
Fairness in Submodular Maximization over a Matroid Constraint

Marwa El Halabi
Samsung - SAIT AI Lab
Montreal

Jakub Tarnawski
Microsoft Research

Ashkan Norouzi-Fard
Google Zurich

Thuy-Duong Vuong
Stanford University

Abstract

Submodular maximization over a matroid constraint is a fundamental problem with various applications in machine learning. Some of these applications involve decision-making over datapoints with sensitive attributes such as gender or race. In such settings, it is crucial to guarantee that the selected solution is fairly distributed with respect to this attribute. Recently, fairness has been investigated in submodular maximization under a cardinality constraint in both the streaming and offline settings, however the more general problem with matroid constraint has only been considered in the streaming setting and only for monotone objectives. This work fills this gap. We propose various algorithms and impossibility results offering different trade-offs between quality, fairness, and generality.

1 INTRODUCTION

Machine learning algorithms are increasingly used in decision-making processes. This can potentially lead to the introduction or perpetuation of bias and discrimination in automated decisions. Of particular concern are sensitive areas such as education, hiring, credit access, bail decisions, and law enforcement (Munoz et al., 2016; White House OSTP, 2022; European Union FRA, 2022). There has been a growing body of work attempting to mitigate these risks by developing *fair* algorithms for fundamental problems including classification (Zafar et al., 2017), ranking (Celis et al., 2018c), clustering (Chierichetti et al., 2017), voting (Celis et al., 2018a), matching (Chierichetti et al., 2019), influence maximization (Tsang et al., 2019), data summarization (Celis et al., 2018b), and many others.

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

In this work, we address fairness in the fundamental problem of submodular maximization over a matroid constraint, in the offline setting. Submodular functions model a diminishing returns property that naturally occurs in a variety of machine learning problems such as active learning (Golovin and Krause, 2011), data summarization (Lin and Bilmes, 2011), feature selection (Das and Kempe, 2011), and recommender systems (El-Arini and Guestrin, 2011). Matroids represent a popular and expressive notion of independence systems that encompasses a broad spectrum of useful constraints, e.g. cardinality, block cardinality, linear independence, and connectivity constraints.

Several definitions of algorithmic fairness have been proposed in the literature, but no universal metric exists. Here we adopt the common notion used in many prior studies (Celis et al., 2018a,b,c; Chierichetti et al., 2017, 2019) that requires a solution to be *fair* with respect to a sensitive attribute such as race or gender. Formally, given a set V of n items, each item is associated with a color c representing a sensitive attribute. Let V_1, \dots, V_C denote the corresponding C disjoint groups of items of the same color. A set $S \subseteq V$ is *fair* if $\ell_c \leq |S \cap V_c| \leq u_c$ for all c , for some chosen lower and upper bounds $\ell_c, u_c \in \mathbb{N}$. This notion subsumes several other fairness definitions, e.g. diversity rules (Cohoon et al., 2013; Biddle, 2006), statistical parity (Dwork et al., 2012), or proportional representation rules (Monroe, 1995; Brill et al., 2017). See Celis et al. (2018a, Section 4) for a more detailed overview.

Without fairness, the problem of maximizing a submodular function over a matroid constraint in the offline setting has been studied extensively. For monotone objectives, a tight $(1 - 1/e)$ -approximation is known (Calinescu et al., 2011b; Feige, 1998), while for non-monotone objectives the best known approximation is 0.401 given in (Buchbinder and Feldman, 2023). An information-theoretic hardness of 0.478 was also shown for the non-monotone case in (Oveis Gharan and Vondrák, 2011).

Fair submodular maximization has been considered under a *cardinality* constraint, in both the offline and

streaming models. Celis et al. (2018a) presented a tight $(1 - 1/e)$ -approximation to the problem in the offline setting with monotone objectives. In the streaming setting, El Halabi et al. (2020) gave a one-pass 0.3178-approximation algorithm in the monotone case, and a one-pass $0.1921q$ -approximation algorithm in the non-monotone case, where $q = 1 - \max_c \frac{\ell_c}{|V_c|}$.¹ They also showed that achieving a better than q -approximation in this case requires $\Omega(n)$ memory.

Fair submodular maximization under a general *matroid* constraint was only studied in the streaming setting, and only for monotone objectives. El Halabi et al. (2023) provided a one-pass $1/2$ -approximation algorithm that uses time and memory exponential in the rank of the matroid; this factor is tight (Feldman et al., 2020). They also gave a polynomial-time two-pass $1/11.656$ -approximation which uses $O(k \cdot C)$ memory and violates the fairness lower bounds by a factor 2.

While some applications necessitate algorithms that use small memory compared to the data size, others can afford a larger memory budget. In this work, we focus on the classic offline setting, which has not yet been studied despite being more basic than streaming. Our results show that it is possible to obtain stronger performance guarantees in this case. We also consider both monotone and non-monotone objectives, which cover a wider range of applications.

1.1 Our Contributions

We present several algorithms and impossibility results for fair matroid submodular maximization (FMSM) in the offline setting, with trade-offs between quality, fairness, and generality. Table 1 summarizes our approximation results and prior ones for this problem.

First, we observe that the streaming algorithms of El Halabi et al. (2023) for monotone objectives apply in the offline setting. This yields a $1/2$ -approximation in *exponential time* with respect to the rank of the matroid. And, with a slight modification, the two-pass algorithm therein achieves a $1/(4 + \epsilon)$ -approximation in polynomial time, with a factor-2 violation of the fairness lower bounds. We extend the latter result to non-monotone objectives, obtaining a $(1 - \beta)/(8 + \epsilon)$ -approximation with a $1/\beta$ violation of the fairness lower bounds, for any $\beta \in [0, 1/2]$ (Theorem 3.1).

It is of course preferable to obtain algorithms that satisfy the fairness constraints exactly. However, we give strong evidence that this is very challenging in general.

¹These approximations are better than the ones stated in (El Halabi et al., 2020). They follow from their results by plugging in the state-of-the-art approximations of Feldman et al. (2021, Theorem 1 and 5) for streaming submodular maximization over a matroid constraint.

A commonly used approach for submodular maximization problems is *relax-and-round*, which consists of first solving a continuous relaxation of the problem obtained using the *multilinear extension* (see definition in Section 2) of the objective over the convex hull $P_{\mathcal{F}}$ of the domain, then rounding the fractional solution. Most studied submodular maximization problems that admit a constant-factor approximation algorithm also admit one based on the relax-and-round approach. Unfortunately, we show that for FMSM, this approach cannot yield a better than $O(1/\sqrt{n})$ -approximation in general (Theorem 3.2), even for monotone objectives.

Yet, if we allow the fairness constraint to be satisfied only *in expectation*, we can adapt this approach to obtain a tight $(1 - 1/e)$ -approximation in expectation for monotone objectives. For non-monotone objectives, the approximation becomes $(1 - r - \epsilon)/4$ where r is the minimum ℓ_∞ -norm of any vector in $P_{\mathcal{F}}$ or its complement. Both algorithms are also guaranteed to only violate the fairness constraints by a multiplicative factor expected to be small in practice (Theorem 3.3).

In fact, the factor r is closely tied to the hardness of non-monotone FMSM. We show that no algorithm can achieve a better than $(1 - r)$ -approximation in sub-exponential time, even without the matroid constraint (Theorem 3.7). In the case of uniform or no matroid, $1 - r$ coincides with the notion of excess ratio $q = 1 - \max_c \frac{\ell_c}{|V_c|}$ studied by El Halabi et al. (2020). As such, this inapproximability result extends the one provided therein for the streaming setting.

While the $(1 - r)$ factor in the approximation for non-monotone FMSM cannot be improved even for the uniform matroid, we show that the $1/4$ factor can be improved to 0.401 in this case (Theorem 3.6). This improves over the $0.1921q$ -approximation of (El Halabi et al., 2020) in the streaming setting.

Finally, we study an interesting subclass of FMSM, which we call decomposable FMSM, where the objective decomposes into submodular functions on the equivalence classes of the matroid (see Definition 2.3) or on the color groups of the fairness constraint. This subclass enables us to consider fairness in the important *submodular welfare problem* (see Section 2). We provide a tight $(1 - 1/e)$ -approximation to decomposable FMSM for monotone objectives and a $(1 - r - \epsilon)/4$ for non-monotone objectives, in expectation, *without* violating the fairness constraint (Corollary 4.4).

1.2 Additional Related Work

Before this paper, non-monotone FMSM was only studied in the offline setting in the special case of cardinality constraint by Yuan and Tang (2023), who obtained a

Table 1: Summary of results for FMSM in the offline setting (only polynomial-time algorithms with respect to n, k and C are included). In the “Fairness Approx.” column, a (θ_1, θ_2) entry means the algorithm’s output satisfies $|S \cap V_c| \in [\theta_1 \ell_c, \theta_2 u_c]$. The parameters $\epsilon > 0, \beta \in [0, 1/2]$ are arbitrary constants, and $r = \min_{x \in P_{\mathcal{F}} \cup (1 - P_{\mathcal{F}})} \|x\|_{\infty}$.

Function	Matroid	Approx. Ratio	Fairness Approx.
Monotone	General	$1 - 1/e$ (Thm. 3.3)	$(1, 1)$ in expectation
Non-monotone	General	$(1 - r - \epsilon)/4$ (Thm. 3.3)	$(1 - \sqrt{3 \ln(2C)/\ell_c}, 1 + \sqrt{3 \ln(2C)/u_c})$
Monotone	General	$1/(4 + \epsilon)$ (El Halabi et al., 2023)	$(1/2, 1)$
Non-monotone	General	$(1 - \beta)/(8 + \epsilon)$ (Thm. 3.1)	$(\beta, 1)$
Monotone	Uniform	$1 - 1/e$ (Celis et al., 2018a)	$(1, 1)$
Non-monotone	Uniform	$0.401(1 - r)$ (Thm. 3.6)	$(1, 1)$
Monotone decomposable	General	$1 - 1/e$ (Thm. 4.4)	$(1, 1)$
Non-monotone decomposable	General	$(1 - r - \epsilon)/4$ (Thm. 4.4)	$(1, 1)$

0.2005-approximation for a specific setting of fairness bounds where $\frac{\ell_c}{|V_c|} = a$ and $\frac{u_c}{|V_c|} = b$ for all c , for some constants $a, b \in [0, 1]$.² Our results recover their guarantee (see Theorem 3.6 and the discussion below it).

Several works have studied other special cases and variants of FMSM. Wang et al. (2021) studied a special case of monotone FMSM in the streaming setting, where fairness lower and upper bounds are equal for each color, i.e., $\ell_c = u_c$ for all c and without any matroid constraint. Tang et al. (2023) and Tang and Yuan (2023) studied a randomized variant of FMSM with cardinality constraint in the offline setting, where fairness constraint only needs to be satisfied in expectation.

A closely related problem to FMSM is submodular maximization over two matroid constraints; FMSM reduces to this problem when $\ell_c = 0$ for all c . Lee et al. (2010) gave a $1/(2 + \epsilon)$ -approximation for this problem for monotone objectives, and a $1/(4 + \epsilon)$ -approximation for non-monotone objectives.

An important special case of decomposable FMSM arises by considering fairness in the submodular welfare problem. Without fairness, this is a well-studied problem for which a tight $(1 - 1/e)$ -approximation is known (Vondrák, 2008; Khot et al., 2005). The submodular welfare problem has many practical applications, including resource allocation in wireless networks, budget allocation in advertising campaigns, and division of resources among multiple stakeholders. In several of these applications, it is important to ensure that the selected allocation is fair. Fairness in the submodular welfare problem has been studied by Benabbou et al. (2021), Aziz et al. (2023), and Sun et al. (2023). How-

ever, these works consider different notions of fairness than we do: fairness is imposed on the value received by each individual in the selected allocation, instead of the size of the allocation assigned to each group of individuals.

In this work, we focus on the case where the color groups are disjoint. The case where groups can overlap was considered in (Celis et al., 2018a) for the special case of monotone FMSM with cardinality constraint. They show that when elements can be assigned to 3 or more colors, even determining feasibility is NP-hard. But if fairness constraints are allowed to be violated, they give a $(1 - 1/e - o(1))$ -approximation algorithm, which satisfies the fairness constraint in *expectation*. Our result for general monotone FMSM (Theorem 3.3) is based on the same tools used for this result.

Fairness in submodular maximization has also been studied under a different notion of fairness (Wang et al., 2024; Tsang et al., 2019; Tang and Yuan, 2023) which imposes a lower bound on the value that each group derives from the solution, where groups are not necessarily subsets of V , and the value is modeled by a monotone submodular function. Both the variant of fair submodular maximization considered in this line of work and our work can be formulated as variants of multi-objective submodular maximization problems (Krause et al., 2008; Udwani, 2018; Chekuri et al., 2010); see Appendix C for details.

