288	880
SUB-PLAY	929	1293	1732	1173	2519	1529
MADDPG						
Victim	92048	106458	147257	98678	131235	115135
Victim-play	828	1172	1549	1115	2201	1373
SUB-PLAY	1577	2556	3682	2304	3785	2781

from this ensemble for deployment. Intuitively, this may prevent the attacker from adapting to a specific policy.

The results on the left side of Table 5 indicate that if attackers have access to all policies, the policy ensemble shows limited defensive effect. For instance, the entry in the first row and first column, -0.07, denotes a decrease of 0.07% in the effectiveness of adversarial policies, which is almost negligible (note that -100% corresponds to the failure of the adversarial policy). This implies that the assumption of freezing the victim’s policy in Section 3.1 can be relaxed.

If attackers have access to only a subset of the victim’s policies (33%), there is a certain degree of reduction in the effectiveness of adversarial policies (the results on the right side of Table 5). The insight derived from this is that the victim could periodically update the ensemble pool to prevent the attacker from adapting to all policies. The victim may also consider increasing the diversity of the policies in the pool [12, 28], making their weaknesses significantly different, which might lead to fluctuations in adversarial policies, preventing them from converging. Policy ensemble can be coupled with a dynamic switching mechanism to enhance the defense approach further, wherein the switching time or policy selection is dynamically changed [30, 39, 41].

Fine-tuning. Continual fine-tuning during deployment may also prevent the attacker from adapting to a specific victim, offering lower training costs than policy ensemble. However, the limited defensive effectiveness of fine-tuning, as shown in Figure 15, suggests that the distance between policies before and after fine-tuning remains close.

6 Discussion

Emphasizing Deployment Details. We not only introduce SUB-PLAY but also reveal that even with partial observations, adversarial policies induce the failure of MARL. Since mitigating the threat of adversarial policies through improvements in the training

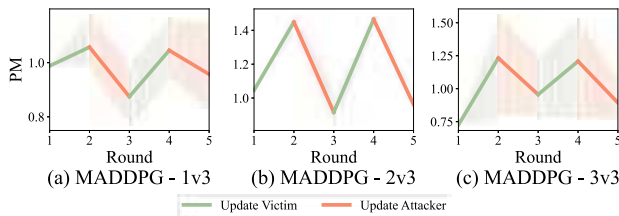


Figure 14: Adversarial retraining results of 5 rounds.

framework is challenging, we propose that defenders prioritize the deployment details of MARL rather than solely focusing on enhancing algorithm performance. Additionally, *SUB-PLAY* can serve as a method to measure the lower bound of MARL performance in adversarial scenarios.

Limitation and Future Work. *SUB-PLAY* still has some limitations. (1) Due to the scarcity of environments that facilitate multi-agent competitive settings with partial observability, our testing is limited to two environments. However, we extensively evaluate *SUB-PLAY* in various settings, including different types of partial observability limitations, multiple scenarios with varying numbers of agents, and two MARL architectures. (2) Although no additional interaction is required, the training costs incurred by *SUB-PLAY* are directly proportional to the number of subgames. Thus, a trade-off exists between training costs and attack performance, which still needs to be addressed. (3) The current method assumes that the attacker engages in multiple interactions with the victim. We plan to adopt offline RL techniques [7, 10] to relax this setup.

Real-world Scenarios. Agents’ restricted perception capabilities and observation permissions give rise to numerous partially observable scenarios in real-world settings. For instance, the anticipated applications of drone swarms and robots in MAS encompass encirclement systems [24], security systems [27], strategic maneuvers [11], and human-robot teams [8]. Nonetheless, their environmental perception is confined by the deployed sensors. For example, the 4D LiDAR L1 on Unitree Go2 has a scanning distance of 30 meters [51], leading to partially observable phenomena when the targets exceed this range or become obstructed.

Potential Damages. The potential damages of *SUB-PLAY* include attaining targeted victories or illicit profits. For instance, these could involve defeating specific opponents in RoboMaster [9], exploiting strategic vulnerabilities in poker AI like Dou Dizhu [65] to gain illegal profits online (given the prevalence of AI in online poker), or bypassing security MAS to jeopardize property and personal safety. To address the potential damages posed by *SUB-PLAY*, we explore potential defense methods in Section 5.7 and provide directions for future research: Compared to costly adversarial retraining and limited defensive performance of fine-tuning, deploying MARL in the form of policy ensembles and increasing the diversity of the policy pool is a more practical and effective approach.

