

Fast Non-Monotone Submodular Maximisation Subject to a Matroid Constraint

Pau Segui-Gasco ^{*†} *Hyo-Sang Shin* [†]

Abstract

In this work we present the first practical $(\frac{1}{\epsilon} - \epsilon)$ -approximation algorithm to maximise a general non-negative submodular function subject to a matroid constraint. Our algorithm is based on combining the decreasing-threshold procedure of Badanidiyuru and Vondrak (SODA 2014) with a smoother version of the measured continuous greedy algorithm of Feldman et al. (FOCS 2011). This enables us to obtain an algorithm that requires $O(\frac{nr^2}{\epsilon^4} (\frac{\bar{d}+d}{d})^2 \log^2(\frac{n}{\epsilon}))$ value oracle calls, where n is the cardinality of the ground set, r is the matroid rank, and $d, \bar{d} \in \mathbb{R}^+$ are the absolute values of the minimum and maximum marginal values that the function f can take i.e.: $-d \leq f_S(i) \leq \bar{d}$, for all $i \in E$ and $S \subseteq E$, where E is the ground set. The additional value oracle calls with respect to the work of Badanidiyuru and Vondrak come from the greater spread in the sampling of the multilinear extension that the possibility of negative marginal values introduce.

1 Introduction

Submodular maximisation problems have drawn a lot of attention recently [1]. This interest is due to a good confluence of theoretical results and practical applicability. Intuitively, a submodular set function is said to be so because it exhibits diminishing marginal returns, i.e.: the marginal value that an element adds to a set decreases as the size of the set increases. This simple property arises naturally in many applications and is what enables good theoretical tractability. It has been used in a variety of application domains, to name but a few: markets [2, 3], influence in networks [4], document summarisation [5], and sensor placement [6, 7]. In particular, this paper focuses on matroid-constrained submodular maximization of a general non-negative submodular function.

Matroids are an incredibly powerful abstraction of independence. They capture seemingly disconnected notions such as linear independence, forests in graphs, traversals, among many others. Interestingly, with linear sum objectives (modular function), they are inextricably linked with the greedy algorithm. If the greedy algorithm is optimal, then there is an implicit matroid; if there is a matroid then the greedy algorithm is optimal [8]. They provide a flexible framework to characterise a variety of relevant constraints such as: partitions, schedules, cardinality, or even rigidity.

*p.seguigasco@cranfield.ac.uk

†Centre for Autonomous and Cyber-Physical Systems, Institute of Aerospace Sciences, School of Aerospace Transport and Manufacturing, Cranfield University, Bedfordshire, UK.

Importantly, the combination of a submodular function and a matroid constraint plays a unifying role for many well-known combinatorial optimisation problems such as: Submodular Welfare (also known as Submodular Task Allocation), Max k -Cover, Max Generalised Assignment, Max Facility Location, and Constrained Max Cut (e.g. Max Bisection) among others.

However, maximising a submodular function subject to a matroid constraint is NP-Hard. Hence much of the research effort has focused on obtaining good approximation algorithms. A classic result by Nemhauser, Wolsey, and Fisher [9] shows that the Greedy Algorithm is a $\frac{1}{2}$ -approximation for non-negative monotone submodular functions. More recently, a fruitful avenue of research was spurred by Vondrak [10, 11] by showing that solving a relaxation of the problem based on the multilinear extension using the Continuous Greedy Algorithm and then rounding the result yields good approximation algorithms. They present their result in the context of non-negative monotone submodular functions to yield a $(1 - \frac{1}{e})$ -approximation. Shortly after, Feldman et al. [12] modified Vondrak’s algorithm to develop the Measured Continuous Greedy Algorithm which supported both non-negative non-monotone and monotone submodular functions. Their algorithm was the first to achieve a $\frac{1}{e}$ -approximation for general non-negative submodular functions subject to a matroid constraint. Both continuous greedy algorithms can find the aforementioned approximation bounds for solving relaxation problems subject to the more general constraint class of down-closed solvable polytopes. An important breakthrough in this field are Contention Resolution Schemes, a rounding framework proposed in [13], because they provide a paradigm for developing rounding schemes for a combination of useful constraints including matroids and knapsacks. Thus the combination of continuous greedy relaxations and Contention Resolution schemes enabled approximation algorithms for many important submodular maximization problems.

However, continuous greedy algorithms are impractically slow. Vondrak’s algorithm required $\Theta(n^8)$ value oracle calls, while Feldman’s Measured Continuous Greedy requires $O(n^6)$ assuming oracle access to the multilinear extension, which needs to be sampled in general, creating additional overhead, possibly around $O(n^3)$ or $O(n^2)$.

To remedy this, Badanidiyuru and Vondrak [14] proposed an efficient $(1 - \frac{1}{e} - \epsilon)$ -approximation algorithm that uses $O(\frac{nr}{\epsilon^4} \log^2 \frac{n}{\epsilon})$ value oracle calls, for non-negative monotone submodular function maximisation subject to a matroid constraint. Where n is the cardinality of the ground set, and r is the rank of the matroid. They achieve this by using a Decreasing-Threshold procedure that enables a reduction of both the number of steps and the number of samples needed at each step of the continuous greedy process.

In the inapproximability front, for matroid constrained non-negative monotone submodular maximisation Feige [15] showed that improving over the $1 - \frac{1}{e}$ threshold is NP-Hard, and so the continuous greedy guarantees are tight. In the non-monotone case, [16] showed that no polynomial time algorithm can achieve an approximation better than 0.478. Closing the gap between $\frac{1}{e}$ and 0.478 remains an important open problem where recent advances have been made: Ene and Nguyen in [17] give a 0.372-approximation, while Feldman et al in [18] improve it to a 0.385-approximation. This improvements are relatively small, recall that $\frac{1}{e} \approx 0.368$, but show that there is room for future improvement. Both algorithms are based on the measured continuous greedy, and would therefore benefit from the techniques presented here.

