
Effective Policy Learning for Multi-Agent Online Coordination Beyond Submodular Objectives

Qixin Zhang^{1*} Yan Sun^{2*} Can Jin⁴ Xikun Zhang¹
Yao Shu³ Puning Zhao⁵ Li Shen⁵ Dacheng Tao^{1†}

¹College of Computing and Data Science, Nanyang Technological University, Singapore

²The University of Sydney ³Hong Kong University of Science and Technology (Guangzhou)

⁴Rutgers University ⁵Sun Yat-sen University

{qixin.zhang, xikun.zhang, dacheng.tao}@ntu.edu.sg; sun9899@uni.sydney.edu.au
yaoshu@hkust-gz.edu.cn; {zhaopn, shenli6}@mail.sysu.edu.cn

Abstract

In this paper, we present two effective policy learning algorithms for multi-agent online coordination (MA-OC) problem. The first one, MA-SPL, not only can achieve the optimal $(1 - \frac{c}{e})$ -approximation guarantee for the MA-OC problem with submodular objectives but also can handle the unexplored α -weakly DR-submodular and (γ, β) -weakly submodular scenarios, where c is the curvature of the investigated submodular functions, α denotes the diminishing-return (DR) ratio and the tuple (γ, β) represents the submodularity ratios. Subsequently, in order to reduce the reliance on the unknown parameters α, γ, β inherent in the MA-SPL algorithm, we further introduce the second online algorithm named MA-MPL. This MA-MPL algorithm is entirely *parameter-free* and simultaneously can maintain the same approximation ratio as the first MA-SPL algorithm. The core of our MA-SPL and MA-MPL algorithms is a novel continuous-relaxation technique termed as *policy-based continuous extension*. Compared with the well-established *multi-linear extension*, a notable advantage of this new *policy-based continuous extension* is its ability to provide a lossless rounding scheme for any set function, thereby enabling us to tackle the challenging weakly submodular objectives. Finally, extensive simulations are conducted to validate the effectiveness of our proposed algorithms.

1 Introduction

Coordinating multiple autonomous agents to cooperatively complete complex tasks in time-varying environments is a significant challenge with extensive applications in machine learning, robotics and control, including target tracking [27, 129, 130, 132], area monitoring [55, 73, 97], multi-path planning [100, 101, 102], mobile sensor placement [65, 67, 95], environmental mapping [6, 75] and task assignment [92, 5]. Critically, with the rapid advancement of self-supervised learning and the development of large Transformers with billions of parameters, large language models (LLMs) are now capable of generating factual and coherent responses to human queries [11, 25, 107, 51]. However, the mixed quality of training data can lead to the generation of undesired responses, presenting a significant challenge. Therefore, how to select a limited number of high-quality input-output examples from multi-source and massive datasets for fine-tuning LLMs remains a critical challenge—a problem that can be formulated as multi-agent collaborative data selection [69, 70, 80, 93]. Moreover, the emergence of foundation models further highlights the necessity of effective

*These authors contributed equally to this work.

†Corresponding author: Dacheng Tao.

cooperation architectures to enable multi-agent AI systems to mutually learn and adapt in dynamic environments [4, 23, 28, 60]. Motivated by these diverse scenarios, thus this paper delves into the Multi-Agent Online Coordination(MA-OC) problem, establishing a theoretical paradigm for real-world embodied systems, LLMs fine-tuning, and intelligent AI agent collaboration.

Prior to this, numerous studies have demonstrated that the utility functions associated with a wide range of multi-agent coordination scenarios often exhibit a *diminishing-return*(DR) property. Specifically, as the number of agents increases, the marginal gain in benefit will tend to decrease. For instance, in area monitoring using a fleet of unmanned aerial vehicles(UAVs), due to the overlaps of sensing ranges, the increment of the total monitored area from adding an additional UAV typically becomes less and less as the team of UAVs expands. Note that the *diminishing-return* property is also known as *submodularity* in mathematics [7, 42]. Consequently, a vast majority of research regarding multi-agent coordination focus on *submodular* objectives [45, 50, 64, 66, 78, 92, 94, 95, 114, 115, 126, 129, 131].

However, recent works [55, 56, 59, 62, 117] observed that there also exist many multi-agent coordination scenarios inducing utility functions that are *close-to-submodular*, but *not strictly submodular*. A notable example is the employment of a swarm of UAVs to track multiple moving objects. In this scenario, each UAV needs to periodically determine its moving direction and speed. Particularly, under the Kalman filter framework, [55, 56, 62] pointed out that the aforementioned trajectory selection problem of UAVs can be formulated as a multi-agent variant of the general online *weakly* submodular maximization problems. It is worth noting that the currently well-known algorithms for weakly submodular maximization are highly dependent on the discrete local search [20, 43, 53, 76, 106]. So far, however, how to extend this local search into online settings still remains an uncharted territory. Thus, almost all existing studies regarding multi-agent coordination with weakly submodular objectives choose to neglect the changing environment and focus on simple offline scenarios [55, 56, 62, 117]. In view of all this, a natural question arises:

Q1: *Is it possible to design an effective online algorithm for MA-OC problem with weakly submodular objectives?*

In addition, we also find that when the objective function is exactly *submodular*, the state-of-the-art algorithms [94, 126] for the MA-OC problem only can guarantee a sub-optimal $(\frac{1-e^{-c}}{c})$ -approximation, which mismatches the best possible $(1 - \frac{c}{e})$ -approximation established in the works [105, 54] for *single-agent* submodular maximization. Here, c denotes the curvature of the investigated submodular functions. Given this drawback, another question comes to our mind, that is,

Q2: *Is it possible to achieve the optimal $(1 - \frac{c}{e})$ -approximation ratio for MA-OC problem with submodular objectives?*

In the subsequent sections of this paper, we will provide an affirmative answer to these two questions by presenting an effective online algorithm named MA-SPL, which not only can achieve the optimal $(1 - \frac{c}{e})$ -approximation for the MA-OC problem with submodular objectives but also can address the previously unexplored weakly submodular scenarios. The core of our proposed MA-SPL algorithm is a novel continuous-relaxation framework named *policy-based continuous extension*, which can efficiently transform the discrete set function maximization problem into a solvable continuous optimization task. Furthermore, compared to the well-established *multi-linear extension* [15], a notable advantage of our proposed policy-based continuous extension is that it can provide a lossless rounding scheme for any set objective function. In contrast, all known lossless rounding schemes for the *multi-linear extension* require the set objective function to be *submodular* [15, 18, 19]. Moreover, to eliminate the dependence of both DR ratio and submodularity ratio inherent in our proposed MA-SPL algorithm, we further present a *parameter-free* online algorithm termed as MA-MPL for the MA-OC problem with general weakly submodular objectives.

In summary, we make the following contributions:

- This paper introduces an innovative continuous-relaxation technique named *policy-based continuous extension* for the general multi-agent coordination problem. Furthermore, we conduct an in-depth exploration of the differentiability, monotonicity and submodularity of our proposed policy-based continuous extension. More importantly, when the investigated set objective function is submodular or weakly submodular, we design three different surrogate functions for our policy-based continuous extension. The stationary points of these three surrogate functions can yield a better approximation guarantee than those of the original policy-based continuous extension itself.

Table 1: Comparison of the different algorithms for T -round MA-OC problem. Note that ‘Approx.’ denotes the obtained approximation ratio, ‘#Com.’ represents the number of communication, ‘#Queries’ denotes the number of queries to the set objective functions, ‘D-Regret’ denotes the dynamic regret bound, ‘Proj-free’ indicates whether the method does not require projection, ‘Para-free’ indicates whether the method does not require prior knowledge of curvature c and parameters α, γ, β , $\mathcal{P}_T$ is the deviation of maximizer sequence, τ is the spectral gap of the weight matrix, $d(G)$ is the diameter of the graph G , $\phi(\gamma, \beta) \triangleq \beta(1 - \gamma) + \gamma^2$ and $\kappa \triangleq \sum_{i=1}^n \kappa_i$.

Method	Utility	Para-free	Proj-free	Graph	Approx.	#Com.	#Queries	D-Regret
OSG [115, 50]	Submodular	✓	✓	Complete	$\left(\frac{1}{1+c}\right)$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\tilde{\mathcal{O}}(\sqrt{\mathcal{P}_T T})$
MA-OSMA [126]	Submodular	✗	✗	Connected	$\left(\frac{1-e^{-c}}{c}\right)$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$
MA-OSEA [126]	Submodular	✗	✓	Connected	$\left(\frac{1-e^{-c}}{c}\right)$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\tilde{\mathcal{O}}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$
MA-SPL	Submodular	✓	✗	Connected	$(1 - \frac{c}{e})$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$
(Algorithm 1)	α -weakly DR-Sub	✗	✗	Connected	$(1 - e^{-\alpha})$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$
	(γ, β) -weakly Sub	✗	✗	Connected	$\left(\frac{\gamma^2(1-e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)}\right)$	$\mathcal{O}(T)$	$\mathcal{O}(\kappa T)$	$\mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$
MA-MPL	α -weakly DR-Sub	✓	✓	Connected	$(1 - e^{-\alpha})$	$\mathcal{O}(T^{3/2})$	$\mathcal{O}(\kappa T^{5/2})$	$\mathcal{O}(d(G)\sqrt{\mathcal{P}_T T})$
(Algorithm 2)	(γ, β) -weakly Sub	✓	✓	Connected	$\left(\frac{\gamma^2(1-e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)}\right)$	$\mathcal{O}(T^{3/2})$	$\mathcal{O}(\kappa T^{5/2})$	$\mathcal{O}(d(G)\sqrt{\mathcal{P}_T T})$

- Building on these surrogate functions, we then propose a novel online algorithm named MA-SPL for the concerned MA-OC problem. Moreover, we also verify that, when the set objective function is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, our proposed MA-SPL can achieve an approximation ratio of $(1 - \frac{c}{e})$, $(1 - e^{-\alpha})$ or $(\frac{\gamma^2(1-e^{-\phi(\gamma, \beta)})}{\beta(1-\gamma)+\gamma^2})$ with a dynamic regret bound of $\mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right)$ to the best comparator in hindsight, respectively. Here, $\mathcal{P}_T$ is the deviation of maximizer sequence, τ is the spectral gap of the network, α denotes the DR ratio, β represents the upper submodularity ratio and γ is the lower submodularity ratio.
- To eliminate the dependence of both DR ratio and submodularity ratio of our MA-SPL algorithm, we next present a *parameter-free* online algorithm named MA-MPL for the MA-OC problem. The cornerstone of this MA-MPL algorithm is a novel inequality between our proposed policy-based continuous extension and the original set objective function. Moreover, when the objective function is monotone α -weakly DR-submodular or (γ, β) -weakly submodular, our proposed MA-MPL algorithm also can enjoy the same approximation ratio with a regret bound of $\mathcal{O}(d(G)\sqrt{\mathcal{P}_T T})$ to the best comparator in hindsight, where $d(G)$ is the diameter of the corresponding communication graph G .
- We conduct numerical experiments to verify the effectiveness of our proposed algorithms.

Related Work. Due to space limitations, the comprehensive literature review is placed in Appendix A. In particular, we present a detailed comparison of our proposed MA-SPL algorithm and MA-MPL algorithm with existing studies on multi-agent online coordination in Table 1.

2 Problem Setup

This section will provide a detailed introduction to multi-agent online coordination(MA-OC) problem.

In MA-OC problem, we generally consider a collection of n distinct agents, indexed by the set $\mathcal{N} \triangleq \{1, \dots, n\}$ and interconnected through an undirected network $G(\mathcal{N}, \mathcal{E})$. Here, $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ represents the possible communication links among agents. Additionally, each agent $i \in \mathcal{N}$ is endowed with a unique set of actions $\mathcal{V}_i \triangleq \{v_{i,1}, \dots, v_{i,\kappa_i}\}$, meaning that these action sets are mutually disjoint, i.e., $\mathcal{V}_i \cap \mathcal{V}_j = \emptyset$ for any $i, j \in \mathcal{N}$. In the process of multi-agent online coordination, at every time spot $t \in [T]$, each agent $i \in \mathcal{N}$ will separately select one action $a_i(t)$ from its own action set $\mathcal{V}_i$. After committing to these choices, the environment will reveal a utility set function f_t defined over the aggregated action space $\mathcal{V} \triangleq \cup_{i \in \mathcal{N}} \mathcal{V}_i$. Then, the agents receive the utility $f_t(\cup_{i \in \mathcal{N}} \{a_i(t)\})$. As a result, the goal of agents is to maximize their cumulative reward as much as possible. Specifically, at each time step $t \in [T]$, we need to address the following set function maximization problem in a multi-agent collaborative manner:

$$\max f_t(S), \quad \text{s.t. } S \subseteq \mathcal{V} \text{ and } |S \cap \mathcal{V}_i| \leq 1, \forall i \in \mathcal{N}. \quad (1)$$

Furthermore, in numerous practical applications regarding MA-OC problem, each agent usually has a limited perceptual range, allowing it to sense only the environmental changes in its immediate

surroundings. For example, in target tracking scenarios involving a swarm of unmanned aerial vehicles(UAVs), each UAV only can perceive the targets within its sensing radius, leaving those outside this range undetected. To model this limitation, several studies on the MA-OC problem [126, 115, 114, 95, 94] adopt a local feedback model. More specifically, after f_t is revealed, each agent $i \in \mathcal{N}$ is only permitted to query a local marginal oracle $\mathcal{Q}_t^i : \mathcal{V}_i \times 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ defined as $\mathcal{Q}_t^i(a, S) \triangleq f_t(a|S) \triangleq (f_t(S \cup \{a\}) - f_t(S))$ for any $a \in \mathcal{V}_i$ and $S \subseteq \mathcal{V}$. This implies that, at each time $t \in [T]$, agents only can receive the marginal evaluations about the actions within their individual action set, rather than the full information of f_t . We also impose this local-feedback constraint in this paper.

Before going into the details, it is crucial to emphasize that, in many real-world scenarios, there exists such an (approximate) local marginal feedback oracle $\mathcal{Q}_t^i$ for each agent $i \in \mathcal{N}$ after f_t is revealed. Typically, in addition to the decision-making process, agents often utilize various off-the-shelf learning algorithms to estimate the marginal contributions of their available actions based on the observed and collected information (See [27]). Moreover, the local information available to one agent is often insufficient for precisely assessing the actions of other agents who are not in close vicinity. Given this fact, confining each agent to the marginal estimations of the actions within its own action set also can further reduce the accumulation of learning errors.

Dynamic ρ -Regret: Generally speaking, the set function maximization problem (1) is NP-hard [83, 35], indicating that no polynomial-time algorithms can solve it optimally. Thus, this paper employs the *dynamic ρ -regret* [21, 61, 103, 115, 126, 134] to measure the performance of our proposed algorithms for MA-OC problem. In *dynamic ρ -regret*, the algorithm is compared against a sequence of local maximizers with scale parameter $\rho \in [0, 1]$, i.e., $R_\rho^*(T) \triangleq \rho \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T f_t(\cup_{i \in \mathcal{N}} \{a_i(t)\})$, where $\mathcal{A}_t^*$ is the optimal solution of problem (1) and $a_i(t)$ is the action chosen by agent i at time t .

3 Preliminaries

In this section, we introduce some basic concepts and the frequently used notations.

Notations. For any positive integer n , the symbol $[n]$ denotes the set $\{1, \dots, n\}$. $\mathbf{0}_p$ and $\mathbf{1}_p$ denote the p -dimensional vector whose all components are 0 and 1, respectively. Moreover, $\|\cdot\|_1$ and $\|\cdot\|_2$ stand for the L_1 norm and L_2 norm for vectors, respectively. We also use Δ_m to represent the standard m -dimensional simplex, that is, $\Delta_m \triangleq \{(x_1, \dots, x_m) | \sum_{i=1}^m x_i \leq 1, x_i \geq 0, \forall i \in [m]\}$.

Submodularity and Curvature. Let $\mathcal{V}$ be a finite ground set and $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ be a set function mapping any subset of $\mathcal{V}$ to a non-negative real number. Then, for any two subsets $S, T \subseteq \mathcal{V}$, we denote by $f(T|S)$ the marginal contribution of adding the elements of T to S , i.e., $f(T|S) \triangleq f(T \cup S) - f(S)$. In particular, when T is a singleton set $\{v\}$, we also use $f(v|S)$ to represent $f(\{v\}|S)$. Therefore, we say a set function f is *submodular* if and only if it satisfies the *diminishing-return property* [85, 42, 41], that is, $f(v|S) \geq f(v|T)$ for any $S \subseteq T \subseteq \mathcal{V}$ and $v \in \mathcal{V} \setminus T$. To precisely characterize the diminishing-return property, [26, 36, 105, 109] introduced the concept of *curvature* for submodular functions, which is defined as $c \triangleq 1 - \min_{S \subseteq \mathcal{V}, v \notin S} \frac{f(S \cup \{v\}) - f(S)}{f(\{v\}) - f(\emptyset)}$.

Monotonicity. A set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ is *monotone* if and only if $f(S) \leq f(T)$ for any $S \subseteq T \subseteq \mathcal{V}$. Moreover, in this paper, we suppose the set function f is *normalized*, that is, $f(\emptyset) = 0$.

Weak Submodularity. A set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ is said to be γ -weakly submodular from below for some $\gamma \in (0, 1]$ if and only if $\sum_{v \in T \setminus S} f(v|S) \geq \gamma(f(T) - f(S))$ for any two subsets $S \subseteq T \subseteq \mathcal{V}$, where γ is called as the *lower submodularity ratio* [29, 30, 20]. Similarly, we also can define the *weak submodularity from above*, that is, a set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ is β -weakly submodular from above for some $\beta \geq 1$ if and only if $\sum_{v \in T \setminus S} f(v|T - \{v\}) \leq \beta(f(T) - f(S))$, $\forall S \subseteq T \subseteq \mathcal{V}$, where β is the *upper submodularity ratio*. When a set function f is both γ -weakly submodular from below and β -weakly submodular from above, we say it is (γ, β) -weakly submodular [106].

Weak DR-submodularity. A set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ is α -weakly DR-submodular for some $\alpha \in (0, 1]$ if and only if $f(v|S) \geq \alpha f(v|T)$ for any two subsets $S \subseteq T \subseteq \mathcal{V}$ and $v \in \mathcal{V} \setminus T$. In particular, α is often called as the *diminishing-return(DR) ratio* [9, 68, 43, 48, 76]. It is worth noting that, from the previous definition of weak submodularity, we can infer that an α -weakly DR-submodular function automatically satisfies the conditions for being $(\alpha, \frac{1}{\alpha})$ -weakly submodular.

4 Policy-based Continuous-Relaxation Framework

Before presenting our proposed algorithms for the MA-OC problem, we firstly explore the offline set function maximization problem (1). In recent years, compared to discrete optimization, the field of continuous optimization has made significant advancements, yielding a broad spectrum of effective algorithmic frameworks and theoretical tools. Consequently, one promising strategy to addressing the set function maximization problem (1) is to convert it into a solvable continuous optimization problem throughout *continuous-relaxation* techniques.

A well-known continuous-relaxation framework is the *multi-linear extension* [15], which was introduced for maximizing submodular set functions. Regrettably, this relaxation framework cannot be directly applied to the general set function maximization problem (1), as most existing lossless rounding schemes for *multi-linear extension*—such as pipage rounding [1], swap rounding [18], and contention resolution [19]—rely heavily on the *submodular* assumption. Note that lossless rounding schemes refer to methods that convert the obtained continuous solution into a feasible discrete solution without any loss in terms of the objective function value. To date, how to losslessly round the *multi-linear extension* of non-submodular set functions, e.g. (γ, β) -weakly submodular and α -weakly DR-submodular functions, still remains an open question [106]. To overcome this hurdle, we will introduce an innovative continuous-relaxation technique in the subsequent part of this section.

4.1 Policy-based Continuous Extension

From the previous description about the MA-OC problem provided in Section 2, we can view the set function maximization problem (1) as a variant of multi-agent cooperative game [2, 98, 112, 120]. Inspired by this viewpoint, we naturally consider whether each agent $i \in \mathcal{N}$ can learn a policy $\pi_i \triangleq (\pi_{i,1}, \dots, \pi_{i,\kappa_i})$ over its individual action space $\mathcal{V}_i \triangleq \{v_{i,1}, \dots, v_{i,\kappa_i}\}$ and then utilizes this policy π_i to make decision, where each $\pi_{i,m}$ represents the probability of agent i taking the action $v_{i,m}$, $\forall m \in [\kappa_i]$. Based on this idea, if letting $a_i \in \mathcal{V}_i \cup \{\emptyset\}$ denote the random action chosen by each policy π_i , $\forall i \in \mathcal{N}$, then we can obtain the following policy-based continuous extension, namely,

Definition 1. If $\pi_i \triangleq (\pi_{i,1}, \dots, \pi_{i,\kappa_i}) \in \Delta_{\kappa_i}$ for any $i \in \mathcal{N}$, then the policy-based continuous extension $F_t : \prod_{i=1}^n \Delta_{\kappa_i} \rightarrow \mathbb{R}_+$ for the set function maximization problem (1) can be defined as:

$$F_t(\pi_1, \dots, \pi_n) \triangleq \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}, \forall i \in \mathcal{N}} \left(f_t \left(\bigcup_{i=1}^n \{a_i\} \right) \prod_{i=1}^n p(a_i | \pi_i) \right), \quad (2)$$

where $p(\cdot | \pi_i)$ is a probability distribution over the set $\mathcal{V}_i \cup \{\emptyset\}$, that is, $p(v_{i,m} | \pi_i) = \pi_{i,m}$, $\forall i \in [\mathcal{N}]$, $\forall m \in [\kappa_i]$ and $p(\emptyset | \pi_i) = 1 - \sum_{m=1}^{\kappa_i} \pi_{i,m}$, $\forall i \in \mathcal{N}$.

Remark 1. It is noteworthy that, in Eq.(2), with the probability $1 - \sum_{m=1}^{\kappa_i} \pi_{i,m}$, the policy π_i will not pick any action from $\mathcal{V}_i$, which means there is a possibility that no action will be chosen, i.e., $\emptyset$.

Remark 2. The definition in Eq.(2) highlights a notable advantage of our proposed policy-based continuous extension: it does not assign probabilities to any subset that violates the constraint of problem (1). As a result, for any set function f_t and any $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, throughout the Definition 1, we can easily generate a subset, i.e., $\bigcup_{i=1}^n \{a_i\}$, that adheres to the constraints of problem (1) while ensuring $\mathbb{E}(f_t(\bigcup_{i=1}^n \{a_i\})) = F_t(\pi_1, \dots, \pi_n)$. In contrast, all known lossless rounding schemes of the multi-linear extension require the function f_t to be submodular [15, 18, 19].

Remark 3. Notably, we observe that the works [96, 128] have introduced two relaxation techniques for lattice submodular and k -submodular functions, both of which are analogous to our policy-based continuous extension F_t . However, it is crucial to emphasize that there exist notable differences between our work and [96, 128]: **a)** The lattice formulation typically requires an ‘order’ relationship among different actions of the same agent. However, in our multi-agent coordination problem, we do not impose any specific order on the decisions. Consequently, the results and algorithms from [96, 128] are not directly applicable to our scenario. **b)** In addition to submodularity, our paper also considers weak submodularity and allows for varying sizes of action sets among agents.

With the policy-based continuous extension F_t defined in Eq.(2), the set function maximization problem (1) can be naturally relaxed into a continuous maximization task, i.e.,

$$\max F_t(\pi_1, \dots, \pi_n), \quad \text{s.t. } \|\pi_i\|_1 \leq 1, \pi_i \in [0, 1]^{\kappa_i}, \forall i \in \mathcal{N}. \quad (3)$$

In order to effectively tackle the policy optimization problem (3), we next investigate the properties of our proposed policy-based continuous extension F_t .

4.2 Properties of Policy-based Continuous Extension

This subsection will focus on characterizing the differentiability, monotonicity and submodularity of our proposed policy-based continuous extension. Specifically, we have the following theorem:

Theorem 1 (Proof in Appendix D). *The policy-based continuous extension $F_t : \prod_{i=1}^n \Delta_{\kappa_i} \rightarrow \mathbb{R}_+$ defined in Eq.(2) satisfies the following properties:*

1): *For any point $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, the first-order derivative of F_t at variable $\pi_{i,m}, \forall i \in \mathcal{N}, \forall m \in [\kappa_i]$, can be expressed as follows:*

$$\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \triangleq \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right),$$

where $a_j \sim \pi_j$ indicates that action a_j is randomly selected from $\mathcal{V}_j \cup \{\emptyset\}$ based on the policy π_j ;

2): *If the set function f_t is monotone, then $\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \geq 0$ for any point $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, $i \in \mathcal{N}$ and $m \in [\kappa_i]$, which means the monotonicity of f_t can be inherited by F_t ;*

3): *If f_t is α -weakly DR-submodular, then F_t is α -weakly continuous DR-submodular [58, 88] over $\prod_{i=1}^n \Delta_{\kappa_i}$, that is, for any two point $(\pi_1^a, \dots, \pi_n^a) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and $(\pi_1^b, \dots, \pi_n^b) \in \prod_{i=1}^n \Delta_{\kappa_i}$, if $\pi_i^a \leq \pi_i^b \forall i \in \mathcal{N}$, we have that $\nabla F_t(\pi_1^a, \dots, \pi_n^a) \geq \alpha \nabla F_t(\pi_1^b, \dots, \pi_n^b)$;*

4): *For any subset S within the constraint of problem (1) and any point $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, when f_t is monotone α -weakly DR-submodular, the following inequality holds:*

$$\alpha \left(f_t(S) - F_t(\pi_1, \dots, \pi_n) \right) \leq \sum_{(i,m): v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n), \quad (4)$$

where $\{(i, m) : v_{i,m} \in S\}$ denotes the set of all indices (i, m) such that $v_{i,m} \in S$. Similarly, when f_t is monotone (γ, β) -weakly submodular, we can show that

$$\left(\gamma^2 f_t(S) - (\beta(1 - \gamma) + \gamma^2) F_t(\pi_1, \dots, \pi_n) \right) \leq \sum_{(i,m): v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n). \quad (5)$$

Remark 4. Part 3) indicates that when f_t is submodular, namely $\alpha = 1$, our policy-based continuous extension F_t belongs to the well-studied continuous DR-submodular functions. However, it is important to highlight that these existing results for continuous DR-submodular maximization [57, 58, 89, 90, 111, 122] cannot be directly extended to our policy-based continuous extension F_t . This is because all of them heavily rely on the inequality $\langle \mathbf{y} - \mathbf{x}, \nabla G(\mathbf{x}) \rangle \geq G(\mathbf{y} \vee \mathbf{x}) - G(\mathbf{x})$ where G is a monotone continuous DR-submodular function and $\vee$ denotes coordinate-wise maximum operation. Note that the previous inequality requires that the domain of G must be closed under the maximum operation $\vee$. However, the domain $\prod_{i=1}^n \Delta_{\kappa_i}$ of our proposed F_t does not meet this requirement.

5 Multi-Agent Policy Learning

This section explores how to utilize the policy-based continuous extension introduced in Section 4 to address our concerned MA-OC problem. Broadly speaking, a notable advantage of continuous-relaxation techniques is that they enable the use of gradient-based methods, such as gradient ascent [8, 72, 86] and Frank-Wolfe method [10, 71], to tackle discrete optimization problems. Moreover, as is well established in the literature [3, 31, 44, 71], under mild conditions, a wide range of gradient-based algorithms can converge to the stationary points of their target objectives. Motivated by these findings, we next investigate the stationary points of our proposed policy-based continuous extension F_t .

