
Conditional Gradient Method for Stochastic Submodular Maximization: Closing the Gap

Aryan Mokhtari
Massachusetts Institute of Technology

Hamed Hassani
University of Pennsylvania

Amin Karbasi
Yale University

Abstract

In this paper, we study the problem of *constrained* and *stochastic* continuous submodular maximization. Even though the objective function is not concave (nor convex) and is defined in terms of an expectation, we develop a variant of the conditional gradient method, called **Stochastic Continuous Greedy**, which achieves a *tight* approximation guarantee. More precisely, for a monotone and continuous DR-submodular function and subject to a *general* convex body constraint, we prove that **Stochastic Continuous Greedy** achieves a $[(1 - 1/e)\text{OPT} - \epsilon]$ guarantee (in expectation) with $\mathcal{O}(1/\epsilon^3)$ stochastic gradient computations. This guarantee matches the known hardness results and closes the gap between deterministic and stochastic continuous submodular maximization. By using stochastic continuous optimization as an interface, we also provide the first $(1 - 1/e)$ tight approximation guarantee for maximizing a *monotone but stochastic* submodular *set* function subject to a general matroid constraint.

1 Introduction

Many procedures in statistics and artificial intelligence require solving non-convex problems, including clustering [Abbasi and Younis, 2007], training deep neural networks [Bengio et al., 2007], and performing Bayesian optimization [Snoek et al., 2012], to name a few. Historically, the focus has been to convexify non-convex objectives; in recent years, there has been significant progress to optimize non-convex func-

tions directly. This direct approach has led to provably good guarantees, for specific problem instances. Examples include latent variable models [Anandkumar et al., 2014], non-negative matrix factorization [Arora et al., 2016], robust PCA [Netrapalli et al., 2014], matrix completions [Ge et al., 2016], and training certain specific forms of neural networks [Mei et al., 2016]. However, it is well known that in general finding the global optimum of a non-convex optimization problem is NP-hard [Murty and Kabadi, 1987]. This computational barrier has mainly shifted the goal of non-convex optimization towards two directions: a) finding an *approximate* local minimum by avoiding saddle points [Ge et al., 2015, Anandkumar and Ge, 2016, Jin et al., 2017, Paternain et al., 2017], or b) characterizing general conditions under which the underlying non-convex optimization is tractable [Hazan et al., 2016].

In this paper, we consider a broad class of non-convex optimization problems that possess special combinatorial structures. More specifically, we focus on constrained maximization of stochastic continuous submodular functions (CSF) that demonstrate diminishing returns, i.e., continuous DR-submodular functions,

$$\max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}) \doteq \max_{\mathbf{x} \in \mathcal{C}} \mathbb{E}_{\mathbf{z} \sim P} [\tilde{F}(\mathbf{x}, \mathbf{z})]. \quad (1)$$

Here, the functions $\tilde{F} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ are stochastic where $\mathbf{x} \in \mathcal{X}$ is the optimization variable, $\mathbf{z} \in \mathcal{Z}$ is a realization of the random variable \mathbf{Z} drawn from a distribution P , and $\mathcal{X} \subset \mathbb{R}_+^n$ is a compact set. Our goal is to maximize the expected value of the random functions $\tilde{F}(\mathbf{x}, \mathbf{z})$ over the convex body $\mathcal{C} \subseteq \mathbb{R}_+^n$. Note that we *only assume* that $F(\mathbf{x})$ is DR-submodular, and *not necessarily* the stochastic functions $\tilde{F}(\mathbf{x}, \mathbf{z})$. We also consider situations where the distribution P is either unknown (e.g., when the objective is given as an implicit stochastic model) or the domain of the random variable \mathbf{Z} is very large (e.g., when the objective is defined in terms of an empirical risk) which makes the cost of computing the expectation very high. In these regimes, stochastic optimization methods, which operate on computationally cheap estimates of gradients, arise as natural solutions. In fact, very recently, it was

shown in [Hassani et al., 2017] that stochastic gradient methods achieve a $(1/2)$ approximation guarantee to Problem (1). The authors also showed that current versions of the conditional gradient method (a.k.a., Frank-Wolfe), such as continuous greedy [Vondrák, 2008] or its close variant [Bian et al.], can perform arbitrarily poorly in stochastic continuous submodular maximization settings.

Our contributions. We provide the first tight $(1 - 1/e)$ approximation guarantee for Problem (1) when the continuous function F is monotone, smooth, DR-submodular, and the constraint set \mathcal{C} is a bounded convex body. To this end, we develop a novel conditional gradient method, called **Stochastic Continuous Greedy (SCG)**, that produces a solution with an objective value larger than $((1 - 1/e)\text{OPT} - \epsilon)$ after $O(1/\epsilon^3)$ iterations while only having access to unbiased estimates of the gradients (here OPT denotes the optimal value of Problem (1)). SCG is also memory efficient in the following sense: in contrast to previously proposed conditional gradient methods in stochastic convex [Hazan and Luo, 2016] and non-convex [Reddi et al., 2016] settings, SCG does not require using a minibatch in each step. Instead it simply averages over the stochastic estimates of the previous gradients.