Our contribution in this work is the first $(\frac{1}{e} - \epsilon)$ -approximation algorithm to maximise a general non-negative submodular function subject to a matroid constraint with a practical time complexity. Our algorithm is based on combining the Decreasing-Threshold procedure of [14] with a smoother version of the measured continuous greedy algorithm of [12]. This allows us to obtain an algorithm that requires $O(\frac{nr^2}{\epsilon^4} (\frac{\bar{d}+d}{d})^2 \log^2(\frac{n}{\epsilon}))$ value oracle calls, where $\underline{d}, \bar{d} \in \mathbb{R}^+$ are the absolute values of

the minimum and maximum marginal values that the function f can take, i.e. $-d \leq f_S(i) \leq \bar{d}$, for all $i \in E$ and $S \subseteq E$. The additional oracle calls with respect to [14] come from the additional spread in the sampling of the multilinear extension that the possibility of negative marginal values introduces.

Definitions Let us now introduce the definitions of a matroid and a submodular function. A function $f : 2^E \rightarrow \mathbb{R}^+$ on a set E is said to be submodular if:

$$f(A) + f(B) \geq f(A \cup B) + f(A \cap B), \text{ for all } A, B \subseteq E. \quad (1)$$

A more intuitive, but equivalent, definition can be formulated in terms of the marginal value added by an element: given $Y, X \subseteq E$ satisfying $X \subseteq Y$ and $x \in E \setminus Y$, then $f(X + x) - f(X) \geq f(Y + x) - f(Y)$. Herein we overload the symbol $+$ ($-$) to use it as shorthand notation for the addition (subtraction) of an element to a set, i.e. $S + i = S \cup \{i\}$ ($S - i = S \setminus \{i\}$). A matroid can be defined as follows [8]: a pair $\mathcal{M} = (E, \mathcal{I})$ is called a matroid if E is a finite set and \mathcal{I} is a nonempty collection of subsets of E satisfying: $\emptyset \in \mathcal{I}$; if $A \in \mathcal{I}$ and $B \subseteq A$, then $B \in \mathcal{I}$; and if $A, B \in \mathcal{I}$ and $|A| < |B|$, then $A + z \in \mathcal{I}$ for some $z \in B \setminus A$.

A matroid base $B \subseteq E$ is a maximally independent set $B \in \mathcal{I}$, that becomes dependent by adding an additional element $e \in E \setminus B$, i.e. $B + e \notin \mathcal{I}$. A key property of matroids is that we can exchange elements between bases, let \mathcal{B} denote the set of all bases of a matroid \mathcal{M} , then the exchange property can be formalised as follows:

Lemma 1.1. (Corollary 39.12A in [8]) *Let $\mathcal{M} = (N, \mathcal{I})$ be a matroid, and $B_1, B_2 \in \mathcal{B}$ be two bases. Then there is a bijection $\phi : B_1 \rightarrow B_2$ such that for every $b \in B_1$ we have $B_1 - b + \phi(b) \in \mathcal{B}$.*

Now we can formally define our target problem:

Given a ground set E , a matroid defined upon it $\mathcal{M} = (E, \mathcal{I})$, and a general non-negative submodular function $f : 2^E \rightarrow \mathbb{R}^+$, find:

$$\max_{S \in \mathcal{I}} f(S). \quad (2)$$

A common approach to solve combinatorial optimisation problems is by using relaxation and rounding. This involves solving a continuous problem, in which we allow fractional solutions, and a rounding procedure which transforms the fractional output back to a discrete solution. Now let us explain the relaxation that is used in continuous greedy algorithms. There are two key elements needed to define the relaxation: the domain that contains fractional solutions, and a function to evaluate fractional solutions. In our problem, the most natural representation for the domain is the matroid polytope $P(\mathcal{M})$, which is the convex hull of the incidence vectors of the independent sets in the matroid, i.e.: $P(\mathcal{M}) = \text{Conv}(\{\mathbf{1}_S : \forall S \in \mathcal{I}\})$ [8]. Here $\mathbf{1}_S$ denotes the incidence vector of a set $S \subseteq E$, which contains a 1 for each element in S and 0 elsewhere, thus $\mathbf{1}_S \in [0, 1]^E$. The function for evaluation of fractional solutions is the multilinear extension: given a set E and a submodular set function $f : 2^E \rightarrow \mathbb{R}^+$, its multilinear extension $F : [0, 1]^E \rightarrow \mathbb{R}^+$ is defined as:

$$F(\mathbf{y}) \triangleq \mathbb{E}[f(R(\mathbf{y}))] = \sum_{S \subseteq E} f(S) \prod_{i \in S} y_i \prod_{j \notin S} (1 - y_j), \quad (3)$$

where $R(\mathbf{y})$ is a random set containing each element $i \in E$ with probability y_i . Additionally, F can be lower bounded in terms of f :

Lemma 1.2. (from [12]) Let $f : 2^E \rightarrow \mathbb{R}^+$ be a submodular function; let $F : [0, 1]^E \rightarrow \mathbb{R}^+$ be the multilinear extension of f ; and let $\mathbf{y} \in [0, 1]^E$ be a vector whose components are bounded by a , i.e. $y_i \leq a, \forall i \in E$. Then, for every $S \subseteq E$, we have that

$$F(\mathbf{y} \vee \mathbf{1}_S) \geq (1 - a)f(S). \quad (4)$$

A key magnitude used in the Measured Continuous Greedy Algorithm is the marginal value of an element $e \in E$ given a fractional solution $\mathbf{y} \in P(\mathcal{M})$, we refer to it by $\Delta F_e(\mathbf{y})$, defined as follows:

$$\Delta F_e(\mathbf{y}) \triangleq \mathbb{E}[f_{R(\mathbf{y})}(e)] = F(\mathbf{y} \vee \mathbf{1}_e) - F(\mathbf{y}). \quad (5)$$

In other words, it is the value that would be created by adding the full element i to the fractional solution \mathbf{y} . Here we have used the common shorthand notation for the marginal value of an element $f_S(e) = f(S + e) - f(S)$.

