Geometric Rescaling Algorithms for Submodular Function Minimization[*]

Daniel Dadush[†1], László A. Végh[‡2], and Giacomo Zambelli[2]

dadush@cwi.nl, {l.vegh,g.zambelli}@lse.ac.uk
[1]Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
[2]London School of Economics and Political Science, London, UK

Abstract

We present a new class of polynomial-time algorithms for submodular function minimization (SFM), as well as a unified framework to obtain strongly polynomial SFM algorithms. Our algorithms are based on simple iterative methods for the minimum-norm problem, such as the conditional gradient and Fujishige-Wolfe algorithms. We exhibit two techniques to turn simple iterative methods into polynomial-time algorithms.

Firstly, we adapt the geometric rescaling technique, which has recently gained attention in linear programming, to SFM and obtain a weakly polynomial bound $O((n^4 \cdot \mathrm{EO} + n^5) \log(nL))$.

Secondly, we exhibit a general combinatorial black-box approach to turn εL-approximate SFM oracles into strongly polynomial exact SFM algorithms. This framework can be applied to a wide range of combinatorial and continuous algorithms, including pseudo-polynomial ones. In particular, we can obtain strongly polynomial algorithms by a repeated application of the conditional gradient or of the Fujishige-Wolfe algorithm. Combined with the geometric rescaling technique, the black-box approach provides an $O((n^5 \cdot \mathrm{EO} + n^6) \log^2 n)$ algorithm.

Finally, we show that one of the techniques we develop in the paper can also be combined with the cutting-plane method of Lee, Sidford, and Wong [29], yielding a simplified variant of their $O(n^3 \log^2 n \cdot \mathrm{EO} + n^4 \log^{O(1)} n)$ algorithm.

1 Introduction

Given a finite set V, a function $f : 2^V \to \mathbb{R}$ is *submodular* if

$$f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y) \quad \forall X, Y \subseteq V. \tag{1}$$

We denote $n := |V|$. Examples include the graph cut function, the coverage function, and the entropy function. Submodularity can be interpreted as a diminishing returns property and is therefore important in economics and game theory. Submodular optimization is widely applied in machine learning and computer vision (see e.g. [1]).

We will assume that the function f is given via an *evaluation oracle*: for every set $S \subseteq V$, we can query the value $f(S)$ in time EO. We will assume throughout that $f(\emptyset) = 0$; this is without loss of generality. In the *submodular function minimization (SFM)* problem, the objective is to find a minimizer of this function:

$$\min_{S \subseteq V} f(S). \tag{SFM}$$

The first weakly polynomial-time algorithm was given by Grötschel, Lovász, and Schrijver in 1981 [19], whereas the first strongly polynomial algorithm was given by the same authors in their 1988 book [20]. Both algorithms used the ellipsoid method. It remained an important goal to find a strongly polynomial combinatorial algorithm; this question was resolved independently by Schrijver [32] and by Iwata, Fleischer, and Fujishige [24] in 2000. Currently the best running time of a combinatorial algorithm is $O(n^5 \cdot \mathrm{EO} + n^6)$ by Orlin [30]. In a recent breakthrough, Lee, Sidford, and Wong [29] provided a $O(n^3 \log^2 n \cdot \mathrm{EO} + n^4 \log^{O(1)} n)$ algorithm based on a new cutting-planes method.

However, the above algorithms do not appear to work well for large scale instances that arise in applications such as speech recognition or image segmentation. A recent line of work has focused on exploiting special structures of specific submodular functions arising in these applications, such as decomposability [12, 13, 27, 35], but for general functions simple iterative algorithms appear to outperform the provably polynomial algorithms [17]. In particular, the Fujishige-Wolfe minimum-norm point algorithm [14, 37] appears to be among the best ones in practice [1, 17], despite the fact that the first pseudo-polynomial running time bound was given as recently as 2014 by Chakrabarty et al. [4].

Our contributions   This paper presents polynomial-time algorithms based on simple iterative methods such as the conditional gradient algorithm or the Fujishige-Wolfe algorithm. We exhibit two different techniques to improve the performance of these algorithms to polynomially bounded. The first technique uses *geometric rescaling*, whereas the second provides a *unified combinatorial framework for strongly polynomial SFM algorithms*. In what follows, we provide an overview of both techniques.

Geometric rescaling has recently gained attention in the context of linear programming. This is a general algorithmic technique to turn simple iterative algorithms to polynomial-time algorithms for LP feasibility, by adaptively changing the scalar product. The first such algorithms were given by Betke [3], and by Dunagan and Vempala [10], and a number of papers have since appear on the subject. We refer the reader to [7] for an overview of the literature.

In this paper we focus on one such algorithm, introduced by the authors in [7] and named there the Full Support Image Algorithm (the same algorithm was obtained independently by Hoberg and Rothvoß [22]). We show how this algorithm can be adapted to (SFM). In Section 6 we will see that the framework we introduce is robust, in the sense that it can be easily adapted to different rescaling algorithms so long as they are applicable to conic problems in the separation oracle model, such as the algorithms in [2, 6, 7, 31] (see for example a recent note of Fujishige [16] showing how an algorithm of Chubanov [6] can be used in this framework). The reason for focusing on the Image Algorithm of [7] is that, within this framework, it provides the best running time bounds for (SFM) among the known algorithms in this class.

We introduce new techniques that enable our Image Algorithm to provide both primal and dual optimal solutions for (SFM). The *sliding technique* is used to obtain a primal optimal solution: we reduce the optimization problem (SFM) to a dynamically changing feasibility problem. The *pull-back technique* enables to identify a dual optimality certificate. Moreover, the same technique allows us also to obtain approximate dual solutions, and it is also applicable in the general LP feasibility setting.

For integer valued submodular functions, our geometric rescaling algorithm finds both primal and dual optimal solutions, in running time $O((n^4 \cdot \mathrm{EO} + n^5) \log(nL))$, where the complexity parameter L denotes the largest norm of a point in the base polytope. This matches the best weakly polynomial guarantees [23, 25] prior to the work of Lee, Sidford, and Wong [29].

Unified combinatorial framework. Building on the geometric rescaling technique, we also obtain a *strongly polynomial* $O((n^5 \cdot \mathrm{EO} + n^6) \log^2 n)$ algorithm. This is obtained from a unified combinatorial framework which allows us to turn any algorithm that can produce a δL-approximate solution to (SFM) in pseudo-polynomial $\mathrm{poly}(n, 1/\delta)$ running time, into an exact strongly polynomial algorithm. More specifically, if we are given an oracle that can produce a set $W \subseteq V$ and a point y in the base polytope such that $f(W) \leq y^-(V) + \delta L_f$ (where $y^-(V)$ is the sum of all negative components of y and L_f is the largest 1-norm of

any point in the base polytope), then with the choice $\delta \in \Theta(1/n^3)$ we can find an optimal solution to (SFM) in roughly $O(n^2)$ calls to the oracle.

Hence, somewhat surprisingly, even pseudo-polynomial time algorithms such as the conditional gradient or the Fujishige-Wolfe algorithm immediately give rise to strongly polynomial time algorithms.

We emphasise that our approach is *black-box*: we explicitly formulate the approximate oracle requirement, and show that any routine fulfilling such requirements provides a strongly polynomial-time algorithm for (SFM). To illustrate our point, we will show that various routines in the literatures, such as [24, 25, 29], satisfy such black-box requirements. We believe that our approach is more modular than the previous works, in the sense that in all previous papers the combinatorial arguments on strongly polynomial progress were intertwined with the details of some "basic" routine.

We can also apply this unified framework to the cutting plane method. Using the general cutting plane algorithm of Lee, Sidford, and Wong [29], we show that our black-box method can recover their running time bound $O(n^3 \log^2 n \cdot \mathrm{EO} + n^4 \log^{O(1)} n)$ for (SFM). This is made possible by the use of the same sliding technique developed for our geometric rescaling algorithm.

The general combinatorial framework is based on maintaining a *ring family* guaranteed to contain all minimizer sets, where the size of the family decreases through the algorithm until a minimizer is found. This technique was introduced by Iwata, Fleischer, and Fujishige [24], and used in multiple subsequent papers, such as Iwata and Orlin [25], and Lee, Sidford, and Wong [29]. We note that this technique ultimately traces back to strongly polynomial algorithms for minimum-cost flows, pioneered by Tardos [36]. Our implementation also adopts a simplified variant of the bucketing technique of [29] that leads to a factor n improvement in the running time compared to the original framework of [24].

The rest of the paper is structured as follows. Section 2 contains definitions and the necessary background, including an overview of the relevant iterative methods. Section 3 presents the weakly polynomial geometric rescaling algorithm to solve SFM, while in Subsection 3.4, we describe the pull-back technique that enables the implementation of the approximate oracle using our geometric rescaling method. Section 4 presents the general framework for strongly polynomial algorithms. Section 5 shows how cutting plane methods, and in particular the cutting plane algorithm in [29], can be used in the strongly polynomial framework. Finally, in Section 6 we discuss variants of the weakly-polynomial geometric rescaling algorithm.

2   Preliminaries

We refer the reader to [15] and [33, Sections 44-45] for the basics of submodular optimization; these references contain all definitions as well as the proof of the results presented next. The survey [1] provides an overview of continuous algorithms for submodular function minimization.

For a vector $z \in \mathbb{R}^V$, we denote by $z(v)$ the component of z relative to $v \in V$, and for a subset $S \subseteq V$ we use the notation $z(S) = \sum_{v \in S} z(v)$. We let $\|z\|_1 := \sum_{v \in V} |z(v)|$ and $\|z\|_2 := (\sum_{v \in V} |z(v)|^2)^{1/2}$ denote the ℓ_1 and ℓ_2-norms, respectively; we will also use the simpler $\|z\| := \|z\|_2$. For a number $a \in \mathbb{R}$, we let $a^+ = \max\{0, a\}$ and $a^- = \min\{0, a\}$; hence, $a = a^+ + a^-$. Similarly, given a vector $z \in \mathbb{R}^V$, we denote $z^+ = (z(v)^+)_{v \in V}$ and $z^- = (z(v)^-)_{v \in V}$.

We denote by $\mathbb{S}^n_{++}$ the set of $n \times n$ symmetric positive definite real matrices. Every $Q \in \mathbb{S}^n_{++}$ defines the scalar product $\langle x, y \rangle_Q \overset{\mathrm{def}}{=} x^\mathsf{T} Q y$, inducing the norm $\|x\|_Q \overset{\mathrm{def}}{=} \sqrt{\langle x, x \rangle_Q}$.

The base polytope and the greedy algorithm   Let f be a submodular function, where we assume as usual that $f(\emptyset) = 0$. The *base polytope* of f is defined as

$$B(f) := \left\{ x \in \mathbb{R}^V : x(S) \leq f(S) \ \forall S \subseteq V, \ x(V) = f(V) \right\}.$$

This polytope $B(f)$ is non-empty for every submodular function f. Its elements are called *bases*, and its vertices are the *extreme bases*. Extreme bases correspond to permutations

of the ground set. More precisely, for any ordering $v_1, v_2, \ldots, v_n$ of the elements of V, the following point is a vertex of $B(f)$, and every vertex is of this form for some ordering:

$$
\begin{aligned}
x(v_1) &:= f(\{v_1\}), \\
x(v_i) &:= f(\{v_1, \ldots, v_i\}) - f(\{v_1, \ldots, v_{i-1}\}) \quad \forall i = 2, \ldots, n.
\end{aligned}
\tag{2}
$$

Furthermore, given a weight function $w : V \to \mathbb{R}$, one can compute an extreme base minimizing $w^\mathsf{T} x$ by the greedy algorithm $\mathrm{GREEDYMIN}(f, w)$ as follows: order the vertices in V so that $w(v_1) \leq w(v_2) \leq \ldots \leq w(v_n)$, and output x defined by (2) as the optimal solution. The value of the minimum-cost is then given by

$$
\min_{x \in B(f)} w^\mathsf{T} x = \sum_{i=1}^{n-1} f(\{v_1, \ldots, v_i\})(w(v_i) - w(v_{i+1})) + f(V)w(v_n).
\tag{3}
$$

The subroutine $\mathrm{GREEDYMIN}(f, w)$ requires $O(n \cdot \mathrm{EO} + n \log n)$ arithmetic operations. If w has several entries with the same value, then there are multiple ways to sort the elements of V in ascending value of w, each ordering potentially giving rise to a different optimal extreme base of $B(f)$. The extreme bases corresponding to the possible tie-breakings are the vertices of the face of $B(f)$ minimizing $w^\mathsf{T} x$.

If $v_1, \ldots, v_n$ is the ordering computed by $\mathrm{GREEDYMIN}(f, w)$, we define

$$
\mathrm{MINSET}(f, w) \overset{\text{def}}{=} \arg\min\{f(S) : S = \{v_1, \ldots, v_i\} \exists i \in [n] \text{ or } S = \emptyset\}.
\tag{4}
$$

A min-max characterization of (SFM) was given by Edmonds:

Theorem 2.1 (Edmonds [11]). *For any submodular function $f : 2^V \to \mathbb{R}$,*

$$
\max\{x^-(V) : x \in B(f)\} = \min\{f(S) : S \subseteq V\}.
\tag{5}
$$

We will often use the following simple consequence. Assume that for some $x \in B(f)$, $S \subseteq V$, and $\varepsilon > 0$, we have $f(S) \leq x^-(V) + \varepsilon$. Then $f(S) \leq f(T) + \varepsilon$ for any $T \subseteq V$.

Complexity parameters When dealing with weakly polynomial time algorithms for (SFM), various complexity parameters have been considered in the literature to measure the running time. All these parameters turn out to be equivalent within an $O(n)$ factor. Nonetheless, in certain parts of the paper different choices will fit more easily, hence we introduce all of them below. We define

$$
L_f \overset{\text{def}}{=} \max\{\|z\|_1 : z \in B(f)\}, \quad L_{f,2} \overset{\text{def}}{=} \max\{\|z\|_2 : z \in B(f)\}, \quad F_f \overset{\text{def}}{=} \max\{|f(S)| : S \subseteq V\}.
$$

Some of our algorithms require the explicit knowledge of these complexity parameters. While these are hard to compute, we introduce easily computable upper-bounds that are essentially equivalent (see Claim 2.2) to the ones above. These upper-bounds are given by the norms of the vector $\alpha : V \to \mathbb{R}$ defined by

$$
\alpha(v) := \max\{|f(\{v\})|, |f(V) - f(V \setminus \{v\})|\}, \quad v \in V.
$$

In the following claim, we highlight that all these parameters are essentially equivalent within an $O(n)$ factor.

Claim 2.2. *For any submodular function $f : 2^V \to \mathbb{R}$ such that $f(\emptyset) = 0$, the following hold.*

(i) $L_{f,2} \leq L_f \leq \sqrt{n} L_{f,2}$.

(ii) $L_f = \Theta(F_f)$.

(iii) For every $z \in B(f)$, $|z(v)| \leq \alpha(v)$ for all $v \in V$. Hence, $L_f \leq \|\alpha\|_1 = \alpha(V)$ and $L_{f,2} \leq \|\alpha\|_2$.

(iv) $\alpha(v) \leq 2F_f$ for all $v \in V$.

Proof. (i) is obvious, for (ii) see e.g. [5, Lemma 5], and also [21, 26], for (iii) see [15, Section 3.3], and (iv) follows immediately from the definition of α. $\qquad\square$

Our running time bounds will contain terms of the form $\log(nZ)$ for various choices of $Z \in \{F_f, L_f, L_{f,2}, \|\alpha\|_1, \|\alpha\|_2\}$. The above lemma shows that all these terms are within a constant factor of one another, hence the specific choice of complexity parameter does not matter.

The minimum-norm point problem Fujishige [14] showed a reduction of (SFM) to the following convex quadratic optimization problem.

Theorem 2.3 (Fujishige [14]). *Let z be the unique optimal solution to*

$$\min\left\{\frac{1}{2}\|x\|_2^2 : x \in B(f)\right\}. \tag{6}$$

Then, the set $S^* = \{v \in V : z(v) < 0\}$ is the inclusionwise minimal minimizer of (SFM), *and $f(S^*) = z^-(V)$. Furthermore, $|f(S^*)| \le \sqrt{n}\|z\|_2$.*

Note that in case of $f(V) = 0$, Theorems 2.1 and 2.3 imply that the minimizer of the 2-norm also minimizes the 1-norm in $B(f)$. Indeed, in this case for any $y \in B(f)$ $y^-(V) = f(V) - y^+(V) = -y^+(V)$, implying $\|y\|_1 = -2y^-(V)$, which gives $\|z\|_1 = -2z^-(V) = -2f(S^*) = -2\max\{x^-(V) : x \in B(f)\} = \min\{\|x\|_1 : x \in B(f)\}$.

Note that $z \in B(f)$ is the minimum norm point if and only if $z^\top x \ge \|z\|_2^2$ for all $x \in B(f)$. An approximate optimal solution to (6) can be converted to an approximate optimal solution to (5), as stated below.

Theorem 2.4 (Bach [1]). *Assume that $z \in B(f)$ satisfies $\|z\|_2^2 \le z^\mathsf{T} x + \varepsilon$ for all $x \in B(f)$, for some $\varepsilon > 0$. Let $S = \mathrm{MINSET}(f, z)$. Then, $f(S) \le z^-(V) + \sqrt{2n\varepsilon}$. Consequently, $f(S) \le f(T) + \sqrt{2n\varepsilon}$ for any $T \subseteq V$.*

The above is a consequence of the more general statement in [1, Proposition 10.5]. Since it might not be immediate to see such implication, for completeness we provide a proof of the previous theorem in the appendix.

2.1   Iterative methods for SFM

Convex optimization algorithms can be naturally applied to SFM, either by solving the quadratic formulation (6), or by minimizing the so-called *Lovász-extension*, which we do not discuss here. We refer the reader to [1] for an overview of such algorithms. Here, we briefly outline two important algorithms based on (6).

The conditional gradient algorithm The conditional gradient, or Frank-Wolfe algorithm, maintains a point $y \in B(f)$, represented as a convex combination $y = \sum_{i=1}^{k} \lambda_i g_i$ of extreme bases. It is initialized with $y = g$ for an arbitrary extreme base g. Every iteration runs $\mathrm{GREEDYMIN}(f, y)$ to obtain an extreme base g'. If $y^\mathsf{T} g' \ge \|y\|_2^2$, then y is the minimum-norm point in $B(f)$, and the algorithm terminates. Otherwise, y is replaced by the minimum-norm point y' on the line segment $[y, g']$. This is also known as the von Neumann algorithm, described by Dantzig [9]. The algorithm was first used to test membership in the base polytope by Sohoni [34]. The standard convergence analysis of the conditional gradient algorithm and Theorem 2.4 provide the following convergence bound (see e.g. [1, Sec 10.8]).

Theorem 2.5. *For any $\delta > 0$, within $O(n/\delta^2)$ iterations the conditional gradient algorithm computes $y \in B(f)$ such that for $S = \mathrm{MINSET}(f, y)$, we have $f(S) \le y^-(V) + O(\delta L_{f,2})$. The total running time is $O((n^2 \cdot \mathrm{EO} + n^2 \log n)/\delta^2)$.*

The Fujishige-Wolfe algorithm Wolfe [37] gave a finite algorithm for finding the minimum-norm point in a polytope given by its vertices; his algorithm can also be interpreted as an active set method [1]. Fujishige adapted Wolfe's algorithm to SFM [14, 17]. We now give a brief sketch of the algorithm; for a more detailed description see [4, 17, 37].