7 Related Work

7.1 RL Security

A substantial body of research is leveraging RL to achieve specific security objectives [16, 44, 56, 61, 67]. Nevertheless, the security

Table 5: The defensive effectiveness of policy ensemble, with values given in percentage (%).

Access	100%			33%		
Scenarios	1v3	2v3	3v3	1v3	2v3	3v3
Uncertainty						
0.00	-0.07	+0.02	-0.04	-2.74	+4.09	-0.89
0.25	+0.02	-0.25	+0.10	-9.86	-13.58	-12.45
0.50	+0.00	-0.02	+0.08	-9.68	-9.14	-17.01
0.75	-0.01	-0.07	+0.04	-15.55	-2.56	+2.85
1.00	+0.00	+0.04	+0.08	-25.78	-0.55	+9.68
Distance						
0.5	-0.09	-0.15	-0.03	-16.17	-7.99	-11.98
1.0	-0.12	-0.12	-0.01	-30.15	-5.65	+0.25
1.5	-0.29	-0.12	-0.02	-20.24	-9.36	-32.69
2.0	-0.13	-0.28	+0.14	-16.01	-20.51	-43.39
Region						
1	-0.08	-0.24	+0.00	-7.99	-37.44	-17.94

dimensions inherent to RL need to be more adequately addressed. This section thoroughly examines security research concerning the three fundamental components of RL.

Reward Manipulation. Unlike the modification of labels in deep learning [55, 64, 68], RL introduces backdoor attacks via reward manipulation [32, 43]. Zhang *et al.* [70] developed a dynamic reward-poisoning attack targeting online RL applications, while Chen *et al.* [6] extended the concept of backdoor attacks to cooperative MASs. Wang *et al.* [57] introduced a unique RL-specific paradigm for backdoor attacks, where the attacker trains a benign policy and a trojan policy, merging them into a backdoor policy by behavior cloning. Guo *et al.* [21] discovered a pseudo-trigger space that can trigger RL backdoors. In response, they proposed PolicyCleanse to perform model detection and backdoor mitigation.

State Manipulation. Inspired by adversarial examples [5, 36, 66], attackers in RL can disrupt the victim by perturbing the environment state. Huang *et al.* [29] applied FGSM [18] to DRL and launched an adversarial attack on the DQN policy in Atari games [45]. Behzadan *et al.* [2] introduced a policy induction attack, where the attacker determines the victim’s actions based on a pre-trained target policy and perturbs the states by FGSM and JSMA [47]. Sun *et al.* [53] proposed a white-box attack called the critical point attack, which strategically explores state-action combinations to identify points with high payoff and inject subtle perturbations during the victim’s deployment phase.

Action Manipulation. The agent’s action determines the agent-environment boundary, allowing attackers to launch attacks by manipulating it. Lee *et al.* [34] proposed two victim manipulation attacks: the myopic action-space attack injects action perturbations based on current observations, while the look-ahead action-space attack considers future steps to maximize the attack’s impact. However, directly manipulating the victim’s actions is impractical.

In contrast, adversarial policies only need to control the attacker’s action. Gleave *et al.* [15] pioneered *Victim-play*, which involves manipulating the adversary’s actions to induce suboptimal decisions from a fixed RL model during deployment. Wu *et al.* [60] incorporated explainable AI techniques into adversarial policy generation, enhancing the stealthiness by launching attacks only when the victim pays attention to them. Guo *et al.* [22] extended *Victim-play* from zero-sum to general-sum environments,

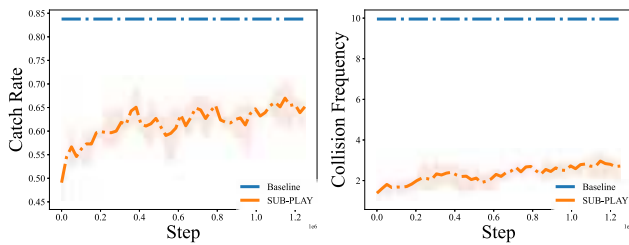


Figure 15: The attack performance varies with the fine-tuning of the victim’s policy.

revealing its potential in assessing the fairness of competitions or systems. Wang *et al.* [59] explored adversarial policies in discrete action scenarios and achieved success against superhuman-level Go AIs, demonstrating that *near-Nash* or *ϵ -equilibrium* policies are exploitable. Guo *et al.* [23] attempted to mitigate the potential threat of adversarial policies from a training perspective and introduced a provable defense called PATROL. The aim is to bring the victim closer to a Nash equilibrium in a two-player competition. Liu *et al.* [38] explored adversarial policy attacks in scenarios where attackers only have partial control over the adversary and proposed adversarial training with two timescales to mitigate the threats posed by adversarial policies.

