Improved Algorithm for Regret Ratio Minimization in Multi-Objective Submodular Maximization

Yanhao Wang[1], **Jiping Zheng**[2,3], **Fanxu Meng**[2]

[1]School of Data Science and Engineering, East China Normal University, Shanghai, China
[2]College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing, China
[3]State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China
yhwang@dase.ecnu.edu.cn, {jzh, mengfanxu}@nuaa.edu.cn

Abstract

Submodular maximization has attracted extensive attention due to its numerous applications in machine learning and artificial intelligence. Many real-world problems require maximizing multiple submodular objective functions at the same time. In such cases, a common approach is to select a representative subset of Pareto optimal solutions with different trade-offs among multiple objectives. To this end, in this paper, we investigate the regret ratio minimization (RRM) problem in multi-objective submodular maximization, which aims to find at most k solutions to best approximate all Pareto optimal solutions w.r.t. any linear combination of objective functions. We propose a novel HS-RRM algorithm by transforming RRM into HITTINGSET problems based on the notions of ϵ-kernel and δ-net, where any α-approximation algorithm for single-objective submodular maximization is used as an oracle. We prove that the maximum regret ratio (MRR) of the output of HS-RRM is bounded by $1-\alpha+O\big((k-d)^{-\frac{2}{d-1}}\big)$, where d is the number of objectives, which improves upon the previous best-known bound of $1-\alpha+O\big((k-d)^{-\frac{1}{d-1}}\big)$ and is nearly asymptotically optimal for any fixed d. Experiments on real-world and synthetic data confirm that HS-RRM achieves lower MRRs than existing algorithms.

Introduction

Submodularity is a property of set functions arising in numerous machine learning and artificial intelligence problems, such as *feature selection* (Krause and Guestrin 2005; Amiridi, Kargas, and Sidiropoulos 2021), *influence maximization* (Kempe, Kleinberg, and Tardos 2003; Becker et al. 2022), *data summarization* (Kumari and Bilmes 2021), and *sensor placement* (Ohsaka and Matsuoka 2021). Given a ground set V, a set function $f : 2^V \to \mathbb{R}$ is said to be *submodular* if $f(X \cup \{v\}) - f(X) \geq f(Y \cup \{v\}) - f(Y)$ for any $X \subset Y \subseteq V$ and $v \in V \setminus Y$, which is well known as the "diminishing returns" property. The *submodular maximization* problem which aims to select a subset of items to maximize f subject to a certain constraint, i.e., $\max_{S\in\mathcal{I}} f(S)$ where $\mathcal{I} \subseteq 2^V$ is the set of all feasible solutions, has been well-studied in combinatorial optimization due to its deep theoretical foundations and broad practical applications. Although many classes of submodular maximization problems are NP-hard (Feige 1998), several polynomial-time approximation algorithms have been proposed for these problems in various settings.

Most existing studies focus on maximizing a single submodular function, while we are often faced with the problems of optimizing multiple possibly conflicting objectives simultaneously. For example, fairness-aware influence maximization problems (Tsang et al. 2019; Fu et al. 2021) aim to spread information fairly among all d demographic groups in a social network, which is modeled by maximizing d submodular functions, each denoting the influence diffusion in one group. As another example, data summarization tasks (Lin and Bilmes 2011; Mirzasoleiman, Badanidiyuru, and Karbasi 2016) are often considered as maximizing two submodular objective functions, which are used to measure *coverage* and *diversity*, respectively. In such cases, a natural problem is finding a set of solutions with different trade-offs among multiple objectives.

A naive approach to finding such "representative" solution sets is to enumerate all *Pareto optimal* solutions (Qian, Yu, and Zhou 2015), i.e., all the solutions attaining an optimal trade-off for at least one combination of multiple objective functions. However, this approach suffers from two limitations. First, computing an optimal solution w.r.t. any given trade-off is infeasible due to the NP-hardness of submodular maximization. Second, the total number of Pareto optimal solutions might be prohibitively large due to the combinatorial nature of subset selection. To tackle these issues, the notion of *regret ratio* has been introduced to measure how well a set of solutions approximates a Pareto optimal solution (Soma and Yoshida 2017; Feng and Qian 2021). Suppose that a trade-off among d objectives is denoted as a nonnegative linear combination of d objective functions $f_1, \ldots, f_d$, i.e., $f_{\boldsymbol{w}}(X) = \sum_{i=1}^{d} w_i f_i(X)$ for a vector $\boldsymbol{w} \in \mathbb{R}_+^d$. The Pareto optimal solution w.r.t. $\boldsymbol{w}$ is the one that maximizes $f_{\boldsymbol{w}}$, i.e., $X^*_{\boldsymbol{w}} = \arg\max_{X\in\mathcal{I}} f_{\boldsymbol{w}}(X)$. Then, for a set $\mathcal{S}$ of solutions, its regret ratio w.r.t. $\boldsymbol{w}$ is defined as $1-\frac{\max_{X\in\mathcal{S}} f_{\boldsymbol{w}}(X)}{f_{\boldsymbol{w}}(X^*_{\boldsymbol{w}})}$. Intuitively, it captures the relative loss of using the best solution in $\mathcal{S}$ to approximate the Pareto optimal one w.r.t. $\boldsymbol{w}$. To approximate all Pareto optimal solutions, we need to find a set of solutions with small regret ratios w.r.t. all possible vectors. This is formalized as the *Regret Ratio Minimization* (RRM) problem in multi-objective submodular maximization (Soma and Yoshida 2017; Feng and Qian 2021), which finds a set $\mathcal{S}$ of at most k solutions to minimize the maximum of the regret ratio for any $\boldsymbol{w} \in \mathbb{R}^d_+$.

Our Results

Although there have been several algorithms proposed for the RRM problem (Soma and Yoshida 2017; Feng and Qian 2021), they still have three drawbacks: (i) they do not work for $k < d$; (ii) they do not provide any theoretical guarantee on the gap between the maximum regret ratio (MRR) of their output and the optimal MRR; (iii) although an upper bound on MRR in the worst case is provided, the bound is not tight and can be further improved. In this paper, we address all the above problems by proposing a novel algorithm called HS-RRM for the RRM problem. The basic idea of HS-RRM is to sample a number of weight vectors based on the geometric notions of *ϵ-kernel* and *δ-net* (Agarwal, Har-Peled, and Varadarajan 2004), to use any given α-approximation algorithm for single-objective submodular maximization to compute a solution w.r.t. each vector, to construct a HITTINGSET instance by checking whether a solution achieves a given regret ratio w.r.t. each vector, and to solve the HITTINGSET instance to find a set of solutions for RRM that "covers" all vectors within the regret ratio. We show that HS-RRM works for any $k \geq 1$ and provides a solution with a maximum regret ratio of $1 - \alpha^2 + \alpha^2 mrr^*_k + O(\lambda)$ when HITTINGSET is solved optimally, where mrr^*_k is the minimum possible MRR by any size-k set of Pareto optimal solutions and $\lambda > 0$ is an error parameter that can be arbitrarily small. In addition, the MRR of the output of HS-RRM has an upper bound of $1-\alpha+O\big((k-d)^{-\frac{2}{d-1}}\big)$, which not only improves upon the best-known bound of $1-\alpha+O\big((k-d)^{-\frac{1}{d-1}}\big)$ (Feng and Qian 2021) but also nearly matches the asymptotically optimal bound of $O\big(k^{-\frac{2}{d-1}}\big)$ for any fixed d. Finally, we experimentally compare HS-RRM with existing algorithms for RRM in (Soma and Yoshida 2017; Feng and Qian 2021) on multi-objective *weighted maximum coverage* and *data summarization* problems. The results confirm that HS-RRM achieves significantly lower maximum regret ratios than existing algorithms in practice.

Additional Related Work

Besides (Soma and Yoshida 2017; Feng and Qian 2021), there are many other studies relevant to this work. The multi-objective submodular maximization (Krause et al. 2008; Udwani 2018; Anari et al. 2019; Fu et al. 2021) problem aims at finding a solution under a specific constraint to maximize the minimum among d submodular functions. It is a special case of RRM when $k = 1$ and the weight vectors are limited to d basis vectors where only one entry is equal to 1 and all others are equal to 0. The submodular maximization under submodular cover (SMSC) problem (Ohsaka and Matsuoka 2021; Gershtein, Milo, and Youngmann 2021) finds a solution to maximize one submodular function subject to that the value of the other submodular function is at least the τ-fraction of its optimum. It can be seen as a special case of RRM when $k = 1$ and $d = 2$. But another difference between SMSC and RRM is that SMSC only considers a given trade-off between both objectives specified by $\tau \in (0, 1)$ instead of all possible trade-offs considered in RRM. The regret ratio minimization problem for linear ranking functions (Nanongkai et al. 2010; Peng and Wong 2014; Cao et al. 2017; Agarwal et al. 2017; Asudeh et al. 2017; Xie et al. 2018; Kumar and Sintos 2018; Storandt and Funke 2019; Wang et al. 2021a,b; Zheng et al. 2022) has been extensively studied in the literature. The algorithms for the linear setting cannot be directly applied to the submodular setting for two reasons. First, they require enumerating all Pareto optimal points for linear functions through the *skyline* operator (Börzsönyi, Kossmann, and Stocker 2001). But, as has been shown, such an enumeration is prohibitively expensive for submodular functions. Second, computing the optimal solution for any linear function is trivial by performing a linear scan over all items, whereas it is NP-hard for submodular functions.

