

Stochastic Conditional Gradient Methods: From Convex Minimization to Submodular Maximization

Aryan Mokhtari

*Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX 78712*

MOKHTARI@AUSTIN.UTEXAS.EDU

Hamed Hassani

*Department of Electrical and Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104*

HASSANI@SEAS.UPENN.EDU

Amin Karbasi

*Department of Electrical Engineering and Computer Science
Yale University
New Haven, CT 06520*

AMIN.KARBASI@YALE.EDU

Editor: Francis Bach

Abstract

This paper considers stochastic optimization problems for a large class of objective functions, including convex and continuous submodular. Stochastic proximal gradient methods have been widely used to solve such problems; however, their applicability remains limited when the problem dimension is large and the projection onto a convex set is computationally costly. Instead, stochastic conditional gradient algorithms are proposed as an alternative solution which rely on (i) Approximating gradients via a simple averaging technique requiring a single stochastic gradient evaluation per iteration; (ii) Solving a linear program to compute the descent/ascent direction. The gradient averaging technique reduces the noise of gradient approximations as time progresses, and replacing projection step in proximal methods by a linear program lowers the computational complexity of each iteration. We show that under convexity and smoothness assumptions, our proposed stochastic conditional gradient method converges to the optimal objective function value at a sublinear rate of $\mathcal{O}(1/t^{1/3})$. Further, for a monotone and continuous DR-submodular function and subject to a *general* convex body constraint, we prove that our proposed method 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. Additionally, we achieve $((1/e)\text{OPT} - \epsilon)$ guarantee after operating on $\mathcal{O}(1/\epsilon^3)$ stochastic gradients for the case that the objective function is continuous DR-submodular but non-monotone and the constraint set is a down-closed convex body. 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 and $(1/e)$ approximation guarantee for the non-monotone case.

Keywords: Stochastic optimization, conditional gradient methods, convex minimization, submodular maximization, gradient averaging, Frank-Wolfe algorithm, greedy algorithm

1. Introduction

Stochastic optimization arises in various applications including wireless communications (Ribeiro, 2010), learning theory (Vapnik, 2013), machine learning (Bottou, 2010), adaptive filters (Haykin, 2008), portfolio selection (Shapiro et al., 2009) to name a few. In this class of problems the goal is to optimize an objective function defined as an expectation over a set of random functions subject to a general convex constraint. In particular, consider an optimization variable $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$ and a random variable $\mathbf{Z} \in \mathcal{Z}$ that together determine the choice of a stochastic function $\tilde{F} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$. The goal is to solve the program

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

where \mathcal{C} is a convex compact set, \mathbf{z} is the realization of the random variable \mathbf{Z} drawn from a distribution P , and $F(\mathbf{x})$ is the expected value of the random functions $\tilde{F}(\mathbf{x}, \mathbf{z})$ with respect to the random variable \mathbf{Z} . In this paper, our focus is on the cases where the function F is either concave or continuous submodular. The first case which considers the problem of maximizing a concave function is equivalent to *stochastic convex minimization*, and the second case in which the goal is to maximize a continuous submodular function is called *stochastic continuous submodular maximization*. Note that the main challenge here is to solve Problem (1) without computing the objective function $F(\mathbf{x})$ or its gradient $\nabla F(\mathbf{x})$ explicitly, since we assume that either the probability distribution P is unknown or the cost of computing the expectation is prohibitive.

In this regime, stochastic algorithms are the method of choice since they operate on computationally cheap stochastic estimates of the objective function gradients. Stochastic variants of the proximal gradient method are perhaps the most popular algorithms to tackle this category of problems both in convex minimization and submodular maximization. However, implementation of proximal methods requires projection onto a convex set at each iteration to enforce feasibility, which could be computationally expensive or intractable in many settings. To avoid the cost of projection, recourse to conditional gradient methods arises as a natural alternative. Unlike proximal algorithms, conditional gradient methods do not suffer from computationally costly projection steps and only require solving linear programs which can be implemented at a significantly lower complexity in many practical applications.

In deterministic convex minimization, in which the exact gradient information is given, conditional gradient method, a.k.a., Frank-Wolfe algorithm (Frank and Wolfe, 1956; Jaggi, 2013), succeeds in lowering the computational complexity of proximal algorithms due to its projection free property. However, in the stochastic regime, where only an estimate of the gradient is available, stochastic conditional gradient methods either diverge or underperform the proximal gradient methods. In particular, in stochastic convex minimization it is known and proven that stochastic conditional gradient methods may not converge to the optimal solution without an increasing batch size, whereas stochastic proximal gradient methods converge to the optimal solution at the sublinear rate of $\mathcal{O}(1/\sqrt{t})$. Hence, the possibility of designing a convergent stochastic conditional gradient method with a small batch size remains unanswered.

In deterministic continuous submodular maximization where objective function is submodular, (neither convex nor concave), a variant of the condition gradient method called,

continuous (Calinescu et al., 2011; Bian et al., 2017b) obtains a $(1 - 1/e)$ approximation guarantee for monotone functions, in contrast the best-known result for proximal gradient methods is a $(1/2)$ approximation guarantee (Hassani et al., 2017). However, in stochastic submodular maximization setting, stochastic variants of the continuous greedy algorithm with a small batch size fail to achieve a constant factor approximation (Hassani et al., 2017), whereas stochastic proximal gradient method recovers the $(1/2)$ approximation obtained by proximal gradient method in deterministic settings (Hassani et al., 2017). It is unknown if one can design a stochastic conditional gradient method that obtains a constant approximation guarantee, ideally $(1 - 1/e)$, for Problem (1).

In this paper, we introduce a stochastic conditional gradient method for solving the generic stochastic optimization Problem (1). The proposed method lowers the noise of gradient approximations through a simple gradient averaging technique which only requires a single stochastic gradient computation per iteration, i.e., the batch size can be as small as 1. The proposed stochastic conditional gradient method improves the best-known convergence guarantees for stochastic conditional gradient methods in both convex minimization and submodular maximization settings. A summary of our theoretical contributions follows.