7.2 Partial Observability in RL

Partial observability limits the attacker’s access to complete environmental information. To address this, existing research suggests two potential methods for the attacker.

Inference. Inference entails using available observations and prior knowledge to complete unobserved content, including environmental [62, 63] and agent inference [20, 48]. (1) Environmental Inference: Partial state information is used to infer the global environment in the spatial dimension. Yang *et al.* [63] proposed a supervised learning-based hallucinator for inferring the environment from current observations, effective in static environments but potentially less suitable for highly dynamic competitions. (2) Agent Inference: Historical interaction is used to infer the unobservable agents in the temporal dimension. Papoudakis *et al.* [48] proposed constructing policies for all agents through representation learning. During deployment, the policies and local observations of the controlled agent are utilized to infer the invisible agents. However, this method relies on the victim’s policy knowledge and is unsuitable for black-box or competitive environments.

Generalization. Enhancing the generalization capability of agents can effectively adapt to the environment’s diversity, dynamics, and unpredictability [33]. Intuitively, this also applies to policy improvements in partially observable scenarios. Ghosh *et al.* [14] demonstrated that partitioning a partially observable task into multiple subtasks can effectively improve the performance of RL policies. *SUB-PLAY* draws on the insight of generalization-based approaches to address the partially observable problem in multi-agent competitive environments, as it applies to dynamic environments and does not rely on additional victim information.

8 Conclusion

This paper proposes *SUB-PLAY*, a novel black-box attack framework in partially observable multi-agent competitive environments. The effectiveness of *SUB-PLAY* in enhancing the attack performance of adversarial policies is showcased through divide-and-conquer strategies and transition dissemination, as evidenced by extensive evaluations conducted across various partially observable limitations and MARL algorithms. Moreover, we examine three potential defense strategies to mitigate the risks associated with *SUB-PLAY*. The evaluation results indicate that policy ensemble is more effective than adversarial retraining and fine-tuning. Future investigations can concentrate on enhancing the diversity of policy pools and implementing mechanisms for dynamic policy switching.

Acknowledgment

We sincerely appreciate the insightful comments from our shepherd and the anonymous reviewers. We would like to extend our gratitude to Xuhong Zhang, Chenghui Shi, Yi Jiang, Yuyou Gan, Guang Yang and Jiawen Wan for their valuable feedback. This work was partially supported by the National Key Research and Development Program of China under grant number 2022YFB3102100.

References

- [1] Trapit Bansal, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. 2018. Emergent Complexity via Multi-Agent Competition. In *ICLR*.
- [2] Vahid Behzadan and Arslan Munir. 2017. Vulnerability of Deep Reinforcement Learning to Policy Induction Attacks. In *International Conference on Machine Learning and Data Mining in Pattern Recognition*.
- [3] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dēbiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with Large Scale Deep Reinforcement Learning. *arXiv* (2019).
- [4] Noam Brown. 2020. Equilibrium Finding for Large Adversarial Imperfect-Information Games. *PhD thesis* (2020).
- [5] Nicholas Carlini and David Wagner. 2017. Towards Evaluating the Robustness of Neural Networks. In *S&P*.
- [6] Yanjiao Chen, Zhicong Zheng, and Xueluan Gong. 2022. MARNet: Backdoor Attacks Against Cooperative Multi-Agent Reinforcement Learning. *IEEE Transactions on Dependable and Secure Computing* (2022).
- [7] Yang Dai, Oubo Ma, Longfei Zhang, Xingxing Liang, Shengchao Hu, Mengzhu Wang, Shouling Ji, Jincui Huang, and Li Shen. 2024. Is Mamba Compatible with Trajectory Optimization in Offline Reinforcement Learning? *arXiv* (2024).
- [8] Ewart J De Visser, Marieke MM Peeters, Malte F Jung, Spencer Kohn, Tyler H Shaw, Richard Pak, and Mark A Neerincx. 2020. Towards a Theory of Longitudinal Trust Calibration in Human-Robot Teams. *International Journal of Social Robotics* (2020).
- [9] DJI. [n. d.]. Robomaster. <https://www.robomaster.com/en-US>.
- [10] Linkang Du, Chen Min, Sun Mingyang, Ji Shouling, Cheng Peng, Chen Jiming, and Zhang Zhikun. 2024. ORL-AUDITOR: Dataset Auditing in Offline Deep Reinforcement Learning. In *NDSS*.
- [11] Rolando Fernandez, Derrick E Asher, Anjon Basak, Piyush K Sharma, Erin G Zaroukian, Christopher D Hsu, Michael R Dorothy, Christopher M Kroninger, Luke Frerichs, John Rogers, et al. 2021. *Multi-Agent Coordination for Strategic Maneuver with a Survey of Reinforcement Learning*. Technical Report. US Army Combat Capabilities Development Command, Army Research Laboratory.
- [12] Wei Fu, Weihua Du, Jingwei Li, Sunli Chen, Jingzhao Zhang, and Yi Wu. 2023. Iteratively Learn Diverse Strategies with State Distance Information. In *NeurIPS*.
- [13] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. 2023. Deep Whole-Body Control: Learning a Unified Policy for Manipulation and Locomotion. In *Conference on Robot Learning*.
- [14] Dibya Ghosh, Jad Rahme, Aviral Kumar, Amy Zhang, Ryan P Adams, and Sergey Levine. 2021. Why Generalization in RL is Difficult: Epistemic Pomdps and Implicit Partial Observability. In *NeurIPS*.
- [15] Adam Gleave, Michael Dennis, Cody Wild, Neel Kant, Sergey Levine, and Stuart Russell. 2020. Adversarial Policies: Attacking Deep Reinforcement Learning. In *ICLR*.