Preliminaries

For a positive integer n, let $[n]$ denote the set of integers $\{1, 2, \ldots, n\}$. Let V denote a finite ground set of size n, i.e., $V = [n]$. We study the set function $f : 2^V \to \mathbb{R}$ defined on the subsets of V. We assume that f is *nonnegative*, i.e., $f(X) \geq 0$ for any $X \subseteq V$, and *submodular*, i.e., $f(X \cup \{v\}) - f(X) \geq f(Y \cup \{v\}) - f(Y)$ for any $X \subset Y \subseteq V$ and $v \in V \setminus Y$. In addition, f is said to be *monotone* if $f(X) \leq f(Y)$ for any $X \subset Y \subseteq V$. We consider the problem of maximizing a submodular function f subject to a specific constraint. Let $\mathcal{I} \subseteq 2^V$ be the set of feasible solutions defined by a given constraint. For instance, the cardinality constraint $r \in \mathbb{Z}^+$ gives a feasible solution set $\mathcal{I} = \{X \subseteq V : |X| \leq r\}$. Other typical choices of constraints for submodular maximization include the *matroid* (Călinescu et al. 2011), *knapsack* (Ene and Nguyen 2019; Tang et al. 2021), and *p-system constraints* (Fisher, Nemhauser, and Wolsey 1978). Formally, the submodular maximization problem asks for a set $X \in \mathcal{I}$ that maximizes $f(X)$ among all sets in $\mathcal{I}$, i.e., $X^* = \arg\max_{X \in \mathcal{I}} f(X)$ and $\mathtt{OPT} = \max_{X \in \mathcal{I}} f(X)$. Note that the results in this work hold for (nonnegative) monotone and non-monotone submodular maximization problems under different constraints.

We define the regret ratio $rr_{f,\mathcal{I}}(X)$ of any $X \in \mathcal{I}$ w.r.t. f as the relative loss between $f(X)$ and the optimum $\mathtt{OPT}$ for maximizing f over $\mathcal{I}$, i.e., $rr_{f,\mathcal{I}}(X) = 1 - \frac{f(X)}{\mathtt{OPT}}$. Intuitively, $rr_{f,\mathcal{I}}(X) \in [0, 1]$ and $rr_{f,\mathcal{I}}(X) = 0$ if and only if $X = X^*$. The notion of regret ratio can be naturally generalized to a set $\mathcal{S} \subseteq \mathcal{I}$ of solutions as $rr_{f,\mathcal{I}}(\mathcal{S}) = 1 - \frac{\max_{X \in \mathcal{S}} f(X)}{\mathtt{OPT}}$. Similarly, $rr_{f,\mathcal{I}}(\mathcal{S}) = 0$ if and only if $X^* \in \mathcal{S}$.

We focus on multi-objective submodular maximization: given d submodular functions $f_1, f_2, \ldots, f_d : 2^V \to \mathbb{R}^+$ on the same feasible set $\mathcal{I}$, we aim at finding the solutions to maximize all d functions simultaneously. Since the objectives might contradict each other, instead of providing a single solution for optimizing all objectives, we turn to find a set of Pareto optimal solutions for different trade-offs among the d objectives. Specifically, each particular trade-off is represented by a nonnegative linear combination of the d functions, i.e., to maximize $f_{\boldsymbol{w}}(X) = \sum_{i=1}^{d} w_i f_i(X)$ for a given weight vector $\boldsymbol{w} = (w_1, \ldots, w_d) \in \mathbb{R}_+^d$. Note that $f_{\boldsymbol{w}}$ is still a nonnegative submodular function. Following the previous work (Qian, Yu, and Zhou 2015; Soma and Yoshida 2017; Feng and Qian 2021), we say X is Pareto optimal if X is optimal w.r.t. $f_{\boldsymbol{w}}(X)$ for any $\boldsymbol{w} \in \mathbb{R}_+^d$. We define the maximum regret ratio (MRR) $mrr_{f_1,\ldots,f_d,\mathcal{I}}(\mathcal{S})$ of a set $\mathcal{S} \subseteq \mathcal{I}$ of feasible solutions w.r.t. $f_1, \ldots, f_d$ as follows:

$$mrr_{f_1,\ldots,f_d,\mathcal{I}}(\mathcal{S}) = \max_{\boldsymbol{w}\in\mathbb{R}_+^d} rr_{f_{\boldsymbol{w}},\mathcal{I}}(\mathcal{S}). \tag{1}$$

Intuitively, the MRR measures to what degree $\mathcal{S}$ can approximate all Pareto optimal solutions in the worst case. The Regret Ratio Minimization (RRM) problem in multi-objective submodular maximization (Soma and Yoshida 2017; Feng and Qian 2021) aims to find a set $\mathcal{S}$ of at most k feasible solutions such that the MRR of $\mathcal{S}$ is minimized, which is formally defined as follows.

Definition 1. *[Regret Ratio Minimization] Given d nonnegative submodular functions $f_1, f_2, \ldots, f_d : 2^V \to \mathbb{R}_+$ and the set $\mathcal{I}$ of feasible sets, find a set $\mathcal{S} \subseteq \mathcal{I}$ of size at most $k \in \mathbb{Z}^+$ to minimize $mrr_{f_1,\ldots,f_d,\mathcal{I}}(\mathcal{S})$, i.e.,*

$$\mathcal{S}^* = \underset{\mathcal{S}\subseteq\mathcal{I}\,:\,|\mathcal{S}|\leq k}{\arg\min}\; mrr_{f_1,\ldots,f_d,\mathcal{I}}(\mathcal{S}).$$

For ease of representation, when the context is clear, the subscript $f_1, \ldots, f_d, \mathcal{I}$ will be omitted from the notation mrr. We denote the optimal MRR for RRM as mrr_k^*, i.e., $mrr_k^* = \min_{\mathcal{S}\subseteq\mathcal{I}\,:\,|\mathcal{S}|\leq k} mrr(\mathcal{S})$. Since the MRRs are scale-invariant (Nanongkai et al. 2010; Feng and Qian 2021), i.e., $rr_{f_{\boldsymbol{w}},\mathcal{I}}(\mathcal{S}) = rr_{f_{t\boldsymbol{w}},\mathcal{I}}(\mathcal{S})$ for any $t > 0$, the RRM problem in Definition 1 is equivalent to that when the weight vectors in Eq. (1) are restricted to unit vectors, i.e., $\mathbb{S}_+^{d-1} = \{\boldsymbol{w} \in \mathbb{R}_+^d : \|\boldsymbol{w}\| = 1\}$. Note that we consider RRM for any positive integer $k \in \mathbb{Z}^+$ instead of only for $k \geq d$ in (Soma and Yoshida 2017; Feng and Qian 2021).

The HS-RRM Algorithm

In this section, we first propose a novel algorithm for Regret Ratio Minimization (RRM) called HS-RRM. Then, we provide a thorough theoretical analysis for HS-RRM.

Algorithm Description

The general idea of our HS-RRM algorithm is to transform the RRM problem into several instances of the HITTINGSET problem, which will subsequently be solved by existing algorithms. Specifically, a set system $\Sigma = (\mathcal{U}, \mathcal{C})$ consists of a universe $\mathcal{U}$ of elements and a collection $\mathcal{C}$ of subsets of $\mathcal{U}$. A set $H \subseteq \mathcal{U}$ is called a hitting set of Σ if $H \cap C \neq \emptyset$ for any $C \in \mathcal{C}$. The HITTINGSET problem asks for a hitting set of the minimum size, which is NP-hard (Karp 1972) in general. The classic greedy algorithm (Feige 1998) is an $(1 + \ln|\mathcal{U}|)$-approximation algorithm for HITTINGSET and no polynomial-time algorithm can achieve an approximation factor of $(1 - o(1)) \ln|\mathcal{U}|$ unless $P = NP$. Nevertheless, in some special cases, HITTINGSET can be solved optimally in polynomial time.