5.1 Stationary Points and Surrogate Functions

At first, we recall the definition of stationary points for maximization problem, that is,

Definition 2. *Given a differentiable function $G : \mathcal{K} \rightarrow \mathbb{R}$ and a domain $\mathcal{C} \subseteq \mathcal{K}$, a point $\mathbf{x} \in \mathcal{C}$ is called as a stationary point for the function G over the domain $\mathcal{C}$ if and only if $\max_{\mathbf{y} \in \mathcal{C}} \langle \mathbf{y} - \mathbf{x}, \nabla G(\mathbf{x}) \rangle \leq 0$.*

Next, we examine the performance of the stationary points of our proposed policy-based continuous extension F_t relative to the maximum value of problem (1). Specifically, we have that

Theorem 2 (Proof in Appendix E). Given a set function $f_t : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$, if $(\pi_1^s, \dots, \pi_n^s)$ is a stationary point of its policy-based continuous extension F_t over the domain $\prod_{i=1}^n \Delta_{\kappa_i}$ and S^* denotes the optimal solution of the corresponding maximization problem (1), then the following inequalities hold:

- 1): When f_t is monotone submodular with curvature $c \in [0, 1]$, $F_t(\pi_1^s, \dots, \pi_n^s) \geq \left(\frac{1}{1+c}\right) f_t(S^*)$;
- 2): When f_t is monotone α -weakly DR-submodular, $F_t(\pi_1^s, \dots, \pi_n^s) \geq \left(\frac{\alpha^2}{1+\alpha^2}\right) f_t(S^*)$;
- 3): When f_t is monotone (γ, β) -weakly submodular, $F_t(\pi_1^s, \dots, \pi_n^s) \geq \left(\frac{\gamma^2}{\beta+\beta(1-\gamma)+\gamma^2}\right) f_t(S^*)$.

Remark 5. It is worth noting that, in a certain sense, the approximation guarantees established in Theorem 2 is *tight*. In Appendix E.3, we will present a simple instance of a submodular function f_t , i.e., $c = \alpha = \gamma = \beta = 1$, whose policy-based continuous extension F_t can attain a $\frac{1}{2}$ -approximation guarantee at a stationary point that is also a local maximum.

Theorem 2 indicates that when the original set function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, the direct application of gradient-based methods on our proposed policy-based continuous extension F_t can ensure an approximation ratio of $\left(\frac{1}{1+c}\right)$, $\left(\frac{\alpha^2}{1+\alpha^2}\right)$ or $\left(\frac{\gamma^2}{\beta+\beta(1-\gamma)+\gamma^2}\right)$ to the problem (1), respectively. However, as shown in [105, 36], the optimal approximation ratio for maximizing a monotone submodular function with curvature c is $(1 - \frac{c}{e})$, which significantly exceeds the previous $\left(\frac{1}{1+c}\right)$ -approximation provided by stationary points of F_t . Similarly, [53] pointed out that the optimal approximation for α -weakly DR-submodular maximization is $(1 - e^{-\alpha})$, which is also greater than the $\left(\frac{\alpha^2}{1+\alpha^2}\right)$ -approximation established in Theorem 2. Motivated by these findings, we naturally wonder whether it is possible to enhance the approximation guarantees of the stationary points of our proposed policy-based continuous extension.

Previously, the *multi-linear extension* of submodular functions also encountered a similar issue, namely, the stationary points of the *multi-linear extension* only can guarantee a *sub-optimal* $\frac{1}{2}$ -approximation [38, 58]. To overcome this drawback, prior studies [39, 40, 122, 126] constructed a novel surrogate function for the *multi-linear extension* and proved that the stationary points of this surrogate function can achieve the optimal $(1 - \frac{1}{e})$ -approximation. Inspired by this idea, we also hope to develop a surrogate function for our proposed policy-based continuous extension F_t such that it can improve the approximation ratios of the stationary points of F_t . Specifically, we have that

Theorem 3 (Proof in Appendix F). Similar to [122, 125, 126], for any given policy-based continuous extension F_t introduced in Definition 1, we consider a surrogate function $F_t^s : \prod_{i=1}^n \Delta_{\kappa_i} \rightarrow \mathbb{R}_+$ whose gradient at each point $\mathbf{x} \in \prod_{i=1}^n \Delta_{\kappa_i}$ is a weighted average of the gradient $\nabla F_t(z * \mathbf{x})$, i.e., $\nabla F_t^s(\mathbf{x}) = \int_0^1 w(z) \nabla F_t(z * \mathbf{x}) dz$ where $w(z)$ is a positive weight function over $[0, 1]$. After elaborately designing the weight function $w(z)$, we can show that:

- 1): When f_t is α -weakly DR-submodular and $w(z) = e^{\alpha(z-1)}$, for any stationary point $(\pi_1^s, \dots, \pi_n^s)$ of the surrogate objective F_t^s over $\prod_{i=1}^n \Delta_{\kappa_i}$, then we have $F_t(\pi_1^s, \dots, \pi_n^s) \geq (1 - e^{-\alpha}) f_t(S^*)$;
- 2): When f_t is (γ, β) -weakly submodular and $w(z) = e^{\phi(\gamma, \beta)(z-1)}$ where $\phi(\gamma, \beta) = \beta(1 - \gamma) + \gamma^2$, for any stationary point $(\pi_1^s, \dots, \pi_n^s)$ of the surrogate objective F_t^s over the domain $\prod_{i=1}^n \Delta_{\kappa_i}$, then we can show that $F_t(\pi_1^s, \dots, \pi_n^s) \geq \left(\frac{\gamma^2(1-e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)}\right) f_t(S^*)$;
- 3): When f_t is submodular with curvature c and $w(z) = e^{z-1}$, for any stationary point $(\pi_1^s, \dots, \pi_n^s)$ of the objective $(F_t^s + \frac{G_t}{e})$ over $\prod_{i=1}^n \Delta_{\kappa_i}$, then we have $F_t(\pi_1^s, \dots, \pi_n^s) \geq (1 - \frac{c}{e}) f_t(S^*)$, where S^* is the optimal subset of problem (1) and $G_t(\pi_1, \dots, \pi_n) \triangleq \sum_{i=1}^n \sum_{m=1}^{\kappa_i} (f_t(v_{i,m} | \mathcal{V} - \{v_{i,m}\})) \pi_{i,m}$.

Remark 6. Note that part 3) of Theorem 3 considers the objective $(F_t^s + \frac{G_t}{e})$ instead of the surrogate F_t^s alone. Here, G_t is a special linear function related to the minimum margin gains of the corresponding submodular objective f_t , that is, $f_t(a | \mathcal{V} - \{a\})$, $\forall a \in \mathcal{V}$. Furthermore, it is important to emphasize that the proof of Theorem 3 is not a parallel copy of the articles [122, 126] regarding the *multi-linear extension*. This is because the works [122, 126] also utilized the same inequality in Remark 4, which requires our domain to be closed under the operation $\vee$. However, the domain of our F_t does not satisfy this condition. As a result, new techniques are required to verify Theorem 3.

Algorithm 1: Multi-Agent Surrogate Policy Learning(MA-SPL)

Input: Time horizon T , weight function $w(z)$, action set $\mathcal{V}_i \triangleq \{v_{i,1}, \dots, v_{i,\kappa_i}\}$, weight matrix $\mathbf{W} \triangleq [w_{ij}]_{n \times n}$, parameters (α, γ, β) and step size $\eta_t, \forall t \in [T]$
// **Policy Initialization (Lines 1-2)**
1 Initialize a policy vector $(\pi_{i,1}(1), \dots, \pi_{i,n}(1))$ for each agent $i \in \mathcal{N}$;
2 Set the policy $\pi_{i,i}(1) \triangleq \frac{1}{\kappa_i} \mathbf{1}_{\kappa_i}$ and $\pi_{i,j}(1) \triangleq \mathbf{0}_{\kappa_j}$ when $i \neq j$ for any $i \in \mathcal{N}$;
3 **for** $t = 1, \dots, T$ **do**
4 **for** each agent $i \in \mathcal{N}$ **do**
5 // **Actions Sampling and Information Exchange (Lines 5-8)**
6 Compute the normalized policy $\mathbf{p}_i(t) \triangleq \frac{\pi_{i,i}(t)}{\|\pi_{i,i}(t)\|_1}$;
7 Utilize the normalized policy $\mathbf{p}_i(t)$ to sample an action $a_i(t)$ from $\mathcal{V}_i$;
8 Agent i executes the sampled action $a_i(t)$;
9 Exchange policy vector $(\pi_{i,1}(t), \dots, \pi_{i,n}(t))$ with the neighboring agent $j \in \mathcal{N}_i$;
10 // **Surrogate Gradient Estimation (Lines 9-16)**
11 Generate a random number $z_i(t)$ from r.v. $\mathcal{Z}$ where $\Pr(\mathcal{Z} \leq z) \triangleq \frac{\int_0^z w(a) da}{\int_0^1 w(a) da}, \forall z \in [0, 1]$;
12 **for** $j = 1, \dots, n$ **do**
13 Utilize the weighted policy $z_i(t) * \pi_{i,j}$ to sample an action $\tilde{a}_j(t) \in \mathcal{V}_i \cup \{\emptyset\}$;
14 Compute the random subset of actions $S_i(t) := \cup_{j \neq i, j \in \mathcal{N}} \{\tilde{a}_j(t)\}$;
15 Estimate the derivatives of the surrogate function based on Theorem 1 and 3, i.e.,
16
$$\frac{\partial F_t^s}{\partial \pi_{i,m}}(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \triangleq d_{i,m}(t) \triangleq \left(\int_{z=0}^1 w(z) dz \right) f_t(v_{i,m} | S_i(t)), \forall m \in [\kappa_i];$$

17 **if** f_t is submodular **then**
18 Update $d_{i,m}(t) = (d_{i,m}(t) + e^{-1} f_t(v_{i,m} | \mathcal{V} - \{v_{i,m}\})), \forall m \in [\kappa_i]$;
19 Aggregate the surrogate gradient estimations $(d_{i,1}(t), \dots, d_{i,\kappa_i}(t))$ as vector $\mathbf{d}_i(t)$;
20 // **Policy Update (Lines 17-19)**
21 Update $\pi_{i,j}(t+1) \triangleq \sum_{k \in \mathcal{N}_i \cup \{i\}} w_{ik} \pi_{k,j}(t)$ for any $j \neq i$ and $j \in \mathcal{N}$;
22 Compute $\mathbf{y}_{i,i}(t+1) \triangleq \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \pi_{j,i}(t) + \eta_t \mathbf{d}_i(t)$;
23 Update $\pi_{i,i}(t+1) \triangleq \arg \min_{\mathbf{b} \in \Delta_{\kappa_i}} \|\mathbf{b} - \mathbf{y}_{i,i}(t+1)\|_2$;

5.2 Multi-Agent Policy Learning via Surrogate Functions

Theorem 3 suggests that, when the original set function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, the direct application of gradient-based methods targeting stationary points to the surrogate function F_t^s or its variant $(F_t^s + \frac{G_t}{e})$ can achieve a tight approximation ratio of $(1 - \frac{c}{e})$, $(1 - e^{-\alpha})$ or $(\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2})$ to the set function maximization problem (1), respectively. Furthermore, recent study [126] developed an effective online algorithm for the multi-agent *submodular* coordination problem based on the well-studied consensus technique [84, 99, 118, 119] and the surrogate functions of *multi-linear extension* [122, 126]. Thus, in order to address the unexplored weakly submodular scenarios and simultaneously achieve the optimal $(1 - \frac{c}{e})$ -approximation for submodular settings, we naturally consider replacing the surrogate functions of *multi-linear extension* in the algorithms of [126] with those of our proposed policy-based continuous extension F_t . Motivated by this idea, we then present a general online algorithm named MA-SPL for MA-OC problem, with details presented in Algorithm 1.

In Algorithm 1, at every time step $t \in [T]$, each agent $i \in \mathcal{N}$ will maintain a local policy vector $(\pi_{i,1}(t), \dots, \pi_{i,n}(t))$. Here, $\pi_{i,i}(t)$ represents the policy being executed by agent i , while $\pi_{i,j}(t), j \neq i$ reflects agent i 's current estimate of the policy $\pi_{j,j}(t)$ being taken by other agent j . After that, each agent i selects an action $a_i(t)$ from $\mathcal{V}_i$ based on the normalized policy $\mathbf{p}_i(t) \triangleq \frac{\pi_{i,i}(t)}{\|\pi_{i,i}(t)\|_1}$ and shares its local policy vector with the neighboring agent $j \in \mathcal{N}_i$, where $\mathcal{N}_i$ denotes the neighbors

of agent i . Then, according to the results of Theorem 1 and Theorem 3, each agent i estimates the first-order partial derivatives of our proposed surrogate function F_t^s at every coordinate $\pi_{i,m}$. Specifically, agent i initially samples a random number $z_i(t)$ from the random variable $\mathcal{Z}$ with distribution $\Pr(\mathcal{Z} \leq z) \triangleq \frac{\int_0^z w(a)da}{\int_0^1 w(a)da}$, $\forall z \in [0, 1]$ and then approximates each $\frac{\partial F_t^s}{\partial \pi_{i,m}}(\pi_{i,1}(t), \dots, \pi_{i,n}(t))$ by $d_{i,m}(t) \triangleq (\int_{z=0}^1 w(z)dz) f_t(v_{i,m} | S_i(t))$ where $S_i(t)$ is a random set generated from the weighted policy vector $z_i(t) * (\pi_{i,1}(t), \dots, \pi_{i,n}(t))$. Furthermore, when the set objective function f_t is *submodular*, Algorithm 1 will further adjust each surrogate gradient estimation $d_{i,m}(t)$ with the minimum marginal contribution $f_t(v_{i,m} | \mathcal{V} - \{v_{i,m}\})$, $\forall m \in [\kappa_i]$ (See Lines 14-15). Finally, each agent i updates its policy $\pi_{i,i}(t)$ throughout a projected ascent along the direction $\mathbf{d}_i(t) \triangleq (d_{i,1}(t), \dots, d_{i,\kappa_i}(t))$.

In sharp contrast with the previous MA-OSMA and MA-OSEA algorithms in [126], the key innovation of Algorithm 1 lies in the use of our proposed policy-based continuous extension and its surrogate functions to update the policy vector in Lines 9-19, rather than relying on the well-studied *multi-linear extension* [15]. Moreover, our proposed policy-based continuous extension naturally aligns with the actions sampling process in Lines 5-7, ensuring that, for any set function f_t , the function value at the executed actions $\cup_{i \in \mathcal{N}} \{a_i(t)\}$ is at least as large as the expected function value $F_t(\pi_{1,1}, \dots, \pi_{n,n})$, namely, $\mathbb{E}(f_t(\cup_{i \in \mathcal{N}} \{a_i(t)\})) \geq F_t(\pi_{1,1}, \dots, \pi_{n,n})$. This is a significant advantage over the *multi-linear extension*, as the latter generally requires the *submodular* assumption to guarantee the losslessness of the action sampling of Lines 5-7 (See Lemma 13 in [126]).

Next, we provide the theoretical analysis for the proposed Algorithm 1. Before that, we introduce some standard assumptions about the communication graph $G(\mathcal{N}, \mathcal{E})$ and the weight matrix $\mathbf{W}$, i.e.,

Assumption 1. *The graph $G(\mathcal{N}, \mathcal{E})$ is connected. Furthermore, the weight matrix $\mathbf{W} \triangleq [w_{ij}]_{n \times n} \in \mathbb{R}_+^{n \times n}$ is symmetric and doubly stochastic, namely, $\mathbf{W}^T = \mathbf{W}$ and $\mathbf{W}\mathbf{1}_n = \mathbf{1}_n$. That is to say, $\tau < 1$ where $\tau = \max(|\lambda_2(\mathbf{W})|, |\lambda_n(\mathbf{W})|)$ is the second largest magnitude of the eigenvalues of $\mathbf{W}$ and $\lambda_i(\mathbf{W})$ denotes the i -th largest eigenvalue of matrix $\mathbf{W}$.*

With this Assumption 1, we then can get the following convergence results for our Algorithm 1, i.e.,

Theorem 4 (Proof in Appendix G). *Under Assumption 1, when each set objective function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, if we set the weight function $w(z)$ according to Theorem 3 and choose $\eta_t = \mathcal{O}(\sqrt{\frac{(1-\tau)P_T}{T}})$ where $P_T \triangleq \sum_{t=2}^T |A_t^* \Delta A_{t-1}^*|$ is the deviation of maximizer sequence and the symbol Δ denotes the symmetric difference, namely, $S \Delta T = (S \setminus T) \cup (T \setminus S)$, then our proposed Algorithm 1 can achieve a dynamic ρ -regret bound of $\mathcal{O}(\sqrt{\frac{P_T T}{1-\tau}})$, that is, $\mathbb{E}(R_\rho^*(T)) \leq \mathcal{O}(\sqrt{\frac{P_T T}{1-\tau}})$, where $\rho = (1 - \frac{c}{\epsilon})$, $\rho = (1 - e^{-\alpha})$ or $\rho = (\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2})$, respectively.*

Remark 7. *To the best of our knowledge, this is the first result that achieves the tight $(1 - \frac{c}{\epsilon})$ -approximation for the MA-OC problem with submodular objectives and simultaneously can tackle the previously unexplored (γ, β) -weakly submodular and α -weakly DR-submodular scenarios.*

Remark 8. *Note that when considering the weakly submodular objectives, Line 9 of Algorithm 1 requires prior knowledge of the ratios α, γ, β to set the weight function $w(z)$. However, in general, accurately computing these parameters will incur exponential computations. To overcome this drawback, we further present a **parameter-free** online algorithm named MA-MPL in Appendix B.*

6 Numerical Experiments

In this section, we validate the effectiveness of our proposed MA-SPL and MA-MPL algorithms via two different multi-target tracking scenarios. Especially in Figure 1(a), 1(b), 1(c), we consider a *submodular* facility-location objective function [115, 126] with different proportions of ‘Random’, ‘Adversarial’ and ‘Polyline’ targets. According to the results in Figure 1(a)-1(c), we can find that the average utility of our MA-SPL can significantly exceed the state-of-the-art MA-OSMA and MA-OSEA algorithms in [126], which is consistent with our Theorem 4. Note that the suffixes in Figure 1(a)-1(d) represent two different choices for communication graphs, where ‘c’ stands for a complete graph and ‘r’ denotes an Erdos-Renyi random graph with average degree 4. In contrast, Figure 1(d)-1(f) adopt a bayesian A-optimal criterion for the target tracking task, which will lead to a α -weakly

DR-submodular utility function [53, 55, 106]. Given the unknown DR ratio α , in Figure 1(e)-1(f), we perform a brute-force 0.1-network search to find the optimal parameter settings for MA-SPL algorithm. Subsequently, we report the best $\alpha = 0.1$ case and the worse $\alpha = 1$ scenario in Figure 1(d). Similarly, from Figure 1(d), we also find that our proposed MA-MPL and MA-SPL can substantially outperform the ‘RANDOM’ baseline, which is in accord with our theoretical findings. Particularly in ‘RANDOM’ baseline, we let each agent i randomly execute an action from its own action set $\mathcal{V}_i$. Due to space limitations, more discussions about experiment setups and results are presented in Appendix C.




Figure 1: A comparison of the running average utility across two distinct target tracking simulations.

7 Conclusion

In this paper, we primarily introduce an effective online policy learning algorithm named MA-SPL for the concerned MA-OC problem. Compared with the state-of-art MA-OSMA and MA-OSEA algorithms [126], our proposed MA-SPL not only can guarantee a tight $(1 - \frac{\epsilon}{e})$ -approximation for MA-OC problem with submodular objectives but also can address the unexplored (γ, β) -weakly submodular and α -weakly DR-submodular scenarios. Subsequently, to eliminate the dependence on the unknown DR ratio and submodularity ratio in our MA-SPL algorithm, we further present a *parameter-free* online algorithm named MA-MPL for the MA-OC problem. The key cornerstone of our MA-SPL and MA-MPL algorithms is a novel continuous relaxation termed as *policy-based continuous extension*. In sharp contrast with the well-studied *multi-linear extension* [15], a notable advantage of this new *policy-based continuous extension* is its ability to provide a lossless rounding scheme for any set function, thereby enabling us to tackle the challenging weakly submodular objective functions.

Acknowledgment

This project is supported by the National Research Foundation, Singapore, under its NRF Professorship Award No. NRF-P2024-001.

References

- [1] Alexander A Ageev and Maxim I Sviridenko. Pipe rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8:307–328, 2004.
- [2] Stefano V Albrecht, Filippas Christianos, and Lukas Schäfer. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024.
- [3] Zeyuan Allen-Zhu. How to make the gradients small stochastically: Even faster convex and nonconvex sgd. *Advances in Neural Information Processing Systems*, 31, 2018.
- [4] Anonymous. Scaling large-language-model-based multi-agent collaboration. In *Submitted to ACL Rolling Review - June 2024*, 2024. under review.
- [5] Gürdal Arslan, Jason R Marden, and Jeff S Shamma. Autonomous vehicle-target assignment: A game-theoretical formulation. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 129(5):584–596, 2007.
- [6] Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Decentralized active information acquisition: Theory and application to multi-robot slam. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4775–4782. IEEE, 2015.

- [7] Francis Bach et al. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends® in machine learning*, 6(2-3):145–373, 2013.
- [8] Dimitri Bertsekas. *Convex optimization algorithms*. Athena Scientific, 2015.
- [9] Ilija Bogunovic, Junyao Zhao, and Volkan Cevher. Robust maximization of non-submodular objectives. In *International Conference on Artificial Intelligence and Statistics*, pages 890–899. PMLR, 2018.
- [10] Gábor Braun, Alejandro Carderera, Cyrille W Combettes, Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Sebastian Pokutta. Conditional gradient methods. *arXiv preprint arXiv:2211.14103*, 2022.
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [12] Niv Buchbinder and Moran Feldman. Constrained submodular maximization via a nonsymmetric technique. *Mathematics of Operations Research*, 44(3):988–1005, 2019.
- [13] Niv Buchbinder and Moran Feldman. Constrained submodular maximization via new bounds for dr-submodular functions. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1820–1831, 2024.
- [14] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1433–1452. SIAM, 2014.
- [15] Grigori Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- [16] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical science*, pages 273–304, 1995.
- [17] Luiz Chamon and Alejandro Ribeiro. Approximate supermodularity bounds for experimental design. *Advances in Neural Information Processing Systems*, 30, 2017.
- [18] Chandra Chekuri, Jan Vondrak, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 575–584, 2010.
- [19] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- [20] Lin Chen, Moran Feldman, and Amin Karbasi. Weakly submodular maximization beyond cardinality constraints: Does randomization help greedy? In *International Conference on Machine Learning*, pages 804–813. PMLR, 2018.
- [21] Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization. In *International Conference on Artificial Intelligence and Statistics*, pages 1896–1905. PMLR, 2018.
- [22] Shengminjie Chen, Donglei Du, Wenguo Yang, Dachuan Xu, and Suixiang Gao. Continuous non-monotone dr-submodular maximization with down-closed convex constraint, 2024.
- [23] Weize Chen, Jiarui Yuan, Chen Qian, Cheng Yang, Zhiyuan Liu, and Maosong Sun. Optima: Optimizing effectiveness and efficiency for LLM-based multi-agent system. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11534–11557, Vienna, Austria, July 2025. Association for Computational Linguistics.

- [24] Yixin Chen, Ankur Nath, Chunli Peng, and Alan Kuhnle. Discretely beyond $1/e$: Guided combinatorial algorithms for submodular maximization. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 108929–108973. Curran Associates, Inc., 2024.
- [25] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [26] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete applied mathematics*, 7(3):251–274, 1984.
- [27] Micah Corah and Nathan Michael. Scalable distributed planning for multi-robot, multi-target tracking. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 437–444. IEEE, 2021.
- [28] Yufan Dang, Chen Qian, Xueheng Luo, Jingru Fan, Zihao Xie, Ruijie Shi, Weize Chen, Cheng Yang, Xiaoyin Che, Ye Tian, et al. Multi-agent collaboration via evolving orchestration. *Advances in Neural Information Processing Systems*, 2025.
- [29] Abhimanyu Das and David Kempe. Submodular meets spectral: greedy algorithms for subset selection, sparse approximation and dictionary selection. In *International Conference on Machine Learning*, pages 1057–1064, 2011.
- [30] Abhimanyu Das and David Kempe. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *Journal of Machine Learning Research*, 19(3):1–34, 2018.
- [31] Yoel Drori and Ohad Shamir. The complexity of finding stationary points with stochastic gradient descent. In *International Conference on Machine Learning*, pages 2658–2667. PMLR, 2020.
- [32] Bin Du, Kun Qian, Christian Claudel, and Dengfeng Sun. Jacobi-style iteration for distributed submodular maximization. *IEEE transactions on automatic control*, 67(9):4687–4702, 2022.
- [33] Marwa El Halabi, Suraj Srinivas, and Simon Lacoste-Julien. Data-efficient structured pruning via submodular optimization. *Advances in Neural Information Processing Systems*, 35:36613–36626, 2022.
- [34] Ethan R Elenberg, Rajiv Khanna, Alexandros G Dimakis, and Sahand Negahban. Restricted strong convexity implies weak submodularity. *The Annals of Statistics*, 46(6B):3539–3568, 2018.
- [35] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [36] Moran Feldman. Guess free maximization of submodular and linear sums. *Algorithmica*, 83(3):853–878, 2021.
- [37] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 570–579. IEEE, 2011.
- [38] Yuval Filmus and Justin Ward. The power of local search: Maximum coverage over a matroid. In *29th Symposium on Theoretical Aspects of Computer Science*, volume 14, pages 601–612. LIPIcs, 2012.
- [39] Yuval Filmus and Justin Ward. A tight combinatorial algorithm for submodular maximization subject to a matroid constraint. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 659–668. IEEE, 2012.

- [40] Yuval Filmus and Justin Ward. Monotone submodular maximization over a matroid via non-oblivious local search. *SIAM Journal on Computing*, 43(2):514–542, 2014.
- [41] Marshall L Fisher, George L Nemhauser, and Laurence A Wolsey. An analysis of approximations for maximizing submodular set functions—ii. In *Polyhedral Combinatorics*, pages 73–87. Springer, 1978.
- [42] Satoru Fujishige. *Submodular functions and optimization*. Elsevier, 2005.
- [43] Khashayar Gatmiry and Manuel Gomez-Rodriguez. Non-submodular function maximization subject to a matroid constraint, with applications. *arXiv preprint arXiv:1811.07863*, 2018.
- [44] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.
- [45] Bahman Gharesifard and Stephen L Smith. Distributed submodular maximization with limited information. *IEEE transactions on control of network systems*, 5(4):1635–1645, 2017.
- [46] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. Diverse sequential subset selection for supervised video summarization. *Advances in neural information processing systems*, 27, 2014.
- [47] Suning Gong, Qingqin Nong, Wenjing Liu, and Qizhi Fang. Parametric monotone function maximization with matroid constraints. *Journal of Global Optimization*, 75(3):833–849, November 2019.
- [48] Suning Gong, Qingqin Nong, Tao Sun, Qizhi Fang, Dingzhu Du, and Xiaoyu Shao. Maximize a monotone function with a generic submodularity ratio. *Theoretical Computer Science*, 853:16–24, 2021.
- [49] Michael Grant and Stephen Boyd. Cvx: Matlab software for disciplined convex programming, version 2.1, 2014.
- [50] David Grimsman, Mohd Shabbir Ali, Joao P Hespanha, and Jason R Marden. The impact of information in distributed submodular maximization. *IEEE Transactions on Control of Network Systems*, 6(4):1334–1343, 2018.
- [51] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [52] Kai Han, Shuang Cui, Benwei Wu, et al. Deterministic approximation for submodular maximization over a matroid in nearly linear time. *Advances in Neural Information Processing Systems*, 33:430–441, 2020.
- [53] Chris Harshaw, Moran Feldman, Justin Ward, and Amin Karbasi. Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In *International Conference on Machine Learning*, pages 2634–2643. PMLR, 2019.
- [54] Nicholas Harvey, Christopher Liaw, and Tasuku Soma. Improved algorithms for online submodular maximization via first-order regret bounds. *Advances in Neural Information Processing Systems*, 33:123–133, 2020.
- [55] Abolfazl Hashemi, Mahsa Ghasemi, Haris Vikalo, and Ufuk Topcu. Submodular observation selection and information gathering for quadratic models. In *International Conference on Machine Learning*, pages 2653–2662. PMLR, 2019.
- [56] Abolfazl Hashemi, Mahsa Ghasemi, Haris Vikalo, and Ufuk Topcu. Randomized greedy sensor selection: Leveraging weak submodularity. *IEEE Transactions on Automatic Control*, 66(1):199–212, 2020.
- [57] Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Zebang Shen. Stochastic conditional gradient++:(non) convex minimization and continuous submodular maximization. *SIAM Journal on Optimization*, 30(4):3315–3344, 2020.