2 PRELIMINARIES

Consider a ground set V with n items and a non-negative set function $f : 2^V \rightarrow \mathbb{R}_+$. We say that f is *submodular* if $f(Y \cup \{e\}) - f(Y) \geq f(X \cup \{e\}) - f(X)$ for any two sets $Y \subseteq X \subseteq V$ and any element $e \in V \setminus X$. Moreover, f is *monotone* if $f(Y) \leq f(X)$ for any two sets $Y \subseteq X \subseteq V$, and *non-monotone* otherwise.

² The approximation given in (Yuan and Tang, 2023) is $0.401/2$ for $a \leq 1/2$ and $0.401/3$ otherwise. However, their analysis can be easily modified to show $0.401/2 = 0.2005$ -approximation for any a .

Throughout the paper, we assume that f is given as a value oracle that computes $f(S)$ given $S \subseteq V$.

Matroids. A *matroid* is a family of sets $\mathcal{I} \subseteq 2^V$ with the following properties:

- if $X \subseteq Y$ and $Y \in \mathcal{I}$, then $X \in \mathcal{I}$;
- if $X, Y \in \mathcal{I}$ with $|X| < |Y|$, then there exists $e \in Y$ such that $X + e \in \mathcal{I}$.

We use $X + e$ for $X \cup \{e\}$. We assume that the matroid is given as an independence oracle. We call a set $X \subseteq V$ *independent* if $X \in \mathcal{I}$ and a *base* if it is also maximal with respect to inclusion. All the bases of a matroid share the same cardinality k , referred to as the *rank* of the matroid. *Partition matroids* are an important class of matroids where V is partitioned into $\bigcup_i G_i$, where each block G_i has an upper bound $k_i \in \mathbb{N}$, and a set X is independent if $|X \cap G_i| \leq k_i$ for all i .

Fair Matroid Submodular Maximization (FMSM). Let V be partitioned into C sets, where V_c denotes items of color c for all $c \in [C] = \{1, \dots, C\}$. Given fairness bounds $(\ell_c, u_c)_{c \in [C]}$ and a matroid $\mathcal{I} \subseteq 2^V$ of rank k , the collections of *fair* sets \mathcal{C} and *feasible* sets \mathcal{F} are defined as:

$$\begin{aligned} \mathcal{C} &= \{S \subseteq V \mid \ell_c \leq |S \cap V_c| \leq u_c \quad \forall c \in [C]\}, \\ \mathcal{F} &= \mathcal{I} \cap \mathcal{C}. \end{aligned}$$

In FMSM, the goal is to find a set $S \in \mathcal{F}$ that maximizes $f(S)$. We use OPT to refer to the optimal value, i.e., $\text{OPT} = \max_{S \in \mathcal{F}} f(S)$. We assume a feasible solution exists, i.e., $\mathcal{F} \neq \emptyset$. An algorithm is an α -approximation to FMSM if it outputs a set $S \in \mathcal{F}$ such that $f(S) \geq \alpha \cdot \text{OPT}$.

We denote the convex hull of indicator vectors of fair sets by $P_{\mathcal{C}} = \text{conv}(\{\mathbf{1}_S \mid S \in \mathcal{C}\})$ (fairness polytope), of independent sets by $P_{\mathcal{I}} = \text{conv}(\{\mathbf{1}_S \mid S \in \mathcal{I}\})$ (matroid polytope), and of feasible sets by $P_{\mathcal{F}} = \text{conv}(\{\mathbf{1}_S \mid S \in \mathcal{F}\})$ (feasible polytope). We also define the complement of $P_{\mathcal{F}}$ as $\mathbf{1} - P_{\mathcal{F}} = \{x \in [0, 1]^n \mid \mathbf{1} - x \in P_{\mathcal{F}}\}$, where $\mathbf{1}$ is the all-ones vector. We recall in the following lemma some useful properties shown by El Halabi et al. (2023, Lemma C.1, Corollary C.4, Theorem C.5).

Lemma 2.1. *The following hold:*

- $P_{\mathcal{C}} = \{x \in [0, 1]^n : \sum_{i \in V_c} x_i \in [\ell_c, u_c] \quad \forall c \in [C]\}$.
- $P_{\mathcal{F}}$ corresponds to the intersection of the matroid and fairness polytopes: $P_{\mathcal{F}} = P_{\mathcal{I}} \cap P_{\mathcal{C}}$.
- $P_{\mathcal{F}}$ is solvable, i.e., linear functions can be maximized over it in polynomial time.

The following lemma from (Buchbinder et al., 2014, Lemma 2.2) will be useful in the non-monotone case.

Lemma 2.2. *Let g be a non-negative submodular function and let B be a random subset of V containing every element of V with probability at most p (not necessarily independently). Then $\mathbb{E}[g(B)] \geq (1 - p)g(\emptyset)$.*

Multilinear extension. An important concept in submodular maximization is the multilinear extension $F : [0, 1]^n \rightarrow \mathbb{R}_+$ of a submodular function f :

$$F(x) = \mathbb{E}[f(R(x))] = \sum_{S \subseteq V} f(S) \prod_{i \in S} x_i \prod_{j \in V \setminus S} (1 - x_j),$$

where $R(x)$ is the set obtained by independently selecting each element $i \in V$ with probability x_i .

Decomposable FMSM. We also study a subclass of the FMSM problem where the objective decomposes into submodular functions on the equivalence classes of the matroid or on the color groups of the fairness constraint. We recall the notion of equivalent elements in a matroid (Chekuri et al., 2010, Definition I.3).

Definition 2.3. Two elements $i, j \in V$ are equivalent in a matroid \mathcal{I} if for any set $S \in \mathcal{I}$ not containing i and j , $S + i \in \mathcal{I}$ if and only if $S + j \in \mathcal{I}$.

This defines an equivalence relation. For example, the equivalence classes of a partition matroid are the partition groups G_i .

Definition 2.4. Let $\mathcal{G} \subseteq 2^V$ be the equivalence classes of the matroid \mathcal{I} . We say that f is a decomposable submodular function over \mathcal{F} if

$$f(S) = \sum_{G \in \mathcal{G}} f_{1,G}(S \cap G) + \sum_{c=1}^C f_{2,c}(S \cap V_c),$$

where $f_{1,G}, f_{2,c}$ are submodular functions.

Submodular welfare problem. One noteworthy example of decomposable submodular functions arises from the submodular welfare problem. In this problem, we are given a set of items B and a set of agents A each with a monotone submodular utility function $w_i : 2^B \rightarrow \mathbb{R}_+$. The goal is to distribute the items among the agents to maximize the social welfare $\sum_{i \in A} w_i(S_i)$, where $(S_i)_{i \in A}$ are the disjoint sets of items assigned to each agent. It is known that this problem can be written as a monotone submodular maximization over a partition matroid constraint by defining the ground set as $V = A \times B$, the objective as $f(S) = \sum_{i \in A} w_i(S_i)$ where $S_i = \{e \in B \mid (i, e) \in S\}$, and using the partition matroid $\mathcal{I} = \{S \subseteq V \mid (\forall e \in B) |S \cap (A \times \{e\})| \leq 1\}$ (Lehmann et al., 2006). In certain applications, it is important to ensure that each group of agents receives

a *fair* allocation of items. To that end, we introduce the *fair submodular welfare problem*, where each agent is assigned exactly one color $c \in [C]$; A_c is the set of agents of color c . Let $V_c = A_c \times B$ and $f_c(S) = \sum_{i \in A_c} w_i(S_i)$, then f_c is monotone submodular and $f(S) = \sum_{c=1}^C f_c(S \cap V_c)$ is a monotone decomposable submodular function over \mathcal{F} . The fair submodular welfare problem is then a special case of monotone decomposable FMSM.

3 GENERAL CASE

In this section, we provide algorithms and impossibility results for FMSM with general submodular objectives.

3.1 Algorithms with Lower Bounds Violation

We start by providing constant-factor approximation algorithms for FMSM that violate the fairness lower bounds by a constant factor. Such an algorithm can be directly obtained in the monotone case from the two-pass streaming algorithm of El Halabi et al. (2023, Section 4). In the first pass, their algorithm finds a feasible set S ; in the second pass, S is extended to an $\alpha/2$ -approximate solution which violates the lower bounds by a factor 2, using an α -approximation algorithm \mathcal{A} for monotone submodular maximization over the intersection of two matroids. Using the state-of-the-art *offline* $1/(2 + \epsilon)$ -approximation algorithm of Lee et al. (2010, Theorem 3.1) as \mathcal{A} , instead of a streaming algorithm, directly yields an improved $1/(4 + \epsilon)$ -approximation ratio.

We adapt this algorithm to non-monotone objectives by randomly dropping an appropriate number of elements from S to balance the loss in objective value and the violation of the fairness lower bounds. We defer the details to Appendix A.

Theorem 3.1. *There exists a polynomial-time algorithm for non-monotone FMSM, which outputs a set S such that (i) $S \in \mathcal{I}$, (ii) $\lfloor \beta \ell_c \rfloor \leq |V_c \cap S| \leq u_c$ for any color $c \in [C]$, and (iii) $\mathbb{E}[f(S)] \geq (1 - \beta)\text{OPT}/(8 + \epsilon)$ for any $\beta \in [0, 1/2]$ and $\epsilon > 0$.*

Proof sketch. We extend the result of El Halabi et al. (2023, Lemma 4.3 and 4.4) to non-monotone objectives; showing that dropping elements from S allows us to only lose a factor $(1 - \beta)$, in expectation, in the objective value by Lemma 2.2, while keeping at least $\lfloor \beta \ell_c \rfloor$ elements from each color. Plugging in the state-of-the-art $1/(4 + \epsilon)$ -approximation algorithm of Lee et al. (2010, Theorem 4.1) as \mathcal{A} directly yields a $(1 - \beta)/(8 + \epsilon)$ -approximation, in expectation. \square

For example, if we accept a factor 2 violation of the fairness lower bounds ($\beta = 1/2$), Theorem 3.1 gives a

$1/(16 + \epsilon)$ -approximation for non-monotone FMSM.

Note that if we use a streaming algorithm for \mathcal{A} , the adapted algorithm becomes a two-pass streaming algorithm, as in the monotone case. In particular, using the state-of-the-art $1/7.464$ -approximation algorithm by Garg et al. (2021, Theorem 19) yields a $(1 - \beta)/14.928$ -approximation in the streaming setting (Theorem A.3).

3.2 Hardness of Rounding

The optimal approximation algorithms for maximizing a submodular function over a matroid constraint in (Calinescu et al., 2011b; Feldman et al., 2011) rely on first approximately solving a continuous relaxation of the problem obtained by replacing f by its *multilinear extension* and the matroid constraint by its convex hull, then rounding the solution, which can be done without any loss of utility using *pipage or swap rounding* (Calinescu et al., 2011b; Chekuri et al., 2010). It is then natural to attempt a similar relax-and-round approach here: first approximately solve the continuous relaxation of FMSM,

$$\max_{x \in P_{\mathcal{F}}} F(x), \quad (1)$$

and then round the obtained solution to a set in \mathcal{F} . Unfortunately, this approach fails in the presence of the fairness constraint! In particular, we show that the integrality gap of Problem (1) is $\Omega(\sqrt{n})$, even for a monotone objective and a partition matroid. Hence, it is not possible to obtain better than $O(1/\sqrt{n})$ -approximation using this approach.

Theorem 3.2. *There is a family of FMSM instances where f is monotone and \mathcal{I} is a partition matroid, for which the integral optimum solution has value 1, but the multilinear extension admits a fractional solution of value $\Omega(\sqrt{n})$.*

Proof sketch. We encode bipartite perfect matching as an instance of \mathcal{F} where \mathcal{I} is a partition matroid. Then we construct an instance where the objective is monotone and the bipartite graph has $\Theta(\sqrt{n})$ perfect matchings each of value 1, while taking $x \in P_{\mathcal{F}}$ as the average of the indicator vectors of the perfect matchings yields $F(x) = \Theta(\sqrt{n})$. \square

3.3 Algorithms with Expected Fairness

Motivated by Theorem 3.2, we adapt the relax-and-round approach to ignore the fairness constraints during the rounding phase. This allows us to use the randomized swap rounding algorithm of Chekuri et al. (2010, Section IV) to obtain an independent-set solution which preserves the same utility as the fractional solution, and which satisfies the fairness constraints *in*

expectation. This solution is further guaranteed, with constant probability, to violate the fairness bounds by at most a multiplicative factor, expected to be small in practice. This approach is inspired from Celis et al. (2018a, Theorem 14), who considered a special case of FMSM with a monotone objective and a uniform matroid, but with possibly non-disjoint color groups. Theorem 3.3 extends their result to any matroid and non-monotone objectives. The non-monotone case is more challenging than the monotone case; we show in Section 3.5 that it is impossible to obtain an approximation better than $(1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty})$ for non-monotone FMSM in sub-exponential time, even for solving the continuous Problem (1).