Connection to Discrete Problems. Even though submodularity has been mainly studied in discrete domains [Fujishige, 2005], many efficient methods for optimizing such submodular set functions rely on continuous relaxations either through a multi-linear extension [Vondrák, 2008] (for maximization) or Lovasz extension [Lovász, 1983] (for minimization). In fact, Problem (1) has a discrete counterpart, recently considered in [Hassani et al., 2017, Karimi et al., 2017]:

$$\max_{S \in \mathcal{I}} f(S) \doteq \max_{S \in \mathcal{I}} \mathbb{E}_{\mathbf{z} \sim P} [\tilde{f}(S, \mathbf{z})], \quad (2)$$

where the functions $\tilde{f} : 2^V \times \mathcal{Z} \rightarrow \mathbb{R}_+$ are stochastic, S is the optimization set variable defined over a ground set V , $\mathbf{z} \in \mathcal{Z}$ is the realization of a random variable \mathbf{Z} drawn from the distribution P , and \mathcal{I} is a general matroid constraint. Since P is unknown, problem (2) cannot be directly solved using the current state-of-the-art techniques. Instead, Hassani et al. [2017] showed that by lifting the problem to the continuous domain (via multi-linear relaxation) and using stochastic gradient methods on a continuous relaxation to reach a solution that is within a factor $(1/2)$ of the optimum. Contemporarily, [Karimi et al., 2017] used a concave relaxation technique to provide a $(1 - 1/e)$ approximation for the class of submodular coverage functions. Our work also closes the gap for maximizing the stochastic submodular set maximization, namely, Problem (2), by providing the first tight $(1 - 1/e)$ approximation

guarantee for general monotone submodular set functions subject to a matroid constraint.

Notation. Lowercase boldface \mathbf{v} denotes a vector and uppercase boldface \mathbf{A} a matrix. We use $\|\mathbf{v}\|$ to denote the Euclidean norm of vector \mathbf{v} . The i -th element of the vector \mathbf{v} is written as v_i and the element on the i -th row and j -th column of the matrix \mathbf{A} is denoted by $A_{i,j}$.

2 Related Work

Maximizing a deterministic submodular set function has been extensively studied. The celebrated result of Nemhauser et al. [1978] shows that a greedy algorithm achieves a $(1 - 1/e)$ approximation guarantee for a monotone function subject to a cardinality constraint. It is also known that this result is tight under reasonable complexity-theoretic assumptions [Feige, 1998]. Recently, variants of the greedy algorithm have been proposed to extend the above result to non-monotone and more general constraints [Feige et al., 2011, Buchbinder et al., 2015, 2014, Feldman et al., 2017]. While discrete greedy algorithms are fast, they usually do not provide the tightest guarantees for many classes of feasibility constraints. This is why continuous relaxations of submodular functions, e.g., the multilinear extension, have gained a lot of interest [Vondrák, 2008, Calinescu et al., 2011, Chekuri et al., 2014, Feldman et al., 2011, Gharan and Vondrák, 2011, Sviridenko et al., 2015]. In particular, it is known that the continuous greedy algorithm achieves a $(1 - 1/e)$ approximation guarantee for monotone submodular functions under a general matroid constraint [Calinescu et al., 2011]. An improved $((1 - e^{-c})/c)$ -approximation guarantee can be obtained if f has curvature c [Vondrák, 2010].

Continuous submodularity naturally arises in many learning applications such as robust budget allocation [Staub and Jegelka, 2017, Soma et al., 2014], online resource allocation [Eghbali and Fazel, 2016], learning assignments [Golovin et al., 2014], as well as Adwords for e-commerce and advertising [Devanur and Jain, 2012, Mehta et al., 2007]. Maximizing a *deterministic* continuous submodular function dates back to the work of Wolsey [1982]. More recently, Chekuri et al. [2015] proposed a multiplicative weight update algorithm that achieves $(1 - 1/e - \epsilon)$ approximation guarantee after $\tilde{O}(n/\epsilon^2)$ oracle calls to gradients of a monotone smooth submodular function F (i.e., twice differentiable DR-submodular) subject to a polytope constraint. A similar approximation factor can be obtained after $\mathcal{O}(n/\epsilon)$ oracle calls to gradients of F for monotone DR-submodular functions subject to a down-closed convex body using the continuous greedy method [Bian et al.]. However, such results require ex-

act computation of the gradients ∇F which is not feasible in Problem (1). An alternative approach is then to modify the current algorithms by replacing gradients $\nabla F(\mathbf{x}_t)$ by their stochastic estimates $\nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t)$; however, this modification may lead to arbitrarily poor solutions as demonstrated in [Hassani et al., 2017]. Another alternative is to estimate the gradient by averaging over a (large) mini-batch of samples at each iteration. While this approach can potentially reduce the noise variance, it increases the computational complexity of each iteration and is not favorable. The work by Hassani et al. [2017] is perhaps the first attempt to solve Problem (1) only by executing stochastic estimates of gradients (without using a large batch). They showed that the stochastic gradient ascent method achieves a $(1/2 - \epsilon)$ approximation guarantee after $O(1/\epsilon^2)$ iterations. Although this work opens the door for maximizing stochastic CSFs using computationally cheap stochastic gradients, it fails to achieve the optimal $(1 - 1/e)$ approximation. To close the gap, we propose in this paper **Stochastic Continuous Greedy** which outputs a solution with function value at least $((1 - 1/e)\text{OPT} - \epsilon)$ after $O(1/\epsilon^3)$ iterations. Notably, our result only requires the expected function F to be monotone and DR-submodular and the stochastic functions \tilde{F} need not be monotone nor DR-submodular. Moreover, in contrast to the result in [Bian et al.], which holds only for down-closed convex constraints, our result holds for any convex constraints.

Our result also has important implications for Problem (2); that is, maximizing a stochastic discrete submodular function subject to a matroid constraint. Since the proposed SCG method works in stochastic settings, we can relax the discrete objective function f in Problem (2) to a continuous function F through the multi-linear extension (note that expectation is a linear operator). Then we can maximize F within a $(1 - 1/e - \epsilon)$ approximation to the optimum value by using only $O(1/\epsilon^3)$ oracle calls to the stochastic gradients of F . Finally, a proper rounding scheme (such as the contention resolution method [Chekuri et al., 2014]) results in a feasible set whose value is a $(1 - 1/e)$ approximation to the optimum set in expectation.

The focus of our paper is on the maximization of stochastic submodular functions. However, there are also very interesting results for minimization of such functions [Staib and Jegelka, 2017, Ene et al., 2017, Chakrabarty et al., 2017, iye].