Like the case for linear functions (Agarwal et al. 2017; Kumar and Sintos 2018; Wang et al. 2021a), the transformation from RRM to HITTINGSET is based on two important notions in computational geometry, namely δ-net and ϵ-kernel (Agarwal, Har-Peled, and Varadarajan 2004; Agarwal et al. 2017). For a parameter $\delta > 0$, a set $\mathcal{N} \subset \mathbb{S}_+^{d-1}$ is called a δ-net of $\mathbb{S}_+^{d-1}$ if there is $\boldsymbol{w} \in \mathcal{N}$ such that $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq \cos(\delta)$ for any $\boldsymbol{v} \in \mathbb{S}_+^{d-1}$. A δ-net of size $O(\delta^{-\frac{1}{d-1}})$ can be computed by drawing a uniform grid on $\mathbb{S}_+^{d-1}$ (Agarwal, Har-Peled, and Varadarajan 2004). Intuitively, a (discrete) δ-net $\mathcal{N}$ can approximate (continuous) $\mathbb{S}_+^{d-1}$ such that RRM is reduced to a finite number of vectors within bounded errors. Moreover, given a set $P \subset \mathbb{R}_+^d$ of vectors, an ϵ-kernel (Agarwal, Har-Peled, and Varadarajan 2004) is a subset $Q \subseteq P$ such that $\max_{\boldsymbol{v}\in Q}\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq (1 - \epsilon) \max_{\boldsymbol{v}\in P}\langle \boldsymbol{v}, \boldsymbol{w} \rangle$ for any $\boldsymbol{w} \in \mathbb{R}_+^d$. An ϵ-kernel of size $O(\epsilon^{-\frac{2}{d-1}})$ can be computed by first generating a δ-net $\mathcal{N}$, where $\delta = O(\sqrt{\epsilon})$, then scaling P and $\mathcal{N}$ properly, and obtaining the nearest neighbors of all $\boldsymbol{w} \in \mathcal{N}$ from P. The goal of ϵ-kernels is to approximate the maximum (i.e., the function values in RRM) w.r.t. all possible weight vectors within bounded errors, which coincides with the notion of MRR as shown in (Agarwal et al. 2017; Cao et al. 2017; Kumar and Sintos 2018; Xie et al. 2018).

Next, we present our HS-RRM algorithm in Algorithm 1. First, an α-approximate solution X_i is computed to maximize each f_i over $\mathcal{I}$ using any α-approximate oracle $\mathcal{A}$ for single-objective submodular maximization where $\alpha \in (0, 1]$. For example, when $\mathcal{I}$ is defined by a cardinality constraint, there are oracles with approximation factors $\alpha = 1 - \frac{1}{e}$ and $\alpha = \frac{1}{e}$ for maximizing monotone and non-monotone submodular objective functions (Krause and Golovin 2014). A function $f'_{\boldsymbol{w}}(X) = \sum_{i=1}^{d} \frac{w_i}{f_i(X_i)} f_i(X)$ w.r.t. each $\boldsymbol{w}$ is defined by rescaling the value of each f_i by $f_i(X_i)$. We will soon show that the MRR of any $\mathcal{S} \subseteq \mathcal{I}$ is not changed by rescaling in our theoretical analysis. We then draw an $\frac{\alpha\lambda}{2d}$-net $\mathcal{N}$, where $\lambda \in (0, 1)$ is an error parameter. The d basis vectors $\mathcal{B} = \{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d\}$, where $\boldsymbol{b}_i$ is the vector with $b_i = 1$ and $b_j = 0$ for any $j \neq i$, are also added to $\mathcal{N}$. For each $\boldsymbol{w} \in \mathcal{N}$, we call $\mathcal{A}$ to get an α-approximate solution $X'_{\boldsymbol{w}}$ to maximize $f'_{\boldsymbol{w}}$ over $\mathcal{I}$. Then, we compute a "base" set $\mathcal{S}_1$ using ϵ-kernels: we draw k uniform vectors from the nonnegative orthant of the sphere $\mathbb{S}_R^{d-1}$ of radius $R = \frac{1+\sqrt{d}}{\alpha}$ as a δ'-net $\mathcal{N}'$ and obtain $\mathcal{S}_1$ by adding the nearest neighbor of each vector $\boldsymbol{v} \in \mathcal{N}'$ from $\mathcal{S}_{\mathcal{N}} = \{X'_{\boldsymbol{w}} : \boldsymbol{w} \in \mathcal{N}\}$.

Then, we perform an iterative procedure for computation via the transformation from RRM to HITTINGSET. The upper and lower bounds τ_h, τ_l of the threshold τ are initialized to 1 and 0, respectively. Moreover, $\mathcal{S}_2$ keeps the set with the smallest MRR found in the iterative procedure so far and is initialized as $\emptyset$. At each iteration, we first compute the threshold τ as $(\tau_h + \tau_l)/2$. We then build a set system Σ based on the α-approximate solutions for the vectors in $\mathcal{N}$. Specifically, the universe $\mathcal{U}$ is equal to $\mathcal{N}$ and the set collection $\mathcal{C}$ consists of the sets for all vectors in $\mathcal{N}$ with respect to the threshold τ. Each set $C_\tau(\boldsymbol{v})$ for $\boldsymbol{v} \in \mathcal{N}$ contains each $\boldsymbol{w} \in \mathcal{N}$ such that $f'_{\boldsymbol{v}}(X'_{\boldsymbol{w}})$ is at least τ-fraction of $f'_{\boldsymbol{v}}(X'_{\boldsymbol{v}})$. After that, any algorithm for HITTINGSET can be used to

Algorithm 1: HS-RRM

Input: Submodular set functions $f_1, \ldots, f_d : 2^V \to \mathbb{R}^+$; feasible solution set $\mathcal{I} \subseteq 2^V$; integer $k \in \mathbb{Z}^+$

Parameter: Error term $\lambda \in (0, 1)$; α-approximate oracle $\mathcal{A}$ for maximizing any submodular function f over $\mathcal{I}$

Output: Set $\mathcal{S} \subseteq \mathcal{I}$ s.t. $|\mathcal{S}| \leq k$

for $i = 1, \ldots, d$ **do**
  Run $\mathcal{A}$ to get an α-approximate $X_i \in \mathcal{I}$ w.r.t. f_i.
end for
Define a function $f'_{\boldsymbol{w}}(X) = \sum_{i=1}^{d} \frac{w_i}{f_i(X_i)} f_i(X)$.
Compute an $\frac{\alpha\lambda}{2d}$-net $\mathcal{N}$ of $\mathbb{S}_+^{d-1}$.
Add the d basis vectors $\mathcal{B} = \{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d\}$ to $\mathcal{N}$.
for all $\boldsymbol{w} \in \mathcal{N}$ **do**
  Run $\mathcal{A}$ to get an α-approximate $X'_{\boldsymbol{w}} \in \mathcal{I}$ w.r.t. $f'_{\boldsymbol{w}}$.
end for
Draw k uniform vectors from the nonnegative orthant of $\mathbb{S}_R^{d-1}$ of radius $R = \frac{1+\sqrt{d}}{\alpha}$ as $\mathcal{N}'$ and initialize $\mathcal{S}_1 \leftarrow \emptyset$.
for all $\boldsymbol{v} \in \mathcal{N}'$ **do**
  $X^* \leftarrow \arg\min_{X \in \{X'_{\boldsymbol{w}} : \boldsymbol{w} \in \mathcal{N}\}} \sum_{i=1}^{d} (v_i - f'_i(X))^2$.
  $\mathcal{S}_1 \leftarrow \mathcal{S}_1 \cup \{X^*\}$.
end for
Initialize $\tau_h \leftarrow 1$, $\tau_l \leftarrow 0$, and $\mathcal{S}_2 \leftarrow \emptyset$.
while $\tau_h - \tau_l \geq \lambda$ **do**
  $\tau \leftarrow (\tau_l + \tau_h)/2$
  for all $\boldsymbol{v} \in \mathcal{N}$ **do**
    $C_\tau(\boldsymbol{v}) \leftarrow \{\boldsymbol{w} \in \mathcal{N} : f'_{\boldsymbol{v}}(X'_{\boldsymbol{w}}) \geq \tau \cdot f'_{\boldsymbol{v}}(X'_{\boldsymbol{v}})\}$
  end for
  $\mathcal{C}_\tau \leftarrow \{C_\tau(\boldsymbol{v}) : \boldsymbol{v} \in \mathcal{N}\}$
  $H \leftarrow$ HittingSet$(\mathcal{N}, \mathcal{C}_\tau)$, $\mathcal{S} \leftarrow \{X'_{\boldsymbol{w}} : \boldsymbol{w} \in H\}$
  if $|\mathcal{S}| \leq k$ **then**
    $\mathcal{S}_2 \leftarrow \mathcal{S}$ if $mrr(\mathcal{S}_2) < mrr(\mathcal{S})$ and $\tau_l \leftarrow \tau$.
  else
    $\tau_h \leftarrow \tau$
  end if
end while
return $\arg\min_{\mathcal{S} \in \{\mathcal{S}_1, \mathcal{S}_2\}} mrr(\mathcal{S})$

find an (exact or approximate) hitting set H on $\Sigma = (\mathcal{N}, \mathcal{C}_\tau)$ and a set $\mathcal{S} = \{X'_{\boldsymbol{w}} : \boldsymbol{w} \in H\}$ for RRM is found based on H. If $|\mathcal{S}| \leq k$, i.e., $\mathcal{S}$ is valid for RRM, we will compare $\mathcal{S}$ with the current $\mathcal{S}_2$ and replace $\mathcal{S}_2$ with $\mathcal{S}$ if the maximum regret ratio of $\mathcal{S}$ is smaller. Also, we will set the lower bound τ_l to the current threshold τ to obtain better results at later iterations. Otherwise, i.e., $|\mathcal{S}| > k$, we will set the upper bound τ_h to the current threshold τ so that a valid result can be found. The above iterative procedure terminates when the difference between the upper and lower bounds is below λ. Finally, the better set with a smaller MRR between $\mathcal{S}_1$ and $\mathcal{S}_2$ is returned as the output for RRM.