- (i) In *stochastic convex minimization*, we propose a stochastic variant of the Frank-Wolfe algorithm which converges to the optimal objective function value at the sublinear rate of $\mathcal{O}(1/t^{1/3})$. In other words, the proposed SFW algorithm achieves an ϵ -suboptimal objective function value after $\mathcal{O}(1/\epsilon^3)$ stochastic gradient evaluations.
- (ii) In *stochastic continuous submodular maximization*, we propose a stochastic conditional gradient method, which can be interpreted as a stochastic variant of the continuous greedy algorithm, that obtains the first tight $(1 - 1/e)$ approximation guarantee, when the function is monotone. For the non-monotone case, the proposed method obtains a $(1/e)$ approximation guarantee. Moreover, for the more general case of γ -weakly DR-submodular monotone maximization the proposed stochastic conditional gradient method achieves a $(1 - e^{-\gamma})$ -approximation guarantee.
- (iii) In *stochastic discrete submodular maximization*, the proposed stochastic conditional gradient method achieves $(1 - 1/e)$ and $(1/e)$ approximation guarantees for monotone and non-monotone settings, respectively, by maximizing the multilinear extension of the stochastic discrete objective function. Further, if the objective function is monotone and has curvature c the proposed stochastic conditional gradient method achieves a $(1 - e^{-c})/c$ approximation guarantee.

We begin the paper by studying the related work on stochastic methods in convex minimization and submodular maximization (Section 2). Then, we proceed by stating stochastic convex minimization problem (Section 3). We introduce a stochastic conditional gradient method which can be interpreted as a stochastic variant of Frank-Wolfe (FW) algorithm for solving stochastic convex minimization problems (Section 3.1). The proposed **Stochastic Frank-Wolfe** method (SFW) differs from the vanilla FW method in replacing gradients by their stochastic approximations evaluated based on averaging over previously observed stochastic gradients. We further analyze the convergence properties of the proposed SFW method (Section 3.2). In particular, we show that the averaging technique in SFW lowers

the noise of gradient approximation (Lemma 1) and consequently the sequence of objective function values generated by SFW converges to the optimal objective function value at a sublinear rate of $\mathcal{O}(1/t^{1/3})$ in expectation (Theorem 3). We complete this result by proving that the sequence of objective function values almost surely converges to the optimal objective function value (Theorem 4).

We then focus on the application of the proposed stochastic conditional gradient method in continuous submodular maximization (Section 4). After defining the notions of submodularity (Section 4.1), we introduce the **Stochastic Continuous Greedy** (SCG) algorithm for solving continuous submodular maximization problems (Section 4.2). The proposed SCG algorithm achieves the optimal $(1 - 1/e)$ -approximation when the expected objective function is DR-submodular, monotone, and smooth (Theorem 7). To be more precise, the expected objective value of the iterates generated by SCG in expectation is not smaller $(1 - 1/e)OPT - \epsilon$ after $\mathcal{O}(1/\epsilon^3)$ stochastic gradient evaluations. Moreover, for the case that the expected function is not DR-submodular but γ weakly DR-submodular, the SCG algorithm obtains an $(1 - e^{-\gamma})$ -approximation guarantee (Theorem 8). We further extend our results to non-monotone setting by introducing the **Non-monotone Stochastic Continuous Greedy** (NMSCG) method. We show that under the assumptions that the expected function is only DR-submodular and smooth NMSCG reaches a $(1/e)$ -approximation guarantee (Theorem 9).

The continuous multilinear extension of discrete submodular functions implies that the results for stochastic continuous DR-submodular maximization can be extended to stochastic discrete submodular maximization. We formalize this connection by introducing the stochastic discrete submodular maximization problem and defining its continuous multilinear extension (Section 5). By leveraging this connection, one can relax the discrete problem to a stochastic continuous submodular maximization, use SCG to solve the relaxed continuous problem within a $(1 - 1/e - \epsilon)$ approximation to the optimum value (Theorem 12), and use a proper rounding scheme (such as the contention resolution method (Chekuri et al., 2014)) to obtain a feasible set whose value is a $(1 - 1/e - \epsilon)$ approximation to the optimum set in expectation. In summary, we show that SCG achieves an $((1 - 1/e)OPT - \epsilon)$ approximation for a generic discrete monotone stochastic submodular maximization problem after $\mathcal{O}(n^{3/2}/\epsilon^3)$ iterations where n is the size of the ground set (Corollary 13). We further prove a $1/e$ approximation guarantee for NMSCG when we maximize a *non-monotone* stochastic submodular set function (Theorem 15). Moreover, if the expected set function has curvature $c \in [0, 1]$, SCG reaches an $(1/c)(1 - e^{-c})$ approximation guarantee (Theorem 17). We finally close the paper by concluding remarks (Section 7).

Notation. Lowercase boldface \mathbf{v} denotes a vector and uppercase boldface \mathbf{A} denotes 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 $\mathbf{A}_{i,j}$.

2. Related Work

In this section we overview the literature on conditional gradient methods in convex minimization as well as submodular maximization and compare our novel theoretical guarantees with the existing results.