An affinely independent set of points $X \in \mathbb{R}^n$ is called a *corral* if the orthogonal projection of 0 onto the affine hull of X is in the relative interior of the convex hull of X. In particular, the optimal solution to the minimum-norm point problem can be obtained by a corral, comprising vertices of the minimal face of the polytope containing the minimum-norm point.

Every *major cycle* of the Fujishige-Wolfe algorithm starts and ends with a corral formed by extreme bases in $B(f)$. The algorithm is initialized with an arbitrary extreme base (note that every singleton set is a corral). Let X be the corral at the beginning of a major cycle,

and let y be the projection of 0 onto the affine hull of X; this can be obtained by a closed-form formula. Let us run GREEDYMIN(f, y) to obtain an extreme base g' minimizing $y^\mathsf{T} g'$. If $y^\mathsf{T} g' \geq \|y\|_2^2$, then the algorithm terminates returning y as the minimum-norm point in $B(f)$. Otherwise, we consider $X' = X \cup \{g'\}$, which is also affinely independent. We set $\bar{x} = y$, and compute y' as the projection of 0 onto the affine hull of X'. If y' is in the relative interior of $\mathrm{conv}(X')$, the major cycle terminates with the new corral X'. Otherwise, we start a *minor cycle*: we replace X' by the extreme points of the minimal face of $\mathrm{conv}(X')$ intersecting the line segment $[\bar{x}, y']$; the new $\bar{x}$ is defined to be the unique intersection point. Minor cycles are repeated until a corral is obtained. Finite convergence is guaranteed since $\|\bar{x}\|_2$ decreases in every major and minor cycle, and the number of corrals is finite. However, a bound on the convergence rate, which we summarize below, was only recently given in [4]. Further improvements on the analysis were observed by Lacoste and Jaggi in [28].

Theorem 2.6 (Chakrabarty et al. [4], Lacoste and Jaggi [28])**.** *For any $\delta > 0$, within $O(n/\delta^2)$ iterations (major and minor cycles) Fujishige-Wolfe algorithm computes $y \in B(f)$ such that, for $S = \mathrm{MINSET}(f, y)$, we have $f(S) \leq y^-(V) + O(\delta L_{f,2})$. The total running time is $O((n^2 \cdot \mathrm{EO} + n^3)/\delta^2)$.*

3 Weakly polynomial algorithm via rescaling

Throughout this section we assume that f is an **integer valued** submodular function. The assumption that f is integer valued is needed in the context of weakly polynomial-time algorithms..

The geometric rescaling algorithm The Full Support Image Algorithm in [7, Section 3.2] is applicable to the following oracle setting. Let $\Sigma \subseteq \mathbb{R}^n$ be non-empty, full dimensional cone; our aim is to find a feasible point in the interior. We are given a separation oracle for $\mathrm{int}(\Sigma)$; that is, for any vector w, the oracle decides whether $w \in \mathrm{int}(\Sigma)$, and if not, it returns a vector z such that $z^\mathsf{T} w \leq 0$ but $z^\mathsf{T} y > 0$ for all $y \in \mathrm{int}(\Sigma)$. Then the algorithm finds a point in $\mathrm{int}(\Sigma)$ in $O(n^3 \log \omega^{-1})$ calls to the separation oracle, where ω is a condition number which we will define in Section 3.3. We remark that the parameter ω can be lower bounded by the *width of the cone* Σ, defined as the radius of the largest ball contained in Σ and centered on the surface of the unit sphere.

Consider now a submodular function f with $f(V) = 0$. Assume we want to decide whether $f(S) \geq 0$ for all $S \subseteq V$, that is, if $S = \emptyset$ is an optimal solution to (SFM). It follows from the definition of the base polytope that $0 \leq f(S)$ for every $S \subseteq V$ if and only if $0 \in B(f)$ (note that $f(V) = 0$ is needed for this equivalence). Consider now the cone

$$\Sigma = \{w \in \mathbb{R}^n : w^\mathsf{T} y \geq 0 \quad \forall y \in B(f)\} \tag{7}$$

(that is, Σ is the negative of the polar cone of $B(f)$).

Remark 3.1. *The cone Σ has a non-empty interior if and only if $0 \notin B(f)$.*

Proof. Observe that $0 \notin B(f)$ if and only if there exists an hyperplane weakly separating 0 and $B(f)$, that is, $0 \notin B(f)$ if and only if there exists $w \in \mathbb{R}^n$ such that $w^\mathsf{T} y > 0$ for all $y \in B(f)$, i.e., there exists w is in the interior of Σ. $\qquad\square$

A separation oracle for $\mathrm{int}(\Sigma)$ is provided by GREEDYMIN, since for every $y \in \mathbb{R}^V$ we have $y \in \mathrm{int}(\Sigma)$ if and only if $\min_{x \in B(f)} y^\mathsf{T} x > 0$. Consequently, if the algorithm does not terminate in the required running time bound, we can conclude that $f(S) \geq 0$ for all $S \subseteq V$. We could use this algorithm in a binary search framework to solve (SFM). When querying $\min_{S \subseteq V} f(S) \geq -\mu$ for a $\mu > 0$, we shift $f(S)$ by $f(S) + \mu$ for every $S \subsetneq V$, $S \neq \emptyset$.

The main drawback of the binary search scheme is that it only provides the optimum value, but does not give either an optimal set S, nor a dual certificate as in Theorem 2.1. Also, the binary search leads to an extra $\log F_f$ factor in the running time.

In this section, we describe a variant of this algorithm, which provides a primal optimal solution, and does not require binary search. This will be achieved by dynamically shifting or "sliding" the function f throughout the algorithm, as explained below. However, the

algorithm does not directly return a dual certificate of optimality. This can be obtained using the pull-back technique introduced in Section 3.4; see also the remark after Theorem 3.14.

We start by describing the sliding framework. Besides the geometric rescaling algorithm described next, this technique will also be useful for devising simple cutting plane algorithms for SFM in Section 5.

Sliding the function Throughout the algorithm, we maintain a value $\mu \in \mathbb{Z}_+$, along with a set W, such that $f(W) = -\mu$. We initialize $\mu = \max\{0, -f(V)\}$, and set $W = \emptyset$ or $W = V$ accordingly. Hence $-\mu$ gives an upper bound on $\min_{S \subseteq V} f(S)$. The algorithm terminates once it concludes that $f(W) = \min_{S \subseteq V} f(S)$ for the current W. We define the function $f_\mu : 2^V \to \mathbb{Z}$ as

$$f_\mu(S) \overset{\text{def}}{=} \begin{cases} 0, & \text{if } S = \emptyset \text{ or } S = V, \\ f(S) + \mu, & \text{otherwise.} \end{cases} \tag{8}$$

This operation is known as the μ-enlargement of the function f (see Fujishige [15, Section 3.1(d)]). The operation has been used in the context of submodular function minimization in [18].

Lemma 3.2. *For a submodular function f and a value $\mu \geq \max\{0, -f(V)\}$, the function f_μ is submodular. If $0 \in B(f_\mu)$, then $-\mu \leq f(S)$ for every $S \subseteq V$. Furthermore, $B(f_\mu) \subseteq B(f_{\mu'})$ whenever $\mu \leq \mu'$.*

Proof. The function f' defined by $f'(S) = f(S) + \mu$ for all $S \subseteq V$ is clearly submodular. We obtain f_μ from f' by decreasing the value of $f'(\emptyset)$ and $f'(V)$; note that the bound on μ guarantees that these are both nonnegative. Submodularity is maintained, since for any choice of X and Y, the right-hand-side in (1) decreases by at least as much as the left-hand-side when replacing f' by f_μ. If $0 \in B(f_\mu)$, then $0 \leq f_\mu(S)$ for any $S \subseteq V$. If $S \notin \{\emptyset, V\}$, then this gives $f(S) \geq -\mu$; the choice of μ guarantees the same for $S = \emptyset$ and $S = V$. For $\mu' \geq \mu$, the containment $B(f_\mu) \subseteq B(f_{\mu'})$ follows, since the constraints $x(S) \leq f_{\mu'}(S)$ are implied by the constraints $x(S) \leq f_\mu(S)$. $\square$

The following Lemma will be used to update the value of μ.

Lemma 3.3. *Consider a value $\mu \geq \max\{0, -f(V)\}$, and let $w : V \to \mathbb{R}$ be a cost function such that $\min\{w^\mathsf{T} x : x \in B(f_\mu)\} > 0$. For $S = \textsc{MinSet}(f_\mu, w)$, we have $f(S) < -\mu$.*

Proof. Let $v_1, \ldots, v_n$ be the ordering of V returned by $\textsc{GreedyMin}(f_\mu, w)$ such that $w(v_1) \leq w(v_2) \leq \ldots \leq w(v_n)$. From (3) and from the fact that $f_\mu(V) = 0$ (by construction) we see that the minimum value of $w^\mathsf{T} x$ over $B(f_\mu)$ can be written as

$$w^\mathsf{T} x = \sum_{i=1}^{n-1} (f(\{v_1, \ldots, v_i\}) + \mu)(w(v_i) - w(v_{i+1})).$$

Since $w^\mathsf{T} x > 0$ and $w(v_i) - w(v_{i+1}) \leq 0$ for $i = 1, \ldots, n-1$, it follows that $f(\{v_1, \ldots, v_i\}) < -\mu$ for some value of i, implying the claim. $\square$

Lemma 3.4. *Consider a value $\mu \geq \max\{0, -f(V)\}$ such that $\mu = -f(W)$ for some $W \subseteq V$. Then, $L_{f_\mu} \leq 4L_f$.*

Proof. For any permutation of the ground set, let g and g' be the corresponding extreme bases in $B(f)$ and in $B(f_\mu)$, respectively. These only differ in the first and last components: respectively by $+\mu$, and by $-\mu - f(V)$. Hence, $\|g'\|_1 \leq \|g\|_1 + 2\mu + |f(V)|$. Note that $\mu \leq L_f$; this is because $\mu = -f(W)$ for a certain set W, and therefore any permutation that starts with the elements of W will give an extreme base of 1-norm at least $|f(W)|$. Similarly, $|f(V)| \leq L_f$. The claim follows. $\square$

Algorithm 1 The sliding von Neumann algorithm

Input: A submodular function $f : 2^V \to \mathbb{Z}$, a value $\mu \geq \max\{0, -f(V)\}$, a set $W \subseteq V$ with $f(W) = -\mu$, a matrix $Q \in \mathbb{S}_{++}^n$, and an $\varepsilon > 0$.

Output:

- A value $\mu' \geq \mu$ and a set $W' \subseteq V$ with $f(W') = -\mu'$,
- bases $g_1, \ldots, g_k \in B(f_{\mu'})$, $x \in \mathbb{R}^k$, $y \in \mathbb{R}^n$ such that $y = \sum_{i=1}^{k} x_i g_i / \|g_i\|_Q$, $\vec{e}^\mathsf{T} x = 1$, $x \geq 0$, and $\|y\|_Q \leq \varepsilon$.

1: Set $\mu' := \mu$, $W' := W$.
2: Pick g_1 as an arbitrary vertex of $B(f_\mu)$. Set $x_1 := 1$, $y := g_1 / \|g_1\|_Q$.
3: Let $k := 2$.
4: **while** $\|y\|_Q > \varepsilon$ **do**
5:     Let $g_k \leftarrow \textsc{GreedyMin}(f_{\mu'}, Qy)$.
6:     **if** $y^\mathsf{T} Q g_k > 0$ **then**   ▷ sliding
7:         $W' := \textsc{MinSet}(f_{\mu'}, Qy)$; $\delta := -f_{\mu'}(W')$; $\mu' := -f(W')$;
8:         Set v_1 and v_n to be the first and last elements of V in increasing order by the weight vector Qy.
9:         $g_k(v_1) := g_k(v_1) + \delta$; $g_k(v_n) := g_k(v_n) - \delta$.
10:     **end if**
11:

$$\lambda := \frac{\left\langle y - \frac{g_k}{\|g_k\|_Q}, y \right\rangle_Q}{\left\| y - \frac{g_k}{\|g_k\|_Q} \right\|_Q^2};$$

12:     $y := (1 - \lambda)y + \lambda g_k / \|g_k\|_Q$;   ▷ min Q-norm point on $[y, g_k/\|g_k\|_Q]$
13:     $x_k := \lambda$;
14:     **for** $i = 1, \ldots, k-1$ **do** $x_i := (1 - \lambda)x_i$
15:     $k := k + 1$
    return μ', W', the vectors $g_1, \ldots, g_k$, x, and y.

3.1 The sliding von Neumann algorithm

The Full Support Image Algorithm of [7] uses the von Neumann algorithm as the basic subroutine. The von Neumann algorithm was described in [8] to find a feasible solution to the system $A^\mathsf{T} y > 0$ for a matrix $A \in \mathbb{R}^{n \times p}$. At every iteration, y is maintained as a convex combination of columns of A normalized by their norm, that is, y is maintained in the convex hull of $a_1/\|a_1\|, \ldots, a_p/\|a_p\|$. Initially $y = a_i/\|a_i\|$ for some arbitrary $i \in [n]$; at any iteration, the algorithm terminates if $A^\mathsf{T} y > 0$, otherwise a column a_k such that $a_k^\mathsf{T} y \leq 0$ is selected, and y is updated to be the point of minimum norm in the line segment $[a_k/\|a_k\|, y]$. The von Neumann algorithm can be seen as a variant of the conditional gradient algorithm for the problem $\min\{\frac{1}{2}\|y\|^2 : y \in \text{conv}(\{a_1/\|a_1\|, \ldots, a_p/\|a_p\|\})\}$, differing in the fact that von Neumann algorithm only needs to decide whether the minimum value of the norm is positive.

Our *sliding von Neumann algorithm* (Algorithm 1) is a modification of the standard von Neumann algorithm, adapted to the context of submodular function minimization. The algorithm is applied to the extreme bases of $B(f)$, in order to decide if there exists a point in the interior of the cone Σ defined in (7). The main differences are the following.

- The algorithm incorporates the adaptive shifting f_μ described previously. In particular, when the current y satisfies $g^\mathsf{T} y > 0$ for all $g \in B(f_\mu)$, the algorithm does not stop, but it determines $S \subseteq V$ with $f_\mu(S) < 0$ as in Lemma 3.3, it updates $\mu := -f(S)$, and resumes from the current point y.

- Rather than maintaining y as a convex combination of the extreme bases of $B(f_\mu)$ normalized by their 2-norms, the algorithm will use a more general norm, defined by a symmetric positive definite matrix Q given as part of the input. This is because the algorithm will be used as a subroutine of Algorithm 2, where the norm will be periodically rescaled.

Lemma 3.5. *Algorithm 1 terminates in $\lceil 1/\varepsilon^2 \rceil$ iterations, returning a correct output.*

Proof. Correctness. At every iteration, the algorithm calls GREEDYMIN($f_{\mu'}, Qy$) to obtain an extreme base $g_k \in B(f_\mu)$ minimizing $y^\mathsf{T} Qx$ over $B(f_\mu)$. If $y^\mathsf{T} Qg_k \leq 0$, then we perform a von Neumann update, that is, we replace y with the minimum Q-norm point on the line segment $\left[y, \frac{g_k}{\|g_k\|_Q} \right]$ (which is given by the choice of λ in line 11).

Consider now the case $y^\mathsf{T} Qg_k > 0$. In this case, at line 7 the current value of μ' is updated to a new value, say μ''. According to Lemma 3.3, the set W' determined at Line 7 satisfies $f_{\mu'}(W') < 0$, hence $\mu'' = -f(W') = \mu' - f_{\mu'}(W') > \mu'$. In particular, observe that $\delta = \mu'' - \mu' > 0$. By definition, for all $S \subseteq V$,

$$f_{\mu''}(S) = \begin{cases} f_{\mu'}(S) + \delta & \text{if } S \neq \emptyset, V \\ 0 & \text{if } S = \emptyset \text{ or } S = V \end{cases} \tag{9}$$

The vector g_k is updated at Line 9. Let us denote by g_k' the updated vector. It follows from (2) and from (9) that g_k' is the extreme base of $B(f_{\mu''})$ defined by the same ordering of the elements $v_1, \ldots, v_n$ of V that defined g_k as an extreme base of $B(f_{\mu'})$.

It follows that g_k' is the solution output by GREEDYMIN($f_{\mu''}, Qy$), hence it is an optimal solution for $\min\{y^\mathsf{T} Qz : z \in B(f_{\mu''})\}$. Since, by the choice of μ'', $f(\{v_1, \ldots, v_i\}) \geq -\mu''$ for all $i \in \{1, \ldots, n\}$, Lemma 3.3 implies that $y^\mathsf{T} Qg_k' \leq 0$.

Since g_k' is in $B(f_{\mu''})$ and $y^\mathsf{T} Qg_k' \leq 0$, we can perform a von Neumann update. Lemma 3.2 implies that $B(f_{\mu''}) \supseteq B(f_{\mu'})$; hence all vectors g_i computed thus far are still contained in $B(f_{\mu'})$. Since the algorithm terminates when $\|y\|_Q \leq \varepsilon$, it follows that the t and the elements $g_1, \ldots, g_k$ returned by the algorithm satisfy the requirements.

Running time. By the standard analysis of von Neumann's algorithm given by Dantzig [8], $1/\|y\|_Q^2$ increases by at least 1 at every update, hence after k iterations $\|y\|_Q \leq 1/\sqrt{k}$. We remark that, even though Dantzig's analysis applies to the case of the 2-norm (i.e. $Q = I$), the case of a norm defined by a generic positive definite matrix Q can be easily reduced to the 2-norm case. Indeed, note that for any vector $z \in \mathbb{R}^n$, $\|z\|_Q = \|Q^{1/2}z\|_2$. In particular, since at every iteration we maintain $y = \sum_{i=1}^k x_i g_i / \|g_i\|_Q$, multiplying on both sides by $Q^{1/2}$ this is equivalent to $Q^{1/2}y = \sum_{i=1}^k x_i (Q^{1/2}g_i)/\|Q^{1/2}g_i\|_2$. Hence the steps of von Neumann algorithm for the norm defined by Q are the same as the steps of the standard von Neumann algorithm applied to the vectors $(Q^{1/2}g)/\|Q^{1/2}g_i\|_2$, $g \in B(f)$, where $Q^{1/2}y$ is maintained as a convex combination of such vectors. Dantzig's analysis implies that $1/\|Q^{1/2}y\|_2^2 = 1/\|y\|_Q^2$ increases by at least 1 at every update. $\square$

3.2 Geometric rescaling algorithm for SFM

In our geometric rescaling algorithm, Algorithm RESCALING-SFM shown in Algorithm 2, we set parameters

$$\varepsilon \stackrel{\text{def}}{=} \frac{1}{20n}, \qquad T \stackrel{\text{def}}{=} 5n \log(nL_{f,2}).$$

In the algorithm below we also maintain a set Γ collecting all the elements of the base polytope generated during the Sliding von Neumann iterations; its main purpose is for clarity during the analysis, as we will need to refer to it.