Theorem 3.3. *There exists a polynomial-time algorithm for FMSM which outputs a solution $S \in \mathcal{I}$ such that $\mathbb{E}[|S \cap V_c|] \in [\ell_c, u_c]$ and $\mathbb{E}[f(S)] \geq \alpha \text{OPT}$, where $\alpha = 1 - 1/e$ if f is monotone and $\alpha = (1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} - \epsilon)/4$ otherwise, for any $\epsilon > 0$. Moreover, the solution satisfies with constant probability the following for all $c \in [C]$,*

$$\left(1 - \sqrt{\frac{3 \ln(2C)}{\ell_c}}\right) \ell_c \leq |S \cap V_c| \leq \left(1 + \sqrt{\frac{3 \ln(2C)}{u_c}}\right) u_c.$$

Proof. We first obtain an α -approximation fractional solution $x \in P_{\mathcal{F}}$ for $\max_{x \in P_{\mathcal{F}}} F(x)$. For monotone objectives, we can use the continuous greedy algorithm of Calinescu et al. (2011a, Section 3.1 and Appendix A), which achieves a $(1 - 1/e)$ -approximation for maximizing the multilinear extension over any integral polytope, with high probability.³ Note that $P_{\mathcal{F}}$ is integral as it is the convex hull of integral points.

For non-monotone objectives, we can use the Frank-Wolfe type algorithm of Du (2022, Sections 3.5 and 4.5) which achieves a $(1 - \min_{x \in P} \|x\|_{\infty} - \epsilon)/4$ -approximation for maximizing the multilinear extension over any polytope P , with high probability — see also (Mualet and Feldman, 2023, Section 3) for an explicit variant of the algorithm and its analysis. To obtain the approximation ratio $(1 - \min_{x \in P_{\mathcal{F}} \cup \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} - \epsilon)/4$, we apply the algorithm of Du (2022) to both Problem (1) and its complement: $\max_{x \in \mathbb{1} - P_{\mathcal{F}}} \bar{F}(x)$, where \bar{F} is the multilinear extension of the complement \bar{f} of f ; $\bar{f}(S) = f(V \setminus S)$. Both the continuous greedy and the Frank-Wolfe type algorithms run in polynomial time if the constraint polytope is solvable, which is the case for $P_{\mathcal{F}}$ and $\mathbb{1} - P_{\mathcal{F}}$ by Lemma 2.1-c.

Next we round the fractional solution x to an independent (but not necessarily feasible) set $S \in \mathcal{I}$ using the randomized swap rounding scheme of Chekuri

³The result in (Calinescu et al., 2011a) is given for matroid polytopes, but it applies more generally for any integral polytope.

et al. (2010, Section IV). The rounded solution is guaranteed to satisfy $\mathbb{E}[|S \cap V_c|] = x(V_c) \in [\ell_c, u_c]$, $\mathbb{E}[f(S)] \geq F(x) \geq \alpha \text{OPT}$ (see Theorem II.1 therein). The rest of the claim follows from the concentration bounds given in (Chekuri et al., 2010, Theorem II.1) and union bound. See details in Appendix B.2. \square

As discussed in (Celis et al., 2018a) — see discussion under Theorem 14 therein — the violation of the fairness constraints in the above theorem is expected to be small, i.e., $\sqrt{3 \ln(2C)}/u_c \leq \sqrt{3 \ln(2C)}/\ell_c \ll 1$, in typical fairness applications. However, in general, this violation can be arbitrarily bad. In particular, when \mathcal{F} encodes a bipartite perfect matching constraint, as in the proof of Theorem 3.2, we have $C = \Theta(\sqrt{n})$ and $u_c = \ell_c = 1$, hence $\sqrt{3 \ln(2C)}/u_c = \sqrt{3 \ln(2C)}/\ell_c \gg 1$.

Note that the approximation ratio for non-monotone objectives in Theorem 3.3 can be as good as $1/4$ in some cases (for example if $\ell_c = 0$ for all c or $u_c = |V_c|$ for all c), but it can also be arbitrarily bad. To illustrate this, we compute below the minimum ℓ_{∞} -norm of $P_{\mathcal{F}}$ and its complement for two special cases of $P_{\mathcal{F}}$. See Appendix B.2 for the full proofs.

Example 3.4 (Uniform matroid). Let \mathcal{I} be the uniform matroid with rank k . We order the color groups such that $\frac{\ell_1}{|V_1|} \leq \dots \leq \frac{\ell_C}{|V_C|}$ and let t be the largest index $t \in [C]$ such that $\frac{\ell_t}{|V_t|} \sum_{c=1}^t |V_c| + \sum_{c=t+1}^C \ell_c \leq k$ ($t \geq 1$ is well-defined since $\mathcal{F} \neq \emptyset$). Then $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = \min\{\max_c \frac{\ell_c}{|V_c|}, 1 - \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}\}$, where $\tau = \frac{k - \sum_{c=t+1}^C \ell_c}{\sum_{c=1}^t |V_c|}$.

Proof sketch. We construct closed-form solutions to $\min_{x \in P_{\mathcal{F}}} \|x\|_{\infty}$ and $\min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty}$ with ℓ_{∞} -norms $\max_c \frac{\ell_c}{|V_c|}$ and $1 - \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}$, respectively. \square

The fairness bounds are often set such that the representation of color groups in a fair set is proportional to their representation in the ground set, i.e., $\ell_c, u_c \propto k|V_c|/n$. If in particular $\ell_c \leq k|V_c|/n$ for all c , we get in Example 3.4 $t = C, \tau = k/n$ and $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} \leq \min\{k/n, 1 - k/n\} \leq 1/2$. The corresponding approximation ratio in Theorem 3.3 for non-monotone objectives is then at least $1/8 - \epsilon$.

Example 3.5 (Bipartite perfect matching). Let \mathcal{F} be the feasible set arising from the bipartite perfect matching problem in the proof of the integrality gap Theorem 3.2. We have $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = 1 - \Theta(1/\sqrt{n})$.

Proof sketch. We show that the point x used in the proof of Theorem 3.2 is a solution of

$$\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} \text{ with } \|x\|_{\infty} = \|\mathbb{1} - x\|_{\infty} = 1 - \Theta(1/\sqrt{n}). \quad \square$$

The corresponding approximation factor for Example 3.5 in Theorem 3.3 for non-monotone objectives is then $O(1/\sqrt{n})$. So relaxing the fairness constraints was not helpful in this case!

Recall that the $(1 - 1/e)$ -approximation for monotone objectives in Theorem 3.3 is tight even without the fairness constraint (Feige, 1998). For non-monotone objectives, the $1/4$ factor is also tight for maximizing the multilinear extension of a submodular function over a general polytope (Muelem and Feldman, 2023, Theorem 5.1); so improving it, if possible, would require more specialized algorithms for solving Problem (1).

3.4 Uniform Matroid Case

In the special case of the uniform matroid, both the fairness violation and the approximation ratio for non-monotone objectives given in Theorem 3.3 can be improved. In particular, recall that a $(1 - 1/e)$ -approximation can be obtained in this case for monotone objectives, without any violation of the fairness constraint (Celis et al., 2018a, Theorem 18). For non-monotone objectives, we show that a $0.401(1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty})$ -approximation can be obtained in expectation, also without any violation of the fairness constraint.⁴

Theorem 3.6. *There exists a polynomial-time algorithm for non-monotone FMSM where \mathcal{I} is a uniform matroid, which achieves $0.401(1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty})$ -approximation in expectation.*

Proof sketch. El Halabi et al. (2020, Theorem 5.2) presented a polynomial-time algorithm with an expected $(1 - \max_c \frac{\ell_c}{|V_c|})\alpha$ -approximation for non-monotone FMSM over a uniform matroid, given any α -approximation algorithm for non-monotone submodular maximization over a matroid constraint.⁵ Plugging in the state-of-the-art 0.401-approximation of Buchbinder and Feldman (2023, Corollary 1.3) then yields a $0.401(1 - \max_c \frac{\ell_c}{|V_c|})$ -approximation in expectation. Recall from Example 3.4 that $\min_{x \in P_{\mathcal{F}}} \|x\|_{\infty} = \max_c \frac{\ell_c}{|V_c|}$.

To obtain the potentially better approximation factor $0.401(1 - \min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty})$, we solve (as in Theorem 3.3) the complement problem $\max_{V \setminus S \in \mathcal{F}} \bar{f}(S)$.

⁴Note that the 0.401 factor in the approximation given in Theorem 3.6 can be improved if a better approximation for non-monotone submodular maximization over a general or a partition matroid is developed in the future.

⁵The result in (El Halabi et al., 2020) is given for the streaming setting, but it also holds offline.

We follow a similar strategy as El Halabi et al. (2020). Namely, we first drop the lower bounds from the constraint. The resulting problem $\max_{S \subseteq V} \{\bar{f}(S) : |S \cap V_c| \leq |V_c| - \ell_c, \forall c \in [C]\}$ is a non-monotone submodular maximization problem over a partition matroid. Hence, a solution S with an expected 0.401-approximation can be obtained for it. Next, we augment S to a feasible solution by carefully sampling enough random elements from each color to satisfy the lower bounds. Each element $e \in V \setminus S$ is added with probability at most \bar{x}_e , where \bar{x} is a solution of $\min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty}$ as in Example 3.4. Finally, we use Lemma 2.2 to bound the loss in value resulting from the additional elements. The sampling step is more involved here than in the original problem because \bar{x} is non-integral. See details in Appendix B.3. \square

Theorem 3.6 recovers the result of Yuan and Tang (2023) for the special case where $\frac{\ell_c}{|V_c|} = a$ and $\frac{u_c}{|V_c|} = b$ for all $c \in [C]$, for some constants $a, b \in [0, 1]$.² Indeed, in that case $t = C$ and $\tau = k/n \geq a$ in Example 3.4, and $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = \min\{a, 1 - \min\{k/n, b\}\}$, so the resulting approximation ratio is at least $0.401 \max\{1 - a, a\} \geq 0.2005$.

More generally, for the typical choice of fairness bounds discussed earlier in Section 3.3, where $\ell_c \leq k|V_c|/n$ for all c , the resulting approximation ratio is also at least $0.401/2 = 0.2005$.

3.5 Non-monotone Hardness

In this section, we show that no approximation ratio better than $(1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty})$ can be obtained for non-monotone FMSM in sub-exponential time, even without the matroid constraint.

Theorem 3.7. *For any r in the form $r = 1 - \frac{1}{t}$ with $t \geq 2$ and any $\epsilon > 0$, a $(1 - r + \epsilon)$ -approximation for FMSM that works for instances where $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} \leq r$ requires an exponential number of value queries. This is also the case for the continuous relaxation of FMSM, $\max_{x \in P_{\mathcal{F}}} F(x)$. Furthermore, this is true even in the two special cases where (i) $C = 2$ and $\mathcal{I} = 2^V$ (no matroid constraint), (ii) $C = 1$ and \mathcal{I} is a partition matroid.*

Proof sketch. The proof follows from the hardness results given in (Vondrák, 2013, Theorem 1.2 and 1.9) for non-monotone submodular maximization over the bases of a matroid, by showing that the hard instance used therein is an instance of both special cases (i) and (ii) of FMSM. See Appendix B.4. \square

Theorem 3.7 implies that any constant-factor approximation for FMSM or its continuous relaxation requires an exponential number of value queries in general.

Note that a uniform matroid with $k = n$ is equivalent to having no matroid constraint. In that case, Example 3.4 shows that $\min_{x \in P_{\mathcal{F}} \cup (1 - P_{\mathcal{F}})} \|x\|_{\infty} = \min\{\max_c \frac{\ell_c}{|V_c|}, 1 - \min_c \frac{u_c}{|V_c|}\}$, since $t = C$ and $\tau = 1$. Hence, Example 3.4 and Theorem 3.7 yield a lower bound of $\max\{1 - \max_c \frac{\ell_c}{|V_c|}, \min_c \frac{u_c}{|V_c|}\}$ for the fairness constraint alone. This matches the lower bound given in (El Halabi et al., 2020, Theorem 5.1) for FMSM in the streaming setting – the bound therein is $1 - \max_c \frac{\ell_c}{|V_c|}$ and is stated for the uniform matroid, but the hard instance used in the proof does not use any matroid and has $\min_c \frac{u_c}{|V_c|} = 0$.

