

Symmetry and Approximability of Submodular Maximization Problems*

Jan Vondrák †

January 29, 2013

Abstract

A number of recent results on optimization problems involving submodular functions have made use of the *multilinear relaxation* of the problem. These results hold typically in the *value oracle* model, where the objective function is accessible via a black box returning $f(S)$ for a given S . We present a general approach to deriving inapproximability results in the value oracle model, based on the notion of *symmetry gap*. Our main result is that for any fixed instance that exhibits a certain symmetry gap in its multilinear relaxation, there is a naturally related class of instances for which a better approximation factor than the symmetry gap would require exponentially many oracle queries. This unifies several known hardness results for submodular maximization, e.g. the optimality of $(1 - 1/e)$ -approximation for monotone submodular maximization under a cardinality constraint, and the impossibility of $(\frac{1}{2} + \epsilon)$ -approximation for unconstrained (nonmonotone) submodular maximization.

As a new application, we consider the problem of maximizing a nonmonotone submodular function over the bases of a matroid. A $(\frac{1}{6} - o(1))$ -approximation has been developed for this problem, assuming that the matroid contains two disjoint bases. We show that the best approximation one can achieve is indeed related to packings of bases in the matroid. Specifically, for any $k \geq 2$, there is a class of matroids of fractional base packing number $\nu = \frac{k}{k-1}$, such that any algorithm achieving a better than $(1 - \frac{1}{\nu}) = \frac{1}{k}$ -approximation for this class would require exponentially many value queries. In particular, there is no constant-factor approximation for maximizing a nonmonotone submodular function over the bases of a general matroid. On the positive side, we present a $\frac{1}{2}(1 - \frac{1}{\nu} - o(1))$ -approximation algorithm assuming fractional base packing number at least ν where $\nu \in (1, 2]$. We also present an improved 0.309-approximation for maximization of a nonmonotone submodular function subject to a matroid independence constraint, improving the previously known factor of $\frac{1}{4} - \epsilon$. For this problem, we obtain a hardness of $(\frac{1}{2} + \epsilon)$ -approximation for any fixed $\epsilon > 0$.

1 Introduction

Submodular set functions are defined by the following condition for all pairs of sets S, T :

$$f(S \cup T) + f(S \cap T) \leq f(S) + f(T),$$

or equivalently by the property that the *marginal value* of any element, $f_S(j) = f(S + j) - f(S)$, satisfies $f_T(j) \leq f_S(j)$, whenever $j \notin T \supset S$. In addition, a set function is called monotone if $f(S) \leq f(T)$ whenever $S \subseteq T$. Throughout this paper, we assume that $f(S)$ is nonnegative.

Submodular functions have been studied in the context of combinatorial optimization since the 1970's, especially in connection with matroids [10, 11, 23, 24, 25, 33, 34, 21, 15]. Submodular functions appear mostly for the following two reasons: (i) submodularity arises naturally in various combinatorial settings, and many algorithmic applications use it either explicitly or implicitly; (ii) submodularity has a natural interpretation as the property of *diminishing returns*, which defines an important class of utility/valuation functions. Submodularity as an abstract concept is both general enough to be useful for applications and

*An extended abstract [32] appeared in IEEE FOCS 2009. This work was done while the author was at Princeton University.

†IBM Almaden Research Center, San Jose, CA 95120 (jvondrak@us.ibm.com).

it carries enough structure to allow strong positive results. A fundamental algorithmic result is that any submodular function can be *minimized* in strongly polynomial time [14, 27].

In contrast to submodular minimization, submodular maximization problems are typically hard to solve exactly. Consider the classical problem of maximizing a monotone submodular function subject to a cardinality constraint, $\max\{f(S) : |S| \leq k\}$. It is known that this problem admits a $(1 - 1/e)$ -approximation [23] and this is optimal for two different reasons: (i) Given only black-box access to $f(S)$, we cannot achieve a better approximation, unless we ask exponentially many value queries [25]. This holds even if we allow unlimited computational power. (ii) In certain special cases where $f(S)$ has a compact representation on the input, it is NP-hard to achieve an approximation better than $1 - 1/e$ [12]. The reason why the hardness threshold is the same in both cases is apparently not well understood.

An optimal $(1 - 1/e)$ -approximation for the problem $\max\{f(S) : |S| \leq k\}$ where f is monotone submodular is achieved by a simple greedy algorithm [23]. This seems to be rather coincidental; for other variants of submodular maximization, such as unconstrained (nonmonotone) submodular maximization [13], monotone submodular maximization subject to a matroid constraint [24, 4, 30], or submodular maximization subject to linear constraints [18, 19], greedy algorithms achieve suboptimal results. A tool which has proven useful in approaching these problems is the *multilinear relaxation*.

Multilinear relaxation. Let us consider a discrete optimization problem $\max\{f(S) : S \in \mathcal{F}\}$, where $f : 2^X \rightarrow R$ is the objective function and $\mathcal{F} \subset 2^X$ is the collection of feasible solutions. In case f is a linear function, $f(S) = \sum_{j \in S} w_j$, it is natural to replace this problem by a linear programming problem. For a general set function $f(S)$, however, a linear relaxation is not readily available. Instead, the following relaxation has been proposed [4, 30, 5]: For $\mathbf{x} \in [0, 1]^X$, let $\hat{\mathbf{x}}$ denote a random vector in $\{0, 1\}^X$ where each coordinate of x_i is rounded independently to 1 with probability x_i and 0 otherwise.¹ We define

$$F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})] = \sum_{S \subseteq X} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j).$$

This is the unique *multilinear polynomial* which coincides with f on $\{0, 1\}$ -vectors. We remark that although we cannot compute the exact value of $F(\mathbf{x})$ for a given $\mathbf{x} \in [0, 1]^X$ (which would require querying 2^n possible values of $f(S)$), we can compute $F(\mathbf{x})$ approximately, by random sampling. Sometimes this causes technical issues, which we also deal with in this paper.

Instead of the discrete problem $\max\{f(S) : S \in \mathcal{F}\}$, we consider a continuous optimization problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$, where $P(\mathcal{F})$ is the convex hull of characteristic vectors corresponding to \mathcal{F} ,

$$P(\mathcal{F}) = \left\{ \sum_{S \in \mathcal{F}} \alpha_S \mathbf{1}_S : \sum_{S \in \mathcal{F}} \alpha_S = 1, \alpha_S \geq 0 \right\}.$$

The reason why the extension $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ is useful for submodular maximization problems is that $F(\mathbf{x})$ has convexity properties that allow one to solve the continuous problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ (within a constant factor) in a number of interesting cases. Moreover, fractional solutions can be often rounded to discrete ones without losing *anything* in terms of the objective value. Then, our ability to solve the multilinear relaxation translates directly into an algorithm for the original problem. In particular, this is true when the collection of feasible solutions forms a matroid.

Pipage rounding was originally developed by Ageev and Sviridenko for rounding solutions in the bipartite matching polytope [2]. The technique was adapted to matroid polytopes by Calinescu et al. [4], who proved that for any submodular function $f(S)$ and any \mathbf{x} in the matroid base polytope $B(\mathcal{M})$, the fractional solution \mathbf{x} can be rounded to a base $B \in \mathcal{B}$ such that $f(B) \geq F(\mathbf{x})$. This approach leads to an optimal $(1 - 1/e)$ -approximation for the Submodular Welfare Problem, and more generally for monotone submodular maximization subject to a matroid constraint [4, 30]. It is also known that the factor of $1 - 1/e$ is optimal for the Submodular Welfare Problem both in the NP framework [17] and in the value oracle model [22]. Under the assumption that the submodular function $f(S)$ has *curvature* c , there is a $\frac{1}{c}(1 - e^{-c})$ -approximation and this is also optimal in the value oracle model [31]. The framework of pipage rounding can be also extended to nonmonotone submodular functions; this presents some additional issues which we discuss in this paper.

¹We denote vectors consistently in boldface (\mathbf{x}) and their coordinates in italics (x_i). We also identify vectors in $\{0, 1\}^n$ with subsets of $[n]$ in a natural way.

For the problem of unconstrained (nonmonotone) submodular maximization, a $2/5$ -approximation was developed in [13]. This algorithm implicitly uses the multilinear relaxation $\max\{F(\mathbf{x}) : \mathbf{x} \in [0, 1]^X\}$. For *symmetric* submodular functions, it is shown in [13] that a uniformly random solution $\mathbf{x} = (1/2, \dots, 1/2)$ gives $F(\mathbf{x}) \geq \frac{1}{2}OPT$, and there is no better approximation algorithm in the value oracle model. Recently, a $1/2$ -approximation was found for unconstrained maximization of a general nonnegative submodular function [3]. This algorithm can be formulated in the multilinear relaxation framework, but also as a randomized combinatorial algorithm.

Using additional techniques, the multilinear relaxation can be also applied to submodular maximization with knapsack constraints ($\sum_{j \in S} c_{ij} \leq 1$). For the problem of maximizing a monotone submodular function subject to a constant number of knapsack constraints, there is a $(1 - 1/e - \epsilon)$ -approximation algorithm for any $\epsilon > 0$ [18]. For maximizing a nonmonotone submodular function subject to a constant number of knapsack constraints, a $(1/5 - \epsilon)$ -approximation was designed in [19].

One should mention that not all the best known results for submodular maximization have been achieved using the multilinear relaxation. The greedy algorithm yields a $1/(k + 1)$ -approximation for monotone submodular maximization subject to k matroid constraints [24]. Local search methods have been used to improve this to a $1/(k + \epsilon)$ -approximation, and to obtain a $1/(k + 1 + 1/(k - 1) + \epsilon)$ -approximation for the same problem with a nonmonotone submodular function, for any $\epsilon > 0$ [19, 20]. For the problem of maximizing a nonmonotone submodular function over the *bases* of a given matroid, local search yields a $(1/6 - \epsilon)$ -approximation, assuming that the matroid contains two disjoint bases [19].

1.1 Our results

Our main contribution (Theorem 1.8) is a general hardness construction that yields inapproximability results in the value oracle model in an automated way, based on what we call the *symmetry gap* for some fixed instance. In this generic fashion, we are able to replicate a number of previously known hardness results (such as the optimality of the factors $1 - 1/e$ and $1/2$ mentioned above), and we also produce new hardness results using this construction (Theorem 1.2). Our construction helps explain the particular hardness thresholds obtained under various constraints, by exhibiting a small instance where the threshold can be seen as the gap between the optimal solution and the best symmetric solution. The query complexity results in [13, 22, 31] can be seen in hindsight as special cases of Theorem 1.8, but the construction in this paper is somewhat different and technically more involved than the previous proofs for particular cases.

Concrete results Before we proceed to describe our general hardness result, we present its implications for two more concrete problems. We also provide closely matching approximation algorithms for these two problems, based on the multilinear relaxation. In the following, we assume that the objective function is given by a value oracle and the feasibility constraint is given by a membership oracle: a value oracle for f returns the value $f(S)$ for a given set S , and a membership oracle for \mathcal{F} answers whether $S \in \mathcal{F}$ for a given set S .

First, we consider the problem of maximizing a nonnegative (possibly nonmonotone) submodular function subject to a *matroid base* constraint. (This generalizes for example the maximum bisection problem in graphs.) We show that the approximability of this problem is related to base packings in the matroid. We use the following definition.

Definition 1.1. For a matroid \mathcal{M} with a collection of bases \mathcal{B} , the *fractional base packing number* is the maximum possible value of $\sum_{B \in \mathcal{B}} \alpha_B$ for $\alpha_B \geq 0$ such that $\sum_{B \in \mathcal{B}: j \in B} \alpha_B \leq 1$ for every element j .

Example. Consider a uniform matroid with bases $\mathcal{B} = \{B \subset [n] : |B| = k\}$. ($[n]$ denotes the set of integers $\{1, 2, \dots, n\}$.) Here, we can take each base with a coefficient $\alpha_B = 1/\binom{n-1}{k-1}$, which satisfies the condition $\sum_{B \in \mathcal{B}: j \in B} \alpha_B \leq 1$ for every element j since every element is contained in $\binom{n-1}{k-1}$ bases. We obtain that the fractional packing number is at least $\binom{n}{k} / \binom{n-1}{k-1} = \frac{n}{k}$. It is also easy to check that the fractional packing number cannot be larger than $\frac{n}{k}$.

Theorem 1.2. For any ν in the form $\nu = \frac{k}{k-1}$, $k \geq 2$, and any fixed $\epsilon > 0$, a $(1 - \frac{1}{\nu} + \epsilon) = \frac{1}{k}$ -approximation for the problem $\max\{f(S) : S \in \mathcal{B}\}$, where $f(S)$ is a nonnegative submodular function, and \mathcal{B} is a collection of bases in a matroid with fractional packing number at least ν , would require exponentially many value queries.

On the other hand, for any $\nu \in (1, 2]$, there is a randomized $\frac{1}{2}(1 - \frac{1}{\nu} - o(1))$ -approximation for the same problem.

In case the matroid contains two disjoint bases ($\nu = 2$), we obtain a $(\frac{1}{4} - o(1))$ -approximation, improving the previously known factor of $\frac{1}{6} - o(1)$ [19]. In the range of $\nu \in (1, 2]$, our positive and negative results are within a factor of 2. For maximizing a submodular function over the bases of a general matroid, we obtain the following.

Corollary 1.3. *For the problem $\max\{f(S) : S \in \mathcal{B}\}$, where $f(S)$ is a nonnegative submodular function, and \mathcal{B} is a collection of bases in a matroid, any constant-factor approximation requires an exponential number of value queries.*

We also consider the problem of maximizing a nonnegative submodular function subject to a matroid independence constraint.

Theorem 1.4. *For any $\epsilon > 0$, a $(\frac{1}{2} + \epsilon)$ -approximation for the problem $\max\{f(S) : S \in \mathcal{I}\}$, where $f(S)$ is a nonnegative submodular function, and \mathcal{I} is a collection of independent sets in a matroid, would require exponentially many value queries.*

On the other hand, there is a randomized $\frac{1}{4}(-1 + \sqrt{5} - o(1)) \simeq 0.309$ -approximation for the same problem.

Our algorithmic result improves a previously known $(\frac{1}{4} - o(1))$ -approximation [19]. The hardness threshold follows from our general result, but also quite easily from [13].

Hardness from the symmetry gap. Now we describe our general hardness result. Consider an instance $\max\{f(S) : S \in \mathcal{F}\}$ which exhibits a certain degree of symmetry. This is formalized by the notion of a *symmetry group* \mathcal{G} . We consider permutations $\sigma \in \mathbf{S}(X)$ where $\mathbf{S}(X)$ is the symmetric group (of all permutations) on the ground set X . We also use σ for the naturally induced mapping of subsets of X : $\sigma(S) = \{\sigma(i) : i \in S\}$. We say that the instance is invariant under $\mathcal{G} \subset \mathbf{S}(X)$, if for any $\sigma \in \mathcal{G}$ and any $S \subseteq X$, $f(S) = f(\sigma(S))$ and $S \in \mathcal{F} \Leftrightarrow \sigma(S) \in \mathcal{F}$. We emphasize that even though we apply σ to sets, it must be derived from a permutation on X . For $\mathbf{x} \in [0, 1]^X$, we define the “symmetrization of \mathbf{x} ” as

$$\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})],$$

where $\sigma \in \mathcal{G}$ is uniformly random and $\sigma(\mathbf{x})$ denotes \mathbf{x} with coordinates permuted by σ .

Erratum: The main hardness result in the conference version of this paper [32] was formulated for an arbitrary feasibility constraint \mathcal{F} , invariant under \mathcal{G} . Unfortunately, this was an error and the theorem does not hold in that form — an algorithm could gather some information from querying \mathcal{F} , and combining this with information obtained by querying the objective function f it could possibly determine the hidden optimal solution. The possibility of gathering information from the membership oracle for \mathcal{F} was neglected in the proof. (The reason for this was probably that the feasibility constraints used in concrete applications of the theorem were very simple and indeed did not provide any information about the optimum.) Nevertheless, to correct this issue, one needs to impose a stronger symmetry constraint on \mathcal{F} , namely the condition that $S \in \mathcal{F}$ depends only on the symmetrized version of S , $\overline{\mathbf{1}}_S = \mathbf{E}_{\sigma \in \mathcal{G}}[\mathbf{1}_{\sigma(S)}]$. This is the case in all the applications of the hardness theorem in [32] and [26] and hence these applications are not affected.

Definition 1.5. *We call an instance $\max\{f(S) : S \in \mathcal{F}\}$ on a ground set X strongly symmetric with respect to a group of permutations \mathcal{G} on X , if $f(S) = f(\sigma(S))$ for all $S \subseteq X$ and $\sigma \in \mathcal{G}$, and $S \in \mathcal{F} \Leftrightarrow S' \in \mathcal{F}$ whenever $\mathbf{E}_{\sigma \in \mathcal{G}}[\mathbf{1}_{\sigma(S)}] = \mathbf{E}_{\sigma \in \mathcal{G}}[\mathbf{1}_{\sigma(S')}]$.*

Example. A cardinality constraint, $\mathcal{F} = \{S \subseteq [n] : |S| \leq k\}$, is strongly symmetric with respect to all permutations, because the condition $S \in \mathcal{F}$ depends only on the symmetrized vector $\overline{\mathbf{1}}_S = \frac{|S|}{n}\mathbf{1}$. Similarly, a partition matroid constraint, $\mathcal{F} = \{S \subseteq [n] : |S \cap X_i| \leq k\}$ for disjoint sets X_i , is also strongly symmetric. On the other hand, consider a family of feasible solution $\mathcal{F} = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 1\}\}$. This family is invariant under a group generated by the cyclic rotation $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 1$. It is not strongly symmetric, because the condition $S \in \mathcal{F}$ does not depend only the symmetrized vector $\overline{\mathbf{1}}_S = (\frac{1}{4}|S|, \frac{1}{4}|S|, \frac{1}{4}|S|, \frac{1}{4}|S|)$; some pairs are feasible and others are not.

Next, we define the symmetry gap as the ratio between the optimal solution of $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$ and the best *symmetric* solution of this problem.

Definition 1.6 (Symmetry gap). Let $\max\{f(S) : S \in \mathcal{F}\}$ be an instance on a ground set X , which is strongly symmetric with respect to $\mathcal{G} \subset \mathbf{S}(X)$. Let $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ be the multilinear extension of $f(S)$ and $P(\mathcal{F}) = \text{conv}\{\mathbf{1}_I : I \in \mathcal{F}\}$ the polytope associated with \mathcal{F} . Let $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$. The symmetry gap of $\max\{f(S) : S \in \mathcal{F}\}$ is defined as $\gamma = \overline{OPT}/OPT$ where

$$OPT = \max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\},$$

$$\overline{OPT} = \max\{F(\bar{\mathbf{x}}) : \mathbf{x} \in P(\mathcal{F})\}.$$

We give examples of computing the symmetry gap in Section 2. Next, we need to define the notion of a *refinement* of an instance. This is a natural way to extend a family of feasible sets to a larger ground set. In particular, this operation preserves the types of constraints that we care about, such as cardinality constraints, matroid independence, and matroid base constraints.

Definition 1.7 (Refinement). Let $\mathcal{F} \subseteq 2^X$, $|X| = k$ and $|N| = n$. We say that $\tilde{\mathcal{F}} \subseteq 2^{N \times X}$ is a refinement of \mathcal{F} , if

$$\tilde{\mathcal{F}} = \left\{ \tilde{S} \subseteq N \times X \mid (x_1, \dots, x_k) \in P(\mathcal{F}) \text{ where } x_j = \frac{1}{n} |\tilde{S} \cap (N \times \{j\})| \right\}.$$

In other words, in the refined instance, each element $j \in X$ is replaced by a set $N \times \{j\}$. We call this set the *cluster* of elements corresponding to j . A set \tilde{S} is in $\tilde{\mathcal{F}}$ if and only if the fractions x_j of the respective clusters that are intersected by \tilde{S} form a vector $\mathbf{x} \in P(\mathcal{F})$, i.e. a convex combination of sets in \mathcal{F} .

Our main result is that the symmetry gap for any strongly symmetric instance translates automatically into hardness of approximation for refined instances. (See Definition 1.5 for the notion of being “strongly symmetric”.) We emphasize that this is a query-complexity lower bound, and hence independent of assumptions such as $P \neq NP$.

Theorem 1.8. Let $\max\{f(S) : S \in \mathcal{F}\}$ be an instance of nonnegative (optionally monotone) submodular maximization, strongly symmetric with respect to \mathcal{G} , with symmetry gap $\gamma = \overline{OPT}/OPT$. Let \mathcal{C} be the class of instances $\max\{\tilde{f}(S) : S \in \tilde{\mathcal{F}}\}$ where \tilde{f} is nonnegative (optionally monotone) submodular and $\tilde{\mathcal{F}}$ is a refinement of \mathcal{F} . Then for every $\epsilon > 0$, any (even randomized) $(1 + \epsilon)\gamma$ -approximation algorithm for the class \mathcal{C} would require exponentially many value queries to $\tilde{f}(S)$.

We remark that the result holds even if the class \mathcal{C} is restricted to instances which are themselves symmetric under a group related to \mathcal{G} (see the discussion in Section 3, after the proofs of Theorem 1.8 and 1.9). On the algorithmic side, submodular maximization seems easier for symmetric instances and in this case we obtain optimal approximation factors, up to lower-order terms (see Section 5).

Our hardness construction yields impossibility results also for solving the continuous problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$. In the case of matroid constraints, this is easy to see, because an approximation to the continuous problem gives the same approximation factor for the discrete problem (by pipage rounding, see Appendix B). However, this phenomenon is more general and we can show that the value of a symmetry gap translates into an inapproximability result for the multilinear optimization problem under any constraint satisfying a symmetry condition.