For the algorithm to solve HittingSet in Algorithm 1, we have different choices in different settings. Firstly, the greedy algorithm that iteratively adds the element covering the most number of uncovered sets until all sets are fully covered is general to solve any HittingSet instance with an approximation factor of $1+\ln|\mathcal{U}|$. Secondly, when a HittingSet instance is very small, its optimal solution can be computed based on the Integer Programming (IP) formulation despite NP-hardness. Thirdly, for a special case of RRM with $d = 2$, the weight vectors can be sorted by their angles from 0 (w.r.t. $\boldsymbol{w} = (1, 0)$) to $\frac{\pi}{2}$ (w.r.t. $\boldsymbol{w} = (0, 1)$) and the memberships between elements and sets can be expressed as sub-intervals of $[0, \frac{\pi}{2}]$ in 1D. Such HittingSet instances are equivalent to the *interval cover* problem (Cormen et al. 2022), for which the greedy algorithm can provide an optimal solution in polynomial time.

Theoretical Analysis

Next, we will analyze our proposed HS-RRM algorithm theoretically. First, we show that the maximum regret ratio of any $\mathcal{S} \subseteq \mathcal{I}$ defined on $f_{\boldsymbol{w}}$ and $f'_{\boldsymbol{w}}$ is identical in Lemma 1.

Lemma 1. *Suppose that* $mrr(\mathcal{S}) = \max_{\boldsymbol{w} \in \mathbb{R}_+^d} rr_{f_{\boldsymbol{w}}, \mathcal{I}}(\mathcal{S})$ *and* $mrr'(\mathcal{S}) = \max_{\boldsymbol{w} \in \mathbb{R}_+^d} rr_{f'_{\boldsymbol{w}}, \mathcal{I}}(\mathcal{S})$*. For any* $\mathcal{S} \subseteq \mathcal{I}$*, it holds that* $mrr(\mathcal{S}) = mrr'(\mathcal{S})$*.*

Proof. For any vector $\boldsymbol{w} = (w_1, \ldots, w_d) \in \mathbb{R}_+^d$, we can find a vector $\boldsymbol{w}' = (w_1 f_1(X_1), \ldots, w_d f_d(X_d)) \in \mathbb{R}_+^d$ such that $rr_{\boldsymbol{w}, \mathcal{I}}(\mathcal{S})$ on $f_{\boldsymbol{w}}$ is equal to $rr'_{\boldsymbol{w}', \mathcal{I}}(\mathcal{S})$ on $f'_{\boldsymbol{w}'}$ and vice versa, because $f_i(X_i) > 0$ for any $i \in [d]$. In this way, we can establish a one-to-one mapping between $\boldsymbol{w}$ and $\boldsymbol{w}'$ via an affine transformation $T : \mathbb{R}_+^d \to \mathbb{R}_+^d$, where $\boldsymbol{w}' = T(\boldsymbol{w}) = \boldsymbol{w} \cdot \boldsymbol{f}$ and $\boldsymbol{f} = (f_1(X_1), \ldots, f_d(X_d))$. Therefore, for the vector $w^* \in \mathbb{R}_+^d$ with the maximum $rr_{f_{\boldsymbol{w}}, \mathcal{I}}(\mathcal{S})$, $T(w^*) \in \mathbb{R}_+^d$ also achieves the same (maximum) $rr_{f'_{\boldsymbol{w}}, \mathcal{I}}(\mathcal{S})$ and naturally $mrr(\mathcal{S}) = mrr'(\mathcal{S})$. $\square$

Given the result of Lemma 1, all our subsequent analyses will be based on f' instead of f. Next, we provide an analysis of the error bound for approximating the MRR with a δ-net $\mathcal{N}$ in Lemma 2.

Lemma 2. *Suppose that an* α*-approximate oracle* $\mathcal{A}$ *is used in HS-RRM. For a set* $\mathcal{N} \cup \mathcal{B} \subset \mathbb{S}_+^{d-1}$ *and a set* $\mathcal{S} \subseteq \mathcal{I}$ *of feasible solutions, if* $\max_{X \in \mathcal{S}} f'_{\boldsymbol{w}}(X) \geq \tau \cdot f'_{\boldsymbol{w}}(X'_{\boldsymbol{w}})$ *for any vector* $\boldsymbol{w} \in \mathcal{N}$*, then* $mrr(\mathcal{S}) \leq 1 - \alpha\tau + \frac{2\delta d}{\alpha}$*.*

Proof. For any $X \in \mathcal{I}$, its values on the normalized functions $f'_i(X) = \frac{f_i(X)}{f_i(X_i)}$ are denoted as a vector $\boldsymbol{f}'(X) = (f'_1(X), \ldots, f'_d(X))$. Thus, $f'_{\boldsymbol{w}}(X) = \langle \boldsymbol{w}, \boldsymbol{f}'(X) \rangle$. Then, due to the Cauchy-Schwarz inequality,

$$\begin{aligned} |f'_{\boldsymbol{w}}(X) - f'_{\boldsymbol{v}}(X)| &= |\langle \boldsymbol{w}, \boldsymbol{f}'(X) \rangle - \langle \boldsymbol{v}, \boldsymbol{f}'(X) \rangle| \\ &= |\langle \boldsymbol{w} - \boldsymbol{v}, \boldsymbol{f}'(X) \rangle| \leq \|\boldsymbol{w} - \boldsymbol{v}\| \cdot \|\boldsymbol{f}'(X)\|. \end{aligned}$$

Since $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq \cos(\delta)$, we have $\|\boldsymbol{w} - \boldsymbol{v}\| = \sqrt{2 - 2\langle \boldsymbol{v}, \boldsymbol{w} \rangle} \leq \sqrt{2 - 2\cos(\delta)} \leq \delta$ based on the properties of trigonometric functions. In addition, we have $\|\boldsymbol{f}'(X)\| \leq \frac{\sqrt{d}}{\alpha}$, as $f'_i(X) \leq \frac{1}{\alpha}$. Thus, it holds that

$$|f'_{\boldsymbol{w}}(X) - f'_{\boldsymbol{v}}(X)| \leq \frac{\delta\sqrt{d}}{\alpha}. \quad (2)$$

Let $X^*_{\boldsymbol{v}} = \arg\max_{X \in \mathcal{I}} f'_{\boldsymbol{v}}(X)$ for any $\boldsymbol{v} \in \mathbb{S}_+^{d-1}$ and $X^*_{\boldsymbol{w}} = \arg\max_{X \in \mathcal{I}} f'_{\boldsymbol{w}}(X)$ for $\boldsymbol{w} \in \mathcal{N}$ such that $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq \cos(\delta)$. We first consider the case of basis vectors $\boldsymbol{b}_i \in \mathcal{B}$. Since there exists $X \in \mathcal{S}$ such that $f'_{\boldsymbol{b}_i}(X) \geq \tau f'_{\boldsymbol{b}_i}(X_i) \geq \alpha\tau f'_{\boldsymbol{b}_i}(X^*_i)$, the regret ratio of $\mathcal{S}$ w.r.t. $\boldsymbol{b}_i$ is at most $1 - \alpha\tau$. For the remaining vectors $\boldsymbol{v} \in \mathbb{S}_+^{d-1}$, we consider two cases separately. In the first case, we assume that $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) \leq \frac{1}{\sqrt{d}}$. Since $\|\boldsymbol{v}\| = 1$, there exists some $i \in [d]$ with $v_i \geq \frac{1}{\sqrt{d}}$. Let us consider $X \in \mathcal{S}$ with $f'_{\boldsymbol{b}_i}(X) \geq \alpha\tau f'_{\boldsymbol{b}_i}(X^*_i)$ for such v_i. We have $f'_{\boldsymbol{v}}(X) \geq \frac{\alpha\tau f'_{\boldsymbol{b}_i}(X^*_i)}{\sqrt{d}} \geq \frac{\alpha\tau}{\sqrt{d}}$. It guarantees that the regret ratio of $\mathcal{S}$ w.r.t. $\boldsymbol{v}$ bounded by X is also at most $1-\alpha\tau$, as $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) \leq \frac{1}{\sqrt{d}}$. In the second case, i.e., $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) > \frac{1}{\sqrt{d}}$, there is $X \in \mathcal{S}$ such that