On the other hand, note that Example 3.5 does not necessarily imply a hardness of $O(1/\sqrt{n})$ for the bipartite perfect matching constraint, since the hard instance in Theorem 3.7 is not an instance of that.

4 DECOMPOSABLE CASE

In this section, we study a special case of FMSM where f is a decomposable submodular function over \mathcal{F} (Definition 2.4). We follow the relax-and-round approach discussed in Section 3.2. Unlike the general case, this approach works here; we show that the integrality gap of Problem (1) is 1 in this case. To that end, we show that the randomized swap rounding algorithm proposed in (Chekuri et al., 2010, Section V) for the intersection of two matroids can be applied to \mathcal{F} .

We achieve this by reducing rounding on \mathcal{F} to rounding on the intersection of two matroids, one with the same equivalence classes (Definition 2.3) as \mathcal{I} and the other with the color groups V_c as equivalence classes. Our reduction uses the following simple observation.

Fact 4.1. *Given any family of sets $\mathcal{J} \subseteq 2^V$, let $\tilde{\mathcal{J}}$ denote the family of extendable sets: $\tilde{\mathcal{J}} = \{S \subseteq V \mid \text{there exists } S' \in \mathcal{J} \text{ such that } S \subseteq S'\}$. Then \mathcal{J} and $\tilde{\mathcal{J}}$ have the same maximal sets.*

When the matroid in \mathcal{F} is the uniform matroid, $\tilde{\mathcal{F}}$ was shown to form a matroid (El Halabi et al., 2020, Lemma 4.1). We can then view any fair set as a base in a matroid, as shown by the following lemma. See Appendix B.5 for the proofs of this section.

Lemma 4.2. *Let $\tilde{\mathcal{C}}_k := \{S \subseteq V \mid \text{there exists } S' \in \mathcal{C}, |S'| \leq k \text{ such that } S \subseteq S'\}$. Then any fair set $S \in \mathcal{C}$ is a base of the matroid $\tilde{\mathcal{C}}_k$ with $k = |S|$ and vice versa.*

We are now ready to give our reduction.

Theorem 4.3. *There exists a polynomial-time rounding scheme which, given $x \in P_{\mathcal{F}}$, returns a feasible set $S \in \mathcal{F}$ that for any decomposable submodular function f over \mathcal{F} satisfies $\mathbb{E}[f(S)] \geq F(x)$.*

Proof sketch. Since $P_{\mathcal{F}}$ is solvable (Lemma 2.1-c), we

can write x as a convex combination of sets in \mathcal{F} in strongly polynomial time (Schrijver, 2003, Theorem 5.15); i.e., $x = \sum_t \alpha_t \mathbf{1}_{I_t}$ for some $I_t \in \mathcal{F}, \alpha_t \geq 0, \sum_t \alpha_t = 1$. Lemma 4.2 implies that every set I_t is in the intersection of the matroids \mathcal{I} and $\tilde{\mathcal{C}}_k$, where k is the size of the largest set I_t . Then x is a valid input for swap rounding. If all sets I_t have the same size k , then the returned rounded solution $S \in \mathcal{I} \cap \tilde{\mathcal{C}}_k$ will also have size k , and thus is a base of $\tilde{\mathcal{C}}_k$. By Lemma 4.2, S is then also a fair set. The claim then follows from (Chekuri et al., 2010, Theorem II.3) by noting that the equivalence classes of $\tilde{\mathcal{C}}_k$ are the color groups V_c , and thus any decomposable submodular function over \mathcal{F} is also decomposable over $\mathcal{I} \cap \tilde{\mathcal{C}}_k$. If sets I_t do not have the same size, we reduce to the equal-sizes case by adding dummy elements. \square

Theorem 4.3 allows us to obtain the same approximation factors as in Theorem 3.3 but without violating the fairness constraint.

Corollary 4.4. *There exists a polynomial-time algorithm for decomposable FMSM which achieves, in expectation, $(1 - 1/e)$ -approximation if f is monotone and $(1 - \min_{x \in P_{\mathcal{F}} \cup (1 - P_{\mathcal{F}})} \|x\|_{\infty} - \epsilon)/4$ -approximation otherwise, for any $\epsilon > 0$.*

Proof. This follows directly from the proof of Theorem 3.3 and Theorem 4.3. Note also that both the continuous greedy and the Frank Wolfe type algorithms used to solve Problem (1) return a fractional solution x which is already a convex combination of sets in \mathcal{F} . So this step can be skipped in the proof of Theorem 4.3. \square

Recall that the fair submodular welfare problem (c.f., Section 2) is a special case of monotone decomposable FMSM. Corollary 4.4 then gives a $(1 - 1/e)$ -approximation for it.

In the special cases where $C = 1$ or \mathcal{I} is the uniform matroid, Corollary 4.4 applies to any submodular objective, since $V_1 = V$ in the first case and all elements are equivalent in the second case, i.e., $\mathcal{G} = \{V\}$ in Definition 2.3. Hence, Corollary 4.4 recovers the result of Celis et al. (2018a, Theorem 18) for monotone objectives. For non-monotone objectives, the approximation guarantee in Corollary 4.4 can be of course improved to the one given in Theorem 3.6.

It follows also that the hardness result in Theorem 3.7 applies to decomposable FMSM too.

Corollary 4.5. *For any r of the form $r = 1 - \frac{1}{t}$ with $t \geq 2$ and any $\epsilon > 0$, a $(1 - r + \epsilon)$ -approximation for decomposable FMSM that works for instances where $\min_{x \in P_{\mathcal{F}} \cup (1 - P_{\mathcal{F}})} \|x\|_{\infty} \leq r$ requires an exponential number of value queries. This is also the case for*

the continuous relaxation of *FMSM*, $\max_{x \in P_{\mathcal{F}}} F(x)$. Furthermore, this is true even in the two special cases where (i) $C = 2$ and $\mathcal{I} = 2^V$ (no matroid constraint), (ii) $C = 1$ and \mathcal{I} is a partition matroid.

Proof. This follows directly from the proof of Theorem 3.7, by noting that any submodular function is decomposable over \mathcal{F} in both special cases; in (i) \mathcal{I} is a n -uniform matroid and in (ii) $C = 1$. \square

5 FUTURE DIRECTIONS

Our work leaves the following tantalizing question open: *is there a constant-factor approximation to monotone FMSM (without fairness violations)?* Even more fundamentally: is there a constant-factor approximation to the problem of maximizing a monotone submodular function over the set of *perfect* matchings in a bipartite graph? On one hand, given our integrality gap in Theorem 3.2 for the multilinear extension, a positive answer would require novel algorithmic techniques. On the other hand, in a plenary SODA talk, Vondrák (2013) stated the following as a fundamental question: “Can we approximate every maximization problem with a monotone submodular objective (up to constant factors) if we can approximate it with a linear objective?” As maximum-weight perfect matching is in P , a superconstant hardness of approximation for the submodular perfect matching problem would imply a negative answer to Vondrák’s question.

Acknowledgements

We thank Rico Zenklusen, Jakab Tardos, and Moran Feldman for helpful discussions.

References

- Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. Computing welfare-maximizing fair allocations of indivisible goods. *European Journal of Operational Research*, 307(2):773–784, 2023. (Cited on 3)
- Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations for matroid rank valuations. *ACM Transactions on Economics and Computation*, 9(4):1–41, 2021. (Cited on 3)
- Dan Biddle. *Adverse impact and test validation: A practitioner’s guide to valid and defensible employment testing*. Gower Publishing, Ltd., 2006. (Cited on 1)
- Markus Brill, Jean-François Laslier, and Piotr Skowron. Multiwinner approval rules as apportionment methods. In *AAAI*, pages 414–420. AAAI Press, 2017. (Cited on 1)
- N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. Submodular maximization with cardinality constraints. SIAM, 2014. (Cited on 4)
- Niv Buchbinder and Moran Feldman. Constrained submodular maximization via a nonsymmetric technique. *Mathematics of Operations Research*, 44(3):988–1005, 2019. (Cited on 18)
- Niv Buchbinder and Moran Feldman. Constrained submodular maximization via new bounds for dr-submodular functions. *arXiv preprint arXiv:2311.01129*, 2023. (Cited on 1, 7)
- Gruia 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, 2011a. (Cited on 6, 16)
- Gruia 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, 2011b. (Cited on 1, 5)
- L. Elisa Celis, Lingxiao Huang, and Nisheeth K. Vishnoi. Multiwinner voting with fairness constraints. In *IJCAI*, pages 144–151. ijcai.org, 2018a. (Cited on 1, 2, 3, 6, 7, 8, 17, 21)
- L. Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. Fair and diverse dpp-based data summarization. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 715–724. PMLR, 2018b. (Cited on 1)
- L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. Ranking with fairness constraints. In *ICALP*, volume 107 of *LIPICs*, pages 28:1–28:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018c. (Cited on 1)
- Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding for matroid polytopes and applications. *arXiv preprint arXiv:0909.4348*, 2009. (Cited on 21)
- Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584. IEEE Computer Society, 2010. (Cited on 3, 4, 5, 6, 8, 17, 20, 21)
- Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *NIPS*, pages 5029–5037, 2017. (Cited on 1)
- Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Matroids, matchings, and fair-

- ness. In *AISTATS*, volume 89 of *Proceedings of Machine Learning Research*, pages 2212–2220. PMLR, 2019. (Cited on 1)
- Joanne McGrath Cohoon, James P. Cohoon, Seth Reichelson, and Selwyn Lawrence. Effective recruiting for diversity. In *FIE*, pages 1123–1124. IEEE Computer Society, 2013. (Cited on 1)
- Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *ICML*, pages 1057–1064. Omnipress, 2011. (Cited on 1)
- Donglei Du. Lyapunov function approach for approximation algorithm design and analysis: with applications in submodular maximization. *arXiv preprint arXiv:2205.12442*, 2022. (Cited on 6, 16)
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012. (Cited on 1)
- Khalid El-Arini and Carlos Guestrin. Beyond keyword search: discovering relevant scientific literature. In *KDD*, pages 439–447. ACM, 2011. (Cited on 1)
- Marwa El Halabi, Slobodan Mitrovic, Ashkan Norouzi-Fard, Jakab Tardos, and Jakub Tarnawski. Fairness in streaming submodular maximization: Algorithms and hardness. In *NeurIPS*, 2020. (Cited on 2, 7, 8, 18, 20)
- Marwa El Halabi, Federico Fusco, Ashkan Norouzi-Fard, Jakab Tardos, and Jakub Tarnawski. Fairness in streaming submodular maximization over a matroid constraint. In *International Conference on Machine Learning*, pages 9150–9171. PMLR, 2023. (Cited on 2, 3, 4, 5, 13, 14, 20)
- European Union FRA. *Bias in Algorithms — Artificial Intelligence and Discrimination*. European Union Agency for Fundamental Rights. Publications Office of the European Union, Luxembourg, 2022. (Cited on 1)
- Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998. (Cited on 1, 7)
- Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 570–579. IEEE, 2011. (Cited on 5)
- Moran Feldman, Ashkan Norouzi-Fard, Ola Svensson, and Rico Zenklusen. The one-way communication complexity of submodular maximization with applications to streaming and robustness. In *STOC*, pages 1363–1374. ACM, 2020. (Cited on 2)
- Moran Feldman, Paul Liu, Ashkan Norouzi-Fard, Ola Svensson, and Rico Zenklusen. Streaming submodular maximization under matroid constraints. *arXiv preprint arXiv:2107.07183*, 2021. (Cited on 2)
- Paritosh Garg, Linus Jordan, and Ola Svensson. Semi-streaming algorithms for submodular matroid intersection. In *IPCO*, volume 12707 of *Lecture Notes in Computer Science*, pages 208–222. Springer, 2021. (Cited on 5, 15)
- Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res.*, 42:427–486, 2011. (Cited on 1)
- HO Hartley and JNK Rao. Sampling with unequal probabilities and without replacement. *The Annals of Mathematical Statistics*, pages 350–374, 1962. (Cited on 19)
- Subhash Khot, Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. In *Internet and Network Economics: First International Workshop, WINE 2005, Hong Kong, China, December 15-17, 2005. Proceedings 1*, pages 92–101. Springer, 2005. (Cited on 3)
- Andreas Krause, H Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust submodular observation selection. *Journal of Machine Learning Research*, 9(Dec):2761–2801, 2008. (Cited on 3)
- Jon Lee, Maxim Sviridenko, and Jan Vondrák. Submodular maximization over multiple matroids via generalized exchange properties. *Mathematics of Operations Research*, 35(4):795–806, 2010. (Cited on 3, 5, 15)
- Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006. (Cited on 4)
- Hui Lin and Jeff A. Bilmes. A class of submodular functions for document summarization. In *ACL*, pages 510–520. The Association for Computer Linguistics, 2011. (Cited on 1)
- Burt L Monroe. Fully proportional representation. *American Political Science Review*, 89(4):925–940, 1995. (Cited on 1)
- Loay Muallem and Moran Feldman. Resolving the approximability of offline and online non-monotone dr-submodular maximization over general convex sets. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 2542–2564. PMLR, 25–27 Apr 2023. (Cited on 6, 7, 16)