3 Continuous Submodularity

We begin by recalling the definition of a submodular set function: A function $f : 2^V \rightarrow \mathbb{R}_+$, defined on the

ground set V , is called submodular if for all subsets $A, B \subseteq V$, we have

$$f(A) + f(B) \geq f(A \cap B) + f(A \cup B).$$

The notion of submodularity goes beyond the discrete domain [Wolsey, 1982, Vondrák, 2007, Bach, 2015]. Consider a continuous function $F : \mathcal{X} \rightarrow \mathbb{R}_+$ where the set \mathcal{X} is of the form $\mathcal{X} = \prod_{i=1}^n \mathcal{X}_i$ and each \mathcal{X}_i is a compact subset of \mathbb{R}_+ . We call the continuous function F submodular if for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ we have

$$F(\mathbf{x}) + F(\mathbf{y}) \geq F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}), \quad (3)$$

where $\mathbf{x} \vee \mathbf{y} := \max(\mathbf{x}, \mathbf{y})$ (component-wise) and $\mathbf{x} \wedge \mathbf{y} := \min(\mathbf{x}, \mathbf{y})$ (component-wise). In this paper, our focus is on differentiable continuous submodular functions with two additional properties: monotonicity and diminishing returns. Formally, a submodular function F is monotone (on the set \mathcal{X}) if

$$\mathbf{x} \leq \mathbf{y} \implies F(\mathbf{x}) \leq F(\mathbf{y}), \quad (4)$$

for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$. Note that $\mathbf{x} \leq \mathbf{y}$ in (4) means that $x_i \leq y_i$ for all $i = 1, \dots, n$. Furthermore, a differentiable submodular function F is called *DR-submodular* (i.e., shows diminishing returns) if the gradients are antitone, namely, for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ we have

$$\mathbf{x} \leq \mathbf{y} \implies \nabla F(\mathbf{x}) \geq \nabla F(\mathbf{y}). \quad (5)$$

When the function F is twice differentiable, submodularity implies that all cross-second-derivatives are non-positive [Bach, 2015], i.e.,

$$\text{for all } i \neq j, \text{ for all } \mathbf{x} \in \mathcal{X}, \quad \frac{\partial^2 F(\mathbf{x})}{\partial x_i \partial x_j} \leq 0, \quad (6)$$

and DR-submodularity implies that all second-derivatives are non-positive [Bian et al.], i.e.,

$$\text{for all } i, j, \text{ for all } \mathbf{x} \in \mathcal{X}, \quad \frac{\partial^2 F(\mathbf{x})}{\partial x_i \partial x_j} \leq 0. \quad (7)$$

4 Stochastic Continuous Greedy

In this section, we introduce our main algorithm, **Stochastic Continuous Greedy** (SCG), which is a stochastic variant of the continuous greedy method to solve Problem (1). We only assume that the expected objective function F is monotone and DR-submodular and the stochastic functions $\tilde{F}(\mathbf{x}, \mathbf{z})$ may not be monotone nor submodular. Since the objective function F is monotone and DR-submodular, continuous greedy algorithm [Bian et al., Calinescu et al., 2011] can be used in principle to solve Problem (1). Note that each update of continuous greedy requires computing

Algorithm 1 Stochastic Continuous Greedy (SCG)

Require: Stepsizes $\rho_t > 0$. Initialize $\mathbf{d}_0 = \mathbf{x}_0 = \mathbf{0}$

- 1: **for** $t = 1, 2, \dots, T$ **do**
- 2: Compute $\mathbf{d}_t = (1 - \rho_t)\mathbf{d}_{t-1} + \rho_t \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t)$;
- 3: Compute $\mathbf{v}_t = \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}} \{\langle \mathbf{d}_t, \mathbf{v} \rangle\}$;
- 4: Update the variable $\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{T} \mathbf{v}_t$;
- 5: **end for**

the gradient of F , i.e., $\nabla F(\mathbf{x}) := \mathbb{E}[\nabla \tilde{F}(\mathbf{x}, \mathbf{z})]$. However, if we only have access to the (computationally cheap) stochastic gradients $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$, then the continuous greedy method will not be directly usable [Hasani et al., 2017]. This limitation is due to the non-vanishing variance of gradient approximations. To resolve this issue, we introduce stochastic version of the continuous greedy algorithm which reduces the noise of gradient approximations via a common averaging technique in stochastic optimization [Ruszczynski, 1980, 2008, Yang et al., 2016, Mokhtari et al., 2017].

Let $t \in \mathbf{N}$ be a discrete time index and ρ_t a given stepsize which approaches zero as t approaches infinity. Our proposed estimated gradient \mathbf{d}_t is defined by the following recursion

$$\mathbf{d}_t = (1 - \rho_t)\mathbf{d}_{t-1} + \rho_t \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t), \quad (8)$$

where the initial vector is defined as $\mathbf{d}_0 = \mathbf{0}$. It can be shown that the averaging technique in (8) reduces the noise of gradient approximation as time increases. More formally, the expected noise of gradient estimation $\mathbb{E}[\|\mathbf{d}_t - \nabla F(\mathbf{x}_t)\|^2]$ approaches zero asymptotically (Lemma 2). This property implies that the gradient estimate \mathbf{d}_t is a better candidate for approximating the gradient $\nabla F(\mathbf{x}_t)$ comparing to the unbiased gradient estimate $\nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t)$ that suffers from a high variance approximation. We therefore define the ascent direction \mathbf{v}_t of SCG as follows

$$\mathbf{v}_t = \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}} \{\langle \mathbf{d}_t, \mathbf{v} \rangle\}, \quad (9)$$

which is a linear objective maximization over the convex set \mathcal{C} . Indeed, if instead of the gradient estimate \mathbf{d}_t we use the exact gradient $\nabla F(\mathbf{x}_t)$ for the updates in (9), the continuous greedy update will be recovered. Here, as in continuous greedy, the initial decision vector is the null vector, $\mathbf{x}_0 = \mathbf{0}$. Further, the stepsize for updating the iterates is equal to $1/T$, and the variable \mathbf{x}_t is updated as

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{T} \mathbf{v}_t. \quad (10)$$

The stepsize $1/T$ and the initialization $\mathbf{x}_0 = \mathbf{0}$ ensure that after T iterations the variable \mathbf{x}_T ends up

in the convex set \mathcal{C} . We should highlight that the convex body \mathcal{C} may not be down-closed or contain $\mathbf{0}$. Nonetheless, the solution \mathbf{x}_T returned by SCG will be a feasible point in \mathcal{C} . The steps of the proposed SCG method are outlined in Algorithm 1.

