A tight combinatorial algorithm for submodular maximization subject to a matroid constraint

Moran Feldman, Yuval Filmus, Roy Schwartz, Justin Ward

January 25, 2017

Abstract

We give a simplified exposition of the algorithm of Filmus and Ward (2014) for maximizing a submodular function subject to a matroid constraint.

1 Introduction

Monotone submodular functions abound in combinatorial optimization. The greedy algorithm gives the optimal approximation ratio, $1-1/e$, for optimizing a monotone submodular function over a cardinality constraint. However, over a general matroid constraint, and even over a partition matroid constraint, it only gives a $1/2$ approximation. The continuous greedy algorithm, due to Calienscu et al. [CCPV11] (see also [FNS11]), gives an optimal $1-1/e$ approximation, but it is based on continuous methods. Filmus and Ward [FW14] gave a purely combinatorial algorithm, based on the paradigm of *non-oblivious local search*.

The conference version of Filmus and Ward [FW12] was simplified considerably in the journal version [FW14]. Here we present a further simplification, due to the first author.

2 Preliminaries

We assume familiarity with the basic definitions, but repeat them here to fix notation.

Basic notation For a set A and an element x, $A+x = A \cup \{x\}$ and $A-x = A \setminus \{x\}$.

A *set function* is a function $f\colon 2^U \to \mathbb{R}$, where U is some finite universe.

For a set $S \subseteq U$, we denote its indicator function (from U to $\{0,1\}$) by 1_S.

Submodular functions A *monotone submodular function* is a set function $f\colon 2^U \to \mathbb{R}$ satisfying the following axioms:

(a) Normalization: $f(\emptyset) = 0$.

(b) Monotonicity: if $A \subseteq B \subseteq U$ then $f(A) \le f(B)$.

(c) Submodularity: for all $A, B \subseteq U$ we have $f(A) + f(B) \ge f(A \cup B) + f(A \cap B)$.

Multilinear extension Given a function $f\colon 2^U \to \mathbb{R}$, we define its multilinear extension $F\colon [0,1]^U \to \mathbb{R}$ as follows: $F(x) = \mathbb{E}[f(S)]$, where S is a random set chosen so that $i \in S$ with probability x_i, independently. We can expand this definition to obtain a multilinear polynomial in the inputs. Note that $F(1_S) = f(S)$, and in this sense F extends f.

Marginals If f is a set function then we define

$$f(A|B) = f(A \cup B) - f(B).$$

Similarly, for its multilinear extension F we define

$$F(x|y) = F(x \vee y) - F(y),$$

where $\vee$ denotes elementwise maximum.

Matroids A *matroid* $\mathcal{M}$ is a non-empty collection of subsets of a finite universe U satisfying the following axioms:

(a) Downward closure: if $\mathcal{M}$ contains a set $A \subseteq U$ then it contains all its subsets.

(b) Exchange: if $A, B \in \mathcal{M}$ and $|A| < |B|$ then there exists $x \in B \setminus A$ such that $A + x \in \mathcal{M}$.

We call the sets in $\mathcal{M}$ *independent sets*, and the inclusion-maximal independent sets we call *bases*. It turns out that all bases in a matroid have the same size, known as the *rank* of the matroid.

Brualdi's lemma [Bru69] shows that if A, B are any two bases then there exists a bijection π from A to B, fixing $A \cap B$, such that $A - x + \pi(x) \in \mathcal{M}$ for all $x \in A$.

3 Algorithm

For the rest of this note, let f be a monotone submodular function, F its multilinear extension, and $\mathcal{M}$ a matroid. Our goal is to find $S \in \mathcal{M}$ which maximizes $f(S)$. We define below a related monotone submodular function g. The non-oblivious local search algorithm uses non-oblivious local search to find a non-oblivious *local optimum*:

Definition 3.1. A set $S \in \mathcal{M}$ is a *local optimum* (with respect to g) if for any $x \in S$ and $y \notin S$ such that $S - x + y \in \mathcal{M}$, we have $g(S) \geq g(S - x + y)$, or equivalently $g(x|S - x) \geq g(y|S - x)$.