Theorem 1.9. Let $\max\{f(S) : S \in \mathcal{F}\}$ be an instance of nonnegative (optionally monotone) submodular maximization, strongly symmetric with respect to \mathcal{G} , with symmetry gap $\gamma = \overline{OPT}/OPT$. Let \mathcal{C} be the class of instances $\max\{\tilde{f}(S) : S \in \tilde{\mathcal{F}}\}$ where \tilde{f} is nonnegative (optionally monotone) submodular and $\tilde{\mathcal{F}}$ is a refinement of \mathcal{F} . Then for every $\epsilon > 0$, any (even randomized) $(1 + \epsilon)\gamma$ -approximation algorithm for the multilinear relaxation $\max\{\tilde{F}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ of problems in \mathcal{C} would require exponentially many value queries to $\tilde{f}(S)$.

Additions to the conference version and follow-up work. An extended abstract of this work appeared in IEEE FOCS 2009 [32]. As mentioned above, the main theorem in [32] suffers from a technical flaw. This does not affect the applications, but the general theorem in [32] is not correct. We provide a corrected version of the main theorem with a complete proof (Theorem 1.8) and we extend this hardness result to the problem of solving the multilinear relaxation (Theorem 1.9).

Subsequently, further work has been done which exploits the symmetry gap concept. In [26], it has been proved using Theorem 1.8 that maximizing a nonnegative submodular function subject to a matroid independence constraint with a factor better than 0.478 would require exponentially many queries. Even in the case of a cardinality constraint, $\max\{f(S) : |S| \leq k\}$ cannot be approximated within a factor better than 0.491 using subexponentially many queries [26]. In the case of a matroid base constraint, assuming that the fractional base packing number is $\nu = \frac{k}{k-1}$ for some $k \geq 2$, there is no $(1 - e^{-1/k} + \epsilon)$ -approximation in the value oracle model [26], improving the hardness of $(1 - \frac{1}{\nu} + \epsilon) = (\frac{1}{k} + \epsilon)$ -approximation from this paper. These applications are not affected by the flaw in [32], and they are implied by the corrected version of Theorem 1.8 here.

Recently [7], it has been proved using the symmetry gap technique that combinatorial auctions with submodular bidders do not admit any truthful-in-expectation $1/m^\gamma$ -approximation, where m is the number of items and $\gamma > 0$ some absolute constant. This is the first nontrivial hardness result for truthful-in-expectation mechanisms for combinatorial auctions; it separates the classes of monotone submodular functions and coverage functions, where a truthful-in-expectation $(1 - 1/e)$ -approximation is possible [6]. The proof is self-contained and does not formally refer to [32].

Moreover, this hardness result for truthful-in-expectation mechanisms as well as the main hardness result in this paper have been converted from the oracle setting to a computational complexity setting [8, 9]. This recent work shows that the hardness of approximation arising from symmetry gap is not limited to instances given by an oracle, but holds also for instances encoded explicitly on the input, under a suitable complexity-theoretic assumption.

Organization. The rest of the paper is organized as follows. In Section 2, we present applications of our main hardness result (Theorem 1.8) to concrete cases, in particular we show how it implies the hardness statements in Theorem 1.4 and 1.2. In Section 3, we present the proofs of Theorem 1.8 and Theorem 1.9. In Section 4, we prove the algorithmic results in Theorem 1.4 and 1.2. In Section 5, we discuss the special case of symmetric instances. In the Appendix, we present a few basic facts concerning submodular functions, an extension of pipage rounding to matroid independence polytopes (rather than matroid base polytopes), and other technicalities that would hinder the main exposition.

2 From symmetry to inapproximability: applications

Before we get into the proof of Theorem 1.8, let us show how it can be applied to a number of specific problems. Some of these are hardness results that were proved previously by an ad-hoc method. The last application is a new one (Theorem 1.2).

Nonmonotone submodular maximization. Let $X = \{1, 2\}$ and for any $S \subseteq X$, $f(S) = 1$ if $|S| = 1$, and 0 otherwise. Consider the instance $\max\{f(S) : S \subseteq X\}$. In other words, this is the Max Cut problem on the graph K_2 . This instance exhibits a simple symmetry, the group of all (two) permutations on $\{1, 2\}$. We get $OPT = F(1, 0) = F(0, 1) = 1$, while $\overline{OPT} = F(1/2, 1/2) = 1/2$. Hence, the symmetry gap is $1/2$.



Figure 1: Symmetric instance for nonmonotone submodular maximization: Max Cut on the graph K_2 . The white set denotes the optimal solution, while \bar{x} is the (unique) symmetric solution.

Since $f(S)$ is nonnegative submodular and there is no constraint on $S \subseteq X$, this will be the case for any refinement of the instance as well. Theorem 1.8 implies immediately the following: any algorithm achieving a $(\frac{1}{2} + \epsilon)$ -approximation for nonnegative (nonmonotone) submodular maximization requires exponentially many value queries (which was previously known [13]). Note that a “trivial instance” implies a nontrivial hardness result. This is typically the case in applications of Theorem 1.8.

The same symmetry gap holds if we impose some simple constraints: the problems $\max\{f(S) : |S| \leq 1\}$ and $\max\{f(S) : |S| = 1\}$ have the same symmetry gap as above. Hence, the hardness threshold of $1/2$

also holds for nonmonotone submodular maximization under cardinality constraints of the type $|S| \leq n/2$, or $|S| = n/2$. This proves the hardness part of Theorem 1.4. This can be derived quite easily from the construction of [13] as well.

Monotone submodular maximization. Let $X = [k]$ and $f(S) = \min\{|S|, 1\}$. Consider the instance $\max\{f(S) : |S| \leq 1\}$. This instance is invariant under all permutations on $[k]$, the symmetric group \mathbf{S}_k . Note that the instance is *strongly symmetric* (Def. 1.5) with respect to \mathbf{S}_k , since the feasibility constraint $|S| \leq 1$ depends only on the symmetrized vector $\bar{\mathbf{1}}_S = (\frac{1}{k}|S|, \dots, \frac{1}{k}|S|)$. We get $OPT = F(1, 0, \dots, 0) = 1$, while $\overline{OPT} = F(1/k, 1/k, \dots, 1/k) = 1 - (1 - 1/k)^k$.

Here, $f(S)$ is monotone submodular and any refinement of \mathcal{F} is a set system of the type $\tilde{\mathcal{F}} = \{S : |S| \leq \ell\}$. Based on our theorem, this implies that any approximation better than $1 - (1 - 1/k)^k$ for monotone submodular maximization subject to a cardinality constraint would require exponentially many value queries. Since this holds for any fixed k , we get the same hardness result for any $\beta > \lim_{k \rightarrow \infty} (1 - (1 - 1/k)^k) = 1 - 1/e$ (which was previously known [25]).

Submodular maximization over matroid bases. Let $X = A \cup B$, $A = \{a_1, \dots, a_k\}$, $B = \{b_1, \dots, b_k\}$ and $\mathcal{F} = \{S : |S \cap A| = 1 \ \& \ |S \cap B| = k - 1\}$. We define $f(S) = \sum_{i=1}^k f_i(S)$ where $f_i(S) = 1$ if $a_i \in S$ & $b_i \notin S$, and 0 otherwise. This instance can be viewed as a Maximum Directed Cut problem on a graph of k disjoint arcs, under the constraint that exactly one arc tail and $k - 1$ arc heads should be on the left-hand side (S). An optimal solution is for example $S = \{a_1, b_2, b_3, \dots, b_k\}$, which gives $OPT = 1$. The symmetry here is that we can apply the same permutation to A and B simultaneously. Again, the feasibility of a set S depends only on the symmetrized vector $\bar{\mathbf{1}}_S$: in fact $S \in \mathcal{F}$ if and only if $\bar{\mathbf{1}}_S = (\frac{1}{k}, \dots, \frac{1}{k}, 1 - \frac{1}{k}, \dots, 1 - \frac{1}{k})$. There is a unique symmetric solution $\bar{\mathbf{x}} = (\frac{1}{k}, \dots, \frac{1}{k}, 1 - \frac{1}{k}, \dots, 1 - \frac{1}{k})$, and $\overline{OPT} = F(\bar{\mathbf{x}}) = \mathbf{E}[f(\hat{\mathbf{x}})] = \sum_{i=1}^k \mathbf{E}[f_i(\hat{\mathbf{x}})] = \frac{1}{k}$ (since each arc appears in the directed cut induced by $\hat{\mathbf{x}}$ with probability $\frac{1}{k^2}$).




Figure 2: Symmetric instance for submodular maximization over matroid bases.

The refined instances are instances of (nonmonotone) submodular maximization over the bases of a matroid, where the ground set is partitioned into $A \cup B$ and we should take a $\frac{1}{k}$ -fraction of A and a $(1 - \frac{1}{k})$ -fraction of B . (This means that the fractional packing number of bases is $\nu = \frac{k}{k-1}$.) Our theorem implies that for this class of instances, an approximation better than $1/k$ is impossible - this proves the hardness part of Theorem 1.2.

Observe that in all the cases mentioned above, the multilinear relaxation is equivalent to the original problem, in the sense that any fractional solution can be rounded without any loss in the objective value. This implies that the same hardness factors apply to solving the multilinear relaxation of the respective problems. In particular, using the last result (for matroid bases), we obtain that the multilinear optimization problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$ does not admit a constant factor for nonnegative submodular functions and matroid base polytopes. (We remark that a $(1 - 1/e)$ -approximation can be achieved for any *monotone* submodular function and any solvable polytope, i.e. polytope over which we can optimize linear functions [30].)

As Theorem 1.9 shows, this holds more generally - any symmetry gap construction gives an inapproximability result for solving the multilinear optimization problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$. This in fact implies limits on what hardness results we can possibly hope for using this technique. For instance, we cannot prove using the symmetry gap that the monotone submodular maximization problem subject to the intersection of k matroid constraints does not admit a constant factor - because we would also prove that the

respective multilinear relaxation does not admit such an approximation. But we know from [30] that a $(1 - 1/e)$ -approximation is possible for the multilinear problem in this case.

Hence, the hardness arising from the symmetry gap is related to the difficulty of solving the multilinear optimization problem rather than the difficulty of rounding a fractional solution. Thus this technique is primarily suited to optimization problems where the multilinear optimization problem captures closely the original discrete problem.

3 From symmetry to inapproximability: proof

The roadmap At a high level, our proof resembles the constructions of [13, 22]. We construct instances based on continuous functions $F(\mathbf{x}), G(\mathbf{x})$, whose optima differ by a gap for which we want to prove hardness. Then we show that after a certain perturbation, the two instances are very hard to distinguish. This paper generalizes the ideas of [13, 22] and brings two new ingredients. First, we show that the functions $F(\mathbf{x}), G(\mathbf{x})$, which are “pulled out of the hat” in [13, 22], can be produced in a natural way from the multilinear relaxation of the respective problem, using the notion of a *symmetry gap*. Secondly, the functions $F(\mathbf{x}), G(\mathbf{x})$ are perturbed in a way that makes them indistinguishable and this forms the main technical part of the proof. In [13], this step is quite simple. In [22], the perturbation is more complicated, but still relies on properties of the functions $F(\mathbf{x}), G(\mathbf{x})$ specific to that application. The construction that we present here (Lemma 3.2) uses the symmetry properties of a fixed instance in a generic fashion.

First, let us present an outline of our construction. Given an instance $\max\{f(S) : S \in \mathcal{F}\}$ exhibiting a symmetry gap γ , we consider two smooth submodular² functions, $F(\mathbf{x})$ and $G(\mathbf{x})$. The first one is the multilinear extension $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$, while the second one is its symmetrized version $G(\mathbf{x}) = F(\bar{\mathbf{x}})$. We modify these functions slightly so that we obtain functions $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$ with the following property: For any vector \mathbf{x} which is close to its symmetrized version $\bar{\mathbf{x}}$, $\hat{F}(\mathbf{x}) = \hat{G}(\mathbf{x})$. The functions $\hat{F}(\mathbf{x}), \hat{G}(\mathbf{x})$ induce instances of submodular maximization on the refined ground sets. The way we define discrete instances based on $\hat{F}(\mathbf{x}), \hat{G}(\mathbf{x})$ is natural, using the following lemma. Essentially, we interpret the fractional variables as fractions of clusters in the refined instance.

Lemma 3.1. *Let $F : [0, 1]^X \rightarrow R$, $N = [n]$, $n \geq 1$, and define $f : 2^{N \times X} \rightarrow R$ so that $f(S) = F(\mathbf{x})$ where $x_i = \frac{1}{n}|S \cap (N \times \{i\})|$. Then*

1. *If $\frac{\partial F}{\partial x_i} \geq 0$ everywhere for each i , then f is monotone.*
2. *If the first partial derivatives of F are absolutely continuous³ and $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$ almost everywhere for all i, j , then f is submodular.*

Proof. First, assume $\frac{\partial F}{\partial x_i} \geq 0$ everywhere for all i . This implies that F is nondecreasing in every coordinate, i.e. $F(\mathbf{x}) \leq F(\mathbf{y})$ whenever $\mathbf{x} \leq \mathbf{y}$. This means that $f(S) \leq f(T)$ whenever $S \subseteq T$.

Next, assume $\frac{\partial F}{\partial x_i}$ is absolutely continuous for each i and $\frac{\partial^2 F}{\partial x_j \partial x_i} \leq 0$ almost everywhere for all i, j . We want to prove that $\frac{\partial F}{\partial x_i}|_{\mathbf{x}} \geq \frac{\partial F}{\partial x_i}|_{\mathbf{y}}$ whenever $\mathbf{x} \leq \mathbf{y}$, which implies that the marginal values of f are nonincreasing.

Let $\mathbf{x} \leq \mathbf{y}$, fix $\delta > 0$ arbitrarily small, and pick \mathbf{x}', \mathbf{y}' such that $\|\mathbf{x}' - \mathbf{x}\| < \delta, \|\mathbf{y}' - \mathbf{y}\| < \delta, \mathbf{x}' \leq \mathbf{y}'$ and on the line segment $[\mathbf{x}', \mathbf{y}']$, we have $\frac{\partial^2 F}{\partial x_j \partial x_i} \leq 0$ except for a set of (1-dimensional) measure zero. If such a pair of points \mathbf{x}', \mathbf{y}' does not exist, it means that there are sets A, B of positive measure such that $\mathbf{x} \in A, \mathbf{y} \in B$ and for any $\mathbf{x}' \in A, \mathbf{y}' \in B$, the line segment $[\mathbf{x}', \mathbf{y}']$ contains a subset of positive (1-dimensional) measure where $\frac{\partial^2 F}{\partial x_j \partial x_i}$ for some j is positive or undefined. This would imply that $[0, 1]^X$ contains a subset of positive measure where $\frac{\partial^2 F}{\partial x_j \partial x_i}$ for some j is positive or undefined, which we assume is not the case.

Therefore, there is a pair of points \mathbf{x}', \mathbf{y}' as described above. We compare $\frac{\partial F}{\partial x_i}|_{\mathbf{x}'}$ and $\frac{\partial F}{\partial x_i}|_{\mathbf{y}'}$ by integrating along the line segment $[\mathbf{x}', \mathbf{y}']$. Since $\frac{\partial F}{\partial x_i}$ is absolutely continuous and $\frac{\partial}{\partial x_j} \frac{\partial F}{\partial x_i} = \frac{\partial^2 F}{\partial x_j \partial x_i} \leq 0$ along this line segment for all j except for a set of measure zero, we obtain $\frac{\partial F}{\partial x_i}|_{\mathbf{x}'} \geq \frac{\partial F}{\partial x_i}|_{\mathbf{y}'}$. This is true for \mathbf{x}', \mathbf{y}' arbitrarily

²“Smooth submodularity” means the condition $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$ for all i, j .

³A function $F : [0, 1]^X \rightarrow R$ is absolutely continuous, if $\forall \epsilon > 0; \exists \delta > 0; \sum_{i=1}^t \|\mathbf{x}_i - \mathbf{y}_i\| < \delta \Rightarrow \sum_{i=1}^t |F(\mathbf{x}_i) - F(\mathbf{y}_i)| < \epsilon$.

close to \mathbf{x}, \mathbf{y} , and hence by continuity of $\frac{\partial F}{\partial x_i}$, we get $\frac{\partial F}{\partial x_i}|_{\mathbf{x}} \geq \frac{\partial F}{\partial x_i}|_{\mathbf{y}}$. This implies that the marginal values of f are nonincreasing. \square

The way we construct $\hat{F}(\mathbf{x}), \hat{G}(\mathbf{x})$ is such that, given a large enough refinement of the ground set, it is impossible to distinguish the instances corresponding to $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$. As we argue more precisely later, this holds because if the ground set is large and labeled in a random way (considering the symmetry group of the instance), a query about a vector \mathbf{x} effectively becomes a query about the symmetrized vector $\bar{\mathbf{x}}$. We would like this property to imply that all queries with high probability fall in the region where $\hat{F}(\mathbf{x}) = \hat{G}(\mathbf{x})$ and the inability to distinguish between \hat{F} and \hat{G} gives the hardness result that we want. The following lemma gives the precise properties of $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$ that we need.

Lemma 3.2. *Consider a function $f : 2^X \rightarrow R_+$ invariant under a group of permutations \mathcal{G} on the ground set X . Let $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$, $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$, and fix any $\epsilon > 0$. Then there is $\delta > 0$ and functions $\hat{F}, \hat{G} : [0, 1]^X \rightarrow R_+$ (which are also symmetric with respect to \mathcal{G}), satisfying:*

1. For all $\mathbf{x} \in [0, 1]^X$, $\hat{G}(\mathbf{x}) = \hat{F}(\bar{\mathbf{x}})$.
2. For all $\mathbf{x} \in [0, 1]^X$, $|\hat{F}(\mathbf{x}) - F(\mathbf{x})| \leq \epsilon$.
3. Whenever $\|\mathbf{x} - \bar{\mathbf{x}}\|^2 \leq \delta$, $\hat{F}(\mathbf{x}) = \hat{G}(\mathbf{x})$ and the value depends only on $\bar{\mathbf{x}}$.
4. The first partial derivatives of \hat{F}, \hat{G} are absolutely continuous.
5. If f is monotone, then $\frac{\partial \hat{F}}{\partial x_i} \geq 0$ and $\frac{\partial \hat{G}}{\partial x_i} \geq 0$ everywhere.
6. If f is submodular, then $\frac{\partial^2 \hat{F}}{\partial x_i \partial x_j} \leq 0$ and $\frac{\partial^2 \hat{G}}{\partial x_i \partial x_j} \leq 0$ almost everywhere.

The proof of this lemma is the main technical part of this paper and we defer it to the end of this section. Assuming this lemma, we first finish the proof of the main theorem. We prove the following.

Lemma 3.3. *Let \hat{F}, \hat{G} be the two functions provided by Lemma 3.2. For a parameter $n \in Z_+$ and $N = [n]$, define two discrete functions $\hat{f}, \hat{g} : 2^{N \times X} \rightarrow R_+$ as follows: Let $\sigma^{(i)}$ be an arbitrary permutation in \mathcal{G} for each $i \in N$. For every set $S \subseteq N \times X$, we define a vector $\xi(S) \in [0, 1]^X$ by*

$$\xi_j(S) = \frac{1}{n} \left| \{i \in N : (i, \sigma^{(i)}(j)) \in S\} \right|.$$

Let us define:

$$\hat{f}(S) = \hat{F}(\xi(S)), \quad \hat{g}(S) = \hat{G}(\xi(S)).$$

In addition, let $\tilde{\mathcal{F}} = \{\tilde{S} : \xi(\tilde{S}) \in P(\mathcal{F})\}$ be a feasibility constraint such that the condition $S \in \mathcal{F}$ depends only on the symmetrized vector $\bar{\mathbf{1}}_S$. Then deciding whether an instance given by value/membership oracles is $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$ or $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$ (even by a randomized algorithm, with any constant probability of success) requires an exponential number of queries.

Proof. Let $\sigma^{(i)} \in \mathcal{G}$ be chosen independently at random for each $i \in N$ and consider the instances $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$, $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$ as described in the lemma. We show that every deterministic algorithm will follow the same computation path and return the same answer on both instances, with high probability. By Yao's principle, this means that every randomized algorithm returns the same answer for the two instances with high probability, for some particular $\sigma^{(i)} \in \mathcal{G}$.

The feasible sets in the refined instance, $S \in \tilde{\mathcal{F}}$, are such that the respective vector $\xi(S)$ is in the polytope $P(\mathcal{F})$. Since the instance is strongly symmetric, the condition $S \in \mathcal{F}$ depends only on the symmetrized vector $\bar{\mathbf{1}}_S$. Hence, the condition $\xi(S) \in P(\mathcal{F})$ depends only on the symmetrized vector $\bar{\xi}(S)$. Therefore, $S \in \tilde{\mathcal{F}} \Leftrightarrow \xi(S) \in P(\mathcal{F}) \Leftrightarrow \bar{\xi}(S) \in P(\mathcal{F})$. We have $\bar{\xi}(S)_j = \mathbf{E}_{\sigma \in \mathcal{G}} \left[\frac{1}{n} \left| \{i \in N : (i, \sigma^{(i)}(\sigma(j))) \in S\} \right| \right]$. The distribution of $\sigma^{(i)} \circ \sigma$ is the same as that of σ , i.e. uniform over \mathcal{G} . Hence, $\bar{\xi}(S)_j$ and consequently the condition $S \in \tilde{\mathcal{F}}$ does not depend on $\sigma^{(i)}$ in any way. Intuitively, an algorithm cannot learn any information about the permutations $\sigma^{(i)}$ by querying the feasibility oracle, since the feasibility condition $S \in \tilde{\mathcal{F}}$ does not depend on $\sigma^{(i)}$ for any $i \in N$.