In general, given an arbitrary non-negative submodular function there is no closed form of multilinear extension that enables us to evaluate it efficiently. The usual way to deal with this is to sample it, and so we need the following concentration inequality that enables us to bound the sampling error:

Lemma 1.3. (Hoeffding Bound, Theorem 2 in [19]) Let X_1, \dots, X_m be independent random variables such that for each i , $a \leq X_i \leq b$, with $a, b \in \mathbb{R}$. Let $\tilde{X} = \frac{1}{m} \sum_{i=1}^m X_i$. Then

$$\Pr[|\tilde{X} - \mathbb{E}(X)| > t] \leq 2e^{-\frac{2t^2}{(b-a)^2}m}.$$

In this work we shall not dwell on the rounding step. Suffice it to say that there exist algorithms, such as Contention Resolution Schemes [20], Pipage-Rounding [21, 11, 22], or Swap-Rounding [13], that given a general non-negative submodular function f , its multilinear extension F , and a point $\mathbf{y} \in P(\mathcal{M})$, find a set $S \in \mathcal{I}$, such that $f(S) \geq F(\mathbf{y})$. The most efficient technique is Swap-Rounding, and its results are used in [14]. However, the swap-rounding results in [13] hinge around Chernoff-like concentration bounds for monotone submodular functions, but in light of the concentration bounds for non-monotone submodular functions in [23], we believe that this technique can be adapted to work with non-monotone submodular functions.

2 Algorithm

The algorithm we present here is based on the Decreasing-Threshold procedure that enabled Badaniyuru and Vondrak [14] to achieve a very efficient algorithm, $O(\frac{nr}{\epsilon^4} \log^2 \frac{n}{\epsilon})$, for maximising a non-negative monotone submodular function subject to a matroid constraint. The continuous greedy algorithm finds an approximate solution to the relaxation, $\mathbf{y} \in P(\mathcal{M})$, by integrating between $t = 0$ to $t = 1$ the differential equation $\frac{d\mathbf{y}}{dt} = \operatorname{argmax}_{v \in P(\mathcal{M})} \sum_{i \in E} \Delta F_i(\mathbf{y})v_i$. The algorithm in [14] is basically the standard continuous greedy algorithm, but uses a much more efficient Decreasing-Threshold procedure to find the maximum marginal improvement direction, i.e. $v_i \forall i \in E$ in the equation above, in each iteration. To enable an approximation algorithm for the non-monotone case Feldman [12] proposed the measured continuous greedy algorithm, which similarly integrates

Algorithm 1: Accelerated Measured Continuous Greedy

Input : $f : 2^E \rightarrow \mathbb{R}^+$, $\epsilon \in [0, 1]$, $\mathcal{I} \subseteq 2^E$.
Output: A point $\mathbf{y} \in P(\mathcal{M})$, such that $F(\mathbf{y}) \geq (\frac{1}{e} - 2\epsilon)f(OPT)$.

// Initialisation
 $\mathbf{y}(0) \leftarrow \mathbf{0}$
 $\delta \leftarrow \epsilon$

// Main Loop
for $t = \{0, \delta, 2\delta, 3\delta, \dots, 1 - \delta\}$ **do**
 $B(t) \leftarrow \text{Decreasing-Threshold}(f, \mathbf{y}(t), \epsilon, \delta, \mathcal{I})$
 for $i \in E$ **do**
 if $i \in B(t)$ **then**
 $y_i(t + \delta) \leftarrow 1 + e^{-\delta}(y_i(t) - 1)$
 else
 $y_i(t + \delta) \leftarrow y_i(t)$

Return: $\mathbf{y}(1)$

Algorithm 2: Decreasing-Threshold

Input : $f : 2^E \rightarrow \mathbb{R}^+$, $\mathbf{y} \in [0, 1]^E$, $\epsilon \in [0, 1]$, $\delta \in [0, 1]$, $\mathcal{I} \subseteq 2^E$.
Output: A set $B \subseteq E$, such that $B \in \mathcal{I}$.

// Initialisation
 $B \leftarrow \emptyset$;
 $\bar{d} \leftarrow \max_{i \in E} f(i)$;
 $\bar{y}' \leftarrow \max_{i \in E} (1 + e^{-\delta}(y_i - 1))$;

// Main Loop
for ($w = \bar{d}$; $w \geq \epsilon \frac{\bar{d}}{r} (1 - \bar{y}')$; $w \leftarrow w(1 - \epsilon)$) **do**
 for $e \in E$ **do**
 $w_e(B, \mathbf{y}) \leftarrow \Delta F_e(\mathbf{y}(B, \delta))$,
 // averaging $O\left(\frac{r^2}{\epsilon^2} \left(\frac{\bar{d} + d}{\bar{d}}\right)^2 \log(|E|)\right)$ iid random samples.
 if $w_e(B, \mathbf{y}) \geq w$ and $B + e \in \mathcal{I}$ **then**
 $B \leftarrow B + e$

Return: B

*Note that the notation $\mathbf{y}(B, \delta)$ means $y_i(B, \delta) = y_i(t)$ for $i \notin B$, and $y_i(B, \delta) = 1 + e^{-\delta}(y_i - 1)$ for $i \in B$.

$\frac{d\mathbf{y}(t)}{dt} = (1 - \mathbf{y}(t)) \odot \operatorname{argmax}_{v \in P(\mathcal{M})} \sum_{i \in E} \Delta F_i(\mathbf{y}) v_i$ (\odot represents element by element multiplication here).