- [58] Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, pages 5841–5851, 2017.
- [59] Syed Talha Jawaaid and Stephen L Smith. Submodularity and greedy algorithms in sensor scheduling for linear dynamical systems. *Automatica*, 61:282–288, 2015.
- [60] Can Jin, Hongwu Peng, Qixin Zhang, Yujin Tang, Dimitris N Metaxas, and Tong Che. Two heads are better than one: Test-time scaling of multi-agent collaborative reasoning. *arXiv preprint arXiv:2504.09772*, 2025.
- [61] Sham M Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 546–555, 2007.
- [62] Ege Can Kaya, Michael Hibbard, Takashi Tanaka, Ufuk Topcu, and Abolfazl Hashemi. Randomized greedy methods for weak submodular sensor selection with robustness considerations. *Automatica*, 171:111984, 2025.
- [63] G. Khashayar and G. R. Manuel. Non-submodular function maximization subject to a matroid constraint, with applications, 2019.
- [64] Solmaz S Kia. Submodular maximization subject to uniform and partition matroids: From theory to practical applications and distributed solutions. *arXiv preprint arXiv:2501.01071*, 2025.
- [65] Andreas Krause, Jure Leskovec, Carlos Guestrin, Jeanne VanBriesen, and Christos Faloutsos. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, 2008.
- [66] Andreas Krause, H Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust submodular observation selection. *Journal of Machine Learning Research*, 9(12), 2008.
- [67] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008.
- [68] Alan Kuhnle, J David Smith, Victoria Crawford, and My Thai. Fast maximization of non-submodular, monotonic functions on the integer lattice. In *International Conference on Machine Learning*, pages 2786–2795. PMLR, 2018.
- [69] Lilly Kumari, Shengjie Wang, Arnav Das, Tianyi Zhou, and Jeff Bilmes. An end-to-end submodular framework for data-efficient in-context learning. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3293–3308, 2024.
- [70] Lilly Kumari, Shengjie Wang, Tianyi Zhou, Nikhil Sarda, Anthony Rowe, and Jeff Bilmes. Bumblebee: Dynamic KV-cache streaming submodular summarization for infinite-context transformers. In *First Conference on Language Modeling*, 2024.
- [71] Simon Lacoste-Julien. Convergence rate of frank-wolfe for non-convex objectives. *arXiv preprint arXiv:1607.00345*, 2016.
- [72] Guanghui Lan. *First-order and stochastic optimization methods for machine learning*, volume 1. Springer, 2020.
- [73] Ruolin Li, Negar Mehr, and Roberto Horowitz. Submodularity of optimal sensor placement for traffic networks. *Transportation research part B: methodological*, 171:29–43, 2023.
- [74] Yucheng Liao, Yuanyu Wan, Chang Yao, and Mingli Song. Improved projection-free online continuous submodular maximization. *arXiv preprint arXiv:2305.18442*, 2023.
- [75] Jun Liu, Lifeng Zhou, Pratap Tokekar, and Ryan K Williams. Distributed resilient submodular action selection in adversarial environments. *IEEE Robotics and Automation Letters*, 6(3):5832–5839, 2021.

- [76] Cheng Lu, Wenguo Yang, Ruiqi Yang, and Suixiang Gao. Maximizing a non-decreasing non-submodular function subject to various types of constraints. *Journal of Global Optimization*, 83(4):727–751, August 2022.
- [77] Piyushi Manupriya, Pratik Jawanpuria, Karthik S. Gurumoorthy, SakethaNath Jagarlapudi, and Bamdev Mishra. Submodular framework for structured-sparse optimal transport. In *Forty-first International Conference on Machine Learning*, 2024.
- [78] Jason R Marden. The role of information in distributed resource allocation. *IEEE Transactions on Control of Network Systems*, 4(3):654–664, 2016.
- [79] Baharan Mirzasoleiman, Stefanie Jegelka, and Andreas Krause. Streaming non-monotone submodular maximization: Personalized video summarization on the fly. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [80] Baharan Mirzasoleiman, Amin Karbasi, Rik Sarkar, and Andreas Krause. Distributed submodular maximization. *The Journal of Machine Learning Research*, 17(1):8330–8373, 2016.
- [81] Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. Decentralized submodular maximization: Bridging discrete and continuous settings. In *International conference on machine learning*, pages 3616–3625. PMLR, 2018.
- [82] Loay Mualem and Moran Feldman. Resolving the approximability of offline and online non-monotone dr-submodular maximization over general convex sets. In *International Conference on Artificial Intelligence and Statistics*, pages 2542–2564. PMLR, 2023.
- [83] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- [84] Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.
- [85] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, 1978.
- [86] Y Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2013.
- [87] Rad Niazadeh, Tim Roughgarden, and Joshua R Wang. Optimal algorithms for continuous non-monotone submodular and dr-submodular maximization. *Journal of Machine Learning Research*, 21(125):1–31, 2020.
- [88] Mohammad Pedramfar and Vaneet Aggarwal. From linear to linearizable optimization: A novel framework with applications to stationary and non-stationary DR-submodular optimization. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [89] Mohammad Pedramfar, Yididiya Y. Nadew, Christopher John Quinn, and Vaneet Aggarwal. Unified projection-free algorithms for adversarial DR-submodular optimization. In *The Twelfth International Conference on Learning Representations*, 2024.
- [90] Mohammad Pedramfar, Christopher John Quinn, and Vaneet Aggarwal. A unified approach for maximizing continuous DR-submodular functions. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [91] Manish Prajapat, Mojmir Mutny, Melanie Zeilinger, and Andreas Krause. Submodular reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024.
- [92] Guannan Qu, Dave Brown, and Na Li. Distributed greedy algorithm for multi-agent task assignment problem with submodular utility functions. *Automatica*, 105:206–215, 2019.
- [93] Akbar Rafiey. Decomposable submodular maximization in federated setting. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21–27, 2024*.

- [94] Navid Rezazadeh and Solmaz S Kia. Distributed strategy selection: A submodular set function maximization approach. *Automatica*, 153:111000, 2023.
- [95] Alexander Robey, Arman Adibi, Brent Schlotfeldt, Hamed Hassani, and George J Pappas. Optimal algorithms for submodular maximization with distributed constraints. In *Learning for Dynamics and Control*, pages 150–162. PMLR, 2021.
- [96] Aytunc Sahin, Yatao Bian, Joachim Buhmann, and Andreas Krause. From sets to multisets: provable variational inference for probabilistic integer submodular models. In *International Conference on Machine Learning*, pages 8388–8397. PMLR, 2020.
- [97] Brent Schlotfeldt, Vasileios Tzoumas, and George J Pappas. Resilient active information acquisition with teams of robots. *IEEE Transactions on Robotics*, 38(1):244–261, 2021.
- [98] Elham Semsar-Kazerooni and Khashayar Khorasani. Multi-agent team cooperation: A game theory approach. *Automatica*, 45(10):2205–2213, 2009.
- [99] Shahin Shahrampour and Ali Jadbabaie. Distributed online optimization in dynamic environments using mirror descent. *IEEE Transactions on Automatic Control*, 63(3):714–725, 2017.
- [100] Guangyao Shi, Lifeng Zhou, and Pratap Tokekar. Robust multiple-path orienteering problem: Securing against adversarial attacks. *IEEE Transactions on Robotics*, 39(3):2060–2077, 2023.
- [101] Amarjeet Singh, Andreas Krause, Carlos Guestrin, William Kaiser, and Maxim Batalin. Efficient planning of informative paths for multiple robots. In *Proceedings of the 20th international joint conference on Artificial intelligence*, pages 2204–2211, 2007.
- [102] Amarjeet Singh, Andreas Krause, and William J Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 1843–1850, 2009.
- [103] Matthew Streeter and Daniel Golovin. An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584, 2008.
- [104] Yan Sun, Qixin Zhang, Zhiyuan Yu, Xikun Zhang, Li Shen, and Dacheng Tao. Maskpro: Linear-space probabilistic learning for strict (n: M)-sparsity on large language models. *arXiv preprint arXiv:2506.12876*, 2025.
- [105] Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4):1197–1218, 2017.
- [106] Theophile Thiery and Justin Ward. Two-sided weak submodularity for matroid constrained optimization and regression. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178, pages 3605–3634. PMLR, 02–05 Jul 2022.
- [107] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [108] Harry L Van Trees. *Detection, estimation, and modulation theory, part I: detection, estimation, and linear modulation theory*. John Wiley & Sons, 2004.
- [109] Jan Vondrák. Submodularity and curvature: The optimal algorithm (combinatorial optimization and discrete algorithms). *RIMS Kokyuroku Bessatsu B*, 23:253–266, 23:253–266, 2010.
- [110] Yuanyu Wan, Lijun Zhang, and Mingli Song. Improved dynamic regret for online frank-wolfe. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 3304–3327. PMLR, 12–15 Jul 2023.

- [111] Zongqi Wan, Jialin Zhang, Wei Chen, Xiaoming Sun, and Zhijie Zhang. Bandit multi-linear dr-submodular maximization and its applications on adversarial submodular bandits. In *International Conference on Machine Learning*, pages 35491–35524. PMLR, 2023.
- [112] Jianhong Wang, Yuan Zhang, Yunjie Gu, and Tae-Kyun Kim. Shaq: Incorporating shapley value theory into multi-agent q-learning. *Advances in Neural Information Processing Systems*, 35:5941–5954, 2022.
- [113] Jun Wang, Zaifu Zhan, Qixin Zhang, Mingquan Lin, Meijia Song, and Rui Zhang. Data-efficient biomedical in-context learning: A diversity-enhanced submodular perspective. *arXiv preprint arXiv:2508.08140*, 2025.
- [114] Zirui Xu, Xiaofeng Lin, and Vasileios Tzoumas. Bandit submodular maximization for multi-robot coordination in unpredictable and partially observable environments. In *Robotics: Science and Systems*, 2023.
- [115] Zirui Xu, Hongyu Zhou, and Vasileios Tzoumas. Online submodular coordination with bounded tracking regret: Theory, algorithm, and applications to multi-robot coordination. *IEEE Robotics and Automation Letters*, 8(4):2261–2268, 2023.
- [116] Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *International Conference on Machine Learning*, pages 449–457. PMLR, 2016.
- [117] Lintao Ye, Zhi-Wei Liu, Ming Chi, and Vijay Gupta. Maximization of nonsubmodular functions under multiple constraints with applications. *Automatica*, 155:111126, 2023.
- [118] Deming Yuan, Alexandre Proutiere, Guodong Shi, et al. Multi-agent online optimization. *Foundations and Trends® in Optimization*, 7(2-3):81–263, 2024.
- [119] Kun Yuan, Qing Ling, and Wotao Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016.
- [120] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of reinforcement learning and control*, pages 321–384, 2021.
- [121] Mingrui Zhang, Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization: From full-information to bandit feedback. In *Advances in Neural Information Processing Systems*, pages 9206–9217, 2019.
- [122] Qixin Zhang, Zengde Deng, Zaiyi Chen, Haoyuan Hu, and Yu Yang. Stochastic continuous submodular maximization: Boosting via non-oblivious function. In *International Conference on Machine Learning*, pages 26116–26134. PMLR, 2022.
- [123] Qixin Zhang, Zengde Deng, Zaiyi Chen, Kuangqi Zhou, Haoyuan Hu, and Yu Yang. Online learning for non-monotone dr-submodular maximization: From full information to bandit feedback. In *International Conference on Artificial Intelligence and Statistics*, pages 3515–3537. PMLR, 2023.
- [124] Qixin Zhang, Zengde Deng, Xiangru Jian, Zaiyi Chen, Haoyuan Hu, and Yu Yang. Communication-efficient decentralized online continuous dr-submodular maximization. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 3330–3339, 2023.
- [125] Qixin Zhang, Zongqi Wan, Zengde Deng, Zaiyi Chen, Xiaoming Sun, Jialin Zhang, and Yu Yang. Boosting gradient ascent for continuous dr-submodular maximization. *arXiv preprint arXiv:2401.08330*, 2024.
- [126] Qixin Zhang, Zongqi Wan, Yu Yang, Li Shen, and Dacheng Tao. Near-optimal online learning for multi-agent submodular coordination: Tight approximation and communication efficiency. In *The Thirteenth International Conference on Learning Representations*, 2025.

- [127] Peng Zhao and Lijun Zhang. Improved analysis for dynamic regret of strongly convex and smooth functions. In *Learning for Dynamics and Control*, pages 48–59. PMLR, 2021.
- [128] Huanjian Zhou, Lingxiao Huang, and Baoxiang Wang. Improved approximation algorithms for k -submodular maximization via multilinear extension. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [129] Lifeng Zhou and Vijay Kumar. Robust multi-robot active target tracking against sensing and communication attacks. *IEEE Transactions on Robotics*, 39(3):1768–1780, 2023.
- [130] Lifeng Zhou and Pratap Tokekar. Sensor assignment algorithms to improve observability while tracking targets. *IEEE Transactions on Robotics*, 35(5):1206–1219, 2019.
- [131] Lifeng Zhou and Pratap Tokekar. Risk-aware submodular optimization for multirobot coordination. *IEEE Transactions on Robotics*, 38(5):3064–3084, 2022.
- [132] Lifeng Zhou, Vasileios Tzoumas, George J Pappas, and Pratap Tokekar. Resilient active target tracking with multiple robots. *IEEE Robotics and Automation Letters*, 4(1):129–136, 2018.
- [133] Junlong Zhu, Qingtao Wu, Mingchuan Zhang, Ruijuan Zheng, and Keqin Li. Projection-free decentralized online learning for submodular maximization over time-varying networks. *Journal of Machine Learning Research*, 22(51):1–42, 2021.
- [134] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning*, pages 928–936, 2003.

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly state the paper's contribution and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Some limitations of our proposed MA-MPL algorithm have been discussed in the Appendix I.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We have clearly stated the required assumptions and an accompanying complete proof in the appendix for each theory result.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We present the detailed experiment setups and results in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.

- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We reveal the codes about our experiments in supplemental materials.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We present the detailed experiment setups and results in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: We repeat each experiments five times and report the average utility in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We present the detailed experiment setups and results in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.

- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research conforms, in every respect, to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Our work is primarily of theoretical nature and has no immediate societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No high risk data or model have been used.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: No existing asset has been used in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new asset is introduced in the paper

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No experiments with human subjects were conducted.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We conducted no experiments with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core methodology of our work does not involve the use of LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Appendices for “Effective Multi-Agent Online Coordination Beyond Submodular Objectives”

Table of Contents

A Literature Review	27
A.1 Submodular Maximization	27
A.2 Non-Submodular Maximization	27
A.3 Multi-Agent Submodular Maximization	28
A.4 Multilinear Extension	29
B Parameter-free Multi-Agent Policy Learning	30
C Additional Experimental Details and Results	32
C.1 Target Tracking with Facility-Location Objective Functions	32
C.2 Target Tracking under Extended Kalman-Filter Framework	34
C.3 More Details on Experimental Setups	36
D Proof of Theorem 1	36
E Proof of Theorem 2	39
E.1 Proof of Part 2) and Part 3) in Theorem 2	39
E.2 Proof of Part 1) in Theorem 2	40
E.3 A Policy-based Continuous Extension with $1/2$ -Approximation Stationary Point	42
F Proof of Theorem 3	43
F.1 Proof of Part 1) and Part 2) in Theorem 3	43
F.2 Proof of Part 3) in Theorem 3	45
G Convergence Analysis of MA-SPL Algorithm	47
H Convergence Analysis of MA-MPL Algorithm	56
I Limitation and Broader Impact	62

A Literature Review

In this section, we aim to provide a comprehensive review of the related literature.

A.1 Submodular Maximization

A set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ is called submodular if and only if it satisfies the diminishing-return property, namely, for any two subsets $A \subseteq B \subseteq \mathcal{V}$ and $v \in \mathcal{V} \setminus B$, $f(A \cup \{v\}) - f(A) \geq f(B \cup \{v\}) - f(B)$. It is widely recognized that the maximization of submodular functions is **NP**-hard, implying that no polynomial-time algorithms can solve it optimally. To overcome this challenge, [85] introduced a greedy algorithm for solving the monotone submodular maximization problem under a cardinality constraint and demonstrated that this greedy algorithm can achieve an approximation ratio of $(1 - e^{-1})$. Subsequently, [35] showed that this $(1 - e^{-1})$ -approximation guarantee is tight for monotone submodular maximization under reasonable complexity-theoretic assumptions. After that, [41] extended the greedy algorithm to the general matroid constraint. However, [41] also pointed out that under such a matroid constraint, the approximation ratio achievable by the greedy algorithm diminishes from $(1 - e^{-1})$ to $1/2$. To achieve the tight $(1 - e^{-1})$ -approximation under matroid constraint, then [15] proposed a continuous greedy algorithm for submodular functions. A key innovation of this continuous greedy algorithm is a novel continuous-relaxation technique termed as the *multi-linear extension*. Furthermore, there has been extensive research dedicated to the non-monotone submodular maximization [12, 13, 19, 22, 24, 37, 52, 82, 87, 123]

A.2 Non-Submodular Maximization

Recently, numerous studies have found that there exist various real-world applications inducing utility functions that are close to submodular, but not strictly submodular. Examples include variable selection [30, 34], data summarization [46, 79, 113], neural network pruning [33, 104], target tracking [55, 56] and sparse optimal transport [77].

Weakly Submodular Maximization. A important class of close-to-submodular functions is known as γ -weakly submodular functions. Specifically, for a set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$, it is called γ -weakly submodular if and only if, for any two subsets $A \subseteq B \subseteq \mathcal{V}$, the following inequality holds: $\sum_{v \in B \setminus A} f(v|A) \geq \gamma(f(B) - f(A))$. The γ -weakly submodular functions were originally introduced by the work [29]. Furthermore, when considering the simple cardinality constraint, [29] also show that the standard greedy algorithm can achieve an approximation ratio of $(1 - e^{-\gamma})$ for the maximization of a γ -weakly submodular function. Subsequently, [53] proved that, for any $\epsilon > 0$, no polynomial-time algorithm can achieve a $(1 - e^{-\gamma} + \epsilon)$ -approximation for the problem of maximizing a γ -weakly submodular function subject to a cardinality constraint. As for more complicated matroid constraints, [20] showed that the residual random greedy method of [14] can achieve an approximation ratio of $\frac{\gamma^2}{(1+\gamma)^2}$ for the problem of maximizing a monotone γ -weakly submodular functions. After that, [63] examined the approximation performance of the standard greedy algorithm on the γ -weakly submodular maximization problem over a matroid constraint, which indicated that the standard greedy algorithm only can offer an approximation ratio of $\frac{0.4\gamma^2}{\sqrt{r\gamma}+1}$ where r is the rank of the matroid. It is important to note that this approximation ratio $\frac{0.4\gamma^2}{\sqrt{r\gamma}+1}$ is not a constant guarantee and highly depends on the matroid rank r . In order to improve the approximation performance of these greedy-based algorithms over matroid constraints, [106] introduced the notion of upper submodularity ratio β and developed a more powerful distorted local-search algorithm for (γ, β) -weakly submodular maximization problem. More importantly, this distorted local search can guarantee a $\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2}$ -approximation for the problem of maximizing a monotone (γ, β) -weakly submodular functions subject to a matroid constraint. Note that, when the (γ, β) -weakly submodular function is closer to being submodular, i.e., $\gamma, \beta \rightarrow 1$, the approximation ratio $\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2}$ will approach the tight $(1 - 1/e)$. Conversely, when $\gamma, \beta \rightarrow 0$, the approximation ratio $\frac{\gamma^2}{(1+\gamma)^2}$ of the residual random greedy method [20] will trend to the sub-optimal $1/4$.

Weakly DR-Submodular Maximization. The other important class of non-submodular functions is known as α -weakly DR-submodular functions, where α is variously referred to as the diminishing-

return(DR) ratio [68], the generalized curvature [9] or the generic submodularity ratio [48]. The work of [63] is the first to explore the problem of maximizing a α -weakly DR-submodular maximization subject to general matroid constraints and proved that the standard greedy algorithm achieves approximation ratios of $\frac{\alpha}{1+\alpha}$ for the matroid-constrained α -weakly DR-submodular maximization. Recently, [47] also showed that the continuous greedy combined with the contention resolution scheme [19] can obtain a sub-optimal approximation ratio of $(\alpha(1-1/e)(1-e^{-\alpha}))$ for the problem of maximizing a monotone α -weakly DR-submodular functions subject to a matroid constraint. To achieve the tight $(1-e^{-\alpha})$ -approximation guarantee, [76] presented a novel distorted local-search method for the problem of maximizing a α -weakly DR-submodular maximization subject to a matroid constraint, which is motivated via the non-oblivious search [40, 39, 38].

A.3 Multi-Agent Submodular Maximization

Multi-Agent Offline Submodular Maximization. Coordinating multiple agents to collaboratively maximize a submodular function is a critical task with numerous applications in machine learning, robot planning and control. A common solution for multi-agent submodular maximization problem heavily depends on the distributed implementation of the classic sequential greedy method [41], which can ensure a $(\frac{1}{1+c})$ -approximation [26] when the submodular function possesses a curvature of $c \in [0, 1]$. However, this distributed greedy algorithm requires each agent to have full access to the decisions of all previous agents, thereby forming a *complete* directed communication graph. Subsequently, several studies [50, 45, 78] have investigated how the topology of the communication network affects the performance of the distributed greedy method. Particularly, [50] pointed out that the worst-case performance of the distributed greedy algorithm will deteriorate in proportion to the size of the largest independent group of agents in the communication graph. In order to overcome these challenges, various studies [94, 95, 32] utilized the *multi-linear extension* [15] to design algorithms for solving multi-agent submodular maximization problem. Specifically, [32] proposed a multi-agent variant of gradient ascent for multi-agent submodular maximization problem and showed that this multi-agent variant of gradient ascent can attain $\frac{1}{2}OPT - \epsilon$ over any connected communication graph where OPT is the optimal value. After that, to achieve the tight $(1-1/e)$ -approximation, [95] developed a multi-agent variant of continuous greedy method [15, 19]. However, this multi-agent continuous greedy [95] requires the exact knowledge of the *multi-linear extension*, which will lead to the exponential number of function evaluations. To tackle this drawback, [94] also proposed a stochastic variant of continuous greedy method [15, 19], which can enjoy $(\frac{1-e^{-c}}{c})OPT - \epsilon$. Here, c is the curvature of the investigated submodular objectives.

According to the hardness result in *centralized* submodular maximization [105], the optimal achievable approximation guarantee for maximizing a submodular function with curvature c is $(1 - \frac{c}{e})$. However, as previously discussed, the state-of-the-art approximation guarantee for multi-agent offline submodular maximization problems is $(\frac{1-e^{-c}}{c})$, which highly mismatches the best possible guarantee of $(1 - \frac{c}{e})$ established in [105]. Thus, if in Algorithm 1 we treat any incoming objective function $f_t, \forall t \in [T]$ as a fixed submodular set objective f , our Algorithm 1 can be naturally transformed into an approximation algorithm with the tight $(1 - \frac{c}{e})$ guarantee for multi-agent offline submodular maximization problems. Similar to [126], we compare this offline version of our Algorithm 1 with existing algorithms for multi-agent offline submodular maximization problems in Table 2.

Multi-Agent Online Submodular Maximization. [115] is the first one to explore the multi-agent submodular maximization problems in time-varying environments. Moreover, [115] also proposed an *online sequence greedy* (OSG) algorithm for multi-agent online submodular maximization problems and proved this OSG algorithm can achieve a sub-optimal $(\frac{1}{1+c})$ -approximation over a *complete* communication graph, where c is the joint curvature of the investigated submodular objectives. Concurrently, [114] extended this OSG algorithm into bandit settings. In order to improve this sub-optimal $(\frac{1}{1+c})$ -approximation guarantee and reduce the rigid requirement of a fully connected communication network of OSG algorithm, [126] utilized the non-oblivious auxiliary functions presented in [122] to design two multi-agent variants of online gradient ascent algorithm, namely, MA-OSMA algorithm and MA-OSEA algorithm, for the multi-agent online submodular maximization problem. Furthermore, [126] also showed that these two MA-OSMA and MA-OSEA algorithms can attain a regret bound of $\tilde{O}(\sqrt{\frac{PrT}{1-\tau}})$ against a $(\frac{1-e^{-c}}{c})$ -approximation to the best comparator in

Table 2: Comparison of the different algorithms for multi-agent offline submodular maximization problems. Note that ‘Approx.’ denotes the obtained approximation result, ‘#Com.’ represents the number of communication, ‘#Queries’ denotes the number of queries to the set objective functions, ‘Proj-free’ indicates whether the method does not require projection, ‘Para-free’ indicates whether the method does not require the knowledge of curvature c , OPT denotes the optimal function value, $d(G)$ is the diameter of the graph G , τ is the spectral gap of the weight matrix, $\kappa \triangleq \sum_{i=1}^n \kappa_i$ and $\alpha(G) \geq 1$ is the number of nodes in the largest independent set in graph G .

Method	Type	Para-free	Proj-free	Approx.	Graph(G)	#Com.	#Queries	Reference
Greedy Method	det.	✓	✓	$(\frac{1}{1+c})OPT$	complete	$\mathcal{O}(n)$	$\mathcal{O}(\kappa n)$	[26]
PGA	sto.	✓	✗	$\frac{1}{2}OPT - \epsilon$	connected	$\mathcal{O}(\frac{1}{(1-\tau)\epsilon^2})$	$\mathcal{O}(\frac{\kappa}{(1-\tau)\epsilon^2})$	[32]
Greedy Method	det.	✓	✓	$(\frac{1}{1+\alpha(G)})OPT$	connected	$\mathcal{O}(n)$	$\mathcal{O}(\kappa n)$	[50, 45]
CDCG	det.	✓	✓	$(1 - \frac{1}{e})OPT - \epsilon$	connected	$\mathcal{O}(\frac{1}{(1-\tau)\epsilon})$	$\mathcal{O}(\frac{\kappa 2^\kappa}{(1-\tau)\epsilon})$	[95]
Distributed-CG	sto.	✓	✓	$(\frac{1-e^{-c}}{c})OPT - \epsilon$	connected	$\mathcal{O}(\frac{d(G)}{\epsilon})$	$\tilde{\mathcal{O}}(\frac{\kappa d^3(G)}{\epsilon^3})$	[94]
MA-OSMA	sto.	✗	✗	$(\frac{1-e^{-c}}{c})OPT - \epsilon$	connected	$\mathcal{O}(\frac{1}{(1-\tau)\epsilon^2})$	$\mathcal{O}(\frac{\kappa}{(1-\tau)\epsilon^2})$	[126]
MA-OSEA	sto.	✗	✓	$(\frac{1-e^{-c}}{c})OPT - \epsilon$	connected	$\mathcal{O}(\frac{1}{(1-\tau)\epsilon^2})$	$\mathcal{O}(\frac{\kappa \log(\frac{1}{\epsilon})}{(1-\tau)\epsilon^2})$	[126]
Algorithm 1	sto.	✓	✗	$(1 - \frac{\epsilon}{e})OPT - \epsilon$	connected	$\mathcal{O}(\frac{1}{(1-\tau)\epsilon^2})$	$\mathcal{O}(\frac{\kappa}{(1-\tau)\epsilon^2})$	Theorem 4

hindsight, where $\mathcal{P}_T$ is the deviation of maximizer sequence, τ is the spectral gap of the network and c is the joint curvature of submodular objectives.

A.4 Multilinear Extension

As almost all state-of-the-art algorithms for multi-agent submodular coordination [94, 95, 126] rely on the *multi-linear extension* of [15], thus, in this subsection, we review the concept of *multi-linear extension* and compare it with our proposed policy-based continuous extension in Section 4.