The main part of the proof is to show that even queries to the objective function are unlikely to reveal any information about the permutations $\sigma^{(i)}$. The key observation is that for any fixed query Q to the objective function, the associated vector $\mathbf{q} = \xi(Q)$ is very likely to be close to its symmetrized version $\bar{\mathbf{q}}$. To see this, consider a query Q . The associated vector $\mathbf{q} = \xi(Q)$ is determined by

$$q_j = \frac{1}{n} |\{i \in N : (i, \sigma^{(i)}(j)) \in Q\}| = \frac{1}{n} \sum_{i=1}^n Q_{ij}$$

where Q_{ij} is an indicator variable of the event $(i, \sigma^{(i)}(j)) \in Q$. This is a random event due to the randomness in $\sigma^{(i)}$. We have

$$\mathbf{E}[Q_{ij}] = \Pr[Q_{ij} = 1] = \Pr_{\sigma^{(i)} \in \mathcal{G}} [(i, \sigma^{(i)}(j)) \in Q].$$

Adding up these expectations over $i \in N$, we get

$$\sum_{i \in N} \mathbf{E}[Q_{ij}] = \sum_{i \in N} \Pr_{\sigma^{(i)} \in \mathcal{G}} [(i, \sigma^{(i)}(j)) \in Q] = \mathbf{E}_{\sigma \in \mathcal{G}} [|\{i \in N : (i, \sigma(j)) \in Q\}|].$$

For the purposes of expectation, the independence of $\sigma^{(1)}, \dots, \sigma^{(n)}$ is irrelevant and that is why we replace them by one random permutation σ . On the other hand, consider the symmetrized vector $\bar{\mathbf{q}}$, with coordinates

$$\bar{q}_j = \mathbf{E}_{\sigma \in \mathcal{G}} [q_{\sigma(j)}] = \frac{1}{n} \mathbf{E}_{\sigma \in \mathcal{G}} [|\{i \in N : (i, \sigma^{(i)}(\sigma(j))) \in Q\}|] = \frac{1}{n} \mathbf{E}_{\sigma \in \mathcal{G}} [|\{i \in N : (i, \sigma(j)) \in Q\}|]$$

using the fact that the distribution of $\sigma^{(i)} \circ \sigma$ is the same as the distribution of σ - uniformly random over \mathcal{G} . Note that the vector \mathbf{q} depends on the random permutations $\sigma^{(i)}$ but the symmetrized vector $\bar{\mathbf{q}}$ does not; this will be also useful in the following. For now, we summarize that

$$\bar{q}_j = \frac{1}{n} \sum_{i=1}^n \mathbf{E}[Q_{ij}] = \mathbf{E}[q_j].$$

Since each permutation $\sigma^{(i)}$ is chosen independently, the random variables $\{Q_{ij} : 1 \leq i \leq n\}$ are independent (for a fixed j). We can apply Chernoff's bound (see e.g. [1], Corollary A.1.7):

$$\Pr \left[\left| \sum_{i=1}^n Q_{ij} - \sum_{i=1}^n \mathbf{E}[Q_{ij}] \right| > a \right] < 2e^{-2a^2/n}.$$

Using $q_j = \frac{1}{n} \sum_{i=1}^n Q_{ij}$, $\bar{q}_j = \frac{1}{n} \sum_{i=1}^n \mathbf{E}[Q_{ij}]$ and setting $a = n\sqrt{\delta/|X|}$, we obtain

$$\Pr \left[|q_j - \bar{q}_j| > \sqrt{\delta/|X|} \right] < 2e^{-2n\delta/|X|}.$$

By the union bound,

$$\Pr[|\mathbf{q} - \bar{\mathbf{q}}|^2 > \delta] \leq \sum_{j \in X} \Pr[|q_j - \bar{q}_j|^2 > \delta/|X|] < 2|X|e^{-2n\delta/|X|}. \quad (1)$$

Note that while δ and $|X|$ are constants, n grows as the size of the refinement and hence the probability is exponentially small in the size of the ground set $N \times X$.

Define $D(\mathbf{q}) = \|\mathbf{q} - \bar{\mathbf{q}}\|^2$. As long as $D(\mathbf{q}) \leq \delta$ for every query issued by the algorithm, the answers do not depend on the randomness of the input. This is because then the values of $\hat{F}(\mathbf{q})$ and $\hat{G}(\mathbf{q})$ depend only on $\bar{\mathbf{q}}$, which is independent of the random permutations $\sigma^{(i)}$, as we argued above. Therefore, assuming that $D(\mathbf{q}) \leq \delta$ for each query, the algorithm will always follow the same path of computation and issue the same sequence of queries \mathcal{S} . (Note that this is just a fixed sequence which can be written down before we started running the algorithm.) Assume that $|\mathcal{S}| = e^{o(n)}$, i.e., the number of queries is subexponential in n . By (1), using a union bound over all $Q \in \mathcal{S}$, it happens with probability $1 - e^{-\Omega(n)}$ that $D(\mathbf{q}) = \|\mathbf{q} - \bar{\mathbf{q}}\|^2 \leq \delta$ for all $Q \in \mathcal{S}$. (Note that δ and $|X|$ are constants here.) Then, the algorithm indeed follows this path of computation and gives the same answer. In particular, the answer does not depend on whether the instance is $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$ or $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$. \square

Proof of Theorem 1.8. Fix an $\epsilon > 0$. Given an instance $\max\{f(S) : S \in \mathcal{F}\}$ strongly symmetric under \mathcal{G} , let $\hat{F}, \hat{G} : [0, 1]^X \rightarrow R$ be the two functions provided by Lemma 3.2. We choose a large number n and consider a refinement $\tilde{\mathcal{F}}$ on the ground set $N \times X$, where $N = [n]$. We define discrete instances of submodular maximization $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$ and $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$. As in Lemma 3.3, for each $i \in N$ we choose a random permutation $\sigma^{(i)} \in \mathcal{G}$. This can be viewed as a random shuffle of the labeling of the ground set before we present it to an algorithm. For every set $S \subseteq N \times X$, we define a vector $\xi(S) \in [0, 1]^X$ by

$$\xi_j(S) = \frac{1}{n} \left| \{i \in N : (i, \sigma^{(i)}(j)) \in S\} \right|.$$

In other words, $\xi_j(S)$ measures the fraction of copies of element j contained in S ; however, for each i the i -copies of all elements are shuffled by $\sigma^{(i)}$. Next, we define

$$\hat{f}(S) = \hat{F}(\xi(S)), \quad \hat{g}(S) = \hat{G}(\xi(S)).$$

We claim that \hat{f} and \hat{g} are submodular (for any fixed ξ). Note that the effect of $\sigma^{(i)}$ is just a renaming (or shuffling) of the elements of $N \times X$, and hence for the purpose of proving submodularity we can assume that $\sigma^{(i)}$ is the identity for all i . Then, $\xi_j(S) = \frac{1}{n} |S \cap (N \times \{j\})|$. Due to Lemma 3.1, the property $\frac{\partial^2 \hat{F}}{\partial x_i \partial x_j} \leq 0$ (almost everywhere) implies that \hat{f} is submodular. In addition, if the original instance was monotone, then $\frac{\partial \hat{F}}{\partial x_j} \geq 0$ and \hat{f} is monotone. The same holds for \hat{g} .

The value of $\hat{g}(S)$ for any feasible solution $S \in \tilde{\mathcal{F}}$ is bounded by $\hat{g}(S) = \hat{G}(\xi(S)) = \hat{F}(\overline{\xi(S)}) \leq \overline{OPT} + \epsilon$. On the other hand, let \mathbf{x}^* denote a point where the optimum of the continuous problem $\max\{\hat{F}(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$ is attained, i.e. $\hat{F}(\mathbf{x}^*) \geq OPT - \epsilon$. For a large enough n , we can approximate the point \mathbf{x}^* arbitrarily closely by a rational vector with n in the denominator, which corresponds to a discrete solution $S^* \in \tilde{\mathcal{F}}$ whose value $\hat{f}(S^*)$ is at least, say, $OPT - 2\epsilon$. Hence, the ratio between the optima of the *discrete* optimization problems $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$ and $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$ can be made at most $\frac{OPT + \epsilon}{OPT - 2\epsilon}$, i.e. arbitrarily close to the symmetry gap $\gamma = \frac{OPT}{OPT - 2\epsilon}$.

By Lemma 3.3, distinguishing the two instances $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$ and $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$, even by a randomized algorithm, requires an exponential number of value queries. Therefore, we cannot estimate the optimum within a factor better than γ . \square

Next, we prove Theorem 1.9 (again assuming Lemma 3.2), i.e. an analogous hardness result for solving the multilinear optimization problem.

Proof of Theorem 1.9. Given a symmetric instance $\max\{f(S) : S \in \mathcal{F}\}$ and $\epsilon > 0$, we construct refined and modified instances $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$, $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\}$, derived from the continuous functions \hat{F}, \hat{G} provided by Lemma 3.2, exactly as we did in the proof of Theorem 1.8. Lemma 3.3 states that these two instances cannot be distinguished using a subexponential number of value queries. Furthermore, the gap between the two modified instances corresponds to the symmetry gap γ of the original instance: $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\} \geq OPT - 2\epsilon$ and $\max\{\hat{g}(S) : S \in \tilde{\mathcal{F}}\} \leq \overline{OPT} + \epsilon = \gamma OPT + \epsilon$.

Now we consider the multilinear relaxations of the two refined instances, $\max\{\tilde{F}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ and $\max\{\tilde{G}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$. Note that \tilde{F}, \tilde{G} , (although related to \hat{F}, \hat{G}) are not exactly the same as the functions \hat{F}, \hat{G} ; in particular, they are defined on a larger (refined) domain. However, we show that the gap between the optima of the two instances remains the same.

First, the value of $\max\{\tilde{F}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ is at least the optimum of the discrete problem, $\max\{\hat{f}(S) : S \in \tilde{\mathcal{F}}\}$, which is at least $OPT - 2\epsilon$ as in the proof of Theorem 1.8. The value of $\max\{\tilde{G}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ can be analyzed as follows. For any fractional solution $\mathbf{x} \in P(\mathcal{F})$, the value of $\tilde{G}(\mathbf{x})$ is the expectation $\mathbf{E}[\hat{g}(\hat{\mathbf{x}})]$, where $\hat{\mathbf{x}}$ is obtained by independently rounding the coordinates of \mathbf{x} to $\{0, 1\}$. Recall that \hat{g} is obtained by discretizing the continuous function \hat{G} (using Lemma 3.1). In particular, $\hat{g}(S) = \hat{G}(\tilde{\mathbf{x}})$ where $\tilde{x}_i = \frac{1}{n} |S \cap (N \times \{i\})|$ is the fraction of the respective cluster contained in S , and $|N| = n$ is the size of each cluster (the refinement parameter). If $\mathbf{1}_S = \hat{\mathbf{x}}$, i.e. S is chosen by independent sampling with probabilities according to \mathbf{x} , then for large n the fractions $\frac{1}{n} |S \cap (N \times \{i\})|$ will be strongly concentrated around their expectation. As \hat{G} is continuous, we get $\lim_{n \rightarrow \infty} \mathbf{E}[\hat{g}(\hat{\mathbf{x}})] = \lim_{n \rightarrow \infty} \mathbf{E}[\hat{G}(\tilde{\mathbf{x}})] = \hat{G}(\mathbf{E}[\tilde{\mathbf{x}}]) = \hat{G}(\bar{\mathbf{x}})$. Here, $\bar{\mathbf{x}}$ is the vector \mathbf{x} projected back to the original ground set X , where the coordinates of each cluster have been

averaged. By construction of the refinement, if $\mathbf{x} \in P(\tilde{\mathcal{F}})$ then $\bar{\mathbf{x}}$ is in the polytope corresponding to the original instance, $P(\mathcal{F})$. Therefore, $\hat{G}(\bar{\mathbf{x}}) \leq \max\{\hat{G}(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\} \leq \gamma OPT + \epsilon$. For large enough n , this means that $\max\{\hat{G}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\} \leq \gamma OPT + 2\epsilon$. This holds for an arbitrarily small fixed $\epsilon > 0$, and hence the gap between the instances $\max\{\tilde{F}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ and $\max\{\hat{G}(\mathbf{x}) : \mathbf{x} \in P(\tilde{\mathcal{F}})\}$ (which cannot be distinguished) can be made arbitrarily close to γ . \square

Hardness for symmetric instances We remark that since Lemma 3.2 provides functions \hat{F} and \hat{G} symmetric under \mathcal{G} , the refined instances that we define are invariant with respect to the following symmetries: permute the copies of each element in an arbitrary way, and permute the classes of copies according to any permutation $\sigma \in \mathcal{G}$. This means that our hardness results also hold for instances satisfying such symmetry properties.

It remains to prove Lemma 3.2. Before we move to the final construction of $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$, we construct as an intermediate step a function $\tilde{F}(\mathbf{x})$ which is helpful in the analysis.

Construction Let us construct a function $\tilde{F}(\mathbf{x})$ which satisfies the following:

- For \mathbf{x} “sufficiently close” to $\bar{\mathbf{x}}$, $\tilde{F}(\mathbf{x}) = G(\mathbf{x})$.
- For \mathbf{x} “sufficiently far away” from $\bar{\mathbf{x}}$, $\tilde{F}(\mathbf{x}) \simeq F(\mathbf{x})$.
- The function $\tilde{F}(\mathbf{x})$ is “approximately” smooth submodular.

Once we have $\tilde{F}(\mathbf{x})$, we can fix it to obtain a smooth submodular function $\hat{F}(\mathbf{x})$, which is still close to the original function $F(\mathbf{x})$. We also fix $G(\mathbf{x})$ in the same way, to obtain a function $\hat{G}(\mathbf{x})$ which is equal to $\hat{F}(\mathbf{x})$ whenever \mathbf{x} is close to $\bar{\mathbf{x}}$. We defer this step until the end.

We define $\tilde{F}(\mathbf{x})$ as a convex linear combination of $F(\mathbf{x})$ and $G(\mathbf{x})$, guided by a “smooth transition” function, depending on the distance of \mathbf{x} from $\bar{\mathbf{x}}$. The form that we use is the following:⁴

$$\tilde{F}(\mathbf{x}) = (1 - \phi(D(\mathbf{x})))F(\mathbf{x}) + \phi(D(\mathbf{x}))G(\mathbf{x})$$

where $\phi : R_+ \rightarrow [0, 1]$ is a suitable smooth function, and

$$D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 = \sum_i (x_i - \bar{x}_i)^2.$$

The idea is that when \mathbf{x} is close to $\bar{\mathbf{x}}$, $\phi(D(\mathbf{x}))$ should be close to 1, i.e. the convex linear combination should give most of the weight to $G(\mathbf{x})$. The weight should shift gradually to $F(\mathbf{x})$ as \mathbf{x} gets further away from $\bar{\mathbf{x}}$. Therefore, we define $\phi(t) = 1$ in a small interval $t \in [0, \delta]$, and $\phi(t)$ tends to 0 as t increases. This guarantees that $\tilde{F}(\mathbf{x}) = G(\mathbf{x})$ whenever $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 \leq \delta$. We defer the precise construction of $\phi(t)$ to Lemma 3.7, after we determine what properties we need from $\phi(t)$. Note that regardless of the definition of $\phi(t)$, $\tilde{F}(\mathbf{x})$ is symmetric with respect to \mathcal{G} , since $F(\mathbf{x})$, $G(\mathbf{x})$ and $D(\mathbf{x})$ are.

Analysis of the construction Due to the construction of $\tilde{F}(\mathbf{x})$, it is clear that when $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$ is small, $\tilde{F}(\mathbf{x}) = G(\mathbf{x})$. When $D(\mathbf{x})$ is large, $\tilde{F}(\mathbf{x}) \simeq F(\mathbf{x})$. The main issue, however, is whether we can say something about the first and second partial derivatives of \tilde{F} . This is crucial for the properties of monotonicity and submodularity, which we would like to preserve. Let us write $\tilde{F}(\mathbf{x})$ as

$$\tilde{F}(\mathbf{x}) = F(\mathbf{x}) - \phi(D(\mathbf{x}))H(\mathbf{x})$$

⁴We remark that a construction analogous to [22] would be $\tilde{F}(\mathbf{x}) = F(\mathbf{x}) - \phi(H(\mathbf{x}))$ where $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$. While this makes the analysis easier in [22], it cannot be used in general. Roughly speaking, the problem is that in general the partial derivatives of $H(\mathbf{x})$ are not bounded in any way by the value of $H(\mathbf{x})$.

where $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$. By differentiating once, we get

$$\frac{\partial \tilde{F}}{\partial x_i} = \frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} - \phi'(D(\mathbf{x})) \frac{\partial D}{\partial x_i} H(\mathbf{x}) \quad (2)$$

and by differentiating twice,

$$\begin{aligned} \frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} &= \frac{\partial^2 F}{\partial x_i \partial x_j} - \phi(D(\mathbf{x})) \frac{\partial^2 H}{\partial x_i \partial x_j} - \phi''(D(\mathbf{x})) \frac{\partial D}{\partial x_i} \frac{\partial D}{\partial x_j} H(\mathbf{x}) \\ &\quad - \phi'(D(\mathbf{x})) \left(\frac{\partial D}{\partial x_j} \frac{\partial H}{\partial x_i} + \frac{\partial^2 D}{\partial x_i \partial x_j} H(\mathbf{x}) + \frac{\partial D}{\partial x_i} \frac{\partial H}{\partial x_j} \right). \end{aligned} \quad (3)$$

The first two terms on the right-hand sides of (2) and (3) are not bothering us, because they form convex linear combinations of the derivatives of $F(\mathbf{x})$ and $G(\mathbf{x})$, which have the properties that we need. The remaining terms might cause problems, however, and we need to estimate them.

Our strategy is to define $\phi(t)$ in such a way that it eliminates the influence of partial derivatives of D and H where they become too large. Roughly speaking, D and H have negligible partial derivatives when \mathbf{x} is very close to $\bar{\mathbf{x}}$. As \mathbf{x} moves away from $\bar{\mathbf{x}}$, the partial derivatives grow but then the behavior of $\phi(t)$ must be such that their influence is suppressed.

We start with the following important claim.⁵

Lemma 3.4. *Assume that $F : [0, 1]^X \rightarrow R$ is differentiable and invariant under a group of permutations of coordinates \mathcal{G} . Let $G(\mathbf{x}) = F(\bar{\mathbf{x}})$, where $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$. Then for any $\mathbf{x} \in [0, 1]^X$,*

$$\nabla G|_{\mathbf{x}} = \nabla F|_{\bar{\mathbf{x}}}.$$

Proof. To avoid confusion, we use \mathbf{x} for the arguments of the functions F and G , and \mathbf{u} , $\bar{\mathbf{u}}$, etc. for points where their partial derivatives are evaluated. To rephrase, we want to prove that for any point \mathbf{u} and any coordinate i , the partial derivatives of F and G evaluated at $\bar{\mathbf{u}}$ are equal: $\frac{\partial G}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \frac{\partial F}{\partial x_i} \Big|_{\mathbf{x}=\bar{\mathbf{u}}}$.

First, consider $F(\mathbf{x})$. We assume that $F(\mathbf{x})$ is invariant under a group of permutations of coordinates \mathcal{G} , i.e. $F(\mathbf{x}) = F(\sigma(\mathbf{x}))$ for any $\sigma \in \mathcal{G}$. Differentiating both sides at $\mathbf{x} = \mathbf{u}$, we get by the chain rule:

$$\frac{\partial F}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \sum_j \frac{\partial F}{\partial x_j} \Big|_{\mathbf{x}=\sigma(\mathbf{u})} \frac{\partial}{\partial x_i} (\sigma(\mathbf{x}))_j = \sum_j \frac{\partial F}{\partial x_j} \Big|_{\mathbf{x}=\sigma(\mathbf{u})} \frac{\partial x_{\sigma(j)}}{\partial x_i}.$$

Here, $\frac{\partial x_{\sigma(j)}}{\partial x_i} = 1$ if $\sigma(j) = i$, and 0 otherwise. Therefore,

$$\frac{\partial F}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \frac{\partial F}{\partial x_{\sigma^{-1}(i)}} \Big|_{\mathbf{x}=\sigma(\mathbf{u})}.$$

Now, if we evaluate the left-hand side at $\bar{\mathbf{u}}$, the right-hand side is evaluated at $\sigma(\bar{\mathbf{u}}) = \bar{\mathbf{u}}$, and hence for any i and any $\sigma \in \mathcal{G}$,

$$\frac{\partial F}{\partial x_i} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} = \frac{\partial F}{\partial x_{\sigma^{-1}(i)}} \Big|_{\mathbf{x}=\bar{\mathbf{u}}}. \quad (4)$$

Turning to $G(\mathbf{x}) = F(\bar{\mathbf{x}})$, let us write $\frac{\partial G}{\partial x_i}$ using the chain rule:

$$\frac{\partial G}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \frac{\partial}{\partial x_i} F(\bar{\mathbf{x}}) \Big|_{\mathbf{x}=\mathbf{u}} = \sum_j \frac{\partial F}{\partial x_j} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} \cdot \frac{\partial \bar{x}_j}{\partial x_i}.$$

We have $\bar{x}_j = \mathbf{E}_{\sigma \in \mathcal{G}}[x_{\sigma(j)}]$, and so

$$\frac{\partial G}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \sum_j \frac{\partial F}{\partial x_j} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} \cdot \frac{\partial}{\partial x_i} \mathbf{E}_{\sigma \in \mathcal{G}}[x_{\sigma(j)}] = \mathbf{E}_{\sigma \in \mathcal{G}} \left[\sum_j \frac{\partial F}{\partial x_j} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} \cdot \frac{\partial x_{\sigma(j)}}{\partial x_i} \right].$$

⁵We remind the reader that ∇F , the gradient of F , is a vector whose coordinates are the first partial derivatives $\frac{\partial F}{\partial x_i}$. We denote by $\nabla F|_{\mathbf{x}}$ the gradient evaluated at \mathbf{x} .