$$\begin{aligned}
f'_{\boldsymbol{v}}(X) &\geq f'_{\boldsymbol{w}}(X) - \frac{\delta\sqrt{d}}{\alpha} \geq \alpha\tau f'_{\boldsymbol{w}}(X^*_{\boldsymbol{w}}) - \frac{\delta\sqrt{d}}{\alpha} \\
&\geq \alpha\tau f'_{\boldsymbol{w}}(X^*_{\boldsymbol{v}}) - \frac{\delta\sqrt{d}}{\alpha} \geq \alpha\tau(f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) - \frac{\delta\sqrt{d}}{\alpha}) - \frac{\delta\sqrt{d}}{\alpha} \\
&\geq \left(\alpha\tau - \frac{2\delta d}{\alpha}\right) \cdot f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}})
\end{aligned}$$

where the first and fourth inequalities are obtained from Eq. (2), the second inequality holds from the condition that $f'_{\boldsymbol{w}}(X) \geq \tau f'_{\boldsymbol{w}}(X'_{\boldsymbol{w}})$ and $X'_{\boldsymbol{w}}$ is α-approximate, the third inequality comes from the optimality of $X^*_{\boldsymbol{w}}$, and the fifth inequality is because $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) > \frac{1}{\sqrt{d}}$. By combining all above results, we have $mrr(\mathcal{S}) \leq 1 - \alpha\tau + \frac{2\delta d}{\alpha}$. □

Next, we give the approximation factor of the output $\mathcal{S}$ of HS-RRM based on Lemma 2 in Theorem 1.

Theorem 1. *The maximum regret ratio $mrr(\mathcal{S})$ of the output $\mathcal{S}$ of HS-RRM is at most $1-\alpha^2+\alpha^2 mrr^*_k+O(\lambda)$, where mrr^*_k is the optimal MRR of any size-k set of Pareto optimal solutions, when an exact solver is used for* HITTINGSET *or $1-\alpha^2+\alpha^2 mrr^*_{k'}+O(\lambda)$, where $k' = O(\frac{k}{\log|\mathcal{N}|})$, when an $O(\log|\mathcal{N}|)$-approximate solver is used for* HITTINGSET.

Proof. Since $\delta = \frac{\alpha\lambda}{2d}$ in Algorithm 1, we have $mrr(\mathcal{S}_2) \leq 1 - \alpha\tau'_l + \lambda$ from Lemma 2, where τ'_l is the value of τ_l after the iterative procedure. Let τ'_h be the value of τ_h after the iterative procedure. Obviously, $\tau'_h \leq \tau'_l + \lambda$. When an exact solver is used for HITTINGSET, we can say that the set system constructed w.r.t. τ'_h does not have any hitting set of size at most k. Or equivalently, for any $\mathcal{S} \subseteq \mathcal{S}_{\mathcal{N}}$ with $|\mathcal{S}| \leq k$, where $\mathcal{S}_{\mathcal{N}} = \{X'_{\boldsymbol{w}} : \boldsymbol{w} \in \mathcal{N}\}$, there exists $\boldsymbol{w} \in \mathcal{N}$ such that $\max_{X\in\mathcal{S}} f'_{\boldsymbol{w}}(X) < \tau'_h \cdot f'_{\boldsymbol{w}}(X'_{\boldsymbol{w}}) \leq \tau'_h \cdot f'_{\boldsymbol{w}}(X^*_{\boldsymbol{w}})$. For any $\boldsymbol{v} \in \mathbb{S}^{d-1}_+$, there exists $\boldsymbol{w} \in \mathcal{N}$ such that $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) \leq f'_{\boldsymbol{w}}(X^*_{\boldsymbol{w}}) + \frac{\delta\sqrt{d}}{\alpha} \leq \frac{f'_{\boldsymbol{w}}(X'_{\boldsymbol{w}})+\delta\sqrt{d}}{\alpha}$ according to Eq. (2) and the fact that $X'_{\boldsymbol{w}}$ is α-approximate. By mapping each $X^*_{\boldsymbol{v}}$ to the corresponding $X'_{\boldsymbol{w}} \in \mathcal{S}_{\mathcal{N}}$ with $f'_{\boldsymbol{w}}(X'_{\boldsymbol{w}}) \geq \alpha f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) - \delta\sqrt{d}$, we find that any size-$k$ set $\mathcal{S}^*$ of Pareto optimal solutions satisfies that $\max_{X\in\mathcal{S}^*} f'_{\boldsymbol{w}}(X) \leq \frac{\tau'_h \cdot f'_{\boldsymbol{w}}(X^*_{\boldsymbol{w}})+\delta\sqrt{d}}{\alpha}$ for some $\boldsymbol{w} \in \mathcal{N}$, which implies that $mrr^*_k \geq 1 - \frac{\tau'_h}{\alpha} - \lambda$. Thus,

$$\begin{aligned}
mrr(\mathcal{S}_2) - mrr^*_k &< \frac{\tau'_h}{\alpha} - \alpha\tau'_l + 2\lambda \\
&\leq \frac{\tau'_l + \lambda - \alpha^2\tau'_l}{\alpha} + 2\lambda = \frac{1-\alpha^2}{\alpha} \cdot \tau'_l + (\frac{\lambda}{\alpha} + 2\lambda) \\
&\leq \frac{1-\alpha^2}{\alpha} \cdot \left(\alpha(1 - mrr^*_k - \lambda) - \lambda\right) + (\frac{\lambda}{\alpha} + 2\lambda) \\
&\leq (1-\alpha^2)(1 - mrr^*_k) + \lambda'
\end{aligned}$$

where $\lambda' = O(\lambda)$ can be arbitrarily small when the value of λ is close to 0. Therefore, we conclude that $mrr(\mathcal{S}_2) \leq 1-\alpha^2+\alpha^2 mrr^*_k+O(\lambda)$. When an $O(\log|\mathcal{N}|)$-approximate solver is used for HITTINGSET, the difference in the analysis is that it only guarantees that $mrr^*_{k'} \geq 1 - \frac{\tau'_h}{\alpha} - \lambda$ where $k' = O(\frac{k}{\log|\mathcal{N}|})$ (Agarwal et al. 2017), while all remaining parts remain unchanged. The conclusion thus becomes $mrr(\mathcal{S}_2) \leq 1 - \alpha^2 + \alpha^2 mrr^*_{k'} + O(\lambda)$. □

As shown by Theorem 1, HS-RRM is a near-optimal algorithm for RRM when α is close to 1, and HITTINGSET is optimally solved. When α is smaller, its approximation factor decreases quadratically with α. When only approximate solutions for HITTINGSET are available, HS-RRM still has a bicriteria approximation for the optimal MRR $mrr^*_{k'}$ of a smaller result size $k' < k$.

Next, we provide an upper bound of the MRR of the output $\mathcal{S}$ of HS-RRM in Theorem 2.

Theorem 2. *The maximum regret ratio $mrr(\mathcal{S})$ of the output $\mathcal{S}$ in Algorithm 1 is at most $1 - \alpha^2 + O(k^{-\frac{2}{d-1}})$ for any fixed dimensionality $d = O(1)$ and approximation factor $\alpha^{-1} = O(1)$. Additionally, in Lines 10–14 of Algorithm 1, if $k > d$ and each X_i for $i \in [d]$ is included in $\mathcal{S}_1$, it will hold that $mrr(\mathcal{S}) \leq 1 - \alpha + O((k-d)^{-\frac{2}{d-1}})$.*