At first, we define $\kappa \triangleq \sum_{i=1}^n \kappa_i$. Furthermore, from the definition of $\mathcal{V}$ in Section 2, we can know that $\kappa = |\mathcal{V}|$ such that we can re-define $\mathcal{V} \triangleq [\kappa] \triangleq \{1, \dots, \kappa\}$. Then, we can show that

Definition 3. For a set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$, we define its *multi-linear extension* as

$$G(\mathbf{x}) = \sum_{\mathcal{A} \subseteq \mathcal{V}} \left(f(\mathcal{A}) \prod_{a \in \mathcal{A}} x_a \prod_{a \notin \mathcal{A}} (1 - x_a) \right) = \mathbb{E}_{\mathcal{R} \sim \mathbf{x}} (f(\mathcal{R})), \quad (6)$$

where $\mathbf{x} = (x_1, \dots, x_\kappa) \in [0, 1]^\kappa$ and $\mathcal{R} \subseteq \mathcal{V}$ is a random set that contains each element $a \in \mathcal{V}$ independently with probability x_a and excludes it with probability $1 - x_a$. We write $\mathcal{R} \sim \mathbf{x}$ to denote that $\mathcal{R} \subseteq \mathcal{V}$ is a random set sampled according to $\mathbf{x}$.

With this *multi-linear extension*, we then can transfer the previous discrete subset selection problem (1) into a continuous maximization which aims at learning the optimal independent probability for each element $a \in \mathcal{V}$, that is,

$$\max_{\mathbf{x}} G(\mathbf{x}), \quad \text{s.t. } \mathbf{x} \in [0, 1]^\kappa \text{ and } \sum_{a \in \mathcal{V}_i} x_a \leq 1, \forall i \in \mathcal{N}. \quad (7)$$

It is important to note that, if we round any point $\mathbf{x}$ included into the constraint of problem (7) by the definition of multi-linear extension, namely, Eq.(6), there is a certain probability that the resulting subset will violate the partition constraint of the subset selection problem (1). Therefore, for *multi-linear extension*, we need to design a specific rounding methods based on the properties of the investigated set objective functions. However, current known lossless rounding schemes for multi-linear extension, such as pipage rounding [1], swap rounding [18] and contention resolution [19], are heavily dependent on the *submodular* assumption. Currently, how to losslessly round the multi-linear extension of *non-submodular* set functions, e.g. (γ, β) -weakly submodular and α -weakly DR-submodular functions, still remains an open question [106]. Conversely, our proposed policy-based continuous extension in Section 4 does not assign probabilities to any subsets that are out of the partition constraint of problem (1), which means that, for any set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ and any given policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, we can, through the Definition 1, easily produce a subset that conforms to the constraint of problem (1) without any loss in terms of the expected function value $F_t(\pi_1, \dots, \pi_n)$.

B Parameter-free Multi-Agent Policy Learning

It is important to note that, when considering the weakly submodular scenarios, the Line 9 of Algorithm 1 requires the prior knowledge of the unknown submodularity ratio (γ, β) or the diminishing-return (DR) ratio α to set the weight function $w(z)$. However, in general, accurately computing these parameters will incur exponential computations. To overcome this drawback, in this section, we explore how to design a *parameter-free* online algorithm for the MA-OC problem with (γ, β) -weakly submodular or α -weakly DR-submodular objectives.

Note that, in part 4) of Theorem 1, we establish a novel relationship between our proposed policy-based continuous extension F_t and its original set function f_t . More specifically, when the f_t is monotone α -weakly DR-submodular or (γ, β) -weakly submodular, the weighted discrepancy between any $f_t(S)$ and its policy-based continuous extension $F_t(\pi_1, \dots, \pi_n)$ can be bounded by some sum of first-order derivatives, i.e., $\sum_{(i,m):v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n)$. It is worth noting that the computation of this sum $\sum_{(i,m):v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n)$ does not rely on the knowledge of α, γ, β . Therefore, we naturally consider whether it is possible to devise a *parameter-free* online algorithm for the MA-OC problem by controlling this aforementioned sum $\sum_{(i,m):v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n)$ to narrow the gap between $F_t(\pi_1, \dots, \pi_n)$ and $f_t(S)$.

Before that, [21, 88, 121, 133] utilized the idea of meta-actions [103] to devise a Meta-Frank-Wolfe algorithm for online submodular maximization problems. The core of this Meta-Frank-Wolfe algorithm lies in iteratively optimizing an upper bound of the gap of the corresponding *multi-linear extension* G_t , i.e., $\langle y, \nabla G_t(x) \rangle \geq G_t(y) - G_t(x)$, which is very similar to our previously discussed idea. Motivated by this finding, we then leverage the idea of meta-actions [103] and the insight from part 4) of Theorem 1 to design a *parameter-free* MA-MPL algorithm for the MA-OC problem with (γ, β) -weakly submodular or α -weakly DR-submodular objectives, as shown in Algorithm 2.

Like the Algorithm 1, the core of our Algorithm 2 is primarily composed of three interleaved components, namely, Policy Update and Information Exchange (Lines 5-10), Actions Sampling (Lines 11-13), Surrogate Gradient Estimation (Lines 9-16) and Batch Gradient Estimation and Linear Oracles Update (Lines 14-21). Firstly, at every time step $t \in [T]$, each agent $i \in \mathcal{N}$ employs K online linear oracles $\{Q_i^{(1)}, \dots, Q_i^{(K)}\}$ to mimic the process of decentralized Meta-Frank-Wolfe [81, 94] for maximizing our proposed policy-based continuous extension F_t . Specifically, during every inner iteration $k \in [K]$, each agent i pushes the i -th local policy vector $\pi_{i,i}^{(k-1)}$ along the direction $\mathbf{v}_i^{(k)} \in \Delta_{\kappa_i}$ provided by the oracle $Q_i^{(k)}$, while maintaining other policies $\pi_{i,j}^{(k-1)}, j \neq i$ unchanged. Then, agent i exchanges the updated policy vector $(\mathbf{y}_{i,1}^{(k)}(t), \dots, \mathbf{y}_{i,n}^{(k)}(t))$ with the neighboring agent $j \in \mathcal{N}_i$ and simultaneously utilizes these received information to initialize the next policy $\pi_{i,m}^{(k)}$, i.e., $\pi_{i,m}^{(k)} \triangleq \max_{j \in \mathcal{N}_i \cup \{i\}} (\mathbf{y}_{j,m}^{(k)}(t)), \forall m \in \mathcal{N}$, where $\mathbf{y}_{i,i}^{(k)}(t) \triangleq \pi_{i,i}^{(k-1)}(t) + \frac{1}{K} \mathbf{v}_i^{(k)}(t)$ and $\mathbf{y}_{i,m}^{(k)}(t) \triangleq \pi_{i,m}^{(k-1)}(t), \forall m \neq i$. After completing all K iterations, each agent i normalizes the final policy vector $\pi_{i,i}^{(K)}(t)$ to select an action $a_i(t)$ from $\mathcal{V}_i$ (See Lines 11-13). Next, in Lines 14-20, each agent i uses a L -batch stochastic estimation $\mathbf{d}_i^{(k)}(t)$ to approximate the first-order partial derivatives of our proposed policy-based continuous extension F_t at every policy vector $(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t))$ and every coordinate $\pi_{i,m}, \forall m \in [\kappa_i]$. Finally, each agent i feeds the obtained L -batch gradient estimation $\mathbf{d}_i^{(k)}(t)$ back to its corresponding linear oracle $Q_i^{(k)}$.

Note that in the process of Algorithm 2, each agent $i \in \mathcal{N}$ only needs to evaluate the margin contributions of the actions within its own action set $\mathcal{V}_i$ (See Line 18).

It is important to highlight that, compared with the previous studies [21, 103, 126, 133], the innovations of our proposed MA-MPL algorithm are threefold: First, we utilize a policy-based continuous extension instead of the well-studied *multi-linear extension* to adjust the linear oracles in Lines 14-21. Second, rather than using a weight matrix to aggregate the received information, we employ the coordinate-wise *maximization operation* to update the policy vector (See Line 10). Third, to reduce communication complexity, we implement the batch gradient estimation in Lines 16-20.

Next, we provide the theoretical analysis for the proposed Algorithm 2. Before that, we introduce some standard assumptions about the linear maximization oracles $Q_i^{(k)}, \forall i \in \mathcal{N}, k \in [K]$, namely,

Algorithm 2: Multi-Agent Meta-Policy Learning(MA-MPL)

Input: Time horizon T , action set $\mathcal{V}_i \triangleq \{v_{i,1}, \dots, v_{i,\kappa_i}\}$, Communication Graph $G(\mathcal{N}, \mathcal{E})$, number of online linear maximization oracles, namely, K and batch size L
// Policy Vector and Online Linear Oracles Initialization (Lines 1-2)

- 1 Initialize a policy vector $(\pi_{i,1}^{(0)}(t), \dots, \pi_{i,n}^{(0)}(t)) \triangleq \mathbf{0}, \forall t \in [T]$ for each agent $i \in \mathcal{N}$;
- 2 Initialize K online linear oracles over Δ_{κ_i} , i.e., $\{Q_i^{(1)}, \dots, Q_i^{(K)}\}$, for each agent $i \in \mathcal{N}$;
- 3 **for** $t = 1, \dots, T$ **do**
- 4 **for** each agent $i \in \mathcal{N}$ **do**
- 5 // Policy Update and Information Exchange (Lines 5-10)
- 6 **for** $k = 1, \dots, K$ **do**
- 7 Set the shared vector $\mathbf{y}_{i,m}^{(k)}(t) \triangleq \pi_{i,m}^{(k-1)}(t)$ for any $m \neq i$ and $m \in \mathcal{N}$;
- 8 Obtain the update direction $\mathbf{v}_i^{(k)}(t) \in \Delta_{\kappa_i}$ from oracle $Q_i^{(k)}$;
- 9 Compute $\mathbf{y}_{i,i}^{(k)}(t) \triangleq \pi_{i,i}^{(k-1)}(t) + \frac{1}{K} \mathbf{v}_i^{(k)}(t)$;
- 10 Exchange the vector $(\mathbf{y}_{i,1}^{(k)}(t), \dots, \mathbf{y}_{i,n}^{(k)}(t))$ with the neighboring agent $j \in \mathcal{N}_i$;
- 11 Set $\pi_{i,m}^{(k)} \triangleq \max_{j \in \mathcal{N}_i \cup \{i\}} (\mathbf{y}_{j,m}^{(k)}(t)), \forall m \in \mathcal{N}$;
- 12 // Actions Sampling (Lines 11-13)
- 13 Compute the normalized policy $\mathbf{p}_i(t) \triangleq \pi_{i,i}^{(K)}(t) / \|\pi_{i,i}^{(K)}(t)\|_1$;
- 14 Utilize the normalized policy $\mathbf{p}_i(t)$ to sample an action $a_i(t)$ from $\mathcal{V}_i$;
- 15 Agent i executes the sampled action $a_i(t)$;
- 16 // Batch Gradient Estimation and Linear Oracles Update(Lines 14-21)
- 17 **for** $k = 1, \dots, K$ **do**
- 18 Set the gradient estimation $\mathbf{d}_i^{(k)}(t) := \mathbf{0}_{\kappa_i}$;
- 19 **for** $l = 1, \dots, L$ **do**
- 20 Utilize the policy $\pi_{i,j}^{(k)}(t)$ to sample an action $\tilde{a}_j(t) \in \mathcal{V}_i \cup \{\emptyset\}$ for any $j \in \mathcal{N}$;
- 21 Estimate the derivatives of F_t based on Theorem 1, i.e.,
$$\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \triangleq g_{i,m}^{(k)}(t) \triangleq f_t(v_{i,m} | \cup_{j \neq i} \{\tilde{a}_j(t)\}), \forall m \in [\kappa_i];$$
- 22 Aggregate the gradient estimations $(g_{i,1}^{(k)}(t), \dots, g_{i,\kappa_i}^{(k)}(t))$ as vector $\mathbf{g}_i^{(k)}(t)$;
- 23 Update $\mathbf{d}_i^{(k)}(t) \triangleq \mathbf{d}_i^{(k)}(t) + \frac{1}{L} \mathbf{g}_i^{(k)}(t)$;
- 24 Feed back the batch gradient estimation $\mathbf{d}_i^{(k)}(t)$ to the linear oracle $Q_i^{(k)}$;

Assumption 2. Each linear maximization oracle $Q_i^{(k)}$ can achieve a dynamic regret of $\mathcal{O}(\sqrt{V_T T})$ where V_T is the variation of any feasible path $(\mathbf{u}_1, \dots, \mathbf{u}_T) \in \prod_{t=1}^T \Delta_{\kappa_i}$, that is to say, $V_T \triangleq \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ for any path $(\mathbf{u}_1, \dots, \mathbf{u}_T) \in \prod_{t=1}^T \Delta_{\kappa_i}$

Remark 9. It is worth noting that there exist several effective and efficient algorithms that can achieve a regret bound of $\mathcal{O}(\sqrt{V_T T})$ for online linear maximization problem, for instance, online Frank-Wolfe [110] and online gradient ascent [116, 127, 134].

Theorem 5 (Proof provided in Appendix H). Under Assumption 2, when the communication graph $G(\mathcal{N}, \mathcal{E})$ is connected and each set function f_t is monotone α -weakly DR-submodular or (γ, β) -weakly submodular, if we set $L = \mathcal{O}(T)$ and $K = \mathcal{O}(\sqrt{T})$, our proposed MA-MPL algorithm can achieve a dynamic ρ -regret bound of $\mathcal{O}(d(G) \sqrt{P_T T})$, that is, $\mathbb{E}(R_\rho^*(T)) \leq \mathcal{O}(d(G) \sqrt{P_T T})$, where $\rho = (1 - e^{-\alpha})$ or $\rho = (\frac{\gamma^2(1 - e^{-(\beta(1-\gamma) + \gamma^2)})}{\beta(1-\gamma) + \gamma^2})$, respectively.

Remark 10. Note that, in Theorem 5, $d(G)$ represents the diameter of graph $G(\mathcal{N}, \mathcal{E})$, i.e., the length of the shortest path between the most distanced nodes. Moreover, $P_T \triangleq \sum_{t=2}^T |\mathcal{A}_t^* \triangle \mathcal{A}_{t-1}^*|$ is the deviation of maximizer sequence and $\triangle$ denotes the symmetric difference.

C Additional Experimental Details and Results

In this section, we test the effectiveness of our proposed Algorithm 1 and Algorithm 2 in two different multi-target tracking scenarios.

C.1 Target Tracking with Facility-Location Objective Functions

In line with the prior studies [115, 126], we consider a 2-dimensional plane where 20 agents are deployed to track 30 moving targets over a duration of 25 seconds, subdivided into $T \triangleq 1250$ discrete iterations. At every iteration, agents must determine their movement direction from “up”, “down”, “left”, “right”, or “diagonally”. Simultaneously, agents also need to adjust their speeds from a predefined set of 5, 10, or 15 units/s. As for targets, we categorize them into three different types: the unpredictable ‘Random’ and the structured ‘Polyline’ as well as the challenging ‘Adversarial’. More specifically, at every iteration, ‘Random’ target will change its movement angle θ randomly from $[0, 2\pi]$ and steers at a random speed between 5 units/s and 10 units/s. In contrast, the ‘Polyline’ target will maintain its trajectory and only behaves like the ‘Random’ target at the specific $\{0, \lfloor \frac{T}{k} \rfloor, 2\lfloor \frac{T}{k} \rfloor, \dots, (k-1)\lfloor \frac{T}{k} \rfloor\}$ -th iteration where T is the predefined total iterations and k is a random number from $\{1, 2, 4\}$. Regarding the ‘Adversarial’ target, it mimics the ‘Random’ targets when all agents are beyond a 20 units. Nevertheless, upon detecting an agent within the 20-unit range, the ‘Adversarial’ target will evade at a speed of 15 units/s for one second, pointing to the direction that maximizes the average distance of all agents.

Generally speaking, every motion of any agent can be characterized by three key parameters, namely, its movement angle θ , speed s and the unique identifier i . With all these three parameters, then the action set $\mathcal{V}_i$ available to each agent $i \in [20]$ can be mathematically represented as:

$$\mathcal{V}_i = \{(\theta, s, i) : s \in \{5, 10, 15\} \text{ units/s}, \theta \in \{\frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi, \dots, 2\pi\}\}, \forall i \in [20],$$

where each tuple (θ, s, i) encodes a specific action of agent i , that is, it will move at a speed of s in the direction of θ . Furthermore, at each time step $t \in [T]$, we denote the location of each target $j \in [30]$ as $o_j(t)$. Similarly, we also utilize the symbol $o_{(\theta, s, i)}(t)$ to represent the new position of agent i after moving from its previous location at time $(t-1)$ with a movement angle θ and speed s .

To enhance the tracking quality, a common strategy is to minimize the distances between agents and targets. Inspired by this idea, many studies [27, 114, 115, 126] naturally consider the following facility-location objective function for agents at every iteration $t \in [T]$, that is,

$$f_t(S) \triangleq \sum_{j=1}^{30} \max_{(\theta, s, i) \in S} \frac{1}{\|o_{(\theta, s, i)}(t) - o_j(t)\|_2},$$

where $\|o_{(\theta, s, i)}(t) - o_j(t)\|_2$ represents the Euclidean distance between the location $o_j(t)$ of target j and the new position $o_{(\theta, s, i)}(t)$ after agent i executing the action (θ, s, i) and S is a subset of the ground action set $\mathcal{V} \triangleq \cup_{i=1}^n \mathcal{V}_i$. It is important to note that the larger the value of $\frac{1}{\|o_{(\theta, s, i)}(t) - o_j(t)\|_2}$ becomes, the closer the action (θ, s, i) drives agent i to the target j .

Considering this facility-location utility set function f_t and the truth that each agent only can execute one decision from $\mathcal{V}_i$ at every time $t \in [T]$, then we can easily transform the aforementioned multi-target tracking problem as a multi-agent online coordination(MA-OC) problem introduced in Section 2. Particularly, at each time $t \in [T]$, we need to tackle the following facility-location utility set function maximization problem in a multi-agent manner, namely,

$$\max f_t(S), \quad \text{s.t. } S \subseteq \mathcal{V} \text{ and } |S \cap \mathcal{V}_i| = 1, \forall i \in \mathcal{N}. \quad (8)$$

Furthermore, numerous studies [91, 115, 126] have verified that this facility-location objective function f_t is *monotone submodular*. As a result, the problem (8) can be equivalently reformulated as the problem (1) in Section 2, i.e.,

$$\max f_t(S), \quad \text{s.t. } S \subseteq \mathcal{V} \text{ and } |S \cap \mathcal{V}_i| \leq 1, \forall i \in \mathcal{N}.$$




Figure 2: Comparison of the running average utility and the running average number of targets within 1 unit as well as the running average distance of Top-5 nearest targets of OSG, MA-OSMA, MA-OSEA, MA-OSEA with our proposed MA-SPL and MA-MPL algorithms on the multi-target tracking scenario with 'Random': 'Adversarial': 'Polyline' = 8:1:1.




Figure 3: Comparison of the running average utility and the running average number of targets within 1 unit as well as the running average distance of Top-5 nearest targets of OSG, MA-OSMA, MA-OSEA, MA-OSEA with our proposed MA-SPL and MA-MPL algorithms on the multi-target tracking scenario with 'Random': 'Adversarial': 'Polyline' = 6:3:1.




Figure 4: Comparison of the running average utility and the running average number of targets within 1 unit as well as the running average distance of Top-5 nearest targets of OSG, MA-OSMA, MA-OSEA, MA-OSEA with our proposed MA-SPL and MA-MPL algorithms on the multi-target tracking scenario with 'Random': 'Adversarial': 'Polyline' = 4:5:1.

In simulations, we initialize the starting positions of all agents and targets randomly within 20-unit radius circle centered at the origin. Furthermore, we consider different proportions of 'Random', 'Polyline', and 'Adversarial' targets. Specifically, we set the proportions of targets as 'Random': 'Adversarial': 'Polyline' = 8:1:1 in Figure 2 and 6:3:1 in Figure 3 as well as 4:5:1 in Figure 4. Like [126], we also use the suffixes to represent two different choices for communication graphs, where 'c' stands for a complete graph and 'r' denotes an Erdos-Renyi random graph with average degree 4. According to the results in Figure 2(a), Figure 3(a) and Figure 4(a), we can find that the average utility of our proposed MA-SPL algorithm can significantly exceed the state-of-the-art MA-OSMA and MA-OSEA algorithms in [126], which is consistent with our Theorem 4, that is to say, our proposed MA-SPL algorithm can achieve a tight $(1 - \frac{\epsilon}{e})$ -approximation guarantee

for submodular objectives, while the MA-OSMA and MA-OSEA in [126] only can guarantee a sub-optimal $(\frac{1-e^{-c}}{c})$ -approximation. Like [126], in Figure 2(b), Figure 3(b) and Figure 4(b), we also compare the average distance from agents to their closest five targets of our proposed MA-SPL and MA-MPL algorithms against OSG [115], MA-OSMA and MA-OSEA algorithms. Similarly, we observe that our proposed MA-SPL can effectively reduce the average distance between agents and targets. Furthermore, from Figure 2(c), Figure 3(c), and Figure 4(c), we also can infer that, compared with OSG, MA-OSMA, and MA-OSEA algorithms, our proposed MA-SPL can make more targets gather within 1 unit of agents. It is worth noting that as the proportion of ‘Adversarial’ targets increases, both the maximum running average utility and the maximum running average number of targets within 1 unit exhibit a downward trend.

C.2 Target Tracking under Extended Kalman-Filter Framework

In previous Appendix C.1, we considered a simplified multi-target tracking model, that is, we assume that the monitoring quality of each moving target j only depends on its nearest agent. However, in many real-world scenarios, due to agents’ different sensing capabilities and varying observation angles, relying *solely* on the information collected by the nearest agent for accurately tracking the state of each target is unrealistic. Instead, we generally need to aggregate the data from multiple agents to comprehensively reconstruct the targets’ behaviors.

Recently, in order to obtain an accurate estimation of every target’s location, [55, 56] employed an extend Kalman-filter(EKF) framework to process the observations of multiple distinct agents. Specifically, let us consider a target tracking task using a swarm of agents(UAVs) equipped with GPS and radar systems. In this scenario, at each time $t \in [T]$, each agent i can measure the range to each target j throughout its radar system and also can obtain its own position information from GPS. Furthermore, like the work [55], we also assume that the range measurements of the radar systems follow a quadratic model, namely,

$$r_{(\theta,s,i) \rightarrow j}(t) = \frac{1}{2} \|o_{(\theta,s,i)}(t) - o_j(t)\|_2^2 + \xi_{(\theta,s,i) \rightarrow j}(t), \quad (9)$$

where the symbol $r_{(\theta,s,i) \rightarrow j}(t)$ denotes the range measurement of agent i to target j after agent i executes the action (θ, s, i) , $o_j(t)$ is the location of target j , $\xi_{(\theta,s,i) \rightarrow j}(t)$ is the additive independent noise and the symbol $o_{(\theta,s,i)}(t)$ represents the new position of agent i after moving from its previous location at time $t - 1$ with a movement angle θ and speed s .

Note that when the new position $o_{(\theta,s,i)}(t)$ and the range measurement $r_{(\theta,s,i) \rightarrow j}(t)$ are known, we can view Equation (9) as a random quadratic experiment of the unknown location $o_j(t)$. In other words, different action (θ, s, i) can lead to distinct random quadratic observation of the unknown parameter $o_j(t)$. Inspired by this perspective, we can formulated the action selection problem in multi-target tracking task as an experimental design problem, which aims at selecting a feasible subset from the whole collection of experiments $\{r_{(\theta,s,i) \rightarrow j}(t) = \frac{1}{2} \|o_{(\theta,s,i)}(t) - o_j(t)\|_2^2 + \xi_{(\theta,s,i) \rightarrow j}(t) \mid (\theta, s, i) \in \mathcal{V}\}$ such that the measurements of these selected sub-experiments can accurately estimate the location of every target. Particularly, under the classical Van Trees’ inequality [108], the work [55] established a lower bound for the covariance matrix associated with the EKF estimator of the location $o_j(t)$. Specifically, at time step $t \in [T]$, if we consider the sub-experiment $S \subseteq \mathcal{V}$, then the lower-bound matrix $\mathbf{B}_S^j(t)$ in the Van Trees’ inequality for the EKF estimate of the unknown location $o_j(t)$ can be expressed as (Please refer to Theorem 2 in [55]):

$$\mathbf{B}_S^j(t) = \left(\sum_{(\theta,s,i) \in S} \frac{1}{\sigma_{(\theta,s,i) \rightarrow j}^2(t)} \left(\mathbf{P} + \mathbf{z}_{(\theta,s,i) \rightarrow j} \mathbf{z}_{(\theta,s,i) \rightarrow j}^T \right) + \mathbf{I}_j(t) \right)^{-1}, \quad (10)$$

where $\mathbf{z}_{(\theta,s,i) \rightarrow j} \triangleq o_{(\theta,s,i)}(t) - o_j(t-1) - \mathbb{E}(o_j(t) - o_j(t-1))$, $\sigma_{(\theta,s,i) \rightarrow j}^2(t) \triangleq \text{Var}(\xi_{(\theta,s,i) \rightarrow j}(t))$, $\mathbf{P} \triangleq \text{Cov}(o_j(t) - o_j(t-1))$ and $\mathbf{I}_j(t)$ is the Fisher information matrix associated with the normalized random gap $(o_j(t) - o_j(t-1) - \mathbb{E}(o_j(t) - o_j(t-1)))$. Note that in Equation (10), the tuple (θ, s, i) represents not only a selected action but also an observation experiment.

With this lower-bound covariance matrix $\mathbf{B}_S^j(t)$, we then can utilize various so-called alphabetical criteria [16, 17] to design a utility function for selecting a superior action set such that the resulting




Figure 5: In Figure 5(a), we compare the running average utility of our proposed MA-SPL and MA-MPL algorithms with that of the ‘RANDOM’ baseline in a multi-target tracking simulation with A-optimal objective functions. In ‘RANDOM’ baseline, each agent i randomly selects an action from its own action set $\mathcal{V}_i$ at every iteration. As for Figure 5(b) and Figure 5(c), we illustrate the impact of different DR ratios $\alpha \in \{0.1, 0.2, \dots, 1\}$ on MA-SPL algorithm. Particularly in Figure 5(b), we considers a complete communication graph among agents, whereas we employs an Erdos-Renyi graph with an average degree of 4 in Figure 5(c). Note that Figure 5(a) only shows the results for the best scenario ($\alpha = 0.1$) and the worst-case scenario ($\alpha = 1$) for our proposed MA-SPL algorithm.

EKF estimation can accurately approximate the location of each target. One commonly used strategy is to employ the A-optimality, i.e., we consider minimizing the trace of the lower-bound covariance matrix $\mathbf{B}_S^j(t)$ or equivalently maximize:

$$f_t(S) = \sum_{j=1}^{30} \left(\text{Tr}(\mathbf{I}_j^{-1}(t)) - \text{Tr}(\mathbf{B}_S^j(t)) \right) \quad (11)$$

where “Tr” is the trace of matrix and we also consider 30 moving targets in simulation.

Thus, under the A-optimality criterion, we can model the agents’ action selection task as a special instance of the multi-agent online coordination problem with the set objective function f_t defined in Equation (11). Furthermore, according to recent studies [53, 55, 106], we can show that the utility function f_t in Equation (11) is monotone α -weakly DR-submodular and (γ, β) -weakly submodular (See Theorem 6 in [55] and Theorem C.2 in [106]).

In our simulations, to simplify the computation of the lower-bound matrix $\mathbf{B}_S^j(t)$, we model the movement of each target as a two-dimensional Brownian motion. Specifically, we set $o_j(t) \triangleq o_j(t-1) + 0.02 * \mathcal{N}(\mathbf{0}_2, \mathbf{I}_2)$ where $\mathbf{I}_2$ is the 2-dimensional identity matrix. Moreover, we assume that the noise $\xi_{(\theta, s, i) \rightarrow j}$ follows an independent normal distribution, i.e., $\xi_{(\theta, s, i) \rightarrow j} \sim \mathcal{N}(0, 0.01)$. As for agents, at every iteration $t \in [T]$, we adjust their speeds from a set of 2, 7, or 12 units/s and simultaneously change their movement directions from “up”, “down”, “left”, “right”, or “diagonally”. As a result, the action set $\mathcal{V}_i$ available to each agent $i \in [20]$ can be mathematically formulated as:

$$\mathcal{V}_i = \{(\theta, s, i) : s \in \{2, 7, 12\} \text{ units/s}, \theta \in \{\frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi, \dots, 2\pi\}\}, \forall i \in [20],$$

where θ denotes the movement angle, s is the speed and i represents the unique identifier.