The advantage of the Decreasing-Threshold procedure is that it requires a smaller number of integration steps and a smaller number of samples per step. Here we apply its key idea to the measured continuous greedy. However, we need to perform two crucial modifications. The first modification is a smoother version of the integration rule in the measured continuous greedy algorithm of Feldman et al. [12]. We use an update step that makes the algorithm more continuous-like. Instead of using the integration step $y(t + \delta) = y(t) + \delta(1 - y(t))$ as proposed by [12], we increment the solution at a rate of $\frac{dy(t)}{dt} = 1 - y(t)$ by using the integration rule: $y(t + \delta) = 1 + e^{-\delta}(y(t) - 1)$, see Algorithm 1. This allows us to smooth the tradeoff between error and running time. The second modification is in the sampling error up to which we estimate the marginal values of the multilinear extension in the Decreasing-Threshold procedure. Badanidiyuru and Vondrak [14] used an additive and multiplicative error bound which enabled them to use only $O(\frac{1}{\epsilon^2} r \log(n))$ samples. However, since our algorithm is meant to work with non-monotone functions, we may have negative marginal values, this increases the number of samples required for two reasons: first, it prevents us from using a multiplicative-additive error bound; and second the marginal values, when sampled, have inherently more spread. To quantify our error we use Hoeffding's concentration inequality, which forces us to use $O(\frac{r^2}{\epsilon^2} (\frac{\bar{d}+d}{d})^2 \log(n))$, where $\underline{d}, \bar{d} \in \mathbb{R}^+$ are the absolute values of the minimum and maximum marginal values that the function f can take, i.e. $-\underline{d} \leq f_S(i) \leq \bar{d}$, for all $i \in E$ and $S \subseteq E$. (Recall that if f were to be monotone we would have $\underline{d} = 0$.) The resulting Decreasing-Threshold procedure is presented in Algorithm 2, and it takes an additional $O(r(\frac{\bar{d}+d}{d})^2)$ value oracle calls. Using only an additive bound instead of an additive and multiplicative introduces an extra $O(r)$ factor in the number of samples required per evaluation of the multilinear extension. While we also incur the additional $O((\frac{\bar{d}+d}{d})^2)$ term to cope with the larger spread that is introduced by the possibility of negative marginal values. This results in an algorithm that finds an $\frac{1}{\epsilon} - \epsilon$ approximation with a number of value oracle calls of $O(\frac{nr^2}{\epsilon^4} (\frac{\bar{d}+d}{d})^2 \log^2(\frac{n}{\epsilon}))$. It is not as efficient as the algorithm for the monotone case, but to the best of our knowledge, this constitutes the first practical $\frac{1}{\epsilon} - \epsilon$ approximation algorithm for general non-negative submodular function subject to a matroid constraint.

3 Algorithm Analysis

We split the analysis of the accelerated measured continuous greedy algorithm in three parts: first we show that the solution produced is feasible, i.e. $\mathbf{y}(1) \in P(\mathcal{M})$; then we prove the $\frac{1}{\epsilon} - \epsilon$ approximation; and finally we study its running time.

Feasibility

Theorem 3.1. *The accelerated measured continuous greedy algorithm produces a feasible fractional solution, i.e., $\mathbf{y}(1) \in P(\mathcal{M})$.*

Proof. We follow the approach used by [12]. We first define a vector \mathbf{x} that coordinate wise upper-bounds $\mathbf{y}(1)$. Then, given that $P(\mathcal{M})$ is down-monotone, we only need to show that \mathbf{x} is in $P(\mathcal{M})$ to show that $\mathbf{y}(1) \in P(\mathcal{M})$. Consider the vector $\mathbf{x} = \delta \sum_{t=0}^{\frac{1}{\delta}-1} \mathbf{1}_{B(\mathbf{y}(t\delta))}$. This is a coordinate-wise upper bound of $\mathbf{y}(1)$ because when $i \in B$, we have that $y_i(t+\delta) - y_i(t) = 1 + e^{-\delta}(y_i(t) - 1) - y_i(t) =$

$(1 - e^{-\delta})(1 - y_i(t)) \leq 1 - e^{-\delta} \leq \delta$, for all $\delta \in [0, 1]$; and when $i \notin B$ $y_i(t + \delta) - y_i(t) = 0$. We now show that \mathbf{x} is in $P(\mathcal{M})$. First, note that by definition $\mathbf{1}_B \in P(\mathcal{M})$. Then, observe that \mathbf{x}/δ is the sum of $\frac{1}{\delta}$ points in $P(\mathcal{M})$, thus $(\mathbf{x}/\delta)/(1/\delta) = \mathbf{x}$ is a convex combination of points in $P(\mathcal{M})$, hence $\mathbf{x} \in P(\mathcal{M})$, and consequently $\mathbf{y}(1) \in P(\mathcal{M})$. \square

In fact, it is possible to find a point that still lies in $P(\mathcal{M})$ even with a $t > 1$, specifically stopping at a value depending on a magnitude called the density of the matroid. This yields tighter approximation bounds that match those found for particular matroids, such as partition matroids. Asymptotically, however, these bounds are the same as those presented here. Therefore, in the aim of simplicity, the analysis is carried out with a stopping time of 1. The interested reader is referred to [12].

Now let us establish a bound to the coordinates of $\mathbf{y}(t)$ that will become useful later in the analysis of the approximation ratio.