5 Convergence Analysis

In this section, we study the convergence properties of SCG for solving Problem (1). To do so, we first assume that the following conditions hold.

Assumption 1. *The Euclidean norm of the elements in the constraint set \mathcal{C} are uniformly bounded, i.e., for all $\mathbf{x} \in \mathcal{C}$ we can write*

$$\|\mathbf{x}\| \leq D. \quad (11)$$

Assumption 2. *The function F is DR-submodular and monotone. Further, its gradients are L -Lipschitz continuous over the set \mathcal{X} , i.e., for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$*

$$\|\nabla F(\mathbf{x}) - \nabla F(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|. \quad (12)$$

Assumption 3. *The variance of the unbiased stochastic gradients $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ is bounded above by σ^2 , i.e., for any vector $\mathbf{x} \in \mathcal{X}$ we can write*

$$\mathbb{E} \left[\|\nabla \tilde{F}(\mathbf{x}, \mathbf{z}) - \nabla F(\mathbf{x})\|^2 \right] \leq \sigma^2, \quad (13)$$

where the expectation is with respect to the randomness of $\mathbf{z} \sim P$.

Due to the initialization step of SCG (starting from $\mathbf{0}$) we need a bound on the furthest feasible solution from $\mathbf{0}$ that we can end up with; and such a bound is guaranteed by Assumption 1. The condition in Assumption 2 ensures that the objective function F is smooth. Note again that $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ may or may not be Lipschitz continuous. Finally, the required condition in Assumption 3 guarantees that the variance of stochastic gradients $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ is bounded by a finite constant $\sigma^2 < \infty$ which is customary in stochastic optimization.

To study the convergence of SCG, we first derive an upper bound for the expected error of gradient approximation $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ in the following lemma.

Lemma 1. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1. If Assumptions 1-3 are satisfied, then the sequence of expected squared gradient errors $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ for the iterates generated by SCG satisfies*

$$\begin{aligned} & \mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \\ & \leq \left(1 - \frac{\rho_t}{2}\right) \mathbb{E}[\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2] + \rho_t^2 \sigma^2 \\ & \quad + \frac{L^2 D^2}{T^2} + \frac{2L^2 D^2}{\rho_t T^2}. \end{aligned} \quad (14)$$

Proof. See Section 9.1. \blacksquare

The result in Lemma 1 showcases that the expected squared error of gradient approximation $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ decreases at each iteration by the factor $(1 - \rho_t/2)$ if the remaining terms on the right hand side of (14) are negligible relative to the term $(1 - \rho_t/2)\mathbb{E}[\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2]$. This condition can be satisfied, if the parameters $\{\rho_t\}$ are chosen properly. We formalize this claim in the following lemma and show that the expected error $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ converges to zero at a sublinear rate of $\mathcal{O}(t^{-2/3})$.

Lemma 2. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1. If Assumptions 1-3 are satisfied and $\rho_t = \frac{4}{(t+8)^{2/3}}$, then for $t = 0, \dots, T$ we have*

$$\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \leq \frac{Q}{(t+9)^{2/3}}, \quad (15)$$

where $Q := \max\{\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2 9^{2/3}, 16\sigma^2 + 3L^2D^2\}$.

Proof. See Section 9.2. \blacksquare

Let us now use the result in Lemma 2 to show that the sequence of iterates generated by SCG reaches a $(1 - 1/e)$ approximation for Problem (1).

Theorem 1. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1. If Assumptions 1-3 are satisfied and $\rho_t = \frac{4}{(t+8)^{2/3}}$, then the expected objective function value for the iterates generated by SCG satisfies the inequality*

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - 1/e)OPT - \frac{2DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (16)$$

where $OPT = \max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x})$.

Proof. Let \mathbf{x}^* be the global maximizer within the constraint set \mathcal{C} . Based on the smoothness of the function F with constant L we can write

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle - \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= F(\mathbf{x}_t) + \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2, \end{aligned} \quad (17)$$

where the equality follows from the update in (10). Since \mathbf{v}_t is in the set \mathcal{C} , it follows from Assumption 1 that the norm $\|\mathbf{v}_t\|^2$ is bounded above by D^2 . Apply

this substitution and add and subtract the inner product $\langle \mathbf{d}_t, \mathbf{v}_t \rangle$ to the right hand side of (17) to obtain

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{v}_t, \mathbf{d}_t \rangle + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2} \\ &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{x}^*, \mathbf{d}_t \rangle + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2}. \end{aligned} \quad (18)$$

Note that the second inequality in (18) holds since based on (9) we can write $\langle \mathbf{x}^*, \mathbf{d}_t \rangle \leq \max_{\mathbf{v} \in \mathcal{X}} \{\langle \mathbf{v}, \mathbf{d}_t \rangle\} = \langle \mathbf{v}_t, \mathbf{d}_t \rangle$. Now add and subtract the inner product $\langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle / T$ to the RHS of (18) to get

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle \\ &\quad + \frac{1}{T} \langle \mathbf{v}_t - \mathbf{x}^*, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2}. \end{aligned} \quad (19)$$

We further have $\langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle \geq F(\mathbf{x}^*) - F(\mathbf{x}_t)$; this follows from monotonicity of F as well as concavity of F along positive directions; see, e.g., [Calinescu et al., 2011]. Moreover, by Young's inequality we can show that the inner product $\langle \mathbf{v}_t - \mathbf{x}^*, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle$ is lower bounded by $-(\beta_t/2)\|\mathbf{v}_t - \mathbf{x}^*\|^2 - (1/2\beta_t)\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2$ for any $\beta_t > 0$. By applying these substitutions into (19) we obtain