Given the unknown diminishing-return(DR) ratio α of our investigated set function f_t , we test 10 different configurations from $\alpha = 0.1$ to $\alpha = 1$ for our proposed MA-SPL and then present the results in Figure 5(b) and Figure 5(c). Particularly in Figure 5(b), we considers a complete communication graph among agents, whereas we employs an Erdos-Renyi graph with an average degree of 4 in Figure 5(c). Subsequently, we compare the results of the best-case scenario ($\alpha = 0.1$) and the worst-case scenario ($\alpha = 1$) of our proposed MA-SPL algorithm with the *parameter-free* MA-MPL algorithm and ‘RANDOM’ baseline in Figure 5(a). According to the curves in Figure 5(a), we can find that the running average utility of our proposed MA-MPL and MA-SPL algorithms can significantly exceed the baseline ‘RANDOM’ algorithm, which is in accord with our Theorem 4 and Theorem 5. Moreover, we also find that the *parameter-free* MA-MPL algorithm can effectively outperform the MA-SPL algorithm associated with a 0.1-network search. It is worth noting that no previous works explore the MA-OC problem with weakly submodular objectives. Thus, we adopt the ‘RANDOM’ algorithm as a baseline in Figure 5(a). In ‘RANDOM’, each agent i randomly selects an action from

its own action set $\mathcal{V}_i$. Like the previous Appendix C.1, we also use the suffixes to represent two different choices for communication graphs in Figure 5(a), where ‘c’ stands for a complete graph and ‘r’ denotes an Erdos-Renyi random graph with average degree 4.

C.3 More Details on Experimental Setups

This subsection discusses some additional details about our experiments.

At first, we describe the parameter configurations about our proposed ‘MA-SPL’, ‘MA-MPL’, ‘MA-OSMA’ and ‘MA-OSEA’. Specifically, we make the following setups:

- In ‘MA-SPL’, namely, the Algorithm 1, we set the step size $\eta_t = \frac{1}{\sqrt{T}}$ and employ a 10 batch of stochastic estimation to approximate the surrogate gradient of our proposed policy-based continuous extension in Section 5.1. As for the projection in Line 19 of Algorithm 1, we utilize the CVX optimization solver [49].
- In ‘MA-MPL’, namely, the Algorithm 2 of Appendix B, we set the batch size $L = 10$ and the number of oracles $K = 15$. As for the online linear maximization oracle, we utilize the online gradient ascent algorithm [116, 134] with step size $\eta = \mathcal{O}(\frac{1}{\sqrt{T}})$.
- In ‘MA-OSMA’, namely Algorithm 1 of [126], we consider the Euclidean distance with $\phi(\mathbf{x}) = \frac{\|\mathbf{x}\|_2^2}{2}$, set the step size $\eta_t = \frac{1}{\sqrt{T}}$ and implement a 10 batch of stochastic estimation to approximate the surrogate gradient of *multi-linear extension*. Similarly, we also use the CVX solver for projection operations.
- In ‘MA-OSEA’, namely Algorithm 2 of [126], we also set the step size $\eta_t = \frac{1}{\sqrt{T}}$ and consider the mixing parameter $\gamma = 1/T^{1.5}$.

As for the communication graph G , we consider two different setups:

- ‘Complete graph’ where we set the weight $w_{ij} = \frac{1}{n}, \forall i, j \in [n]$ where $n = |\mathcal{N}|$ is the number of agents.
- ‘Erdos-Renyi random graph with average degree 4’ where if the edge (i, j) is an edge of the graph, we set $w_{ij} \triangleq 1/(1 + \max(d_i, d_j))$ where d_i is the degree of agent $i \in \mathcal{N}$ and when (i, j) is not an edge of the graph and $i \neq j$, we consider $w_{ij} = 0$. Finally, we set $w_{ii} \triangleq 1 - \sum_{j \in \mathcal{N}_i} w_{ij}$ where $\mathcal{N}_i$ is the neighboring nodes of agent i .

Furthermore, for all curves related to MA-SPL, MA-MPL, MA-OSMA, MA-OSEA and OSG, we repeat these algorithms **five runs** and then report the average result. Note that, in Figure 2(a), Figure 3(a), Figure 4(a) and Figure 5(a), the running average utility at any time t is defined as $(\sum_{t_1 \in [t]} \frac{f_t(\cup_{i \in [n]} \{a_i(t_1)\})}{t})$ where $a_i(t_1)$ is the action chosen by agent $i \in [n]$ at time t_1 .

D Proof of Theorem 1

In this section, we prove the Theorem 1.

1): From Definition 1, we have that

$$\begin{aligned} & \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \\ &= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \in \mathcal{N}} \left(f_t(\cup_{j=1}^n \{a_j\}) \frac{\partial \left(\prod_{j=1}^n p(a_j | \pi_j) \right)}{\partial \pi_{i,m}} \right) \\ &= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \in \mathcal{N}} \left(f_t(\cup_{j=1}^n \{a_j\}) \frac{\partial p(a_i | \pi_i)}{\partial \pi_{i,m}} \prod_{j \neq i, j \in \mathcal{N}} p(a_j | \pi_j) \right). \end{aligned} \quad (12)$$

Note that $p(v_{i,m} | \pi_i) = \pi_{i,m}, \forall i \in [n], \forall m \in [\kappa_i]$ and $p(\emptyset | \pi_i) = 1 - \sum_{m=1}^{\kappa_i} \pi_{i,m}, \forall i \in \mathcal{N}$. Therefore, we have $\frac{\partial p(a_i | \pi_i)}{\partial \pi_{i,m}} = 1$ when $a_i = v_{i,m}$, $\frac{\partial p(a_i | \pi_i)}{\partial \pi_{i,m}} = -1$ when $a_i = \emptyset$ and $\frac{\partial p(a_i | \pi_i)}{\partial \pi_{i,m}} = 0$ when

$a_i \notin \{v_{i,m}, \emptyset\}$. Then, according to Eq.(12), we can get the following equality:

$$\begin{aligned}
& \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \\
&= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \in \mathcal{N}} \left(f_t(\cup_{j=1}^n \{a_j\}) \frac{\partial p(a_i|\pi_i)}{\partial \pi_{i,m}} \prod_{j \neq i, j \in \mathcal{N}} p(a_j|\pi_j) \right) \\
&= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(\sum_{a_i \in \mathcal{V}_i} \left(f_t(\cup_{j=1}^n \{a_j\}) \frac{\partial p(a_i|\pi_i)}{\partial \pi_{i,m}} \prod_{j \neq i, j \in \mathcal{N}} p(a_j|\pi_j) \right) \right) \\
&= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(\left(f_t(\{v_{i,m}\} \cup (\cup_{j \neq i} \{a_j\})) - f_t(\cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \prod_{j \neq i, j \in \mathcal{N}} p(a_j|\pi_j) \right) \\
&= \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \prod_{j \neq i, j \in \mathcal{N}} p(a_j|\pi_j) \right) \\
&= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right), \tag{13}
\end{aligned}$$

where $a_j \sim \pi_j$ indicates that action a_j is randomly selected from $\mathcal{V}_j \cup \{\emptyset\}$ based on the policy π_j .

2): When f_t is monotone, we can know that $f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \geq 0$ such that $\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \geq 0$. In other words, for any $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, we have $\nabla F_t(\pi_1, \dots, \pi_n) \geq \mathbf{0}$, namely, for any two point $(\pi_1^a, \dots, \pi_n^a) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and $(\pi_1^b, \dots, \pi_n^b) \in \prod_{i=1}^n \Delta_{\kappa_i}$, if $\pi_i^a \leq \pi_i^b$ for any $i \in \mathcal{N}$, we have $F_t(\pi_1^a, \dots, \pi_n^a) \leq F_t(\pi_1^b, \dots, \pi_n^b)$. So F_t is monotone.

3): For any fixed two policy vector $(\pi_1^a, \dots, \pi_n^a) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and $(\pi_1^b, \dots, \pi_n^b) \in \prod_{i=1}^n \Delta_{\kappa_i}$ where $\pi_i^a \leq \pi_i^b, \forall i \in \mathcal{N}$, we consider an unified sampling strategy to generate actions. Before that, we set each $\pi_i^a = (\pi_{i,1}^a, \dots, \pi_{i,\kappa_i}^a)$ and $\pi_i^b = (\pi_{i,1}^b, \dots, \pi_{i,\kappa_i}^b), \forall i \in \mathcal{N}$. At first, we transfer the sampling process according to policy π_i^a to a uniform random variable $X_i \in [0, 1]$, namely,

$$a(X_i, \pi_i^a) \triangleq \begin{cases} v_{i,1} & \text{If } X_i \in [0, \pi_{i,1}^a) \\ v_{i,m} & \text{If } X_i \in [\sum_{k=1}^{m-1} \pi_{i,k}^a, \sum_{k=1}^m \pi_{i,k}^a) \text{ for any integer } m \in [2, \kappa_i] \\ \emptyset & \text{If } X_i \geq \sum_{k=1}^{\kappa_i} \pi_{i,k}^a \end{cases} \tag{14}$$

When X_i is an uniform random variable over range $[0, 1]$, it is easy to check that the $a(X_i, \pi_i^a)$ follows the same law as the policy π_i^a , that is, $\Pr(a(X_i, \pi_i^a) = v_{i,m}) = \pi_{i,m}^a$ and $\Pr(a(X_i, \pi_i^a) = \emptyset) = 1 - \sum_{m=1}^{\kappa_i} \pi_{i,m}^a$ where the symbol 'Pr' denotes the probability.

Similarly, we also can transfer the sampling process according to policy π_i^b to two independent uniform random variables $X_i, Y \in [0, 1]$, namely,

$$a(X_i, Y, \pi_i^a, \pi_i^b) \triangleq \begin{cases} a(X_i, \pi_i^a) & \text{If } X_i < \sum_{k=1}^{\kappa_i} \pi_{i,k}^a \\ v_{i,1} & \text{If } X_i \geq \sum_{k=1}^{\kappa_i} \pi_{i,k}^a \text{ and } Y \in [0, \frac{\pi_{i,1}^b - \pi_{i,1}^a}{1 - \sum_{k=1}^{\kappa_i} \pi_{i,k}^a}) \\ v_{i,m} & \text{If } X_i \geq \sum_{k=1}^{\kappa_i} \pi_{i,k}^a \text{ and } Y \in [\frac{\sum_{k=1}^{m-1} (\pi_{i,k}^b - \pi_{i,k}^a)}{1 - \sum_{k=1}^{\kappa_i} \pi_{i,k}^a}, \frac{\sum_{k=1}^m (\pi_{i,k}^b - \pi_{i,k}^a)}{1 - \sum_{k=1}^{\kappa_i} \pi_{i,k}^a}), \forall m \in [2, \kappa_i] \\ \emptyset & \text{If } X_i \geq \sum_{k=1}^{\kappa_i} \pi_{i,k}^a \text{ and } Y \geq \frac{\sum_{k=1}^{\kappa_i} (\pi_{i,k}^b - \pi_{i,k}^a)}{1 - \sum_{k=1}^{\kappa_i} \pi_{i,k}^a} \end{cases} \tag{15}$$

From part **1**) and fixed a sequence of independent uniform variables $X_i, Y_i, \forall i \in \mathcal{N}$, we can have that

$$\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1^a, \dots, \pi_n^a) = \mathbb{E}_{X_i} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, \pi_i^a)\}) \right),$$

and

$$\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1^b, \dots, \pi_n^b) = \mathbb{E}_{(X_i, Y_i)} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, Y_i, \pi_i^a, \pi_i^b)\}) \right).$$

From Eq.(14) and Eq.(15), we can know that $\{a(X_i, \pi_i^a)\} \subseteq \{a(X_i, Y_i, \pi_i^a, \pi_i^b)\}$ such that $\cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, \pi_i^a)\} \subseteq \cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, Y_i, \pi_i^a, \pi_i^b)\}$. Then, from the definition of α -DR submodularity, we have that $f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, \pi_i^a)\}) \geq \alpha f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a(X_i, Y_i, \pi_i^a, \pi_i^b)\})$, so $\nabla F_t(\pi_1^a, \dots, \pi_n^a) \geq \alpha \nabla F_t(\pi_1^b, \dots, \pi_n^b)$.

4): For any subset S within the constraint of problem (1) and any point $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, when f_t is monotone (γ, β) -weakly monotone submodular, we firstly can show that for any two actions $a_1, a_2 \in \mathcal{V}$ and the subset $B \subseteq \mathcal{V}$, we have that

$$\gamma(f_t(a_1 | B \cup \{a_2\}) + f_t(a_2 | B)) = \gamma(f_t(B \cup \{a_1, a_2\}) - f_t(B)) \leq f_t(a_1 | B) + f_t(a_2 | B),$$

such that

$$f_t(a_1 | B) \geq \gamma f_t(a_1 | B \cup \{a_2\}) - (1 - \gamma) f_t(a_2 | B). \quad (16)$$

Therefore, we can show that

$$\begin{aligned} & \sum_{(i,m): v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \\ &= \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\gamma f_t(v_{i,m} | \cup_{j \in \mathcal{N}} \{a_j\}) - (1 - \gamma) f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right), \end{aligned} \quad (17)$$

where the first equality follows from part **1**) of Theorem 1 and the second inequality comes from the Eq.(16). Then, from the γ -weakly submodularity,

$$\begin{aligned} & \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \gamma \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\gamma f_t(S | \cup_{j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \gamma (f_t(S) - F_t(\pi_1, \dots, \pi_n)), \end{aligned} \quad (18)$$

where the final inequality follows from the monotonicity.

Furthermore, from the β -weakly upper submodularity, we also have,

$$\begin{aligned} & \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\leq \sum_{i \in \mathcal{N}} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\leq \beta \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - f_t(\emptyset) \right) \\ &= \beta F_t(\pi_1, \dots, \pi_n), \end{aligned} \quad (19)$$

where the second inequality follows from the definition of β -upper submodularity.

Merging Eq.(19) and Eq.(18) into Eq.(17), we have that

$$\begin{aligned} & \sum_{(i,m): v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \\ &\geq \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\gamma f_t(v_{i,m} | \cup_{j \in \mathcal{N}} \{a_j\}) - (1 - \gamma) f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq (\gamma^2 (f_t(S) - F_t(\pi_1, \dots, \pi_n)) - (1 - \gamma) \beta F_t(\pi_1, \dots, \pi_n)) \\ &= \gamma^2 f_t(S) - ((1 - \gamma) \beta + \gamma^2) F_t(\pi_1, \dots, \pi_n). \end{aligned}$$

As for f_t is monotone α -weakly DR-submodular, we also can show that

$$\begin{aligned} \sum_{(i,m):v_{i,m} \in S} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) &= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\sum_{(i,m):v_{i,m} \in S} f_t(v_{i,m} \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \alpha \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(S \mid \cup_{j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \alpha (f_t(S) - F_t(\pi_1, \dots, \pi_n)), \end{aligned}$$

where the first inequality follows from the α -DR submodularity and the final inequality comes from the monotonicity of f_t .

E Proof of Theorem 2

In this section, we verify the Theorem 2.

Before that, we firstly suppose that the symbol $\mathbf{1}_S$ represents the indicator function over the subset $S \subseteq \mathcal{V}$, namely, for any action $v_{i,m} \in S$, the vector $\mathbf{1}_S$ sets the corresponding probability for this action $v_{i,m}$ as 1. Note that, when S satisfies the constraints of problem (1), namely, $|S \cap \mathcal{V}_i| \leq 1$, we can infer that $\mathbf{1}_S \in \prod_{i=1}^n \Delta_{\kappa_i}$.

E.1 Proof of Part 2) and Part 3) in Theorem 2

With the previously defined symbol $\mathbf{1}_S$, we can rewrite the 4) of Theorem 1 as:

i): when f_t is monotone α -weakly DR-submodular, the following inequality holds:

$$\alpha (f_t(S) - F_t(\pi_1, \dots, \pi_n)) \leq \langle \mathbf{1}_S, \nabla F_t(\pi_1, \dots, \pi_n) \rangle;$$

ii): when f_t is monotone (γ, β) -weakly submodular, we can show that

$$(\gamma^2 f_t(S) - (\beta(1 - \gamma) + \gamma^2) F_t(\pi_1, \dots, \pi_n)) \leq \langle \mathbf{1}_S, \nabla F_t(\pi_1, \dots, \pi_n) \rangle.$$

In order to prove Theorem 2, we next show the relationship between $\langle (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle$ and $F_t(\pi_1, \dots, \pi_n)$.

From Eq.(13), if $\pi_i = (\pi_{i,1}, \dots, \pi_{i,\kappa_i}), \forall i \in \mathcal{N}$, we can show

$$\begin{aligned} &\langle (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} \pi_{i,m} \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_1, \dots, \pi_n) \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} \pi_{i,m} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} \left(\pi_{i,m} \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(f_t(v_{i,m} \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \prod_{j \neq i, j \in \mathcal{N}} p(a_j \mid \pi_j) \right) \right) \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} \left(p(v_{i,m} \mid \pi_i) \sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(f_t(v_{i,m} \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \prod_{j \neq i, j \in \mathcal{N}} p(a_j \mid \pi_j) \right) \right) \quad (20) \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} \left(\sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(f_t(v_{i,m} \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) p(v_{i,m} \mid \pi_i) * \prod_{j \neq i, j \in \mathcal{N}} p(a_j \mid \pi_j) \right) \right) \\ &= \sum_{i=1}^n \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}} \left(\sum_{a_j \in \mathcal{V}_j \cup \{\emptyset\}, \forall j \neq i} \left(f_t(a_i \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \prod_{j \in \mathcal{N}} p(a_j \mid \pi_j) \right) \right) \\ &= \sum_{i=1}^n \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i \mid \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right), \end{aligned}$$

where the second equality follows from the **1**) of Theorem 1, the third equality comes from the final equality of the Eq.(13), the fourth equality follows from $p(v_{i,m}|\pi_i) = \pi_{i,m}$ and the final equality from the fact that the random element a_i is drawn from the policy π_i .

Therefore, if the original set function f_t is β -weakly submodular from above, we can show that

$$\begin{aligned} & \langle (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \\ &= \sum_{i=1}^n \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\sum_{i=1}^n f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\leq \beta \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - f_t(\emptyset) \right) \\ &= \beta F_t(\pi_1, \dots, \pi_n). \end{aligned} \quad (21)$$

Therefore, by merging Eq.(21) and the **4**) of Theorem 1, when the set function f_t is monotone and (γ, β) -weakly submodular, for any subset S within the constraint of problem (1) and $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, we have that

$$\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \geq \gamma^2 f_t(S) - (\beta + \beta(1 - \gamma) + \gamma^2) F_t(\pi_1, \dots, \pi_n).$$

Similarly, from the definition of α -weakly DR-submodular function, we know that α -weakly DR-submodular function automatically satisfies the conditions for being $\frac{1}{\alpha}$ -weakly submodular from above, we also have that, when f_t is monotone α -weakly DR-submodular,

$$\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \geq \alpha f_t(S) - (\alpha + \frac{1}{\alpha}) F_t(\pi_1, \dots, \pi_n).$$

Therefore, when $(\pi_1^s, \dots, \pi_n^s)$ is a stationary point of F_t over the domain $\prod_{i=1}^n \Delta_{\kappa_i}$, we have

i): f_t is monotone (γ, β) -weakly submodular, for any S within the constraint of problem (1),

$$\gamma^2 f_t(S) - (\beta + \beta(1 - \gamma) + \gamma^2) F_t(\pi_1^s, \dots, \pi_n^s) \leq \langle \mathbf{1}_S - (\pi_1^s, \dots, \pi_n^s), \nabla F_t(\pi_1^s, \dots, \pi_n^s) \rangle \leq 0.$$

In other words,

$$F_t(\pi_1^s, \dots, \pi_n^s) \geq \frac{\gamma^2}{\beta + \beta(1 - \gamma) + \gamma^2} f_t(S^*),$$

where S^* is the optimal subset of problem (1).

We get the **3**) in Theorem 2.

Similarly, **ii):** when f_t is monotone α -weakly DR-submodular, for any S within the constraint of problem (1),

$$\alpha f_t(S) - (\alpha + \frac{1}{\alpha}) F_t(\pi_1^s, \dots, \pi_n^s) \leq \langle \mathbf{1}_S - (\pi_1^s, \dots, \pi_n^s), \nabla F_t(\pi_1^s, \dots, \pi_n^s) \rangle \leq 0.$$

In other words,

$$F_t(\pi_1^s, \dots, \pi_n^s) \geq \frac{\alpha^2}{1 + \alpha} f_t(S^*).$$

where S^* is the optimal subset of problem (1).

We get the **2**) in Theorem 2.

E.2 Proof of Part 1) in Theorem 2

In this subsection, we prove the part **1**) in Theorem 2.

From Eq.(17), we have

$$\langle \mathbf{1}_S, \nabla F_t(\pi_1, \dots, \pi_n) \rangle = \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right).$$

From Eq.(20), we have

$$\langle (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle = \sum_{i=1}^n \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right).$$

Therefore, we have that

$$\begin{aligned} & \langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \\ &= \sum_{(i,m): v_{i,m} \in S} \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) - \sum_{i=1}^n \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\sum_{(i,m): v_{i,m} \in S} f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) - \sum_{i=1}^n f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\sum_{v_{i,m} \in (S \setminus \cup_{j \in \mathcal{N}} \{a_j\})} f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) - \sum_{a_i \in (\cup_{j \in \mathcal{N}} \{a_j\} \setminus S)} f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right). \end{aligned} \quad (22)$$

When f_t is monotone submodular, we have

$$\sum_{v_{i,m} \in (S \setminus \cup_{j \in \mathcal{N}} \{a_j\})} f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \geq f_t(\cup_{j \in \mathcal{N}} \{a_j\} \cup S) - f_t(\cup_{j \in \mathcal{N}} \{a_j\}). \quad (23)$$

Furthermore, from the definition of curvature c , we also have

$$f_t(\cup_{j \in \mathcal{N}} \{a_j\} \cup S) - f_t(S) \geq (1 - c) \left(f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - f_t(\cup_{j \in \mathcal{N}} \{a_j\} \cap S) \right), \quad (24)$$

where this inequality follows from $(\cup_{j \in \mathcal{N}} \{a_j\} \cup S) \setminus S = (\cup_{j \in \mathcal{N}} \{a_j\}) \setminus (\cup_{j \in \mathcal{N}} \{a_j\} \cap S)$.

Similarly, from the submodularity, we have

$$\sum_{a_i \in (\cup_{j \in \mathcal{N}} \{a_j\} \setminus S)} f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \leq \left(f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - f_t(\cup_{j \in \mathcal{N}} \{a_j\} \cap S) \right), \quad (25)$$

where the inequality comes from that $(\cup_{j \in \mathcal{N}} \{a_j\} \cap S) \subseteq (\cup_{j \neq i, j \in \mathcal{N}} \{a_j\})$ if $a_i \in \cup_{j \in \mathcal{N}} \{a_j\} \setminus S$.

Merging Eq.(23), Eq.(24) and Eq.(25) into Eq.(22), we can show that

$$\begin{aligned} & \langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t(\pi_1, \dots, \pi_n) \rangle \\ &= \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(\sum_{v_{i,m} \in S \setminus \cup_{j \in \mathcal{N}} \{a_j\}} f_t(v_{i,m} | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) - \sum_{a_i \in \cup_{j \in \mathcal{N}} \{a_j\} \setminus S} f_t(a_i | \cup_{j \neq i, j \in \mathcal{N}} \{a_j\}) \right) \\ &\geq \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(S) - f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - c \left(f_t(\cup_{j \in \mathcal{N}} \{a_j\}) - f_t(\cup_{j \in \mathcal{N}} \{a_j\} \cap S) \right) \right) \\ &\geq \mathbb{E}_{a_j \sim \pi_j, \forall j \in \mathcal{N}} \left(f_t(S) - (1 + c) f_t(\cup_{j \in \mathcal{N}} \{a_j\}) \right) \\ &= f_t(S) - (1 + c) F_t(\pi_1, \dots, \pi_n). \end{aligned}$$

Thus, if $(\pi_1^s, \dots, \pi_n^s)$ is a stationary point of F_t over the domain $\prod_{i=1}^n \Delta_{\kappa_i}$, when f_t is monotone submodular with curvature c , we have, for any S in the constraint of problem (1), we have

$$f_t(S) - (1 + c) F_t(\pi_1^s, \dots, \pi_n^s) \leq \langle \mathbf{1}_S - (\pi_1^s, \dots, \pi_n^s), \nabla F_t(\pi_1^s, \dots, \pi_n^s) \rangle \leq 0.$$

So,

$$F_t(\pi_1^s, \dots, \pi_n^s) \geq \frac{1}{1 + c} f_t(S^*),$$

where S^* is the optimal subset of problem (1).

E.3 A Policy-based Continuous Extension with 1/2-Approximation Stationary Point

In this subsection we consider a special set function.

At first, let the universe set U consist of $n - 1$ elements $\{x_1, \dots, x_{n-1}\}$ and $n - k$ elements $\{y_1, \dots, y_{n-k}\}$, all of weight 1, and $n - 1$ elements $\{\epsilon_1, \dots, \epsilon_{n-1}\}$ of arbitrarily small weight $\epsilon > 0$. Then, we define two different types of sets namely, A_i and A_{i+n} for any $i \in [n]$, that is to say ,

$$\begin{aligned} A_i &\triangleq \{\epsilon_i\} \text{ for } 1 \leq i \leq n - 1, & A_n &\triangleq \{x_1, \dots, x_{n-1}\}, \\ A_{n+i} &\triangleq \{x_i\} \text{ for } 1 \leq i \leq n - 1, & A_{2n} &\triangleq \{y_1, \dots, y_{n-k}\}. \end{aligned}$$

After that, we define a coverage set function $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}_+$ over these $2n$ distinct set $\{A_1, \dots, A_n, A_{n+1}, \dots, A_{2n}\}$ where $\mathcal{V} = [2n]$. Specifically, for any subset $S \subseteq \mathcal{V}$,

$$f(S) \triangleq \sum_{v \in \bigcup_{i \in S} A_i} w(v), \quad (26)$$

where $w(v)$ is the weight of element v .

Moreover, we consider a partition constraint that contains at most one of $\{A_i, A_{n+i}\}$ for any $i \in [n]$. If we set $\mathcal{V}_i \triangleq \{i, i + n\}$ and $\mathcal{V} \triangleq \bigcup_{i \in [n]} \mathcal{V}_i \triangleq [2n]$, we naturally obtain the following coverage maximization problem:

$$\max_{S \subseteq \mathcal{V}} f(S) \text{ s.t. } |S \cap \mathcal{V}_i| \leq 1 \quad \forall i \in [n]. \quad (27)$$

Note that this problem (27) is a special case of the concern problem (1). From the result of [38], we know that the coverage function f in Equation (26) is a submodular set function, namely, $\alpha = \beta = \gamma = 1$.

A key feature of the coverage maximization problem (27) is that [38] found that the standard greedy [85] will be stuck at a local maximum subset $\{A_1, \dots, A_n\}$ where $S = [n]$ and $f(S) = (1 + \epsilon)n$. In contrast, when ϵ is very small, the optimal subset for the problem (27) is $\{A_{n+1}, \dots, A_{2n}\}$ where $S = \{n + 1, \dots, 2n\}$ and $f(S) = 2n - k - 1$. Note that $\lim_{n \rightarrow \infty} \lim_{\epsilon \rightarrow 0} \lim_{k \rightarrow 0} \frac{(1 + \epsilon)n}{2n - k - 1} = \frac{1}{2}$. Motivated by this finding of [38], we also can show that the point $\mathbf{1}_{[n]} \triangleq (\underbrace{1, \dots, 1}_n, 0, \dots, 0)$ is a

local stationary point of the policy-based continuous extension of the set function f in Equation (26). More specifically, we have the following theorem:

Theorem 6. *The point $\mathbf{1}_{[n]}$ is a stationary point of the policy-based continuous extension F of the set function f in Equation (26). Moreover, we can show $\frac{F(\mathbf{1}_{[n]})}{f(\{n+1, \dots, 2n\})} = \frac{(1 + \epsilon)n}{2n - k - 1} \rightarrow \frac{1}{2}$.*

Remark 11. *This theorem indicates that when the objective set function is submodular, namely, $c = \alpha = \gamma = \beta = 1$, the approximation guarantees established in Theorem 2 is tight.*

Proof. At first, for any $i \in [n]$, we assume $\pi_i = (\pi_{i,1}, \pi_{i,2})$. Then, according to Definition 1, we have that the policy-based continuous extension F of the set function f in Equation (26) can be formulated as:

$$F(\pi_1, \dots, \pi_n) \triangleq \sum_{i=1}^n \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}} \left(f\left(\bigcup_{i=1}^n \{a_i\}\right) \prod_{i=1}^n p(a_i | \pi_i) \right), \quad (28)$$

where $p(i | \pi_i) = \pi_{i,1}$, $p(n + i | \pi_i) = \pi_{i,2}$ and $p(\emptyset | \pi_i) = 1 - \pi_{i,1} - \pi_{i,2}$.