Convex minimization. The problem of minimizing a stochastic convex function subject to a convex constraint has been tackled by many researchers and many approaches have been reported in the literature. Projected stochastic gradient descent (PSGD) and stochastic variants of Frank-Wolfe algorithm are among the most popular approaches. PSGD updates the iterates by descending through the negative direction of stochastic gradient with a proper stepsize and projecting the resulted point onto the feasible set (Robbins and Monro, 1951; Nemirovski and Yudin, 1978; Nemirovskii et al., 1983). It is known that in convex and smooth settings, the iterates of PSGD converge at a rate of $\mathcal{O}(1/\sqrt{t})$, where t is the number of iterations. Although stochastic gradient computation is inexpensive, the cost of projection step can be high (Fujishige and Isotani, 2011) or intractable (Collins et al., 2008). In such cases projection-free conditional gradient methods, a.k.a., Frank-Wolfe algorithm, are more practical (Frank and Wolfe, 1956; Jaggi, 2013). The online Frank-Wolfe algorithm proposed by Hazan and Kale (2012) requires $\mathcal{O}(1/\epsilon^4)$ stochastic gradient evaluations to reach ϵ suboptimal objective function value, i.e., $\mathbb{E}[f(\mathbf{x}) - OPT] \leq \epsilon$ under the assumption that the objective function is convex and has bounded gradients. The stochastic Frank-Wolfe studied by Hazan and Luo (2016) obtains an improved complexity of $\mathcal{O}(1/\epsilon^3)$ under the assumptions that the expected objective function is smooth (has Lipschitz gradients) and Lipschitz continuous (the gradients are bounded). More importantly, the stochastic Frank-Wolfe algorithm in (Hazan and Luo, 2016) requires an increasing batch size b as time progresses, i.e., $b = \mathcal{O}(t)$. In this paper, we propose a stochastic variant of conditional gradient method which achieves the complexity of $\mathcal{O}(1/\epsilon^3)$ under milder assumptions (only requires smoothness of the expected function) and operates with a fixed batch size, e.g., $b = 1$.

Submodular maximization. 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; Mirzasoleiman et al., 2016; 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 the objective function has curvature c (Vondrák, 2010). Moreover, the best-known result for the case that the objective function has curvature c is shown by (Sviridenko et al., 2015) which is a $((1 - c)/e)$ -approximation guarantee. The problem of maximizing submodular functions also has been studied for the non-monotone case, and constant approximation guarantees have been established (Feldman et al., 2011; Buchbinder et al., 2015; Ene and Nguyen, 2016; Buchbinder and Feldman, 2016).

Ref.	setting	assumptions	batch	rate	complexity
Jaggi (2013)	det.	smooth	—	$\mathcal{O}(1/t)$	—
Hazan and Kale (2012)	stoch.	smooth, bounded grad.	$\mathcal{O}(t)$	$\mathcal{O}(1/t^{1/2})$	$\mathcal{O}(1/\epsilon^4)$
Hazan and Luo (2016)	stoch.	smooth, bounded grad.	$\mathcal{O}(t^2)$	$\mathcal{O}(1/t)$	$\mathcal{O}(1/\epsilon^3)$
This paper	stoch.	smooth, bounded var.	$\mathcal{O}(1)$	$\mathcal{O}(1/t^{1/3})$	$\mathcal{O}(1/\epsilon^3)$

Table 1: Convergence guarantees of conditional gradient (FW) methods for convex minimization

Ref.	setting	function	const.	utility	complexity
Chekuri et al. (2015)	det.	mon.smooth sub.	poly.	$(1 - 1/e)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^2)$
Bian et al. (2017b)	det.	mon. DR-sub.	cvx-down	$(1 - 1/e)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon)$
Bian et al. (2017a)	det.	non-mon. DR-sub.	cvx-down	$(1/e)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon)$
Hassani et al. (2017)	det.	mon. DR-sub.	convex	$(1/2)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon)$
Hassani et al. (2017)	stoch.	mon. DR-sub.	convex	$(1/2)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^2)$
Hassani et al. (2017)	stoch.	mon. weak DR-sub.	convex	$\frac{\gamma^2}{1+\gamma^2}\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^2)$
This paper	stoch.	mon. DR-sub.	convex	$(1 - 1/e)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^3)$
This paper	stoch.	weak DR-sub.	convex	$(1 - e^{-\gamma})\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^3)$
This paper	stoch.	non-mon. DR-sub.	convex	$(1/e)\text{OPT} - \epsilon$	$\mathcal{O}(1/\epsilon^3)$

Table 2: Convergence guarantees for continuous DR-submodular function maximization

Continuous submodularity naturally arises in many learning applications such as robust budget allocation (Staib 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 a $(1 - 1/e - \epsilon)$ approximation guarantee after $\tilde{\mathcal{O}}(n/\epsilon^2)$ 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., 2017b). However, such results require exact computation of the gradients ∇F which is not feasible in Problem (14). 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 (14) 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 $\mathcal{O}(1/\epsilon^2)$ iterations. Although this work opens the door for maximizing stochastic continuous submodular functions using computa-

Ref.	setting	function	constraint	approx	method
Nemhauser and Wolsey (1981)	det.	mon. sub.	cardinality	$1 - 1/e$	disc. greedy
Nemhauser and Wolsey (1981)	det.	mon. sub.	matroid	$1/2$	disc. greedy
Calinescu et al. (2011)	det.	mon. sub.	matroid	$1 - 1/e$	con. greedy
Hassani et al. (2017)	stoch.	mon. sub.	matroid	$1/2$	SGA
This paper	stoch.	mon. sub.	matroid	$1 - 1/e$	SCG
This paper	stoch.	mon. sub.	matroid	$(1 - 1/e^c)/c$	SCG
This paper	stoch.	sub.	matroid	$1/e$	NMSCG

Table 3: Convergence guarantees for submodular set function maximization

tionally 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., 2017b), which holds for down-closed convex constraints, our result holds for any convex constraints. For non-monotone DR-submodular functions, we also propose the non-monotone stochastic continuous greedy (NMSCG) method that achieves a solution with function value at least $((1/e)\text{OPT} - \epsilon)$ after at most $O(1/\epsilon^3)$ iterations. Crucially, the feasible set in this case should be down-closed or otherwise no constant approximation guarantee is possible (Chekuri et al., 2014).