Algorithm RESCALING-SFM is the adaptation of the Full Support Image Algorithm to our submodular setting, using the sliding von Neumann algorithm. We need to modify the algorithm and its analysis to reflect that the feasible region keeps changing due to the updates to the value of μ.

The value μ keeps increasing during the algorithm; it is updated within the sliding von Neumann subroutine. We also maintain a set W with $f(W) = -\mu$. The algorithm stops after T rescalings. At this point, we conclude from a volumetric argument that the current W is the minimizer of f. We show the following running time bound.

Algorithm 2 Rescaling-SFM

Input: A submodular function $f : 2^V \to \mathbb{Z}$.
Output: A set W minimizing $f(W)$.
1: Set $Q := I_n$, $R := I_n$. Set $\Gamma := \emptyset$.
2: Set $\mu := \max\{0, -f(V)\}$.
3: **if** $f(V) < 0$ **then** $W := V$, **else** $W := \emptyset$.
4: **for** $i = 1, \dots, T$ **do**
5:      Call SLIDING VON NEUMANN$(f, \mu, W, Q, \varepsilon)$ to obtain the new values
        of μ and W, and vectors $g_1, \dots, g_k, x, y$.
6:      **If** $y = 0$, then **stop**; **return** W
7:      **Else rescale**

$$R := \frac{1}{(1+\varepsilon)^2} \left(R + \sum_{i=1}^{k} \frac{x_i}{\|g_i\|_Q^2} g_i g_i^\mathsf{T} \right); \quad Q := R^{-1}. \quad (10)$$

8:      $\Gamma := \Gamma \cup \{g_1, \dots, g_k\}$;
     return W.

Theorem 3.6. *Algorithm* RESCALING-SFM *finds an optimal solution to* (SFM) *in time* $O((n^4 \cdot \mathrm{EO} + n^5) \log(nL_{f,2}))$.

Note that, the definition of T requires knowing the value of $L_{f,2}$; we can replace it by the bound $\|\alpha\|_2$ as in Section 2. As noted there, this changes the overall running time bound only by a constant factor.

Remark 3.7. *The rescaling formula* (10) *uses the denominator* $(1+\varepsilon)^2$, *whereas in [7] the denominator* $1 + \varepsilon$ *is used instead. This slightly different definition is needed in the proof of Lemma 3.16 in Section 3.4. Nevertheless, the analysis in [7] goes through essentially unchanged by choosing, as we did here, $\varepsilon = 1/(20n)$ instead of the choice $\varepsilon = 1/(11n)$ made in [7].*

In the analysis of Algorithm 2, we will refer to several lemmas from [7]. While the proofs are essentially given in [7], they differ slightly from the current context due to the different choice of rescaling, mentioned in Remark 3.7, and the different notation. For completeness we will provide proofs in the appendix.

3.3 Analysis

Let us define the ellipsoid

$$E(R) \overset{\text{def}}{=} \{x \in \mathbb{R}^n : x^\mathsf{T} R x \leq 1\}.$$

Further, let

$$\Sigma_\mu \overset{\text{def}}{=} \{w \in \mathbb{R}^n : w^\mathsf{T} x \geq 0 \quad \forall x \in B(f_\mu)\}, \qquad F_\mu \overset{\text{def}}{=} \Sigma_\mu \cap \mathbb{B}^n, \quad (11)$$

where $\mathbb{B}^n$ denotes the unit ball centered at the origin in $\mathbb{R}^n$.

Σ_μ is the set of normal vectors of hyperplanes that weakly separate 0 from $B(f_\mu)$. A vector in the interior of Σ_μ gives a strong separation, and verifies that $0 \notin B(f_\mu)$. This in turn implies that $f_\mu(S) < 0$ for some set $S \subseteq V$, and thus the minimum value of f is strictly less than the current estimate $-\mu$.

The main ideas of the analysis are showing that *(a)* the ellipsoid $E(R)$ contains the set F_μ at every iteration (Lemma 3.10), and that *(b)* the volume of $E(R)$ keeps decreasing by a constant factor at every rescaling (Lemma 3.12). For an integer valued f, one can lower bound the volume in terms of n and $L_{f,2}$, assuming that F_μ has a nonempty interior. Hence,

at termination one can conclude that the interior of F_μ is empty, which implies that $f_\mu \geq 0$, or equivalently, the minimum value of the function is $-\mu$ for the current μ.

The analysis below provides a slightly different argument than the volume analysis, by bounding the Q-norm of the bases used during the algorithm. This will be needed for the "pull-back" argument for finding a dual certificate of optimality in Section 3.4.

Clearly, GREEDYMIN(f_μ, w) can be used as a separation oracle for Σ_μ. Further, if $\mu' \geq \mu$, then by Lemma 3.2 $B(f_{\mu'}) \supseteq B(f_\mu)$, which implies that $\Sigma_{\mu'} \subseteq \Sigma_\mu$ and $F_{\mu'} \subseteq F_\mu$ by definition.

As in [7], for a convex set $X \subset \mathbb{R}^n$ and a vector $a \in \mathbb{R}^n$, we define the width

$$\text{width}_X(a) \stackrel{\text{def}}{=} \max\{a^\mathsf{T} z \colon z \in X\}.$$

Further, we define the condition number

$$\omega_\mu \stackrel{\text{def}}{=} \min_{x \in B(f_\mu) \setminus \{0\}} \frac{\text{width}_{F_\mu}(x)}{\|x\|_2}.$$

A key estimate for the running time analysis is the following.

Lemma 3.8. *Let f be an integer valued submodular function, and let $\mu \in \mathbb{Z}$ satisfy $\min_{S \subseteq V} f(S) < -\mu \leq \min\{0, f(V)\}$. Then*

$$\omega_\mu \geq \frac{1}{4nL_{f,2}}.$$

Proof. Note that, by Claim 2.2 and Lemma 3.4, $L_{f_\mu,2} \leq L_{f_\mu} \leq 4L_f \leq 4\sqrt{n}L_{f,2}$, hence $\|x\|_2 \leq 4\sqrt{n}L_{f,2}$ for every $x \in B(f_\mu)$. The claim follows by showing

$$\text{width}_{F_\mu}(x) \geq 1/\sqrt{n}. \tag{12}$$

To prove this, we note that the assumption of the lemma implies $0 \notin B(f_\mu)$. Let z denote the minimum norm point in $B(f_\mu)$, and let $\hat{z} = z/\|z\|_2$. Since z is the minimum norm point in $B(f_\mu)$, we have $z^\mathsf{T} x \geq \|z\|_2^2$ for all $x \in B(f_\mu)$. Thus for all $x \in B(f_\mu)$

$$\hat{z}^\mathsf{T} x \geq \|z\|_2,$$

and, in particular, this implies $\hat{z} \in F_\mu$. By Theorem 2.3, if S is the minimizer of f_μ, $f_\mu(S) \leq -1$, because f is integer valued and μ is a nonnegative integer such that $f(S) < -\mu$. It follows that $1 \leq |f_\mu(S)| \leq \sqrt{n}\|z\|_2$, thus $\hat{z}^\mathsf{T} x \geq 1/\sqrt{n}$. Since $\hat{z} \in F_\mu$, this provides the bound on $\text{width}_{F_\mu}(x)$ for every $x \in B(f_\mu)$. $\square$

Lemma 3.9. *At every stage of Algorithm RESCALING-SFM, $\Gamma \subseteq B(f_\mu)$. Furthermore, after t rescalings, the current matrix R is of the form*

$$R = \alpha I_n + \sum_{g \in \Gamma} \gamma_g \frac{gg^\mathsf{T}}{\|g\|_2}, \tag{13}$$

where $\alpha = \frac{1}{(1+\varepsilon)^{2t}}$ and $\gamma_g > 0$ for all $g \in \Gamma$.

Proof. To see that $\Gamma \subseteq B(f_\mu)$ throughout the algorithm, it suffices to observe that when a new element g is included in Γ, it is an extreme base for $B(f_\mu)$ for the current value of μ, so in particular it is an element of $B(f_\mu)$, and that throughout the algorithm $B(f_\mu)$ can only become inclusionwise larger, as the value of μ never decreases. The second part of the statement follows immediately from the recursive construction of R as defined in (10). $\square$

Lemmas 3.10, 3.11, and 3.13 below follow essentially from [7], thus we will often refer to [7] in the proofs. To enable the reader to check the arguments in [7], we first explain how the notation differs between the two papers. Specifically, in

Lemma 3.10 ([7, Lemma 10])**.** *Throughout Algorithm RESCALING-SFM, $F_\mu \subseteq E(R)$ holds.*

Proof. At initialization we have $F_\mu \subseteq \mathbb{B}^n = E(R)$, since initially $R = I_n$. During the algorithm the value of μ can increase (in the sliding von Neumann algorithm) or the matrix R can be updated (rescaling). In the former case, as noted above, F_μ becomes smaller (with respect to set inclusion) as μ increases in value, hence the property is maintained in these steps.

The proof that $F_\mu \subseteq E(R)$ whenever R is updated is essentially identical to the proof of Lemma 10 in [7]. For simplicity we provide a proof in Lemma A.1 in Appendix A.2. $\qquad\square$

Lemma 3.11. *Throughout Algorithm* RESCALING-SFM, $\|x\|_Q \geq \omega_\mu \|x\|_2$ *must hold for every* $x \in B(f) \setminus \{0\}$.

Proof. Since $F_\mu \subseteq E(R)$ by Lemma 3.10, for every $x \in B(f) \setminus \{0\}$ we have $\text{width}_{F_\mu}(x) \leq \text{width}_{E(R)}(x)$. Furthermore $\text{width}_{E(R)}(x) = \|x\|_Q$ (this is an easy fact, see [7, Lemma 2.15]). The statement now follows from the definition of ω_μ. $\qquad\square$

Lemma 3.12 ([7, Lemma 11]). *The determinant of* R *increases at least by a factor* $16/9$ *at every rescaling.*

Proof. Proof in Appendix A.2. $\qquad\square$

Lemma 3.13 ([7, Lemma 4.11]). *At any stage of Algorithm* RESCALING-SFM, *there exists a point* $\bar{g} \in \Gamma$ *such that*

$$\|\bar{g}\|_Q \sqrt{\det(R)^{1/n} - 1} \leq \|\bar{g}\|_2.$$

Proof. Proof in Appendix A.2. $\qquad\square$

Proof of Theorem 3.6. One possibility is that the algorithm terminates at step 6 because $y = 0$. In this case, observe that $y = 0$ is a conic combination of elements $g_1, \ldots, g_k \in B(f_\mu)$, so it is also a convex combination of those elements, thus showing that $0 \in B(f_\mu)$. It follows that $\min_S f_\mu(S) = 0$, hence $\min_S f(S) = -\mu = f(W)$, showing that W is a minimizer.

Otherwise, the algorithm performs $T = 5n \log(nL_{f,2})$ rescalings. Lemma 3.12 shows that after T rescalings, $\det(R) \geq (16/9)^T$. It follows from Lemma 3.13 that, after T rescalings, there exists a point $g_h \in B(f_\mu)$ with $\|g_h\|_Q < \|g_k\|_2/(4nL_{f,2})$. By Lemma 3.11 we conclude that $\omega_\mu \leq \|g_k\|_Q/\|g_k\|_2 < 1/(4nL_{f,2})$. Noting that μ is maintained to be an integer throughout the execution of the algorithm, Lemma 3.8 implies that $\min_{S \subseteq V} f(S) = -\mu$. Since the algorithm maintains a set W with $f(W) = -\mu$, we can conclude that W is a minimizer for (SFM). This shows that the algorithm terminates with a correct solution.

The algorithm calls the sliding von Neumann subroutine $T = O(n \log(nL_{f,2}))$ times; by Lemma 3.5 each call takes at most $\lceil 1/\varepsilon^2 \rceil = O(n^2)$ iterations. At the kth iteration of von Neumann, it takes time $O(n \cdot \text{EO} + n \log n)$ to run GREEDYMIN and time $O(k)$ to update the coefficients $x_1, \ldots, x_k$. These give a bound of $O(n^3 \cdot \text{EO} + n^4)$ for each sliding von Neumann subroutine.

Further, every rescaling has to compute $O(n^2)$ outer products $g_i g_i^\mathsf{T}$, add their weighted sum to R, and compute $Q = R^{-1}$. The computation is dominated by computing the outer products, which take altogether $O(n^4)$ time. Hence the iterations between two subsequent rescalings take time $O(n^3 \cdot \text{EO} + n^4)$, yielding the claimed complexity bound. $\qquad\square$

3.4 Finding approximate solutions via the geometric rescaling algorithm

The main purpose of this section is to show how RESCALING-SFM can be used to derive approximate solutions to (SFM) in time that depends logarithmically on the desired precision. This will also lead to a strongly polynomial-time based on RESCALING-SFM, as explained in Section 4. Furthermore, as a bi-product, we will also be able to generate dual certificates of optimality for (SFM) from the output of RESCALING-SFM.

Theorem 3.14. *Setting* $T = O(n \log(n\delta^{-1}))$ *in Algorithm 2, from its output one can compute a set* $W \subseteq V$ *and a point* $y \in B(f)$, *expressed as a convex combination of* $O(n)$ *extreme bases of* $B(f)$, *such that* $f(W) \leq y^-(V) + \delta L_f$. *In particular,* $f(W) \leq \min_{S \subseteq V} f(S) + \delta L_f$. *The overall running time is* $O((n^4 \cdot \text{EO} + n^5) \log(n\delta^{-1}))$.

To prove the above theorem we will use a *"pull-back"* technique. Recall that in RESCALING-SFM, we keep modifying the matrix Q defining the scalar product. Lemmas 3.12 and 3.13 guarantee that after t rescalings, we can identify a vector $g \in B(f_\mu)$ that has a small Q-norm for the current Q, and the bound decreases geometrically with t. Our key technical claim, Lemma 3.16, shows a constructive way to identify a vector $v \in B(f_\mu)$ with $\|v\|_2 \leq \|g\|_Q$. Provided a vector v with small 2-norm (and thus small 1-norm), we can easily satisfy the requirements of APPROX-SFM, using the following lemma.

Lemma 3.15. *Let $\mu \geq \max\{0, -f(V)\}$ and $W \subseteq V$ such that $f(W) = -\mu$. Let $\Gamma \subseteq B(f_\mu)$ finite and $\{\mu_g\}_{g \in \Gamma} \subseteq [0, \mu]$ be such that g is an extreme base of $B(f_{\mu_g})$ for all $g \in \Gamma$. Given $v = \sum_{g \in \Gamma} \lambda_g g$ where $\lambda_g \geq 0$, $g \in \Gamma$, and $\sum_{g \in \Gamma} \lambda_g = 1$, in time $O(n|\Gamma|)$ we can compute $y \in B(f)$, given as a convex combination of $|\Gamma|$ extreme bases of $B(f)$, such that*

$$f(W) \leq y^-(V) + \frac{\|v\|_1}{2}.$$

Proof. For all $g \in \Gamma$, let $\tilde{g}$ be the extreme base of $B(f)$ defined by the same ordering which defined the extreme base g of $B(f_{\mu_g})$. Define $y := \sum_{g \in \Gamma} \lambda_g \tilde{g}$.

Observe that, given $g \in \Gamma$, if $v_1, \ldots, v_n$ is the ordering defining g, then $\tilde{g}(v_1) = g(v_1) - \mu_g$, $\tilde{g}(v_j) = g(v_j)$ for $j = 2, \ldots, n-1$, and $\tilde{g}(v_n) = g(v_n) + \mu_g + f(V)$. Thus, computing $\tilde{g}$ for all $g \in \Gamma$ requires time $O(|\Gamma|)$ and computing y requires time $O(n|\Gamma|)$. Furthermore, we have that $\|\tilde{g} - g\|_1 = 2\mu_g + f(V) \leq 2\mu + f(V)$. This implies that

$$\|y\|_1 = \|\sum_{g \in \Gamma} \lambda_g \tilde{g}\|_1 \leq \|\sum_{g \in \Gamma} \lambda_g g\|_1 + \sum_{g \in \Gamma} \lambda_g \|\tilde{g} - g\|_1 \leq \|v\|_1 + 2\mu + f(V) = \|v\|_1 - 2f(W) + f(V).$$

Since $\|y\|_1 = f(V) - 2y^-(V)$, the above implies

$$f(W) \leq \frac{\|v\|_1 + f(V) - \|y\|_1}{2} = y^-(V) + \frac{\|v\|_1}{2}.$$

$\square$

Our next lemma enables pulling back a vector with small Q-norm to a vector with no larger 2-norm. This is done gradually, by pulling back at each rescaling of RESCALING-SFM. The lemma is not specific to the context of submodular function minimization. In our application, the columns of the matrix A will be the bases used in the current iteration of the sliding von Neumann algorithm. We also note that this technique is applicable to the general Full Support Image Algorithm in [7], enabling to find approximate solutions as well as dual certificates of infeasibility.

Lemma 3.16. *Let $A \in \mathbb{R}^{n \times p}$, $R \in \mathbb{S}_{++}^n$, $Q = R^{-1}$, and $\varepsilon > 0$. Let $x \in \mathbb{R}_+^p$ such that $y := \sum_{i=1}^p x_i \frac{a_i}{\|a_i\|_Q}$ satisfies $\|y\|_Q \leq \varepsilon$. Define*

$$R' \stackrel{\text{def}}{=} \frac{1}{(1+\varepsilon)^2} \left(R + \sum_{i=1}^p \frac{x_i}{\|a_i\|_Q^2} a_i a_i^\mathsf{T} \right), \tag{14}$$

and $Q' \stackrel{\text{def}}{=} (R')^{-1}$. For every $v \in \mathbb{R}^n$, there exists $\nu \in \mathbb{R}_+^p$ such that $\|v + A\nu\|_Q \leq \|v\|_{Q'}$. Moreover, such a vector ν can be computed in time $O(n^2 p)$.

Proof. For any given $v \in \mathbb{R}^n$, we define $u \stackrel{\text{def}}{=} \frac{1}{(1+\varepsilon)^2} RQ'v$ and let

$$\beta \stackrel{\text{def}}{=} \max_{i \in [p]} \frac{\langle a_i, u \rangle_Q}{\|a_i\|_Q}, \quad \nu_i \stackrel{\text{def}}{=} \frac{x_i}{\|a_i\|_Q} \left(\beta - \frac{\langle a_i, u \rangle_Q}{\|a_i\|_Q} \right) \quad \text{for } i \in [p]. \tag{15}$$

We will show that the statement is satisfied by the choice of $\nu \in \mathbb{R}_+^p$ defined above. These values can be clearly computed in $O(n^2 p)$ time.