From part 1) of Theorem 1, we also can show that,

$$\begin{aligned} \frac{\partial F}{\partial \pi_{i,1}}(\mathbf{1}_{[n]}) &= \epsilon \text{ for } 1 \leq i \leq n - 1, & \frac{\partial F}{\partial \pi_{n,1}}(\mathbf{1}_{[n]}) &= n - 1, \\ \frac{\partial F}{\partial \pi_{i,2}}(\mathbf{1}_{[n]}) &= 0 \text{ for } 1 \leq i \leq n - 1, & \frac{\partial F}{\partial \pi_{n,2}}(\mathbf{1}_{[n]}) &= n - k. \end{aligned}$$

As a result, for any $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_2$, we have

$$\begin{aligned}
& \langle (\pi_1, \dots, \pi_n) - \mathbf{1}_{[n]}, \nabla F(\mathbf{1}_{[n]}) \rangle \\
&= \sum_{i=1}^n \left(\pi_{i,1} \frac{\partial F}{\partial \pi_{i,1}}(\mathbf{1}_{[n]}) + \pi_{i,2} \frac{\partial F}{\partial \pi_{i,2}}(\mathbf{1}_{[n]}) \right) - \sum_{i=1}^n \frac{\partial F}{\partial \pi_{i,1}}(\mathbf{1}_{[n]}) \\
&= \sum_{i=1}^n \left(\pi_{i,1} \frac{\partial F}{\partial \pi_{i,1}}(\mathbf{1}_{[n]}) + \pi_{i,2} \frac{\partial F}{\partial \pi_{i,2}}(\mathbf{1}_{[n]}) \right) - (1 + \epsilon)(n - 1) \\
&= \epsilon \sum_{i=1}^{n-1} \pi_{i,1} + (n - 1)\pi_{n,1} + (n - k)\pi_{n,2} - (1 + \epsilon)(n - 1) \\
&= \epsilon \sum_{i=1}^{n-1} \pi_{i,1} + (n - 1)(\pi_{n,1} + \pi_{n,2}) + ((n - k) - (n - 1))\pi_{n,2} - (1 + \epsilon)(n - 1) \\
&= \epsilon \sum_{i=1}^{n-1} (\pi_{i,1} - 1) + ((n - k) - (n - 1))\pi_{i,2} - (n - 1)(1 - \pi_{n,1} - \pi_{n,2}) \\
&= \epsilon \sum_{i=1}^{n-1} (\pi_{i,1} - 1) + (1 - k)\pi_{i,2} - (n - 1)(1 - \pi_{n,1} - \pi_{n,2}) \leq 0,
\end{aligned}$$

where the final inequality follows from $1 - k \leq 0$ and $1 - \pi_{n,1} - \pi_{n,2} \geq 0$. As a result, from Definition 2, we can know that the point $\mathbf{1}_{[n]}$ is a stationary point of the policy-based continuous extension F . Note that $\mathcal{F}(\mathbf{1}_{[n]}) = f([n]) = (1 + \epsilon)(n - 1)$ such that $\frac{F(\mathbf{1}_{[n]})}{f(\{n+1, \dots, 2n\})} = \frac{(1+\epsilon)n}{2n-k-1}$. Particularly when $k \rightarrow 0$, $\epsilon \rightarrow 0$ and $n \rightarrow \infty$, $\frac{F(\mathbf{1}_{[n]})}{f(\{n+1, \dots, 2n\})} = \frac{(1+\epsilon)n}{2n-k-1} \rightarrow \frac{1}{2}$. $\square$

F Proof of Theorem 3

In this section, we cut the proof of Theorem 3 in two distinct subsections. Specifically, Appendix F.1 provides the proof for parts 1) and 2), while Appendix F.2 addresses the part 3).

F.1 Proof of Part 1) and Part 2) in Theorem 3

Before verifying the parts 1) and 2), we firstly prove the following theorem:

Theorem 7. Given a monotone set function f_t , for its policy-based continuous extension F_t introduced in Definition 1, if we consider a surrogate function $F_t^s : \prod_{i=1}^n \Delta_{\kappa_i} \rightarrow \mathbb{R}_+$ whose gradient at each point $\mathbf{x} \in \prod_{i=1}^n \Delta_{\kappa_i}$ is a weighted average of the gradient $\nabla F_t(z * \mathbf{x})$, namely, $\nabla F_t^s(\mathbf{x}) \triangleq \int_0^1 w(z) \nabla F_t(z * \mathbf{x}) dz$ where $w(z)$ is a positive weight function over $[0, 1]$, we have that:

i): When f_t is α -weakly DR-submodular and $w(z) = e^{\alpha(z-1)}$, for any policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1), then the following inequality holds:

$$\begin{aligned}
& \langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t^s(\pi_1, \dots, \pi_n) \rangle \\
&= \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \int_0^1 e^{\alpha(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \\
&\geq (1 - e^{-\alpha}) f_t(S) - F_t(\pi_1, \dots, \pi_n);
\end{aligned} \tag{29}$$

ii): When f_t is (γ, β) -weakly submodular and $w(z) = e^{\phi(\gamma, \beta)(z-1)}$ where $\phi(\gamma, \beta) = \beta(1 - \gamma) + \gamma^2$, for any policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1),

then the following inequality holds:

$$\begin{aligned}
& \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla F_t^s(\pi_1, \dots, \pi_n) \right\rangle \\
&= \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) \, dz \right\rangle \\
&\geq \left(\frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \right) f_t(S) - F_t(\pi_1, \dots, \pi_n).
\end{aligned} \tag{30}$$

Proof. For **i)**: At first, we verify the following relationship:

$$\begin{aligned}
& \left\langle (\pi_1, \dots, \pi_n), \int_0^1 e^{\alpha(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) \, dz \right\rangle \\
&= \int_0^1 e^{\alpha(z-1)} \left\langle (\pi_1, \dots, \pi_n), \nabla F_t(z * (\pi_1, \dots, \pi_n)) \right\rangle dz \\
&= \int_0^1 e^{\alpha(z-1)} \frac{dF_t(z * (\pi_1, \dots, \pi_n))}{dz} dz \\
&= \int_0^1 e^{\alpha(z-1)} dF_t(z * (\pi_1, \dots, \pi_n)) \\
&= e^{\alpha(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) \Big|_{z=0}^{z=1} - \int_0^1 F_t(z * (\pi_1, \dots, \pi_n)) d(e^{\alpha(z-1)}) \\
&= F_t(\pi_1, \dots, \pi_n) - \alpha \int_0^1 e^{\alpha(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) \, dz,
\end{aligned}$$

where the final equality follows from $F_t(\mathbf{0}) = 0$.

Then, we have the following inequality:

$$\begin{aligned}
& \left\langle \mathbf{1}_S, \int_0^1 e^{\alpha(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) \, dz \right\rangle \\
&= \int_0^1 e^{\alpha(z-1)} \left\langle \mathbf{1}_S, \nabla F_t(z * (\pi_1, \dots, \pi_n)) \right\rangle dz \\
&\geq \alpha \int_0^1 e^{\alpha(z-1)} (f_t(S) - F_t(z * (\pi_1, \dots, \pi_n))) \, dz \\
&= \left(\alpha \int_0^1 e^{\alpha(z-1)} dz \right) f_t(S) - \alpha \int_0^1 e^{\alpha(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) \, dz \\
&= (1 - e^{-\alpha}) f_t(S) - \alpha \int_0^1 e^{\alpha(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) \, dz,
\end{aligned}$$

where the first inequality follows from the part **4)** of Theorem 1.

As a result, we have

$$\left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \int_0^1 e^{\alpha(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) \, dz \right\rangle \geq (1 - e^{-\alpha}) f_t(S) - F_t(\pi_1, \dots, \pi_n).$$

As for **ii)**: Similarly, when f_t is monotone (γ, β) -weakly submodular and $w(z) = e^{\phi(\gamma, \beta)(z-1)}$ where $\phi(\gamma, \beta) = \beta(1 - \gamma) + \gamma^2$, we have

$$\begin{aligned}
& \left\langle (\pi_1, \dots, \pi_n), \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \\
&= \int_0^1 e^{\phi(\gamma, \beta)(z-1)} dF_t(z * (\pi_1, \dots, \pi_n)) \\
&= e^{\phi(\gamma, \beta)(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) \Big|_{z=0}^{z=1} - \phi(\gamma, \beta) \int_0^1 e^{\phi(\gamma, \beta)(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) dz \\
&= F_t(\pi_1, \dots, \pi_n) - \phi(\gamma, \beta) \int_0^1 e^{\phi(\gamma, \beta)(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) dz.
\end{aligned}$$

Then, we also have

$$\begin{aligned}
& \left\langle \mathbf{1}_S, \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \\
&= \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \left\langle \mathbf{1}_S, \nabla F_t(z * (\pi_1, \dots, \pi_n)) \right\rangle dz \\
&\geq \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \left(\gamma^2 f_t(S) - \phi(\gamma, \beta) F_t(z * (\pi_1, \dots, \pi_n)) \right) dz \\
&= \left(\gamma^2 \int_0^1 e^{\phi(\gamma, \beta)(z-1)} dz \right) f_t(S) - \phi(\gamma, \beta) \int_0^1 e^{\phi(\gamma, \beta)(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) dz \\
&= \left(\frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \right) f_t(S) - \phi(\gamma, \beta) \int_0^1 e^{\phi(\gamma, \beta)(z-1)} F_t(z * (\pi_1, \dots, \pi_n)) dz,
\end{aligned}$$

where the first inequality follows from the part 4) of Theorem 1.

As a result, we have

$$\left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \int_0^1 e^{\phi(\gamma, \beta)(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \geq \left(\frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \right) f_t(S) - F_t(\pi_1, \dots, \pi_n).$$

□

From Definition 2, we can show that, when the policy vector $(\pi_1^s, \dots, \pi_n^s)$ is the stationary point of the surrogate objective F_t^s over domain $\prod_{i=1}^n \Delta_{\kappa_i}$ and S satisfies the constraint of problem (1), due to $\mathbf{1}_S \in \prod_{i=1}^n \Delta_{\kappa_i}$, the following inequality holds:

$$\left\langle \mathbf{1}_S - (\pi_1^s, \dots, \pi_n^s), \int_0^1 w(z) \nabla F_t(z * (\pi_1^s, \dots, \pi_n^s)) dz \right\rangle \leq 0,$$

where $w(z)$ is the corresponding weight function.

Then, following Theorem 7, we can show that when f_t is monotone α -weakly DR-submodular and $w(z) = e^{\alpha(z-1)}$, for any policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1),

$$F_t(\pi_1^s, \dots, \pi_n^s) \geq (1 - e^{-\alpha}) f_t(S).$$

So we get the result of part 1) of Theorem 3. Similarly, from ii) of Theorem 7, we also can achieve the part 2) of Theorem 3.

F.2 Proof of Part 3) in Theorem 3

In this subsection, we prove 3) in Theorem 3.

At first, we define a modular function g_t over $\mathcal{V}$, namely, for any subset $S \subseteq \mathcal{V}$, $g_t(S) \triangleq \sum_{a \in S} f_t(a | \mathcal{V} \setminus \{a\})$. Then, when f_t is monotone submodular, we can show that the set function $(f_t - g_t)$ is also a monotone submodular function. That is,

Theorem 8. For a monotone submodular function $f_t : \mathcal{V} \rightarrow \mathbb{R}_+$, we define a modular function $g_t(S) \triangleq \sum_{a \in S} f_t(a|\mathcal{V} \setminus \{a\})$, $\forall S \subseteq \mathcal{V}$, then we also can show that the set function $(f_t - g_t)$ is monotone submodular and non-negative.

Proof. Because the set function $(f_t - g_t)$ is the sum of a submodular function f_t and a modular function $-g_t$, $(f_t - g_t)$ must be submodular [105]. Then, we show the monotonicity of the set function $(f_t - g_t)$. At first, for any subset $S \subseteq \mathcal{V}$ and any action $a \in \mathcal{V} \setminus S$, we have that

$$(f_t - g_t)(a|S) = f_t(a|S) - g_t(a|S) = f_t(a|S) - f_t(a|\mathcal{V} \setminus \{a\}) \geq 0, \quad (31)$$

where the final inequality follows from the submodularity of f_t .

From Eq.(31), we can know that $(f_t - g_t)$ is monotone. Moreover, due to $f_t(\emptyset) = g_t(\emptyset) = 0$, $(f_t - g_t)$ is non-negative. $\square$

Next, we show the policy-based continuous extension of our aforementioned modular function g_t .

Theorem 9. Given a monotone submodular function $f_t : \mathcal{V} \rightarrow \mathbb{R}_+$, we define its related modular function as $g_t(S) \triangleq \sum_{a \in S} f_t(a|\mathcal{V} \setminus \{a\})$, $\forall S \subseteq \mathcal{V}$. Then we can show that the policy-based continuous extension of g_t is as follows: $G_t(\pi_1, \dots, \pi_n) \triangleq \sum_{i=1}^n \sum_{m=1}^{\kappa_i} (f_t(v_{i,m}|\mathcal{V} - \{v_{i,m}\}))\pi_{i,m}$ if $\pi_i = (\pi_{i,1}, \dots, \pi_{i,\kappa_i}) \in \Delta_{\kappa_i}$, $\forall i \in \mathcal{N}$.

Proof. By the definition of g_t and Definition 1, we can show that

$$\begin{aligned} G_t(\pi_1, \dots, \pi_n) &= \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}, \forall i \in \mathcal{N}} \left(g_t(\cup_{i=1}^n \{a_i\}) \prod_{i=1}^n p(a_i|\pi_i) \right) \\ &= \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}, \forall i \in \mathcal{N}} \left(\left(\sum_{i=1}^n f_t(a_i|\mathcal{V} \setminus \{a_i\}) \right) \prod_{i=1}^n p(a_i|\pi_i) \right) \\ &= \sum_{i=1}^n \sum_{a_i \in \mathcal{V}_i \cup \{\emptyset\}} f_t(a_i|\mathcal{V} \setminus \{a_i\}) p(a_i|\pi_i) \\ &= \sum_{i=1}^n \sum_{m=1}^{\kappa_i} (f_t(v_{i,m}|\mathcal{V} - \{v_{i,m}\}))\pi_{i,m}, \end{aligned}$$

where the final equality follows from $\mathcal{V}_i = \{v_{i,1}, \dots, v_{i,\kappa_i}\}$, $\forall i \in \mathcal{N}$. $\square$

From Theorem 8, we know that, when f_t is monotone submodular, so is $(f_t - g_t)$. Then, if we apply Theorem 7 to the monotone submodular function $(f_t - g_t)$, we can have the following relationship:

Lemma 1. Given a monotone submodular set function f_t , we define a modular function $g_t(S) \triangleq \sum_{a \in S} f_t(a|\mathcal{V} \setminus \{a\})$, $\forall S \subseteq \mathcal{V}$. then, for any policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1), the following inequality holds:

$$\begin{aligned} &\left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \int_0^1 e^{z-1} \nabla (F_t - G_t)(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \\ &\geq (1 - 1/e) (f_t - g_t)(S) - (F_t - G_t)(\pi_1, \dots, \pi_n), \end{aligned} \quad (32)$$

where F_t and G_t is the corresponding policy-based continuous extension of f_t and g_t , respectively.

Remark 12. Note that the submodular objective function is also 1-weakly DR-submodular function. Thus, the results of Lemma 1 come from Theorem 7 and Theorem 8.

Moreover, from the definition of G_t in Theorem 9, we know that G_t is a linear function over $\prod_{i=1}^n \Delta_{\kappa_i}$ such that

$$\left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla G_t(\pi_1, \dots, \pi_n) \right\rangle = G_t(\mathbf{1}_S) - G_t(\pi_1, \dots, \pi_n) = g_t(S) - G_t(\pi_1, \dots, \pi_n). \quad (33)$$

Merging Eq.(32) and Eq.(33), we get the following theorem:

Theorem 10. Given a monotone submodular set function f_t with curvature c , for its policy-based continuous extension F_t introduced in Definition 1, if we consider a surrogate function $F_t^s : \prod_{i=1}^n \Delta_{\kappa_i} \rightarrow \mathbb{R}_+$ whose gradient at each point $\mathbf{x} \in \prod_{i=1}^n \Delta_{\kappa_i}$ is a weighted average of the gradient $\nabla F_t(z * \mathbf{x})$, namely, $\nabla F_t^s(\mathbf{x}) \triangleq \int_0^1 e^{(z-1)} \nabla F_t(z * \mathbf{x}) dz$, we have that, for any policy vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1), the following inequality holds:

$$\begin{aligned} & \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla \left(F_t^s + \frac{G_t}{e} \right) (\pi_1, \dots, \pi_n) \right\rangle \\ &= \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \frac{\nabla G_t(\pi_1, \dots, \pi_n)}{e} + \int_0^1 e^{\alpha(z-1)} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \quad (34) \\ &\geq \left(1 - \frac{c}{e} \right) f_t(S) - F_t(\pi_1, \dots, \pi_n). \end{aligned}$$

Proof. According to Eq.(32) and Eq.(33), we have that

$$\begin{aligned} & \left\langle \mathbf{1}_S - (\pi_1, \dots, \pi_n), \nabla G_t(\pi_1, \dots, \pi_n) + \int_0^1 e^{z-1} \nabla (F_t - G_t)(z * (\pi_1, \dots, \pi_n)) dz \right\rangle \\ &\geq (1 - 1/e) f_t(S) + \frac{g_t(S)}{e} - F_t((\pi_1, \dots, \pi_n)) \quad (35) \\ &\geq (1 - 1/e) f_t(S) + \frac{(1-c)f_t(S)}{e} - F_t((\pi_1, \dots, \pi_n)) \\ &= (1 - c/e) f_t(S) - F_t((\pi_1, \dots, \pi_n)), \end{aligned}$$

where the final inequality follows from the definition of curvature c (See Lemma 2.1. in [105]).

Moreover, due to that G_t is a linear function, we have

$$\begin{aligned} & \int_0^1 e^{z-1} \nabla (F_t - G_t)(z * (\pi_1, \dots, \pi_n)) dz \\ &= \int_0^1 e^{z-1} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz - \int_0^1 e^{z-1} dz \nabla G_t(\pi_1, \dots, \pi_n) \quad (36) \\ &= \int_0^1 e^{z-1} \nabla F_t(z * (\pi_1, \dots, \pi_n)) dz - (1 - 1/e) \nabla G_t(\pi_1, \dots, \pi_n), \end{aligned}$$

where the first equality follows from $\nabla G_t(z * (\pi_1, \dots, \pi_n)) = \nabla G_t(\pi_1, \dots, \pi_n)$.

Merging Eq.(35) and Eq.(36), we get the Eq.(34). $\square$

From the result of Theorem 10, for any stationary point $(\pi_1^s, \dots, \pi_n^s)$ of the objective $(F_t^s + \frac{G_t}{e})$ over $\prod_{i=1}^n \Delta_{\kappa_i}$ and any subset S within the constraint of problem (1), then we have the following inequality:

$$0 \geq \left\langle \mathbf{1}_S - (\pi_1^s, \dots, \pi_n^s), \nabla \left(F_t^s + \frac{G_t}{e} \right) (\pi_1^s, \dots, \pi_n^s) \right\rangle \geq \left(1 - \frac{c}{e} \right) f_t(S) - F_t(\pi_1^s, \dots, \pi_n^s).$$

In other words, $F_t(\pi_1^s, \dots, \pi_n^s) \geq (1 - \frac{c}{e}) f_t(S^*)$ where S^* is the optimal subset of problem (1).

G Convergence Analysis of MA-SPL Algorithm

In this section, we verify Theorem 4.

Before that, we firstly show that the surrogate gradient estimations in Line 16 of Algorithm 1, namely, $\mathbf{d}_i(t)$ is bounded and our proposed policy-based continuous extension F_t is smooth.

Lemma 2. If we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq V, e \in V \setminus S, t \in [T]} (f_t(e|S))$, then we have $\|\mathbf{d}_i(t)\|_2 \leq \sqrt{\kappa_i} (\int_0^1 w(z) dz + \frac{1}{e}) M$ where $\mathbf{d}_i(t)$ is the surrogate gradient estimations in Line 16 of Algorithm 1.

Proof. From Lines 13-15, we know that $d_{i,m}(t) \triangleq (\int_{z=0}^1 w(z)dz) f_t(v_{i,m}|S_i(t))$ or $d_{i,m}(t) \triangleq ((\int_{z=0}^1 w(z)dz) f_t(v_{i,m}|S_i(t)) + e^{-1} f_t(v_{i,m}|\mathcal{V} - \{v_{i,m}\}))$. Thus, we have $|d_{i,m}(t)| \leq (\int_0^1 w(z)dz + \frac{1}{e})M$ for any $i \in \mathcal{N}$ and $m \in [\kappa_i]$ such that $\|\mathbf{d}_i(t)\|_2 = \sqrt{\sum_{m=1}^{\kappa_i} |d_{i,m}(t)|^2} \leq \sqrt{\kappa_i}(\int_0^1 w(z)dz + \frac{1}{e})M$. $\square$

Before investigating the smoothness of our proposed policy-based continuous extension F_t , we firstly show its second-order partial derivative, i.e.,

Lemma 3. For any vector $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$, the second-order derivative of F_t at variables $\pi_{i_1, m_1}, \pi_{i_2, m_2}, \forall i_1 \in \mathcal{N}, \forall m_1 \in [\kappa_{i_1}], \forall i_2 \in \mathcal{N}, \forall m_2 \in [\kappa_{i_2}]$, can be defined as:

i): When $i_1 = i_2$, $\frac{\partial^2 F_t}{\partial \pi_{i_1, m_1} \partial \pi_{i_1, m_2}}(\pi_1, \dots, \pi_n) = 0$;

ii): When $i_1 \neq i_2$,

$$\begin{aligned} & \frac{\partial^2 F_t}{\partial \pi_{i_1, m_1} \partial \pi_{i_2, m_2}}(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \\ &= \mathbb{E}_{a_i \sim \pi_i, \forall i \in \mathcal{N}} \left(f_t(v_{i_1, m_1} | (\cup_{i \notin \{i_1, i_2\}} \{a_i\}) \cup \{v_{i_2, m_2}\}) - f_t(v_{i_1, m_1} | \cup_{i \notin \{i_1, i_2\}} \{a_i\}) \right). \end{aligned}$$

Remark 13. From the Lemma 3, if we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$, then we can show that $|\frac{\partial^2 F_t}{\partial \pi_{i_1, m_1} \partial \pi_{i_2, m_2}}(\pi_{i,1}(t), \dots, \pi_{i,n}(t))| \leq M$.

Lemma 4. If we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$, then we can show that the Hessian matrix of the policy-based continuous extension F_t of the previous monotone set function f_t satisfies that,

$$\|\nabla^2 F_t(\pi_1, \dots, \pi_n)\|_{2,\infty}^2 \leq M^2 \sum_{i=1}^n \kappa_i, \quad (37)$$

where $(\pi_1, \dots, \pi_n) \in \prod_{i=1}^n \Delta_{\kappa_i}$.

Remark 14. For any matrix $A \in \mathbb{R}^{n \times n}$, the $(2, \infty)$ -norm of A is defined as $\|A\|_{2,\infty} = \sup\{\|A\mathbf{x}\|_\infty : \mathbf{x} \in \mathbb{R}^n, \|\mathbf{x}\|_2 = 1\}$ where $\|\cdot\|_2$ denotes the L_2 norm.

Proof. From the definition of the norm $\|\cdot\|_{2,\infty}$, we can show that

$$\|\nabla^2 F_t(\pi_1, \dots, \pi_n)\|_{2,\infty}^2 = \max_{j \in \mathcal{N}} \|\nabla^2 F_t(\pi_1, \dots, \pi_n)[j, :]\|_2^2 \leq M^2 \sum_{i=1}^n \kappa_i,$$

where $\nabla^2 F_t(\pi_1, \dots, \pi_n)[j, :]$ is the j -th line of the Hessian matrix $\nabla^2 F_t(\pi_1, \dots, \pi_n)$ and the final inequality follows from Remark 13. $\square$

From Lemma 4, we can show that our proposed policy-based continuous extension F_t is $(M \sum_{i=1}^n \kappa_i)$ -smooth, namely,

Lemma 5. If we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$, then we can show that our proposed policy-based continuous extension F_t is $(M \sum_{i=1}^n \kappa_i)$ -smooth.

Proof. For any two point $(\pi_1^a, \dots, \pi_n^a) \in \prod_{i=1}^n \Delta_{\kappa_i}$ and $(\pi_1^b, \dots, \pi_n^b) \in \prod_{i=1}^n \Delta_{\kappa_i}$, if $\pi_i^a \leq \pi_i^b$ for any $i \in \mathcal{N}$, the following inequality holds:

$$\begin{aligned} & \|\nabla F_t(\pi_1^a, \dots, \pi_n^a) - \nabla F_t(\pi_1^b, \dots, \pi_n^b)\|_2 \\ &= \left\| \int_0^1 \nabla^2 F_t(\pi_1^a + \lambda \pi_1^b, \dots, \pi_n^a + \lambda \pi_n^b) d\lambda (\pi_1^a - \pi_1^b, \dots, \pi_n^a - \pi_n^b) \right\|_2 \\ &\leq \left(\sum_{i=1}^n \kappa_i \right)^{1/2} \left\| \int_0^1 \nabla^2 F_t(\pi_1^a + \lambda \pi_1^b, \dots, \pi_n^a + \lambda \pi_n^b) (\pi_1^a - \pi_1^b, \dots, \pi_n^a - \pi_n^b) d\lambda \right\|_\infty \\ &\leq M \left(\sum_{i=1}^n \kappa_i \right) \|(\pi_1^a, \dots, \pi_n^a) - (\pi_1^b, \dots, \pi_n^b)\|_2, \end{aligned}$$

where the first inequality comes from $\|\mathbf{x}\|_2 \leq \sqrt{n} \|\mathbf{x}\|_\infty$ for any n -dimensional vector $\mathbf{x}$ and the final inequality follows from the Eq.(37). Thus, F_t is $(M \sum_{i=1}^n \kappa_i)$ -smooth. $\square$

With Lemma 2 and Lemma 5, we next prove Theorem 4. At first, we show that each policy vector $(\pi_{i,1}(t), \dots, \pi_{i,n}(t)), \forall i \in \mathcal{N}, \forall t \in [T]$ of Algorithm 1 is included in the constraint $\prod_{i=1}^n \Delta_{\kappa_i}$. That is, we have the following lemma:

Lemma 6. In Algorithm 1, Under Assumption 1, $(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \in \prod_{i=1}^n \Delta_{\kappa_i}$.