In reality a local optimum cannot be found efficiently, and instead the algorithm finds an *approximate local optimum*, in which we are only guaranteed that $g(S) \geq (1-\epsilon)g(S - x + y)$. We refer the reader to the original paper [FW14] for the details; the analysis here is only slightly affected, and we leave the necessary changes to the reader.

The function g is defined as follows:

$$g(A) = \int_0^1 \frac{e^{p-1}}{p} F(p \cdot 1_A)\, dp. \tag{$*$}$$

The reader might be worried that this isn't well-defined, since $\frac{e^{p-1}}{p}$ blows up as $p \to 0$. Fortunately, this is not the case, essentially since f is normalized:

Lemma 3.1. *The function g is well-defined.*

Proof. Recall that $F(p \cdot 1_A) = \mathbb{E}[f(S)]$, where S is a random subset of A in which each element is found with probability p. In particular, $\Pr[S \neq \emptyset] \leq |A|p$. Since $f(\emptyset) = 0$, we deduce that $F(p \cdot 1_A) \leq p|A|f(A)$. Therefore

$$\int_0^1 \frac{e^{p-1}}{p} F(p \cdot 1_A)\, dp \leq \int_0^1 e^{p-1}|A|f(A)\, dp \leq (1 - 1/e)|A|f(A).$$

It follows that $g(A)$ is well-defined. □

The marginals of g have a particularly simple formula:

Lemma 3.2. *Suppose that $x \notin A$. Then*

$$g(x|A) = \int_0^1 e^{p-1} F(1_x | p \cdot 1_A)\, dp.$$

Proof. Since F is multilinear, we have $F(p \cdot 1_{A+x}) - F(p \cdot 1_A) = p(F(p \cdot 1_A + 1_x) - F(p \cdot 1_A))$. Therefore

$$\begin{aligned} g(x|A) = g(A+x) - g(A) &= \int_0^1 \frac{e^{p-1}}{p}(F(p \cdot 1_{A+x}) - F(p \cdot 1_A))\, dp \\ &= \int_0^1 e^{p-1}(F(p \cdot 1_A + 1_x) - F(p \cdot 1_A))\, dp \\ &= \int_0^1 e^{p-1} F(1_x | p \cdot 1_A)\, dp. \end{aligned}$$ □

Incidentally, this formula gives another proof of Lemma 3.1.

As promised, the function g is also monotone submodular:

Lemma 3.3. *The function g is monotone submodular.*

Proof. The formula directly implies that $g(\emptyset) = 0$, and Lemma 3.2 implies that g is monotone. To show submodularity, it suffices to show that $g(x|A) \geq g(x|B)$ whenever $A \subseteq B$ and $x \notin B$. This follows from Lemma 3.2 together with the inequality $F(1_x|p \cdot 1_A) \geq F(1_x|p \cdot 1_B)$, which follows from the submodularity of f. □

4 Analysis

Let S be a local optimum, and let O be a global optimum, that is, an optimal solution for the optimization problem (in fact the analysis works for any set $O \in \mathcal{M}$). Our goal is to bound $f(S)/f(O)$ from below. Brualdi's lemma (mentioned in the preliminaries) shows that there is a mapping $\pi \colon S \to O$ which fixes $S \cap O$ and satisfies $S - x + \pi(x)$ for all $x \in S$. The local optimality constraints imply, in particular, that for all $x \in S$:

$$g(x|S - x) \geq g(\pi(x)|S - x). \tag{†}$$

The proof now proceeds by giving a lower bound and an upper bound on $\sum_{x \in S} g(x|S - x)$.