Again, $\frac{\partial x_{\sigma(j)}}{\partial x_i} = 1$ if $\sigma(j) = i$ and 0 otherwise. Consequently, we obtain

$$\frac{\partial G}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} = \mathbf{E}_{\sigma \in \mathcal{G}} \left[\frac{\partial F}{\partial x_{\sigma^{-1}(i)}} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} \right] = \frac{\partial F}{\partial x_i} \Big|_{\mathbf{x}=\bar{\mathbf{u}}}$$

where we used Eq. (4) to remove the dependence on $\sigma \in \mathcal{G}$. \square

Observe that the symmetrization operation $\bar{\mathbf{x}}$ is idempotent, i.e. $\bar{\bar{\mathbf{x}}} = \bar{\mathbf{x}}$. Because of this, we also get $\nabla G|_{\bar{\mathbf{x}}} = \nabla F|_{\bar{\mathbf{x}}}$. Note that $G(\bar{\mathbf{x}}) = F(\bar{\mathbf{x}})$ follows from the definition, but it is not obvious that the same holds for gradients, since their definition involves points where $G(\mathbf{x}) \neq F(\mathbf{x})$. For second partial derivatives, the equality no longer holds, as can be seen from a simple example such as $F(x_1, x_2) = 1 - (1 - x_1)(1 - x_2)$, $G(x_1, x_2) = 1 - (1 - \frac{x_1+x_2}{2})^2$.

Next, we show that the functions $F(\mathbf{x})$ and $G(\mathbf{x})$ are very similar in the close vicinity of the region where $\bar{\mathbf{x}} = \mathbf{x}$. Recall our definitions: $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$, $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$. Based on Lemma 3.4, we know that $H(\bar{\mathbf{x}}) = 0$ and $\nabla H|_{\bar{\mathbf{x}}} = 0$. In the following lemmas, we present bounds on $H(\mathbf{x})$, $D(\mathbf{x})$ and their partial derivatives.

Lemma 3.5. *Let $f : 2^X \rightarrow [0, M]$ be invariant under a permutation group \mathcal{G} . Let $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$, $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$ and $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$ where $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ and $G(\mathbf{x}) = F(\bar{\mathbf{x}})$. Then*

1. $|\frac{\partial^2 H}{\partial x_i \partial x_j}| \leq 8M$ everywhere, for all i, j ;
2. $\|\nabla H(\mathbf{x})\| \leq 8M|X|\sqrt{D(\mathbf{x})}$;
3. $|H(\mathbf{x})| \leq 8M|X| \cdot D(\mathbf{x})$.

Proof. First, let us get a bound on the second partial derivatives. Assuming without loss of generality $x_i = x_j = 0$, we have⁶

$$\frac{\partial^2 F}{\partial x_i \partial x_j} = \mathbf{E}[f(\hat{\mathbf{x}} \vee (\mathbf{e}_i + \mathbf{e}_j)) - f(\hat{\mathbf{x}} \vee \mathbf{e}_i) - f(\hat{\mathbf{x}} \vee \mathbf{e}_j) + f(\hat{\mathbf{x}})]$$

(see [30]). Consequently,

$$\left| \frac{\partial^2 F}{\partial x_i \partial x_j} \right| \leq 4 \max |f(S)| = 4M.$$

It is a little bit more involved to analyze $\frac{\partial^2 G}{\partial x_i \partial x_j}$. Since $G(\mathbf{x}) = F(\bar{\mathbf{x}})$ and $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$, we get by the chain rule:

$$\frac{\partial^2 G}{\partial x_i \partial x_j} = \sum_{k, \ell} \frac{\partial^2 F}{\partial x_k \partial x_\ell} \frac{\partial \bar{x}_k}{\partial x_i} \frac{\partial \bar{x}_\ell}{\partial x_j} = \mathbf{E}_{\sigma, \tau \in \mathcal{G}} \left[\sum_{k, \ell} \frac{\partial^2 F}{\partial x_k \partial x_\ell} \frac{\partial x_{\sigma(k)}}{\partial x_i} \frac{\partial x_{\tau(\ell)}}{\partial x_j} \right].$$

It is useful here to use the Kronecker symbol, $\delta_{i,j}$, which is 1 if $i = j$ and 0 otherwise. Note that $\frac{\partial x_{\sigma(k)}}{\partial x_i} = \delta_{i, \sigma(k)} = \delta_{\sigma^{-1}(i), k}$, etc. Using this notation, we get

$$\begin{aligned} \frac{\partial^2 G}{\partial x_i \partial x_j} &= \mathbf{E}_{\sigma, \tau \in \mathcal{G}} \left[\sum_{k, \ell} \frac{\partial^2 F}{\partial x_k \partial x_\ell} \delta_{\sigma^{-1}(i), k} \delta_{\sigma^{-1}(j), \ell} \right] = \mathbf{E}_{\sigma, \tau \in \mathcal{G}} \left[\frac{\partial^2 F}{\partial x_{\sigma^{-1}(i)} \partial x_{\tau^{-1}(j)}} \right], \\ \left| \frac{\partial^2 G}{\partial x_i \partial x_j} \right| &\leq \mathbf{E}_{\sigma, \tau \in \mathcal{G}} \left[\left| \frac{\partial^2 F}{\partial x_{\sigma^{-1}(i)} \partial x_{\tau^{-1}(j)}} \right| \right] \leq 4M \end{aligned}$$

and therefore

$$\left| \frac{\partial^2 H}{\partial x_i \partial x_j} \right| = \left| \frac{\partial^2 F}{\partial x_i \partial x_j} - \frac{\partial^2 G}{\partial x_i \partial x_j} \right| \leq 8M.$$

⁶ $\mathbf{x} \vee \mathbf{y}$ denotes the coordinate-wise maximum, $(\mathbf{x} \vee \mathbf{y})_i = \max\{x_i, y_i\}$ and $\mathbf{x} \wedge \mathbf{y}$ denotes the coordinate-wise minimum, $(\mathbf{x} \wedge \mathbf{y})_i = \min\{x_i, y_i\}$.

Next, we estimate $\frac{\partial H}{\partial x_i}$ at a given point \mathbf{u} , depending on its distance from $\bar{\mathbf{u}}$. Consider the line segment between $\bar{\mathbf{u}}$ and \mathbf{u} . The function $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$ is C_∞ -differentiable, and hence we can apply the mean value theorem to $\frac{\partial H}{\partial x_i}$: There exists a point $\tilde{\mathbf{u}}$ on the line segment $[\bar{\mathbf{u}}, \mathbf{u}]$ such that

$$\frac{\partial H}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} - \frac{\partial H}{\partial x_i} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} = \sum_j \frac{\partial^2 H}{\partial x_j \partial x_i} \Big|_{\mathbf{x}=\tilde{\mathbf{u}}} (u_j - \bar{u}_j).$$

Recall that $\frac{\partial H}{\partial x_i} \Big|_{\mathbf{x}=\bar{\mathbf{u}}} = 0$. Applying the Cauchy-Schwartz inequality to the right-hand side, we get

$$\left(\frac{\partial H}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} \right)^2 \leq \sum_j \left(\frac{\partial^2 H}{\partial x_j \partial x_i} \Big|_{\mathbf{x}=\tilde{\mathbf{u}}} \right)^2 \|\mathbf{u} - \bar{\mathbf{u}}\|^2 \leq (8M)^2 |X| \|\mathbf{u} - \bar{\mathbf{u}}\|^2.$$

Adding up over all $i \in X$, we obtain

$$\|\nabla H(\mathbf{u})\|^2 = \sum_i \left(\frac{\partial H}{\partial x_i} \Big|_{\mathbf{x}=\mathbf{u}} \right)^2 \leq (8M|X|)^2 \|\mathbf{u} - \bar{\mathbf{u}}\|^2.$$

Finally, we estimate the growth of $H(\mathbf{u})$. Again, by the mean value theorem, there is a point $\tilde{\mathbf{u}}$ on the line segment $[\bar{\mathbf{u}}, \mathbf{u}]$, such that

$$H(\mathbf{u}) - H(\bar{\mathbf{u}}) = (\mathbf{u} - \bar{\mathbf{u}}) \cdot \nabla H(\tilde{\mathbf{u}}).$$

Using $H(\bar{\mathbf{u}}) = 0$, the Cauchy-Schwartz inequality and the above bound on ∇H ,

$$(H(\mathbf{u}))^2 \leq \|\nabla H(\tilde{\mathbf{u}})\|^2 \|\mathbf{u} - \bar{\mathbf{u}}\|^2 \leq (8M|X|)^2 \|\tilde{\mathbf{u}} - \bar{\mathbf{u}}\|^2 \|\mathbf{u} - \bar{\mathbf{u}}\|^2.$$

Clearly, $\|\tilde{\mathbf{u}} - \bar{\mathbf{u}}\| \leq \|\mathbf{u} - \bar{\mathbf{u}}\|$, and therefore

$$|H(\mathbf{u})| \leq 8M|X| \cdot \|\mathbf{u} - \bar{\mathbf{u}}\|^2.$$

□

Lemma 3.6. For the function $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$, we have

1. $\nabla D = 2(\mathbf{x} - \bar{\mathbf{x}})$, and therefore $\|\nabla D\| = 2\sqrt{D(\mathbf{x})}$.
2. For all i, j , $|\frac{\partial^2 D}{\partial x_i \partial x_j}| \leq 2$.

Proof. Let us write $D(\mathbf{x})$ as

$$D(\mathbf{x}) = \sum_i (x_i - \bar{x}_i)^2 = \sum_i \mathbf{E}_{\sigma \in \mathcal{G}} [x_i - x_{\sigma(i)}] \mathbf{E}_{\tau \in \mathcal{G}} [x_i - x_{\tau(i)}].$$

Taking the first partial derivative,

$$\frac{\partial D}{\partial x_j} = 2 \sum_i \mathbf{E}_{\sigma \in \mathcal{G}} [x_i - x_{\sigma(i)}] \frac{\partial}{\partial x_j} \mathbf{E}_{\tau \in \mathcal{G}} [x_i - x_{\tau(i)}].$$

As before, we have $\frac{\partial x_i}{\partial x_j} = \delta_{ij}$. Using this notation, we get

$$\begin{aligned} \frac{\partial D}{\partial x_j} &= 2 \sum_i \mathbf{E}_{\sigma \in \mathcal{G}} [x_i - x_{\sigma(i)}] \mathbf{E}_{\tau \in \mathcal{G}} [\delta_{ij} - \delta_{\tau(i), j}] \\ &= 2 \sum_i \mathbf{E}_{\sigma, \tau \in \mathcal{G}} [(x_i - x_{\sigma(i)}) (\delta_{ij} - \delta_{i, \tau^{-1}(j)})] \\ &= 2 \mathbf{E}_{\sigma, \tau \in \mathcal{G}} [x_j - x_{\sigma(j)} - x_{\tau^{-1}(j)} + x_{\sigma(\tau^{-1}(j))}]. \end{aligned}$$

Since the distributions of $\sigma(j)$, $\tau^{-1}(j)$ and $\sigma(\tau^{-1}(j))$ are the same, we obtain

$$\frac{\partial D}{\partial x_j} = 2 \mathbf{E}_{\sigma \in \mathcal{G}} [x_j - x_{\sigma(j)}] = 2(x_j - \bar{x}_j)$$

and

$$\|\nabla D\|^2 = \sum_j \left| \frac{\partial D}{\partial x_j} \right|^2 = 4 \sum_j (x_j - \bar{x}_j)^2 = 4D(\mathbf{x}).$$

Finally, the second partial derivatives are

$$\frac{\partial^2 D}{\partial x_i \partial x_j} = 2 \frac{\partial}{\partial x_i} (x_j - \bar{x}_j) = 2 \frac{\partial}{\partial x_i} \mathbf{E}_{\sigma \in \mathcal{G}} [x_j - x_{\sigma(j)}] = 2 \mathbf{E}_{\sigma \in \mathcal{G}} [\delta_{ij} - \delta_{i, \sigma(j)}]$$

which is clearly bounded by 2 in the absolute value. \square

Now we come back to $\tilde{F}(\mathbf{x})$ and its partial derivatives. Recall equations (2) and (3). The problematic terms are those involving $\phi'(D(\mathbf{x}))$ and $\phi''(D(\mathbf{x}))$. Using our bounds on $H(\mathbf{x})$, $D(\mathbf{x})$ and their derivatives, however, we notice that $\phi'(D(\mathbf{x}))$ always appears with factors on the order of $D(\mathbf{x})$ and $\phi''(D(\mathbf{x}))$ appears with factors on the order of $(D(\mathbf{x}))^2$. Thus, it is sufficient if $\phi(t)$ is defined so that we have control over $t\phi'(t)$ and $t^2\phi''(t)$. The following lemma describes the function that we need.

Lemma 3.7. *For any $\alpha, \beta > 0$, there is $\delta \in (0, \beta)$ and a function $\phi : R_+ \rightarrow [0, 1]$ with an absolutely continuous first derivative such that*

1. For $t \leq \delta$, $\phi(t) = 1$.
2. For $t \geq \beta$, $\phi(t) < e^{-1/\alpha}$.
3. For all $t \geq 0$, $|t\phi'(t)| \leq 4\alpha$.
4. For almost all $t \geq 0$, $|t^2\phi''(t)| \leq 10\alpha$.

Proof. First, observe that if we prove the lemma for some particular value $\beta > 0$, we can also prove it for any other value $\tilde{\beta} > 0$, by modifying the function as follows: $\tilde{\phi}(t) = \phi(\beta t / \beta')$. This corresponds to a scaling of the parameter t by β' / β . Observe that then $|t\tilde{\phi}'(t)| = |\frac{\beta}{\beta'} t \phi'(\beta t / \beta')| \leq 4\alpha$ and $|t^2\tilde{\phi}''(t)| = |(\frac{\beta}{\beta'})^2 t^2 \phi''(\beta t / \beta')| \leq 10\alpha$, so the conditions are still satisfied.

Therefore, we can assume without loss of generality that $\beta > 0$ is a value of our choice, for example $\beta = e^{1/(2\alpha^2)} + 1$. If we want to prove the result for a different value of β , we can just scale the argument t and the constant δ by $\beta / (e^{1/(2\alpha^2)} + 1)$; the bounds on $t\phi'(t)$ and $t^2\phi''(t)$ still hold.

We can assume that $\alpha \in (0, \frac{1}{8})$ because for larger α , the statement only gets weaker. As we argued, we can assume WLOG that $\beta = e^{1/(2\alpha^2)} + 1$. We set $\delta = 1$ and $\delta_2 = 1 + (1 + \alpha)^{-1/2} \leq 2$. (We remind the reader that in general these values will be scaled depending on the actual value of β .) We define the function as follows:

1. $\phi(t) = 1$ for $t \in [0, \delta]$.
2. $\phi(t) = 1 - \alpha(t - 1)^2$ for $t \in [\delta, \delta_2]$.
3. $\phi(t) = (1 + \alpha)^{-1-\alpha}(t - 1)^{-2\alpha}$ for $t \in [\delta_2, \infty)$.

Let's verify the properties of $\phi(t)$. For $t \in [0, \delta]$, we have $\phi'(t) = \phi''(t) = 0$. For $t \in [\delta, \delta_2]$, we have

$$\phi'(t) = -2\alpha(t - 1), \quad \phi''(t) = -2\alpha,$$

and for $t \in [\delta_2, \infty)$,

$$\begin{aligned} \phi'(t) &= -2\alpha(1 + \alpha)^{-1-\alpha} (t - 1)^{-2\alpha-1}, \\ \phi''(t) &= 2\alpha(1 + 2\alpha)(1 + \alpha)^{-1-\alpha} (t - 1)^{-2\alpha-2}. \end{aligned}$$

First, we check that the values and first derivatives agree at the breakpoints. For $t = \delta = 1$, we get $\phi(1) = 1$ and $\phi'(1) = 0$. For $t = \delta_2 = 1 + (1 + \alpha)^{-1/2}$, we get $\phi(\delta_2) = (1 + \alpha)^{-1}$ and $\phi'(\delta_2) = -2\alpha(1 + \alpha)^{-1/2}$. Next, we need to check is that $\phi(t)$ is very small for $t \geq \beta$. The function is decreasing for $t > \beta$, therefore it is enough to check $t = \beta = e^{1/(2\alpha^2)} + 1$:

$$\phi(\beta) = (1 + \alpha)^{-1-\alpha}(\beta - 1)^{-2\alpha} \leq (\beta - 1)^{-2\alpha} = e^{-1/\alpha}.$$

The derivative bounds are satisfied trivially for $t \in [0, \delta]$. For $t \in [\delta, \delta_2]$, using $t \leq \delta_2 = 1 + (1 + \alpha)^{-1/2}$,

$$|t\phi'(t)| = t \cdot 2\alpha(t - 1) \leq 2\alpha(1 + (1 + \alpha)^{-1/2})(1 + \alpha)^{-1/2} \leq 4\alpha$$

and using $\alpha \in (0, \frac{1}{8})$,

$$|t^2\phi''(t)| = t^2 \cdot 2\alpha \leq 2\alpha(1 + (1 + \alpha)^{-1/2})^2 \leq 8\alpha.$$

For $t \in [\delta_2, \infty)$, using $t - 1 \geq (1 + \alpha)^{-1/2}$,

$$\begin{aligned} |t\phi'(t)| &= t \cdot \frac{2\alpha}{(1 + \alpha)^{1+\alpha}} (t - 1)^{-2\alpha-1} = \frac{2\alpha}{1 + \alpha} \left((1 + \alpha)(t - 1)^2 \right)^{-\alpha} \frac{t}{t - 1} \\ &\leq \frac{2\alpha}{1 + \alpha} \cdot \frac{t}{t - 1} \leq \frac{2\alpha}{1 + \alpha} \cdot \frac{1 + (1 + \alpha)^{-1/2}}{(1 + \alpha)^{-1/2}} = 2\alpha \cdot \frac{1 + (1 + \alpha)^{-1/2}}{(1 + \alpha)^{1/2}} \leq 4\alpha \end{aligned}$$

and finally, using $\alpha \in (0, \frac{1}{8})$,

$$\begin{aligned} |t^2\phi''(t)| &= t^2 \cdot \frac{2\alpha(1 + 2\alpha)}{(1 + \alpha)^{1+\alpha}} (t - 1)^{-2\alpha-2} = \frac{2\alpha(1 + 2\alpha)}{1 + \alpha} \left((1 + \alpha)(t - 1)^2 \right)^{-\alpha} \left(\frac{t}{t - 1} \right)^2 \\ &\leq \frac{2\alpha(1 + 2\alpha)}{1 + \alpha} \cdot \left(\frac{1 + (1 + \alpha)^{-1/2}}{(1 + \alpha)^{-1/2}} \right)^2 \leq 8\alpha(1 + 2\alpha) \leq 10\alpha. \end{aligned}$$

□

Using Lemmas 3.5, 3.6 and 3.7, we now prove bounds on the derivatives of $\tilde{F}(\mathbf{x})$.