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \frac{1}{T} (F(\mathbf{x}^*) - F(\mathbf{x}_t)) - \frac{LD^2}{2T^2} \\ &\quad - \frac{1}{2T} \left(\beta_t \|\mathbf{v}_t - \mathbf{x}^*\|^2 + \frac{\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2}{\beta_t} \right). \end{aligned} \quad (20)$$

Replace $\|\mathbf{v}_t - \mathbf{x}^*\|^2$ by its upper bound $4D^2$ and compute the expected value of (20) to write

$$\begin{aligned} \mathbb{E}[F(\mathbf{x}_{t+1})] &\geq \mathbb{E}[F(\mathbf{x}_t)] + \frac{1}{T} \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_t)] \\ &\quad - \frac{1}{2T} \left[4\beta_t D^2 + \frac{\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]}{\beta_t} \right] - \frac{LD^2}{2T^2}. \end{aligned} \quad (21)$$

Substitute $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ by its upper bound $Q/(t+9)^{2/3}$ according to the result in (15). Further, set $\beta_t = (Q^{1/2})/(2D(t+9)^{1/3})$ and regroup the resulted expression to obtain

$$\begin{aligned} \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_{t+1})] &\leq \left(1 - \frac{1}{T}\right) \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_t)] \\ &\quad + \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \frac{LD^2}{2T^2}. \end{aligned} \quad (22)$$

By applying the inequality in (22) recursively for $t =$

$0, \dots, T-1$ we obtain

$$\begin{aligned} \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_T)] &\leq \left(1 - \frac{1}{T}\right)^T (F(\mathbf{x}^*) - F(\mathbf{x}_0)) \\ &+ \sum_{t=0}^{T-1} \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \sum_{t=0}^{T-1} \frac{LD^2}{2T^2}. \end{aligned} \quad (23)$$

Simplifying the terms on the right hand side (23) leads to the expression

$$\begin{aligned} \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_T)] &\leq \frac{1}{e}(F(\mathbf{x}^*) - F(\mathbf{x}_0)) + \frac{2DQ^{1/2}}{T^{1/3}} + \frac{LD^2}{2T}. \end{aligned} \quad (24)$$

Here, we use the fact that $F(\mathbf{x}_0) \geq 0$, and hence the expression in (24) can be simplified to

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - 1/e)F(\mathbf{x}^*) - \frac{2DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (25)$$

and the claim in (16) follows. \blacksquare

The result in Theorem 1 shows that the sequence of iterates generated by SCG, which only has access to a noisy unbiased estimate of the gradient at each iteration, is able to achieve the optimal approximation bound $(1 - 1/e)$, while the error term vanishes at a sublinear rate of $\mathcal{O}(T^{-1/3})$.

6 Discrete Submodular Maximization

According to the results in Section 5, the SCG method achieves in expectation a $(1 - 1/e)$ -optimal solution for Problem (1). The focus of this section is on extending this result into the discrete domain and showing that SCG can be applied for maximizing a stochastic submodular *set* function f , namely Problem (2), through the multilinear extension of the function f . To be more precise, in lieu of solving the problem in (2) one can solve the continuous optimization problem

$$\max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}), \quad (26)$$

where F is the multilinear extension of the function f defined as

$$F(\mathbf{x}) = \sum_{S \subset V} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j), \quad (27)$$

and the convex set $\mathcal{C} = \text{conv}\{1_I : I \in \mathcal{I}\}$ is the matroid polytope [Calinescu et al., 2011]. Note that in (27), x_i denotes the i -th element of the vector \mathbf{x} .

Indeed, the continuous greedy algorithm is able to solve the program in (26); however, each iteration of the method is computationally costly due to gradient $\nabla F(\mathbf{x})$ evaluations. Instead, Badanidiyuru and

Vondrák [2014] and Chekuri et al. [2015] suggested approximating the gradient using a sufficient number of samples from f . This mechanism still requires access to the set function f multiple times at each iteration, and hence is not feasible for solving Problem (2). The idea is then to use a stochastic (unbiased) estimate for the gradient ∇F . In Appendix 9.3, we provide a method to compute an unbiased estimate of the gradient using n samples from $\hat{f}(S_i, \mathbf{z})$, where $\mathbf{z} \sim P$ and S_i 's, $i = 1, \dots, n$, are carefully chosen sets. Indeed, the stochastic gradient ascent method proposed in [Hassani et al., 2017] can be used to solve the multilinear extension problem in (26) using unbiased estimates of the gradient at each iteration. However, the stochastic gradient ascent method fails to achieve the optimal $(1 - 1/e)$ approximation. Further, the work in [Karimi et al., 2017] achieves a $(1 - 1/e)$ approximation solution only when each $\hat{f}(\cdot, \mathbf{z})$ is a coverage function. Here, we show that SCG achieves the first $(1 - 1/e)$ tight approximation guarantee for the discrete stochastic submodular Problem (2). More precisely, we show that SCG finds a solution for (26), with an expected function value that is at least $(1 - 1/e)\text{OPT} - \epsilon$, in $\mathcal{O}(1/\epsilon^3)$ iterations. To do so, we first show in the following lemma that the difference between any coordinates of gradients of two consecutive iterates generated by SCG, i.e., $\nabla_j F(\mathbf{x}_{t+1}) - \nabla_j F(\mathbf{x}_t)$ for $j \in \{1, \dots, n\}$, is bounded by $\|\mathbf{x}_{t+1} - \mathbf{x}_t\|$ multiplied by a factor which is independent of the problem dimension n .