Proof. A ϵ'-kernel of size k is built by running Lines 10–14 of Algorithm 1. According to the analysis of (Agarwal, Har-Peled, and Varadarajan 2004), we have $k = O(\epsilon'^{-\frac{2}{d-1}})$ and thus equivalently $\epsilon' = O(k^{-\frac{2}{d-1}})$. The property of ϵ'-kernel ensures that $\max_{X\in\mathcal{S}_1} f'_{\boldsymbol{v}}(X) \geq (1-\epsilon') \max_{X\in\mathcal{S}_{\mathcal{N}}} f'_{\boldsymbol{v}}(X)$ for any $\boldsymbol{v} \in \mathbb{S}^{d-1}_+$. By applying a similar procedure as given in the proof of Lemma 2 using Eq. (2), we further get $\max_{X\in\mathcal{S}_{\mathcal{N}}} f'_{\boldsymbol{v}}(X) \geq f'_{\boldsymbol{v}}(X'_{\boldsymbol{w}}) \geq \alpha f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) - \delta\sqrt{d} - \frac{\delta\sqrt{d}}{\alpha}$, where $\boldsymbol{w} \in \mathcal{N}$ with $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq \cos(\delta)$. In the case of $\boldsymbol{v} = \boldsymbol{b}_i$, since $1 \leq f'_{\boldsymbol{b}_i}(X^*_{\boldsymbol{b}_i}) \leq \frac{1}{\alpha}$, we have

$$\max_{X\in\mathcal{S}_{\mathcal{N}}} f'_{\boldsymbol{b}_i}(X) \geq \left(\alpha^2 - (1+\alpha)\delta\sqrt{d}\right) \cdot f'_{\boldsymbol{b}_i}(X^*_{\boldsymbol{b}_i}).$$

Thus, $rr_{f_{\boldsymbol{b}_i},\mathcal{I}}(\mathcal{S}_1) \leq 1-\alpha^2+O(k^{-\frac{2}{d-1}})$ when $\alpha^{-1} = O(1)$. The same regret ratio holds for any vector $\boldsymbol{v} \in \mathbb{S}^{d-1}_+$ with $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) \leq \frac{1}{\sqrt{d}}$, as indicated in Lemma 2. For any vector $\boldsymbol{v} \in \mathbb{S}^{d-1}_+$ with $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) > \frac{1}{\sqrt{d}}$, $\max_{X\in\mathcal{S}_{\mathcal{N}}} f'_{\boldsymbol{v}}(X) \geq \left(\alpha - \frac{2\delta d}{\alpha}\right) \cdot f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}})$ and thus the regret ratio is $1-\alpha+O(k^{-\frac{2}{d-1}})$. To sum up, we have $mrr(\mathcal{S}_1) \leq 1 - \alpha^2 + O(k^{-\frac{2}{d-1}})$.

When $k > d$, it is possible to improve the above bound on $mrr(\mathcal{S}_1)$ by using a similar method as (Soma and Yoshida 2017; Feng and Qian 2021). In Lines 10–14 of Algorithm 1, we first add the d solutions $X_1, \ldots, X_d$ to $\mathcal{S}_1$ and then draw $k-d$ uniform vectors for $\mathbb{S}^{d-1}_R$ for ϵ'-kernel computation. By applying the above adaptation, the regret ratio for $\boldsymbol{b}_i$ and $\boldsymbol{v} \in \mathbb{S}^{d-1}_+$ with $f'_{\boldsymbol{v}}(X^*_{\boldsymbol{v}}) \leq \frac{1}{\sqrt{d}}$ is reduced to $1-\alpha$. Nevertheless, the regret ratios for remaining vectors increase to $1 - \alpha + O((k-d)^{-\frac{2}{d-1}})$. Accordingly, we have $mrr(\mathcal{S}) \leq 1 - \alpha + O((k-d)^{-\frac{2}{d-1}})$. When $k \gg d$, it can be simply treated as $mrr(\mathcal{S}) \leq 1 - \alpha + O(k^{-\frac{2}{d-1}})$. □

By constructing an instance of the RRM problem, we give a lower bound on the MRR of any size-k set of feasible solutions in Theorem 3.

Theorem 3. *For any $d, k \in \mathbb{Z}^+$, there exists an instance of the RRM problem such that any size-k set $\mathcal{S} \subseteq \mathcal{I}$ has a maximum regret ratio $\Omega(k^{-\frac{2}{d-1}})$.*

Proof. We construct such an RRM instance by generalizing (Nanongkai et al. 2012, Theorem 8) and (Soma and Yoshida 2017, Theorem 11). Let the ground set V be a δ-net $\mathcal{N}$ of $\mathbb{S}_+^{d-1}$. Then, the d objective functions are defined as $f_i(S) = \max_{\boldsymbol{v} \in S} v_i$ for all $i \in [d]$ and $S \subseteq V$. Obviously, $f_i : 2^V \to \mathbb{R}^+$ is a nonnegative, monotone, and submodular function. Next, the set of feasible solutions is defined by a cardinality constraint $r = 1$, i.e., $\mathcal{I} = \{S \subseteq V : |S| = 1\}$. Since V is a δ-net, for any $\boldsymbol{w} \in \mathbb{S}_+^{d-1}$, there exists $\boldsymbol{v} \in V$ such that $f_{\boldsymbol{w}}(\{\boldsymbol{v}\}) = \langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq \cos(\delta)$. To find a set $\mathcal{S}$ of feasible solutions with $mrr(\mathcal{S}) \leq \epsilon$, we should guarantee that $\max_{X \in \mathcal{S}} f_{\boldsymbol{w}}(X) \geq (1-\epsilon)\cos(\delta)$ for any $\boldsymbol{w} \in \mathbb{S}_+^{d-1}$. To ensure $\langle \boldsymbol{v}, \boldsymbol{w} \rangle \geq (1-\epsilon)\cos(\delta)$, we require that

$$\|\boldsymbol{v} - \boldsymbol{w}\| \leq \sqrt{2 - 2(1-\epsilon)\cos(\delta)} \leq \sqrt{2 - 2\cos(\delta) + 2\epsilon}$$

This implies that it should find a set $\mathcal{S}$ such that there exists $\boldsymbol{v} \in \mathcal{S}$ with $\|\boldsymbol{v} - \boldsymbol{w}\| \leq \sqrt{2 - 2\cos(\delta) + 2\epsilon}$ for every $\boldsymbol{w} \in \mathbb{S}_+^{d-1}$ to guarantee $mrr(\mathcal{S}) \leq \epsilon$. Note that the area of $\mathbb{S}_+^{d-1}$ is $\frac{dC_d}{2^d}$, where C_d is a constant depending only on d. We hope to find a set of hyperspheres of radius $R = \sqrt{2 - 2\cos(\delta) + 2\epsilon}$ to cover $\mathbb{S}_+^{d-1}$ so that $mrr(\mathcal{S}) \leq \epsilon$. Since the area each hypersphere can cover is at most $dC_d R^{d-1}$, the number of hyperspheres needed is at least

$$|\mathcal{S}| \geq \frac{dC_d 2^{-d}}{dC_d R^{d-1}} = \frac{1}{2} \cdot \big(\frac{1}{2R}\big)^{d-1}.$$

Suppose that $\delta \ll \varepsilon$. When $|\mathcal{S}| \leq k$, $\epsilon = \Omega(k^{-\frac{2}{d-1}})$. For the RRM instance as constructed above, the MRR of any size-k set of feasible solutions is $\Omega(k^{-\frac{2}{d-1}})$. □

By combining the results of Theorems 2 and 3, we conclude that HS-RRM is nearly asymptotically optimal for RRM on any fixed dimensionality $d = O(1)$, where the gap lies in the approximation factor α for submodular maximization, which cannot be bridged unless $P = NP$.

Finally, we analyze the complexity of the HS-RRM algorithm. We assume that the time complexity of calling the oracle $\mathcal{A}$ is $t(\mathcal{A})$. To the best of our knowledge, we have $t(\mathcal{A}) = \Omega(n)$, where $n = |V|$, for almost all approximate oracles. We also consider $d = O(1)$ and $\alpha^{-1} = O(1)$. The time complexity of computing the δ-net is $O(\delta^{-\frac{1}{d-1}})$ (denoted as $t(\mathcal{N})$ for short). Thus, the time to compute $X'_{\boldsymbol{w}}$ for each $\boldsymbol{w} \in \mathcal{N}$ is $O(t(\mathcal{A})t(\mathcal{N}))$. Then, the time for computing the ϵ'-kernel and $\mathcal{S}_1$ is $O(kt(\mathcal{N}))$ when a trivial linear scan is used for nearest neighbor search. The `while` loop has at most $O(\log(\frac{1}{\lambda}))$ iterations. At each iteration, the time to build the set system is $O(t^2(\mathcal{N}))$. When the greedy algorithm is used for HITTINGSET, the time complexity is $O(kt^2(\mathcal{N}))$. To sum them up, the overall time complexity of HS-RRM is $O(t(\mathcal{A})t(\mathcal{N}) + k\log(\frac{1}{\lambda})t^2(\mathcal{N}))$.