First, we observe that, by substituting the definitions of R' and u, we obtain

$$v = R'QRQ'v = u + \sum_{i=1}^p x_i \frac{\langle a_i, u \rangle_Q}{\|a_i\|_Q} \frac{a_i}{\|a_i\|_Q},$$

which, from the definition of ν and β, implies that

$$v + A\nu = u + \beta y. \tag{16}$$

Next, notice that

$$
\begin{aligned}
\|v\|_{Q'} & = \sqrt{v^{\mathsf{T}} Q' R' Q' v} = \frac{1}{(1+\varepsilon)} \left((v^{\mathsf{T}} Q') R(Q'v) + v^{\mathsf{T}} Q' \left(\sum_{i=1}^{p} \frac{x_i}{\|a_i\|_Q^2} a_i a_i^{\mathsf{T}} \right) Q'v \right)^{\frac{1}{2}} \\
& \geq \frac{1}{(1+\varepsilon)} \left((v^{\mathsf{T}} Q' R) Q(RQ'v) \right)^{\frac{1}{2}} = (1+\varepsilon) \|u\|_Q.
\end{aligned}
$$

From the above and observing that $|\beta| \leq \|u\|_Q$, from the definition of β, we have

$$\|v + A\nu\|_Q \leq \|u\|_Q + |\beta| \|y\|_Q \leq (1+\varepsilon)\|u\|_Q \leq \|v\|_{Q'},$$

where the first inequality follows from (16) and the triangle inequality. $\qquad\square$

Remark 3.17. *If a vector y is expressed as a convex combination of κ extreme bases, then in time $O(n^2\kappa)$ we can reduce the number of extreme bases in the convex combination to $O(n)$ by Carathéodory's theorem.*

We are ready to prove Theorem 3.14, showing how APPROX-SFM can be implemented using RESCALING-SFM.

Proof of Theorem 3.14. For the second part of the statement, note that if we are given $y \in B(f)$ and $W \subset V$ such that $f(W) \leq y^-(V) + \delta L_f$, it follows from Theorem 2.1 that $f(W) \leq \min_{S \subseteq V} f(S) + \delta$.

We now describe how to find such a vector y and set W. Run algorithm RESCALING-SFM(f), setting the limit on the number of rescalings to a number $T = cn\log(n\delta^{-1})$ for some constant c to be specified later. At the end of the execution, we identified a value $\mu > 0$ and a set $W \subseteq V$ such that $f(W) = -\mu$. Let Γ be the set of points in $B(f_\mu)$ maintained throughout the algorithm.

By Lemma 3.13, for an appropriate choice of c, after T rescalings there exists $\bar{g} \in \Gamma$ such that

$$\|\bar{g}\|_Q \leq \frac{2\delta}{4\sqrt{n}} \|\bar{g}\|_2.$$

Let $\hat{g} = \bar{g}/\|\bar{g}\|_2$. The running time of RESCALING-SFM(f) with the above choice of T is $O((n^4\mathrm{EO} + n^5)\log(n\delta^{-1}))$. Note also that $|\Gamma| \in O(n^3\log(n\delta^{-1}))$, thus finding $\bar{g}$ requires time $O(n^5\log(n\delta^{-1}))$ to compute the Q-norms of all elements of Γ.

By applying Lemma 3.16 for T times (considering the rescaling matrices used in the algorithm in reverse order), we can find a vector $\nu \in \mathbb{R}_+^\Gamma$ such that $\|\hat{g} + \sum_{g \in \Gamma} \nu_g g\|_2 \leq \|\hat{g}\|_Q$. Recall that each rescaling matrix is defined by at most n^2 vectors in Γ, therefore each application of Lemma 3.16 requires time $O(n^4)$ (assuming that the matrices Q and R used at every rescaling are saved in memory so we do not need to recompute them). Thus, overall, the time required to compute ν is $O(n^5\log(n\delta^{-1}))$.

Define $\alpha = 1 + \|\bar{g}\|_2 \cdot \sum_{i=1}^{h} \nu_g$, and $\lambda \in \mathbb{R}_+^\Gamma$ by

$$
\lambda_g = \begin{cases} \dfrac{\|\bar{g}\|_2 \nu_g}{\alpha}, & g \in \Gamma \setminus \{\bar{g}\} \\[2mm] \dfrac{1 + \|\bar{g}\|_2 \nu_{\bar{g}}}{\alpha}, & g = \bar{g} \end{cases}
$$

Define $v := \sum_{g \in \Gamma} \lambda_g g$. Observe that $\sum_{g \in \Gamma} \lambda_g = 1$, thus $v \in B(f_\mu)$. Computing v requires time $O(n^4\log(n\delta^{-1}))$, since we need to sum $|\Gamma|$ n-dimensional vectors.

Furthermore,

$$\|v\|_1 \leq \sqrt{n}\|v\|_2 = \sqrt{n}\frac{\|\bar{g}\|_2}{\alpha}\left\|\hat{g} + \sum_{g \in \Gamma} \nu_g g\right\|_2 \leq \sqrt{n} L_{f_\mu, 2}\|\hat{g}\|_Q \leq 2\delta L_f,$$

where the last inequality follows from the fact that $L_{f_\mu,2} \le L_{f_\mu} \le 4L_f$ by Lemma 3.4. By Lemma 3.15, in time $O(n^4 \log(n\delta^{-1}))$ we can compute $y \in B(f)$ satisfying $f(W) \le y^-(V) + \frac{\|v\|_1}{2} \le y^-(V) + \delta L_f$.

Finally, since y is expressed as a convex combination of $|\Gamma| \in O(n^3 \log(n\delta^{-1}))$ extreme bases, by Remark 3.17 we can express y as a convex combination of $O(n)$ extreme bases in time $O(n^5 \log(n\delta^{-1}))$. $\qquad\square$

Remark 3.18. The bound $O(n^5 \log(n\delta^{-1}))$ for computing ν in the above proof was assuming $O(n^2)$ time for computing Q-scalar products $\langle g, u\rangle_Q$. We note that this can be easily improved by a factor n: we can assume that Qg was precomputed and stored during RESCALING-SFM for all bases g used during the sequence of rescalings. Indeed, it was necessary to compute the norms $\|g\|_Q$ in the sliding von Neumann algorithm. Thus, the bound improves to $O(n^4 \log(n\delta^{-1}))$; however, this does not change the overall running time estimate.

Finding a dual certificate in Rescaling-SFM For an integer valued f, a pair W and y satisfying the requirements of APPROX-SFM$(f, 1/L_f)$ are an optimal pair of primal and dual solutions as in Theorem 2.1. Hence the algorithm of Theorem 3.14 for $\delta = 1/L_f$ provides a dual certificate of optimality in time $O((n^4 \cdot \mathrm{EO} + n^5) \log(nL_f))$, the same as the complexity bound as in Theorem 3.6 (using that $L_f \le \sqrt{n}L_{f,2}$).

Approximate infeasibility certificates for conic linear programs Consider the feasibility problem for conic linear programs of the form

$$A^\mathsf{T} z > 0, \tag{17}$$

where A is a $p \times n$ matrix of rank n, whose columns $a_1, \dots, a_p$ are assumed to have unit ℓ_2-norm. It is well known that (17) is feasible if and only if 0 is not contained in the convex hull of $a_1, \dots, a_p$. Alternatively, one can say that (17) is infeasible if and only if $\min\{\|v\|_2 : v \in \mathrm{conv}(A)\} = 0$. For any $\delta > 0$, Lemma 3.16 allows us to compute in time $O(n^2p^2 \log(\delta^{-1}))$, via the Image Algorithm of [7], either a solution to (17), or a δ-approximate certificate of infeasibility, that is, a point $v \in \mathrm{conv}(A)$ such that $\|v\|_2 \le \delta$.

Just as Algorithm 2 in this paper, the Image Algorithm maintains matrices $R, Q \in \mathbb{S}^n_{++}$, $Q = R^{-1}$, which are updated at every rescaling. Initially $R = Q = I_n$. Between each rescaling, the algorithm uses von Neumann's method to compute, in at most $\lceil \varepsilon^{-2}\rceil$ iterations, a vector $x \in \mathbb{R}^p_+$, $\bar{e}^\mathsf{T} x = 1$, such that the point $y := \sum_{i=1}^p x_i \frac{a_i}{\|a_i\|_Q}$ satisfies either $A^\mathsf{T} Qy > 0$ or $\|y\|_Q \le \varepsilon$, for $\varepsilon := 1/(20n)$. In the former case, the algorithm stops since $z := Qy$ is a feasible solution to (17). In the latter case, R is replaced with the matrix R' given in (14), and $Q = R^{-1}$ is recomputed. After $T = O(n \log(\delta^{-1}))$ rescalings, by Lemmas 3.12 and 3.13 there exists a column a_k, $k \in [p]$, such that $\|a_k\|_Q \le \delta$. By applying Lemma 3.16 for T times (considering the rescaling matrices used in the algorithm in reverse order), we can find a vector $\nu \in \mathbb{R}^p_+$ such that $\|a_k + A\nu\|_2 \le \|a_k\|_Q \le \delta$. If we define

$$v = \frac{a_k + A\nu}{1 + \|\nu\|_1},$$

then $v \in \mathrm{conv}(A)$ and $\|v\|_2 \le \delta/(1 + \|\nu\|_1) \le \delta$. Each call of von Neumann algorithm requires $O(np^2)$ arithmetic operations (see proof of Theorem 3.2 in [7]), whereas each application of Lemma 3.16 requires $O(n^2p)$ operations. It follows that the overall running time of the Image algorithm is $O(n^2p^2 \log(\delta^{-1}))$.

4 Strongly polynomial algorithms

In this section, we provide a general scheme to convert an approximate SFM algorithm to a strongly polynomial one. We assume that the SFM algorithm is provided via the following oracle.

> Oracle Approx-SFM
> **Input:** A submodular function $f : 2^V \to \mathbb{R}$ and $\delta > 0$.
> **Output:** A set $W \subseteq V$, and a vector $y \in B(f)$ such that
>
> $$f(W) \leq y^-(V) + \delta L_f.$$
>
> Further, y is given as a convex combination of $O(n)$ bases of $B(f)$.

Recall from Theorem 2.1 that $f(S) \geq y^-(V)$ for any $S \subseteq V$. Consequently, the set W returned by the oracle is within δL_f from the optimal solution to (SFM).

The output of the oracle will be used via the following lemma.

Lemma 4.1. *Let* $y \in B(f)$ *and* $W \subseteq V$ *such that* $f(W) \leq y^-(V) + \delta L_f$. *If* $y(v) < -\delta L_f$, *then* v *must be contained in every minimizer of* f.

Proof. Let $S \subseteq V \setminus \{v\}$. Then $f(S) \geq y(S) \geq y^-(V \setminus \{v\}) \geq f(W) - y(v) - \delta L_f > f(W)$. This shows that S cannot be an optimal solution to (SFM). $\qquad\square$

For a $y \in B(f)$ returned by the approximate oracle and an element v as in the lemma, minimizing f can be reduced to minimizing the contraction $f' : 2^{V \setminus \{v\}} \to \mathbb{Z}$, defined as $f'(S) \overset{\text{def}}{=} f(S \cup \{v\}) - f(\{v\})$.

The existence of such a v will not always be guaranteed. If we cannot immediately reduce the size of the ground set by contraction, we will still be able to identify structural properties of optimal solutions using the next simple lemma:

Lemma 4.2. *Let* $y \in B(f)$, $U \subseteq V$, *and* $v \in V \setminus U$. *Assume that* $y(v) > -y^-(V \setminus U)$. *Then any minimizer to* (SFM) *that contains* v *must contain some element of* U.

Proof. Let $S \subseteq V \setminus U$, $v \in S$. Then $f(S) \geq y(v) + y^-(V \setminus U) > 0 = f(\emptyset)$, hence S cannot be a minimizer. $\qquad\square$

During the algorithm we will maintain a partial order $\preceq$, where $u \preceq v$ means that every minimizer to (SFM) containing v must also contain u. New relations will be derived using the lemma above. To guarantee either a contraction or a new relation in every iteration, we need to call the oracle Approx-SFM on a modified version $f^\downarrow$ instead of the original f, introduced in the next subsection.

Section 4.1 introduces the machinery and notation. The strongly polynomial scheme, Algorithm 3, is given in Section 4.2. It repeatedly calls Approx-SFM with $\delta = 1/(3n^3)$ and contracts elements of the ground set, or learns new relations concerning the optimal solutions to (SFM). Theorem 4.7 gives the running time bound that is polynomial in n and the running time of the approximate oracle calls. Section 4.3 presents an enhanced version of the algorithm with improved running times.

Approximate oracle implementations The oracle Approx-SFM has been used implicitly or explicitly in several papers in the literature:

- Our geometric rescaling algorithm (Algorithm 2), in time $O((n^4 \cdot \text{EO} + n^5) \log(n\delta^{-1}))$, as explained in Theorem 3.14 in Section 3.4.

- The conditional gradient method. Theorem 2.5 asserts that in time $O((n^2 \cdot \text{EO} + n^2 \log n)\delta^{-2})$, we can obtain a suitable y as a convex combination of $O(n/\delta^2)$ extreme bases. According to Remark 3.17, we need $O(n^3/\delta^2)$ arithmetic operations to reduce the support to $O(n)$ extreme bases. This gives an overall running time of $O((n^2 \cdot \text{EO} + n^3)\delta^{-2})$.

- The Fujishige-Wolfe algorithm in time $O((n^2 \cdot \text{EO} + n^4)\delta^{-2})$ (Theorem 2.6).

- The Iwata-Fleischer-Fujishige weakly polynomial scaling algorithm [24], in time $O(n^5 \cdot EO \log(n\delta^{-1}))$.[1]

---

[1] The authors give both a weakly and a strongly polynomial algorithm; here we are referring to the weakly polynomial algorithm of Section 3 in [24]. The observation we reported here is explicitly stated by the authors in the last paragraph of Section 3 in [24].

- The Iwata-Orlin weakly polynomial algorithm [25], in time $O((n^4 \cdot \text{EO} + n^5) \log(n\delta^{-1}))$.[2]
- the Sidford-Lee-Wong cutting plane method in $O(n^2 \cdot \text{EO} \log(n\delta^{-1}) + n^3 \log^{O(1)}(n\delta^{-1}))$; this is shown in Section 5.

4.1 Ring families and the structure of optimal solutions

From every call to APPROX-SFM, our algorithm will either find an element $v \in V$ as in Lemma 4.1, or at least one pair (v, u) such that every optimal solution to (SFM) containing v must also contain u; these pairs will be derived using Lemma 4.2. In the former case, we can simply reduce the size of the ground set by contracting such an element. To make use of the pairs (v, u) obtained in the latter case, we use the following framework, first introduced by Iwata, Fleischer, and Fujishige [24], and subsequently used in several other algorithms, e.g. in [25, 29].

Directed graph and partial order We maintain a directed graph $D = (V, F)$, with the property that if $(u, v) \in F$, and $S \subseteq V$ is an optimal solution to (SFM) with $u \in S$, then $v \in S$. In other words, for every minimizer $S \subseteq V$ to (SFM), $\delta_F^+(S) = \emptyset$.

We can assume that D is a acyclic. Indeed, given the set of vertices C of any strongly connected component of D, every minimizer of f must either contain C or be disjoint from C, hence we can contract all vertices of C into a single new node v_C. That is, we replace V with $V' := V \cup \{v_C\} \setminus C$, and f with the function $f' : V' \to \mathbb{R}$ defined by

$$f'(S) \stackrel{\text{def}}{=} \begin{cases} f(S), & \text{if } S \subseteq V' \setminus \{v_C\}, \\ f(S \cup C \setminus \{v_C\}), & \text{if } v_C \in S \subseteq V'. \end{cases} \tag{18}$$

The acyclic graph $D = (V, F)$ defines a partial order $\preceq_F$. We have $u \preceq_F v$ if there exists a directed path in F from v to u (so, in particular, $v \preceq_F v$ for all $v \in V$). We say that an ordering of the vertices is *consistent* with $\preceq_F$, if u is ordered before v whenever $u \preceq_F v$. We use the simpler notation $u \preceq v$ if F is clear from the context.

The *transitive closure* of the acyclic graph $D = (V, F)$ is the graph $\bar{D} = (V, \bar{F})$, where $(u, v) \in \bar{F}$ if and only if there is a directed path from u to v in F. D is said to be *transitive* if it coincides with its transitive closure. Clearly D and its transitive closure define the same partial order $\preceq_F$. We maintain the digraph $D = (V, F)$ transitive throughout the algorithm.

Ring families In terms of the partial order $\preceq$, every optimal solution to (SFM) is *a lower ideal* in the poset $(V, \preceq)$. Let $\mathcal{F}$ denote the set of all lower ideals, namely, $S \in \mathcal{F}$ if $v \in S$ and $u \preceq v$ implies $u \in S$. Equivalently, in terms of the digraph $D = (V, F)$, we let

$$\mathcal{F} := \{X \subseteq V \ : \ \delta_F^+(X) = \emptyset\}. \tag{19}$$

Thus, the family $\mathcal{F}$ contains all minimizers. The family $\mathcal{F} \subseteq 2^V$ is a *ring family*, that is, for every $X, Y \subseteq V$, if $X, Y \in \mathcal{F}$ then $X \cap Y, X \cup Y \in \mathcal{F}$.

We note that one can define submodular functions over ring families instead of the entire 2^V. Submodular function minimization over ring families has been well-studied and can be reduced to submodular function minimization over 2^V, see [15, Section 3] and [33, Chapter 49].

The following definitions and results are similar to those in [33, Section 49.3]. For any $v \in V$, let

$$v^\downarrow \stackrel{\text{def}}{=} \{u \in V \ : \ u \preceq v\}, \quad v^\uparrow \stackrel{\text{def}}{=} \{u \in V \ : \ v \preceq u\}$$

Note that $v^\downarrow$ is the unique minimal element of $\mathcal{F}$ containing v, and $V \setminus v^\uparrow$ is the unique maximal element of $\mathcal{F}$ not containing v. If $D = (V, F)$ is transitive, then for all $v \in V$
$v^\downarrow = \{v\} \cup \{u \in V \ : \ (v, u) \in F\}$ and $v^\uparrow = \{v\} \cup \{u \in V \ : \ (u, v) \in F\}$.

---

[2]This is not explicitly stated in [25]. However, the analysis in that paper shows that, in time $O((n^4 \cdot \text{EO} + n^5) \log(n\delta^{-1}))$, one obtains a set $W \subseteq V$ and a point $x \in B(f)$ such that $x(W) = f(W)$, $x(v) \geq 0$ for all $v \in V \setminus W$, and $\Phi(x) := \sum_{v \in W}(x^+(v))^2 \leq \delta^2 L_f^2/n$. This implies that $f(W) = x^-(W) + x^+(W) \leq x^-(V) + \sqrt{n\Phi(x)} \leq x^-(V) + \delta L_f$.

For every $X \subseteq V$, we define

$$X^{\downarrow} \overset{\text{def}}{=} \bigcup_{v \in X} v^{\downarrow}, \quad X^{\uparrow} \overset{\text{def}}{=} \bigcup_{v \in X} v^{\uparrow}$$

that is, $X^{\downarrow}$ is the unique minimal element of $\mathcal{F}$ containing X and $V \setminus X^{\uparrow}$ is the unique maximal element of $\mathcal{F}$ disjoint from X.