Lemma 3. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1 with iterates \mathbf{x}_t , and recall the definition of the multilinear extension function F in (27). If we define r as the rank of the matroid \mathcal{I} and $m_f \triangleq \max_{i \in \{1, \dots, n\}} f(i)$, then*

$$|\nabla_j F(\mathbf{x}_{t+1}) - \nabla_j F(\mathbf{x}_t)| \leq m_f \sqrt{r} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|, \quad (28)$$

holds for $j = 1, \dots, n$.

Proof. See Section 9.4. \blacksquare

The result in Lemma 3 states that in an *ascent direction* of SCG, the gradient is $m_f \sqrt{r}$ -Lipschitz continuous. Here, m_f is the maximum marginal value of the function f and r is the rank for the matroid. Using the result of Lemma 3 and a coordinate-wise analysis, the bounds in Theorem 1 can be improved and specified for the case of multilinear extension maximization problem as we show in the following theorem.

Theorem 2. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1. Recall the definition of the multilinear extension function F in (27) and the definitions of r and m_f in Lemma 3. Further, set the averaging parameter as $\rho_t = 4/(t+8)^{2/3}$. If Assumptions 1 and 3 hold, then the iterate \mathbf{x}_T generated*

by SCG satisfies the inequality

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - 1/e)OPT - \frac{2DK}{T^{1/3}} - \frac{m_f \sqrt{r} D^2}{2T}, \quad (29)$$

where $K := \max\{\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|9^{1/3}, 4\sigma + \sqrt{3r}m_f D\}$.

Proof. The proof is similar to the proof of Theorem 1. The main difference is to write the analysis coordinate-wise and replace L by $m_f \sqrt{r}$, as shown in Lemma 3. For more details, check Section 9.5 in the supplementary material. ■

The result in Theorem 2 indicates that the sequence of iterates generated by SCG achieves a $(1 - 1/e)OPT - \epsilon$ approximation guarantee. Note that the constants on the right hand side of (29) are independent of n , except K which depends on σ . It can be shown that, in the worst case, the variance σ depends on the size of the ground set n and the variance of the stochastic functions $\tilde{f}(\cdot, \mathbf{z})$.

Let us now explain how the variance of the stochastic gradients of F relates to the variance of the marginal values of f . Consider a generic submodular set function g and its multilinear extension G . It can be shown

$$\nabla_j G(\mathbf{x}) = G(\mathbf{x}; x_j = 1) - G(\mathbf{x}; x_j = 0). \quad (30)$$

Hence, from submodularity we have $\nabla_j G(\mathbf{x}) \leq g(\{j\})$. Using this simple fact we can deduce that

$$\mathbb{E} \left[\|\nabla \tilde{F}(\mathbf{x}, \mathbf{z}) - \nabla F(\mathbf{x})\|^2 \right] \leq n \max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]. \quad (31)$$

Therefore, the constant σ can be upper bounded by

$$\sigma \leq \sqrt{n} \max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]^{1/2}. \quad (32)$$

As a result, we have the following guarantee for SCG in the case of multilinear functions.

Corollary 1. *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 1. Suppose the conditions in Theorem 2 are satisfied. Then, the sequence of iterates generated by SCG achieves a $(1 - 1/e)OPT - \epsilon$ solution after $\mathcal{O}(n^{3/2}/\epsilon^3)$ iterations.*

Proof. According to the result in Theorem 2, SCG reaches a $(1 - 1/e)OPT - \mathcal{O}(n^{1/2}/T^{1/3})$ solution after T iterations. Therefore, to achieve a $((1 - 1/e)OPT - \epsilon)$ approximation, $\mathcal{O}(n^{3/2}/\epsilon^3)$ iterations are required. ■

7 Numerical Experiments

In our experiments, we consider a movie recommendation application [Stan et al., 2017] consisting of N users and n movies. Each user i has a user-specific utility

function $f(\cdot, i)$ for evaluating sets of movies. The goal is to find a set of k movies such that in expectation over users' preferences it provides the highest utility, i.e., $\max_{|S| \leq k} f(S)$, where $f(S) \doteq \mathbb{E}_{i \sim P}[f(S, i)]$. This is an instance of the (discrete) stochastic submodular maximization problem defined in (2). For simplicity, we assume f has the form of an empirical objective function, i.e. $f(S) = \frac{1}{N} \sum_{i=1}^N f(S, i)$. In other words, the distribution P is assumed to be uniform on the integers between 1 and N . The continuous counterpart of this problem is to consider the the multilinear extension $F(\cdot, i)$ of any function $f(\cdot, i)$ and solve the problem in the continuous domain as follows. Let $F(\mathbf{x}) = \mathbb{E}_{i \sim \mathcal{D}}[F(\mathbf{x}, i)]$ for $x \in [0, 1]^n$ and define the constraint set $\mathcal{C} = \{\mathbf{x} \in [0, 1]^n : \sum_{i=1}^n x_i \leq k\}$. The discrete and continuous optimization formulations lead to the same optimal value [Calinescu et al., 2011]:

$$\max_{S: |S| \leq k} f(S) = \max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}). \quad (33)$$

Therefore, by running SCG we can find a solution in the continuous domain that is at least $1 - 1/e$ approximation to the optimal value. By rounding that fractional solution (for instance via randomized Pipage rounding [Calinescu et al., 2011]) we obtain a set whose utility is at least $1 - 1/e$ of the optimum solution set of size k . We note that randomized Pipage rounding does not need access to the value of f . We also remark that each iteration of SCG can be done very efficiently in $\mathcal{O}(n)$ time (the argmax step reduces to finding the largest k elements of a vector of length n). Therefore, such approach easily scales to big data scenarios where the size of the data set N (e.g. number of users) or the number of items n (e.g. number of movies) are large.