Lemma 3.8. *Let $\tilde{F}(\mathbf{x}) = (1 - \phi(D(\mathbf{x})))F(\mathbf{x}) + \phi(D(\mathbf{x}))G(\mathbf{x})$ where $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$, $f : 2^X \rightarrow [0, M]$, $G(\mathbf{x}) = F(\bar{\mathbf{x}})$, $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2$ are as above and $\phi(t)$ is provided by Lemma 3.7 for a given $\alpha > 0$. Then, whenever $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$,*

$$\frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} \leq 512M|X|\alpha.$$

If, in addition, $\frac{\partial F}{\partial x_i} \geq 0$, then

$$\frac{\partial \tilde{F}}{\partial x_i} \geq -64M|X|\alpha.$$

Proof. We have $\tilde{F}(\mathbf{x}) = F(\mathbf{x}) - \phi(D(\mathbf{x}))H(\mathbf{x})$ where $H(\mathbf{x}) = F(\mathbf{x}) - G(\mathbf{x})$. By differentiating once, we get

$$\frac{\partial \tilde{F}}{\partial x_i} = \frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} - \phi'(D(\mathbf{x})) \frac{\partial D}{\partial x_i} H(\mathbf{x}),$$

i.e.

$$\left| \frac{\partial \tilde{F}}{\partial x_i} - \left(\frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} \right) \right| = \left| \phi'(D(\mathbf{x})) \frac{\partial D}{\partial x_i} H(\mathbf{x}) \right|.$$

By Lemma 3.5 and 3.6, we have $|\frac{\partial D}{\partial x_i}| = 2|x_i - \bar{x}_i| \leq 2$ and $|H(\mathbf{x})| \leq 8M|X| \cdot D(\mathbf{x})$. Therefore,

$$\left| \frac{\partial \tilde{F}}{\partial x_i} - \left(\frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} \right) \right| \leq 16M|X|D(\mathbf{x}) \cdot \left| \phi'(D(\mathbf{x})) \right|.$$

By Lemma 3.7, $|D(\mathbf{x})\phi'(D(\mathbf{x}))| \leq 4\alpha$, and hence

$$\left| \frac{\partial \tilde{F}}{\partial x_i} - \left(\frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} \right) \right| \leq 64M|X|\alpha.$$

Assuming that $\frac{\partial F}{\partial x_i} \geq 0$, we also have $\frac{\partial G}{\partial x_i} \geq 0$ (see Lemma 3.4) and therefore, $\frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} = (1 - \phi(D(\mathbf{x}))) \frac{\partial F}{\partial x_i} + \phi(D(\mathbf{x})) \frac{\partial G}{\partial x_i} \geq 0$. Consequently,

$$\frac{\partial \tilde{F}}{\partial x_i} \geq -64M|X|\alpha.$$

By differentiating \tilde{F} twice, we obtain

$$\begin{aligned} \frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} &= \frac{\partial^2 F}{\partial x_i \partial x_j} - \phi(D(\mathbf{x})) \frac{\partial^2 H}{\partial x_i \partial x_j} - \phi''(D(\mathbf{x})) \frac{\partial D}{\partial x_i} \frac{\partial D}{\partial x_j} H(\mathbf{x}) \\ &\quad - \phi'(D(\mathbf{x})) \left(\frac{\partial D}{\partial x_j} \frac{\partial H}{\partial x_i} + \frac{\partial^2 D}{\partial x_i \partial x_j} H(\mathbf{x}) + \frac{\partial D}{\partial x_i} \frac{\partial H}{\partial x_j} \right). \end{aligned}$$

Again, we use Lemma 3.5 and 3.6 to bound $|H(\mathbf{x})| \leq 8M|X|D(\mathbf{x})$, $|\frac{\partial H}{\partial x_i}| \leq 8M|X|\sqrt{D(\mathbf{x})}$, $|\frac{\partial^2 H}{\partial x_i \partial x_j}| \leq 8M$, $|\frac{\partial D}{\partial x_i}| \leq 2\sqrt{D(\mathbf{x})}$ and $|\frac{\partial^2 D}{\partial x_i \partial x_j}| \leq 2$. We get

$$\left| \frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} - \frac{\partial^2 F}{\partial x_i \partial x_j} + \phi(D(\mathbf{x})) \frac{\partial^2 H}{\partial x_i \partial x_j} \right| \leq 32M|X| \left| D^2(\mathbf{x})\phi''(D(\mathbf{x})) \right| + 48M|X| \left| D(\mathbf{x})\phi'(D(\mathbf{x})) \right|$$

Observe that $\phi'(D(\mathbf{x}))$ appears with $D(\mathbf{x})$ and $\phi''(D(\mathbf{x}))$ appears with $(D(\mathbf{x}))^2$. By Lemma 3.7, $|D(\mathbf{x})\phi'(D(\mathbf{x}))| \leq 4\alpha$ and $|D^2(\mathbf{x})\phi''(D(\mathbf{x}))| \leq 10\alpha$. Therefore,

$$\left| \frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} - \left(\frac{\partial^2 F}{\partial x_i \partial x_j} - \phi(D(\mathbf{x})) \frac{\partial^2 H}{\partial x_i \partial x_j} \right) \right| \leq 320M|X|\alpha + 192M|X|\alpha = 512M|X|\alpha.$$

If $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$ for all i, j , then also $\frac{\partial^2 G}{\partial x_i \partial x_j} = \mathbf{E}_{\sigma, \tau \in \mathcal{G}} \left[\frac{\partial^2 F}{\partial x_{\sigma^{-1}(i)} \partial x_{\tau^{-1}(j)}} \right] \leq 0$ (see the proof of Lemma 3.5). Also, $\frac{\partial^2 F}{\partial x_i \partial x_j} - \phi(D(\mathbf{x})) \frac{\partial^2 H}{\partial x_i \partial x_j} = \phi(D(\mathbf{x})) \frac{\partial^2 F}{\partial x_i \partial x_j} + (1 - \phi(D(\mathbf{x}))) \frac{\partial^2 G}{\partial x_i \partial x_j} \leq 0$. We obtain

$$\frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} \leq 512M|X|\alpha.$$

□

Finally, we can finish the proof of Lemma 3.2.

Proof of Lemma 3.2. Let $\epsilon > 0$ and $f : 2^X \rightarrow [0, M]$. We choose $\beta = \frac{\epsilon}{16M|X|}$, so that $|H(\mathbf{x})| = |F(\mathbf{x}) - G(\mathbf{x})| \leq \epsilon/2$ whenever $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 \leq \beta$ (due to by Lemma 3.5, which states that $|H(\mathbf{x})| \leq 8M|X|D(\mathbf{x})$). Also, let $\alpha = \frac{\epsilon}{2000M|X|^3}$. For these values of $\alpha, \beta > 0$, let $\delta > 0$ and $\phi : R_+ \rightarrow [0, 1]$ be provided by Lemma 3.7. We define

$$\tilde{F}(\mathbf{x}) = (1 - \phi(D(\mathbf{x})))F(\mathbf{x}) + \phi(D(\mathbf{x}))G(\mathbf{x}).$$

Lemma 3.8 provides bounds on the first and second partial derivatives of $\tilde{F}(\mathbf{x})$. Finally, we modify $\tilde{F}(\mathbf{x})$ so that it satisfies the required conditions (submodularity and optionally monotonicity). For that purpose, we add a suitable multiple of the following function:

$$J(\mathbf{x}) = |X|^2 + 3|X| \sum_{i \in X} x_i - \left(\sum_{i \in X} x_i \right)^2.$$

We have $0 \leq J(\mathbf{x}) \leq 3|X|^2$, $\frac{\partial J}{\partial x_i} = 3|X| - 2 \sum_{i \in X} x_i \geq |X|$. Further, $\frac{\partial^2 J}{\partial x_i \partial x_j} = -2$. Note also that $J(\bar{\mathbf{x}}) = J(\mathbf{x})$, since $J(\mathbf{x})$ depends only on the sum of all coordinates $\sum_{i \in X} x_i$. To make $\tilde{F}(\mathbf{x})$ submodular and optionally monotone, we define:

$$\hat{F}(\mathbf{x}) = \tilde{F}(\mathbf{x}) + 256M|X|\alpha J(\mathbf{x}),$$

$$\hat{G}(\mathbf{x}) = G(\mathbf{x}) + 256M|X|\alpha J(\mathbf{x}).$$

We verify the properties of $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$:

1. For any $\mathbf{x} \in P(\mathcal{F})$, we have

$$\begin{aligned} \hat{G}(\mathbf{x}) &= G(\mathbf{x}) + 256M|X|\alpha J(\mathbf{x}) \\ &= F(\bar{\mathbf{x}}) + 256M|X|\alpha J(\bar{\mathbf{x}}) \\ &= \hat{F}(\bar{\mathbf{x}}). \end{aligned}$$

2. When $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 \geq \beta$, Lemma 3.7 guarantees that $0 \leq \phi(D(\mathbf{x})) < e^{-1/\alpha} \leq \alpha$ and

$$\begin{aligned} |\hat{F}(\mathbf{x}) - F(\mathbf{x})| &\leq \phi(D(\mathbf{x}))|G(\mathbf{x}) - F(\mathbf{x})| + 256M|X|\alpha J(\mathbf{x}) \\ &\leq \alpha M + 768M|X|^3\alpha \\ &\leq \epsilon \end{aligned}$$

using $0 \leq F(\mathbf{x}), G(\mathbf{x}) \leq M$, $|X| \leq J(\mathbf{x}) \leq 3|X|^2$ and $\alpha = \frac{\epsilon}{2000M|X|^3}$.

When $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 < \beta$, we chose the value of β so that $|G(\mathbf{x}) - F(\mathbf{x})| < \epsilon/2$ and so by the above,

$$\begin{aligned} |\hat{F}(\mathbf{x}) - F(\mathbf{x})| &\leq \phi(D(\mathbf{x}))|G(\mathbf{x}) - F(\mathbf{x})| + 256M|X|\alpha J(\mathbf{x}) \\ &\leq \epsilon/2 + 768M|X|^3\alpha \\ &\leq \epsilon. \end{aligned}$$

3. Due to Lemma 3.7, $\phi(t) = 1$ for $t \in [0, \delta]$. Hence, whenever $D(\mathbf{x}) = \|\mathbf{x} - \bar{\mathbf{x}}\|^2 \leq \delta$, we have $\tilde{F}(\mathbf{x}) = G(\mathbf{x}) = F(\bar{\mathbf{x}})$, which depends only on $\bar{\mathbf{x}}$. Also, we have $\hat{F}(\mathbf{x}) = \hat{G}(\mathbf{x}) = F(\bar{\mathbf{x}}) + 256M|X|\alpha J(\mathbf{x})$ and again, $J(\mathbf{x})$ depends only on $\bar{\mathbf{x}}$ (in fact, only on the average of all coordinates of \mathbf{x}). Therefore, $\hat{F}(\mathbf{x})$ and $\hat{G}(\mathbf{x})$ in this case depend only on $\bar{\mathbf{x}}$.

4. The first partial derivatives of \hat{F} are given by the formula

$$\frac{\partial \hat{F}}{\partial x_i} = \frac{\partial F}{\partial x_i} - \phi(D(\mathbf{x})) \frac{\partial H}{\partial x_i} - \phi'(D(\mathbf{x})) \frac{\partial D}{\partial x_i} H(\mathbf{x}) + 256M|X|\alpha \frac{\partial J}{\partial x_i}.$$

The functions F, H, D, J are infinitely differentiable, so the only possible issue is with ϕ . By inspecting our construction of ϕ (Lemma 3.7), we can see that it is piecewise infinitely differentiable, and ϕ' is continuous at the breakpoints. Therefore, it is also absolutely continuous. This implies that $\frac{\partial \hat{F}}{\partial x_i}$ is absolutely continuous.

The function $\hat{G}(\mathbf{x}) = F(\bar{\mathbf{x}}) + 256M|X|\alpha J(\mathbf{x})$ is infinitely differentiable, so its first partial derivatives are also absolutely continuous.

5. Assuming $\frac{\partial F}{\partial x_i} \geq 0$, we get $\frac{\partial \hat{F}}{\partial x_i} \geq -64M|X|\alpha$ by Lemma 3.8. Using $\frac{\partial J}{\partial x_i} \geq |X|$, we get $\frac{\partial \hat{F}}{\partial x_i} = \frac{\partial \tilde{F}}{\partial x_i} + 256M|X|\alpha J(\mathbf{x}) \geq 0$. The same holds for $\frac{\partial \hat{G}}{\partial x_i}$ since $\frac{\partial G}{\partial x_i} \geq 0$.

6. Assuming $\frac{\partial^2 F}{\partial x_i \partial x_j} \leq 0$, we get $\frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} \leq 512M|X|\alpha$ by Lemma 3.8. Using $\frac{\partial^2 J}{\partial x_i \partial x_j} = -2$, we get $\frac{\partial^2 \hat{F}}{\partial x_i \partial x_j} = \frac{\partial^2 \tilde{F}}{\partial x_i \partial x_j} + 256M|X|\alpha \frac{\partial^2 J}{\partial x_i \partial x_j} \leq 0$. The same holds for $\frac{\partial^2 \hat{G}}{\partial x_i \partial x_j}$ since $\frac{\partial^2 G}{\partial x_i \partial x_j} \leq 0$.

□

This concludes the proofs of our main hardness results.

4 Algorithms using the multilinear relaxation

Here we turn to our algorithmic results. First, we discuss the problem of maximizing a submodular (but not necessarily monotone) function subject to a matroid independence constraint.

4.1 Matroid independence constraint

Consider the problem $\max\{f(S) : S \in \mathcal{I}\}$, where \mathcal{I} is the collection of independent sets in a matroid \mathcal{M} . We design an algorithm based on the multilinear relaxation of the problem, $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{M})\}$. Our algorithm can be seen as "continuous local search" in the matroid polytope $P(\mathcal{M})$, constrained in addition by the box $[0, t]^X$ for some fixed $t \in [0, 1]$. The intuition is that this forces our local search to use fractional solutions that are more fuzzy than integral solutions and therefore less likely to get stuck in a local optimum. On the other hand, restraining the search space too much would not give us much freedom in searching for a good fractional point. This leads to a tradeoff and an optimal choice of $t \in [0, 1]$ which we leave for later.

The matroid polytope is defined as $P(\mathcal{M}) = \text{conv}\{\mathbf{1}_I : I \in \mathcal{I}\}$, or equivalently [10] as $P(\mathcal{M}) = \{\mathbf{x} \geq 0 : \forall S; \sum_{i \in S} x_i \leq r_{\mathcal{M}}(S)\}$, where $r_{\mathcal{M}}(S)$ is the rank function of \mathcal{M} . We define

$$P_t(\mathcal{M}) = P(\mathcal{M}) \cap [0, t]^X = \{\mathbf{x} \in P(\mathcal{M}) : \forall i; x_i \leq t\}.$$

We consider the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P_t(\mathcal{M})\}$. We remind the reader that $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$ denotes the multilinear extension. Our algorithm works as follows.

Fractional local search in $P_t(\mathcal{M})$

(given $t = \frac{r}{q}$, $r \leq q$ integer)

1. Start with $\mathbf{x} := (0, 0, \dots, 0)$. Fix $\delta = 1/q$.
2. If there is $i, j \in X$ and a direction $\mathbf{v} \in \{\mathbf{e}_j, -\mathbf{e}_i, \mathbf{e}_j - \mathbf{e}_i\}$ such that $\mathbf{x} + \delta\mathbf{v} \in P_t(\mathcal{M})$ and $F(\mathbf{x} + \delta\mathbf{v}) > F(\mathbf{x})$, set $\mathbf{x} := \mathbf{x} + \delta\mathbf{v}$ and repeat.
3. If there is no such direction \mathbf{v} , apply pipage rounding to \mathbf{x} and return the resulting solution.

Notes. The procedure as presented here would not run in polynomial time. A modification which runs in polynomial time is that we move to a new solution only if $F(\mathbf{x} + \delta\mathbf{v}) > F(\mathbf{x}) + \frac{\delta}{\text{poly}(n)}OPT$ (where we first get a rough estimate of OPT using previous methods). For simplicity, we analyze the variant above and finally discuss why we can modify it without losing too much in the approximation factor. We also defer the question of how to estimate the value of $F(\mathbf{x})$ to the end of this section.

For $t = 1$, we have $\delta = 1$ and the procedure reduces to discrete local search. However, it is known that discrete local search alone does not give any approximation guarantee. With additional modifications, an algorithm based on discrete local search achieves a $(\frac{1}{4} - o(1))$ -approximation [19].

Our version of fractional local search avoids this issue and leads directly to a good fractional solution. Throughout the algorithm, we maintain \mathbf{x} as a linear combination of q independent sets such that no element appears in more than r of them. A local step corresponds to an add/remove/switch operation preserving this condition.

Finally, we use pipage rounding to convert a fractional solution into an integral one. As we show in Lemma B.3 (in the Appendix), a modification of the technique from [4] can be used to find an integral solution without any loss in the objective function.

Theorem 4.1. *The fractional local search algorithm for any fixed $t \in [0, \frac{1}{2}(3 - \sqrt{5})]$ returns a solution of value at least $(t - \frac{1}{2}t^2)OPT$, where $OPT = \max\{f(S) : S \in \mathcal{I}\}$.*

We remark that for $t = \frac{1}{2}(3 - \sqrt{5})$, we would obtain a $\frac{1}{4}(-1 + \sqrt{5}) \simeq 0.309$ -approximation, improving the factor of $\frac{1}{4}$ [19]. This is not a rational value, but we can pick a rational t arbitrarily close to $\frac{1}{2}(3 - \sqrt{5})$. For values $t > \frac{1}{2}(3 - \sqrt{5})$, our analysis does not yield a better approximation factor.

First, we discuss properties of the point found by the fractional local search algorithm.

Lemma 4.2. *The outcome of the fractional local search algorithm x is a "fractional local optimum" in the following sense. (All the partial derivatives are evaluated at x .)*

- For any i such that $\mathbf{x} - \delta \mathbf{e}_i \in P_t(\mathcal{M})$, $\frac{\partial F}{\partial x_i} \geq 0$.
- For any j such that $\mathbf{x} + \delta \mathbf{e}_j \in P_t(\mathcal{M})$, $\frac{\partial F}{\partial x_j} \leq 0$.
- For any i, j such that $\mathbf{x} + \delta(\mathbf{e}_j - \mathbf{e}_i) \in P_t(\mathcal{M})$, $\frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0$.

Proof. We use the property (see [4]) that along any direction $\mathbf{v} = \pm \mathbf{e}_i$ or $\mathbf{v} = \mathbf{e}_i - \mathbf{e}_j$, the function $F(\mathbf{x} + \lambda \mathbf{v})$ is a convex function of λ . Also, observe that if it is possible to move from \mathbf{x} in the direction of \mathbf{v} by any nonzero amount, then it is possible to move by $\delta \mathbf{v}$, because all coordinates of \mathbf{x} are integer multiples of δ and all the constraints also have coefficients which are integer multiples of δ . Therefore, if $\frac{dF}{d\lambda} > 0$ and it is possible to move in the direction of \mathbf{v} , we would get $F(\mathbf{x} + \delta \mathbf{v}) > F(\mathbf{x})$ and the fractional local search would continue.

If $\mathbf{v} = -\mathbf{e}_i$ and it is possible to move along $-\mathbf{e}_i$, we get $\frac{dF}{d\lambda} = -\frac{\partial F}{\partial x_i} \leq 0$. Similarly, if $\mathbf{v} = \mathbf{e}_j$ and it is possible to move along \mathbf{e}_j , we get $\frac{dF}{d\lambda} = \frac{\partial F}{\partial x_j} \leq 0$. Finally, if $\mathbf{v} = \mathbf{e}_j - \mathbf{e}_i$ and it is possible to move along $\mathbf{e}_j - \mathbf{e}_i$, we get $\frac{dF}{d\lambda} = \frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0$. \square

We refer to the following exchange property for matroids (which follows easily from [28], Corollary 39.12a; see also [19]).

Lemma 4.3. *If $I, C \in \mathcal{I}$, then for any $j \in C \setminus I$, there is $\pi(j) \subseteq I \setminus C$, $|\pi(j)| \leq 1$, such that $I \setminus \pi(j) + j \in \mathcal{I}$. Moreover, the sets $\pi(j)$ are disjoint (each $i \in I \setminus C$ appears at most once as $\pi(j) = \{i\}$).*

Using this, we prove a lemma about fractional local optima which generalizes Lemma 2.2 in [19].

Lemma 4.4. *Let \mathbf{x} be the outcome of fractional local search over $P_t(\mathcal{M})$. Let $C \in \mathcal{I}$ be any independent set. Let $C' = \{i \in C : x_i < t\}$. Then*

$$2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{1}_{C'}) + F(\mathbf{x} \wedge \mathbf{1}_C).$$

Note that for $t = 1$, the lemma reduces to $2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{1}_C) + F(\mathbf{x} \wedge \mathbf{1}_C)$ (similar to Lemma 2.2 in [19]). For $t < 1$, however, it is necessary to replace C by C' in the first expression, which becomes apparent in the proof. The reason is that we do not have any information on $\frac{\partial F}{\partial x_i}$ for coordinates where $x_i = t$.

Proof. Let $C \in \mathcal{I}$ and assume $\mathbf{x} \in P_t(\mathcal{M})$ is a local optimum. Since $\mathbf{x} \in P(\mathcal{M})$, we can decompose it into a convex linear combination of vertices of $P(\mathcal{M})$, $\mathbf{x} = \sum_{I \in \mathcal{I}} x_I \mathbf{1}_I$ where $\sum x_I = 1$. By the smooth submodularity of $F(\mathbf{x})$ (see [30]),

$$F(\mathbf{x} \vee \mathbf{1}_{C'}) - F(\mathbf{x}) \leq \sum_{j \in C'} (1 - x_j) \frac{\partial F}{\partial x_j} = \sum_{j \in C'} \sum_{I: j \notin I} x_I \frac{\partial F}{\partial x_j} = \sum_I x_I \sum_{j \in C' \setminus I} \frac{\partial F}{\partial x_j}.$$

All partial derivatives here are evaluated at \mathbf{x} . On the other hand, also by submodularity,

$$F(\mathbf{x}) - F(\mathbf{x} \wedge \mathbf{1}_C) \geq \sum_{i \notin C} x_i \frac{\partial F}{\partial x_i} = \sum_{i \notin C} \sum_{I: i \in I} x_I \frac{\partial F}{\partial x_i} = \sum_I x_I \sum_{i \in I \setminus C} \frac{\partial F}{\partial x_i}.$$

To prove the lemma, it remains to prove the following.

Claim. Whenever $x_I > 0$, $\sum_{j \in C' \setminus I} \frac{\partial F}{\partial x_j} \leq \sum_{i \in I \setminus C} \frac{\partial F}{\partial x_i}$.

Proof: For any $I \in \mathcal{I}$, we can apply Lemma 4.3 to get a mapping π such that $I \setminus \pi(j) + j \in \mathcal{I}$ for any $j \in C \setminus I$. Now, consider $j \in C' \setminus I$, i.e. $j \in C \setminus I$ and $x_j < t$.