- [16] Vasudev Gohil, Hao Guo, Satwik Patnaik, and Jeyavijayan Rajendran. 2022. AT-TRITION: Attacking Static Hardware Trojan Detection Techniques Using Reinforcement Learning. In *CCS*.
- [17] Micah Goldblum, Dimitris Tsipras, Chulin Xie, Xinyun Chen, Avi Schwarzschild, Dawn Song, Aleksander Mądry, Bo Li, and Tom Goldstein. 2022. Dataset Security for Machine Learning: Data Poisoning, Backdoor Attacks, and Defenses. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *ICLR*.
- [19] Sven Gronauer and Klaus Diepold. 2022. Multi-Agent Deep Reinforcement Learning: A Survey. *Artificial Intelligence Review* (2022).
- [20] Pengjie Gu, Mengchen Zhao, Jianye Hao, and Bo An. 2022. Online Ad Hoc Teamwork under Partial Observability. In *ICLR*.
- [21] Junfeng Guo, Ang Li, Lixu Wang, and Cong Liu. 2023. PolicyCleanse: Backdoor Detection and Mitigation for Competitive Reinforcement Learning. In *ICCV*.
- [22] Wenbo Guo, Xian Wu, Sui Huang, and Xinyu Xing. 2021. Adversarial Policy Learning in Two-Player Competitive Games. In *ICML*.
- [23] Wenbo Guo, Xian Wu, Lun Wang, Xinyu Xing, and Dawn Song. 2023. PATROL: Provable Defense against Adversarial Policy in Two-player Games. In *USENIX Security*.
- [24] Ahmed T Hafez, Anthony J Marasco, Sidney N Givigi, Mohamad Iskandarani, Shahram Yousefi, and Camille Alain Rabbath. 2015. Solving Multi-UAV Dynamic Encirclement via Model Predictive Control. *IEEE Transactions on Control Systems Technology* (2015).
- [25] Ping He, Yifan Xia, Xuhong Zhang, and Shouling Ji. 2023. Efficient Query-Based Attack against ML-Based Android Malware Detection under Zero Knowledge Setting. In *CCS*.
- [26] Johannes Heinrich, Marc Lanctot, and David Silver. 2015. Fictitious Self-Play in Extensive-Form Games. In *ICML*.
- [27] Péter Miksa Hell and Péter János Varga. 2019. Drone Systems for Factory Security and Surveillance. *Interdisciplinary Description of Complex Systems: INDECS* (2019).
- [28] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. Deep Reinforcement Learning that Matters. In *AAAI*.
- [29] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. 2017. Adversarial Attacks on Neural Network Policies. *arXiv* (2017).
- [30] Yi Jiang, Chenghui Shi, Oubo Ma, Youliang Tian, and Shouling Ji. 2023. Text Laundering: Mitigating Malicious Features Through Knowledge Distillation of Large Foundation Models. In *Inscrypt*.
- [31] Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladen Koltun, and Davide Scaramuzza. 2023. Champion-Level Drone Racing Using Deep Reinforcement Learning. *Nature* (2023).
- [32] Panagiota Kiourti, Kacper Wardega, Susmit Jha, and Wenchao Li. 2020. TrojDRL: Evaluation of Backdoor Attacks on Deep Reinforcement Learning. In *57th ACM/IEEE Design Automation Conference (DAC)*.
- [33] Robert Kirk, Amy Zhang, Edward Grefenstette, and Tim Rocktäschel. 2023. A Survey of Zero-Shot Generalisation in Deep Reinforcement Learning. *Journal of Artificial Intelligence Research* (2023).
- [34] Xian Yeow Lee, Sambit Ghadai, Kai Liang Tan, Chinmay Hegde, and Soumik Sarkar. 2020. Spatiotemporally Constrained Action Space Attacks on Deep Reinforcement Learning Agents. In *AAAI*.