Our result also has important implications for the problem of 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 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 when the functions are monotone. 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 expectation¹. Using the same procedure we can also prove a $(1/e)$ approximation guarantee in expectation for the non-monotone case. Additionally, when the set function f is monotone and has a curvature $c < 1$ – check (41) for the definition of the curvature – we show that the approximation factor can be improved from $(1 - 1/e)$ to $((1 - 1/e^c)/c)$.

3. Stochastic Convex Minimization

Many problems in machine learning can be reduced to the minimization of a convex objective function defined as an expectation over a set of random functions. In this section, our focus

1. For the ease of presentation, and in the discrete setting, we only present our results for the matroid constraint. However, our stochastic continuous algorithms can provide constant factor approximations for any constrained submodular maximization setting where an efficient and loss-less rounding scheme exists. It includes, for example, knapsack constraints among many others.

is on developing a stochastic variant of the conditional gradient method, a.k.a., Frank-Wolfe, which can be applied to solve general stochastic convex minimization problems. To make the notation consistent with other works on stochastic convex minimization, instead of solving Problem (1) for the case that the objective function F is concave, we assume that F is convex and intend to minimize it subject to a convex set \mathcal{C} . Therefore, the goal is to solve the program

$$\min_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}) := \min_{\mathbf{x} \in \mathcal{C}} \mathbb{E}_{\mathbf{z} \sim P} [\tilde{F}(\mathbf{x}, \mathbf{z})]. \quad (2)$$

A canonical subset of problems having this form is support vector machines, least mean squares, and logistic regression.

In this section we assume that only the expected (average) function F is convex and smooth, and the stochastic functions \tilde{F} may not be convex nor smooth. Since the expected function F is convex, descent methods can be used to solve the program in (2). However, computing the gradients (or Hessians) of the function F precisely requires access to the distribution P , which may not be feasible in many applications. To be more specific, we are interested in settings where the distribution P is either unknown or evaluation of the expected value is computationally prohibitive. In this regime, stochastic gradient descent methods, which operate on stochastic approximations of the gradients, are the mostly used alternatives. In the following section, we aim to develop a stochastic variant of the Frank-Wolfe method which converges to an optimal solution of (2), while it only requires access to a single stochastic gradient $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ at each iteration.

3.1. Stochastic Frank-Wolfe Method

In this section, we introduce a stochastic variant of the Frank-Wolfe method to solve Problem (2). Assume that at each iteration we have access to the stochastic gradient $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ which is an unbiased estimate of the gradient $\nabla F(\mathbf{x})$. It is known that a naive stochastic implementation of Frank-Wolfe (replacing gradient $\nabla F(\mathbf{x})$ by $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$) might diverge, due to non-vanishing variance of gradient approximations. To resolve this issue, we introduce a stochastic version of the Frank-Wolfe algorithm which reduces the noise of gradient approximations via a common averaging technique in stochastic optimization (Ruszczyński, 1980, 2008; Yang et al., 2016; Mokhtari et al., 2017).

By letting $t \in \mathbf{N}$ be a discrete time index and ρ_t a given stepsize which approaches zero as t grows, the proposed biased gradient estimation \mathbf{d}_t is defined by the recursion

$$\mathbf{d}_t = (1 - \rho_t)\mathbf{d}_{t-1} + \rho_t \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t), \quad (3)$$

where the initial vector is given as $\mathbf{d}_0 = \mathbf{0}$. We will show that the averaging technique in (3) reduces the noise of gradient approximation as time increases. More formally, the expected noise of the gradient estimation $\mathbb{E} [\|\mathbf{d}_t - \nabla F(\mathbf{x}_t)\|^2]$ approaches zero asymptotically. This property implies that the biased 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 descent direction \mathbf{v}_t of our proposed **Stochastic Frank-Wolfe** (SFW) method as the solution of the linear program

$$\mathbf{v}_t \in \operatorname{argmin}_{\mathbf{v} \in \mathcal{C}} \{\mathbf{d}_t^T \mathbf{v}\}. \quad (4)$$

Algorithm 1 Stochastic Frank-Wolfe (SFW)

Require: Stepsizes $\rho_t > 0$ and $\gamma_t > 0$. Initialize $\mathbf{d}_0 = \mathbf{0}$ and choose $\mathbf{x}_0 \in \mathcal{C}$

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

As in the traditional FW method, the updated variable \mathbf{x}_{t+1} is a convex combination of \mathbf{v}_t and the iterate \mathbf{x}_t

$$\mathbf{x}_{t+1} = (1 - \gamma_{t+1})\mathbf{x}_t + \gamma_{t+1}\mathbf{v}_t, \tag{5}$$

where γ_t is a proper positive stepsize. Note that each iteration of the proposed SFW method only requires a single stochastic gradient computation, unlike the methods in (Hazan and Luo, 2016; Reddi et al., 2016) which an require increasing number of stochastic gradient evaluations as the number of iterations t grows.

The proposed SFW is summarized in Algorithm 1. The core steps are Steps 2-5 which follow the updates in (3)-(5). The initial variable \mathbf{x}_0 can be any feasible vector in the convex set \mathcal{C} and the initial gradient estimate is set to be the null vector $\mathbf{d}_0 = \mathbf{0}$. The sequence of positive parameters ρ_t and γ_t should be diminishing at proper rates as we describe in the following convergence analysis section.

3.2. Convergence Analysis

In this section we study the convergence rate of the proposed SFW method for solving the constraint convex program in (2). To do so, we first assume that the following conditions hold.