In our experiments, we use the following baselines:

- (i) Stochastic Continuous Greedy (SCG) with $\rho_t = \frac{1}{2}t^{-2/3}$ and mini-batch size B . The details for computing an unbiased estimator for the gradient ∇F are given in Section 9.3 in the supplementary material.
- (ii) Stochastic Gradient Ascent (SGA) of [Hassani et al., 2017]: with stepsize $\mu_t = c/\sqrt{t}$ and mini-batch size B .
- (iii) Frank-Wolfe (FW) variant of [Bian et al., Calinescu et al., 2011]: with parameter T for the total number of iterations and batch size B (we further let $\alpha = 1, \delta = 0$, see Algorithm 1 in [Bian et al.] or the continuous greedy method of [Calinescu et al., 2011] for more details).
- (iv) Batch-mode Greedy (Greedy): by running the vanilla greedy algorithm (in the discrete domain) in the following way. At each round of the algorithm (for selecting a new element), B random

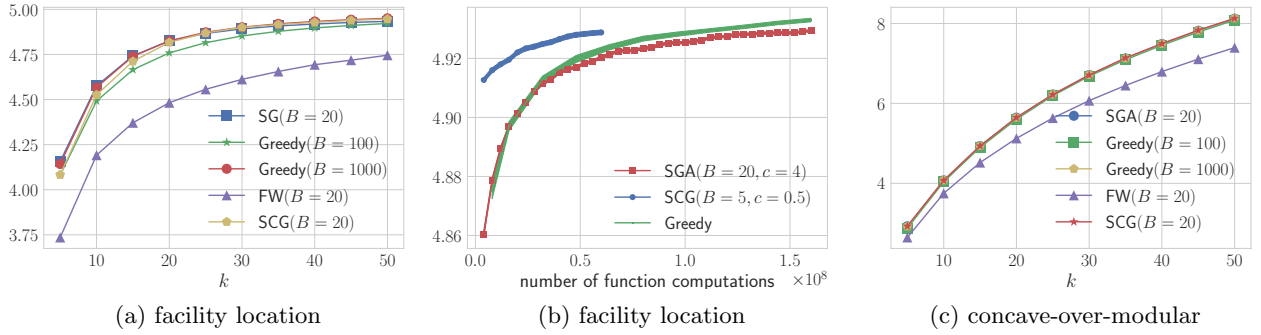


Figure 1: Comparison of the performances of SG, Greedy, FW, and SCG in a movie recommendation application. Fig. 1a illustrates the performance of the algorithms in terms of the facility-location objective value w.r.t. the cardinality constraint size k after $T = 2000$ iterations. Fig. 1b compares the considered methods in terms of runtime (for a fixed $k = 40$) by illustrating the facility location objective function value vs. the number of (simple) function evaluations. Fig. 1c demonstrates the concave-over-modular objective function value vs. the size of the cardinality constraint k after running the algorithms for $T = 2000$ iterations.

users are picked and the function f is estimated by the average over of the B selected users.

To run the experiments we use the MovieLens data set. It consists of 1 million ratings (from 1 to 5) by $N = 6041$ users for $n = 4000$ movies. Let $r_{i,j}$ denote the rating of user i for movie j (if such a rating does not exist we assign $r_{i,j}$ to 0). In our experiments, we consider two well motivated objective functions. The first one is called “facility location” where the valuation function by user i is defined as $f(S, i) = \max_{j \in S} r_{i,j}$. In words, the way user i evaluates a set S is by picking the highest rated movie in S . Thus, the objective function is

$$f_{\text{fac}}(S) = \frac{1}{N} \sum_{i=1}^N \max_{j \in S} r_{i,j}. \quad (34)$$

In our second experiment, we consider a different user-specific valuation function which is a concave function composed with a modular function, i.e., $f(S, i) = (\sum_{j \in S} r_{i,j})^{1/2}$. Again, by considering the uniform distribution over the set of users, we obtain

$$f_{\text{con}}(S) = \frac{1}{N} \sum_{i=1}^N \left(\sum_{j \in S} r_{i,j} \right)^{1/2}. \quad (35)$$

Figure 1 depicts the performance of different algorithms for the two proposed objective functions. As Figures 1a and 1c show, the FW algorithm needs a higher mini-batch size to be comparable to SCG. Note that a smaller batch size leads to less computational effort (under the same value for B, T , the computational complexity of FW, SGA, SCG is almost the same). Figure 1b shows the performance of the algorithms

with respect to the number of times the (simple) functions (i.e., $f(\cdot, i)$ ’s) are evaluated. Note that the total number of (simple) function evaluations for SGA and SGA is nBT , where T is the number of iterations. Also, for Greedy the total number of evaluations is nkB . This further shows that SCG has a better computational complexity requirement w.r.t. SGA as well as the Greedy algorithm (in the discrete domain).

8 Conclusion

In this paper, we provided the first tight approximation guarantee for maximizing a stochastic monotone DR-submodular function subject to a general convex body constraint. We developed **Stochastic Continuous Greedy** that achieves a $[(1-1/e)\text{OPT} - \epsilon]$ guarantee (in expectation) with $\mathcal{O}(1/\epsilon^3)$ stochastic gradient computations. We also demonstrated that our continuous algorithm can be used to provide the first $(1-1/e)$ tight approximation guarantee for maximizing a *monotone but stochastic* submodular set function subject to a general matroid constraint. We believe that our results provide an important step towards unifying discrete and continuous submodular optimization in stochastic settings.

Acknowledgements

This work was done while A. Mokhtari was visiting the Simons Institute for the Theory of Computing, and his work was partially supported by the DIMACS/Simons Collaboration on Bridging Continuous and Discrete Optimization through NSF grant #CCF-1740425. The work of A. Karbasi was supported by DARPA Young Faculty Award (D16AP00046) and AFOSR YIP (FA9550-18-1-0160).