- [35] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-Agent Reinforcement Learning in Sequential Social Dilemmas. In *AAMAS*.
- [36] Zhuohang Li, Yi Wu, Jian Liu, Yingying Chen, and Bo Yuan. 2020. AdvPulse: Universal, Synchronization-free, and Targeted Audio Adversarial Attacks via Subsecond Perturbations. In *CCS*.
- [37] Michael L Littman. 1994. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In *Machine learning proceedings 1994*.
- [38] Xiangyu Liu, Souradip Chakraborty, Yanhao Sun, and Furong Huang. 2024. Rethinking Adversarial Policies: A Generalized Attack Formulation and Provable Defense in RL. In *ICLR*.
- [39] Xuejiao Liu, Oubo Ma, Wei Chen, Yingjie Xia, and Yuxuan Zhou. 2022. HDRS: A Hybrid Reputation System with Dynamic Update Interval for Detecting Malicious Vehicles in VANETS. *IEEE Transactions on Intelligent Transportation Systems* (2022).
- [40] 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. In *NIPS*.
- [41] Oubo Ma, Xuejiao Liu, and Yingjie Xia. 2023. ABM-V: An Adaptive Backoff Mechanism for Mitigating Broadcast Storm in VANETS. *IEEE Transactions on Vehicular Technology* (2023).
- [42] Oubo Ma, Yuwen Pu, Linkang Du, Yang Dai, Ruo Wang, Xiaolei Liu, Yingcai Wu, and Shouling Ji. 2024. SUB-PLAY: Adversarial Policies against Partially Observed Multi-Agent Reinforcement Learning Systems. *arXiv* (2024).
- [43] Yuzhe Ma, Xuezhou Zhang, Wen Sun, and Jerry Zhu. 2019. Policy Poisoning in Batch Reinforcement Learning and Control. In *NeurIPS*.
- [44] Suman Maiti, Anjana Balabhaskara, Sunandan Adhikary, Ipsita Koley, and Soumyajit Dey. 2023. Targeted Attack Synthesis for Smart Grid Vulnerability Analysis. In *CCS*.
- [45] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-Level Control through Deep Reinforcement Learning. *Nature* (2015).
- [46] Thanh Thi Nguyen, Ngoc Duy Nguyen, and Saeid Nahavandi. 2020. Deep Reinforcement Learning for Multiagent Systems: A Review of Challenges, Solutions, and Applications. *IEEE Transactions on Cybernetics* (2020).
- [47] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. 2016. The Limitations of Deep Learning in Adversarial Settings. In *EuroS&P*.
- [48] Georgios Papoudakis, Filippos Christianos, and Stefano Albrecht. 2021. Agent Modelling under Partial Observability for Deep Reinforcement Learning. In *NeurIPS*.
- [49] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *The Journal of Machine Learning Research* (2020).
- [50] Kui Ren, Qian Wang, Cong Wang, Zhan Qin, and Xiaodong Lin. 2019. The Security of Autonomous Driving: Threats, Defenses, and Future Directions. *Proc. IEEE* (2019).
- [51] Unitree Robotics. [n. d.]. New Creature of Embodied AI Unitree Go2. <https://www.unitree.com/go2>.
- [52] David Silver, Satinder Singh, Doina Precup, and Richard S Sutton. 2021. Reward is Enough. *Artificial Intelligence* (2021).
- [53] Jianwen Sun, Tianwei Zhang, Xiaofei Xie, Lei Ma, Yan Zheng, Kangjie Chen, and Yang Liu. 2020. Stealthy and Efficient Adversarial Attacks against Deep Reinforcement Learning. In *AAAI*.