Assumption 1 *The convex set \mathcal{C} is bounded with diameter D , i.e., for all $\mathbf{x}, \mathbf{y} \in \mathcal{C}$ we can write*

$$\|\mathbf{x} - \mathbf{y}\| \leq D. \tag{6}$$

Assumption 2 *The expected function F is convex. Moreover, its gradients ∇F are L -Lipschitz continuous over the set \mathcal{C} , i.e., for all $\mathbf{x}, \mathbf{y} \in \mathcal{C}$*

$$\|\nabla F(\mathbf{x}) - \nabla F(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|. \tag{7}$$

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 all random variables \mathbf{z} and vectors $\mathbf{x} \in \mathcal{C}$ we can write*

$$\mathbb{E} \left[\|\nabla \tilde{F}(\mathbf{x}, \mathbf{z}) - \nabla F(\mathbf{x})\|^2 \right] \leq \sigma^2. \tag{8}$$

Assumption 1 is standard in constrained convex optimization and is implied by the fact that the set \mathcal{C} is convex and compact. The condition in Assumption 2 ensures that the objective function F is smooth on the set \mathcal{C} . Note that here we only assume that the (average) function F has Lipschitz continuous gradients, and the stochastic gradients $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ may not be Lipschitz continuous. Finally, the required condition in Assumption 3 is customary

in stochastic optimization and guarantees that the variance of stochastic gradients $\nabla\tilde{F}(\mathbf{x}, \mathbf{z})$ is bounded by a finite constant $\sigma^2 < \infty$.

To study the convergence rate of SFW, we first derive an upper bound on the error of gradient approximation $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2$ in the following lemma.

Lemma 1 *Consider the proposed Stochastic Frank-Wolfe (SFW) method outlined in Algorithm 1. If the conditions in Assumptions 1-3 hold, the sequence of squared gradient errors $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2$ satisfies*

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \mid \mathcal{F}_t] \leq \left(1 - \frac{\rho_t}{2}\right) \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 + \rho_t^2 \sigma^2 + \frac{2L^2 D^2 \gamma_t^2}{\rho_t}, \quad (9)$$

where \mathcal{F}_t is a sigma-algebra measuring all sources of randomness up to step t .

Proof Check Appendix A. ■

The result in Lemma 1 shows that squared error of gradient approximation $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2$ decreases in expectation at each iteration by the factor $(1 - \rho_t/2)$ if the remaining terms on the right hand side of (9) are negligible relative to the term $(1 - \rho_t/2)\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2$. This condition can be satisfied, if the parameters ρ_t and γ_t are properly chosen. This observation verifies our intuition that the noise of the stochastic gradient approximation diminishes as the number of iterations increases.

In the following lemma, we derive an upper bound on the suboptimality $F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)$ which depends on the norm of gradient error $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|$.

Lemma 2 *Consider the proposed Stochastic Frank-Wolfe (SFW) method outlined in Algorithm 1. If the conditions in Assumptions 1-3 are satisfied, the suboptimality $F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)$ satisfies*

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*) \leq (1 - \gamma_{t+1})(F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \gamma_{t+1} D \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| + \frac{LD^2 \gamma_{t+1}^2}{2}. \quad (10)$$

Proof Check Appendix B. ■

Based on Lemma 2, the suboptimality $F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)$ approaches zero if the error of gradient approximation $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|$ converges to zero sufficiently fast and the last term $(LD^2 \gamma_{t+1}^2)/2$ is summable, i.e., $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$. In the special case of zero approximation error, which is equivalent to the case that we have access to the expected gradient $\nabla F(\mathbf{x}_t)$, by setting $\gamma_t = \mathcal{O}(1/t)$ it can be shown that the suboptimality $F(\mathbf{x}_t) - F(\mathbf{x}^*)$ converges to zero at the sublinear rate of $\mathcal{O}(1/t)$. Therefore, the result in Lemma 2 is consistent with the analysis of Frank-Wolfe method (Jaggi, 2013).

In the following theorem, by using the results in Lemmas 1 and 2, we establish an upper bound on the expected suboptimality $\mathbb{E}[F(\mathbf{x}_t) - F(\mathbf{x}^*)]$.

Theorem 3 *Consider the proposed Stochastic Frank-Wolfe (SFW) method outlined in Algorithm 1. Suppose the conditions in Assumptions 1-3 are satisfied. If we set $\gamma_t = 2/(t+8)$*

and $\rho_t = 4/(t+8)^{2/3}$, then the expected suboptimality $\mathbb{E}[F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)]$ is bounded above by

$$\mathbb{E}[F(\mathbf{x}_t) - F(\mathbf{x}^*)] \leq \frac{Q}{(t+9)^{1/3}}, \quad (11)$$

where the constant Q is given by

$$Q := \max \left\{ 9^{1/3}(F(\mathbf{x}_0) - F(\mathbf{x}^*)), \frac{LD^2}{2} + 2D \max \left\{ 3\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|, (16\sigma^2 + 2L^2D^2)^{1/2} \right\} \right\}. \quad (12)$$

Proof Check Appendix C. ■

The result in Theorem 3 indicates that the expected suboptimality $\mathbb{E}[F(\mathbf{x}_t) - F(\mathbf{x}^*)]$ of the iterates generated by the SFW method converges to zero at least at a sublinear rate of $\mathcal{O}(1/t^{1/3})$. In other words, it shows that to achieve the expected suboptimality $\mathbb{E}[F(\mathbf{x}_t) - F(\mathbf{x}^*)] \leq \epsilon$, the number of required stochastic gradients (sample gradients) to reach this accuracy is $\mathcal{O}(1/\epsilon^3)$.

To complete the convergence analysis of SFW we also prove that the sequence of the objective function values $F(\mathbf{x}_t)$ converges to the optimal value $F(\mathbf{x}^*)$ almost surely. This result is formalized in the following theorem.