The lower bound and $f^{\downarrow}$ Let us define the lower bound function

$$\ell(v) \overset{\text{def}}{=} f((V \setminus v^{\uparrow}) \cup \{v\}) - f(V \setminus v^{\uparrow}),$$

as in [15, (3.95)]. Observe that $(V \setminus v^{\uparrow}) \cup \{v\} \in \mathcal{F}$ for every $v \in V$. It is easy to see that $z \geq \ell$ for every extreme basis $z \in B(f)$ defined by an ordering consistent with $\preceq$. The next claim shows that ℓ yields lower bounds on the function values of f. In particular, if $\ell(v) > 0$, then no minimizer of (SFM) contains v.

Claim 4.3. *For every $X, Y \in \mathcal{F}$ with $X \subseteq Y$, we have $\ell(Y \setminus X) + f(X) \leq f(Y)$. In particular, if we let $P = \{v \in V : \ell(v) > 0\}$, we have that $f(X) > f(X \setminus P^{\uparrow})$ for all $X \in \mathcal{F}$.*

Proof. For the first part of the claim, let us take the elements of $Y \setminus X$ in a consistent order with $\preceq$ as $z_1, \ldots, z_r$. Then, $Z_i = X \cup \{z_1, \ldots, z_i\} \in \mathcal{F}$ for each $i \in [r]$, and $Z_i \subseteq (V \setminus z_i^{\uparrow}) \cup \{z_i\}$. Submodularity implies $f(Z_i) - f(Z_{i-1}) \geq \ell(z_i)$. The claim follows by adding up all these inequalities.

The second part of the claim follows from first part of the statement if we choose $Y = X \setminus P^{\uparrow}$ (note that $X \setminus P^{\uparrow} \in \mathcal{F}$ since it is the intersection of X and $V \setminus P^{\uparrow}$, which are both elements of $\mathcal{F}$). $\qquad\square$

Next, we define a submodular function $f^{\downarrow} : 2^V \to \mathbb{Z}$ as

$$f^{\downarrow}(X) \overset{\text{def}}{=} f(X^{\downarrow}) - \ell^-(X^{\downarrow} \setminus X).$$

As shown in Lemma 4.4 below, this function is submodular, $f^{\downarrow}(S) = f(S)$ for every $S \in \mathcal{F}$, and $f^{\downarrow}(S) \geq f(S^{\downarrow})$ for all $S \subseteq V$. Consequently, $f^{\downarrow}$ takes its minimum value on sets in $\mathcal{F}$, and if all minimizers in f are contained in $\mathcal{F}$, then minimizing $f^{\downarrow}$ is equivalent to minimizing f.

An important advantage of using $f^{\downarrow}$ instead of f is that $z \geq \ell^-$ for all bases $z \in B(f^{\downarrow})$ (Claim 4.5), and therefore the complexity parameter $L_{f^{\downarrow}}$ is bounded by $|\ell^-(V)|$ (Claim 4.6). This can be arbitrarily smaller than L_f, and therefore ApproX-SFM$(f^{\downarrow}, \delta)$ returns a higher accuracy solution than ApproX-SFM(f, δ).

We note that the definition of $f^{\downarrow}$ only uses the function values of f on $\mathcal{F}$. In fact, similar definitions can be used to reduce submodular function minimization on ring families to minimizing submodular functions on 2^V, see e.g. [33, Section 49.3].

Lemma 4.4. *The function $f^{\downarrow}$ is submodular on 2^V with $f^{\downarrow}(S) \geq f(S^{\downarrow})$ for all $S \subseteq V$ and $f^{\downarrow}(S) = f(S)$ for every $S \in \mathcal{F}$. Consequently, minimizing f on the ring family $\mathcal{F}$ is equivalent to minimizing $f^{\downarrow}$ on 2^V. The complexity of GreedyMin$(f^{\downarrow}, w)$ can be bounded by $O(n \cdot \text{EO} + n^2)$, where EO is the complexity to evaluating f.*

Proof of Lemma 4.4. To prove that $f^{\downarrow}$ is submodular, we show that $f^{\downarrow} = b$ where $b : 2^V \to \mathbb{Z}$ is defined by

$$b(X) \overset{\text{def}}{=} \min\{f(Y) - \ell^-(Y \setminus X) : X \subseteq Y, Y \in \mathcal{F}\}.$$

Observe that, by Claim 4.3, $\ell^-(Y) \leq f(Y)$ for every $Y \in \mathcal{F}$. In particular, b is the so called *contraction of f by vector ℓ^-*, defined in [15, (3.23)], where it is shown that b is submodular.

Let us show that $b = f^{\downarrow}$. By definition, $b(X) \leq f^{\downarrow}(X)$ for any $X \subseteq V$. Consider now $Y \in \mathcal{F}$ such that $Y \supsetneq X^{\downarrow}$. Again by Claim 4.3, we have $f(Y) \geq f(X^{\downarrow}) + \ell^-(Y \setminus X^{\downarrow})$; this implies that $X^{\downarrow}$ is the minimizer in the definition of b, therefore $f^{\downarrow} = b$.

The nonpositivity of ℓ^- gives that $f^{\downarrow}(S) \geq f(S^{\downarrow})$ for all $S \subsetneq V$; it is clear that $f^{\downarrow}(S) = f(S)$ for all $S \in \mathcal{F}$, $S \neq V$. Regarding the complexity of GreedyMin, one needs to compute the values of $f^{\downarrow}(\{v_1, \ldots, v_i\})$ for every $i \in [n]$ for a given order of the vertices;

thus, we need to find the sets $S_i = \{v_1, \ldots, v_i\}^\downarrow$. When moving from i to $i+1$, we need to compute $S_{i+1} = S_i \cup v_i^\downarrow$, which can be done in $O(n)$ time because F is maintained as a transitive digraph. Adding the ℓ^- values also take $O(n)$ time for each set. Hence, we obtain an overhead $O(n^2)$ over the $O(n \cdot \mathrm{EO})$ oracle queries and $O(n \log n)$ time for sorting the ground set. $\qquad\square$

Claim 4.5. *For every $v \in V$, $f^\downarrow(V) - f^\downarrow(V \setminus \{v\}) \in \{\ell^-(v), \ell(v)\}$. In particular, $y(v) \geq \ell^-(v)$ for every $y \in B(f^\downarrow)$.*

Proof. If $V \setminus \{v\} \notin \mathcal{F}$, then $(V \setminus \{v\})^\downarrow = V$, therefore $f^\downarrow(V) - f^\downarrow(V \setminus \{v\}) = f(V) - f(V) + \ell^-(v)$ by definition of $f^\downarrow$. If $V \setminus \{v\} \in \mathcal{F}$, then $v^\uparrow = \{v\}$, therefore $f^\downarrow(V) - f^\downarrow(V \setminus \{v\}) = f(V) - f(V \setminus \{v\}) = f((V \setminus v^\uparrow) \cup \{v\}) - f(V \setminus v^\uparrow) = \ell(v)$. For the last part, note that for any extreme base g of $B(f^\downarrow)$, $g(v) = f^\downarrow(S) - f^\downarrow(S \setminus \{v\})$ for some $S \subseteq V$ containing v, and by submodularity $f^\downarrow(S) - f^\downarrow(S \setminus \{v\}) \geq f^\downarrow(V) - f^\downarrow(V \setminus \{v\}) \geq \ell^-(v)$. $\qquad\square$

We remark that the second part of the Claim 4.5 also follows from [15, Theorem 3.6]. The next claim shows that $|\ell^-(V)|$ is a good approximation of the complexity parameter $L_f^\downarrow$.

Claim 4.6. *Assuming that $f(V) \leq 0$, we have $|\ell^-(V)|/|V| \leq L_{f^\downarrow} \leq 2|\ell^-(V)|$.*

Proof. By Claim 4.5, $y(v) \geq \ell^-(v)$ for every $v \in V$ and $y \in B(f^\downarrow)$, therefore $y^-(V) \geq \ell^-(V)$. Together with $f^\downarrow(V) \leq 0$, this implies $\|y\|_1 = f^\downarrow(V) - 2y^-(V) \leq -2\ell^-(V)$.

For the lower bound, let us choose $v \in V$ with lowest value of $\ell^-(v)$. Thus, $|\ell^-(v)| \geq |\ell^-(V)|/|V|$. Consider any extreme base g of $B(f^\downarrow)$ from an order where v comes last. Then by the first part of Claim 4.5 $g(v) \in \{\ell^-(v), \ell(v)\}$, hence $|g(v)| \geq |\ell^-(v)|$, which implies $L_{f^\downarrow} \geq \|g\|_1 \geq |g(v)| \geq |\ell^-(v)| \geq |\ell^-(V)|/|V|$. $\qquad\square$

4.2 The basic strongly polynomial scheme

Algorithm 3 The basic strongly polynomial algorithm

Input: A submodular function $f : 2^V \to \mathbb{R}$ with $f(V) \leq 0$, and $\delta > 0$.
Output: An optimal solution to (SFM).
 1: Initialize $F := \emptyset$, $T := \emptyset$.
 2: **while** $\ell^-(V) < 0$ **do**
 3:     Call Approx-SFM$(f^\downarrow, \delta)$ to obtain W and $y \in B(f^\downarrow)$, represented as a
        convex combination $y = \sum_{i=1}^k x_i g_i$.
 4:     **for** $z \in V$ such that $f(V \setminus z^\uparrow) > -|V| \cdot y^-(V \setminus z^\uparrow)$ **do**
 5:         **for** $i = 1, \ldots, k$ **do**, $g_i' \leftarrow$ Push$(g_i, f^\downarrow, z^\uparrow)$.
 6:     $y' := \sum_{i=1}^k x_i g_i'$.
 7:     **for** $v \in V \setminus z^\uparrow$ such that $y'(v) > -y'^-(V \setminus z^\uparrow)$ **do**
 8:         add arc (v, z) to F.
 9:     **for** $v \in V$ such that $y(v) < 2\ell^-(V)\delta$ **do** $\triangleright$ contraction
10:         Replace f by $f(S \cup v^\downarrow) - f(v^\downarrow)$ on the ground set $V := V \setminus v^\downarrow$.
11:         Set $f(V) := \min\{0, f(V)\}$.
12:         Set $T := T \cup v^\downarrow$.
13:     Contract all strongly connected components of F to single nodes using
        (18).
14:     Replace F by its transitive closure
        return the pre-image of T in the original ground set.

Algorithm 3 builds a ring family $\mathcal{F}$ represented by a directed graph F with the property that $\mathcal{F}$ contains all optimal solutions to (SFM).

The subroutine $\text{PUSH}(g, h, S)$ is defined for a submodular function $h : 2^V \to \mathbb{Z}$, an extreme basis $g \in B(h)$ given by an ordering of the elements of V, and a subset $S \subseteq V$. It returns another extreme basis $g' \in B(g)$ by moving the set S backward in the ordering defining g. That is, every element of $V \setminus S$ will precede every element in S, and the relative ordering of the elements inside $V \setminus S$ as well as those inside S remain unchanged.

In Theorem 4.7, we show that for $\delta = 1/(3n^3)$, the algorithm terminates within n^2 iterations. Let $\text{AO}^\downarrow(f, \delta)$ denote the maximum of $AO(f^\downarrow, \delta)$, where $f^\downarrow$ ranges over all possible choices of ring families $\mathcal{F}$.

Recall that Lemma 4.4 provides an upper bound $O(n^2 \cdot \text{EO} + n^2)$ on the running time of $\text{GREEDY-MIN}(f^\downarrow, w)$, in contrast to the running time bound $O(n^2 \cdot \text{EO} + n \log n)$ on $\text{GREEDY-MIN}(f, w)$. Thus, $\text{AO}^\downarrow(f, \delta)$ can be upper bounded by the worst case running time bound on $AO(f, \delta)$ plus $O(n/\log n)$ times the worst case bound on the number of calls to the greedy algorithm in $AO(f, \delta)$.

Theorem 4.7. *Using $\delta = 1/(3n^3)$, Algorithm 3 finds the unique minimal optimal solution to* (SFM) *in time $O(n^2(\text{AO}^\downarrow(f, 1/(3n^3)) + n^4 \cdot \text{EO} + n^5)$.*

Using the bounds from Theorems 2.5 and 2.6, we obtain $O(n^{10} \cdot \text{EO} + n^{11})$ using the conditional gradient algorithm, and $O(n^{11} \cdot \text{EO} + n^{12})$ using the Fujishige-Wolfe algorithm. While these running times are high degree polynomials, we emphasize that they can be obtained by repeated applications of simple iterative methods, without using any form of scaling.

Theorem 3.14 gives a running time $O((n^6 \cdot \text{EO} + n^7) \log n)$ using the RESCALING-SFM algorithm. In Section 4.3, we give an enhanced version of the algorithm with running time $O((n^5 \cdot \text{EO} + n^6) \log^2 n)$.

Let us now give an overview of Algorithm 3. The ground set is reduced by contracting elements that must be included in every optimal solution. The set T represents the current set of contracted elements. Thus, the submodular function at the current stage will be defined as $f(S \cup T) - f(T)$ for the original input function f, with the possible exception of $f(V)$. Therefore, the complexity of evaluating the current f is still EO. We will use n below to denote the size of the original ground set V.

Once $\ell(v) \geq 0$ for all $v \in V$, we conclude that $S = \emptyset$ is the minimizer of the current function; we return T as the unique minimal minimizer of the original function. Otherwise, as long as $\ell^-(V) < 0$, each main iteration calls the oracle $\text{APPROX-SFM}(f^\downarrow, \delta)$. Two types of contractions are used. All cycles in F can be contracted to single elements using the operation (18), since an optimal solution can contain either all or no element of a cycle (line 13). The other type of contraction (in line 10) reduces the size of the ground set by eliminating elements that must be contained in every optimal solution.

The other main step of the algorithm is adding new arcs to F. The following lemma shows the validity of these steps and that either of these operations should occur in every iteration.

Lemma 4.8. *The following properties are valid throughout Algorithm 3.*

(i) *All minimizers to* (SFM) *are contained in the ring family $\mathcal{F}$.*

(ii) *If $\ell^-(V) = 0$ then $f(Y) \geq 0$ for all $Y \subseteq V$.*

(iii) *Every arc (v, z) added to F in line 8 satisfies the property that every minimizer that contains v must also contain z.*

(iv) *Every $v \in V$ contracted in line 10 must be contained in all minimizers of* (SFM).

(v) *If $\delta \leq 1/(3n^3)$, then every iteration either contracts an element or adds a new arc to F.*

Proof. We show all properties by induction on the number of iterations of the algorithm. Part (i) is trivially true at the beginning, since $\mathcal{F} = \emptyset$. Part (iii) guarantees that it remains true whenever F is extended (and thus $\mathcal{F}$ is restricted to a smaller ring family).

For (ii), Claim 4.3 shows $f(Y) \geq \ell(Y) \geq \ell^-(V)$ for any $Y \in \mathcal{F}$. Our assumption is $\ell^-(V) = 0$, that is $\ell \geq 0$, and thus $f(Y) \geq 0$ for all $Y \in \mathcal{F}$. By (i), it follows that $f(Y) \geq 0$ on every set $S \subseteq V$. For the rest of the proof, we assume that $\ell^-(V) < 0$.

For (iii), Lemma 4.2 implies that for v selected in line 7, every minimizer of $f^\downarrow$ that contains v, must also contain some element of $z^\uparrow$. Using (i), and the definition of $f^\downarrow$, the minimizers of f are also minimizers of $f^\downarrow$, and are contained in $\mathcal{F}$. Thus, if S minimizes $f(S)$ and $S \cap z^\uparrow \neq \emptyset$, then $z \in S$. It follows that every minimizer of f containing v must also contain z, therefore the new arc (v, z) is valid.

For (iv), consider a v such that $y(v) < 2\ell^-(V)\delta$ as in line 10. Lemma 4.1 and Claim 4.6 imply that v is contained in every minimizer of $f^\downarrow$. By (i), $v^\downarrow$ must be contained in all elements of $v^\downarrow$.

We now turn to the proof of (v). Note that, if the algorithm enters the while loop when $|V| = 1$, say $V = \{v\}$, then $y(v) = f(V) = \ell(\{v\}) = \ell^-(V) < 2\ell^-(V)\delta$, so the algorithm contracts v in line 10, and subsequently terminates. Thus, we can assume that $|V| \geq 2$. The argument is as follows: we note that if no element is contracted in line 10, then the condition in line 9 yields

$$y(v) \geq 2\ell^-(V)\delta \quad \forall v \in V. \tag{20}$$

From this property, we will derive that the condition in line 4 is true for at least one $z \in V$. Finally, we will show that in this case, the condition in line 7 is true for some $v \in V \setminus z^\uparrow$.

In more detail, if no element is contracted in line 10, then from (20) we get the lower bound

$$f^\downarrow(S) \geq y(S) \geq y^-(V) \geq 2|V| \cdot \ell^-(V)\delta \quad \forall S \subseteq V. \tag{21}$$

Since $f^\downarrow(S) = f(S)$ for $S \in \mathcal{F}$ by the definition of $f^\downarrow$, and $\mathcal{F}$ contains all minimizers of f by (i), we have that

$$f(S) \geq 2|V| \cdot \ell^-(V)\delta \quad \forall S \subseteq V. \tag{22}$$

Next, we show that if (21) holds, then at least one $z \in V$ satisfies

$$f(V \setminus z^\uparrow) > |V| \cdot |y^-(V)|, \tag{23}$$

a bound which is slightly stronger than the condition $f(V \setminus z^\uparrow) > -|V| \cdot y^-(V \setminus z^\uparrow)$ in line 4. Select $z \in V$ such that $\ell(z)$ is the most negative possible. In particular, $\ell(z) \leq \ell^-(V)/|V|$. By (22), we have

$$\frac{\ell^-(V)}{|V|} \geq \ell(z) = f((V \setminus z^\uparrow) \cup \{z\}) - f(V \setminus z^\uparrow) \geq 2|V| \cdot \ell^-(V)\delta - f(V \setminus z^\uparrow). \tag{24}$$

Consequently,

$$f(V \setminus z^\uparrow) \geq |V| \cdot |\ell^-(V)| \cdot \left(\frac{1}{|V|^2} - 2\delta\right).$$

From the assumption $\delta \leq 1/(3n^3) \leq 1/(3|V|^3)$ we obtain $1/|V|^2 - 2\delta > 2|V|\delta$ since $|V| \geq 2$. Therefore (23) follows since

$$f(V \setminus z^\uparrow) > 2|V|^2 \cdot |\ell^-(V)|\delta \geq |V| \cdot |y^-(V)|.$$

In the final inequality we used (21) (recalling that $\ell^-(V)$ and $y^-(V)0$ are non-positive).

It is left to show that if $z \in V$ satisfies the condition in line 4, then at least one $v \in V \setminus z^\uparrow$ will be selected in line 7. By construction and from the fact that $V \setminus z^\uparrow \in \mathcal{F}$, for every $i \in [k]$ we have $g_i'(V \setminus z^\uparrow) = f^\downarrow(V \setminus z^\uparrow) = f(V \setminus z^\uparrow)$, and $g_i'(u) \geq g_i(u)$ for every $u \in V \setminus z^\uparrow$. It follows that $y'(V \setminus z^\uparrow) = f(V \setminus z^\uparrow)$ and $y'(u) \geq y(u)$ for all $u \in V \setminus z^\uparrow$.