If $\pi(j) = \emptyset$, is possible to move from \mathbf{x} in the direction of \mathbf{e}_j , because $I + j \in \mathcal{I}$ and hence we can replace I by $I + j$ (or at least we can do this for some nonzero fraction of its coefficient) in the linear combination. Because $x_j < t$, we can move by a nonzero amount inside $P_t(\mathcal{M})$. By Lemma 4.2, $\frac{\partial F}{\partial x_j} \leq 0$.

Similarly, if $\pi(j) = \{i\}$, it is possible to move in the direction of $\mathbf{e}_j - \mathbf{e}_i$, because I can be replaced by $I \setminus \pi(j) + i$ for some nonzero fraction of its coefficient. By Lemma 4.2, in this case $\frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0$.

Finally, for any $i \in I$ we have $x_i > 0$ and therefore we can decrease x_i while staying inside $P_t(\mathcal{M})$. By Lemma 4.2, we have $\frac{\partial F}{\partial x_i} \geq 0$ for all $i \in I$. This means

$$\sum_{j \in C' \setminus I} \frac{\partial F}{\partial x_j} \leq \sum_{j \in C' \setminus I: \pi(j) = \emptyset} \frac{\partial F}{\partial x_j} + \sum_{j \in C' \setminus I: \pi(j) = \{i\}} \frac{\partial F}{\partial x_i} \leq \sum_{i \in I \setminus C} \frac{\partial F}{\partial x_i}$$

using the inequalities we derived above, and the fact that each $i \in I \setminus C$ appears at most once in $\pi(j)$. This proves the Claim, and hence the Lemma. \square

Now we are ready to prove Theorem 4.1.

Proof. Let \mathbf{x} be the outcome of the fractional local search over $P_t(\mathcal{M})$. Define $A = \{i : x_i = t\}$. Let C be the optimum solution and $C' = C \setminus A = \{i \in C : x_i < t\}$. By Lemma 4.4,

$$2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{1}_{C'}) + F(\mathbf{x} \wedge \mathbf{1}_C).$$

First, let's analyze $F(\mathbf{x} \wedge \mathbf{1}_C)$. We apply Lemma A.4 (in the Appendix), which states that $F(\mathbf{x} \wedge \mathbf{1}_C) \geq \mathbf{E}[f(T(\mathbf{x} \wedge \mathbf{1}_C))]$. Here, $T(\mathbf{x} \wedge \mathbf{1}_C)$ is a random threshold set corresponding to the vector $\mathbf{x} \wedge \mathbf{1}_C$, i.e.

$$T(\mathbf{x} \wedge \mathbf{1}_C) = \{i : (\mathbf{x} \wedge \mathbf{1}_C)_i > \lambda\} = \{i \in C : x_i > \lambda\} = T(\mathbf{x}) \cap C$$

where $\lambda \in [0, 1]$ is uniformly random. Equivalently,

$$F(\mathbf{x} \wedge \mathbf{1}_C) \geq \mathbf{E}[f(T(\mathbf{x}) \cap C)].$$

Due to the definition of a threshold set, with probability t we have $\lambda < t$ and $T(\mathbf{x})$ contains $A = \{i : x_i = t\} = C \setminus C'$. Then, $f(T(\mathbf{x}) \cap C) + f(C') \geq f(C)$ by submodularity. We conclude that

$$F(\mathbf{x} \wedge \mathbf{1}_C) \geq t(f(C) - f(C')). \quad (5)$$

Next, let's analyze $F(\mathbf{x} \vee \mathbf{1}_{C'})$. We consider the ground set partitioned into $X = C \cup \bar{C}$, and we apply Lemma A.5 (in the Appendix). (Here, \bar{C} denotes $X \setminus C$, the complement of C inside X .) We get

$$F(\mathbf{x} \vee \mathbf{1}_{C'}) \geq \mathbf{E}[f((T_1(\mathbf{x} \vee \mathbf{1}_{C'}) \cap C) \cup (T_2(\mathbf{x} \vee \mathbf{1}_{C'}) \cap \bar{C}))].$$

The random threshold sets look as follows: $T_1(\mathbf{x} \vee \mathbf{1}_{C'}) \cap C = (T_1(\mathbf{x}) \cup C') \cap C$ is equal to C with probability t , and equal to C' otherwise. $T_2(\mathbf{x} \vee \mathbf{1}_{C'}) \cap \bar{C} = T_2(\mathbf{x}) \cap \bar{C}$ is empty with probability $1 - t$. (We ignore the contribution when $T_2(\mathbf{x}) \cap \bar{C} \neq \emptyset$.) Because T_1 and T_2 are independently sampled, we get

$$F(\mathbf{x} \vee \mathbf{1}_{C'}) \geq t(1 - t)f(C) + (1 - t)^2 f(C').$$

Provided that $t \in [0, \frac{1}{2}(3 - \sqrt{5})]$, we have $t \leq (1 - t)^2$. Then, we can write

$$F(\mathbf{x} \vee \mathbf{1}_{C'}) \geq t(1 - t)f(C) + tf(C'). \quad (6)$$

Combining equations (5) and (6), we get

$$F(\mathbf{x} \vee \mathbf{1}_{C'}) + F(\mathbf{x} \wedge \mathbf{1}_C) \geq t(f(C) - f(C')) + t(1 - t)f(C) + tf(C') = (2t - t^2)f(C).$$

Therefore,

$$F(\mathbf{x}) \geq \frac{1}{2}(F(\mathbf{x} \vee \mathbf{1}_{C'}) + F(\mathbf{x} \wedge \mathbf{1}_C)) \geq (t - \frac{1}{2}t^2)f(C).$$

Finally, we apply the pipage rounding technique which does not lose anything in terms of objective value (see Lemma B.3). \square

Technical remarks In each step of the algorithm, we need to estimate values of $F(\mathbf{x})$ for given $\mathbf{x} \in P_t(\mathcal{M})$. We accomplish this by using the expression $F(\mathbf{x}) = \mathbf{E}[f(R(\mathbf{x}))]$ where $R(\mathbf{x})$ is a random set associated with \mathbf{x} . By standard bounds, if the values of $f(S)$ are in a range $[0, M]$, we can achieve accuracy $M/\text{poly}(n)$ using a polynomial number of samples. We use the fact that $OPT \geq \frac{1}{n}M$ (see Lemma C.1 in the Appendix) and therefore we can achieve $OPT/\text{poly}(n)$ additive error in polynomial time.

We also relax the local step condition: we move to the next solution only if $F(\mathbf{x} + \delta\mathbf{v}) > F(\mathbf{x}) + \frac{\delta}{\text{poly}(n)}OPT$ for a suitable polynomial in n . This way, we can only make a polynomial number of steps. When we terminate, the local optimality conditions (Lemma 4.2) are satisfied within an additive error of $OPT/\text{poly}(n)$, which yields a polynomially small error in the approximation bound.

4.2 Matroid base constraint

Let us move on to the problem $\max\{f(S) : S \in \mathcal{B}\}$ where \mathcal{B} are the bases of a matroid. For a fixed $t \in [0, 1]$, let us consider an algorithm which can be seen as local search inside the base polytope $B(\mathcal{M})$, further constrained by the box $[0, t]^X$. The matroid base polytope is defined as $B(\mathcal{M}) = \text{conv}\{\mathbf{1}_B : B \in \mathcal{B}\}$ or equivalently [10] as $B(\mathcal{M}) = \{\mathbf{x} \geq 0 : \forall S \subseteq X; \sum_{i \in S} x_i \leq r_{\mathcal{M}}(S), \sum_{i \in X} x_i = r_{\mathcal{M}}(X)\}$, where $r_{\mathcal{M}}$ is the matroid rank function of \mathcal{M} . Finally, we define

$$B_t(\mathcal{M}) = B(\mathcal{M}) \cap [0, t]^X = \{\mathbf{x} \in B(\mathcal{M}) : \forall i \in X; x_i \leq t\}.$$

Observe that $B_t(\mathcal{M})$ is nonempty if and only if there is a convex linear combination $\mathbf{x} = \sum_{B \in \mathcal{B}} \xi_B \mathbf{1}_B$ such that $x_i \in [0, t]$ for all i . This is equivalent to saying that there is a linear combination (a fractional base packing) $\mathbf{x}' = \sum_{B \in \mathcal{B}} \xi'_B \mathbf{1}_B$ such that $x_i \in [0, 1]$ and $\sum \xi'_B \geq \frac{1}{t}$, in other words the fractional base packing number is $\nu \geq \frac{1}{t}$. Since the optimal fractional packing of bases in a matroid can be found efficiently (see Corollary 42.7a in [28], Volume B), we can find efficiently the minimum $t \in [\frac{1}{2}, 1]$ such that $B_t(\mathcal{M}) \neq \emptyset$. Then, our algorithm is the following.

Fractional local search in $B_t(\mathcal{M})$

(given $t = \frac{r}{q}$, $r \leq q$ integer)

1. Let $\delta = \frac{1}{q}$. Assume that $\mathbf{x} \in B_t(\mathcal{M})$; adjust \mathbf{x} (using pipage rounding) so that each x_i is an integer multiple of δ . In the following, this property will be maintained.
2. If there is a direction $\mathbf{v} = \mathbf{e}_j - \mathbf{e}_i$ such that $\mathbf{x} + \delta\mathbf{v} \in B_t(\mathcal{M})$ and $F(\mathbf{x} + \delta\mathbf{v}) > F(\mathbf{x})$, then set $\mathbf{x} := \mathbf{x} + \delta\mathbf{v}$ and repeat.
3. If there is no such direction \mathbf{v} , apply pipage rounding to \mathbf{x} and return the resulting solution.

Notes. We remark that the starting point can be found as a convex linear combination of q bases, $x = \frac{1}{q} \sum_{i=1}^q \mathbf{1}_{B_i}$, such that no element appears in more than r of them, using matroid union techniques (see Theorem 42.9 in [28]). In the algorithm, we maintain this representation. The local search step corresponds to switching a pair of elements in one base, under the condition that no element is used in more than r bases at the same time. For now, we ignore the issues of estimating $F(\mathbf{x})$ and stopping the local search within polynomial time. We discuss this at the end of this section.

Finally, we use pipage rounding to convert the fractional solution \mathbf{x} into an integral one of value at least $F(\mathbf{x})$ (Lemma B.1 in the Appendix). Note that it is not necessarily true that any of the bases in a convex linear combination $\mathbf{x} = \sum \xi_B \mathbf{1}_B$ achieves the value $F(\mathbf{x})$.

Theorem 4.5. *If there is a fractional packing of $\nu \in [1, 2]$ bases in \mathcal{M} , then the fractional local search algorithm with $t = \frac{1}{\nu}$ returns a solution of value at least $\frac{1}{2}(1 - t) OPT$.*

For example, assume that \mathcal{M} contains two disjoint bases B_1, B_2 (which is the case considered in [19]). Then, the algorithm can be used with $t = \frac{1}{2}$ and we obtain a $(\frac{1}{4} - o(1))$ -approximation, improving the $(\frac{1}{6} - o(1))$ -approximation from [19]. If there is a fractional packing of more than 2 bases, our analysis still gives only a $(\frac{1}{4} - o(1))$ -approximation. If the dual matroid \mathcal{M}^* admits a better fractional packing of bases, we can consider the problem $\max\{f(\bar{S}) : S \in \mathcal{B}^*\}$ which is equivalent. For a uniform matroid, $\mathcal{B} = \{B : |B| = k\}$, the fractional base packing number is either at least 2 or the same holds for the dual

matroid, $\mathcal{B}^* = \{B : |B| = n - k\}$ (as noted in [19]). Therefore, we get a $(\frac{1}{4} - o(1))$ -approximation for any uniform matroid. The value $t = 1$ can be used for any matroid, but it does not yield any approximation guarantee.

Analysis of the algorithm We turn to the properties of fractional local optima. We will prove that the point \mathbf{x} found by the fractional local search algorithm satisfies the following conditions that allow us to compare $F(\mathbf{x})$ to the actual optimum.

Lemma 4.6. *The outcome of the fractional local search algorithm \mathbf{x} is a “fractional local optimum” in the following sense.*

- For any i, j such that $\mathbf{x} + \delta(\mathbf{e}_j - \mathbf{e}_i) \in B_t(\mathcal{M})$,

$$\frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0.$$

(The partial derivatives are evaluated at \mathbf{x} .)

Proof. Similarly to Lemma 4.2, observe that the coordinates of \mathbf{x} are always integer multiples of δ , therefore if it is possible to move from \mathbf{x} in the direction of $\mathbf{v} = \mathbf{e}_j - \mathbf{e}_i$ by any nonzero amount, then it is possible to move by $\delta\mathbf{v}$. We use the property that for any direction $\mathbf{v} = \mathbf{e}_j - \mathbf{e}_i$, the function $F(\mathbf{x} + \lambda\mathbf{v})$ is a convex function of λ [4]. Therefore, if $\frac{dF}{d\lambda} > 0$ and it is possible to move in the direction of \mathbf{v} , we would get $F(\mathbf{x} + \delta\mathbf{v}) > F(\mathbf{x})$ and the fractional local search would continue. For $\mathbf{v} = \mathbf{e}_j - \mathbf{e}_i$, we get

$$\frac{dF}{d\lambda} = \frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0.$$

□

We refer to the following exchange property for matroid bases (see [28], Corollary 39.21a).

Lemma 4.7. *For any $B_1, B_2 \in \mathcal{B}$, there is a bijection $\pi : B_1 \setminus B_2 \rightarrow B_2 \setminus B_1$ such that $\forall i \in B_1 \setminus B_2; B_1 - i + \pi(i) \in \mathcal{M}$.*

Using this, we prove a lemma about fractional local optima analogous to Lemma 2.2 in [19].

Lemma 4.8. *Let \mathbf{x} be the outcome of fractional local search over $B_t(\mathcal{M})$. Let $C \in \mathcal{B}$ be any base. Then there is $\mathbf{c} \in [0, 1]^X$ satisfying*

- $c_i = t$, if $i \in C$ and $x_i = t$
- $c_i = 1$, if $i \in C$ and $x_i < t$
- $0 \leq c_i \leq x_i$, if $i \notin C$

such that

$$2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{c}) + F(\mathbf{x} \wedge \mathbf{c}).$$

Note that for $t = 1$, we can set $\mathbf{c} = \mathbf{1}_C$. However, in general we need this more complicated formulation. Intuitively, \mathbf{c} is obtained from \mathbf{x} by raising the variables $x_i, i \in C$ and decreasing x_i for $i \notin C$. However, we can only raise the variables $x_i, i \in C$, where x_i is below the threshold t , otherwise we do not have any information about $\frac{\partial F}{\partial x_i}$. Also, we do not necessarily decrease all the variables outside of C to zero.

Proof. Let $C \in \mathcal{B}$ and assume $\mathbf{x} \in B_t(\mathcal{M})$ is a fractional local optimum. We can decompose \mathbf{x} into a convex linear combination of vertices of $B(\mathcal{M})$, $\mathbf{x} = \sum \xi_B \mathbf{1}_B$. By Lemma 4.7, for each base B there is a bijection $\pi_B : B \setminus C \rightarrow C \setminus B$ such that $\forall i \in B \setminus C; B - i + \pi_B(i) \in \mathcal{B}$.

We define $C' = \{i \in C : x_i < t\}$. The reason we consider C' is that if $x_i = t$, there is no room for an exchange step increasing x_i , and therefore Lemma 4.6 does not give any information about $\frac{\partial F}{\partial x_i}$. We construct the vector \mathbf{c} by starting from \mathbf{x} , and for each B swapping the elements in $B \setminus C$ for their image under π_B ,

provided it is in C' , until we raise the coordinates on C' to $c_i = 1$. Formally, we set $c_i = 1$ for $i \in C'$, $c_i = t$ for $i \in C \setminus C'$, and for each $i \notin C$, we define

$$c_i = x_i - \sum_{B:i \in B, \pi_B(i) \in C'} \xi_B.$$

In the following, all partial derivatives are evaluated at \mathbf{x} . By the smooth submodularity of $F(\mathbf{x})$ (see [5]),

$$F(\mathbf{x} \vee \mathbf{c}) - F(\mathbf{x}) \leq \sum_{j:c_j > x_j} (c_j - x_j) \frac{\partial F}{\partial x_j} = \sum_{j \in C'} (1 - x_j) \frac{\partial F}{\partial x_j} = \sum_B \sum_{j \in C' \setminus B} \xi_B \frac{\partial F}{\partial x_j} \quad (7)$$

because $\sum_{B:j \notin B} \xi_B = 1 - x_j$ for any j . On the other hand, also by smooth submodularity,

$$F(\mathbf{x}) - F(\mathbf{x} \wedge \mathbf{c}) \geq \sum_{i:c_i < x_i} (x_i - c_i) \frac{\partial F}{\partial x_i} = \sum_{i \notin C} (x_i - c_i) \frac{\partial F}{\partial x_i} = \sum_{i \notin C} \sum_{B:i \in B, \pi_B(i) \in C'} \xi_B \frac{\partial F}{\partial x_i}$$

using our definition of c_i . In the last sum, for any nonzero contribution, we have $\xi_B > 0$, $i \in B$ and $j = \pi_B(i) \in C'$, i.e. $x_j < t$. Therefore it is possible to move in the direction $\mathbf{e}_j - \mathbf{e}_i$ (we can switch from B to $B - i + j$). By Lemma 4.6,

$$\frac{\partial F}{\partial x_j} - \frac{\partial F}{\partial x_i} \leq 0.$$

Therefore, we get

$$F(\mathbf{x}) - F(\mathbf{x} \wedge \mathbf{c}) \geq \sum_{i \notin C} \sum_{B:i \in B, j = \pi_B(i) \in C'} \xi_B \frac{\partial F}{\partial x_j} = \sum_B \sum_{i \in B \setminus C: j = \pi_B(i) \in C'} \xi_B \frac{\partial F}{\partial x_j}. \quad (8)$$

By the bijective property of π_B , this is equal to $\sum_B \sum_{j \in C' \setminus B} \xi_B \frac{\partial F}{\partial x_j}$. Putting (7) and (8) together, we get $F(\mathbf{x} \vee \mathbf{c}) - F(\mathbf{x}) \leq F(\mathbf{x}) - F(\mathbf{x} \wedge \mathbf{c})$. \square

Now we are ready to prove Theorem 4.5.

Proof. Assuming that $B_t(\mathcal{M}) \neq \emptyset$, we can find a starting point $\mathbf{x}_0 \in B_t(\mathcal{M})$. From this point, we reach a fractional local optimum $\mathbf{x} \in B_t(\mathcal{M})$ (see Lemma 4.6). We want to compare $F(\mathbf{x})$ to the actual optimum; assume that $OPT = f(C)$.

As before, we define $C' = \{i \in C : x_i < t\}$. By Lemma 4.8, we know that the fractional local optimum satisfies:

$$2F(\mathbf{x}) \geq F(\mathbf{x} \vee \mathbf{c}) + F(\mathbf{x} \wedge \mathbf{c}) \quad (9)$$

for some vector \mathbf{c} such that $c_i = t$ for all $i \in C \setminus C'$, $c_i = 1$ for $i \in C'$ and $0 \leq c_i \leq x_i$ for $i \notin C'$.

First, let's analyze $F(\mathbf{x} \vee \mathbf{c})$. We have

- $(\mathbf{x} \vee \mathbf{c})_i = 1$ for all $i \in C'$.
- $(\mathbf{x} \vee \mathbf{c})_i = t$ for all $i \in C \setminus C'$.
- $(\mathbf{x} \vee \mathbf{c})_i \leq t$ for all $i \notin C$.

We apply Lemma A.5 to the partition $X = C \cup \bar{C}$. We get

$$F(\mathbf{x} \vee \mathbf{c}) \geq \mathbf{E}[f((T_1(\mathbf{x} \vee \mathbf{c}) \cap C) \cup (T_2(\mathbf{x} \vee \mathbf{c}) \cap \bar{C}))]$$

where $T_1(\mathbf{x})$ and $T_2(\mathbf{x})$ are independent threshold sets. Based on the information above, $T_1(\mathbf{x} \vee \mathbf{c}) \cap C = C$ with probability t and $T_1(\mathbf{x} \vee \mathbf{c}) \cap C = C'$ otherwise. On the other hand, $T_2(\mathbf{x} \vee \mathbf{c}) \cap \bar{C} = \emptyset$ with probability at least $1 - t$. These two events are independent. We conclude that on the right-hand side, we get $f(C)$ with probability at least $t(1 - t)$, or $f(C')$ with probability at least $(1 - t)^2$:

$$F(\mathbf{x} \vee \mathbf{c}) \geq t(1 - t)f(C) + (1 - t)^2f(C'). \quad (10)$$

Turning to $F(\mathbf{x} \wedge \mathbf{c})$, we see that

- $(\mathbf{x} \wedge \mathbf{c})_i = x_i$ for all $i \in C'$.
- $(\mathbf{x} \wedge \mathbf{c})_i = t$ for all $i \in C \setminus C'$.
- $(\mathbf{x} \wedge \mathbf{c})_i \leq t$ for all $i \notin C$.