- Cecilia Munoz, Smith Megan, and DJ Patil. *Big data: A report on algorithmic systems, opportunity, and civil rights*. Executive Office of the President. The White House, Washington, DC, 2016. (Cited on 1)
- Shayan Oveis Gharan and Jan Vondrák. Submodular maximization by simulated annealing. In *SODA*, pages 1098–1116. SIAM, 2011. (Cited on 1)
- Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media, 2003. (Cited on 8, 20)
- Ankang Sun, Bo Chen, and Xuan Vinh Doan. Equitability and welfare maximization for allocating indivisible items. *Autonomous Agents and Multi-Agent Systems*, 37(1):8, 2023. (Cited on 3)
- Shaojie Tang and Jing Yuan. Beyond submodularity: a unified framework of randomized set selection with group fairness constraints. *Journal of Combinatorial Optimization*, 45(4):102, 2023. (Cited on 3, 21)
- Shaojie Tang, Jing Yuan, and Twumasi Mensah-Boateng. Achieving long-term fairness in submodular maximization through randomization. *arXiv preprint arXiv:2304.04700*, 2023. (Cited on 3)
- Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, and Yair Zick. Group-fairness in influence maximization. In *IJCAI*, pages 5997–6005. ijcai.org, 2019. (Cited on 1, 3, 21)
- Rajan Udhwani. Multi-objective maximization of monotone submodular functions with cardinality constraint. *Advances in Neural Information Processing Systems*, 31, 2018. (Cited on 3, 21)
- Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 67–74, 2008. (Cited on 3)
- Jan Vondrák. Symmetry and approximability of submodular maximization problems. *SIAM Journal on Computing*, 42(1):265–304, 2013. (Cited on 7, 19)
- Jan Vondrák. Submodular functions and their applications, 2013. URL <https://theory.stanford.edu/~jvondrak/data/SODA-pleinary-talk.pdf>. Plenary talk at the ACM-SIAM Symposium on Discrete Algorithms (SODA) 2013. (Cited on 9)
- Yanhao Wang, Francesco Fabbri, and Michael Mathioudakis. Fair and representative subset selection from data streams. In *Proceedings of the Web Conference 2021*, pages 1340–1350, 2021. (Cited on 3)
- Yanhao Wang, Yuchen Li, Francesco Bonchi, and Ying Wang. Balancing utility and fairness in submodular maximization. In *Proceedings of the 27th International Conference on Extending Database Technology, EDBT*, pages 1–14, 2024. (Cited on 3, 21)
- White House OSTP. *Blueprint for an AI Bill of Rights*. White House Office of Science and Technology Policy, Washington, DC, 2022. (Cited on 1)
- Jing Yuan and Shaojie Tang. Group fairness in non-monotone submodular maximization. *Journal of Combinatorial Optimization*, 45(3):88, 2023. (Cited on 2, 3, 7)
- Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *AISTATS*, volume 54 of *Proceedings of Machine Learning Research*, pages 962–970. PMLR, 2017. (Cited on 1)

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:

- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Fairness in Submodular Maximization over a Matroid Constraint: Supplementary Materials

A Algorithms for non-monotone FMSM with lower bounds violation

In this section, we present a constant-factor approximation algorithm for FMSM, where the function f is non-monotone, which violates the fairness lower bounds by a constant factor. The algorithm is a simple adaptation of the two-pass streaming algorithm in (El Halabi et al., 2023), with both a streaming and offline variant. The first pass of our algorithm is the same (Algorithm 1). It collects a maximal independent set $I_c \in \mathcal{I}$ for each color *independently* and then constructs a feasible solution S in $\cup_c I_c$.

Algorithm 1 FAIR-RESERVOIR in (El Halabi et al., 2023)

- 1: $I_c \leftarrow \emptyset$ for all $c = 1, \dots, C$
 - 2: **for** each element e on the stream **do**
 - 3: Let c be the color of e
 - 4: **If** $I_c + e \in \mathcal{I}$ **then** $I_c \leftarrow I_c + e$
 - 5: $S \leftarrow$ a max-cardinality subset of $\bigcup_c I_c$ in $\mathcal{I} \cap \mathcal{I}_C$, where $\mathcal{I}_C = \{S \subseteq V \mid |V_c \cap S| \leq \ell_c \ \forall c \in [C]\}$
 - 6: **Return** S
-

FAIR-RESERVOIR runs in polynomial time, uses $O(k \cdot C)$ memory, and is guaranteed to return a feasible solution (El Halabi et al., 2023, Theorem 4.2).⁶ Step 5 consists in invoking any offline polynomial-time max-cardinality matroid intersection algorithm (note that \mathcal{I}_C is a partition matroid).

The second pass of our algorithm is very similar to (El Halabi et al., 2023) (Algorithm 2). The feasible set S output by FAIR-RESERVOIR is again divided into sets S_1 and S_2 . But instead of assigning elements deterministically to S_1 or S_2 in such a way that each lower bound is violated by at most a factor two, we randomly assign, for each $c \in [C]$, $\lfloor \beta \ell_c \rfloor$ elements from $S \cap V_c$ to S_1 , and $\lfloor \beta \ell_c \rfloor$ from $(S \setminus S_1) \cap V_c$ to S_2 , where $\beta \in [0, 1/2]$ is a user-defined parameter which controls the trade-off between fairness and objective value. This is possible since S is feasible so it has at least ℓ_c elements of each color c . Note that both S_1 and S_2 are independent in \mathcal{I} and satisfy $\lfloor \beta \ell_c \rfloor = |S_i \cap V_c| \leq \ell_c \leq u_c$.

The goal of the second pass is to extend S_1 and S_2 into high-value independent sets while still satisfying the fairness upper bounds and the relaxed lower bounds. To that end, similarly as (El Halabi et al., 2023), we define the following matroids:

$$\begin{aligned}
 \mathcal{I}^C &= \{X \subseteq V \mid |X \cap V_c| \leq u_c \ \forall c \in [C]\}, \\
 \mathcal{I}_1 &= \{X \subseteq V \mid X \cup S_1 \in \mathcal{I}\}, \\
 \mathcal{I}_2 &= \{X \subseteq V \mid X \cup S_2 \in \mathcal{I}\}.
 \end{aligned} \tag{2}$$

Then we use any algorithm \mathcal{A} that maximizes a *non-monotone* submodular function over two matroid constraints, and run two copies of it in parallel, one with matroids $\mathcal{I}^C, \mathcal{I}_1$ and one with matroids $\mathcal{I}^C, \mathcal{I}_2$. Let R_1, R_2 be their respective outputs. We add to R_i as many elements as necessary from S_i to satisfy the relaxed lower bounds, and return the better of the two resulting solutions.

⁶In the offline setting, we do not need to collect the sets I_c ; we can directly compute a feasible set in V . These sets are needed to maintain a low-memory usage $O(k \cdot C)$ in the streaming setting.

Algorithm 2 RANDOMIZED-FAIR-STREAMING

```

1: Input: Set  $S$  from FAIR-RESERVOIR, routine  $\mathcal{A}$ , and  $\beta \in [0, 1/2]$ 
2:  $S_1 \leftarrow \emptyset, S_2 \leftarrow \emptyset$ 
3: for each color  $c$  do
4:   Let  $B_1^c, B_2^c$  be sets of  $\lfloor \beta \ell_c \rfloor$  random elements from  $S \cap V_c$  and  $(S \cap V_c) \setminus B_1^c$ , respectively
5:    $S_1 \leftarrow S_1 \cup B_1^c$ 
6:    $S_2 \leftarrow S_2 \cup B_2^c$ 
7: Define matroids  $\mathcal{I}^C, \mathcal{I}_1, \mathcal{I}_2$  as in Equation (2)
8: Run two copies of  $\mathcal{A}$ , one for  $\mathcal{I}^C, \mathcal{I}_1$  and one for  $\mathcal{I}^C, \mathcal{I}_2$ , and let  $R_1$  and  $R_2$  be their outputs
9: for  $i = 1, 2$  do
10:   $S'_i \leftarrow R_i$ 
11:  for  $e$  in  $S_i$  do
12:    Let  $c$  be the color of  $e$ 
13:    if  $|S'_i \cap V_c| < u_c$  then  $S'_i \leftarrow S'_i + e$ 
14: Return  $S' = \arg \max(f(S'_1), f(S'_2))$ 
    
```

RANDOMIZED-FAIR-STREAMING only differs from the second-pass algorithm in (El Halabi et al., 2023, Algorithm 2) by the way elements in S are assigned to S_1 and S_2 , and the requirement that \mathcal{A} works for non-monotone objectives. Let us analyze its performance.

Lemma A.1. *The output S' of RANDOMIZED-FAIR-STREAMING satisfies (i) $S' \in \mathcal{I}$, (ii) $|S' \cap V_c| \leq u_c$, and (iii) $|S' \cap V_c| \geq \lfloor \beta \ell_c \rfloor$ for any color $c \in [C]$.*

Proof. The proof follows in a similar way to (El Halabi et al., 2023, Lemma 4.3). For $i \in \{1, 2\}$, we have:

- (i) By the definition of \mathcal{I}_i in Equation (2) the set R_i output by \mathcal{A} satisfies $R_i \cup S_i \in \mathcal{I}$. Hence, $S'_i \in \mathcal{I}$ by downward-closedness since $S'_i \setminus R_i \subseteq S_i$.
- (ii) Since $R_i \in \mathcal{I}^C$ and the elements added to it in the for loop on Line 11 never violate the upper bounds, we have $S'_i \in \mathcal{I}^C$.
- (iii) By construction S_i has at least $\lfloor \beta \ell_c \rfloor$ elements of each color c . And for any color c such that $|S'_i \cap V_c| < u_c$, all the elements in $S_i \cap V_c$ are added to S'_i . Hence, $|S'_i \cap V_c| \geq |S_i \cap V_c| \geq \lfloor \beta \ell_c \rfloor$.

□

Lemma A.2. *Assume that \mathcal{A} is an α -approximation algorithm for non-monotone submodular maximization over the intersection of two matroids. Then the output S' of RANDOMIZED-FAIR-STREAMING satisfies $\mathbb{E}[f(S)] \geq \frac{(1-\beta)\alpha}{2} \text{OPT}$.*

Proof. We first give a lower bound on $\max(f(R_1), f(R_2))$ in the same way as in (El Halabi et al., 2023, Lemma 4.4). Let O be an optimal solution of FMSM. From Lemma 2.1 in (El Halabi et al., 2023), we know that O can be partitioned into two sets O_1, O_2 such that $O_1 \cup S_1 \in \mathcal{I}$ and $O_2 \cup S_2 \in \mathcal{I}$. Therefore $O_i \in \mathcal{I}^C \cap \mathcal{I}_i$ for $i \in \{1, 2\}$ by definition of \mathcal{I}_i and downward-closedness of \mathcal{I}^C . Since \mathcal{A} is an α -approximation algorithm, we have that

$$\begin{aligned}
 \max(f(R_1), f(R_2)) &\geq \frac{f(R_1) + f(R_2)}{2} \\
 &\geq \frac{\alpha}{2} (f(O_1) + f(O_2)) \\
 &\geq \frac{\alpha}{2} (f(O_1 \cup O_2) + f(O_1 \cap O_2)) \\
 &\geq \frac{\alpha}{2} f(O) = \frac{\alpha}{2} \text{OPT}. \tag{3}
 \end{aligned}$$

It remains to compare $f(R_i)$ with $f(S'_i)$ for $i \in \{1, 2\}$. To that end we define two non-negative submodular functions g_1, g_2 as follows:

$$g_i(W) = f(W \cup R_i) \quad \text{for all } W \subseteq S \text{ and } i \in \{1, 2\}.$$