Theorem 4 *Consider the proposed Stochastic Frank-Wolfe (SFW) method outlined in Algorithm 1. Suppose that the conditions in Assumptions 1-3 are satisfied. If we choose ρ_t and γ_t such that (i) $\sum_{t=0}^{\infty} \rho_t = \infty$, (ii) $\sum_{t=0}^{\infty} \rho_t^2 < \infty$, (iii) $\sum_{t=0}^{\infty} \gamma_t = \infty$, and (iv) $\sum_{t=0}^{\infty} (\gamma_t^2 / \rho_t) < \infty$, then the suboptimality $F(\mathbf{x}_t) - F(\mathbf{x}^*)$ converges to zero almost surely, i.e.,*

$$\lim_{t \rightarrow \infty} F(\mathbf{x}_t) - F(\mathbf{x}^*) \stackrel{\text{a.s.}}{=} 0. \quad (13)$$

Proof Check Appendix D. ■

Theorem 4 provides almost sure convergence of the sequence of objective function value $F(\mathbf{x}_t)$ to the optimal value $F(\mathbf{x}^*)$. In other words it shows that the sequence of the objective function values $F(\mathbf{x}_t)$ converges to $F(\mathbf{x}^*)$ with probability 1. Indeed, a valid set of choices for γ_t and ρ_t to satisfy the required conditions in Theorem 4 are $\gamma_t = \mathcal{O}(1/t)$ and $\rho_t = \mathcal{O}(1/t^{2/3})$.

4. Stochastic Continuous Submodular Maximization

In the previous section, we focused on the convex setting, but what if the objective function F is not convex? In this section, 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 that demonstrate diminishing returns, i.e., continuous DR-submodular functions,

$$\max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}) \doteq \max_{\mathbf{x} \in \mathcal{C}} \mathbb{E}_{\mathbf{z} \sim P} [\tilde{F}(\mathbf{x}, \mathbf{z})]. \quad (14)$$

As before, 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 $\mathcal{X} \in \mathbb{R}_+^n$ is a compact set. In this section, we assume that $F(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathcal{X}$.

Our goal is to maximize the expected value of the random functions $\tilde{F}(\mathbf{x}, \mathbf{z})$ over the convex body $\mathcal{C} \subseteq \mathcal{X}$. 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 (14). In Section 3 of (Hassani et al., 2017), the authors also showed that if we simply substitute gradients by stochastic gradients in the update of conditional gradient methods (a.k.a., Frank-Wolfe), such as continuous greedy (Vondrák, 2008) or its close variant (Bian et al., 2017b), the resulted method can perform arbitrarily poorly in stochastic continuous submodular maximization settings. Our goal in this section is to design a stable stochastic variant of conditional gradient method to solve Problem (14) up to a constant factor.

4.1. Preliminaries

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

$$F(\mathbf{x}) + F(\mathbf{y}) \geq F(\mathbf{x} \vee \mathbf{y}) + F(\mathbf{x} \wedge \mathbf{y}), \tag{15}$$

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). Further, a submodular function F is monotone (on the set \mathcal{X}) if

$$\mathbf{x} \leq \mathbf{y} \implies F(\mathbf{x}) \leq F(\mathbf{y}), \tag{16}$$

for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$. Note that $\mathbf{x} \leq \mathbf{y}$ in (16) 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

$$\mathbf{x} \leq \mathbf{y} \implies \nabla F(\mathbf{x}) \geq \nabla F(\mathbf{y}). \tag{17}$$

See (Eghbali and Fazel, 2016) for related definitions and check (Bilmes and Bai, 2017) for more examples of DR-submodular functions. When the function F is twice differentiable,

submodularity implies that all cross-second-derivatives are non-positive (Bach, 2015), i.e.,

$$\text{for all } i \neq j, \text{ for all } \mathbf{x} \in \mathcal{X}, \quad \frac{\partial^2 F(\mathbf{x})}{\partial x_i \partial x_j} \leq 0, \quad (18)$$

and DR-submodularity implies that all second-derivatives are non-positive (Bian et al., 2017b), i.e.,

$$\text{for all } i, j, \text{ for all } \mathbf{x} \in \mathcal{X}, \quad \frac{\partial^2 F(\mathbf{x})}{\partial x_i \partial x_j} \leq 0. \quad (19)$$

4.2. Stochastic Continuous Greedy

We proceed to introduce, **Stochastic Continuous Greedy** (SCG), which is a stochastic variant of the continuous greedy method (Vondrák, 2008) to solve Problem (14). 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 (Calinescu et al., 2011; Bian et al., 2017b) can be used in principle to solve Problem (14). Note that each update of continuous greedy requires computing 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 (Hassani et al., 2017). This limitation is due to the non-vanishing variance of gradient approximations. To resolve this issue, we use the gradient averaging technique in Section 3.1. As in SFW, we define the estimated gradient \mathbf{d}_t by the recursion

$$\mathbf{d}_t = (1 - \rho_t)\mathbf{d}_{t-1} + \rho_t \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t), \quad (20)$$

where ρ_t is a positive stepsize and the initial vector is defined as $\mathbf{d}_0 = \mathbf{0}$. We therefore define the ascent direction \mathbf{v}_t of our proposed SCG method as follows

$$\mathbf{v}_t \in \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}} \{\mathbf{d}_t^T \mathbf{v}\}, \quad (21)$$

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 (21), 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

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{T} \mathbf{v}_t. \quad (22)$$

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 would like to highlight that the convex body \mathcal{C} may not be down-closed or contain $\mathbf{0}$. Nonetheless, the final iterate \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 2. Note that the major difference between SFW in Algorithm 1 and SCG in Algorithm 2 is in Step 4 where the variable \mathbf{x}_{t+1} is computed. In SFW, \mathbf{x}_{t+1} is a convex combination of \mathbf{x}_t and \mathbf{v}_t , while in SCG \mathbf{x}_{t+1} is computed by moving from \mathbf{x}_t towards the direction \mathbf{v}_t with the stepsize $1/T$.