We apply Lemma A.5 to $X = C \cup \bar{C}$.

$$F(\mathbf{x} \wedge \mathbf{c}) \geq \mathbf{E}[f((T_1(\mathbf{x} \wedge \mathbf{c}) \cap C) \cup (T_2(\mathbf{x} \wedge \mathbf{c}) \cap \bar{C}))].$$

With probability t , $T_1(\mathbf{x} \wedge \mathbf{c}) \cap C$ contains $C \setminus C'$ (and maybe some elements of C'). In this case, $f(T_1(\mathbf{x} \wedge \mathbf{c}) \cap C) \geq f(C) - f(C')$ by submodularity. Also, $T_2(\mathbf{x} \wedge \mathbf{c}) \cap \bar{C}$ is empty with probability at least $1 - t$. Again, these two events are independent. Therefore, $F(\mathbf{x} \wedge \mathbf{c}) \geq t(1 - t)(f(C) - f(C'))$. If $f(C') > f(C)$, this bound is vacuous; otherwise, we can replace $t(1 - t)$ by $(1 - t)^2$, because $t \geq 1/2$. In any case,

$$F(\mathbf{x} \wedge \mathbf{c}) \geq (1 - t)^2(f(C) - f(C')). \quad (11)$$

Combining (9), (10) and (11),

$$F(\mathbf{x}) \geq \frac{1}{2}(F(\mathbf{x} \vee \mathbf{c}) + F(\mathbf{x} \wedge \mathbf{c})) \geq \frac{1}{2}(t(1 - t)f(C) + (1 - t)^2f(C)) = \frac{1}{2}(1 - t)f(C).$$

□

Technical remarks Again, we have to deal with the issues of estimating $F(\mathbf{x})$ and stopping the local search in polynomial time. We do this exactly as we did at the end of Section 4.1. One issue to be careful about here is that if $f : 2^X \rightarrow [0, M]$, our estimates of $F(\mathbf{x})$ are within an additive error of $M/\text{poly}(n)$. If the optimum value $OPT = \max\{f(S) : S \in \mathcal{B}\}$ is very small compared to M , the error might be large compared to OPT which would be a problem. The optimum could in fact be very small in general. But it holds that if \mathcal{M} contains no loops and co-loops (which can be eliminated easily), then $OPT \geq \frac{1}{n^2}M$ (see Appendix C). Then, our sampling errors are on the order of $OPT/\text{poly}(n)$ which yields a $1/\text{poly}(n)$ error in the approximation bound.

5 Approximation for symmetric instances

We can achieve a better approximation assuming that the instance exhibits a certain symmetry. This is the same kind of symmetry that we use in our hardness construction (Section 3) and the hard instances exhibit the same symmetry as well. It turns out that our approximation in this case matches the hardness threshold up to lower order terms.

Similar to our hardness result, the symmetries that we consider here are permutations of the ground set X , corresponding to permutations of coordinates in R^X . We start with some basic properties which are helpful in analyzing symmetric instances.

Lemma 5.1. *Assume that $f : 2^X \rightarrow R$ is invariant with respect to a group of permutations \mathcal{G} and $F(\mathbf{x}) = \mathbf{E}[f(\hat{\mathbf{x}})]$. Then for any symmetrized vector $\bar{\mathbf{c}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{c})]$, $\nabla F|_{\bar{\mathbf{c}}}$ is also symmetric w.r.t. \mathcal{G} . I.e., for any $\tau \in \mathcal{G}$,*

$$\tau(\nabla F|_{\mathbf{x}=\bar{\mathbf{c}}}) = \nabla F|_{\mathbf{x}=\bar{\mathbf{c}}}.$$

Proof. Since $f(S)$ is invariant under \mathcal{G} , so is $F(\mathbf{x})$, i.e. $F(\mathbf{x}) = F(\tau(\mathbf{x}))$ for any $\tau \in \mathcal{G}$. Differentiating both sides at $\mathbf{x} = \mathbf{c}$, we get by the chain rule:

$$\left. \frac{\partial F}{\partial x_i} \right|_{\mathbf{x}=\mathbf{c}} = \sum_j \left. \frac{\partial F}{\partial x_j} \right|_{\mathbf{x}=\tau(\mathbf{c})} \frac{\partial}{\partial x_i} (\tau(x))_j = \sum_j \left. \frac{\partial F}{\partial x_j} \right|_{\mathbf{x}=\tau(\mathbf{c})} \frac{\partial x_{\tau(j)}}{\partial x_i}.$$

Here, $\frac{\partial x_{\tau(j)}}{\partial x_i} = 1$ if $\tau(j) = i$, and 0 otherwise. Therefore,

$$\left. \frac{\partial F}{\partial x_i} \right|_{\mathbf{x}=\mathbf{c}} = \left. \frac{\partial F}{\partial x_{\tau^{-1}(i)}} \right|_{\mathbf{x}=\tau(\mathbf{c})}.$$

Note that $\tau(\bar{\mathbf{c}}) = \mathbf{E}_{\sigma \in \mathcal{G}}[\tau(\sigma(\mathbf{c}))] = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{c})] = \bar{\mathbf{c}}$ since the distribution of $\tau \circ \sigma$ is equal to the distribution of σ . Therefore,

$$\left. \frac{\partial F}{\partial x_i} \right|_{\mathbf{x}=\bar{\mathbf{c}}} = \left. \frac{\partial F}{\partial x_{\tau^{-1}(i)}} \right|_{\mathbf{x}=\bar{\mathbf{c}}}$$

for any $\tau \in \mathcal{G}$. □

Next, we prove that the ‘‘symmetric optimum’’ $\max\{F(\bar{\mathbf{x}}) : \mathbf{x} \in P(\mathcal{F})\}$ gives a solution which is a local optimum for the original instance $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$. (As we proved in Section 3, in general we cannot hope to find a better solution than the symmetric optimum.)

Lemma 5.2. *Let $f : 2^X \rightarrow R$ and $\mathcal{F} \subset 2^X$ be invariant with respect to a group of permutations \mathcal{G} . Let $\overline{OPT} = \max\{F(\bar{\mathbf{x}}) : \mathbf{x} \in P(\mathcal{F})\}$ where $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$, and let \mathbf{x}_0 be the symmetric point where \overline{OPT} is attained ($\bar{\mathbf{x}}_0 = \mathbf{x}_0$). Then \mathbf{x}_0 is a local optimum for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$, in the sense that $(\mathbf{x} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} \leq 0$ for any $\mathbf{x} \in P(\mathcal{F})$.*

Proof. Assume for the sake of contradiction that $(\mathbf{x} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} > 0$ for some $\mathbf{x} \in P(\mathcal{F})$. We use the symmetric properties of f and \mathcal{F} to show that $(\bar{\mathbf{x}} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} > 0$ as well. Recall that $\mathbf{x}_0 = \bar{\mathbf{x}}_0$. We have

$$(\bar{\mathbf{x}} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0}] = \mathbf{E}_{\sigma \in \mathcal{G}}[(\mathbf{x} - \mathbf{x}_0) \cdot \sigma^{-1}(\nabla F|_{\mathbf{x}_0})] = (\mathbf{x} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} > 0$$

using Lemma 5.1. Hence, there would be a direction $\bar{\mathbf{x}} - \mathbf{x}_0$ along which an improvement can be obtained. But then, consider a small $\delta > 0$ such that $\mathbf{x}_1 = \mathbf{x}_0 + \delta(\bar{\mathbf{x}} - \mathbf{x}_0) \in P(\mathcal{F})$ and also $F(\mathbf{x}_1) > F(\mathbf{x}_0)$. The point \mathbf{x}_1 is symmetric ($\bar{\mathbf{x}}_1 = \mathbf{x}_1$) and hence it would contradict the assumption that $F(\mathbf{x}_0) = \overline{OPT}$. □

5.1 Submodular maximization over independent sets in a matroid

Let us derive an optimal approximation result for the problem $\max\{f(S) : S \in \mathcal{I}\}$ under the assumption that the instance is ‘‘element-transitive’’.

Definition 5.3. *For a group \mathcal{G} of permutations on X , the orbit of an element $i \in X$ is the set $\{\sigma(i) : \sigma \in \mathcal{G}\}$. \mathcal{G} is called element-transitive, if the orbit of any element is the entire ground set X .*

In this case, we show that it is easy to achieve an optimal $(\frac{1}{2} - o(1))$ -approximation for a matroid independence constraint.

Theorem 5.4. *Let $\max\{f(S) : S \in \mathcal{I}\}$ be an instance symmetric with respect to an element-transitive group of permutations \mathcal{G} . Let $\overline{OPT} = \max\{F(\bar{\mathbf{x}}) : \mathbf{x} \in P(\mathcal{M})\}$ where $\bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$. Then $\overline{OPT} \geq \frac{1}{2}OPT$.*

Proof. Let $OPT = f(C)$. By Lemma 5.2, $\overline{OPT} = F(\mathbf{x}_0)$ where \mathbf{x}_0 is a local optimum for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{M})\}$. This means it is also a local optimum in the sense of Lemma 4.2, with $t = 1$. By Lemma 4.4,

$$2F(\mathbf{x}_0) \geq F(\mathbf{x}_0 \vee \mathbf{1}_C) + F(\mathbf{x}_0 \wedge \mathbf{1}_C).$$

Also, $\mathbf{x}_0 = \bar{\mathbf{x}}_0$. As we are dealing with an element-transitive group of symmetries, this means all the coordinates of \mathbf{x}_0 are equal, $\mathbf{x}_0 = (\xi, \xi, \dots, \xi)$. Therefore, $\mathbf{x}_0 \vee \mathbf{1}_C$ is equal to 1 on C and ξ outside of C . By Lemma A.4 (in the Appendix),

$$F(\mathbf{x}_0 \vee \mathbf{1}_C) \geq (1 - \xi)f(C).$$

Similarly, $\mathbf{x}_0 \wedge \mathbf{1}_C$ is equal to ξ on C and 0 outside of C . By Lemma A.4,

$$F(\mathbf{x}_0 \wedge \mathbf{1}_C) \geq \xi f(C).$$

Combining the two bounds,

$$2F(\mathbf{x}_0) \geq F(\mathbf{x}_0 \vee \mathbf{1}_C) + F(\mathbf{x}_0 \wedge \mathbf{1}_C) \geq (1 - \xi)f(C) + \xi f(C) = f(C) = OPT.$$

□

Since all symmetric solutions $\mathbf{x} = (\xi, \xi, \dots, \xi)$ form a 1-parameter family, and $F(\xi, \xi, \dots, \xi)$ is a concave function, we can search for the best symmetric solution (within any desired accuracy) by binary search. By standard techniques, we get the following.

Corollary 5.5. *There is a $(\frac{1}{2} - o(1))$ -approximation ("brute force" search over symmetric solutions) for the problem $\max\{f(S) : S \in \mathcal{I}\}$ for instances symmetric under an element-transitive group of permutations.*

The hard instances for submodular maximization subject to a matroid independence constraint correspond to refinements of the Max Cut instance for the graph K_2 (Section 2). It is easy to see that such instances are element-transitive, and it follows from Section 3 that a $(\frac{1}{2} + \epsilon)$ -approximation for such instances would require exponentially many value queries. Therefore, our approximation for element-transitive instances is optimal.

5.2 Submodular maximization over bases

Let us come back to the problem of submodular maximization over the bases of matroid. The property that \overline{OPT} is a local optimum with respect to the original problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{F})\}$ is very useful in arguing about the value of \overline{OPT} . We already have tools to deal with local optima from Section 4.2. Here we prove the following.

Lemma 5.6. *Let $B(\mathcal{M})$ be the matroid base polytope of \mathcal{M} and $\mathbf{x}_0 \in B(\mathcal{M})$ a local maximum for the submodular maximization problem $\max\{F(\mathbf{x}) : \mathbf{x} \in B(\mathcal{M})\}$, in the sense that $(\mathbf{x} - \mathbf{x}_0) \cdot \nabla F|_{\mathbf{x}_0} \leq 0$ for any $\mathbf{x} \in B(\mathcal{M})$. Assume in addition that $\mathbf{x}_0 \in [s, t]^X$. Then*

$$F(\mathbf{x}_0) \geq \frac{1}{2}(1 - t + s) \cdot OPT.$$

Proof. Let $OPT = \max\{f(B) : B \in \mathcal{B}\} = f(C)$. We assume that $\mathbf{x}_0 \in B(\mathcal{M})$ is a local optimum with respect to any direction $\mathbf{x} - \mathbf{x}_0$, $\mathbf{x} \in B(\mathcal{M})$, so it is also a local optimum with respect to the fractional local search in the sense of Lemma 4.8, with $t = 1$. The lemma implies that

$$2F(\mathbf{x}_0) \geq F(\mathbf{x} \vee \mathbf{1}_C) + F(\mathbf{x} \wedge \mathbf{1}_C).$$

By assumption, the coordinates of $\mathbf{x} \vee \mathbf{1}_C$ are equal to 1 on C and at most t outside of C . With probability $1 - t$, a random threshold in $[0, 1]$ falls between t and 1, and Lemma A.4 (in the Appendix) implies that

$$F(\mathbf{x} \vee \mathbf{1}_C) \geq (1 - t) \cdot f(C).$$

Similarly, the coordinates of $\mathbf{x} \wedge \mathbf{1}_C$ are 0 outside of C , and at least s on C . A random threshold falls between 0 and s with probability s , and Lemma A.4 implies that

$$F(\mathbf{x} \wedge \mathbf{1}_C) \geq s \cdot f(C).$$

Putting these inequalities together, we get

$$2F(\mathbf{x}_0) \geq F(\mathbf{x} \vee \mathbf{1}_C) + F(\mathbf{x} \wedge \mathbf{1}_C) \geq (1 - t + s) \cdot f(C).$$

□

Totally symmetric instances. The application we have in mind here is a special case of submodular maximization over the bases of a matroid, which we call *totally symmetric*.

Definition 5.7. *We call an instance $\max\{f(S) : S \in \mathcal{F}\}$ totally symmetric with respect to a group of permutations \mathcal{G} , if both $f(S)$ and \mathcal{F} are invariant under \mathcal{G} and moreover, there is a point $\mathbf{c} \in P(\mathcal{F})$ such that $\mathbf{c} = \bar{\mathbf{x}} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{x})]$ for every $\mathbf{x} \in P(\mathcal{F})$. We call \mathbf{c} the center of the instance.*

Note that this is indeed stronger than just being invariant under \mathcal{G} . For example, an instance on a ground set $X = X_1 \cup X_2$ could be symmetric with respect to any permutation of X_1 and any permutation of X_2 . For any $\mathbf{x} \in P(\mathcal{F})$, the symmetric vector $\bar{\mathbf{x}}$ is constant on X_1 and constant on X_2 . However, in a totally symmetric instance, there should be a unique symmetric point.

Bases of partition matroids A canonical example of a totally symmetric instance is as follows. Let $X = X_1 \cup X_2 \cup \dots \cup X_m$ and let integers k_1, \dots, k_m be given. This defines a partition matroid $\mathcal{M} = (X, \mathcal{B})$, whose bases are

$$\mathcal{B} = \{B : \forall j; |B \cap X_j| = k_j\}.$$

The associated matroid base polytope is

$$B(\mathcal{M}) = \{\mathbf{x} \geq 0 : \forall j; \sum_{i \in X_j} x_i = k_j\}.$$

Let \mathcal{G} be a group of permutations such that the orbit of each element $i \in X_j$ is the entire set X_j . This implies that for any $\mathbf{x} \in B(\mathcal{M})$, $\bar{\mathbf{x}}$ is the same vector \mathbf{c} , with coordinates $k_j/|X_j|$ on X_j . If $f(S)$ is also invariant under \mathcal{G} , we have a totally symmetric instance $\max\{f(S) : S \in \mathcal{B}\}$.

Example: welfare maximization To present a more concrete example, consider $X_j = \{a_{j1}, \dots, a_{jk}\}$ for each $j \in [m]$, a set of bases $\mathcal{B} = \{B : \forall j; |B \cap X_j| = 1\}$, and an objective function in the form $f(S) = \sum_{i=1}^k v(\{j : a_{ji} \in S\})$, where $v : 2^{[m]} \rightarrow R_+$ is a submodular function. This is a totally symmetric instance, which captures the welfare maximization problem for combinatorial auctions where each player has the same valuation function. (Including element a_{ji} in the solution corresponds to allocating item j to player i ; see [5] for more details.) We remark that here we consider possibly nonmonotone submodular functions, which is not common for combinatorial auctions; nevertheless the problem still makes sense.

We show that for such instances, the center point achieves an improved approximation.

Theorem 5.8. *Let $\max\{f(S) : S \in \mathcal{B}\}$ be a totally symmetric instance. Let the fractional packing number of bases be ν and the fractional packing number of dual bases ν^* . Then the center point \mathbf{c} satisfies*

$$F(\mathbf{c}) \geq \left(1 - \frac{1}{2\nu} - \frac{1}{2\nu^*}\right) OPT.$$

Recall that in the general case, we get a $\frac{1}{2}(1 - 1/\nu - o(1))$ -approximation (Theorem 4.5). By passing to the dual matroid, we can also obtain a $\frac{1}{2}(1 - 1/\nu^* - o(1))$ -approximation, so in general, we know how to achieve a $\frac{1}{2}(1 - 1/\max\{\nu, \nu^*\} - o(1))$ -approximation. For totally symmetric instances where $\nu = \nu^*$, we improve this to the optimal factor of $1 - 1/\nu$.

Proof. Since there is a unique center $\mathbf{c} = \bar{\mathbf{x}}$ for any $\mathbf{x} \in B(\mathcal{M})$, this means this is also the symmetric optimum $F(\mathbf{c}) = \max\{F(\bar{\mathbf{x}}) : \mathbf{x} \in B(\mathcal{M})\}$. Due to Lemma 5.2, \mathbf{c} is a local optimum for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in B(\mathcal{M})\}$.

Because the fractional packing number of bases is ν , we have $c_i \leq 1/\nu$ for all i . Similarly, because the fractional packing number of dual bases (complements of bases) is ν^* , we have $1 - c_i \leq 1/\nu^*$. This means that $\mathbf{c} \in [1 - 1/\nu^*, 1/\nu]$. Lemma 5.6 implies that

$$2F(\mathbf{c}) \geq \left(1 - \frac{1}{\nu} + 1 - \frac{1}{\nu^*}\right) OPT.$$

□

Corollary 5.9. *Let $\max\{f(S) : S \in \mathcal{B}\}$ be an instance on a partition matroid where every base takes at least an α -fraction of each part, at most a $(1 - \alpha)$ -fraction of each part, and the submodular function $f(S)$ is invariant under a group \mathcal{G} where the orbit of each $i \in X_j$ is X_j . Then, the center point $\mathbf{c} = \mathbf{E}_{\sigma \in \mathcal{G}}[\sigma(\mathbf{1}_B)]$ (equal for any $B \in \mathcal{B}$) satisfies $F(\mathbf{c}) \geq \alpha \cdot OPT$.*

Proof. If the orbit of any element $i \in X_j$ is the entire set X_j , it also means that $\sigma(i)$ for a random $\sigma \in \mathcal{G}$ is uniformly distributed over X_j (by the transitive property of \mathcal{G}). Therefore, symmetrizing any fractional vector $\mathbf{x} \in B(\mathcal{M})$ gives the same vector $\bar{\mathbf{x}} = \mathbf{c}$, where $c_i = k_j/|X_j|$ for $i \in X_j$. Also, our assumptions mean that the fractional packing number of bases is $1/(1 - \alpha)$, and the fractional packing number of dual bases is also $1/(1 - \alpha)$. Due to Lemma 5.8, the center \mathbf{c} satisfies $F(\mathbf{c}) \geq \alpha \cdot OPT$. □

The hard instances for submodular maximization over matroid bases that we describe in Section 2 are exactly of this form (see the last paragraph of Section 2, with $\alpha = 1/k$). There is a unique symmetric solution, $x = (\alpha, \alpha, \dots, \alpha, 1 - \alpha, 1 - \alpha, \dots, 1 - \alpha)$. The fractional base packing number for these matroids is $\nu = 1/(1 - \alpha)$ and Theorem 1.8 implies that any $(\alpha + \epsilon) = (1 - 1/\nu + \epsilon)$ -approximation for such matroids would require exponentially many value queries. Therefore, our approximation in this special case is optimal.

Acknowledgment

The author would like to thank Jon Lee and Maxim Sviridenko for helpful discussions.

References

- [1] N. ALON AND J. SPENCER. The probabilistic method, 3rd edition, John Wiley & Sons, Hoboken, New Jersey, 2008.
- [2] A. AGEEV AND M. SVIRIDENKO. *Pipage rounding: a new method of constructing algorithms with proven performance guarantee*, J. of Combinatorial Optimization 8, 307–328, 2004.
- [3] N. BUCHBINDER, M. FELDMAN, J. NAOR AND R. SCHWARTZ. *A tight linear time (1/2)-approximation for unconstrained submodular maximization*, Proc. of 53th IEEE FOCS, 2012.
- [4] G. CALINESCU, C. CHEKURI, M. PÁL AND J. VONDRÁK. *Maximizing a submodular set function subject to a matroid constraint*, Proc. of 12th IPCO, 182–196, 2007.
- [5] G. CALINESCU, C. CHEKURI, M. PÁL AND J. VONDRÁK. *Maximizing a submodular set function subject to a matroid constraint*, SIAM Journal on Computing 40:6, special section on ACM STOC 2008, 1740–1766, 2011.
- [6] S. DUGHMI, T. ROUGHGARDEN AND Q. YAN. *From convex optimization to randomized mechanisms: towards optimal combinatorial auctions with submodular bidders*, Proc. of 43rd ACM STOC, 149–158, 2011.
- [7] S. DUGHMI AND J. VONDRÁK. *Limitations of randomized mechanisms for combinatorial auctions*, Proc. of 52nd IEEE FOCS, 502–511, 2011.
- [8] S. DOBZINSKI AND J. VONDRÁK. *From query complexity to computational complexity*, Proc. of 44th ACM STOC, 1107–1116, 2012.
- [9] S. DOBZINSKI AND J. VONDRÁK. *The computational complexity of truthfulness in combinatorial auctions*, Proc. of 13th ACM EC, 405–422, 2012.
- [10] J. EDMONDS. *Matroids, submodular functions and certain polyhedra*, Combinatorial Structures and Their Applications, 69–87, 1970.
- [11] J. EDMONDS. *Matroids and the greedy algorithm*, Math. Programming 1, 127–136, 1971.
- [12] U. FEIGE. *A threshold of $\ln n$ for approximating Set Cover*, Journal of the ACM 45, 634–652, 1998.
- [13] U. FEIGE, V. MIRROKNI AND J. VONDRÁK. *Maximizing non-monotone submodular functions*, SIAM Journal on Computing 40:4, 1133–1153, 2011. Conference version in 48th IEEE FOCS, 461–471, 2007.
- [14] L. FLEISCHER, S. FUJISHIGE AND S. IWATA. *A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions*, Journal of the ACM 48:4, 761–777, 2001.
- [15] A. FRANK. *Matroids and submodular functions*, Annotated Bibliographies in Combinatorial Optimization, 65–80, 1997.
- [16] S. FUJISHIGE. *Canonical decompositions of symmetric submodular systems*, Discrete Applied Mathematics 5, 175–190, 1983.