Lemma 3.2. *At time $0 \leq t \leq 1$, we have that:*

$$y_i(t) \leq 1 - e^{-t}, \quad \forall i \in E. \quad (6)$$

Proof. Consider the recurrence $g(n+1) = 1 + e^{-\delta}(1 - g(n))$, with $g(0) = 0$. This recurrence has the following solution: $g(n) = 1 - e^{-n\delta}$. Now, in our algorithm t is incremented linearly, so the number of iteration, n , and t are related by $t = n\delta$. At t , all the coordinates of \mathbf{y} either, they stay constant, i.e. $y_i(t + \delta) = y_i(t)$, when $i \notin B$, or increase, i.e. $y_i(t + \delta) = 1 + e^{-\delta}(y_i(t) - 1)$, when $i \in B$. Hence, given that $g(n)$ is non-decreasing, the recurrence g is an upper bound because it corresponds to incrementing in each and every iteration. Consequently, at t , $y(t) \leq g(\frac{t}{\delta}) = 1 - e^{-t}$. \square

Approximation Ratio First we show the gain that the algorithm makes in a single step, and then use this to build a recurrence relation that yields the approximation ratio. But before we do that, we bound the error introduced by sampling:

Corollary 3.3. *Given a non-negative submodular function $f : 2^E \rightarrow \mathbb{R}^+$, and a point $\mathbf{y} \in [0, 1]^E$; let $\underline{d}, \bar{d} \in \mathbb{R}^+$ be the minimum and maximum marginal values of f , such that $-\underline{d} \leq f_S(j) \leq \bar{d}$ for all $S \subseteq E$ and $j \in E$; let R_1, R_2, \dots, R_m be iid samples drawn from $R(\mathbf{y})$, let $w_j(\mathbf{y}) = \frac{1}{m} \sum_{i=1}^m f_{R_i}(j)$; and let $f(OPT) = \max_{S \in \mathcal{I}} f(S)$. Then,*

$$\Pr(|w_j(\mathbf{y}) - \Delta F_j(\mathbf{y})| \geq \beta f(OPT)) \leq 2e^{-2m\beta^2 \left(\frac{\bar{d}}{\bar{d} + \underline{d}}\right)^2}.$$

Proof. Immediate application from the Hoeffding bound in Lemma 1.3, noting that $f(OPT) = \max_{S \in \mathcal{I}} f(S) \geq \max_{e \in E} f(\{e\}) = \bar{d}$. \square

Now we can present the improvement made by the algorithm in a single step:

Lemma 3.4. *Let OPT be an optimal solution. Given a fractional solution \mathbf{y} , the Decreasing-Threshold produces a set B such that, with $\mathbf{y}' = \mathbf{1} + e^{-\mathbf{1}_B \delta} \odot (\mathbf{y} - \mathbf{1})$, we have:*

$$F(\mathbf{y}') - F(\mathbf{y}) \geq (1 - e^{-\delta}) \left((1 - 4\epsilon)(1 - \bar{y}') f(OPT) - F(\mathbf{y}') \right) \quad (7)$$

Proof. This proof follows closely the proof of Claim 4.1 in [14] but with several modifications to avoid assuming monotonicity, namely: the stopping threshold, the sampling error, and the increment bound. Assume that the Decreasing-Threshold procedure returns a sequence of r elements $B = \{b_1, b_2, \dots, b_r\}$, indexed in the order in which they were chosen. Let $O = \{o_1, o_2, \dots, o_r\}$ be an optimal solution indexed as per the exchange property of the matroids in Lemma 1.1 such that $\phi(b_i) = o_i$. Additionally, let B_i and O_i denote the first i elements of B and O respectively, i.e. B_i is the sequence in which the elements have been added to B in algorithm 2 up until the i th element was added. If the procedure returns fewer than r elements or the optimal solution contained fewer than r elements, formally we just add dummy elements with value 0, so that $|B| = r$ and $|O| = r$.

Now let us bound the marginal values of the elements selected by the Decreasing-Threshold procedure with respect to those in the optimal solution. Recall that $\mathbf{y}(S, \delta)$ is the notation that we use to refer to the point such that $y_k(S, \delta) = y_k$ if $k \notin S$ and $y_k(S, \delta) = 1 + e^{-\delta}(y_k - 1)$ if $k \in S$. When b_i is selected, let w be the current threshold, hence $w_{b_i}(\mathbf{y}(B_{i-1}, \delta)) \geq w$. At this point o_i is a candidate element, thus we have one of two situations depending on whether the procedure has finished: if the threshold has not dropped below $\frac{\epsilon}{r}d(1 - \bar{y}')$, the value of $w_{o_i}(\mathbf{y}(B_{i-1}, \delta))$ must be below the threshold in the previous iteration, i.e. $w_{o_i}(\mathbf{y}(B_{i-1}, \delta)) \leq \frac{w}{(1-\epsilon)}$ (otherwise it would have been chosen already); conversely, if the procedure has terminated, b_i is a dummy element with value 0, and the value of $w_{o_i}(\mathbf{y}(B_{i-1}, \delta))$ is below the stopping threshold, i.e. $w_{o_i}(\mathbf{y}(B_{i-1}, \delta)) \leq \epsilon \frac{\bar{d}}{r}(1 - \bar{y}')$. Consequently, we can relate the marginal value estimate of b_i to that of o_i :

$$w_{b_i}(\mathbf{y}(B_{i-1}, \delta)) \geq (1 - \epsilon)w_{o_i}(\mathbf{y}(B_{i-1}, \delta)) - \epsilon \frac{\bar{d}}{r}(1 - \bar{y}').$$

Note that when b_i is selected we have that $B_{i-1} + b_i \in \mathcal{I}$, and by the definition of o_i (i.e. the matroid exchange property) $B_{i-1} + o_i \in \mathcal{I}$.