Algorithm 2 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 $\mathbf{v}_t \in \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}} \{\mathbf{d}_t^T \mathbf{v}\}$;
 - 4: Update the variable $\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{T} \mathbf{v}_t$;
 - 5: **end for**
-

Remark 5 *Note that the update schemes for the convex case and submodular case are slightly different. In the convex case, we follow the update in (5), while in the submodular setting, we use the update in (22) to compute the new iterate. This modification is essential to achieve a $(1 - 1/e)$ -approximation guarantee for the monotone continuous submodular setting. To better highlight this point, note that if we use an update rule of the form (5) for the submodular setting it could lead to a $1/2$ -approximation. In other words, by starting from a feasible point and using the update rule in (5), SCG will eventually converge to a stationary point of the problem, when \mathbf{v}_t is computed according to the conditional gradient update. Converging to a stationary point in the convex setting is indeed desirable as any first-order stationary point in the convex setting is an optimal solution. However, in the continuous submodular setting, converging to a stationary point might lead to a $1/2$ -approximation guarantee. In fact, in (Hassani et al., 2017) it has been shown that for continuous submodular functions the value of objective function for a local maximum could be at most $1/2$ of the optimal objective function value. We refer the reader to (Hassani et al., 2017) for the example that shows the existence of such local maximum. By starting from the origin $\mathbf{0}$ (which is not a stationary point) and following the update rule in (22), SCG finds a path in the feasible set that without converging to a stationary point leads to an approximation factor of $1 - 1/e$.*

We proceed to study the convergence properties of our proposed SCG method for solving Problem (14). To do so, we first assume that the following conditions hold.

Assumption 4 *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*

$$\|\mathbf{x}\| \leq D. \tag{23}$$

Assumption 5 *The function F is DR-submodular and monotone. Further, its gradients are L -Lipschitz continuous over the set \mathcal{X} , i.e., for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$*

$$\|\nabla F(\mathbf{x}) - \nabla F(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|. \tag{24}$$

Assumption 6 *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*

$$\mathbb{E} \left[\|\nabla \tilde{F}(\mathbf{x}, \mathbf{z}) - \nabla F(\mathbf{x})\|^2 \right] \leq \sigma^2, \tag{25}$$

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

Due to the initialization step of SCG (i.e., 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 4. The condition in Assumption 5 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 6 guarantees that the variance of stochastic gradients $\nabla\tilde{F}(\mathbf{x}, \mathbf{z})$ is bounded by a finite constant $\sigma^2 < \infty$. Note that Assumptions 5-6 are stronger than Assumptions 2-3 since they ensure smoothness and bounded gradients for all points $\mathbf{x} \in \mathcal{X}$ and not only for the feasible points $\mathbf{x} \in \mathcal{C}$.

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

Lemma 6 *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 2. If Assumptions 4-6 are satisfied and $\rho_t = \frac{4}{(t+8)^{2/3}}$, then for $t = 0, \dots, T$ we have*

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \leq \frac{Q}{(t+9)^{2/3}}, \quad (26)$$

where $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 2L^2D^2\}$.

Proof Check Appendix E. ■

Let us now use the result of Lemma 6 to show that the sequence of iterates generated by SCG reaches a $(1 - 1/e)$ approximation guarantee for Problem (14).

Theorem 7 *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 2. If Assumptions 4-6 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*

$$\mathbb{E} [F(\mathbf{x}_T)] \geq (1 - 1/e)OPT - \frac{15DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (27)$$

where $OPT = \max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x})$ and $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 2L^2D^2\}$.

Proof Check Appendix F. ■

The result in Theorem 7 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 $\mathcal{O}(T^{-1/3})$.

4.3. Weak Submodularity

In this section, we extend our results to a more general case where the expected objective function F is weakly-submodular. A continuous function F is γ -weakly DR-submodular if

$$\gamma = \inf_{\mathbf{x}, \mathbf{y} \in \mathcal{X}, \mathbf{x} \leq \mathbf{y}} \inf_{i \in [n]} \frac{[\nabla F(\mathbf{x})]_i}{[\nabla F(\mathbf{y})]_i}, \quad (28)$$

Algorithm 3 Non-monotone Stochastic Continuous Greedy (NMSCG)

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 $\mathbf{v}_t \in \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}, \mathbf{v} \leq \bar{\mathbf{u}} - \mathbf{x}_t} \{\mathbf{d}_t^T \mathbf{v}\}$;
 - 4: Update the variable $\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{T} \mathbf{v}_t$;
 - 5: **end for**
-

where $[\mathbf{a}]_i$ denotes the i -th element of vector \mathbf{a} . In the following theorem, we prove that the proposed SCG method achieves a $(1 - e^{-\gamma})$ approximation guarantee when the expected function F is monotone and weakly DR-submodular with parameter γ .

Theorem 8 Consider *Stochastic Continuous Greedy (SCG)* outlined in Algorithm 2. If Assumptions 4-6 are satisfied and the function F is γ -weakly DR-submodular, then for $\rho_t = \frac{4}{(t+8)^{2/3}}$ the expected objective function value of the iterates generated by SCG satisfies the inequality

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - e^{-\gamma})OPT - \frac{15DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (29)$$

where $OPT = \max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x})$ and $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 2L^2D^2\}$.

Proof Check Appendix G. ■

4.4. Non-monotone Continuous Submodular Maximization

In this section, we aim to extend the results for the proposed **Stochastic Continuous Greedy** algorithm to maximize *non-monotone* stochastic DR-submodular. The problem formulation of interest is similar to Problem (14) except the facts that the objective function $F : \mathcal{X} \rightarrow \mathbb{R}_+$ may not be monotone and the set \mathcal{X} is a bounded box. To be more precise, we aim to solve the program

$$\max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}) \doteq \max_{\mathbf{x} \in \mathcal{C}} \mathbb{E}_{\mathbf{z} \sim P}[\tilde{F}(\mathbf{x}, \mathbf{z})], \quad (30)$$

where $F : \mathcal{X} \rightarrow \mathbb{R}$ is continuous DR-submodular, $\mathcal{X} = \prod_{i=1}^n \mathcal{X}_i$, each $\mathcal{X}_i = [u_i, \bar{u}_i]$ is a bounded interval, and the convex set \mathcal{C} is a subset of $\mathcal{X} = [\underline{\mathbf{u}}, \bar{\mathbf{u}}]$, where $\underline{\mathbf{u}} = [u_1; \dots; u_n]$ and $\bar{\mathbf{u}} = [\bar{u}_1; \dots; \bar{u}_n]$. In this section, we propose the first stochastic conditional gradient method for solving the stochastic non-monotone maximization problem in (30). In this section, we further assume that the convex set \mathcal{C} is down-closed and $\mathbf{0} \in \mathcal{C}$.