Assume that $f(V \setminus z^\uparrow) > -|V| \cdot y^-(V \setminus z^\uparrow)$, as in the condition in line 4. It follows that

$$y'(V \setminus z^\uparrow) = f(V \setminus z^\uparrow) > -|V| \cdot y^-(V \setminus z^\uparrow) \geq -|V| \cdot y'^-(V \setminus z^\uparrow).$$

This in turn implies the existence of $v \in V \setminus z^\uparrow$ such that $y'(v) > -y'^-(V \setminus z^\uparrow)$ in line 7. $\square$

Proof of Theorem 4.7. Lemma 4.8 justifies the contraction steps and the addition of new arcs to F, and shows that the number of main iterations is at most n^2. Let us also note that after every contraction, we decrease the value of $f(V)$ if it becomes positive (that is, if $f(V) > f(v^\uparrow)$ before the contraction of v). This operation clearly maintains submodularity.

It is also safe in the sense that it may not lead to an incorrect output with respect to the original function. Indeed, note that at termination the algorithm returns the current set of T, which are elements that must be contained in every minimizer of the original function. Hence, the algorithm outputs the unique minimal solution to (SFM). On the other hand, if $f(V)$ was ever decreased, then we decrease it to the same value as $f(\emptyset)$. Therefore it can never become the unique minimizer. If the algorithm terminates with the entire ground set V, then it follows that $f(V)$ was never decreased during the algorithm.

Let us now estimate the running time. Besides the calls to APPROX-SFM, the running time is dominated by the operations $\text{PUSH}(g_i', f^\downarrow, z^\uparrow)$, which altogether require $O(n^2 \cdot \text{EO} + n^3)$ for every iteration, and this is required $O(n^2)$ times. Every time an arc (u, v) is added to F, recomputing the transitive closure requires to add arcs from u and all its predecessors to v and all its descendants. This requires $O(n^2)$ operations per added arc, so it requires $O(n^4)$ operations overall, which is within the stated running time bound. $\qquad\square$

4.3   Speeding up the algorithm

The algorithm described in the previous section needs to identify $O(n^2)$ arcs in F. In the worst case, each iteration may only identify a single arc, resulting in $O(n^2)$ calls to APPROX-SFM.

On the other hand, if we were able to guarantee that $|\ell^-(z)|$ is within a factor $O(n^b)$ from $|\ell^-(V)|$ for a constant fraction of all $z \in V$ for some constant $b \geq 1$, the analysis in the proof of Lemma 4.8 implies that for $\delta = 1/O(n^{b+2})$ we would guarantee $f(V \setminus z^\uparrow) \geq -ny^-(V \setminus z^\uparrow)$ for all such $z \in V$. Thus, after running APPROX-SFM$(f^\downarrow, 1/O(n^{b+2}))$, we could extend F by $\Theta(n)$ new arcs.

If this property held in all iterations, then $O(n)$ calls to APPROX-SFM would suffice. However, the number of $z \in V$ with $|\ell^-(z)|$ value "close" to $|\ell^-(V)|$ can be $o(n)$. To deal with this situation, we apply the *bucketing* technique of Lee, Sidford, and Wong [29]. Instead of the entire V, we restrict our function in every iteration to a suitably chosen $\bar{V} \subseteq V$, and run APPROX-SFM restricted to this set with $\delta = n^{-O(\log n)}$. We will obtain $\Theta(|\bar{V}|)$ new arcs in this iteration. Thus, if APPROX-SFM has running time $O((|\bar{V}|^4 \cdot \text{EO} + |\bar{V}|^5) \log^2 n)$, then the amortized cost of extending F by an arc will be $O((n^3 \cdot \text{EO} + n^4) \log^2 n)$.

We note that this improvement is only applicable if $\text{AO}(f, \delta)$ depends logarithmically on $1/\delta$. Since δ can be quasi-polynomial, the conditional gradient or Fujishige-Wolfe methods would not even be polynomial in this framework.

Another speedup technique, also adapted from [29], enables to save on the running time of recomputing the extreme bases in line 6 of Algorithm 3. We can identify the new arc (z, v) by recomputing only one of the g_i's instead of all of them, at the expense of requiring a higher accuracy $2\kappa\delta$ instead of δ. Here, κ denotes an upper bound on the number of extreme bases in the convex combination, hence $\kappa \in O(n)$ by assumption.

The following lemma adapts the argument in Section 15.4.1 in [29].

Lemma 4.9. *Let $f : 2^V \to \mathbb{Z}$ be a submodular function, $\mathcal{F}$ a ring family containing all minimizers of f, and $f^\downarrow : 2^V \to \mathbb{Z}$ be the corresponding function defined by f and $\mathcal{F}$. Then in $O(n \cdot \text{EO})$ time we can find a nonempty subset $\bar{V} \subseteq V$ and a positive integer $b = O(\log n)$, such that*

- *For every $z \in V \setminus \bar{V}$, we have $\ell(z) > 2\ell^-(V)/(4\kappa n)^{4b}$.*
- *There exist at least $|\bar{V}|/2$ distinct $z \in \bar{V}$ such that $\ell(z) \leq 2\ell^-(V)/(4\kappa n)^{4b-4}$.*

Proof. Let us define $V^t \overset{\text{def}}{=} \{z \in V : \ell(z) \leq 2\ell^-(V)/(4\kappa n)^{4t}\}$ for $t = 1, 2, \ldots$. Note that $V^1 \subseteq V^2 \subseteq \ldots$, because $\ell^-(V) \leq 0$. Furthermore, $V^1 \neq \emptyset$, as it contains z with the smallest $\ell(z)$ value. Let b be the smallest value such that $|V^b| \leq 2|V^{b-1}|$. Thus, $b = O(\log n)$, and choosing $\bar{V} = V^b$ satisfies both requirements. $\qquad\square$

For the set $\bar{V}$ and value b as in the lemma, let $\bar{f} : 2^{\bar{V}} \to \mathbb{Z}$ denote the restriction of $f^\downarrow$ to the ground set $\bar{V}$, and let us set

$$\bar{\delta} \overset{\text{def}}{=} \frac{1}{(4\kappa n)^{4b}}, \qquad \delta \overset{\text{def}}{=} \frac{2n^2 + 1}{(4\kappa n)^{4b}}, \tag{25}$$

Let us call APPROX-SFM$(\bar{f}, \bar{\delta})$ to obtain the vector $\bar{y} \in B(\bar{f})$ defined as a convex combination of extreme bases $\bar{g}_1, \ldots, \bar{g}_k \in B(\bar{f})$ with $k \leq \kappa$, and a set $W \subseteq \bar{V}$ such that $\bar{f}(W) \leq \bar{y}^-(\bar{V}) + \bar{\delta} L_{\bar{f}}$.

Let us now extend $\bar{y} \in \mathbb{R}^{\bar{V}}$ to $y \in \mathbb{R}^V$ as follows. For $v \in \bar{V}$, we let $y(v) = \bar{y}(v)$. Then, consider an arbitrary order $v_1, \ldots, v_{n-|\bar{V}|}$ of $V \setminus \bar{V}$, and set $y(v_j) := f^\downarrow(\bar{V} \cup \{v_1, \ldots, v_j\}) - f^\downarrow(\bar{V} \cup \{v_1, \ldots, v_{j-1}\})$. Let us also define $g_1, \ldots, g_k \in \mathbb{R}^V$, by $g_i(v) = \bar{g}_i(v)$ for $v \in \bar{V}$, $g_i(v) = y(v)$ for $v \in V \setminus \bar{V}$ $(i = 1, \ldots, k)$. Note that, by definition, $g_1, \ldots, g_k$ are extreme bases of $B(f^\downarrow)$, and y is a convex combination of $g_1, \ldots, g_k$.

Lemma 4.10. *For the vector y and set W as above, we have that $y \in B(f^\downarrow)$, and $f^\downarrow(W) \leq y^-(V) + \delta L_{f^\downarrow}$.*

Proof. By definition $f^\downarrow(W) = \bar{f}(W)$ and $L_{\bar{f}} \leq L_{f^\downarrow}$, because $\bar{f}$ is a restriction of $f^\downarrow$. Therefore,

$$f^\downarrow(W) \leq \bar{y}^-(\bar{V}) + \bar{\delta} L_{f^\downarrow}.$$

Claim 4.11. $y(v) \geq \ell^-(v)$ *for every* $v \in V \setminus \bar{V}$.

Proof. If $v = v_j$, then $y(v) = f^\downarrow(\bar{V} \cup \{v_1, \ldots, v_j\}) - f^\downarrow(\bar{V} \cup \{v_1, \ldots, v_{j-1}\}) \geq f^\downarrow(V) - f^\downarrow(V \setminus \{v\})$ by submodularity. Further, $f^\downarrow(V) - f^\downarrow(V \setminus \{v\}) \geq \ell^-(v)$ by Claim 4.5. $\qquad \square$

We have $y^-(V) = \bar{y}^-(\bar{V}) + y^-(V \setminus \bar{V})$. By the choice of $\bar{V}$, we have

$$\ell^-(v) \geq 2\ell^-(V)/(4\kappa n)^{4b} = 2\ell^-(V)\bar{\delta} \quad \forall v \in V \setminus \bar{V}.$$

Using the claim above, we get $y^-(V \setminus \bar{V}) \geq 2n\ell^-(V)\bar{\delta}$. Thus, $y^-(V) \geq \bar{y}^-(\bar{V}) + 2n\ell^-(V)\bar{\delta} \geq \bar{y}^-(\bar{V}) - 2n^2\bar{\delta} L_{f^\downarrow}$. Here, the last inequality used the lower bound in Claim 4.6. Consequently,

$$f^\downarrow(W) \leq y^-(V) + (2n^2 + 1)\bar{\delta} L_{f^\downarrow} = y^-(V) + \delta L_{f^\downarrow}.$$

$\square$

This proof shows that we can implement APPROX-SFM(f, δ) by calling APPROX-SFM$(\bar{f}, \bar{\delta})$, and adding the remaining $V \setminus \bar{V}$ elements by $O(n)$ value oracle queries for $f^\downarrow$, and $O(n\kappa)$ arithmetic operations. This gives a running time overhead $O(n \cdot \text{EO} + n^2)$.

We make two modifications to Algorithm 3 as follows. Firstly, in every iteration, we compute $\bar{V}$ and b as in Lemma 4.9, and use this modified implementation of APPROX-SFM with δ as defined in (25).

Secondly, we modify the selection of z as in line 4 to

$$f(V \setminus z^\uparrow) > -2k|V| \cdot y^-(V \setminus z^\uparrow). \tag{26}$$

In line 6, we change the computation of y' as follows. We first compute $x_i(f(V \setminus z^\uparrow) - g_i(V \setminus z^\uparrow))$ for all $i \in [k]$; w.l.o.g., assume the maximum is taken for $i = 1$. We compute $g_1' \leftarrow \text{PUSH}(g_1, f^\downarrow, z^\uparrow)$ and define

$$y' := y + x_1(g_1' - g_1).$$

The rest of the algorithm remains unchanged.

Theorem 4.12. *The above described modification of Algorithm 3 finds an optimal solution to* (SFM) *in time* $O(n \cdot \text{AO}^\downarrow(f, n^{-O(\log n)}) + n^3 \cdot \text{EO} + n^4)$. *Using the implementation with* RESCALING-SFM, *the running time is* $O((n^5 \cdot \text{EO} + n^6) \log^2 n)$.

Proof. Note that properties (i)-(iv) of Lemma 4.8 are maintained throughout the algorithm, just as before. We need to prove the following analogue of Lemma 4.8(v).

Claim 4.13. *After every call of the approximation oracle, either at least one node is contracted at line 10, or at least $\frac{1}{2}|\bar{V}|$ new arcs are added to F.*

Proof. We will first show that, if no node is contracted at line 10, then (26) holds for at least half of the elements z of $\bar{V}$. Subsequently, we will show that, for every z satisfying (26), there exists $v \in V \setminus z^{\uparrow}$ satisfying the condition in line 7.

If no node is contracted at line 10, then as in the proof of Lemma 4.8(v), we can assume that $y(v) \geq 2\ell^-(V)\delta$ for all $v \in V$; further, (21) and (22) hold. By Lemma 4.9 and our choice of $\bar{V}$, at least half of the elements z of $\bar{V}$ satisfy

$$\ell(z) \leq 2\ell^-(V)/(4\kappa n)^{4b-4} = 2\ell^-(V)\delta(4\kappa n)^4/(2n^2 + 1).$$

Consider any such z. As in (24), the assumption (22) implies that

$$f(V \setminus z^{\uparrow}) \geq 2|V| \cdot \ell^-(V)\delta - \ell(z) \geq 2|V|^2 \cdot |\ell^-(V)|\delta \left(\frac{(4\kappa n)^4}{|V|^2(2n^2 + 1)} - \frac{1}{|V|} \right).$$

Using (21), $2|V| \cdot |\ell^-(V)|\delta \geq |y^-(V)|$, and it is easy to see that the expression in the brackets is $> 2\kappa \geq 2k$ for $n \geq 2$. Thus,

$$f(V \setminus z^{\uparrow}) \geq 2k|V| \cdot |y^-(V)|,$$

implying (26).

Let us now show that for any z satisfying (26), there exists a $v \in V \setminus z^{\uparrow}$ such that $y'(v) > -y'^-(V \setminus z)$. Recall our assumption that $i = 1$ maximizes $x_i(f(V \setminus z^{\uparrow}) - g_i(V \setminus z^{\uparrow}))$, and therefore

$$x_1(f(V \setminus z^{\uparrow}) - g_1(V \setminus z^{\uparrow})) \geq \frac{1}{k} \left(f(V \setminus z^{\uparrow}) - y(V \setminus z^{\uparrow}) \right).$$

Also note that g_1' is defined so that $g_1'(V \setminus z^{\uparrow}) = f^{\downarrow}(V \setminus z^{\uparrow}) = f(V \setminus z^{\uparrow})$. We get

$$
\begin{aligned}
y'(V \setminus z^{\uparrow}) &= y(V \setminus z^{\uparrow}) + x_1(g_1'(V \setminus z^{\uparrow}) - g_1(V \setminus z^{\uparrow})) \\
&= y(V \setminus z^{\uparrow}) + x_1(f(V \setminus z^{\uparrow}) - g_1(V \setminus z^{\uparrow})) \\
&\geq \frac{k-1}{k} y(V \setminus z^{\uparrow}) + \frac{1}{k} f(V \setminus z^{\uparrow})
\end{aligned}
$$

Using (26), we obtain

$$y'(V \setminus z^{\uparrow}) > -|V| \cdot y^-(V \setminus z^{\uparrow}) \geq -|V| \cdot y'^-(V \setminus z^{\uparrow}),$$

guaranteeing the existence of $v \in V \setminus z^{\uparrow}$ such that $y'(v) > -y'^-(V \setminus z^{\uparrow})$ in line 7. $\qquad\square$

The running time of APPROX-SFM$(\bar{f}, \bar{\delta})$ is AO$(\bar{f}, \bar{\delta})$. There are at most n iterations where a node gets contracted; the total cost of the oracle calls in these iterations can be bounded by nAO$^{\downarrow}(f, \bar{\delta})$.

Consider now the iterations when no nodes get contracted. In these iterations, the amortized cost of an oracle call per new arc is 2AO$(\bar{f}, \bar{\delta})/|\bar{V}|$. Since AO depends at least linearly on $|\bar{V}|$, this can be upper bounded by AO$^{\downarrow}(f, \bar{\delta})/|V|$. Hence, the total time of the oracle calls is $O(n \cdot$ AO$(f, \bar{\delta}))$.

After every call, there is an overhead $O(n \cdot \text{EO} + n^2)$, totalling $O(n^3 \cdot \text{EO} + n^4)$. For every arc identified, it takes $O(\text{EO} + \kappa)$ to identify which g_i' needs to be computed, and it takes $O(n \cdot \text{EO} + n^2)$ time to compute this g_i'; this takes $O(n^3 \cdot \text{EO} + n^4)$ time overall. As in the proof of Theorem 4.7, recomputing the transitive closures requires $O(n^4)$ operations over the entire execution of the algorithm.

Finally, if we consider APPROX-SFM provided by algorithm RESCALING-SFM as in Theorem 3.14, we obtain a running time bound $O((n^5 \cdot \text{EO} + n^6) \log^2 n)$. $\qquad\square$

5 Cutting plane method

The current best cutting plane method for finding a point in a convex set $C \subseteq \mathbb{R}^n$ provided by a separation oracle is due to Lee, Sidford, and Wong [29]. Assume that, for $R > \varepsilon > 0$, C is contained in a ball of radius R centered at the origin and contains some ball of radius ε. Let $\kappa := nR/\varepsilon$. In general, cutting plane methods maintain at each iteration k a "simple"

convex set $K^{(k)}$ such that $C \subseteq K^{(k)}$, and select a candidate point $x^{(k)} \in K^{(k)}$. With a call to the separation oracle for C, the method either terminates if $x^{(k)} \in \text{int}(C)$, or else it generates a valid inequality for C that weakly separates $x^{(k)}$ from $\text{int}(C)$, and uses this inequality to generate a new relaxation $K^{(k+1)}$ with smaller volume and a new candidate point $x^{(k+1)}$.

Lee, Sidford, and Wong's method [29] maintains $K^{(k)}$ as the intersection of $O(n)$ valid inequalities, and guarantees that the volume of $K^{(k)}$ decreases by a constant factor at every separation oracle call. This ensures that the number of oracle calls is bounded by $O(n \log \kappa)$. The overall running time is $O(n \cdot \text{SO} \log \kappa + n^3 \log^{O(1)} \kappa)$, where SO is the complexity of an (exact) separation oracle.

By contrast, the central-cut ellipsoid method maintains $K^{(k)}$ to be an ellipsoid, where the volume of $K^{(k)}$ decreases by a constant factor every n separation oracle calls (see [20], Lemma 3.2.10), ensuring a $O(n^2 \log \kappa)$ bound on the number of oracle calls. The overall running time is $O((n^2 \cdot \text{SO} + n^4) \log \kappa)^3$

In Part III [29], Lee, Sidford, and Wong apply their cutting plane method to submodular function minimization, and obtain a strongly polynomial running time bound of $O(n^3 \log^2 n \cdot \text{EO} + n^4 \log^{O(1)} n)$, which is currently the best (see [29, Section 15.4]). This is obtained by combining their cutting plane algorithm with an improved version of the combinatorial framework of ring families; one of their important new contributions is the bucketing technique we also use in Section 4.

In this section, we present an alternative way of applying their cutting plane method to SFM. We prove the same running time bound in a substantially simplified way. Firstly, instead of using the Lovász extension as in [20] and in [29], we apply the cutting plane method to find a feasible solution in F_μ, as defined in (11). We use the sliding technique as in Section 3 for the cutting plane algorithm. Secondly, we employ the combinatorial framework in a black-box manner, by implementing APPROX-SFM via the Lee-Sidford-Wong algorithm. The combinatorial interpretation of the certificate returned by the cutting plane method turns out to be much easier than in [29].

Weakly polynomial algorithm Let us start by exhibiting a weakly polynomial $O(n^2 \log(nL_{f,2}) \cdot \text{EO} + n^3 \log^{O(1)}(nL_{f,2}))$ algorithm for SFM, which is the same as the running time in [29]. We use a slight modification of the cutting plane algorithm [29, Section 6.4, Algorithm 2].