Now we need to bound the error incurred by sampling. From Lemma 3.3 we can sample the marginal values up to an additive error of $\beta = \frac{\epsilon}{r}f(OPT)(1 - \bar{y}')$ by taking the average of $O\left(\frac{r^2}{\epsilon^2} \left(\frac{\bar{d}+d}{d}\right)^2 \log(|E|)\right)$ samples with high probability (i.e. with a bad estimate probability decreasing with $\frac{1}{|E|}$). (We do not have a term $O(\frac{1}{1-\bar{y}'})$ in the number of samples because our stopping time, $t = 1$, and the update rule enforce (see Lemma 3.2) that $(1 - \bar{y}') \geq \frac{1}{e}$). Thus, we can write:

$$\begin{aligned} w_{b_i}(\mathbf{y}(B_{i-1}, \delta)) &\leq \Delta F_{b_i}(\mathbf{y}(B_{i-1}, \delta)) + \frac{\epsilon}{r}f(OPT)(1 - \bar{y}') \\ w_{o_i}(\mathbf{y}(B_{i-1}, \delta)) &\geq \Delta F_{o_i}(\mathbf{y}(B_{i-1}, \delta)) - \frac{\epsilon}{r}f(OPT)(1 - \bar{y}') \end{aligned}$$

which, with $\bar{d} \leq f(OPT)$, can be combined with the prior bound to yield:

$$\Delta F_{b_i}(\mathbf{y}(B_{i-1}, \delta)) \geq (1 - \epsilon)\Delta F_{o_i}(\mathbf{y}(B_{i-1}, \delta)) - 3\frac{\epsilon}{r}f(OPT)(1 - \bar{y}'). \quad (8)$$

Then, we can bound the improvement that the Decreasing-Threshold procedure obtains:

$$\begin{aligned}
& F(\mathbf{y}') - F(\mathbf{y}) \\
&= \sum_{i=1}^r (F(\mathbf{y}(B_i, \delta)) - F(\mathbf{y}(B_{i-1}, \delta))) \\
&= \sum_{i=1}^r (y'_{b_i} - y_{b_i}) \frac{\partial F}{\partial y_{b_i}} \Big|_{\mathbf{y}=\mathbf{y}(B_{i-1}, \delta)} \\
&= \sum_{i=1}^r (1 - e^{-\delta})(1 - y_{b_i}) \frac{\partial F}{\partial y_{b_i}} \Big|_{\mathbf{y}=\mathbf{y}(B_{i-1}, \delta)} \\
&= (1 - e^{-\delta}) \sum_{i=1}^r \Delta F_{b_i}(\mathbf{y}(B_{i-1}, \delta)) \\
&\geq (1 - e^{-\delta}) \sum_{i=1}^r \left((1 - \epsilon) \Delta F_{o_i}(\mathbf{y}(B_{i-1}, \delta)) - 3 \frac{\epsilon}{r} f(OPT)(1 - \bar{y}') \right) \\
&= (1 - e^{-\delta}) \left((1 - \epsilon) \sum_{i=1}^r \left(\Delta F_{o_i}(\mathbf{y}(B_{i-1}, \delta)) \right) - 3\epsilon f(OPT)(1 - \bar{y}') \right) \\
&\geq (1 - e^{-\delta}) \left((1 - \epsilon) (F(\mathbf{y}' \vee \mathbf{1}_{OPT}) - F(\mathbf{y}')) - 3\epsilon f(OPT)(1 - \bar{y}') \right) \\
&\geq (1 - e^{-\delta}) \left((1 - 4\epsilon)(1 - \bar{y}') f(OPT) - F(\mathbf{y}') \right).
\end{aligned}$$

The second equality comes from the the multilinearity of F . With the update step $y' = 1 + e^{-\delta}(y - 1)$, we have that the increment is $y' - y = (1 - e^{-\delta})(1 - y)$, which gives the third inequality. The fourth equality is by definition of ΔF_e . The fifth inequality is by the bound in equation 8, the sixth by submodularity, and the last one by Lemma 1.2. \square

Finally, we can use the above result to build a recurrence relation that yields the $\frac{1}{e} - \epsilon$ approximation ratio.

Theorem 3.5. *The accelerated measured continuous greedy algorithm returns a point $\mathbf{y}^* \in P(\mathcal{M})$, such that:*

$$F(\mathbf{y}^*) \geq \left(\frac{1}{e} - 2\epsilon \right) f(OPT). \quad (9)$$

Proof. From Lemma 3.4 we have that:

$$F(\mathbf{y}(t + \delta)) - F(\mathbf{y}(t)) \geq (1 - e^{-\delta}) \left((1 - 4\epsilon)(1 - \bar{y}(t + \delta)) f(OPT) - F(\mathbf{y}(t + \delta)) \right)$$

We can now use the bound on the value of the coordinates of \mathbf{y} in Lemma 3.2, we have that $y_i(t) \leq 1 - e^{-t} \quad \forall i \in E$, hence we can write:

$$F(\mathbf{y}(t + \delta)) - F(\mathbf{y}(t)) \geq (1 - e^{-\delta}) \left((1 - 4\epsilon)e^{-(\delta+t)} f(OPT) - F(\mathbf{y}(t + \delta)) \right). \quad (10)$$

Consider the recurrence relation $a(n+1) - a(n) = k_1(k_2 \exp(-(n+1)\epsilon) - a(n+1))$, which, with $a(0) = a_0$, has the following solution

$$a(n) = \frac{\left(\frac{1}{k_1+1}\right)^n (a_0(-k_1 + e^\epsilon - 1) + k_1 k_2) - k_1 k_2 e^{-n\epsilon}}{-k_1 + e^\epsilon - 1}. \quad (11)$$

This recurrence is equivalent to equation 3 if we set $k_1 = (1 - e^{-\delta})$, $k_2 = (1 - 4\epsilon)$, $n = \frac{t}{\delta}$, and $\delta = \epsilon$. Thus, substituting and simplifying assuming that $F(\mathbf{0}) \geq 0$ (due to the non-negativity of f), we have the following lower bound on the value of the solution $\mathbf{y}(t)$ for any time $t \in [0, 1]$:

$$F(\mathbf{y}(t)) \geq \frac{(1 - e^{-\epsilon})(1 - 4\epsilon)e^{\epsilon(-(\frac{t}{\epsilon}+1))} (e^{\epsilon(\frac{t}{\epsilon}+1)} (\frac{1}{2-e^{-\epsilon}})^{t/\epsilon} - e^\epsilon)}{e^{-\epsilon} + e^\epsilon - 2} f(OPT).$$