Lemma 4.1 (Lower bound)**.** *We have*

$$\sum_{x\in S} g(x|S-x) \geq \left(1-\frac{1}{e}\right) f(O) - \int_0^1 e^{p-1} F(p\cdot 1_S)\,dp.$$

Proof. Let $x \in S$. When $x \notin O$, submodularity of g and Lemma 3.2 imply that

$$g(x|S-x) \geq g(\pi(x)|S-x) \geq g(\pi(x)|S) = \int_0^1 e^{p-1} F(1_{\pi(x)}|p\cdot 1_S)\,dp.$$

We can reach the same conclusion when $x \in O$, with a bit more work: Lemma 3.2, the multilinearity of F, and monotonicity of f imply that

$$\begin{aligned} g(x|S-x) &= \int_0^1 e^{p-1} F(1_x|p\cdot 1_{S-x})\,dp \\ &= \int_0^1 e^{p-1} \frac{F(1_x|p\cdot 1_S)}{1-p}\,dp \geq \int_0^1 e^{p-1} F(1_x|p\cdot 1_S)\,dp. \end{aligned}$$

This is indeed the same inequality as in the preceding case, since $x = \pi(x)$.

Summing the inequality for all $x \in S$, we obtain

$$\sum_{x\in S} g(x|S-x) \geq \int_0^1 e^{p-1} \sum_{x\in S} F(1_{\pi(x)}|p\cdot 1_S)\,dp \geq \int_0^1 e^{p-1} F(1_O|p\cdot 1_S)\,dp,$$

using submodularity of f. Monotonicity of f implies that $F(1_O|p\cdot 1_S) \geq f(O) - F(p\cdot 1_S)$, hence

$$\sum_{x\in S} g(x|S-x) \geq \int_0^1 e^{p-1}(f(O) - F(p\cdot 1_S))\,dp.$$

This implies the stated formula, since $\int_0^1 e^{p-1}\,dp = e^{p-1}|_0^1 = 1-1/e$. □

Lemma 4.2 (Upper bound)**.** *We have*

$$\sum_{x\in S} g(x|S-x) \leq f(S) - \int_0^1 e^{p-1} F(p\cdot 1_S)\,dp.$$

Proof. Let $x \in S$. Lemma 3.2 and multilinearity of g imply that

$$g(x|S-x) = \int_0^1 e^{p-1} F(1_x|p\cdot 1_{S-x})\,dp = \int_0^1 e^{p-1} \partial_x F(p\cdot 1_S)\,dp,$$

where ∂_x denotes partial derivative with respect to the coordinate corresponding to x. Summing this over all $x \in S$, we obtain

$$\sum_{x\in S} g(x|S-x) = \int_0^1 e^{p-1} \langle \nabla F(p\cdot 1_S), 1_S\rangle\,dp,$$

where ∇F denotes the gradient of F.

We now use integration by parts, in the form

$$\int_0^1 a(p)b'(p)\,dp = a(1)b(1) - a(0)b(0) - \int_0^1 a'(p)b(p)\,dp.$$

In our application, $a(p) = e^{p-1}$ and $b'(p) = \langle \nabla F(p \cdot 1_S), 1_S \rangle$, so that $a'(p) = e^{p-1}$ and

$$b(p) = \int_0^p b'(q)\,dq = \int_0^p \langle \nabla F(q \cdot 1_S), 1_S \rangle\,dq = F(p \cdot 1_S),$$

using the normalization of F.

Since $a(0)b(0) = 0$ and $a(1)b(1) = f(S)$, integration by parts yields

$$\sum_{x \in S} g(x|S - x) = f(S) - \int_0^1 e^{p-1} F(p \cdot 1_S)\,dp. \qquad \square$$

Combining both bounds, we obtain our main theorem.

Theorem 4.3. *We have*

$$f(S) \geq \left(1 - \frac{1}{e}\right) f(O).$$

References

[Bru69] Richard A. Brualdi. Comments on bases in dependence structures. *Bulletin of the Australian Mathematical Society*, 1(02):161–167, 1969.

[CCPV11] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

[FNS11] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2011)*, pages 570–579, 2011.

[FW12] Yuval Filmus and Justin Ward. A tight combinatorial algorithm for submodular maximization subject to a matroid constraint. In *53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, pages 659–668, 2012.

[FW14] Yuval Filmus and Justin Ward. A tight combinatorial algorithm for submodular maximization subject to a matroid constraint. *SIAM Journal on Computing*, 43:514–542, 2014.