Experiments

In this section, we experimentally compare HS-RRM with four existing algorithms for RRM: COORDINATE and POLYTOPE in (Soma and Yoshida 2017), and RRMS and RRMS* in (Feng and Qian 2021). We performed the experiments on multi-objective *weighted maximum coverage* and *data summarization* problems. Our experiments were conducted on a server with an Intel® Xeon® W-2123 3.60GHz processor and 32GB memory running Ubuntu 18.04. All the algorithms were implemented in Python 3. We adapted the implementation of the four existing algorithms by Feng and Qian (2021). Our code and data are publicly available at https://github.com/yhwang1990/code-rrm-release. Note that computing the maximum regret ratio (MRR) of a set $\mathcal{S}$ for RRM exactly is infeasible since the number of weight vectors to consider is infinite and submodular maximization is NP-hard. Therefore, we estimate a lower bound of $mrr(\mathcal{S})$ by sampling a validation set of 1,000 functions from $\mathbb{S}_+^{d-1}$ and computing the maximum of regret ratios of $\mathcal{S}$ found among the validation set.

Multi-Objective Weighted Maximum Coverage

Given a universe U of elements, where each element $u \in U$ has a weight $w(u) \geq 0$, and a collection $V = \{S_1, \ldots, S_n\}$ of subsets of U, the weighted maximum coverage (WMC) problem aims to find a subset X of size at most r from V such that the sum of the weights of elements covered by X is maximized, i.e., $\arg\max_{X \subseteq V : |X| \leq r} \sum_{u \in \bigcup_{S \in X} S} w(u)$. The objective function of WMC is nonnegative, monotone, and submodular. It is well known that WMC is an NP-hard problem (Karp 1972). In the multi-objective WMC, each element $u \in U$ is associated with a d-dimensional weight vector $\boldsymbol{w}(u) = (w_1(u), \ldots, w_d(u))$, and each objective function $f_i(X)$ is defined as $\sum_{u \in \bigcup_{S \in X} S} w_i(u)$ accordingly. The greedy algorithm (Nemhauser, Wolsey, and Fisher 1978) with an approximation factor $\alpha = 1 - \frac{1}{e}$ is used as the oracle for WMC in all algorithms.

We use a real-world dataset *email-Eu-core* from SNAP[1], which is a directed graph with 1,004 vertices and 25,571 edges, for our experiments. We set U as the set of all vertices and build a set S for each vertex that contains itself and its (out)neighbors. The weight vector of a vertex is an indicator of which community it belongs to among d communities returned by the asynchronous fluid algorithm (Parés et al. 2017). The feasible solution set $\mathcal{I}$ is defined by a cardinality constraint $r = 10$. Since HS-RRM, RRMS and RRMS* are randomized algorithms for $d > 2$, we run each of them ten times independently and report the average and standard deviation of the estimated MRRs. The estimated MRRs of five algorithms we compare with varying the size of the solutions k when the number of objective functions $d = 2, 5$ and varying d when $k = 10, 25$ are shown in the first row of Fig. 1. Note that COORDINATE always outputs a set of $k = d$ solutions. Thus, its MRR is not changed w.r.t. k. We observe that HS-RRM outperforms all other algorithms in estimated MRRs for different values of k's and d's. This verifies the

[1] http://snap.stanford.edu/data/#email

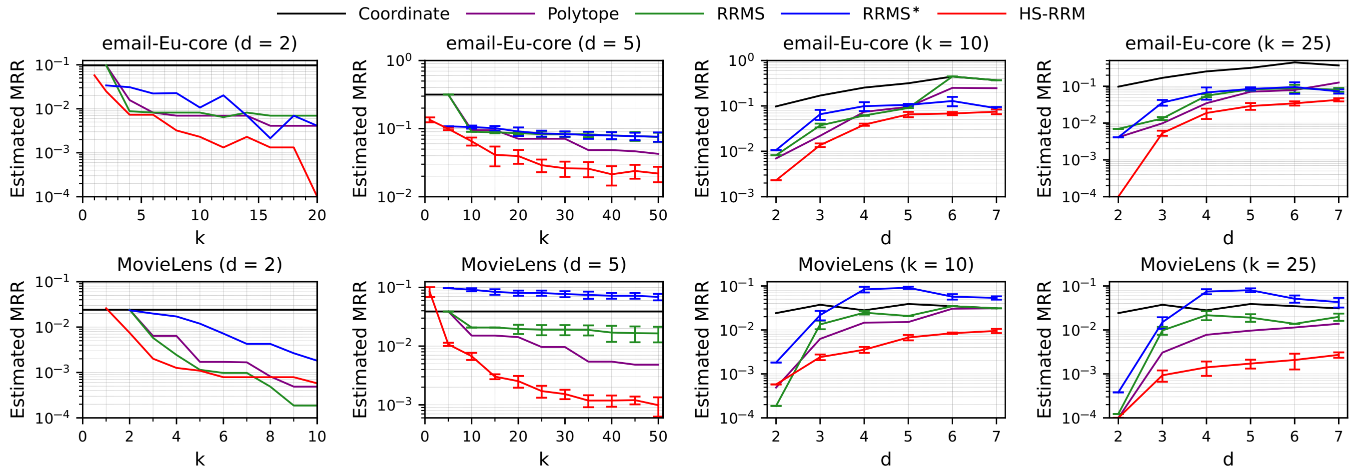

Figure 1: Estimated maximum regret ratios of different algorithms on multi-objective weighted maximum coverage and data summarization with varying the size of the solutions k for $d = 2, 5$ and the number of objective functions d for $k = 10, 25$.

theoretical result that HS-RRM has the lowest worst-case MRR among all known algorithms. In addition, HS-RRM is the only algorithm that provides a valid set of solutions when $k < d$ (e.g., $k = 1$ in Fig. 1). Finally, we also find that the MRR of each algorithm decreases with k and increases with d. This is as expected from Theorem 2 and the intuitions that the MRR is reduced by adding more solutions to $\mathcal{S}$ and RRM is harder for higher dimensionality. But this trend is not strictly followed due to the error in MRR estimation and the randomness in weight vector sampling.

Multi-Objective Data Summarization

Let V be a ground set of n items. The goal of data summarization is to select a subset X of representative items from V. Following several existing studies (Lin and Bilmes 2011; Mirzasoleiman, Badanidiyuru, and Karbasi 2016; Soma and Yoshida 2017) on summarization problems, we consider two possibly conflicting measures of representativeness, namely *coverage* and *diversity*. Specifically, a set X of items is regarded as achieving high coverage if, for any item $v_i \in V$, the items in X are highly similar to v_i. Moreover, the diversity of X is measured by the dissimilarities between items within X. Here, we consider that V consists of a set of movies, each of which is represented as a vector acquired from user rating data by any low-rank matrix factorization method (Koren, Bell, and Volinsky 2009). The similarity s_{ij} of two movies v_i and v_j is computed as the cosine similarity $\cos(\boldsymbol{v}_i, \boldsymbol{v}_j)$ of their representing vectors. Formally, we adopt the objective function in (Mirzasoleiman, Badanidiyuru, and Karbasi 2016) for summarization as follows:

$$f(X) = \sum_{v_i \in V} \sum_{v_j \in X} s_{ij} - \sum_{v_i \in X} \sum_{v_j \in X} s_{ij}.$$

It is known that $f(X)$ is nonnegative, submodular, but non-monotone. We extend the summarization problem to be multi-objective by dividing V into d subsets $V_1, \ldots, V_d$ and defining an objective function f_i on each V_i instead of V. The random greedy algorithm (Buchbinder et al. 2014) with an approximation factor $\alpha = \frac{1}{e}$ is used as the oracle in all algorithms.

In our experiments, we use the MovieLens dataset[2] that contains 100,000 ratings from 943 users on 1,682 movies. We set the dimensionality of user and movie vectors to 25. We run the k-means clustering (Lloyd 1982; Arthur and Vassilvitskii 2007) to partition the movies into d subsets. Like WMC, the feasible solution set $\mathcal{I}$ is defined by $r = 10$, and each algorithm runs ten times independently. The estimated MRRs of five algorithms with varying the size of the solutions k when the number of objective functions $d = 2, 5$ and varying d when $k = 10, 25$ are shown in the second row of Fig. 1. Similar to the results of multi-objective WMC, we also observe that HS-RRM outperforms all other algorithms in terms of estimated MRRs in most cases, which further confirms its theoretical advantage. In a few cases (e.g., k ranges from 8 to 10 when $d = 2$), the MRRs of the output of HS-RRM are slightly higher than those of POLYTOPE and RRMS. This is because we set $\lambda = 10^{-3}$ for HS-RRM when $d = 2$. Consequently, it only adds random solutions to $\mathcal{S}$ when the estimated MRR has reached 10^{-3}. A feasible approach to solving this problem is to use a smaller value of λ. Nevertheless, HS-RRM shows more apparent advantages over all other algorithms than those for WMC in most cases.

Conclusion

In this paper, we investigated the regret ratio minimization (RRM) problem in multi-objective submodular maximization. We proposed an improved algorithm called HS-RRM that not only achieved tighter theoretical bounds but also demonstrated significantly better empirical performance in real-world applications than existing algorithms. Especially, we proved that HS-RRM was a nearly asymptotically optimal algorithm for RRM on any fixed d. An interesting future work is to generalize the RRM problem to weakly or non-submodular objective functions.