So, when the algorithm ends at $t = 1$, we have:

$$F(\mathbf{y}(1)) \geq \frac{(1 - e^{-\epsilon})(1 - 4\epsilon)e^{\epsilon(-(\frac{1}{\epsilon}+1))} (e^{\epsilon(\frac{1}{\epsilon}+1)} (\frac{1}{2-e^{-\epsilon}})^{1/\epsilon} - e^\epsilon)}{e^{-\epsilon} + e^\epsilon - 2} f(OPT).$$

We can find a more intuitive version of this bound by observing that, clearly, for $0 \leq \epsilon \leq 1$:

$$\frac{1}{e} - \frac{(1 - e^{-\epsilon})(1 - 4\epsilon)e^{\epsilon(-(\frac{1}{\epsilon}+1))} (e^{\epsilon(\frac{1}{\epsilon}+1)} (\frac{1}{2-e^{-\epsilon}})^{1/\epsilon} - e^\epsilon)}{e^{-\epsilon} + e^\epsilon - 2} \leq \frac{5}{e}\epsilon \leq 2\epsilon \quad (12)$$

Hence:

$$F(\mathbf{y}(1)) \geq \left(\frac{1}{e} - 2\epsilon\right) f(OPT). \quad (13)$$

Finally, from Lemma 3.2 we have that $\mathbf{y}(1) \in P(\mathcal{M})$. □

Running Time As it is common in the submodular maximisation literature, we quantify the running time in terms of value oracle calls to the submodular function and matroid independence oracle calls. We analyse the running time of the algorithm in two steps: first we study the running time of the Decreasing-Threshold procedure, and then that of the continuous greedy.

Lemma 3.6. *The Decreasing-Threshold procedure makes*

$O\left(\frac{|E|r^2}{\epsilon^3} \log(|E|) \log\left(\frac{r}{\epsilon}\right) \left(\frac{\bar{d}+d}{d}\right)^2\right)$ value oracle calls, and $O\left(\frac{|E|}{\epsilon} \log\frac{r}{\epsilon}\right)$ independence oracle calls.

Proof. First, the number of values that the threshold takes in the Decreasing-Threshold procedure to reach the stopping threshold is, considering that the term $(1 - \bar{y}) \geq \frac{1}{e}$ due to Lemma 3.2, $O\left(\frac{\log\frac{r}{\epsilon}}{\log(1-\epsilon)}\right)$. Second, for each threshold value, the algorithm performs $O(|E|)$ estimates of ΔF_ϵ , and $O(|E|)$ calls to the independence oracle. Therefore, the number of independence oracle calls is the number of calls per threshold step multiplied by the number of threshold steps, i.e.:

$$O\left(\frac{|E|}{\epsilon} \log\frac{r}{\epsilon}\right), \quad (14)$$

which has been simplified noting that $\frac{\log \frac{r}{\epsilon}}{\log 1-\epsilon} \leq \frac{1}{\epsilon} \log(\frac{r}{\epsilon})$.

Now, each estimate of ΔF_ϵ requires $O\left(\frac{r^2}{\epsilon^2} \left(\frac{\bar{d}+d}{d}\right)^2 \log(|E|)\right)$ samples. Hence, we can conclude that the number of value oracle calls is

$$O\left(\frac{|E|r^2}{\epsilon^3} \log(|E|) \log\left(\frac{r}{\epsilon}\right) \left(\frac{\bar{d}+d}{d}\right)^2\right).$$

□

We can now quantify the running time of the whole algorithm.

Theorem 3.7. *The accelerated measured continuous greedy algorithm makes*

$O\left(\frac{|E|r^2}{\epsilon^4} \log(|E|) \log\left(\frac{r}{\epsilon}\right) \left(\frac{\bar{d}+d}{d}\right)^2\right)$ *value oracle calls, and* $O\left(\frac{|E|}{\epsilon^2} \log\left(\frac{r}{\epsilon}\right)\right)$ *independence oracle calls.*

Proof. Considering that the number of steps of the procedure, with $\delta = \epsilon$, is $\frac{1}{\epsilon}$. The number of calls to both oracles is simply the number of calls by the Decreasing-Threshold procedure (Algorithm 2) in Lemma 3.6 multiplied by the number of steps in the continuous greedy algorithm (Algorithm 1). □

4 Discussion and Future Work

In this paper we have presented a $\frac{1}{\epsilon} - \epsilon$ -approximation algorithm for general non-negative submodular function maximisation that requires $O\left(\frac{nr^2}{\epsilon^4} \left(\frac{\bar{d}+d}{d}\right)^2 \log^2\left(\frac{n}{\epsilon}\right)\right)$ value oracle calls. This is the fastest $\frac{1}{\epsilon}$ -approximation algorithm currently available, which enables the use of general (non-monotone) matroid-constrained submodular maximisation for many applications for which existing algorithms were implausibly slow. We think this is of great significance because there has been a recent surge of interest for applying submodular maximisation in fields where large problem instances are paramount, such as Machine Learning [24, 25, 5], particularly in the field of summarisation where non-monotone submodular functions are natural [26, 27]. Our algorithm is slower than the one presented for the monotone case in [14] by $O\left(r \left(\frac{\bar{d}+d}{d}\right)^2\right)$ due to the inability to sample the marginal values of the multilinear extension up to an additive and multiplicative bound. If we could, then we would reduce the additional value oracle calls required to achieve an algorithm with the same running time, we believe this might be possible.

A future avenue of research would be to combine our work with the very interesting results in [28], where an efficient algorithm is proposed to allow the trade-off of value oracle calls and matroid independence calls for non-negative monotone submodular functions, to enable query trade-off for general non-negative submodular functions. Another interesting path is to combine the more continuous-like measured continuous greedy update step that we present here with the acceleration techniques for strong submodular functions presented in [29] to produce an adaptive step algorithm. This way, in each step we could use a large δ that extended to the boundary of the region of validity of the set B , instead of taking a δ that is small enough to satisfy the worst case. Another obvious improvement on the algorithms presented here would be to combine the ideas from the Lazy Greedy Algorithm [30] to adaptively change the decrement of the threshold in the Decreasing-Threshold procedure.