References

- A. A. Abbasi and M. Younis. A survey on clustering algorithms for wireless sensor networks. *Computer communications*, 30(14):2826–2841, 2007.
- A. Anandkumar and R. Ge. Efficient approaches for escaping higher order saddle points in non-convex optimization. In *Proceedings of the 29th Conference on Learning Theory*, pages 81–102, 2016.
- A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Tegarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15: 2773–2832, 2014.
- S. Arora, R. Ge, R. Kannan, and A. Moitra. Computing a nonnegative matrix factorization - provably. *SIAM J. Comput.*, 45(4):1582–1611, 2016.
- F. Bach. Submodular functions: from discrete to continuous domains. *arXiv preprint arXiv:1511.00394*, 2015.
- A. Badanidiyuru and J. Vondrák. Fast algorithms for maximizing submodular functions. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1497–1514, 2014.
- Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. In *Advances in Neural Information Processing Systems 19*, pages 153–160, 2007.
- A. A. Bian, B. Mirzasoleiman, J. M. Buhmann, and A. Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*.
- N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1433–1452, 2014.
- N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. *SIAM J. Comput.*, 44(5): 1384–1402, 2015.
- G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- D. Chakrabarty, Y. T. Lee, A. Sidford, and S. C. Wong. Subquadratic submodular function minimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1220–1231, 2017.
- C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- C. Chekuri, T. Jayram, and J. Vondrák. On multiplicative weight updates for concave and submodular function maximization. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 201–210. ACM, 2015.
- N. R. Devanur and K. Jain. Online matching with concave returns. In *Proceedings of the 44th Symposium on Theory of Computing Conference*, pages 137–144, 2012.
- R. Eghbali and M. Fazel. Designing smoothing functions for improved worst-case competitive ratio in online optimization. In *Advances in Neural Information Processing Systems 29*, pages 3279–3287, 2016.
- A. Ene, H. L. Nguyen, and L. A. Végh. Decomposable submodular function minimization: Discrete and continuous. In *Advances in Neural Information Processing Systems 30*, pages 2874–2884, 2017.
- U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- U. Feige, V. S. Mirrokni, and J. Vondrak. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- M. Feldman, J. Naor, and R. Schwartz. A unified continuous greedy algorithm for submodular maximization. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 570–579, 2011.
- M. Feldman, C. Harshaw, and A. Karbasi. Greed is good: Near-optimal submodular maximization via greedy optimization. In *Proceedings of the 30th Conference on Learning Theory*, pages 758–784, 2017.
- S. Fujishige. *Submodular functions and optimization*, volume 58. Elsevier, 2005.
- R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Proceedings of The 28th Conference on Learning Theory*, pages 797–842, 2015.
- R. Ge, J. D. Lee, and T. Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems 29*, pages 2973–2981, 2016.
- S. O. Gharan and J. Vondrák. Submodular maximization by simulated annealing. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1098–1116, 2011.
- D. Golovin, A. Krause, and M. Streeter. Online submodular maximization under a matroid constraint with application to learning assignments. *arXiv preprint arXiv:1407.1082*, 2014.
- S. H. Hassani, M. Soltanolkotabi, and A. Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems 30*, pages 5843–5853, 2017.
- E. Hazan and H. Luo. Variance-reduced and projection-free stochastic optimization. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1263–1271, 2016.
- E. Hazan, K. Y. Levy, and S. Shalev-Shwartz. On graduated optimization for stochastic non-convex problems. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1833–1841, 2016.
- C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1724–1732, 2017.
- M. R. Karimi, M. Lucic, S. H. Hassani, and A. Krause. Stochastic submodular maximization: The case of coverage functions. In *Advances in Neural Information Processing Systems 30*, pages 6856–6866, 2017.
- L. Lovász. Submodular functions and convexity. In *Mathematical Programming The State of the Art*, pages 235–257. Springer, 1983.

- A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22, 2007.
- S. Mei, Y. Bai, and A. Montanari. The landscape of empirical risk for non-convex losses. *arXiv preprint arXiv:1607.06534*, 2016.
- A. Mokhtari, A. Koppel, G. Scutari, and A. Ribeiro. Large-scale nonconvex stochastic optimization by doubly stochastic successive convex approximation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4701–4705, 2017.
- K. G. Murty and S. N. Kabadi. Some np-complete problems in quadratic and nonlinear programming. *Mathematical programming*, 39(2):117–129, 1987.
- G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, 1978.
- P. Netrapalli, U. N. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain. Non-convex robust PCA. In *Advances in Neural Information Processing Systems 27*, pages 1107–1115, 2014.
- S. Paternain, A. Mokhtari, and A. Ribeiro. A second order method for nonconvex optimization. *arXiv preprint arXiv:1707.08028*, 2017.
- S. J. Reddi, S. Sra, B. Póczos, and A. Smola. Stochastic Frank-Wolfe methods for nonconvex optimization. In *54th Annual Allerton Conference on Communication, Control, and Computing*, pages 1244–1251. IEEE, 2016.
- A. Ruszczyński. Feasible direction methods for stochastic programming problems. *Mathematical Programming*, 19(1):220–229, 1980.
- A. Ruszczyński. A merit function approach to the subgradient method with averaging. *Optimisation Methods and Software*, 23(1):161–172, 2008.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems 25*, pages 2951–2959, 2012.
- T. Soma, N. Kakimura, K. Inaba, and K. Kawarabayashi. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *Proceedings of the 31th International Conference on Machine Learning*, pages 351–359, 2014.
- M. Staib and S. Jegelka. Robust budget allocation via continuous submodular functions. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3230–3240, 2017.
- S. Stan, M. Zadimoghaddam, A. Krause, and A. Karbasi. Probabilistic submodular maximization in sub-linear time. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3241–3250, 2017.
- M. Sviridenko, J. Vondrák, and J. Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1134–1148, 2015.
- J. Vondrák. Submodularity in combinatorial optimization. 2007.
- J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 67–74. ACM, 2008.
- J. Vondrák. Submodularity and curvature : The optimal algorithm (combinatorial optimization and discrete algorithms). *RIMS Kokyuroku Bessatsu*, B23:253–266, 2010.
- L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.
- Y. Yang, G. Scutari, D. P. Palomar, and M. Pesavento. A parallel decomposition method for nonconvex stochastic multi-agent optimization problems. *IEEE Transactions on Signal Processing*, 64(11):2949–2964, 2016.