- [54] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- [55] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. 2019. Neural Cleanse: Identifying and Mitigating Backdoor Attacks in Neural Networks. In *S&P*.
- [56] Jinghan Wang, Chengyu Song, and Heng Yin. 2021. Reinforcement Learning-based Hierarchical Seed Scheduling for Greybox Fuzzing. In *NDSS*.
- [57] Lun Wang, Zaynah Javed, Xian Wu, Wenbo Guo, Xinyu Xing, and Dawn Song. 2021. BACKDOORL: Backdoor Attack against Competitive Reinforcement Learning. In *IJCAI*.
- [58] Shiyong Wang, Jiafu Wan, Daqiang Zhang, Di Li, and Chunhua Zhang. 2016. Towards Smart Factory for Industry 4.0: A Self-Organized Multi-Agent System with Big Data Based Feedback and Coordination. *Computer Networks* (2016).
- [59] Tony Tong Wang, Adam Gleave, Tom Tseng, Kellin Pelrine, Nora Belrose, Joseph Miller, Michael D Dennis, Yawen Duan, Viktor Pogrebniak, Sergey Levine, et al. 2023. Adversarial Policies Beat Superhuman Go AIs. In *ICML*.
- [60] Xian Wu, Wenbo Guo, Hua Wei, and Xinyu Xing. 2021. Adversarial Policy Training against Deep Reinforcement Learning. In *USENIX Security*.
- [61] Yingjie Xia, Xuejiao Liu, Jing Ou, and Oubo Ma. 2023. RLID-V: Reinforcement Learning-Based Information Dissemination Policy Generation in VANETS. *IEEE Transactions on Intelligent Transportation Systems* (2023).
- [62] Zhiwei Xu, Yunpeng Bai, Dapeng Li, Bin Zhang, and Guoliang Fan. 2022. SIDE: State Inference for Partially Observable Cooperative Multi-Agent Reinforcement Learning. In *AAMAS*.
- [63] Yichen Yang, Jeevana Priya Inala, Osbert Bastani, Yewen Pu, Armando Solar-Lezama, and Martin Rinard. 2021. Program Synthesis Guided Reinforcement Learning for Partially Observed Environments. In *NeurIPS*.
- [64] Yuanshun Yao, Huiying Li, Haitao Zheng, and Ben Y Zhao. 2019. Latent Backdoor Attacks on Deep Neural Networks. In *CCS*.
- [65] Yang You, Liangwei Li, Baisong Guo, Weiming Wang, and Cewu Lu. 2020. Combinatorial Q-Learning for Dou Di Zhu. In *AAAI*.
- [66] Honggang Yu, Kaichen Yang, Teng Zhang, Yun-Yun Tsai, Tsung-Yi Ho, and Yier Jin. 2020. CloudLeak: Large-Scale Deep Learning Models Stealing Through Adversarial Examples. In *NDSS*.
- [67] Jiahao Yu, Wenbo Guo, Qi Qin, Gang Wang, Ting Wang, and Xinyu Xing. 2023. AIRS: Explanation for Deep Reinforcement Learning based Security Applications. In *USENIX Security*.
- [68] Heng Zhang, Jun Gu, Zhikun Zhang, Linkang Du, Yongmin Zhang, Yan Ren, Jian Zhang, and Hongran Li. 2023. Backdoor Attacks against Deep Reinforcement Learning based Traffic Signal Control Systems. *Peer-to-Peer Networking and Applications* (2023).
- [69] Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Başar. 2021. Finite-Sample Analysis for Decentralized Batch Multiagent Reinforcement Learning with Networked Agents. *IEEE Trans. Automat. Control* (2021).
- [70] Xuezhou Zhang, Yuzhe Ma, Adish Singla, and Xiaojin Zhu. 2020. Adaptive Reward-Poisoning Attacks against Reinforcement Learning. In *ICML*.
- [71] Xin Zhou, Xiangyong Wen, Zhepei Wang, Yuman Gao, Haojia Li, Qianhao Wang, Tiankai Yang, Haojian Lu, Yanjun Cao, Chao Xu, et al. 2022. Swarm of Micro Flying Robots in the Wild. *Science Robotics* (2022).