- [17] S. KHOT, R. LIPTON, E. MARKAKIS AND A. MEHTA. *Inapproximability results for combinatorial auctions with submodular utility functions*, Algorithmica 52:1, 3–18, 2008.
- [18] A. KULIK, H. SHACHNAI AND T. TAMIR. *Maximizing submodular set functions subject to multiple linear constraints*, Proc. of 20th ACM-SIAM SODA, 545–554, 2009.
- [19] J. LEE, V. MIRROKNI, V. NAGARAJAN AND M. SVIRIDENKO. *Maximizing non-monotone submodular functions under matroid and knapsack constraints*, Proc. of 41st ACM STOC, 323–332, 2009.
- [20] J. LEE, M. SVIRIDENKO AND J. VONDRÁK. *Submodular maximization over multiple matroids via generalized exchange properties*, Math. of Operations Research 35, 795–806, 2010, conference version in APPROX 2009, 244–257.
- [21] L. LOVÁSZ. *Submodular functions and convexity*, A. Bachem et al., editors, Mathematical Programming: The State of the Art, 235–257, 1983.
- [22] V. MIRROKNI, M. SCHAPIRA AND J. VONDRÁK. *Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions*, Proc. of 9th ACM EC, 70–77, 2008.
- [23] G. L. NEMHAUSER, L. A. WOLSEY AND M. L. FISHER. *An analysis of approximations for maximizing submodular set functions I*, Mathematical Programming 14, 265–294, 1978.
- [24] M. L. FISHER, G. L. NEMHAUSER AND L. A. WOLSEY. *An analysis of approximations for maximizing submodular set functions II*, Mathematical Programming Study 8, 73–87, 1978.
- [25] G. L. NEMHAUSER AND L. A. WOLSEY. *Best algorithms for approximating the maximum of a submodular set function*, Math. Oper. Research 3:3, 177–188, 1978.
- [26] S. OVEIS GHARAN AND J. VONDRÁK. *Submodular maximization by simulated annealing*, Proc. of 22nd ACM-SIAM SODA, 1098–1116, 2011.
- [27] A. SCHRIJVER. *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, Journal of Combinatorial Theory, Series B 80, 346–355, 2000.
- [28] A. SCHRIJVER. *Combinatorial optimization - polyhedra and efficiency*. Springer, 2003.
- [29] M. SVIRIDENKO. *A note on maximizing a submodular set function subject to a knapsack constraint*, Operations Research Letters 32, 41–43, 2004.
- [30] J. VONDRÁK. *Optimal approximation for the submodular welfare problem in the value oracle model*, Proc. of 40th ACM STOC, 67–74, 2008.
- [31] J. VONDRÁK. *Submodularity and curvature: the optimal algorithm*, RIMS Kokyuroku Bessatsu B23, Kyoto, 253–266, 2010.
- [32] J. VONDRÁK. *Symmetry and approximability of submodular maximization problems*, Proc. of 50th IEEE FOCS, 251–270, 2009.
- [33] L. WOLSEY. *An analysis of the greedy algorithm for the submodular set covering problem*, Combinatorica 2, 385–393, 1982.
- [34] L. WOLSEY. *Maximizing real-valued submodular functions: Primal and dual heuristics for location problems*, Math. of Operations Research 7, 410–425, 1982.

A Submodular functions and their extensions

In this appendix, we present a few basic facts concerning submodular functions $f(S)$ and their continuous extensions. By $f_A(S)$, we denote the marginal value of S w.r.t. A , $f_A(S) = f(A \cup S) - f(A)$. We also use $f_A(i) = f(A + i) - f(A)$. The notation $A + i$ is shorthand for $A \cup \{i\}$. Similarly, we write $A - i$ to denote $A \setminus \{i\}$.

Definition A.1. *The multilinear extension of a function $f : 2^X \rightarrow R$ is a function $F : [0, 1]^X \rightarrow R$ where*

$$F(\mathbf{x}) = \sum_{S \subseteq X} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j).$$

Definition A.2. *The Lovász extension of a function $f : 2^X \rightarrow R$ is a function $\tilde{F} : [0, 1]^X \rightarrow R$ such that*

$$\tilde{F}(\mathbf{x}) = \sum_{i=0}^n (x_{\pi(i)} - x_{\pi(i+1)}) f(\{\pi(j) : 1 \leq j \leq i\})$$

where $\pi : [n] \rightarrow X$ is a bijection such that $x_{\pi(1)} \geq x_{\pi(2)} \geq \dots \geq x_{\pi(n)}$; we interpret $x_{\pi(0)}, x_{\pi(n+1)}$ as $x_{\pi(0)} = 1, x_{\pi(n+1)} = 0$.

A useful way to view the multilinear extension is that we sample a random set $R(\mathbf{x})$, where each element i appears independently with probability x_i , and we take $F(\mathbf{x}) = \mathbf{E}[f(R(\mathbf{x}))]$. The Lovász extension can be viewed similarly, by sampling a random set in a correlated fashion.

Definition A.3. *For a vector $\mathbf{x} \in [0, 1]^X$, we define the “random threshold set” $T(\mathbf{x})$ by taking a uniformly random $\lambda \in [0, 1]$, and setting*

$$T(\mathbf{x}) = \{i \in X : x_i > \lambda\}.$$

Assuming that $x_1 \geq x_2 \geq \dots \geq x_n, x_0 = 1, x_n = 0$, it is easy to see that the Lovász extension of $f(S)$ is equal to

$$\tilde{F}(\mathbf{x}) = \sum_{i=0}^n (x_i - x_{i+1}) f([i]) = \mathbf{E}[f(T(\mathbf{x}))].$$

It is known that the Lovász extension of a submodular function is always convex [21], which is not true for the multilinear extension. We prove that the Lovász extension is always upper-bounded by the multilinear extension; this lemma appears quite basic but it has not been published to our knowledge.

Lemma A.4. *Let $F(\mathbf{x})$ denote the multilinear extension and let $\tilde{F}(\mathbf{x})$ denote the Lovász extension of a submodular function $f : 2^X \rightarrow R$. Then*

$$F(\mathbf{x}) \geq \tilde{F}(\mathbf{x}).$$

Proof. Let $R(\mathbf{x})$ be a random set where elements are sampled independently with probabilities x_i , and let $T(\mathbf{x})$ be the random threshold set. I.e., we want to prove $\mathbf{E}[f(R(\mathbf{x}))] \geq \mathbf{E}[f(T(\mathbf{x}))]$. We can assume WLOG that the elements are ordered so that $x_1 \geq x_2 \geq \dots \geq x_n$. We also let $x_0 = 1$ and $x_{n+1} = 0$. Then,

$$\begin{aligned} \mathbf{E}[f(R(\mathbf{x}))] &= f(\emptyset) + \sum_{k=1}^n \mathbf{E}[f(R(\mathbf{x}) \cap [k]) - f(R(\mathbf{x}) \cap [k-1])] \\ &= f(\emptyset) + \sum_{k=1}^n \mathbf{E}[f_{R(\mathbf{x}) \cap [k-1]}(R(\mathbf{x}) \cap \{k\})] \end{aligned}$$

and by submodularity,

$$\begin{aligned}
\mathbf{E}[f(R(\mathbf{x}))] &\geq f(\emptyset) + \sum_{k=1}^n \mathbf{E}[f_{[k-1]}(R(\mathbf{x}) \cap \{k\})] \\
&= f(\emptyset) + \sum_{k=1}^n x_k f_{[k-1]}(k) \\
&= f(\emptyset) + \sum_{k=1}^n x_k (f([k]) - f([k-1])) \\
&= \sum_{k=0}^n (x_k - x_{k+1}) f([k]) \\
&= \mathbf{E}[f(T(\mathbf{x}))].
\end{aligned}$$

□

A refinement of this lemma says that we can also consider a partition of the ground set and apply an independent threshold set on each part. This gives a certain hybrid between the multilinear and Lovász extensions.

Lemma A.5. *For any partition $X = X_1 \cup X_2$,*

$$F(\mathbf{x}) \geq \mathbf{E}[f((T_1(\mathbf{x}) \cap X_1) \cup (T_2(\mathbf{x}) \cap X_2))]$$

where $T_1(\mathbf{x})$ and $T_2(\mathbf{x})$ are two independently random threshold sets for \mathbf{x} .

Intuitively, sampling from Y independently with probability p is not worse than taking all of Y with probability p or none of it with probability $1 - p$. This follows from statements proved in [13], but we give a proof here for completeness.

Proof. $F(\mathbf{x}) = \mathbf{E}[f(R(\mathbf{x}))]$ where $R(\mathbf{x})$ is sampled independently with probabilities x_i . Let's condition on $R(\mathbf{x}) \cap X_2 = R_2$. This restricts the remaining randomness to X_1 , and we get $f(R(\mathbf{x})) = f(R_2) + g(R(\mathbf{x}))$, where $g(S) = f_{R_2}(S \cap X_1)$ is again submodular. By Lemma A.4, $\mathbf{E}[g(R(\mathbf{x}))] \geq \mathbf{E}[g(T_1(\mathbf{x}))]$ where $T_1(\mathbf{x})$ is a random threshold set. Hence,

$$\begin{aligned}
\mathbf{E}[f(R(\mathbf{x})) \mid R(\mathbf{x}) \cap X_2 = R_2] &= f(R_2) + \mathbf{E}[g(R(\mathbf{x}))] \\
&\geq f(R_2) + \mathbf{E}[g(T_1(\mathbf{x}))] = \mathbf{E}[f(R_2 \cup (T_1(\mathbf{x}) \cap X_1))].
\end{aligned}$$

By randomizing $R_2 = R(\mathbf{x}) \cap X_2$, we get

$$\mathbf{E}[f(R(\mathbf{x}))] \geq \mathbf{E}[f((R(\mathbf{x}) \cap X_2) \cup (T_1(\mathbf{x}) \cap X_1))].$$

Repeating the same process, conditioning on $T_1(\mathbf{x})$ and applying Lemma A.4 to $R(\mathbf{x}) \cap X_2$, we get

$$\mathbf{E}[f(R(\mathbf{x}))] \geq \mathbf{E}[f((T_2(\mathbf{x}) \cap X_2) \cup (T_1(\mathbf{x}) \cap X_1))].$$

□

B Pipage rounding for nonmonotone submodular functions

The pipage rounding technique [2, 4, 5] starts with a point in the base polytope $\mathbf{y} \in B(\mathcal{M})$ and produces an integral solution $S \in \mathcal{I}$ (in fact, a base) of expected value $\mathbf{E}[f(S)] \geq F(\mathbf{y})$. We recall the procedure here, in its randomized form [5].

Subroutine **HitConstraint**(\mathbf{y}, i, j):
 Denote $\mathcal{A} = \{A \subseteq X : i \in A, j \notin A\}$;
 Find $\delta = \min_{A \in \mathcal{A}} (r_{\mathcal{M}}(A) - y(A))$
 and an optimal $A \in \mathcal{A}$;
 If $y_j < \delta$ then $\{\delta \leftarrow y_j, A \leftarrow \{j\}\}$;
 $y_i \leftarrow y_i + \delta, y_j \leftarrow y_j - \delta$;
 Return (\mathbf{y}, A) .

Algorithm **PipageRound**(\mathcal{M}, \mathbf{y}):
 While (\mathbf{y} is not integral) do
 $T \leftarrow X$;
 While (T contains fractional variables) do
 Pick $i, j \in T$ fractional;
 $(\mathbf{y}^+, A^+) \leftarrow \mathbf{HitConstraint}(\mathbf{y}, i, j)$;
 $(\mathbf{y}^-, A^-) \leftarrow \mathbf{HitConstraint}(\mathbf{y}, j, i)$;
 $p \leftarrow \|\mathbf{y}^+ - \mathbf{y}\| / \|\mathbf{y}^+ - \mathbf{y}^-\|$;
 With probability p , $\{\mathbf{y} \leftarrow \mathbf{y}^-, T \leftarrow T \cap A^-\}$;
 Else $\{\mathbf{y} \leftarrow \mathbf{y}^+, T \leftarrow T \cap A^+\}$;
 EndWhile
 EndWhile
 Output \mathbf{y} .

The application in [4] was to monotone submodular functions, but as the authors mention, monotonicity is not used anywhere in the analysis. The technique as described in [5] yields the following.

Lemma B.1. *The pipage rounding technique, given a membership oracle for a matroid $\mathcal{M} = (X, \mathcal{I})$, a value oracle for a submodular function $f : 2^X \rightarrow R_+$, and \mathbf{y} in the base polytope $B(\mathcal{M})$, returns a random base B of value $\mathbf{E}[f(B)] \geq F(\mathbf{y})$.*

Monotonicity is used in [4] only to argue that a fractional solution can be assumed to lie in the base polytope without loss of generality. Therefore, if we are working with the base polytope (as in Section 4.2), we can use the pipage rounding technique without any modification.

If we are working with nonmonotone submodular functions and the matroid polytope, we have to do some additional adjustments to make sure that we do not lose anything when rounding a fractional solution. We proceed as follows. The following procedure takes a point $\mathbf{x} \in P(\mathcal{M})$ and while there is a fractional coordinate, it either pushes it to its maximum possible value, or makes it zero and removes it from the problem.

Algorithm **Adjust**(\mathcal{M}, \mathbf{x}):
 While ($\mathbf{x} \notin B(\mathcal{M})$) do
 If (there is i and $\delta > 0$ such that $\mathbf{x} + \delta \mathbf{e}_i \in P(\mathcal{M})$) do
 Let $x_{max} = x_i + \max\{\delta : \mathbf{x} + \delta \mathbf{e}_i \in P(\mathcal{M})\}$;
 Let $p = x_i / x_{max}$;
 With probability p , $\{x_i \leftarrow x_{max}\}$;
 Else $\{x_i \leftarrow 0\}$;
 EndIf
 If (there is i such that $x_i = 0$) do
 Delete i from \mathcal{M} and remove the i -coordinate from x .
 EndWhile
 Output $(\mathcal{M}, \mathbf{x})$.

Lemma B.2. *Given $\mathbf{x} \in P(\mathcal{M})$ and a submodular function $f(S)$ with its extension $F(\mathbf{x})$, the procedure **Adjust**(\mathcal{M}, \mathbf{x}) yields a restricted matroid \mathcal{M}' and a point $\mathbf{y} \in B(\mathcal{M}')$ such that $\mathbf{E}[F(\mathbf{y})] \geq F(\mathbf{x})$.*

Proof. If $\mathbf{x} \in P(\mathcal{M})$, there is always a point $\mathbf{z} \in B(\mathcal{M})$ dominating \mathbf{x} in every coordinate (see Corollary 40.2h in [28]). Therefore, if $\mathbf{x} \notin B(\mathcal{M})$, there is a coordinate which can be increased while staying in $P(\mathcal{M})$. In each step, when we choose randomly between increasing and decreasing x_i , the objective function is linear and the expectation of $F(\mathbf{x})$ is preserved. Hence, the process forms a martingale. At the end, $\mathbf{E}[F(\mathbf{y})]$ is equal to the initial value $F(\mathbf{x})$.

As long as we increase variables, each variable is increased to its maximum value and cannot be increased twice. Therefore, after at most n steps we either reach a point in $B(\mathcal{M})$ or make some variable x_i equal to 0. A variable equal to 0 is removed, which can be repeated at most n times. Hence, the process terminates in $O(n^2)$ time. \square

For a given $\mathbf{x} \in P(\mathcal{M})$, we run $(\mathcal{M}', \mathbf{y}) := \mathbf{Adjust}(\mathcal{M}, \mathbf{x})$, followed by $\mathbf{PipageRound}(\mathcal{M}', \mathbf{y})$. The outcome is a base in the restricted matroid where some elements have been deleted, i.e. an independent set in the original matroid. We call this the *extended pipage rounding* procedure. Lemma B.2 together with Lemma B.1 gives the following.

Lemma B.3. *The extended pipage rounding procedure, given a membership oracle for a matroid $\mathcal{M} = (X, \mathcal{I})$, a value oracle for a submodular function $f : 2^X \rightarrow R_+$, and \mathbf{x} in the matroid polytope $P(\mathcal{M})$, yields a random independent set $S \in \mathcal{I}$ of value $\mathbf{E}[f(S)] \geq F(\mathbf{x})$.*

C Solutions of nontrivial value

In this section, we prove that solutions of nontrivial value always exist for the problems of maximizing a nonmonotone submodular function subject to a matroid independence or matroid base constraint. We can assume that our matroid does not contain any loops (which can be removed beforehand), and in the case of a matroid base constraint no co-loops either (since co-loops participate in every solution and can be contracted). The bounds that we prove here are needed to ensure that our sampling errors can be made negligible compared to the value of the optimum.

Lemma C.1. *Let $\mathcal{M} = (X, \mathcal{I})$ be a matroid without loops (elements which are never in an independent set), $|X| = n$, and let $f : 2^X \rightarrow R_+$ be a (nonmonotone) submodular function such that $\max_{S \subseteq X} f(S) = M$. Let $OPT = \max\{f(I) : I \in \mathcal{I}\}$. Then*

$$OPT \geq \frac{1}{n}M.$$

Proof. We consider only solutions of size $|I| \leq 1$, which are independent because our matroid does not contain any loops. Let $f(I^*) = \max\{f(I) : |I| \leq 1\}$. By submodularity and nonnegativity of f , the value of any nonempty set can be estimated as

$$f(S) \leq \sum_{j \in S} f(\{j\}) \leq n f(I^*).$$

By definition, we also have $f(\emptyset) \leq f(I^*)$. Therefore, $M = \max f(S) \leq n f(I^*)$. \square

Lemma C.2. *Let $\mathcal{M} = (X, \mathcal{I})$ be a matroid without loops (elements which are never in a base) and co-loops (elements which are in every base), $|X| = n$ and let \mathcal{B} denote the bases of \mathcal{M} . Let $f : 2^X \rightarrow R_+$ be a (nonmonotone) submodular function such that $\max_{S \subseteq X} f(S) = M$. Let $OPT = \max\{f(B) : B \in \mathcal{B}\}$. Then*

$$OPT \geq \frac{1}{n^2}M.$$

Proof. Let us find $B = \{b_1, \dots, b_k\}$ greedily, by including in each step an element which has maximum possible marginal value $f_{\{b_1, \dots, b_\ell\}}(b_{\ell+1})$ among all elements that preserve independence. We claim that $f(B) \geq \frac{1}{n^2}M$.

The first element b_1 is the element of maximum value $\max_{i \in X} f(\{i\})$ (because there are no loops, all elements are eligible). Let $M = f(S^*)$. By submodularity,

$$f(S^*) \leq \sum_{i \in S^*} f(\{i\}) \leq n f(\{b_1\}).$$

Therefore, $f(\{b_1\}) \geq \frac{1}{n}M$. However, additional elements can decrease this value. Consider the last element b_k that we added to B and assume that $f(B) < \frac{1}{n^2}M$. Due to our greedy procedure and submodularity, the marginal values of successive elements keep decreasing, and therefore

$$f_{B \setminus \{b_k\}}(b_k) \leq \frac{1}{n-1}(f(B) - f(\{b_1\})) < -\frac{1}{n^2}M.$$

Since there are no coloops in \mathcal{M} , adding b_k was not our only choice - otherwise, we would have $\text{rank}(X \setminus \{b_k\}) < \text{rank}(X)$ which means exactly that b_k is a coloop. So there is another element b'_k that would form a base $\{b_1, \dots, b_{k-1}, b'_k\}$. The reason we did not select b'_k was that $f_{B \setminus \{b_k\}}(b'_k) \leq f_{B \setminus \{b_k\}}(b_k) < -\frac{1}{n^2}M$. By submodularity, if we add both elements, we obtain a set $\tilde{B} = \{b_1, \dots, b_k, b'_k\}$ such that

$$f(\tilde{B}) = f(B) + f_B(b'_k) \leq f(B) + f_{B \setminus \{b_k\}}(b'_k) < \frac{1}{n^2}M - \frac{1}{n^2}M = 0$$

which is a contradiction with the nonnegativity of f . □

We remark that this bound is tight up to a constant factor. Consider a complete bipartite directed graph on $X = X_1 \cup X_2$, $|X_1| = |X_2| = n/2$. The submodular function $f(S)$ is the directed cut function, expressing the number of arcs from S to \bar{S} . The matroid \mathcal{M} has bases that contain exactly one element from X_1 and $n-1$ elements from X_2 . It is easy to see that any base B has value $f(B) = 1$, while $f(X_1) = n^2/4$.