For any $i \in \{1, 2\}$, we show that the set $S'_i \setminus R_i$ of elements added to S'_i in the for loop on Line 11 contains each element of S with probability at most β . By construction of S_1, S_2 , we have for any color c and any $e \in S \cap V_c$,

$$P(e \in (S'_1 \setminus R_1) \cap V_c) \leq P(e \in S_1 \cap V_c) = \frac{\lfloor \beta \ell_c \rfloor}{|S \cap V_c|} \leq \frac{\lfloor \beta \ell_c \rfloor}{\ell_c} \leq \beta$$

and

$$\begin{aligned} P(e \in (S'_2 \setminus R_2) \cap V_c) &\leq P(e \in S_2 \cap V_c) \\ &= P(e \notin S_1 \cap V_c) P(e \in S_2 \cap V_c \mid e \notin S_1 \cap V_c) \\ &= \left(1 - \frac{\lfloor \beta \ell_c \rfloor}{|S \cap V_c|}\right) \frac{\lfloor \beta \ell_c \rfloor}{|S \cap V_c| - \lfloor \beta \ell_c \rfloor} \\ &= \frac{\lfloor \beta \ell_c \rfloor}{|S \cap V_c|} \\ &\leq \frac{\lfloor \beta \ell_c \rfloor}{\ell_c} \leq \beta. \end{aligned}$$

We apply Lemma 2.2 to submodular functions g_1 and g_2 . We get that

$$\mathbb{E}[f(S'_i)] = \mathbb{E}[g_i(S'_i \setminus R_i)] \geq (1 - \beta)g_i(\emptyset) = (1 - \beta)f(R_i).$$

Combining that with Equation (3), we get

$$\begin{aligned} \mathbb{E}[f(S')] &= \mathbb{E}[\max(f(S'_1), f(S'_2))] \\ &\geq \max(\mathbb{E}[f(S'_1)], \mathbb{E}[f(S'_2)]) \\ &\geq (1 - \beta) \max(f(R_1), f(R_2)) \\ &\geq \frac{(1 - \beta)\alpha}{2} \text{OPT}. \end{aligned}$$

□

In the offline setting, we can use the state-of-the-art $1/(4 + \epsilon)$ -approximation algorithm of Lee et al. (2010, Theorem 4.1) as \mathcal{A} . Theorem 3.1 then follows from Lemma A.1 and Lemma A.2.

Theorem 3.1. *There exists a polynomial-time algorithm for non-monotone FMSM, which outputs a set S such that (i) $S \in \mathcal{I}$, (ii) $\lfloor \beta \ell_c \rfloor \leq |V_c \cap S| \leq u_c$ for any color $c \in [C]$, and (iii) $\mathbb{E}[f(S)] \geq (1 - \beta)\text{OPT}/(8 + \epsilon)$ for any $\beta \in [0, 1/2]$ and $\epsilon > 0$.*

In the streaming setting, we can use the state-of-the-art $1/7.464$ -approximation algorithm of Garg et al. (2021, Theorem 19) as \mathcal{A} . Theorem A.3 then also follows from Lemma A.1 and Lemma A.2.

Theorem A.3. *There exists a polynomial-time two-pass streaming algorithm for non-monotone FMSM, which uses $O(k \cdot C)$ memory, and outputs a set S such that (i) $S \in \mathcal{I}$, (ii) $\lfloor \beta \ell_c \rfloor \leq |V_c \cap S| \leq u_c$ for any color $c \in [C]$, and (iii) $\mathbb{E}[f(S)] \geq (1 - \beta)\text{OPT}/14.928$ for any $\beta \in [0, 1/2]$.*

B Missing proofs

In this section, we present the proofs that were omitted from the main text. We restate claims for convenience.

B.1 Proofs of Section 3.2

Theorem 3.2. *There is a family of FMSM instances where f is monotone and \mathcal{I} is a partition matroid, for which the integral optimum solution has value 1, but the multilinear extension admits a fractional solution of value $\Omega(\sqrt{n})$.*

Proof. We will encode bipartite perfect matching as an intersection of matroid and fairness constraints. Given a bipartite graph $(A \cup B, E)$, we set up a partition matroid to enforce that every vertex in A has at most one

adjacent edge in the solution, and fairness constraints to enforce that every vertex in B has exactly one adjacent edge in the solution (we assign color b to every edge (a, b)). We then have $k = |A| = |B| = C$.

To define an instance, we then need to describe a bipartite graph and a monotone submodular function defined on its edges. For parameters $t, s \geq 1$ we define the graph as follows: there are two special vertices 1 and 2; and for $i = 1, \dots, t$ we add a path from 1 to 2 of length $2s + 1$ to the graph (thus inserting $2s$ new vertices). Note that this graph is bipartite. Let O_i be the set of the $s + 1$ “odd” edges on the i -th path. We define the submodular objective function as $f(S) = \sum_{i=1}^t \mathbb{1}_{S \cap O_i \neq \emptyset}$.

Note that the graph has exactly t perfect matchings. Indeed, the choice of mate for vertex 1 determines the entire perfect matching, which must then be the union of O_i (for some i) and the even edges from all the other paths. In particular, for each i there is only one perfect matching that intersects O_i . Therefore the integral optimum is 1.

However, if we take x to be the average of the indicator vectors of the t matchings, then we get $x \in P_{\mathcal{F}}$ and

$$F(x) = \sum_{i=1}^t \left(1 - \prod_{e \in O_i} (1 - x_e) \right) = t \cdot \left(1 - \left(1 - \frac{1}{t} \right)^{s+1} \right).$$

If we set $s = t$, then $F(x) = \Omega(t)$, which is $\Omega(\sqrt{n})$ as $n = |E| = \Theta(s \cdot t)$. \square

B.2 Proofs of Section 3.3

We start by giving the full proof of Theorem 3.3.

Theorem 3.3. *There exists a polynomial-time algorithm for FMSM which outputs a solution $S \in \mathcal{I}$ such that $\mathbb{E}[|S \cap V_c|] \in [\ell_c, u_c]$ and $\mathbb{E}[f(S)] \geq \alpha \text{OPT}$, where $\alpha = 1 - 1/e$ if f is monotone and $\alpha = (1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} - \epsilon)/4$ otherwise, for any $\epsilon > 0$. Moreover, the solution satisfies with constant probability the following for all $c \in [C]$,*

$$\left(1 - \sqrt{\frac{3 \ln(2C)}{\ell_c}} \right) \ell_c \leq |S \cap V_c| \leq \left(1 + \sqrt{\frac{3 \ln(2C)}{u_c}} \right) u_c.$$

Proof. We first obtain an α -approximation fractional solution $x \in P_{\mathcal{F}}$ for $\max_{x \in P_{\mathcal{F}}} F(x)$. For monotone objectives, we can use the continuous greedy algorithm of Calinescu et al. (2011a, Section 3.1 and Appendix A), which achieves a $(1 - 1/e)$ -approximation for maximizing the multilinear extension over any integral polytope, both in expectation and with high probability.³ Note that $P_{\mathcal{F}}$ is integral as it is the convex hull of integral points.

For non-monotone objectives, we can use the Frank-Wolfe type algorithm of Du (2022, Sections 3.5 and 4.5) which achieves a $(1 - \min_{x \in P} \|x\|_{\infty} - \epsilon)/4$ -approximation for maximizing the multilinear extension over any polytope P , both in expectation and with high probability — see also (Mualet and Feldman, 2023, Section 3) for an explicit variant of the algorithm and its analysis. Applying this algorithm to Problem (1) then yields a $(1 - \min_{x \in P_{\mathcal{F}}} \|x\|_{\infty} - \epsilon)/4$ -approximation.

To obtain the potentially better approximation ratio $(1 - \min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} - \epsilon)/4$, we solve the complement of Problem (1): $\max_{x \in \mathbb{1} - P_{\mathcal{F}}} \bar{F}(x)$, where \bar{F} is the multilinear extension of the complement \bar{f} of f ; $\bar{f}(S) = f(V \setminus S)$. It is easy to verify that the complement of a submodular function is also submodular. So we can indeed use the Frank-Wolfe type algorithm on the complement problem, which yields $(1 - \min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} - \epsilon)/4$ approximation for it. Given an α -approximate solution \bar{x} for the complement problem, $\mathbb{1} - \bar{x}$ is an α -approximate solution to the original problem. To see this, we show that $\bar{F}(x) = F(\mathbb{1} - x)$:

$$\begin{aligned} \bar{F}(x) &= \sum_{S \subseteq V} \bar{f}(S) \prod_{i \in S} x_i \prod_{j \in V \setminus S} (1 - x_j) \\ &= \sum_{S' \subseteq V} f(S') \prod_{i \in V \setminus S'} x_i \prod_{j \in S'} (1 - x_j) \\ &= F(\mathbb{1} - x). \end{aligned}$$

Hence, $\max_{x \in P_{\mathcal{F}}} F(x) = \max_{x \in P_{\mathcal{F}}} \bar{F}(\mathbb{1} - x) = \max_{x' \in \mathbb{1} - P_{\mathcal{F}}} \bar{F}(x')$. Running the Frank-Wolfe type algorithm on both the original problem and its complement and picking the better solution then yields a $(1 - \min_{x \in P_{\mathcal{F}} \cup \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} - \epsilon)/4$ -approximation to Problem (1).

Both the continuous greedy and the Frank-Wolfe type algorithms run in polynomial time if the constraint polytope is solvable, which is the case for $P_{\mathcal{F}}$ and $\mathbb{1} - P_{\mathcal{F}}$ by Lemma 2.1-c.

Next we round the fractional solution x to an independent (but not necessarily feasible) set $S \in \mathcal{I}$ using the randomized swap rounding scheme of Chekuri et al. (2010, Section IV). The rounded solution is guaranteed to satisfy (see Theorem II.1 therein) $\mathbb{E}[|S \cap V_c|] = x(V_c) \in [\ell_c, u_c]$, $\mathbb{E}[f(S)] \geq F(x) \geq \alpha \text{OPT}$, and for any $c \in [C]$, $\delta_1 \geq 0$, and $\delta_2 \in [0, 1]$,

$$\begin{aligned} \mathbb{P}[|S \cap V_c| \geq (1 + \delta_1)u_c] &\leq \left(\frac{e^{\delta_1}}{(1 + \delta_1)^{(1 + \delta_1)}} \right)^{u_c} \\ &\leq e^{-0.38u_c\delta_1^2} \end{aligned} \quad \text{by (Celis et al., 2018a, Lemma 15)}$$

$$\text{and } \mathbb{P}[|S \cap V_c| \leq (1 - \delta_2)\ell_c] \leq e^{-\ell_c\delta_2^2/2}.$$

Then, by choosing $\delta_1 = \sqrt{3 \ln(2C)/u_c}$ and $\delta_2 = \sqrt{3 \ln(2C)/\ell_c}$, we get by union bound that S satisfies $|S \cap V_c| \in [(1 - \sqrt{3 \ln(2C)/\ell_c})\ell_c, (1 + \sqrt{3 \ln(2C)/u_c})u_c]$ for all $c \in [C]$, with probability at least $1 - 1/(2C)^{1.14} - 1/(2C)^{1.5} \geq 1 - 1/2^{1.14} - 1/2^{1.5} > 0.19$.

Remark B.1. If f is monotone, we can further guarantee that for any $\delta > 0$, the solution in Theorem 3.3 satisfies with constant probability the following for all $c \in [C]$,

$$\left(1 - \sqrt{\frac{3 \ln(2C/(1 - e^{-\delta^2/16}))}{\ell_c}} \right) \ell_c \leq |S \cap V_c| \leq \left(1 + \sqrt{\frac{3 \ln(2C/(1 - e^{-\delta^2/16}))}{u_c}} \right) u_c$$

and $f(S) \geq (1 - 1/e - \delta) \text{OPT}$.

Proof. If f is monotone, we can apply the lower-tail concentration bound in (Chekuri et al., 2010, Theorem II.2). To that end, we assume without loss of generality that for any $i \in V$, there exists a feasible set $S' \in \mathcal{F}$ containing i , otherwise we can simply ignore such element. Note that we can easily check if an element i satisfies this assumption, by maximizing a linear function over \mathcal{F} , which assigns zero weights to all elements except i and a non-zero weight to i . This can be done efficiently by Lemma 2.1-c. Under this assumption, the marginal values of f are in $[0, \text{OPT}]$, since for any $i \in V$, we have $f(i | S) \leq f(i) \leq f(S') \leq \text{OPT}$ by submodularity and monotonicity. Scaling f by $1/\text{OPT}$, we obtain

$$\begin{aligned} \mathbb{P}[f(S) \leq (1 - \delta)F(x)] &\leq e^{-F(x)\delta^2/(8\text{OPT})} \\ &\leq e^{-\delta^2/16}, \end{aligned}$$

for any $\delta > 0$, since $F(x) \geq (1 - 1/e)\text{OPT}$ with high probability. Hence, S is guaranteed to satisfy $f(S) \geq (1 - \delta)F(x) \geq (1 - 1/e - \delta)\text{OPT}$ with probability $1 - e^{-\delta^2/16}$.