Proof. We prove this theorem by induction. At first, from the Line 2 in Algorithm 1, we know that $\mathbf{x}_{i,j}(1) \in \Delta_{\kappa_j}$ for any $i, j \in \mathcal{N}$, so $(\pi_{i,1}(1), \dots, \pi_{i,n}(1)) \in \prod_{i=1}^n \Delta_{\kappa_i}, \forall i \in \mathcal{N}$. Then, if for some $t \in [T]$, $(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \in \prod_{i=1}^n \Delta_{\kappa_i}$. From Line 17, we know that $\pi_{i,m}(t+1) := \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \pi_{j,m}(t)$ for any $m \neq i$. Then, due to Assumption 1, we have $\sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} = 1$ and $w_{ij} \geq 0$ such that $\pi_{i,m}(t+1) \in \Delta_{\kappa_m}, \forall m \neq i$. Moreover, from Line 19, we also know that $\text{upper}_{i,i}(t+1) \in \Delta_{\kappa_i}$. Thus, $(\pi_{i,1}(t+1), \dots, \pi_{i,n}(t+1)) \in \prod_{i=1}^n \Delta_{\kappa_i}$. $\square$

Before going into the proof of Theorem 4, we define some commonly used symbols, namely,

$$\begin{aligned} \bar{\pi}_{\cdot,j}(t) &\triangleq \frac{\sum_{i=1}^n \pi_{i,j}(t)}{n}, \quad \pi_{\cdot,j}^{\text{cate}}(t) \triangleq [\pi_{1,j}(t); \pi_{2,j}(t); \dots; \pi_{n,j}(t)] \in \mathbb{R}^{n \times \kappa_j}, \quad \forall j \in \mathcal{N}; \\ \mathbf{z}_{i,m}(t) &\triangleq \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \pi_{j,m}(t), \quad \forall i, m \in \mathcal{N}; \\ \bar{\mathbf{z}}_{\cdot,j}(t) &\triangleq \frac{\sum_{i=1}^n \mathbf{z}_{i,j}(t)}{n}, \quad \mathbf{z}_{\cdot,j}^{\text{cate}}(t) \triangleq [\mathbf{z}_{1,j}(t); \mathbf{z}_{2,j}(t); \dots; \mathbf{z}_{n,j}(t)] \in \mathbb{R}^{n \times \kappa_j}, \quad \forall j \in \mathcal{N}; \\ \mathbf{r}_{i,j}(t) &\triangleq \pi_{i,j}(t+1) - \mathbf{z}_{i,j}(t), \quad \mathbf{r}_{\cdot,j}^{\text{cate}}(t) \triangleq [\mathbf{r}_{1,j}(t); \mathbf{r}_{2,j}(t); \dots; \mathbf{r}_{n,j}(t)] \in \mathbb{R}^{n \times \kappa_j}, \quad \forall j \in \mathcal{N}. \end{aligned}$$

With these symbols, we can verify that

Lemma 7. $\mathbb{E}(\|\mathbf{r}_{i,j}(t)\|_2) = \mathbb{E}(\|\pi_{i,j}(t+1) - \mathbf{z}_{i,j}(t)\|_2) \leq \sqrt{\kappa_i} (\int_0^1 w(z) dz + \frac{1}{e}) \eta_t M$ where M is the maximum marginal contribution of each monotone set function f_t .

Proof. When $i \neq j$, from Line 17, we know that $\pi_{i,j}(t+1) = \mathbf{z}_{i,j}(t)$ such that $\mathbf{r}_{i,j}(t) = \mathbf{0}_{\kappa_j}$ and $\mathbb{E}(\|\mathbf{r}_{i,j}(t)\|_2) = 0 \leq \sqrt{\kappa_i} (\int_0^1 w(z) dz + \frac{1}{e}) \eta_t M$. As for $i = j$, from Line 18 and Line 19, we have

$$\pi_{i,i}(t+1) := \arg \min_{\mathbf{b} \in \Delta_{\kappa_i}} \|\mathbf{b} - (\mathbf{z}_{i,i}(t) + \eta_t \mathbf{d}_i(t))\|_2.$$

Note that $\mathbf{z}_{i,i}(t) \in \Delta_{\kappa_j}$. Thus, we have

$$\|\pi_{i,i}(t+1) - \mathbf{z}_{i,i}(t)\|_2 \leq \|\mathbf{z}_{i,i}(t) - (\mathbf{z}_{i,i}(t) + \eta_t \mathbf{d}_i(t))\|_2 = \eta_t \|\mathbf{d}_i(t)\|_2 \leq \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_t M, \quad (38)$$

where the final inequality follows from Lemma 2. $\square$

Lemma 8. Under Assumption 1, for any $t \in [T]$ and $i \in \mathcal{N}$, we have that

$$\sum_{i \in \mathcal{N}} \mathbb{E}(\|\pi_{i,j}(t+1) - \bar{\pi}_{\cdot,j}(t+1)\|) \leq \sum_{m=1}^t \sqrt{n} \beta^{t-m} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_m M,$$

$$\sum_{i \in \mathcal{N}} \mathbb{E}(\|z_{i,j}(t+1) - \bar{z}_{\cdot,j}(t+1)\|) \leq \sum_{m=1}^t \sqrt{n} \beta^{t-m} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_m M,$$

where $\tau \triangleq \max(|\lambda_2(\mathbf{W})|, |\lambda_n(\mathbf{W})|)$ is the second largest magnitude of the eigenvalues of the weight matrix $\mathbf{W}$.

Proof. From the definition of $\mathbf{r}_{i,j}(t)$, we can show that

$$\mathbf{x}_{i,j}(t+1) = \mathbf{r}_{i,j}(t) + \mathbf{z}_{i,j}(t) = \mathbf{r}_{i,j}(t) + \sum_{k \in \mathcal{N}_i \cup \{i\}} w_{ik} \mathbf{x}_{k,j}. \quad (39)$$

As a result, from the Eq.(39), we also have that

$$\begin{aligned} \pi_{\cdot,j}^{cate}(t+1) &= \mathbf{r}_{\cdot,j}^{cate}(t) + (\mathbf{W} \otimes \mathbf{I}_{\kappa_j}) \pi_{\cdot,j}^{cate}(t) \\ &= \sum_{m=1}^t (\mathbf{W} \otimes \mathbf{I}_{\kappa_j})^{t-m} \mathbf{r}_{\cdot,j}^{cate}(m) \\ &= \sum_{m=1}^t (\mathbf{W}^{t-m} \otimes \mathbf{I}_{\kappa_j}) \mathbf{r}_{\cdot,j}^{cate}(m), \end{aligned} \quad (40)$$

where the symbol $\otimes$ denotes the Kronecker product.

If we also define $\bar{\pi}_{\cdot,j}^{cate}(t) = [\bar{\pi}_{\cdot,j}(t); \bar{\pi}_{\cdot,j}(t); \dots; \bar{\pi}_{\cdot,j}(t)] \in \mathbb{R}^{n\kappa_j}$ and from the Eq.(40), we also have that

$$\begin{aligned} \bar{\pi}_{\cdot,j}^{cate}(t) &= \left(\frac{\mathbf{1}_n \mathbf{1}_n^T}{n} \otimes \mathbf{I}_{\kappa_j} \right) \pi_{\cdot,j}^{cate}(t) \\ &= \sum_{m=1}^t \left(\frac{\mathbf{1}_n \mathbf{1}_n^T}{n} \otimes \mathbf{I}_{\kappa_j} \right) \mathbf{r}_{\cdot,j}^{cate}(m). \end{aligned} \quad (41)$$

Then, from the Eq.(40) and Eq.(41), we have that, for any $i \in \mathcal{N}$,

$$\pi_{i,j}^{cate}(t+1) - \bar{\pi}_{i,j}^{cate}(t) = \sum_{m=1}^t \sum_{j \in \mathcal{N}_i \cup \{i\}} \left([\mathbf{W}^{t-m}]_{ij} - \frac{1}{N} \right) \mathbf{r}_{\cdot,j}^{cate}(m). \quad (42)$$

Eq.(42) indicates that

$$\begin{aligned} \mathbb{E}(\|\pi_{i,j}^{cate}(t+1) - \bar{\pi}_{i,j}^{cate}(t)\|) &= \mathbb{E} \left(\left\| \sum_{m=1}^t \sum_{j \in \mathcal{N}_i \cup \{i\}} \left([\mathbf{W}^{t-m}]_{ij} - \frac{1}{n} \right) \mathbf{r}_{\cdot,j}^{cate}(m) \right\| \right) \\ &\leq \mathbb{E} \left(\sum_{m=1}^t \sum_{j \in \mathcal{N}_i \cup \{i\}} \left| \left([\mathbf{W}^{t-m}]_{ij} - \frac{1}{n} \right) \right| * \|\mathbf{r}_{\cdot,j}^{cate}(m)\|_2 \right) \\ &\leq \sum_{m=1}^t \sum_{j \in \mathcal{N}_i \cup \{i\}} \left| \left([\mathbf{W}^{t-m}]_{ij} - \frac{1}{n} \right) \right| \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_m M \\ &\leq \sum_{m=1}^t \sqrt{n} \tau^{t-m} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_m M, \end{aligned}$$

where the second inequality comes from Lemma 7 and the final inequality follows from $\sum_{j \in \mathcal{N}_i \cup \{i\}} |[\mathbf{W}^{t-m}]_{ij} - \frac{1}{n}| \leq \sqrt{n} \tau^{t-m}$ (See Proposition 1 in [84]). Due to $\mathbf{z}_{i,j}(t+1) = \sum_{k \in \mathcal{N}_i \cup \{i\}} w_{ik} \pi_{k,j}(t+1)$ we also can have $\mathbb{E}(\|z_{i,j}(t+1) - \bar{z}_{\cdot,j}(t+1)\|_2) \leq \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \mathbb{E}(\|\pi_{i,j}(t+1) - \bar{\pi}_{\cdot,j}(t+1)\|_2) \leq \sum_{m=1}^t \sqrt{n} \tau^{t-m} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \eta_m M. \quad \square$

Lemma 9. When each objective function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, we set the ratio $\rho = (1 - \frac{c}{e}), (1 - e^{-\alpha}), (\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2})$, respectively. Moreover, we use the symbol M to denote the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$. Then, under Assumption 1, if each set function f_t is bounded $\forall t \in [T]$, the following inequality holds for Algorithm 1:

$$\begin{aligned} \rho \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T F_t(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) &\leq 5M^2 n^{\frac{5}{2}} \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{t=1}^T \sum_{m=1}^t \tau^{t-m} \eta_m \\ &+ n\sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M \sum_{t=1}^T \frac{\eta_t}{2} + \sum_{t=1}^T \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\pi_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right). \end{aligned}$$

where $\mathcal{A}_t^*$ is the optimal solution of problem (1) and $\tau \triangleq \max(|\lambda_2(\mathbf{W})|, |\lambda_n(\mathbf{W})|)$ is the second largest magnitude of the eigenvalues of the weight matrix $\mathbf{W}$.

Proof. In order to unify the proofs in different setting, i.e., submodular objective with curvature c , α -weakly DR-submodular objective and (γ, β) -weakly submodular objective, we define some auxiliary symbols. At first, when each objective function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, we set $\rho = (1 - \frac{c}{e}), (1 - e^{-\alpha}), (\frac{\gamma^2(1-e^{-(\beta(1-\gamma)+\gamma^2)})}{\beta(1-\gamma)+\gamma^2})$, respectively. Moreover, compared to weakly submodular scenarios, Theorem 3 designs a different auxiliary function for submodular scenarios. Therefore, when each f_t is monotone submodular with curvature c or weakly submodular, we define $F_t^a = F_t^s$ or $F_t^a = (F_t^s + \frac{G_t}{e})$, respectively. Note that the F_t^s is the surrogate function considered in Theorem 3 for different settings and G_t is a linear function defined as $G_t(\pi_1, \dots, \pi_n) \triangleq \sum_{i=1}^n \sum_{m=1}^{\kappa_i} (f_t(v_{i,k}|\mathcal{V} - \{v_{i,m}\}))\pi_{i,m}$.

From Theorem 7 and Theorem 10, we have that

$$\begin{aligned} \rho f_t(\mathcal{A}_t^*) - F_t(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) &\leq \left\langle \nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)), \mathbf{1}_{\mathcal{A}_t^*} - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\rangle \\ &= \underbrace{\left\langle \nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) - \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}, \mathbf{1}_{\mathcal{A}_t^*} - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\rangle}_{\textcircled{1}} \\ &\quad + \underbrace{\left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}, \mathbf{1}_{\mathcal{A}_t^*} - (\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right\rangle}_{\textcircled{2}} \\ &\quad + \underbrace{\left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}, (\pi_{i,1}(t), \dots, \pi_{i,1}(t)) - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\rangle}_{\textcircled{3}} \end{aligned} \tag{43}$$

where the symbol $\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}$ is the projection over the policy π_i , i.e., $\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}$ represents a $(\sum_{i=1}^n \kappa_i)$ -dimensional vector that only keeps the first-order derivative at variable $\pi_{i,m}$, $\forall m \in [\kappa_i]$ and set other coordinates to 0, that is to say,

$$\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i} \triangleq \left(\dots, 0, \dots, \underbrace{\frac{\partial F_t}{\partial \pi_{i,1}}(\mathbf{x}), \dots, \frac{\partial F_t}{\partial \pi_{i,\kappa_i}}(\mathbf{x})}_{\kappa_i}, \dots, 0, \dots \right),$$

where $\mathbf{x} \triangleq (\pi_{i,1}(t), \dots, \pi_{i,1}(t))$.

For ①, we have

$$\begin{aligned}
& \left\langle \nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) - \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i}, \mathbf{1}_{\mathcal{A}_t^*} - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\rangle \\
& \leq \left\| \nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) - \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i} \right\|_2 \left\| \mathbf{1}_{\mathcal{A}_t^*} - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\|_2 \\
& \leq \sum_{i \in \mathcal{N}} \left(\left\| \left[\nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right]_{\pi_i} - \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i} \right\|_2 \right) \left(\left\| \mathbf{1}_{\mathcal{A}_t^*} \right\|_2 + \left\| (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\|_2 \right) \\
& \leq 2n \sum_{i \in \mathcal{N}} \left(\left\| \nabla F_t^a(\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) - \nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\|_2 \right) \\
& = 2n \sum_{i \in \mathcal{N}} \left(\left\| \int_{z=1}^1 w(z) \nabla F_t(z * \bar{\pi}_{\cdot,1}(t), \dots, z * \bar{\pi}_{\cdot,n}(t)) dz - \int_{z=1}^1 w(z) \nabla F_t(z * \pi_{i,1}(t), \dots, z * \pi_{i,n}(t)) dz \right\|_2 \right) \\
& \leq 2Mn \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) z dz \right) \sum_{i \in \mathcal{N}} \left(\left\| (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\|_2 \right) \\
& \leq 2Mn \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) z dz \right) \sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}} \left\| \pi_{i,j}(t) - \bar{\pi}_{\cdot,j}(t) \right\|_2 \\
& \leq 2M^2 n^{\frac{5}{2}} \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) z dz \right) \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \sum_{m=1}^t \tau^{t-m} \eta_m,
\end{aligned} \tag{44}$$

where the fourth inequality follows Lemma 5 and the final inequality from Lemma 8.

For ③, under Lemma 2 and Lemma 8, we have,

$$\begin{aligned}
& \left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i}, (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) - (\bar{\pi}_{\cdot,1}(t), \dots, \bar{\pi}_{\cdot,n}(t)) \right\rangle \\
& \leq \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \left\| \pi_{i,j}(t) - \bar{\pi}_{\cdot,j}(t) \right\|_2 \\
& \leq \sqrt{\kappa_i} M^2 n^{3/2} \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{m=1}^t \tau^{t-m} \eta_m.
\end{aligned} \tag{45}$$

As for ②, we have,

$$\begin{aligned}
& \mathbb{E} \left(\left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i}, \mathbf{1}_{\mathcal{A}_t^*} - (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\rangle \right) \\
& = \mathbb{E} \left(\mathbb{E} \left(\left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i}, \mathbf{1}_{\mathcal{A}_t^*} - (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\rangle \middle| \pi_{i,j}(t), \forall i, j \in \mathcal{N} \right) \right) \\
& = \mathbb{E} \left(\left\langle \mathbb{E} \left(\sum_{i \in \mathcal{N}} \left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right]_{\pi_i} \middle| \pi_{i,j}(t), \forall i, j \in \mathcal{N} \right), \mathbf{1}_{\mathcal{A}_t^*} - (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\rangle \right) \\
& = \mathbb{E} \left(\left\langle (\mathbf{d}_1(t), \dots, \mathbf{d}_n(t)), \mathbf{1}_{\mathcal{A}_t^*} - (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\rangle \right) \\
& = \underbrace{\mathbb{E} \left(\left\langle (\mathbf{d}_1(t), \dots, \mathbf{d}_n(t)), \mathbf{1}_{\mathcal{A}_t^*} - (\mathbf{z}_{i,1}(t), \dots, \mathbf{z}_{i,n}(t)) \right\rangle \right)}_{\textcircled{4}} \\
& + \underbrace{\mathbb{E} \left(\left\langle (\mathbf{d}_1(t), \dots, \mathbf{d}_n(t)), (\mathbf{z}_{i,1}(t), \dots, \mathbf{z}_{i,n}(t)) - (\pi_{i,1}(t), \dots, \pi_{i,n}(t)) \right\rangle \right)}_{\textcircled{5}}.
\end{aligned} \tag{46}$$

For ⑤, we have that,

$$\begin{aligned}
\textcircled{5} &\leq \mathbb{E} \left(\|(\mathbf{d}_1(t), \dots, \mathbf{d}_n(t))\|_2 \|(\mathbf{z}_{i,1}(t), \dots, \mathbf{z}_{i,n}(t)) - (\boldsymbol{\pi}_{i,1}(t), \dots, \boldsymbol{\pi}_{i,n}(t))\|_2 \right) \\
&\leq \sqrt{n\kappa_i} M \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \mathbb{E} \left(\|(\mathbf{z}_{i,1}(t), \dots, \mathbf{z}_{i,n}(t)) - (\boldsymbol{\pi}_{i,1}(t), \dots, \boldsymbol{\pi}_{i,n}(t))\|_2 \right) \\
&\leq \sqrt{n\kappa_i} M \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \mathbb{E} \left(\|(\mathbf{z}_{i,1}(t) - \bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \mathbf{z}_{i,n}(t) - \bar{\boldsymbol{\pi}}_{\cdot,n}(t))\|_2 \right) \\
&\quad + \sqrt{n\kappa_i} M \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \mathbb{E} \left(\|(\boldsymbol{\pi}_{i,1}(t) - \bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \boldsymbol{\pi}_{i,n}(t) - \bar{\boldsymbol{\pi}}_{\cdot,n}(t))\|_2 \right) \\
&\leq \sqrt{n\kappa_i} M \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \mathbb{E} \left(\|\boldsymbol{\pi}_{i,j}(t) - \bar{\boldsymbol{\pi}}_{\cdot,j}(t)\|_2 + \|\mathbf{z}_{i,j}(t) - \bar{\boldsymbol{\pi}}_{\cdot,j}(t)\|_2 \right) \\
&\leq 2n^2 \sqrt{\kappa_i} M^2 \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{m=1}^t \tau^{t-m} \eta_m
\end{aligned} \tag{47}$$

For ④: At first, from the Line 19 of Algorithm 1, we know that, for any $\mathbf{x} \in \Delta_{\kappa_i}$, we have

$$\|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{x}\|_2 \leq \|\mathbf{z}_{i,i}(t) + \eta_t \mathbf{d}_i(t) - \mathbf{x}\|_2.$$

If we slightly abuse the notation $\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}$ to denote κ_i -dimensional indicator vector over $(\mathcal{A}_t^* \cap \mathcal{V}_i)$.

Note that $\mathbf{1}_{\mathcal{A}_t^*}$ is a $(\sum_{i=1}^n \kappa_i)$ -dimensional indicator vector over $\mathcal{A}_t^*$. Then, we have

$$\begin{aligned}
&\|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \\
&\leq \|\mathbf{z}_{i,i}(t) + \eta_t \mathbf{d}_i(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \\
&= \|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 + 2\eta_t \langle \mathbf{d}_i(t), \mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} \rangle + \eta_t^2 \|\mathbf{d}_i(t)\|_2^2.
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\langle \mathbf{d}_i(t), \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{z}_{i,i}(t) \rangle \\
&\leq \frac{\eta_t}{2} \|\mathbf{d}_i(t)\|_2^2 + \frac{1}{2\eta_t} \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \\
&\leq \frac{\eta_t}{2} \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M + \frac{1}{2\eta_t} \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right),
\end{aligned}$$

where the final inequality follows from Lemma 2.

Note that

$$\begin{aligned}
&\mathbb{E} \left(\left\langle (\mathbf{d}_1(t), \dots, \mathbf{d}_n(t)), \mathbf{1}_{\mathcal{A}_t^*} - (\mathbf{z}_{i,1}(t), \dots, \mathbf{z}_{i,1}(t)) \right\rangle \right) \\
&= \sum_{i=1}^n \mathbb{E} \left(\langle \mathbf{d}_i(t), \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{z}_{i,i}(t) \rangle \right) \\
&\leq \frac{\eta_t}{2} n \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M + \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right).
\end{aligned} \tag{48}$$

Merging Eq.(48), Eq.(47), Eq.(46), Eq.(45), Eq.(44) into Eq.(43), we finally have

$$\begin{aligned}
\rho f_t(\mathcal{A}_t^*) - F_t(\bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \bar{\boldsymbol{\pi}}_{\cdot,n}(t)) &\leq 5M^2 n^{\frac{5}{2}} \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{m=1}^t \beta^{t-m} \eta_m \\
&\quad + \frac{\eta_t}{2} n \sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M + \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right).
\end{aligned}$$

Therefore, we get Lemma 9. $\square$

Next, we prove an upper bound of $\sum_{t=1}^T \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right)$, that is to say,

Lemma 10. Under Assumption 1, we have that

$$\sum_{t=1}^T \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\pi_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \leq \frac{4n}{\eta_{T+1}} + \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E} \left(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \right),$$

where $\mathcal{A}_t^*$ is the optimal solution of problem (1) and we slightly abuse the notation $\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}$ to denote κ_i -dimensional indicator vector over $(\mathcal{A}_t^* \cap \mathcal{V}_i)$. Note that $\mathbf{1}_{\mathcal{A}_t^*}$ is a $(\sum_{i=1}^n \kappa_i)$ -dimensional indicator vector over $\mathcal{A}_t^*$.

Proof.

$$\begin{aligned} & \sum_{t=1}^T \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\pi_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \\ &= \underbrace{\sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_t} \mathbb{E}(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2) - \frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2^2) \right)}_{\textcircled{1}} \\ &+ \underbrace{\sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2^2 - \|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2) \right)}_{\textcircled{2}} \\ &+ \underbrace{\sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\pi_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2) \right)}_{\textcircled{3}} \\ &+ \underbrace{\sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \mathbb{E}(\|\pi_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2)}_{\textcircled{4}}. \end{aligned}$$

Firstly, we have $\textcircled{1} \leq \frac{\sum_{i=1}^n \mathbb{E}(\|\mathbf{z}_{i,i}(1) - \mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}\|_2^2)}{\eta_1} \leq \frac{4n}{\eta_1}$.

Moreover, we have $\textcircled{4} \leq 4n \left(\frac{1}{\eta_{T+1}} - \frac{1}{\eta_1} \right)$.

As for $\textcircled{2}$, we have

$$\begin{aligned} & \sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2^2 - \|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2) \right) \\ &= \sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2^2 + 2 \langle \mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}, \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)} \rangle) \right) \\ &\leq \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E}(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2), \end{aligned}$$

where the final inequality from $\langle \mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}, \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)} \rangle \leq \|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2 * \|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \leq 2\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2$.

Then, we have

$$\begin{aligned}
\textcircled{3} &= \sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E} \left(\|\mathbf{z}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \right) \\
&= \sum_{t=1}^T \sum_{i=1}^n \left(\frac{1}{\eta_{t+1}} \mathbb{E} \left(\left\| \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \boldsymbol{\pi}_{j,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} \right\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \right) \\
&\leq \sum_{t=1}^T \left(\frac{1}{\eta_{t+1}} \mathbb{E} \left(\sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}_i \cup \{i\}} \left(w_{ij} \|\boldsymbol{\pi}_{j,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) - \sum_{i \in \mathcal{N}} \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \right) \\
&= \sum_{t=1}^T \left(\frac{1}{\eta_{t+1}} \mathbb{E} \left(\sum_{i \in \mathcal{N}} \left(\sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ji} \right) \|\boldsymbol{\pi}_{j,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) - \sum_{i \in \mathcal{N}} \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \\
&= 0,
\end{aligned}$$

where the first inequality follows from the convexity of $\|\cdot\|_2^2$, the third equality is from $w_{ij} = w_{ji}$ and the final equality follows from Assumption 1.

We finally get

$$\sum_{t=1}^T \frac{1}{2\eta_t} \sum_{i=1}^n \left(\|\mathbf{z}_{i,i}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 - \|\boldsymbol{\pi}_{i,i}(t+1) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}\|_2^2 \right) \leq \frac{4n}{\eta_{T+1}} + \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E} \left(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \right).$$

□

As a result, we can show that

Lemma 11. Under Assumption 1, when each set objective function f_t is monotone submodular with curvature c , α -weakly DR-submodular or (γ, β) -weakly submodular, if we set the weight function $w(z)$ according to Theorem 3, we have that

$$\begin{aligned}
\rho \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E}(f_t(\cup_{i=1}^n \{a_i(t)\})) &\leq 6M^2 n^{\frac{5}{2}} \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{t=1}^T \sum_{m=1}^t \beta^{t-m} \eta_m \\
&+ n\sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M \sum_{t=1}^T \frac{\eta_t}{2} + \frac{4n}{\eta_{T+1}} + \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E} \left(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \right).
\end{aligned}$$

Proof. Merging Lemma 10 into Lemma 9, we can get that

$$\begin{aligned}
\rho \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T F_t(\bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \bar{\boldsymbol{\pi}}_{\cdot,n}(t)) &\leq 5M^2 n^{\frac{5}{2}} \left(\sum_{i=1}^n \kappa_i \right) \left(\int_0^1 w(z) dz + \frac{1}{e} \right)^2 \sum_{t=1}^T \sum_{m=1}^t \tau^{t-m} \eta_m \\
&+ n\sqrt{\kappa_i} \left(\int_0^1 w(z) dz + \frac{1}{e} \right) M \sum_{t=1}^T \frac{\eta_t}{2} + \frac{4n}{\eta_{T+1}} + \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E} \left(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \right).
\end{aligned}$$

From Theorem 1, we also can show that $|F_t(\mathbf{x}) - F_t(\mathbf{y})| \leq nM\|\mathbf{x} - \mathbf{y}\|_2$ for any $t \in [T]$ such that we have

$$\begin{aligned}
&|F_t(\bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \bar{\boldsymbol{\pi}}_{\cdot,n}(t)) - F_t(\boldsymbol{\pi}_{1,1}(t), \boldsymbol{\pi}_{2,2}(t), \dots, \boldsymbol{\pi}_{n,n}(t))| \\
&\leq nM\|(\bar{\boldsymbol{\pi}}_{\cdot,1}(t), \dots, \bar{\boldsymbol{\pi}}_{\cdot,n}(t)) - (\boldsymbol{\pi}_{1,1}(t), \dots, \boldsymbol{\pi}_{n,n}(t))\|_2 \\
&= nM \sum_{i=1}^n \|\bar{\boldsymbol{\pi}}_{\cdot,i}(t) - \boldsymbol{\pi}_{i,i}(t)\|_2 \leq nM \sum_{j=1}^n \sum_{i=1}^n \|\bar{\boldsymbol{\pi}}_{\cdot,i}(t) - \boldsymbol{\pi}_{j,i}(t)\|_2 \leq n^{5/2} M^2 \left(\int_0^1 w(z) dz + \frac{1}{e} \right) \sum_{m=1}^t \tau^{t-m} \eta_m,
\end{aligned}$$

where the final inequality follows from Lemma 8.