[2]https://grouplens.org/datasets/movielens/

Acknowledgements

This work was supported by the National Natural Science Foundation of China under grant No. 62202169. We thank anonymous reviewers for their constructive comments to help improve this work.

References

Agarwal, P. K.; Har-Peled, S.; and Varadarajan, K. R. 2004. Approximating extent measures of points. *J. ACM*, 51(4): 606–635.

Agarwal, P. K.; Kumar, N.; Sintos, S.; and Suri, S. 2017. Efficient Algorithms for k-Regret Minimizing Sets. In *16th International Symposium on Experimental Algorithms (SEA 2017)*, 7:1–7:23.

Amiridi, M.; Kargas, N.; and Sidiropoulos, N. D. 2021. Information-Theoretic Feature Selection via Tensor Decomposition and Submodularity. *IEEE Trans. Signal Process.*, 69: 6195–6205.

Anari, N.; Haghtalab, N.; Naor, S.; Pokutta, S.; Singh, M.; and Torrico, A. 2019. Structured Robust Submodular Maximization: Offline and Online Algorithms. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics (AISTATS '19)*, 3128–3137.

Arthur, D.; and Vassilvitskii, S. 2007. k-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*, 1027–1035.

Asudeh, A.; Nazi, A.; Zhang, N.; and Das, G. 2017. Efficient Computation of Regret-Ratio Minimizing Set: A Compact Maxima Representative. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD '17)*, 821–834.

Becker, R.; D'Angelo, G.; Ghobadi, S.; and Gilbert, H. 2022. Fairness in Influence Maximization through Randomization. *J. Artif. Intell. Res.*, 73: 1251–1283.

Börzsönyi, S.; Kossmann, D.; and Stocker, K. 2001. The Skyline Operator. In *Proceedings of the 17th International Conference on Data Engineering (ICDE '01)*, 421–430.

Buchbinder, N.; Feldman, M.; Naor, J.; and Schwartz, R. 2014. Submodular Maximization with Cardinality Constraints. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '14)*, 1433–1452.

Călinescu, G.; Chekuri, C.; Pál, M.; and Vondrák, J. 2011. Maximizing a Monotone Submodular Function Subject to a Matroid Constraint. *SIAM J. Comput.*, 40(6): 1740–1766.

Cao, W.; Li, J.; Wang, H.; Wang, K.; Wang, R.; Wong, R. C.-W.; and Zhan, W. 2017. k-Regret Minimizing Set: Efficient Algorithms and Hardness. In *20th International Conference on Database Theory (ICDT 2017)*, 11:1–11:19.

Cormen, T. H.; Leiserson, C. E.; Rivest, R. L.; and Stein, C., eds. 2022. *Introduction to Algorithms, Fourth Edition*. Cambridge, MA, USA: The MIT Press.

Ene, A.; and Nguyen, H. L. 2019. A Nearly-Linear Time Algorithm for Submodular Maximization with a Knapsack Constraint. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, 53:1–53:12.

Feige, U. 1998. A Threshold of ln n for Approximating Set Cover. *J. ACM*, 45(4): 634–652.

Feng, C.; and Qian, C. 2021. Multi-Objective Submodular Maximization by Regret Ratio Minimization with Theoretical Guarantee. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12302–12310.

Fisher, M. L.; Nemhauser, G. L.; and Wolsey, L. A. 1978. An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral Combinatorics: Dedicated to the memory of D.R. Fulkerson*, 73–87.

Fu, X.; Bhatt, R. R.; Basu, S.; and Pavan, A. 2021. Multi-Objective Submodular Optimization with Approximate Oracles and Influence Maximization. In *2021 IEEE International Conference on Big Data (Big Data)*, 328–334.

Gershtein, S.; Milo, T.; and Youngmann, B. 2021. Multi-Objective Influence Maximization. In *Advances in Database Technology - EDBT 2021, 24th International Conference on Extending Database Technology, Nicosia, Cyprus, March 23 - 26, Proceedings*, 145–156.

Karp, R. M. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, 85–103.

Kempe, D.; Kleinberg, J. M.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '03)*, 137–146.

Koren, Y.; Bell, R. M.; and Volinsky, C. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8): 30–37.

Krause, A.; and Golovin, D. 2014. Submodular Function Maximization. In Bordeaux, L.; Hamadi, Y.; and Kohli, P., eds., *Tractability: Practical Approaches to Hard Problems*, 71–104. Cambridge, UK: Cambridge University Press.

Krause, A.; and Guestrin, C. 2005. Near-optimal Nonmyopic Value of Information in Graphical Models. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*, 324–331.

Krause, A.; McMahan, H. B.; Guestrin, C.; and Gupta, A. 2008. Robust Submodular Observation Selection. *J. Mach. Learn. Res.*, 9(93): 2761–2801.

Kumar, N.; and Sintos, S. 2018. Faster Approximation Algorithm for the k-Regret Minimizing Set and Related Problems. In *Proceedings of the Twentieth Workshop on Algorithm Engineering and Experiments (ALENEX '18)*, 62–74.

Kumari, L.; and Bilmes, J. 2021. Submodular Span, with Applications to Conditional Data Summarization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12344–12352.

Lin, H.; and Bilmes, J. 2011. A Class of Submodular Functions for Document Summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 510–520.

Lloyd, S. P. 1982. Least squares quantization in PCM. *IEEE Trans. Inf. Theory*, 28(2): 129–136.

Mirzasoleiman, B.; Badanidiyuru, A.; and Karbasi, A. 2016. Fast Constrained Submodular Maximization: Personalized Data Summarization. In *Proceedings of the 33nd International Conference on Machine Learning (ICML '16)*, 1358–1367.

Nanongkai, D.; Lall, A.; Das Sarma, A.; and Makino, K. 2012. Interactive Regret Minimization. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data (SIGMOD '12)*, 109–120.

Nanongkai, D.; Sarma, A. D.; Lall, A.; Lipton, R. J.; and Xu, J. J. 2010. Regret-Minimizing Representative Databases. *Proc. VLDB Endow.*, 3(1): 1114–1124.

Nemhauser, G. L.; Wolsey, L. A.; and Fisher, M. L. 1978. An analysis of approximations for maximizing submodular set functions—I. *Math. Program.*, 14(1): 265–294.

Ohsaka, N.; and Matsuoka, T. 2021. Approximation algorithm for submodular maximization under submodular cover. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI '21)*, 792–801.

Parés, F.; Gasulla, D. G.; Vilalta, A.; Moreno, J.; Ayguadé, E.; Labarta, J.; Cortés, U.; and Suzumura, T. 2017. Fluid Communities: A Competitive, Scalable and Diverse Community Detection Algorithm. In *Complex Networks & Their Applications VI*, 229–240.

Peng, P.; and Wong, R. C.-W. 2014. Geometry approach for k-regret query. In *2014 IEEE 30th International Conference on Data Engineering (ICDE '14)*, 772–783.

Qian, C.; Yu, Y.; and Zhou, Z. 2015. Subset Selection by Pareto Optimization. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, 1774–1782.

Soma, T.; and Yoshida, Y. 2017. Regret Ratio Minimization in Multi-Objective Submodular Function Maximization. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI '17)*, 905–911.

Storandt, S.; and Funke, S. 2019. Algorithms for Average Regret Minimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 1600–1607.

Tang, J.; Tang, X.; Lim, A.; Han, K.; Li, C.; and Yuan, J. 2021. Revisiting Modified Greedy Algorithm for Monotone Submodular Maximization with a Knapsack Constraint. *Proc. ACM Meas. Anal. Comput. Syst.*, 5(1): 08:1–08:22.

Tsang, A.; Wilder, B.; Rice, E.; Tambe, M.; and Zick, Y. 2019. Group-Fairness in Influence Maximization. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 5997–6005.

Udwani, R. 2018. Multi-objective Maximization of Monotone Submodular Functions with Cardinality Constraint. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, 9513–9524.

Wang, Y.; Li, Y.; Wong, R. C.-W.; and Tan, K.-L. 2021a. A Fully Dynamic Algorithm for k-Regret Minimizing Sets. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 1631–1642.

Wang, Y.; Mathioudakis, M.; Li, Y.; and Tan, K.-L. 2021b. Minimum Coresets for Maxima Representation of Multidimensional Data. In *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '21)*, 138–152.

Xie, M.; Wong, R. C.-W.; Li, J.; Long, C.; and Lall, A. 2018. Efficient k-Regret Query Algorithm with Restriction-free Bound for any Dimensionality. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*, 959–974.

Zheng, J.; Wang, Y.; Wang, X.; and Ma, W. 2022. Continuous k-Regret Minimization Queries: A Dynamic Coreset Approach. *IEEE Trans. Knowl. Data Eng.* Forthcoming.