Then we again use the concentration bounds in (Chekuri et al., 2010, Theorem II.1) but now choosing $\delta_1 = \sqrt{\frac{3 \ln(2C/(1 - e^{-\delta^2/16}))}{u_c}}$ and $\delta_2 = \sqrt{\frac{3 \ln(2C/(1 - e^{-\delta^2/16}))}{\ell_c}}$. We get by union bound that S satisfies all the bounds in the claim with probability at least $(1 - e^{-\delta^2/16})(1 - 1/(2C)^{1.14} - 1/(2C)^{1.5}) > 0.19(1 - e^{-\delta^2/16})$. \square

\square

Next, we prove the statements of Example 3.4 and Example 3.5.

Example 3.4 (Uniform matroid). Let \mathcal{I} be the uniform matroid with rank k . We order the color groups such that $\frac{\ell_1}{|V_1|} \leq \dots \leq \frac{\ell_C}{|V_C|}$ and let t be the largest index $t \in [C]$ such that $\frac{\ell_t}{|V_t|} \sum_{c=1}^t |V_c| + \sum_{c=t+1}^C \ell_c \leq k$ ($t \geq 1$ is well-defined since $\mathcal{F} \neq \emptyset$). Then $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = \min\{\max_c \frac{\ell_c}{|V_c|}, 1 - \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}\}$, where

$$\tau = \frac{k - \sum_{c=t+1}^C \ell_c}{\sum_{c=1}^t |V_c|}.$$

Proof. First we compute $\min_{x \in P_{\mathcal{F}}} \|x\|_{\infty}$. For any $x \in P_{\mathcal{F}}$, we have $\ell_c \leq \sum_{i \in V_c} x_i \leq \|x\|_{\infty} |V_c|$ for all $c \in [C]$, hence $\|x\|_{\infty} \geq \max_c \frac{\ell_c}{|V_c|}$. We show that this lower bound is achieved at a feasible vector $x \in P_{\mathcal{F}}$. Let $x \in [0, 1]^n$ such that $x_i = \frac{\ell_c}{|V_c|}$ for all $i \in V_c$ and $c \in [C]$, then $\|x\|_{\infty} = \max_c \frac{\ell_c}{|V_c|}$ and $x \in P_{\mathcal{F}}$ since $\sum_{i \in V_c} x_i = \ell_c \in [\ell_c, u_c]$ and $\sum_{i \in V} x_i = \sum_c \ell_c \leq k$ by the assumption that $\mathcal{F} \neq \emptyset$. Hence $\min_{x \in P_{\mathcal{F}}} \|x\|_{\infty} = \max_c \frac{\ell_c}{|V_c|}$.

Next we compute $\min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} = \min_{x \in P_{\mathcal{F}}} \|\mathbb{1} - x\|_{\infty} = 1 - \max_{x \in P_{\mathcal{F}}} \min_{i \in V} x_i$. Let $x \in [0, 1]^n$ such that

$$x_i = \begin{cases} \frac{\ell_c}{|V_c|} & \text{for all } i \in V_c, c > t \\ \min\{\tau, \frac{u_c}{|V_c|}\} & \text{for all } i \in V_c, c \leq t. \end{cases}$$

We argue that $x \in P_{\mathcal{F}}$. We have $x(V_c) = \ell_c \in [\ell_c, u_c]$ for all $c > t$ and $x(V_c) = \min\{u_c, \tau|V_c|\}$ for all $c \leq t$. By definition of t and τ , we have $\tau \geq \frac{\ell_t}{|V_t|}$. Hence $x(V_c) \geq \min\{u_c, \frac{\ell_t}{|V_t|}|V_c|\} \geq \ell_c$ for all $c \leq t$. By definition of τ , we also have $x(V) \leq \tau \sum_{c=1}^t |V_c| + \sum_{c=t+1}^C \ell_c = k$. We show next that $\min_{i \in V} x_i = \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}$. This holds trivially if $t = C$. Otherwise, by definition of $t < C$, we have

$$k \leq \frac{\ell_{t+1}}{|V_{t+1}|} \sum_{c=1}^{t+1} |V_c| + \sum_{c=t+2}^C \ell_c = \frac{\ell_{t+1}}{|V_{t+1}|} \sum_{c=1}^t |V_c| + \sum_{c=t+1}^C \ell_c,$$

hence $\tau \leq \frac{\ell_{t+1}}{|V_{t+1}|} = \min_{i \in \cup_{c>t} V_c} x_i$. Thus $\min_{i \in V} x_i$ is attained at some color $c \leq t$. This proves that $\max_{x \in P_{\mathcal{F}}} \min_{i \in V} x_i \geq \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}$. To prove the upper bound, we observe that for any $x' \in P_{\mathcal{F}}$, we have $\min_{i \in V} x'_i |V_c| \leq x'(V_c) \leq u_c$ for every $c \in [C]$ and $k \geq x'(V) \geq \sum_{c=t+1}^C \ell_c + \min_{i \in V} x'_i \sum_{c=1}^t |V_c|$. Thus by definition of τ , we get

$$\min_{i \in V} x'_i \leq \min\{\tau, \min_{c \in [C]} \frac{u_c}{|V_c|}\} \leq \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}.$$

Hence, $\min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty} = 1 - \min\{\tau, \min_{c \leq t} \frac{u_c}{|V_c|}\}$ achieved at $\bar{x} = \mathbb{1} - x$. \square

Example 3.5 (Bipartite perfect matching). Let \mathcal{F} be the feasible set arising from the bipartite perfect matching problem in the proof of the integrality gap Theorem 3.2. We have $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = 1 - \Theta(1/\sqrt{n})$.

Proof. We prove that $\min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty} = 1 - 1/t$ (recall that $t = \Theta(\sqrt{n})$ in Theorem 3.2). The point x used in the proof of Theorem 3.2 puts $1/t$ value on the “odd” edges, and $1 - 1/t$ value on the “even” edges (each of which belongs to all perfect matchings but one). Thus it has $\|x\|_{\infty} = \|\mathbb{1} - x\|_{\infty} = 1 - 1/t$. This proves the \leq direction. For the \geq direction, consider any $x \in P_{\mathcal{F}}$. The point x must satisfy $\sum_{e \in \delta(1)} x_e = 1$, where $\delta(1)$ is the set of edges incident on the special vertex 1. As $|\delta(1)| = t$, we must have $x_e \leq 1/t$ for some $e \in \delta(1)$. Then consider the path from vertex 1 to 2 containing e , and let e' be the next edge after e on that path. The common endpoint of e and e' has degree two, thus $x_{e'} = 1 - x_e \geq 1 - 1/t$. Hence $\min\{\|x\|_{\infty}, \|\mathbb{1} - x\|_{\infty}\} \geq 1 - 1/t$. \square

B.3 Proofs of Section 3.4

In this section, we expand on some details that were omitted from the proof of Theorem 3.6.

Theorem 3.6. *There exists a polynomial-time algorithm for non-monotone FMSM where \mathcal{I} is a uniform matroid, which achieves $0.401(1 - \min_{x \in P_{\mathcal{F}} \cup (\mathbb{1} - P_{\mathcal{F}})} \|x\|_{\infty})$ -approximation in expectation.*

Proof. El Halabi et al. (2020, Theorem 5.2) presented a polynomial-time algorithm with an expected $(1 - \max_c \frac{\ell_c}{|V_c|})\alpha$ -approximation for non-monotone FMSM over a uniform matroid, given any α -approximation algorithm for non-monotone submodular maximization over a matroid constraint.⁵ Plugging in the state-of-the-art 0.401-approximation of Buchbinder and Feldman (2019, Theorem 1.1) then yields a $0.401(1 - \max_c \frac{\ell_c}{|V_c|})$ -approximation in expectation. Recall from Example 3.4 that $\min_{x \in P_{\mathcal{F}}} \|x\|_{\infty} = \max_c \frac{\ell_c}{|V_c|}$.

To obtain the potentially better approximation factor $0.401(1 - \min_{x \in \mathbb{1} - P_{\mathcal{F}}} \|x\|_{\infty})$, we solve (as in Theorem 3.3) the complement problem $\max_{V \setminus S \in \mathcal{F}} \bar{f}(S)$, where $\bar{f}(S) = f(V \setminus S)$. We follow a similar strategy as El Halabi et al. (2020). Namely, we first drop the lower bounds from the constraint. The resulting problem $\max_{S \subseteq V} \{\bar{f}(S) : |S \cap V_c| \leq |V_c| - \ell_c, \forall c \in [C]\}$ is a non-monotone submodular maximization problem over a partition matroid. Hence, a solution S with an expected 0.401-approximation can be obtained for it.

Next, we augment the solution S to a feasible one that we will denote S^+ . Let \bar{x} be a solution of $\min_{x \in \mathbf{1} - P_{\mathcal{F}}} \|x\|_{\infty}$, which has a closed form as shown in Example 3.4. For each color c , we sample a set $B_c \subseteq V_c$ of elements so that:

1. $\mathbb{P}[e \in B_c] = \bar{x}_e$ for all $e \in V_c$,
2. $|B_c|$ is $\lfloor \bar{x}(V_c) \rfloor$ or $\lceil \bar{x}(V_c) \rceil$ for all $c \in [C]$,
3. $\sum_c |B_c|$ is $\lfloor \bar{x}(V) \rfloor$ or $\lceil \bar{x}(V) \rceil$.

This can be done as follows, inspired by (Hartley and Rao, 1962, Section 2.1). Assign consecutive disjoint intervals on the real line to elements, of length \bar{x}_e for each $e \in V$, grouped by colors; e.g., if $V_1 = \{e_1, e_2, \dots\}$, the first two intervals would be $[0, \bar{x}_{e_1})$ and $[\bar{x}_{e_1}, \bar{x}_{e_1} + \bar{x}_{e_2})$. Now sample a random offset $\alpha \in [0, 1)$, and define B_c to be those elements of V_c whose interval contains a point in $\mathbb{Z} + \alpha$. Property 1 follows directly. For property 2, note that the union of intervals of all $e \in V_c$ is also an interval (as elements of the same color are grouped together), of length $\bar{x}(V_c)$; its intersection with $\mathbb{Z} + \alpha$ must thus be of size either $\lfloor \bar{x}(V_c) \rfloor$ or $\lceil \bar{x}(V_c) \rceil$, and every point in this intersection adds exactly one new element to B_c since $\bar{x}_e \leq 1$ for all e . Property 3 follows likewise.

Then, for every color such that $|S \cap V_c| < |B_c|$, add any $|B_c| - |S \cap V_c|$ elements from $B_c \setminus S$ to S . This obtains S^+ . We argue that the augmented set S^+ is feasible (for the complement problem $\max_{V \setminus S \in \mathcal{F}} \bar{f}(S)$). Note that $|B_c| \in \{\lfloor \bar{x}(V_c) \rfloor, \lceil \bar{x}(V_c) \rceil\} \subseteq [|V_c| - u_c, |V_c| - \ell_c]$, since $\bar{x} \in \mathbf{1} - P_{\mathcal{F}}$ and the bounds are integral. For the lower bounds, we have $|S^+ \cap V_c| \geq |B_c| \geq |V_c| - u_c$. For the upper bounds, note that $|S^+ \cap V_c| = \max\{|S \cap V_c|, |B_c|\} \leq |V_c| - \ell_c$, since by definition of S , $|S \cap V_c| \leq |V_c| - \ell_c$. Moreover, $|S^+| \geq \sum_c |B_c| \geq \lfloor \bar{x}(V) \rfloor \geq n - k$, since $\bar{x} \in \mathbf{1} - P_{\mathcal{F}}$ and $n - k$ is an integer.

To bound the loss in value resulting from the additional elements, we make use of Lemma 2.2. Note that for any $e \in V_c$, $\mathbb{P}[e \in S^+ \setminus S] \leq \mathbb{P}[e \in B_c] = \bar{x}_e \leq \|\bar{x}\|_{\infty}$. Lemma 2.2 and the definitions of S and \bar{x} then imply that

$$\mathbb{E}[\bar{f}(S^+)] \geq (1 - \|\bar{x}\|_{\infty}) \mathbb{E}[\bar{f}(S)] \geq 0.401(1 - \min_{x \in \mathbf{1} - P_{\mathcal{F}}} \|x\|_{\infty}) \max_{V \setminus S \in \mathcal{F}} \bar{f}(S).$$

Finally, taking the complement $V \setminus S^+$ yields a $0.401(1 - \min_{x \in \mathbf{1} - P_{\mathcal{F}}} \|x\|_{\infty})$ -approximation to the original problem, in expectation. \square

B.4 Proofs of Section 3.5

In this section, we give the full proof of Theorem 3.7.