Moreover, from the monotonicity of F_t (See 2) in Theorem 1) and Lines 5-8, we also can infer that

$$\mathbb{E}(f_t(\cup_{i=1}^n \{a_i(t)\})) = F_t(\mathbf{p}_1(t), \dots, \mathbf{p}_n(t)) \geq F_t(\boldsymbol{\pi}_{1,1}(t), \boldsymbol{\pi}_{2,2}(t), \dots, \boldsymbol{\pi}_{n,n}(t)),$$

where the final inequality comes from $\mathbf{p}_i(t) := \frac{\pi_{i,i}(t)}{\|\pi_{i,i}(t)\|_1} \geq \pi_{i,i}(t)$. Therefore, we have that

$$\begin{aligned} \rho \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E}(f_t(\cup_{i=1}^n \{a_i(t)\})) &\leq 6M^2 n^{\frac{5}{2}} (\sum_{i=1}^n \kappa_i) (\int_0^1 w(z) dz + \frac{1}{e})^2 \sum_{t=1}^T \sum_{m=1}^t \tau^{t-m} \eta_m \\ &+ n\sqrt{\kappa_i} (\int_0^1 w(z) dz + \frac{1}{e}) M \sum_{t=1}^T \frac{\eta_t}{2} + \frac{4n}{\eta_{T+1}} + \sum_{t=1}^T \sum_{i=1}^n \frac{6}{\eta_{t+1}} \mathbb{E}(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2). \end{aligned}$$

□

From Lemma 11, if we set the step size $\eta_t = \mathcal{O}(\sqrt{\frac{(1-\tau)\mathcal{P}_T}{T}})$ where $\mathcal{P}_T \triangleq \sum_{t=1}^T |\mathcal{A}_{t+1}^* \triangle \mathcal{A}_t^*|$ and $\triangle$ denotes the symmetric difference, we can show that

$$\begin{aligned} \rho \sum_{t=1}^T F_t(\mathbf{1}_{\mathcal{A}_t^*}) - \sum_{t=1}^T \mathbb{E}(f_t(\cup_{i=1}^n \{a_i(t)\})) \\ \leq \mathcal{O}(\sum_{t=1}^T \sum_{m=1}^t \tau^{t-m} \eta_m) + \mathcal{O}(\sum_{t=1}^T \eta_t) + \mathcal{O}(\frac{1}{\eta_{T+1}}) + \mathcal{O}(\sum_{t=1}^T \sum_{i=1}^n \frac{1}{\eta_{t+1}} \mathbb{E}(\|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2)) \\ = \mathcal{O}\left(\sqrt{\frac{\mathcal{P}_T T}{1-\tau}}\right), \end{aligned}$$

where the final equality follows from $\mathcal{P}_T \triangleq \sum_{t=1}^T \sum_{i=1}^n \|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_2 \triangleq \sum_{t=1}^T \sum_{i=1}^n \|\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)}\|_1 \triangleq \sum_{t=1}^T |\mathcal{A}_{t+1}^* \triangle \mathcal{A}_t^*|$.

H Convergence Analysis of MA-MPL Algorithm

In this section, we prove Theorem 5.

At first, we verify that each policy vector $(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \in \prod_{i=1} \Delta_{\kappa_i}$, namely, we have the following lemma:

Lemma 12. *When the communication graph $G(\mathcal{V}, \mathcal{E})$ is connected, we can show that $(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \in \prod_{i=1} \Delta_{\kappa_i}$ for any $i \in \mathcal{N}$ and $0 \leq k \leq K$ as well as $t \in [T]$.*

Proof. From the Line 1 of Algorithm 2, we can know that $(\pi_{i,1}^{(0)}(t), \dots, \pi_{i,n}^{(0)}(t)) := \mathbf{0}, \forall t \in [T]$ such that $(\pi_{i,1}^{(0)}(t), \dots, \pi_{i,n}^{(0)}(t)) \in \prod_{i=1} \Delta_{\kappa_i}$. Moreover, from Lines 8, for a fixed $t \in [T]$, we know that only agent $i \in \mathcal{N}$ can change the policy $\pi_{i,i}^{(k)}(t)$. Furthermore, from Line 10, we also can infer that, for other agent $j \neq i$, its policy $\pi_{j,i}^{(k)}(t)$ is copy of some past iteration of $\pi_{i,i}^{(k)}(t)$, namely, there exists a $k_1 \leq k$ such that $\pi_{j,i}^{(k)}(t) = \pi_{i,i}^{(k_1)}(t)$. According to the Line 8 and Line 10, we have $\pi_{i,i}^{(k)}(t) = \pi_{i,i}^{(k-1)}(t) + \frac{1}{K} \mathbf{v}_i^{(k)}(t)$, so we have $\pi_{i,i}^{(k)}(t) = \frac{1}{K} \sum_{m=1}^k \mathbf{v}_i^{(m)}(t) \in \Delta_{\kappa_i}$ for any $k \in [K]$, $t \in [T]$ and $i \in \mathcal{N}$. Therefore, we have $(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \in \prod_{i=1} \Delta_{\kappa_i}$ for any $i \in \mathcal{N}$ and $k \in [K]$ as well as $t \in [T]$. □

Note that, from Lines 11-13, each agent $i \in \mathcal{N}$ only uses the policy $\pi_{i,i}^{(K)}(t)$ to make decision, so the real policy vector taken by all agents is $(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(K)}(t))$. Motivated by this finding, we next investigate some relationships between the policy vector $(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t))$ and $(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t))$. That is,

Lemma 13. *When the communication graph $G(\mathcal{V}, \mathcal{E})$ is connected, if we set $\kappa \triangleq \sum_{i=1}^n \kappa_i$ and utilize the symbol $\mathbf{1}_\kappa$ to denote the κ -dimensional vector whose all elements are 1, we can show that*

$$0 \leq \frac{1}{n} \left\langle \mathbf{1}_\kappa, (\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - (\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \right\rangle \leq \frac{d(G)}{K},$$

where $d(G)$ is the diameter of a graph, that is to say, the length of the shortest path between the most distanced nodes.

Proof. From Lines 6-10, we can show that, for fixed $k \in [K]$ and $t \in [T]$, each $\pi_{i,j}^{(k)}(t) = \pi_{j,j}^{(k_j)}(t)$ where $k_j \in [k - d(G), k]$. Moreover, according to the Line 8 and Line 10, we have $\pi_{j,j}^{(k)}(t) = \pi_{j,j}^{(k-1)}(t) + \frac{1}{K} \mathbf{v}_j^{(k)}(t)$, so we have $\pi_{j,j}^{(k)}(t) = \frac{1}{K} \sum_{m=1}^k \mathbf{v}_j^{(m)}(t)$ for any $k \in [K]$, $t \in [T]$ and $i \in \mathcal{N}$. Therefore, we have that $\pi_{j,j}^{(k)}(t) - \pi_{i,j}^{(k)}(t) = \frac{1}{K} \sum_{m=k_j}^k \mathbf{v}_j^{(m)}(t)$ such that $0 \leq \langle \mathbf{1}_{\kappa_j}, \pi_{j,j}^{(k)}(t) - \pi_{i,j}^{(k)}(t) \rangle \leq \frac{k-k_j}{K} \leq \frac{d(G)}{K}$. So we have the result of Lemma 13. $\square$

Next, we investigate the relationships between $(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t))$ and $(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t))$.

Lemma 14. When the communication graph $G(\mathcal{V}, \mathcal{E})$ is connected, we can show that

$$\begin{aligned} (\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - (\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) &= \frac{1}{K} (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)); \\ 0 \leq \frac{1}{n} \langle \mathbf{1}_{\kappa}, (\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - (\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \rangle &\leq \frac{1}{K}, \end{aligned}$$

where $\kappa \triangleq \sum_{i=1}^n \kappa_i$ and $d(G)$ is the diameter of a graph, i.e., the length of the shortest path between the most distanced nodes.

Proof. According to the Line 8 and Line 10, we have $\pi_{j,j}^{(k)}(t) = \pi_{j,j}^{(k-1)}(t) + \frac{1}{K} \mathbf{v}_j^{(k)}(t)$. Therefore, we get result of Lemma 14. $\square$

With Lemma 13 and Lemma 14, we next give an upper bound about the error between the gradient $\nabla F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t))$ and $\nabla F_t(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t))$.

Lemma 15. When the communication graph $G(\mathcal{V}, \mathcal{E})$ is connected, if we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$, we can show that

$$\left\| \nabla F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - \nabla F_t(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \right\|_2 \leq \frac{\sqrt{\kappa n M d(G)}}{K}$$

where $\kappa = \sum_{i=1}^n \kappa_i$ and $d(G)$ is the diameter of a graph, i.e., the length of the shortest path between the most distanced nodes.

Proof. At first, we investigate the error between the first-order derivative $\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t))$ and the first-order derivative $\frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t))$ for some $i \in \mathcal{N}$ and $m \in [\kappa_i]$.

$$\begin{aligned} &\left| \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \right| \\ &= \left| \left\langle \left(\int_0^1 \frac{\partial^2 F_t}{\partial \pi_{i,m} \partial \pi_{1,1}}, \dots, \frac{\partial^2 F_t}{\partial \pi_{i,m} \partial \pi_{n,\kappa_n}} \right) (\mathbf{x}), \left(\pi_{1,1}^{(k)}(t) - \pi_{i,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t) - \pi_{i,n}^{(k)}(t) \right) \right\rangle \right| \\ &\leq M \left| \left\langle \mathbf{1}_{\kappa}, \left(\pi_{1,1}^{(k)}(t) - \pi_{i,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t) - \pi_{i,n}^{(k)}(t) \right) \right\rangle \right| \leq \frac{n M d(G)}{K}, \end{aligned}$$

where the first equality comes from $\mathbf{x} = (\lambda \pi_{1,1}^{(k)}(t) + (1 - \lambda) \pi_{i,1}^{(k)}(t), \dots, \lambda \pi_{n,n}^{(k)}(t) + (1 - \lambda) \pi_{i,n}^{(k)}(t))$ for some $\lambda \in [0, 1]$, the first inequality follows from Remark 13 and the final inequality comes from Lemma 13. Then, we have that

$$\begin{aligned} &\left\| \nabla F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - \nabla F_t(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \right\|_2 \\ &\leq \sqrt{\kappa} \max_{i,m} \left| \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - \frac{\partial F_t}{\partial \pi_{i,m}}(\pi_{i,1}^{(k)}(t), \dots, \pi_{i,n}^{(k)}(t)) \right| \leq \frac{\sqrt{\kappa n M d(G)}}{K}, \end{aligned}$$

where $\kappa = \sum_{i=1}^n \kappa_i$ and $d(G)$ is the diameter of a graph, i.e., the length of the shortest path between the most distanced nodes. $\square$

Now, we prove Theorem 5.

Lemma 16. When the communication graph $G(\mathcal{V}, \mathcal{E})$ is connected, if we set M as the maximum marginal contribution of each monotone set function f_t , namely, $M \triangleq \max_{S \subseteq \mathcal{V}, e \in \mathcal{V} \setminus S, t \in [T]} (f_t(e|S))$, we can show that,

i): if f_t is a monotone α -weakly submodular function:

$$\begin{aligned} & (1 - e^{-\alpha}) \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha} + \frac{2\sqrt{\kappa}n^2 Md(G)}{\alpha} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\alpha} \frac{T}{\sqrt{L}} \\ & + \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T \left\langle \mathbf{d}_i^{(k)}(t), \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{v}_i^{(k)}(t) \right\rangle; \end{aligned}$$

ii): if f_t is a monotone (γ, β) -weakly submodular:

$$\begin{aligned} & \frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\phi(\gamma, \beta)} + \frac{2\sqrt{\kappa}n^2 Md(G)}{\phi(\gamma, \beta)} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\phi(\gamma, \beta)} \frac{T}{\sqrt{L}} \\ & - \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T \left\langle \mathbf{d}_i^{(k)}(t), \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{v}_i^{(k)}(t) \right\rangle, \end{aligned}$$

where where we slightly abuse the notation $\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}$ to denote κ_i -dimensional indicator vector over $(\mathcal{A}_t^* \cap \mathcal{V}_i)$. Note that $\mathbf{1}_{\mathcal{A}_t^*}$ is a $(\sum_{i=1}^n \kappa_i)$ -dimensional indicator vector over $\mathcal{A}_t^*$ and $\phi(\gamma, \beta) = \beta(1 - \gamma) + \gamma^2$. Note the $\sum_{i=1}^n \sum_{t=1}^T \left\langle \mathbf{d}_i^{(k)}(t), \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} - \mathbf{v}_i^{(k)}(t) \right\rangle$ is the dynamic regret of linear maximization oracle $Q_i^{(k)}$ over the competitive sequence $(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)})$. Therefore, we also rewrite the previous two results as:

i): if f_t is a monotone α -weakly submodular function:

$$\begin{aligned} & (1 - e^{-\alpha}) \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha} + \frac{2\sqrt{\kappa}n^2 Md(G)}{\alpha} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\alpha} \frac{T}{\sqrt{L}} \\ & + \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \mathcal{R}^{Q_i^{(k)}} \left(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)} \right); \end{aligned}$$

ii): if f_t is a monotone (γ, β) -weakly submodular:

$$\begin{aligned} & \frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\phi(\gamma, \beta)} + \frac{2\sqrt{\kappa}n^2 Md(G)}{\phi(\gamma, \beta)} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\phi(\gamma, \beta)} \frac{T}{\sqrt{L}} \\ & - \frac{1}{K} \sum_{k=1}^K \mathcal{R}^{Q_i^{(k)}} \left(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)} \right), \end{aligned}$$

where $\mathcal{R}^{Q_i^{(k)}} \left(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)} \right)$ is the dynamic regret of linear maximization oracle $Q_i^{(k)}$ over the competitive sequence $(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)})$.

Proof.

$$\begin{aligned}
& F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \\
& \geq \left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), (\pi_{1,1}^{(k)}(t) - \pi_{1,1}^{(k-1)}(t), \dots, (\pi_{1,1}^{(k)}(t) - \pi_{n,n}^{(k-1)}(t))) \right\rangle \\
& \quad - \frac{M(\sum_{i=1}^n \kappa_i)}{2} \left\| (\pi_{1,1}^{(k)}(t) - \pi_{1,1}^{(k-1)}(t), \dots, (\pi_{1,1}^{(k)}(t) - \pi_{n,n}^{(k-1)}(t))) \right\|_2^2 \\
& \geq \frac{1}{K} \left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) \right\rangle - \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} \\
& = \frac{1}{K} \left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), \mathbf{1}_{\mathcal{A}_t^*} \right\rangle - \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} \\
& \quad + \frac{1}{K} \left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \\
& = \frac{1}{K} \underbrace{\left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), \mathbf{1}_{\mathcal{A}_t^*} \right\rangle}_{\textcircled{1}} - \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} \\
& \quad + \frac{1}{K} \underbrace{\left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) - \sum_{i \in \mathcal{N}} \left[\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t)) \right]_{\pi_i}, (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle}_{\textcircled{2}} \\
& \quad + \frac{1}{K} \underbrace{\left\langle \sum_{i \in \mathcal{N}} \left[\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t)) \right]_{\pi_i} - (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle}_{\textcircled{3}} \\
& \quad + \frac{1}{K} \left\langle (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle,
\end{aligned}$$

where the first inequality follows the Lemma 5, namely, our proposed policy-based continuous extension F_t is $(M \sum_{i=1}^n \kappa_i)$ -smooth and the symbol $\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}$ is the projection over the policy π_i , namely, $\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i}$ represents a $(\sum_{i=1}^n \kappa_i)$ -dimensional vector that only keeps the first-order derivative at variable $\pi_{i,m}$, $\forall m \in [\kappa_i]$ and set other coordinates to 0, that is to say,

$$\left[\nabla F_t^a(\pi_{i,1}(t), \dots, \pi_{i,1}(t)) \right]_{\pi_i} \triangleq \left(\dots, 0, \dots, \underbrace{\frac{\partial F_t}{\partial \pi_{i,1}}(\mathbf{x}), \dots, \frac{\partial F_t}{\partial \pi_{i,\kappa_i}}(\mathbf{x})}_{\kappa_i}, \dots, 0, \dots \right),$$

where $\mathbf{x} \triangleq (\pi_{i,1}(t), \dots, \pi_{i,1}(t))$.

For $\textcircled{1}$, according to $\textcircled{4}$ of Theorem 1, when f_t is monotone α -weakly DR-submodular,

$$\left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \geq \alpha \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right);$$

As for the setting that f_t is monotone (γ, β) -weakly submodular, we also have

$$\left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)), \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \geq \gamma^2 f_t(\mathcal{A}_t^*) - (\beta(1-\gamma) + \gamma^2) F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)).$$

For ②, according to Lemma 15, we have

$$\begin{aligned}
& \left| \left\langle \nabla F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) - \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i}, (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \right| \\
&= \left| \left\langle \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i}, (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \right| \\
&\leq \left\| \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} \right\|_2 \left\| (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\|_2 \\
&\leq 2n \sum_{i \in \mathcal{N}} \left\| [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} \right\|_2 \\
&\leq 2n \sum_{i \in \mathcal{N}} \left\| \nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t)) - \nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t)) \right\|_2 \leq \frac{2\sqrt{\kappa}n^2Md(G)}{K}.
\end{aligned}$$

As for ③,

$$\begin{aligned}
& \mathbb{E} \left| \left\langle \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \right| \\
&\leq \mathbb{E} \left(\left\| \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)) \right\|_2 \left\| (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\|_2 \right) \\
&\leq 2n \mathbb{E} \left(\left\| \sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i} - (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)) \right\|_2 \right) \\
&\leq 2n \sqrt{\mathbb{E} \left(\frac{1}{L} \left\| (\mathbf{g}_1^{(k)}(t), \dots, \mathbf{g}_n^{(k)}(t)) \right\|_2^2 \right)} \\
&\leq 2n \sqrt{\frac{\kappa M^2}{L}},
\end{aligned}$$

where the third inequality follows Lines 15-20, namely, $(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t))$ is L -batch stochastic gradient for $\sum_{i \in \mathcal{N}} [\nabla F_t(\pi_{i,1}^{(k-1)}(t), \dots, \pi_{i,n}^{(k-1)}(t))]_{\pi_i}$ and the final inequality follows from Remark 13 and $\kappa = \sum_{i=1}^n \kappa_i$.

Therefore, when f_t is monotone α -weakly DR-submodular, we have that

$$\begin{aligned}
& \mathbb{E} \left(F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) \\
&\geq \frac{\alpha}{K} \mathbb{E} \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) - \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} - \frac{2\sqrt{\kappa}n^2Md(G)}{K^2} - \frac{2nM}{K} \sqrt{\frac{\kappa}{L}} \\
&+ \frac{1}{K} \left\langle (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle.
\end{aligned}$$

Then, we can have that

$$\begin{aligned}
& \mathbb{E} \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\
&\leq \left(1 - \frac{\alpha}{K} \right) \mathbb{E} \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) + \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} + \frac{2\sqrt{\kappa}n^2Md(G)}{K^2} + \frac{2nM}{K} \sqrt{\frac{\kappa}{L}} \\
&- \frac{1}{K} \left\langle (\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t)), (\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t)) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle.
\end{aligned}$$

As a result,

$$\begin{aligned}
& \mathbb{E} \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\
& \leq \left(1 - \frac{\alpha}{K} \right)^K \mathbb{E} \left(f_t(\mathcal{A}_t^*) - F_t(\pi_{1,1}^{(0)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) \\
& \quad - \frac{1}{K} \sum_{k=1}^K \left\langle \left(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t) \right), \left(\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t) \right) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \\
& \quad + \sum_{m=1}^K \left(1 - \frac{\alpha}{K} \right)^{K-m} \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} + \frac{2\sqrt{\kappa}n^2Md(G)}{K^2} + \frac{2nM}{K} \sqrt{\frac{\kappa}{L}} \right) \\
& \leq \left(1 - \frac{\alpha}{K} \right)^K \mathbb{E} (f_t(\mathcal{A}_t^*)) + \frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha K} + \frac{2\sqrt{\kappa}n^2Md(G)}{\alpha K} + \frac{2nM}{\alpha} \sqrt{\frac{\kappa}{L}} \\
& \quad - \frac{1}{K} \sum_{k=1}^K \left\langle \left(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t) \right), \left(\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t) \right) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle \\
& \leq e^{-\alpha} \mathbb{E} (f_t(\mathcal{A}_t^*)) + \frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha K} + \frac{2\sqrt{\kappa}n^2Md(G)}{\alpha K} + \frac{2nM}{\alpha} \sqrt{\frac{\kappa}{L}} \\
& \quad - \frac{1}{K} \sum_{k=1}^K \left\langle \left(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t) \right), \left(\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t) \right) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle,
\end{aligned}$$

where the second inequality follows from $\sum_{m=1}^K \left(1 - \frac{\alpha}{K} \right)^{K-m} \leq \frac{K}{\alpha}$, $F_t(\mathbf{0}) = 0$ and the final inequality from $(1 - \frac{\alpha}{K})^K \leq e^{-\alpha}$ when $K \geq 3$.

So we can get the following result,

$$\begin{aligned}
& (1 - e^{-\alpha}) \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\
& \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha} + \frac{2\sqrt{\kappa}n^2Md(G)}{\alpha} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\alpha} \frac{T}{\sqrt{L}} \\
& \quad - \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T \left\langle \mathbf{d}_i^{(k)}(t), \mathbf{v}_i^{(k)}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} \right\rangle
\end{aligned}$$

where we slightly abuse the notation $\mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)}$ to denote κ_i -dimensional indicator vector over $(\mathcal{A}_t^* \cap \mathcal{V}_i)$. Note that $\mathbf{1}_{\mathcal{A}_t^*}$ is a $(\sum_{i=1}^n \kappa_i)$ -dimensional indicator vector over $\mathcal{A}_t^*$.

Similarly, when f_t is monotone (γ, β) -weakly submodular, we also have

$$\begin{aligned}
& \mathbb{E} \left(F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) - F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) \\
& \geq \frac{1}{K} \mathbb{E} \left(\gamma^2 f_t(\mathcal{A}_t^*) - \phi(\gamma, \beta) F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) - \frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} - \frac{2\sqrt{\kappa}n^2Md(G)}{K^2} - \frac{2nM}{K} \sqrt{\frac{\kappa}{L}} \\
& \quad + \frac{1}{K} \left\langle \left(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t) \right), \left(\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t) \right) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle,
\end{aligned}$$

where $\phi(\gamma, \beta) = \beta(1 - \gamma) + \gamma^2$.

Then, we have

$$\begin{aligned}
& \mathbb{E} \left(\gamma^2 f_t(\mathcal{A}_t^*) - \phi(\gamma, \beta) F_t(\pi_{1,1}^{(k)}(t), \dots, \pi_{n,n}^{(k)}(t)) \right) \\
& \leq \left(1 - \frac{\phi(\gamma, \beta)}{K} \right) \mathbb{E} \left(\gamma^2 f_t(\mathcal{A}_t^*) - \phi(\gamma, \beta) F_t(\pi_{1,1}^{(k-1)}(t), \dots, \pi_{n,n}^{(k-1)}(t)) \right) \\
& \quad + \phi(\gamma, \beta) \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2K^2} + \frac{2\sqrt{\kappa}n^2Md(G)}{K^2} + \frac{2nM}{K} \sqrt{\frac{\kappa}{L}} \right) \\
& \quad - \frac{\phi(\gamma, \beta)}{K} \left\langle \left(\mathbf{d}_1^{(k)}(t), \dots, \mathbf{d}_n^{(k)}(t) \right), \left(\mathbf{v}_1^{(k)}(t), \dots, \mathbf{v}_n^{(k)}(t) \right) - \mathbf{1}_{\mathcal{A}_t^*} \right\rangle.
\end{aligned}$$

As a result, we can get the following result,

$$\begin{aligned} & \frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(K)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\phi(\gamma, \beta)} + \frac{2\sqrt{\kappa}n^2Md(G)}{\phi(\gamma, \beta)} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\phi(\gamma, \beta)} \frac{T}{\sqrt{L}} - \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T \left\langle \mathbf{d}_i^{(k)}(t), \mathbf{v}_i^{(k)}(t) - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} \right\rangle, \end{aligned}$$

□

As a result, when $K = \sqrt{T}$ and $L = T$, if Assumption 2 holds, that is, each linear maximization oracle $Q_i^{(k)}$ can achieve the a regret bound of $\mathcal{O}(\sqrt{V_T T})$ where V_T is the variation of any feasible path $(\mathbf{u}_1, \dots, \mathbf{u}_T)$ where $\mathbf{u}_t \in \Delta_{\kappa_i}, \forall t \in [T]$, that is, $V_T \triangleq \sum_{t=2}^T \|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2$ for any path $(\mathbf{u}_1, \dots, \mathbf{u}_T) \in \prod_{t=1}^T \Delta_{\kappa_i}$, then we can have the following results:

i): when f_t is a monotone α -weakly submodular function, Algorithm 2 achieves:

$$\begin{aligned} & (1 - e^{-\alpha}) \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(K)}(t)) \right) \\ & \leq \left(\frac{nM(\sum_{i=1}^n \kappa_i)}{2\alpha} + \frac{2\sqrt{\kappa}n^2Md(G)}{\alpha} \right) \frac{T}{K} + \frac{2nM\sqrt{\kappa}}{\alpha} \frac{T}{\sqrt{L}} + \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \mathcal{R}^{Q_i^{(k)}} \left(\mathbf{1}_{(\mathcal{A}_1^* \cap \mathcal{V}_i)}, \dots, \mathbf{1}_{(\mathcal{A}_T^* \cap \mathcal{V}_i)} \right) \\ & \leq \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \mathcal{O} \left(\sqrt{T \sum_{t=1}^{T-1} \left\| \mathbf{1}_{(\mathcal{A}_{t+1}^* \cap \mathcal{V}_i)} - \mathbf{1}_{(\mathcal{A}_t^* \cap \mathcal{V}_i)} \right\|_2} \right) + \mathcal{O} \left(d(G)\sqrt{T} \right) \\ & \leq \sum_{i=1}^n \mathcal{O} \left(\sqrt{T \sum_{t=1}^{T-1} \left\| \mathbf{1}_{\mathcal{A}_{t+1}^*} - \mathbf{1}_{\mathcal{A}_t^*} \right\|_2} \right) + \mathcal{O} \left(d(G)\sqrt{T} \right) \\ & \leq \mathcal{O} \left(d(G) \sqrt{T \sum_{t=1}^{T-1} \left\| \mathbf{1}_{\mathcal{A}_{t+1}^*} - \mathbf{1}_{\mathcal{A}_t^*} \right\|_2} \right) = \mathcal{O} \left(d(G)\sqrt{T\mathcal{P}_T} \right), \end{aligned}$$

where the third inequality follows from the concavity of $\sqrt{\cdot}$ function and the final equality comes from $\mathcal{P}_T \triangleq \sum_{t=1}^T |\mathcal{A}_{t+1}^* \triangle \mathcal{A}_t^*| \triangleq \sum_{t=1}^{T-1} \|\mathbf{1}_{\mathcal{A}_t^*} - \mathbf{1}_{\mathcal{A}_{t+1}^*}\|_1 \triangleq \sum_{t=1}^{T-1} \|\mathbf{1}_{\mathcal{A}_t^*} - \mathbf{1}_{\mathcal{A}_{t+1}^*}\|_2$.

ii): Similarly, when f_t is a monotone (γ, β) -weakly submodular:

$$\begin{aligned} & \frac{\gamma^2(1 - e^{-\phi(\gamma, \beta)})}{\phi(\gamma, \beta)} \sum_{t=1}^T f_t(\mathcal{A}_t^*) - \sum_{t=1}^T \mathbb{E} \left(F_t(\pi_{1,1}^{(K)}(t), \dots, \pi_{n,n}^{(K)}(t)) \right) \\ & \leq \mathcal{O} \left(d(G) \sqrt{T \sum_{t=1}^{T-1} \left\| \mathbf{1}_{\mathcal{A}_{t+1}^*} - \mathbf{1}_{\mathcal{A}_t^*} \right\|_2} \right) = \mathcal{O} \left(d(G)\sqrt{T\mathcal{P}_T} \right). \end{aligned}$$

I Limitation and Broader Impact

In this work, in order to eliminate the dependence on the unknown DR ratio and submodularity ratio in our proposed MA-SPL algorithm, we introduce a *parameter-free* online algorithm named MA-MPL for the MA-OC problem. However, this new MA-MPL algorithm typically incurs greater communication complexity, as shown in Table 1. Notably, recent studies [74, 124] employed a blocking procedure from [121] to reduce the number of communication in the decentralized online Frank-Wolfe algorithm [133]. Since our MA-SPL algorithm also can be viewed a variant of decentralized online Frank-Wolfe algorithm [133], we believe that this blocking strategy is a promising technique to help reduce the communication complexity of our proposed MA-MPL algorithm. We plan to explore this in future work. Furthermore, this work focuses on theoretically exploring MA-OC problem with (γ, β) -weakly submodular and α -weakly DR-submodular objectives. So we do not foresee any form of negative social impact induced by our work.