We start with $\mu = \max\{0, -f(V)\}$, and maintain a set $W \subseteq V$ with $f(W) = -\mu$ throughout. The algorithm seeks a point in $\text{int}(F_\mu)$, and the initial relaxation $K^{(0)}$ is the hypercube centered at the origin of side length $2\sqrt{n}$. For the current iterate $x^{(k)}$, GREEDYMIN$(f_\mu, x^{(k)})$ is used as the separation oracle for $\text{int}(F_\mu)$, which returns an extreme base g of $B(f_\mu)$. If $g^\mathsf{T} x^{(k)} > 0$, then $x^{(k)} \in \text{int}(F_\mu)$, thus $x^{(k)}$ is feasible. In this case, instead of terminating, we modify the value of μ as in the sliding von Neumann algorithm. That is, we set $W = \text{MINSET}(f_\mu, x^{(k)})$, and set the new value $\mu' = -f(W)$. From Lemma 3.3, we see that $x^{(k)} \notin \text{int}(F_{\mu'})$. Thus, we can continue with adding a new cutting plane. Note that $F_{\mu'} \subseteq F_\mu$ if $\mu' > \mu$, hence the current relaxation $K^{(k)}$ remains valid, because $F_{\mu'} \subseteq F_\mu \subseteq K^{(k)}$. (Again, this is similar to the sliding objective technique, although we are changing all constraints of the polytope simultaneously.) When $-\mu$ is the minimum value of f, L_μ has no points in the interior, therefore we stop when the volume of the current relaxation becomes too small.

In this setting, we have $\text{SO} = n \cdot \text{EO} + n \log n$. For every value of μ, $F_\mu \subseteq \mathbb{B}^n$ by definition, and Lemma 3.8 implies that, as long as $\min_{S \subseteq V} f(S) < -\mu$, F_μ contains a ball of radius $1/(4nL_{f,2})$. Hence, $\kappa = O(nL_{f,2})$, giving the desired running time bound.

Let us note that the framework just described does not depend on the specifics of how the relaxations $K^{(k)}$ or the candidate point $x^{(k)}$ are constructed, but it can be applied to any cutting plane method. For example, for the central-cut ellipsoid method, the above framework gives an $O((n^3 \cdot \text{EO} + n^4) \log(nL_{f,2}))$ algorithm for submodular function minimization. Indeed, as previously mentioned, the running time of the ellipsoid method is $O((n^2 \cdot \text{SO} + n^4) \log \kappa)$, and as before $\text{SO} = n \cdot \text{EO} + n \log n$ and $\kappa = O(nL_{f,2})$. Interestingly, even such a simple and direct use of the standard ellipsoid method, compared to

---

[3]While not explicitly stated in [20], this running time bound follows immediately from the analysis.

the usual approach of minimizing the Lovász extension, provides a running time that is a factor n lower than any weakly-polynomial SFM-algorithm known prior to the work of Lee-Sidford-Wong [29].

Strongly polynomial algorithm Let us now show an $O(n^2 \log(n\delta^{-1})\text{EO} + n^3 \log^{O(1)}(n\delta^{-1}))$ implementation of APPROX-SFM(f, δ) using the Lee-Sidford-Wong cutting plane method. We use Theorem 31 from [29]. For $K = F_\mu$ (for any value of μ), by definition $F_\mu \subseteq \mathbb{B}^n \subseteq \mathbb{B}^n_\infty(1)$, that is, $R = 1$. Recall, as described above, that we slide μ every time we find a feasible solution in F_μ. The following lemma shows that the algorithm always returns a thin direction as follows.

Theorem 5.1 ([29, Theorem 31]). *For any $\varepsilon \in [0, 1]$, in expected time $O(n \log(n/\varepsilon) \cdot \text{SO} + n^3 \log^{O(1)}(n/\varepsilon))$, the (sliding) cutting plane method returns a value μ, and constraints $a_i^\mathsf{T} x \geq b_i$ for $i \in [h]$, where $h = O(n)$, $\|a_i\|_2 = 1$, which are all valid for F_μ. Each of these constraint is either an original box constraint, that is $x_j \geq -1$ or $-x_j \geq -1$, or an inequality returned by the separation oracle. Let P denote the intersection of these hyperplanes.*

Further, we obtain non-negative numbers $t_1, t_2, t_3, \ldots, t_h$ with $t_1 = 1$, and a point $x^* \in P$, which satisfy the following:

(a) $\|x^*\|_2 \leq 3\sqrt{n}$,

(b) $\left\|\sum_{i=1}^h t_i a_i\right\|_2 = O(\sqrt{n}\varepsilon \log(1/\varepsilon))$,

(c) $a_1^\mathsf{T} x^* - b_1 \leq \varepsilon$,

(d) $\left(\sum_{i=2}^h t_i a_i\right)^\mathsf{T} x^* - \sum_{i=2}^h t_i b_i \leq O(\sqrt{n}\varepsilon \log(1/\varepsilon))$.

The output certifies that the region $P \cap \mathbb{B}^n_\infty(1)$ has small width in the direction of a_1. Indeed, let $\bar{a} := \sum_{i=2}^h t_i a_i$ and $\bar{b} := \sum_{i=2}^h t_i b_i$. By Cauchy-Schwartz and (b), for all $x \in \mathbb{R}^n$, $|(a_1 + \bar{a})^\mathsf{T} x| \leq \|x\| O(\sqrt{n}\varepsilon \log(1/\varepsilon))$, so $b_1 + \bar{b} \leq (a_1 + \bar{a})^\mathsf{T} x^* \leq O(n\varepsilon \log(1/\varepsilon))$. By (c) and (d), $-b_1 - \bar{b} \leq -(a_1 + \bar{a})^\mathsf{T} x^* + \varepsilon + O(\sqrt{n}\varepsilon \log(1/\varepsilon)) \leq O(n\varepsilon \log(1/\varepsilon))$. This shows $|b_1 + \bar{b}| = O(n\varepsilon \log(1/\varepsilon))$. It follows that, for every $x \in P \cap \mathbb{B}^n_\infty(1)$,

$$b_1 \leq a_1^\mathsf{T} x \leq a_1^\mathsf{T} x + \bar{a}^\mathsf{T} x - \bar{b} = (a_1 + \bar{a}^\mathsf{T})x - \bar{b} - b_1 + b_1 \leq O(n\varepsilon \log(1/\varepsilon)) + b_1.$$

We show that for an appropriately chosen ε, this can be used to implement APPROX-SFM(f, δ).

Lemma 5.2. *For an appropriate ε such that $\delta = \Omega(n^{3/2}\varepsilon \log(1/\varepsilon))$, from the output of the cutting plane method we can obtain W and y as required for APPROX-SFM(f, δ), that is, $f(W) \leq y^-(V) + \delta L_f$.*

Proof. Let $[h] = I_b \cup I_s$, where I_b is the set of indices i such that $a_i^\mathsf{T} x_i \geq b_i$ is a box constraint, and I_s is the set of indices corresponding to constraints from the separation oracle. Each constraint in I_s is of the form $a_i = g_i/\|g_i\|_2$ and $b_i = 0$, where g_i is an extreme base of $B(f_{\mu_i})$, where $\mu_i \leq \mu$ was the value of μ at the time when this cutting plane was added. The lemma will easily follow from the next claim.

Claim 5.3. *The index 1 is in I_s, and $\left\|\sum_{i \in I_s} t_i a_i\right\|_2 = O(n\varepsilon \log(1/\varepsilon))$.*

Proof. First, we show that $1 \in I_s$. For a contradiction, assume that $1 \in I_b$, that is, $a_1 = e_j$ or $a_1 = -e_j$ for some $j \in [n]$ and $b_1 = -1$. As noted above, $|b_1 + \bar{b}| = O(n\varepsilon \log(1/\varepsilon))$; hence, $\bar{b} > 0$ follows (for small enough ε). This is a contradiction, since $b_i = -1$ for all $i \in I_b$, and $b_i = 0$ for all $i \in I_s$.

Thus, $1 \in I_s$, and therefore $b_1 = 0$. Thus, $|\bar{b}| = O(n\varepsilon \log(1/\varepsilon))$. Again, this implies that $\sum_{i \in I_b} t_i = O(n\varepsilon \log(1/\varepsilon))$. Together with $\left\|\sum_{i \in I_b} t_i a_i + \sum_{i \in I_s} t_i a_i\right\|_2 = O(\sqrt{n}\varepsilon \log(1/\varepsilon))$ from (b), we get that $\left\|\sum_{i \in I_s} t_i a_i\right\|_2 = O(n\varepsilon \log(1/\varepsilon))$, as required. $\square$

Let $v = \left(\sum_{i \in I_s} \frac{t_i}{\|g_i\|_2} g_i \right) / \left(\sum_{i \in I_s} \frac{t_i}{\|g_i\|_2} \right)$. Since $1 \in I_s$, we have $\sum_{i \in I_s} \frac{t_i}{\|g_i\|_2} \geq \frac{1}{L_{f,2}} \geq \frac{1}{L_f}$. Hence, it follows that

$$\|v\|_1 \leq \sqrt{n} \|v\|_2 \leq L_f \sqrt{n} \left\| \sum_{i \in I_s} t_i a_i \right\|_2 = O(L_f n^{3/2} \varepsilon \log(1/\varepsilon)) \leq 2\delta L_f.$$

Then, Lemma 3.15 is applicable to provide the certificate for APPROX-SFM(f, δ). Note that the set W with $f(W) = -\mu$ has been maintained during the cutting plane algorithm. $\qquad \square$

Combining Lemma 5.2 with Theorem 4.12, and noting that $\kappa = O(n)$, we obtain the running time bound $O(n^3 \log^2(n)\text{EO} + n^4 \log^{O(1)}(n))$.

Comparison to the Lee-Sidford-Wong SFM algorithm Let us now compare the above approach to the SFM algorithm described in [29, Part III]. We employ the same cutting plane method, and a common framework is using ring families; our bucketing argument has been adapted from [29].

Their combinatorial framework is more complex than ours: upper bounds analogous to the lower bounds $\ell(z)$ are needed, and accordingly, their algorithm identifies both outgoing and incoming arcs, as well as removes elements which cannot be contained in any minimizer. The simple trick that enables us to work only with lower bounds, and identify only incoming arcs is repeatedly truncating the value of $f(V)$; thus, we can bound $L_{f^\downarrow}$ in terms of $\ell^-(V)$, as in Claim 4.6.

Our black-box approach clearly separates the combinatorial argument from the cutting plane method, which is used only inside the oracle. In contrast, these two ingredients cannot be clearly separated in [29]. They use the cutting plane method for the formulation using the Lovász extension and do not use sliding. Then, they transform the cutting plane certificate to identify a small norm convex combination in the base polytope. This is analogous to, but substantially more complicated than our Lemma 5.2. In particular, it is not always possible to identify such a combination, since the constraints of the feasible region can have large coefficients. In such cases, these large coefficients can be used to fix some of the variables to 0 and 1, and hence make progress in terms of the ring family. In contrast, the certificate from our sliding cutting plane algorithm on F_μ can be straightforwardly translated in Lemma 5.2 to satisfy the requirements of the approximate oracle.

6 Variants of the geometric rescaling algorithm

The framework of Algorithm 2 is fairly general, in the sense that both the first-order method used to generate short convex combinations of normalized vectors of $B(f_\mu)$ and the rescaling used to update the matrix $Q \in \mathbb{S}_{++}^n$ can be replaced with other alternatives. Here we discuss some of these variants.

6.1 Replacing von Neumann with Fujishige-Wolfe

Within Algorithm 2, the role of Algorithm 1 is to determine a point y with $\|y\|_Q \in O(1/n)$ such that y is a convex combination of points of the form $g/\|g\|_Q$, $g \in B(f_\mu)$. Any algorithm that can produce such output in time polynomial in n can be used in place of von Neumann algorithm. In particular, Fujishige-Wolfe can be adapted to the rescaling setting of Algorithm 2.

As for von Neumann's algorithm, the only modifications that are required are the following. First of all, given a matrix $Q \in \mathbb{S}_{++}^n$, we use the Q-norm and Q-scalar product to normalize the elements of $B(f_\mu)$. We maintain a set X of affinely independent elements of the form $g/\|g\|_Q$, $g \in B(f_\mu)$. At every major cycle – that is, when X is a corral – the algorithm computes the projection y of the origin to the affine hull of X, which belongs to the relative interior of $\text{conv}(X)$ because X is a corral. At any major cycle we are only interested in knowing whether or not $Qy \in \text{int}(\Sigma_\mu)$. Therefore, we call GREEDYMIN(f_μ, Qy), obtaining a minimizer $g \in B(f_\mu)$ as output. If $\langle g, y \rangle_Q \leq 0$, then we set $X' := X \cup \{g/\|g\|_Q\}$. If $\langle g, y \rangle_Q > 0$ (that is, $Qy \in \text{int}(\Sigma_\mu)$), then we determine $W \subset V$ such that $f(W) < -\mu$

and slide f by setting $\mu := -f(W)$, just as in line 6 of the sliding-von Neumann algorithm (Algorithm 1). We update g as in line 9, and then set $X' = X \cup \{g/\|g\|_Q\}$. In both cases, we proceed with either another major cycle, if X' is still a corral, or a minor cycle if X' is not a corral. Note, in particular, that in both cases we only increase the set X when $\langle g, y \rangle_Q \leq 0$. The next statement follows immediately from [4, Theorem 4] and from the same arguments used in the proof of Lemma 3.5.

Theorem 6.1. *Given a value* $\mu \geq \max\{0, -f(V)\}$, *a set* $W \subseteq V$ *with* $f(W) = -\mu$, *a matrix* $Q \in \mathbb{S}^n_{++}$, *and an* $\varepsilon > 0$, *within* $O(1/\varepsilon^2)$ *iterations (major and minor cycles), Wolfe's algorithm computes a value* $\mu' \geq \mu$ *and a set* $W' \subseteq V$ *with* $f(W') = -\mu'$, *bases* $g_1, \ldots, g_k \in B(f_{\mu'})$, $x \in \mathbb{R}^k$, $y \in \mathbb{R}^n$ *such that* $k \leq n$, $y = \sum_{i=1}^k x_i g_i / \|g_i\|_Q$, $\vec{e}^\mathsf{T} x = 1$, $x \geq 0$, *and* $\|y\|_Q \leq \varepsilon$.

In particular, within Algorithm 2, we need to fix $\varepsilon = 1/(20n)$, which implies that the number of iterations required for each call of Fujishige-Wolfe is $O(n^2)$. Each iteration requires at most one call to GREEDYMIN, needing $O(n)$ oracle calls, plus $O(n^2)$ arithmetic operations (the iteration complexity is dominated by the computation of the projection of the origin onto the affine hull of X). Given that the number of rescalings needed is $O(n \log(n L_{f,2}))$, the total running time of Algorithm 2 where the sliding-von Neumann algorithm is replaced with Fujishige-Wolfe is $O((n^4 \cdot \mathrm{EO} + n^5) \log(n L_{f,2}))$.

6.2 Rank-1 rescalings

The multi-rank rescaling (10) used in Algorithm 2 can be replaced by a rank-1 update. Here we discuss two possible such updates. In both cases, the analysis relies on the following lemma, which is an immediate consequence of [22, Lemmas 4 and 5], when applied within the setting of submodular function minimization.

Lemma 6.2 (Hoberg and Rothvoß [22]). *Let* $\mu \geq \max\{0, -f(V)\}$ *and* $Q \in \mathbb{S}^n_{++}$ *such that, for* $R = Q^{-1}$, $F_\mu \subseteq E(R)$. *Let* $y = \sum_{i=1}^k x_i g_i / \|g_i\|_Q$ *where* $g_1, \ldots, g_k \in B(f_\mu)$, $x \in \mathbb{R}^k_+$, $\vec{e}^\mathsf{T} x = 1$, *and assume that there is a set* $I \subset [k]$ *such that the vector* $z := \sum_{i \in I} x_i \frac{g_i}{\|g_i\|_Q}$ *satisfies*

$$\frac{\|y\|_Q}{\|z\|_Q} \leq \frac{1}{3\sqrt{n}}. \tag{27}$$

Define

$$R' := \frac{1}{(1 + 27n)} \left(R + 3 \frac{z z^\mathsf{T}}{\|z\|_Q^2} \right). \tag{28}$$

Then $F_\mu \subseteq E(R')$ *and* $\det(R') \geq (9/4) \det(R)$.

Next we show two ways of using SLIDING VON NEUMANN to produce vectors y and z as in the theorem above. In particular, the rescaling defined by (28) can be used within Algorithm 2 in place of the multi-rank rescaling (10). Note that the analysis of the algorithm is then identical to that of Algorithm 2 (simply replace Lemmas 3.10 and 3.12 by Lemma 6.2 in the analysis), hence we can determine a minimizer for f within $O(n \log(n L_{f,2}))$ rescalings of the form (28).

Betke's rescaling Betke [3] proposed the following rescaling. Within Algorithm 2, fix $\varepsilon = 1/(3(n+1)\sqrt{n})$. Assume that the vector $y = \sum_{i=1}^k x_i g_i / \|g_i\|_Q$ returned by each call of SLIDING VON NEUMANN$(f, \mu, W, Q, \varepsilon)$ is expressed as a convex combination of at most $n + 1$ terms (this is not instantly guaranteed by the sliding von-Neumann algorithm, but it can be done a-posteriori by Carathéodory's theorem). Since $\sum_{i=1}^k x_i = 1$, $x_i \geq 0$ for all $i \in [k]$, it follows that there exists $h \in [k]$ such that $x_h \geq 1/k \geq 1/(n+1)$. The set $I := \{h\}$ satisfies condition (27) in Lemma 6.2; indeed, in this case $z = x_h \frac{g_h}{\|g_h\|_Q}$, and we have

$$\frac{\|y\|_Q}{\|z\|_Q} = \frac{\|y\|_Q}{x_h} \leq \frac{1/(3(n+1)\sqrt{n})}{1/(n+1)} = \frac{1}{3\sqrt{n}}.$$

Observe that, with Betke's rescaling, the sliding-von Neumann algorithm requires $\lceil \varepsilon^{-2} \rceil = O(n^3)$ iterations to produce such a vector. Recall that each iteration of von-Neumann algorithm requires time $O(n \cdot E + n^2)$, for a total running time of $O(n^4 \cdot EO + n^5)$ for each

call to sliding-von Neumann. The vector y returned is expressed as a convex combination of $O(n^3)$ vectors of the form $g/\|g\|_Q$, $g \in B(f_\mu)$, hence we need $O(n^5)$ operation to express the vector y as a convex combination of at most $n+1$ such terms. Finally, the total number of rescalings is $O(n \log(nL_{f,2}))$. This means that the variant of Algorithm 2 using Betke's rescaling instead of (10) requires time $O((n^5 \cdot EO + n^6) \log(nL_{f,2}))$, as opposed to the $O((n^4 \cdot EO + n^5) \log(nL_{f,2}))$ ensured by the multi-rank update (10).