Theorem 3.7. *For any r in the form $r = 1 - \frac{1}{t}$ with $t \geq 2$ and any $\epsilon > 0$, a $(1 - r + \epsilon)$ -approximation for FMSM that works for instances where $\min_{x \in P_{\mathcal{F}} \cup (\mathbf{1} - P_{\mathcal{F}})} \|x\|_{\infty} \leq r$ requires an exponential number of value queries. This is also the case for the continuous relaxation of FMSM, $\max_{x \in P_{\mathcal{F}}} F(x)$. Furthermore, this is true even in the two special cases where (i) $C = 2$ and $\mathcal{I} = 2^V$ (no matroid constraint), (ii) $C = 1$ and \mathcal{I} is a partition matroid.*

Proof. Vondrák (2013, Theorem 1.2, 1.9 and Section 2) showed that any algorithm achieving a better than $(1 - r)$ -approximation for the problem of maximizing a non-monotone submodular function over the *bases* of a matroid \mathcal{I} , where $\min_{x \in P_{\mathcal{I}}: \sum_i x_i = k} \|x\|_{\infty} \leq r$, requires exponentially many value queries. This holds also for the continuous relaxation of the problem (see Theorem 1.9 and Section 2 therein).

We show that the hard instance used in the proofs of the lower bounds in (Vondrák, 2013) is a special case of FMSM with $\min_{x \in P_{\mathcal{F}} \cup (\mathbf{1} - P_{\mathcal{F}})} \|x\|_{\infty} \leq r$. In particular, the hard instance uses a ground set $V = A \cup B$, where A and B are two disjoint sets, each of size $t \times m$, for some large number m , and a partition matroid base constraint $\mathcal{B} = \{S : |S \cap A| = m \text{ and } |S \cap B| = (t - 1)m\}$ (see Section 2, Theorems 1.8 and 1.9 therein).

We can express this matroid base constraint via \mathcal{F} in two ways. First, any matroid base constraint is a special case of \mathcal{F} where $C = 1$, $V = V_1$, and $\ell_1 = u_1 = k$ (the rank of the matroid), i.e., the fairness constraint reduces to $|S| = k$. In the case of \mathcal{B} , we have $k = t \times m$ and $\min_{x \in P_{\mathcal{F}} \cup (\mathbf{1} - P_{\mathcal{F}})} \|x\|_{\infty} = \min_{x \in P_{\mathcal{I}}: \sum_i x_i = k} \|x\|_{\infty} = 1 - \frac{1}{t}$. Alternatively, we can express \mathcal{B} with a fairness constraint alone, with two color groups $V_1 = A$, $V_2 = B$, $\ell_1 = u_1 = m$, and $\ell_2 = u_2 = (t - 1)m$. \square

B.5 Proofs of Section 4

In this section, we prove Lemma 4.2 and Theorem 4.3.

Lemma 4.2. *Let $\tilde{\mathcal{C}}_k := \{S \subseteq V \mid \text{there exists } S' \in \mathcal{C}, |S'| \leq k \text{ such that } S \subseteq S'\}$. Then any fair set $S \in \mathcal{C}$ is a base of the matroid $\tilde{\mathcal{C}}_k$ with $k = |S|$ and vice versa.*

Proof. Note that any fair set is a maximal set of the k -truncation of \mathcal{C} given by $\mathcal{C}_k := \{S \in \mathcal{C} \mid |S| \leq k\}$ with $k = |S|$. The family \mathcal{C}_k is a special case of \mathcal{F} where \mathcal{I} is the k -uniform matroid, so the corresponding family of extendable set $\tilde{\mathcal{C}}_k$ is a matroid (El Halabi et al., 2020, Lemma 4.1). The claim then follows from Fact 4.1. \square

Theorem 4.3. *There exists a polynomial-time rounding scheme which, given $x \in P_{\mathcal{F}}$, returns a feasible set $S \in \mathcal{F}$ that for any decomposable submodular function f over \mathcal{F} satisfies $\mathbb{E}[f(S)] \geq F(x)$.*

Proof. The randomized rounding scheme of Chekuri et al. (2010, Section V) takes as input a fractional solution $x' \in P_{\mathcal{I}_1} \cap P_{\mathcal{I}_2}$, where \mathcal{I}_1 and \mathcal{I}_2 are two matroids, expressed as a convex combination of sets in $\mathcal{I}_1 \cap \mathcal{I}_2$, and returns a set $S' \in \mathcal{I}_1 \cap \mathcal{I}_2$ such that $\mathbb{E}[f'(S')] \geq F'(x')$ for any submodular function f' which decomposes over the equivalence classes of \mathcal{I}_1 or \mathcal{I}_2 , i.e., $f'(S) = \sum_{G \in \mathcal{G}_1} f'_{1,G}(S \cap G) + \sum_{G \in \mathcal{G}_2} f'_{2,G}(S \cap G)$, where $f'_{1,G}, f'_{2,G}$ are submodular functions and $\mathcal{G}_1, \mathcal{G}_2$ are the respective families of equivalence classes of \mathcal{I}_1 and \mathcal{I}_2 (see Theorem II.3 therein). If the sets in the input convex combination are all of the same size k , then the returned set will have size k too⁷.

Given $x \in P_{\mathcal{F}}$, since $P_{\mathcal{F}}$ is solvable (Lemma 2.1-c), we can write x as a convex combination of sets in \mathcal{F} in strongly polynomial time (Schrijver, 2003, Theorem 5.15); i.e., $x = \sum_t \alpha_t \mathbb{1}_{I_t}$ for some $I_t \in \mathcal{F}, \alpha_t \geq 0, \sum_t \alpha_t = 1$.

Let k be the size of the largest set I_t . Lemma 4.2 implies that every set I_t is in the intersection of the matroids \mathcal{I} and $\tilde{\mathcal{C}}_k$, where $\tilde{\mathcal{C}}_k$ is defined as in the lemma. Then x is already a valid input for swap rounding. If all sets I_t have the same size, then the returned rounded solution $S \in \mathcal{I} \cap \tilde{\mathcal{C}}_k$ is a base of $\tilde{\mathcal{C}}_k$, and thus also a fair set by Lemma 4.2. So S is feasible. Otherwise, S is not necessarily feasible. In that case, we first need to add dummy elements to make all sets I_t of equal size.

Let E be a set of k new dummy elements (not in V). For every t , we add enough dummy elements from E to I_t to obtain a set I'_t of size k . We also extend the matroid \mathcal{I} to $\mathcal{I}^+ = \{S \subseteq V \cup E \mid S \setminus E \in \mathcal{I}\}$ and the collection of fair sets \mathcal{C} to $\mathcal{C}^+ = \{S \subseteq V \cup E \mid S \setminus E \in \mathcal{C}\}$. It is easy to verify that \mathcal{I}^+ is also a matroid. Observe that \mathcal{C}^+ is also a collection of fair sets over the extended ground set, since $\mathcal{C}^+ = \{S \subseteq V \cup E \mid |S \cap V_c| \in [\ell_c, u_c] \forall c \in [C], |S \cap E| \in [0, |E|]\}$, so Lemma 4.2 applies to it. Moreover, $\tilde{\mathcal{C}}_k^+ = \{S \subseteq V \cup E \mid \text{there exists } S' \in \mathcal{C}^+, |S'| \leq k \text{ such that } S \subseteq S'\}$ is a matroid of rank k , since any set in $\tilde{\mathcal{C}}_k^+$ has size at most k and there exists $I_t \in \tilde{\mathcal{C}}_k^+$ with size k for some t .

Let $x' = \sum_t \alpha_t \mathbb{1}_{I'_t}$. Since every set I'_t is in the intersection of the matroids \mathcal{I}^+ and $\tilde{\mathcal{C}}_k^+$, x' is a valid input to swap rounding. And since all I'_t have size k , the returned rounded solution $S' \in \mathcal{I}^+ \cap \tilde{\mathcal{C}}_k^+$ will also have size k . Hence, S' is a base of $\tilde{\mathcal{C}}_k^+$, and thus $S' \in \mathcal{C}^+$ by Lemma 4.2. Removing the dummy elements from S' then yields a feasible solution $S = S' \setminus E \in \mathcal{I} \cap \mathcal{C}$.

It remains to show that the solution S preserves the value of $F(x)$ for any decomposable submodular function f over \mathcal{F} . Define $f' : 2^{V \cup E} \rightarrow \mathbb{R}_+$ as $f'(S) = f(S \setminus E)$, then its multilinear extension $F' : [0, 1]^{n+k} \rightarrow \mathbb{R}_+$ is given by $F'(x') = F(x'_V)$ where $x'_V \in [0, 1]^n$ is the vector corresponding to the entries of x' in V . We observe that f' decomposes over the equivalence classes of \mathcal{I}^+ or $\tilde{\mathcal{C}}_k^+$, with $f'_{1,G}(S) = f_{1,G}(S \setminus E)$ and $f'_{2,c}(S) = f_{2,c}(S \setminus E)$, where $f_{1,G}, f_{2,c}$ are the functions in the decomposition of f (Definition 2.4). One can verify that $f', f'_{1,G}, f'_{2,c}$ are submodular since $f, f_{1,G}, f_{2,c}$ are submodular. The solution S returned by swap rounding then satisfies $\mathbb{E}[f'(S')] = \mathbb{E}[f(S)] \geq F'(x') = F(x)$. \square

It is worth mentioning that El Halabi et al. (2023, Appendix C.2) have already shown a reduction of FMSM to submodular maximization over the intersection of two matroid bases. However, their result cannot be used to prove Theorem 4.3, since the matroids in their reduction have different equivalence classes than the ones of \mathcal{I} and the color groups. Using their reduction in Theorem 4.3 would require considering a much more restricted class of

⁷This is not explicitly stated in (Chekuri et al., 2010, Section V), but can be easily deduced from the description of the merging operation therein.

decomposable submodular functions (ones which are completely separable over V , or over subgroups of colors obtained by splitting each color group into two).

C Relation to multi-objective submodular maximization

In this section, we discuss the connection of FMSM to multi-objective submodular maximization, where given m submodular functions f_i , the goal is to maximize their minimum $\min_{i \in [m]} f_i(S)$.

We can write FMSM as a multi-objective submodular maximization problem over two matroids:

$$\max_{S \subseteq V} \left\{ \min_{i \in [C+1]} f_i(S) : S \in \mathcal{I}, |S \cap V_c| \leq u_c \quad \forall c \in [C] \right\}, \quad (4)$$

where $f_{C+1} = f / \text{OPT}$ and $f_c = |S \cap V_c| / \ell_c$ for all $c \in [C]$. Recall that OPT is the optimal value of FMSM, which can be guessed. A solution S of Problem (4) is then also an optimal solution to FMSM. To the best of our knowledge, there are no existing work on multi-objective submodular maximization over two matroids.

In the special cases where there are no upper bounds ($u_c = |V_c|$) or \mathcal{I} is the uniform matroid, and f is monotone, one can apply the result of (Chekuri et al., 2009, Theorem 7.2) on multi-objective submodular maximization over a single matroid to obtain a $(1 - 1/e)$ -approximation for monotone FMSM. But the resulting runtime will be exponential in $\sum_c \ell_c$ as the algorithm requires guessing the ℓ_c elements belonging to the optimal solution for each group c , by brute force enumeration. Recall though that when \mathcal{I} is the uniform matroid, a polynomial time $(1 - 1/e)$ -approximation for monotone FMSM was provided in (Celis et al., 2018a, Theorem 18) under the assumption that the groups are disjoint, which is not assumed in (Chekuri et al., 2010). Celis et al. (2018a) also provided a $(1 - 1/e)$ -approximation algorithm for the overlapping groups case, but its runtime is also exponential in C (see Theorem 20 therein).

Finally, if there are no upper bounds ($u_c = |V_c|$), \mathcal{I} is the uniform matroid, and f is monotone, then the constraint in Problem (4) becomes a simple cardinality constraint. In this case, a $(1 - 1/e - o(1))$ -approximation for monotone FMSM can be obtained in polynomial-time if $k = o(n)$, even with overlapping groups (Udwani, 2018; Tsang et al., 2019).

The line of work (Wang et al., 2024; Tsang et al., 2019; Tang and Yuan, 2023) which studied submodular maximization under a different notion of fairness can all be formulated as a monotone multi-objective submodular maximization problem over a cardinality constraint.