References

- [1] A. Krause and D. Golovin, “Submodular function maximization,” *Tractability: Practical Approaches to Hard Problems* **3** no. 19, (2012) 8.
- [2] B. Lehmann, D. Lehmann, and N. Nisan, “Combinatorial auctions with decreasing marginal utilities,” in *Proceedings of the 3rd ACM conference on Electronic Commerce*, pp. 18–28, ACM, 2001.
- [3] S. Dughmi, T. Roughgarden, and M. Sundararajan, “Revenue submodularity,” in *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 243–252, ACM, 2009.
- [4] D. Kempe, J. Kleinberg, and É. Tardos, “Maximizing the spread of influence through a social network,” in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, ACM, 2003.
- [5] H. Lin and J. Bilmes, “Multi-document summarization via budgeted maximization of submodular functions,” in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 912–920, Association for Computational Linguistics, 2010.
- [6] A. Krause, J. Leskovec, C. Guestrin, J. VanBriesen, and C. Faloutsos, “Efficient sensor placement optimization for securing large water distribution networks,” *Journal of Water Resources Planning and Management* **134** no. 6, (2008) 516–526.
- [7] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, “Cost-effective outbreak detection in networks,” in *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 420–429, ACM, 2007.
- [8] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics. Springer Berlin Heidelberg, 2003.
- [9] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, “An analysis of approximations for maximizing submodular set functionsI,” *Mathematical Programming* **14** no. 1, (Dec, 1978) 265–294.
- [10] J. Vondrak, “Optimal approximation for the submodular welfare problem in the value oracle model,” in *Proceedings of the fortieth annual ACM symposium on Theory of computing - STOC 08*, p. 67. ACM Press, New York, New York, USA, May, 2008.
- [11] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, “Maximizing a Monotone Submodular Function Subject to a Matroid Constraint,” *SIAM Journal on Computing* **40** no. 6, (Jan, 2011) 1740–1766.
- [12] M. Feldman, J. Naor, and R. Schwartz, “A Unified Continuous Greedy Algorithm for Submodular Maximization,” in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 570–579. IEEE, Oct, 2011.
- [13] J. Vondrák, C. Chekuri, and R. Zenklusen, “Submodular function maximization via the multilinear relaxation and contention resolution schemes,” in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 783–792, ACM, 2011.

- [14] A. Badanidiyuru and J. Vondrák, “Fast algorithms for maximizing submodular functions,” in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1497–1514, Society for Industrial and Applied Mathematics. 2014.
- [15] U. Feige, “A threshold of $\ln n$ for approximating set cover,” *Journal of the ACM (JACM)* **45** no. 4, (1998) 634–652.
- [16] S. O. Gharan and J. Vondrák, “Submodular maximization by simulated annealing,” in *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pp. 1098–1116, Society for Industrial and Applied Mathematics. 2011.
- [17] A. Ene and H. L. Nguyen, “Constrained submodular maximization: Beyond $1/e$,” in *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pp. 248–257, IEEE. 2016.
- [18] N. Buchbinder and M. Feldman, “Constrained Submodular Maximization via a Non-symmetric Technique,” [arXiv:1611.03253](https://arxiv.org/abs/1611.03253).
- [19] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association* **58** no. 301, (1963) 13–30.
- [20] C. Chekuri, J. Vondrák, and R. Zenklusen, “Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes,” *SIAM Journal on Computing* **43** no. 6, (Nov, 2014) 1831–1879.
- [21] A. Ageev and M. Sviridenko, “Pipage Rounding: A New Method of Constructing Algorithms with Proven Performance Guarantee,” *Journal of Combinatorial Optimization* **8** no. 3, (Sep, 2004) 307–328.
- [22] J. Vondrák, “Symmetry and approximability of submodular maximization problems,” *SIAM Journal on Computing* **42** no. 1, (2013) 265–304.
- [23] J. Vondrak, “A note on concentration of submodular functions,” [arXiv:1005.2791](https://arxiv.org/abs/1005.2791).
- [24] B. Mirzasoleiman, A. Badanidiyuru, and A. Karbasi, “Fast constrained submodular maximization: Personalized data summarization,” in *ICLM’16: Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 2016.
- [25] J. Bilmes and W. Bai, “Deep Submodular Functions,” [arXiv:1701.08939](https://arxiv.org/abs/1701.08939).
- [26] S. Tschischek, R. K. Iyer, H. Wei, and J. A. Bilmes, “Learning mixtures of submodular functions for image collection summarization,” in *Advances in neural information processing systems*, pp. 1413–1421. 2014.
- [27] A. Dasgupta, R. Kumar, and S. Ravi, “Summarization Through Submodularity and Dispersion.,” in *ACL (1)*, pp. 1014–1022. 2013.
- [28] N. Buchbinder, M. Feldman, and R. Schwartz, “Comparing Apples and Oranges: Query Tradeoff in Submodular Maximization,” [arXiv:1410.0773](https://arxiv.org/abs/1410.0773).
- [29] Z. Wang, B. Moran, X. Wang, and Q. Pan, “An accelerated continuous greedy algorithm for maximizing strong submodular functions,” *Journal of Combinatorial Optimization* **30** no. 4, (2015) 1107–1124.
- [30] M. Minoux, “Accelerated greedy algorithms for maximizing submodular set functions,” in *Optimization Techniques*, pp. 234–243. Springer-Verlag, Berlin/Heidelberg, 1978.

2017-03-21

Fast non-monotone submodular maximisation subject to a matroid constraint

Segui-Gasco, Pau

Pau Segui-Gasco and Hyo-Sang Shin. Fast non-monotone submodular maximisation subject to a matroid constraint

<https://arxiv.org/abs/1703.06053>

Downloaded from Cranfield Library Services E-Repository