Hoberg and Rothvoß rescaling Hoberg and Rothvoß [22] provide a randomized selection rule for the vector z in the rank-1 rescaling (28). The rule is based on the following.

Lemma 6.3 (Hoberg and Rothvoß [22, Section 2.1]). *Let $v_1, \ldots, v_k \in \mathbb{R}^n$ be such that $\sum_{i=1}^k \|v_i\|_2 = 1$. Let $u \in \mathbb{R}^n$, $\|u\|_2 = 1$, be chosen uniformly at random, and let $I = \{i \in [k] : u^{\mathsf{T}} v_i \geq 0\}$. Then with constant probability*

$$\left\| \sum_{i \in I} v_i \right\|_2 \geq \frac{1}{4\sqrt{\pi n}}. \tag{29}$$

Within Algorithm 2, fix $\varepsilon \leq 1/(12\sqrt{\pi}n)$. Given the vector $y = \sum_{i=1}^k x_i g_i/\|g_i\|_Q$ returned by a call of SLIDING VON NEUMANN$(f, \mu, W, Q, \varepsilon)$, we let $v_i = x_i Q^{1/2} g_i/\|g_i\|_Q$ for $i = 1, \ldots, k$; by construction, $\sum_{i=1}^k \|v_i\|_2 = 1$. We can randomly sample the set I as in Lemma 6.3 that satisfies the requirements of Lemma 6.2; we repeat to sample I until condition (29) is satisfied. At this point, the vector $z = \sum_{i \in I} x_i \frac{g_i}{\|g_i\|_Q}$ satisfies condition (27), because

$$\frac{\|y\|_Q}{\|z\|_Q} \leq 4\sqrt{\pi n} \cdot \varepsilon \leq \frac{1}{3\sqrt{n}}.$$

Observe that each call to SLIDING VON NEUMANN$(f, \mu, W, Q, \varepsilon)$ requires $\lceil \varepsilon^{-2} \rceil = O(n^2)$ iterations, for a total time of $O(n^3 \cdot EO + n^4)$ per call, just as for Algorithm 2. Checking if the random set I satisfies condition (29) takes time $O(mk)$, where k is the number of vectors in the convex combination defining y, which is bounded by the number $O(n^2)$ of iterations of von Neumann. The expected number of times we need to generate a random set I before condition (29) is verified is constant, hence in expected $O(n^3)$ arithmetic operation we can compute the rescaling direction z after each call SLIDING VON NEUMANN$(f, \mu, W, Q, \varepsilon)$. Hence the overall expected running time of the algorithm is $O((n^4 \cdot EO + n^5) \log(nL_{f,2}))$, just as for Algorithm 2.

6.3 Rank-1 pullback

The pullback framework described in Section 3.4 can also be adapted to the use of rank-1 updates of the form (28). All arguments in Section 3.4 proceed without any modification, provided that we replace Lemma 3.16 (used for the multi-rank rescaling (10)) with the following analogous statement for the rank-1 rescaling (28).

Lemma 6.4. *Let $A \in \mathbb{R}^{n \times p}$, $R \in \mathbb{S}_{++}^n$, and $Q = R^{-1}$. Let $x \in \mathbb{R}_+^p$ and $I \subseteq [p]$ such that the vectors $y := \sum_{i=1}^p x_i \frac{a_i}{\|a_i\|_Q}$ and $z := \sum_{i \in I} x_i \frac{a_i}{\|a_i\|_Q}$ satisfy*

$$\frac{\|y\|_Q}{\|z\|_Q} \leq \gamma.$$

Define

$$R' := \frac{1}{(1 + 3\gamma^2)} \left(R + 3 \frac{zz^{\mathsf{T}}}{\|z\|_Q^2} \right),$$

and $Q' = (R')^{-1}$. For every $v \in \mathbb{R}^n$, there exists $\mu \in \mathbb{R}_+^p$ such that $\|v + A\mu\|_Q \leq \|v\|_{Q'}$.

Proof. First note that

$$Q' := (1 + 3\gamma^2) \left(Q - \frac{3}{4} \frac{Qzz^{\mathsf{T}}Q^{\mathsf{T}}}{\|z\|_Q^2} \right).$$

Let $v \in \mathbb{R}^n$, and define

$$u = \frac{RQ'v}{1 + 3\gamma^2},$$

Note that

$$v = R'QRQ'v = u + 3\frac{\langle z, u \rangle_Q}{\|z\|_Q^2} z.$$

Thus, if we define

$$\mu_i = \begin{cases} 3\frac{x_i}{\|a_i\|_Q} \frac{|\langle z, u \rangle_Q|}{\|z\|_Q^2} & i \notin I \\ \max\left\{0, -6\frac{x_i}{\|a_i\|_Q} \frac{\langle z, u \rangle_Q}{\|z\|_Q^2}\right\} & i \in I \end{cases}$$

we obtain

$$v + A\mu = u + 3\frac{|\langle z, u \rangle_Q|}{\|z\|_Q^2} y.$$

Note that $\mu \geq 0$ and

$$\begin{aligned} \|v\|_{Q'}^2 &= v^T Q' R(QR'Q)RQ'v = (1 + 3\gamma^2)u^\mathsf{T}\left(Q + 3\frac{Qzz^\mathsf{T}Q}{\|z\|_Q^2}\right)u \\ &= \|u\|_Q^2(1 + 3\gamma^2)(1 + 3\alpha^2), \end{aligned}$$

where we define $\alpha := \frac{\langle z, u \rangle_Q}{\|z\|_Q \|u\|_Q}$. On the other hand

$$\|v + A\mu\|_Q \leq \|u\|_Q + 3\frac{|\langle z, u \rangle_Q|}{\|z\|_Q^2}\|y\|_Q = \|u\|_Q(1 + 3|\alpha|\frac{\|y\|_Q}{\|z\|_Q}) \leq \|u\|_Q(1 + 3|\alpha|\gamma).$$

The statement now follows by noting that $(1 + 3\gamma^2)(1 + 3\alpha^2) - (1 + 3|\alpha|\gamma)^2 = 3(|\alpha| - \gamma)^2 \geq 0$. $\qquad\square$

Acknowledgments.

A   Proofs of technical lemmas

A.1   Proofs of Section 2

Proof of Lemma 2.4. Assume w.l.o.g that $z(v_1) \leq z(v_2) \leq \ldots \leq z_(v_n)$, and denote $S_i = \{v_1, \ldots, v_i\}$. The hypothesis of the statement can be written as $-\min_{x \in B(f)} z^\top(x - z) \leq \varepsilon$. If we let $f' : 2^V \mapsto \mathbb{R}$ be the submodular function defined by $f'(S) = f(S) - z(S)$, observing that $x \in B(f)$ if and only if $x - z \in B(f')$, the previous expression implies $-\min_{y \in B(f')} z^\top y \leq \varepsilon$. By (3), and observing that $f'(V) = 0$, this can be written as

$$\begin{aligned} \varepsilon &\geq \sum_{i=1}^{n-1}(f(S_i) - z(S_i))(z(v_{i+1}) - z(v_i)) \\ &= \int_{-\infty}^{+\infty} f(\{v : z(v) \leq t\}) - z(\{v : z(v) \leq t\})\,dt \\ &\geq \int_{-\eta}^{+\eta} f(\{v : z(v) \leq t\}) - z(\{v : z(v) \leq t\})\,dt, \end{aligned}$$

where the last inequality holds trivially for every $\eta > 0$ since $f(S) - z(S) \geq 0$ for all $S \subseteq V$ because $z \in B(f)$, and where the equality holds because

$$\{v \,:\, z(v) \le t\} = \begin{cases} \emptyset & \text{if } t < z(v_1) \\ S_i & \text{if } z(v_i) \le t < z(v_{i+1}) \\ V & \text{if } t \ge z(v_n). \end{cases}$$

It follows from the above that, for every $\eta > 0$, there exists $t \in [-\eta, \eta]$ such that $f(\{v \,:\, z(v) \le t\}) - z(\{v \,:\, z(v) \le t\}) \le \varepsilon/2\eta$. Setting $S := \{v \,:\, z(v) \le t\}$, this implies

$$f(S) - z^-(V) \le z^+(S) - z^-(V \setminus S) + \varepsilon/2\eta \le n|t| + \varepsilon/2\eta \le n\eta + \varepsilon/2\eta.$$

Choosing $\eta = \sqrt{\varepsilon/2n}$ gives $f(S) - z^-(V) \le \sqrt{2n\varepsilon}$, as required. $\qquad\square$

A.2 Proofs of Section 3.3

Lemma A.1. *The property that $F_\mu \subseteq E(R)$ is preserved whenever Algorithm* RESCALING-SFM *performs a rescaling.*

Proof. Assume that at a given step of the algorithm $F_\mu \subseteq E(R)$ holds, and we rescale R to R', where

$$R' = \frac{1}{(1+\varepsilon)^2} \left(R + \sum_{i=1}^k \frac{x_i}{\|g_i\|_Q^2} g_i g_i^\mathsf{T} \right)$$

for the vector x returned by the sliding von Neumann algorithm.

We show $F_\mu \subseteq E(R')$. Consider an arbitrary point $z \in F_\mu$; then $g_i^\mathsf{T} z \ge 0$ for all $i \in [k]$ and, by the induction hypothesis, $\|z\|_R^2 \le 1$ because $z \in E(R)$.

Recall that, in the algorithm, the vector $y = \sum_{i=1}^k x_i \frac{g_i}{\|g_i\|_Q}$ satisfies $\|y\|_Q \le \varepsilon$. By the Cauchy-Schwartz inequality, we have $y^\mathsf{T} z = y^\mathsf{T} Q^{1/2} Q^{-1/2} z \le \|y\|_Q \|z\|_R \le \varepsilon$, and similarly $g_i^\mathsf{T} z \le \|g_i\|_Q \|z\|_R \le \|g_i\|_Q$ for every $i \in [k]$. We then have

$$
\begin{aligned}
\|z\|_{R'}^2 &= \frac{1}{(1+\varepsilon)^2} z^\mathsf{T} \left(R + \sum_{i=1}^k \frac{x_i}{\|g_i\|_Q^2} g_i g_i^\mathsf{T} \right) z = \frac{1}{(1+\varepsilon)^2} \left(\|z\|_R^2 + \sum_{i=1}^k x_i \left(\frac{g_i^\mathsf{T} z}{\|g_i\|_Q} \right)^2 \right) \\
&\le \frac{1}{(1+\varepsilon)^2} \left(1 + \sum_{i=1}^k x_i \frac{g_i^\mathsf{T} z}{\|g_i\|_Q} \right) = \frac{1 + y^\mathsf{T} z}{(1+\varepsilon)^2} \le 1,
\end{aligned}
$$

where the first inequality follows from the facts that $\|z\|_R \le 1$, $x \ge 0$, and $0 \le g_i^\mathsf{T} z \le \|g_i\|_Q$ for all $i \in [k]$, while the second follows from $y^\mathsf{T} z \le \varepsilon$. Consequently, $z \in E(R')$, completing the proof. $\qquad\square$

Proof of Lemma 3.12. Let R and R' denote the matrix before and after the rescaling. Let $X = \sum_{i=1}^k x_i g_i g_i^\mathsf{T} / \|g_i\|_Q^2$; hence $R' = (R + X)/(1+\varepsilon)^2$. The ratio of the two determinants is

$$\frac{\det(R')}{\det(R)} = \frac{\det(R + X)}{(1+\varepsilon)^{2n} \det(R)} = \frac{\det\left(I_n + R^{-1/2} X R^{-1/2} \right)}{(1+\varepsilon)^{2n}}$$

Now $R^{-1/2} = Q^{1/2}$, and $Q^{1/2} X Q^{1/2}$ is a positive semidefinite matrix. From the above, the inequality $\det(I_n + M) \ge 1 + \operatorname{tr}(M)$ for every $M \in \mathbb{S}_{++}^n$, and the linearity of the trace, we get

$$\frac{\det(R')}{\det(R)} \ge \frac{1 + \operatorname{tr}(Q^{1/2} X Q^{1/2})}{(1+\varepsilon)^{2n}} = \frac{1}{(1+\varepsilon)^{2n}} \left(1 + \sum_{i=1}^k \frac{x_i}{\|g_i\|_Q^2} \operatorname{tr}(Q^{1/2} g_i g_i^\mathsf{T} Q^{1/2}) \right).$$

Finally, $\operatorname{tr}(Q^{1/2} g_i g_i^\mathsf{T} Q^{1/2}) = \operatorname{tr}(g_i^\mathsf{T} Q g_i) = \|g_i\|_Q^2$. Therefore we conclude

$$\frac{\det(R')}{\det(R)} \ge \frac{1 + \sum_{i=1}^k x_i}{(1+\varepsilon)^{2n}} = \frac{2}{(1+\varepsilon)^{2n}} \ge \frac{16}{9},$$

where the last inequality follows from $\varepsilon = \frac{1}{20n}$. $\qquad\square$

Proof of Lemma 3.13. Let $\bar{g} = \arg\min_{g \in \Gamma} \|g\|_Q / \|g\|_2$. Then, by (13),

$$\frac{\|\bar{g}\|_Q^2}{\|\bar{g}\|_2^2} \sum_{g \in \Gamma} \gamma_g \leq \sum_{g \in \Gamma} \gamma_g \frac{\|g\|_Q^2}{\|g\|_2^2} = \sum_{g \in \Gamma} \gamma_g \frac{g^\mathsf{T} Q g}{\|g\|_2^2} = \operatorname{tr}\left(Q \sum_{g \in \Gamma} \gamma_g \frac{g g^\mathsf{T}}{\|g\|_2^2} \right) \tag{30}$$

$$= \operatorname{tr}(Q(R - \alpha I_n)) = \operatorname{tr}(I_n - \alpha Q) = n - \alpha \operatorname{tr}(Q) < n .$$

The final inequality holds since Q is positive definite.

Note from (13) that $\operatorname{tr}(R) = \alpha n + \sum_{g \in \Gamma} \gamma_g \leq n + \sum_{g \in \Gamma} \gamma_g$. The latter and the inequality $\det(M)^{1/n} \leq \operatorname{tr}(M)/n$ for every $M \in \mathbb{S}_{++}^n$ imply that $\sum_{g \in \Gamma} \gamma_g \geq \operatorname{tr}(R) - n \geq n(\det(R)^{1/n} - 1)$. The statement now follows from (30). $\qquad\square$

References

[1] F. Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends in Machine Learning*, 6(2-3):145–373, 2013.

[2] A. Belloni, R. M. Freund, and S. Vempala. An efficient rescaled perceptron algorithm for conic systems. *Mathematics of Operations Research*, 34(3):621–641, 2009.

[3] U. Betke. Relaxation, new combinatorial and polynomial algorithms for the linear feasibility problem. *Discrete & Computational Geometry*, 32(3):317–338, 2004.

[4] D. Chakrabarty, P. Jain, and P. Kothari. Provable submodular minimization using Wolfe's algorithm. In *Advances in Neural Information Processing Systems (NIPS)*, pages 802–809, 2014.

[5] D. Chakrabarty, Y. T. Lee, A. Sidford, and S. C.-w. Wong. Subquadratic submodular function minimization. In *ACM Symposium on Theory of Computing (STOC)*, pages 1220–1231, 2017.

[6] S. Chubanov. A polynomial algorithm for linear feasibility problems given by separation oracles. http://www.optimization-online.org/DB_HTML/2017/01/5838.html, 2017.

[7] D. Dadush, L. A. Végh, and G. Zambelli. Rescaling algorithms for linear conic feasibility. *Mathematics of Operations Research*, 2019. (to appear).

[8] G. B. Dantzig. Converting a converging algorithm into a polynomially bounded algorithm. Technical report, Stanford University, 1991.

[9] G. B. Dantzig. An ε-precise feasible solution to a linear program with a convexity constraint in $1/\varepsilon^2$ iterations independent of problem size. Technical report, Technical Report 92-5, Stanford University, 1992.

[10] J. Dunagan and S. Vempala. A simple polynomial-time rescaling algorithm for solving linear programs. *Mathematical Programming*, 114(1):101–114, 2008.

[11] J. Edmonds. Submodular functions, matroids, and certain polyhedra. *Edited by G. Goos, J. Hartmanis, and J. van Leeuwen*, 11, 1970.

[12] A. R. Ene and H. L. Nguyen. Random coordinate descent methods for minimizing decomposable submodular functions. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 2015.

[13] A. R. Ene, H. L. Nguyen, and L. A. Végh. Decomposable submodular function minimization: Discrete and continuous. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.

[14] S. Fujishige. Lexicographically optimal base of a polymatroid with respect to a weight vector. *Mathematics of Operations Research*, 5(2):186–196, 1980.

[15] S. Fujishige. *Submodular functions and optimization*, volume 58. Elsevier, 2005.

[16] S. Fujishige. A note on submodular function minimization by Chubanov's LP algorithm. http://www.optimization-online.org/DB_HTML/2017/09/6217.html, 2017.

[17] S. Fujishige and S. Isotani. A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, 7(1):3–17, 2011.

[18] S. Fujishige and S. Iwata. A descent method for submodular function minimization. *Mathematical Programming*, 92(2):387–390, Apr 2002.

[19] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, Jun 1981.

[20] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 1988.

[21] E. Hazan and S. Kale. Online submodular minimization. *Journal of Machine Learning Research*, 13(Oct):2903–2922, 2012.

[22] R. Hoberg and T. Rothvoß. An improved deterministic rescaling for linear programming algorithms. In *International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 267–278, 2017.

[23] S. Iwata. A faster scaling algorithm for minimizing submodular functions. *SIAM Journal on Computing*, 32(4):833–840, 2003.

[24] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM (JACM)*, 48(4):761–777, 2001.

[25] S. Iwata and J. B. Orlin. A simple combinatorial algorithm for submodular function minimization. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1230–1237, 2009.

[26] S. Jegelka and J. A. Bilmes. Online submodular minimization for combinatorial structures. In *Proceedings of the 28th International Conference on Machine Learning (ICML11)*, pages 345–352, 2011.

[27] S. Jegelka, H. Lin, and J. A. Bilmes. On fast approximate submodular minimization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 460–468, 2011.

[28] S. Lacoste-Julien and M. Jaggi. On the global linear convergence of frank-wolfe optimization variants. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, pages 496–504, Cambridge, MA, USA, 2015. MIT Press.

[29] Y. T. Lee, A. Sidford, and S. C. Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. *CoRR*, abs/1508.04874, 2015.

[30] J. B. Orlin. A faster strongly polynomial time algorithm for submodular function minimization. *Mathematical Programming*, 118(2):237–251, 2009.

[31] J. Peña and N. Soheili. A deterministic rescaled perceptron algorithm. *Mathematical Programming*, 155(1-2):497–510, 2016.

[32] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355, 2000.

[33] A. Schrijver. *Combinatorial optimization - Polyhedra and Efficiency*. Springer, 2003.

[34] M. A. Sohoni. Membership in submodular and other polyhedra. Technical Report TR-102-92, Indian Institute of Technology, Bombay, India, 1992.

[35] P. Stobbe and A. Krause. Efficient minimization of decomposable submodular functions. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.

[36] É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3):247–255, 1985.

[37] P. Wolfe. Finding the nearest point in a polytope. *Mathematical Programming*, 11(1):128–149, 1976.