We introduce a variant of the **Stochastic Continuous Greedy** method that achieves a $(1/e)$ -approximation guarantee for Problem (30). The proposed **Non-monotone Stochastic Continuous Greedy (NMSCG)** method is inspired by the unified measured continuous greedy algorithm in (Feldman et al., 2011) and the Frank-Wolfe method in (Bian et al., 2017a) for non-monotone *deterministic* continuous submodular maximization. The steps

of NMSCG are summarized in Algorithm 3. The stochastic gradient update (Step 2) and the update of the variable (Step 4) are identical to the ones for SCG in Section 4.2. The main difference between NMSCG and SCG is in the computation of the ascent direction \mathbf{v}_t . In particular, in NMSCG the ascent direction vector \mathbf{v}_t is obtained by solving the linear program

$$\mathbf{v}_t \in \operatorname{argmax}_{\mathbf{v} \in \mathcal{C}, \mathbf{v} \leq \bar{\mathbf{u}} - \mathbf{x}_t} \{\mathbf{d}_t^T \mathbf{v}\}, \quad (31)$$

which differs from (21) by having the extra constraint $\mathbf{v} \leq \bar{\mathbf{u}} - \mathbf{x}_t$. This extra condition is added to ensure that the solution does not grow aggressively, since in non-monotone case dramatic growth of the solution may lead to poor performance. In NMSCG, the initial variable is $\mathbf{x}_0 = \mathbf{0}$, which is a legitimate initialization as we assume that the convex set \mathcal{C} is down-closed. In the following theorem, we establish a $1/e$ - guarantee for NMSCG.

Theorem 9 *Consider Non-monotone Stochastic Continuous Greedy (NMSCG) outlined in Algorithm 3 with the averaging parameter $\rho_t = 4/(t+8)^{2/3}$. If Assumptions 4 and 6 hold and the gradients ∇F are L -Lipschitz continuous and the convex set \mathcal{C} is down-closed, then the iterate \mathbf{x}_T generated by NMSCG satisfies the inequality*

$$\mathbb{E}[F(\mathbf{x}_T)] \geq e^{-1}F(\mathbf{x}^*) - \frac{15DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (32)$$

where $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 2L^2D^2\}$.

Proof See Section H. ■

The result in Theorem 9 states that the sequence of iterates generated by NMSCG achieves a $((1/e)OPT - \epsilon)$ approximation guarantee after $\mathcal{O}(1/\epsilon^3)$ stochastic gradient computations.

5. Stochastic Discrete Submodular Maximization

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 Lovas extension (Lovász, 1983) (for minimization). In fact, Problem (14) has a discrete counterpart, recently considered in (Hassani et al., 2017; Karimi et al., 2017):

$$\max_{S \in \mathcal{I}} f(S) \doteq \max_{S \in \mathcal{I}} \mathbb{E}_{\mathbf{z} \sim P}[\tilde{f}(S, \mathbf{z})], \quad (33)$$

where the function $f : 2^V \rightarrow \mathbb{R}_+$ is submodular, the function $\tilde{f} : 2^V \times \mathcal{Z} \rightarrow \mathbb{R}_+$ is an arbitrary mapping from the product of 2^V and the domain of the random variable \mathbf{Z} to \mathbb{R} , 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 (33) 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 closes the gap for maximizing the stochastic submodular set maximization, namely, Problem (33), by providing the first tight $(1 - 1/e)$ approximation guarantee for general monotone submodular set functions subject to a matroid constraint. For the non-monotone case, we obtain a $(1/e)$ approximation guarantee. We, further, show that a $((1 - e^{-c})/c)$ -approximation guarantee can be achieved when the function f has curvature c .

According to the results in Section 4, SCG achieves in expectation a $(1 - 1/e)$ -optimal solution for Problem (14) when the function F is monotone and DR-submodular, and NMSCG obtains $(1/e)$ -optimal solution for the non-monotone case. The focus of this section is on extending these results into the discrete domain and showing that SCG and NMSCG can be used to maximize a stochastic submodular *set* function f , namely Problem (33), through the multilinear extension of the function f . To be more precise, in lieu of solving the program in (33) one can solve the continuous optimization problem

$$\max_{\mathbf{x} \in \mathcal{C}} F(\mathbf{x}) = \max_{\mathbf{x} \in \mathcal{C}} \sum_{S \subset V} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j), \quad (34)$$

where F is the multilinear extension of the function f and the convex set $\mathcal{C} = \text{conv}\{1_I : I \in \mathcal{I}\}$ is the matroid polytope (Calinescu et al., 2011) which is down-closed (note that in (34), x_i denotes the i -th element of the vector \mathbf{x}). The fractional solution of the program (34) can then be rounded into a feasible discrete solution without any loss (in expectation) in objective value by methods such as randomized PIPAGE ROUNDING (Calinescu et al., 2011). Note that randomized PIPAGE ROUNDING requires $O(n)$ computational complexity (Karimi et al., 2017) for the uniform matroid and $O(n^2)$ complexity for general matroids.

Indeed, the conventional continuous greedy algorithm is able to solve the program in (34); however, each iteration of the method is computationally costly due to gradient $\nabla F(\mathbf{x})$ evaluations. Instead, Feldman et al. (2011) and Calinescu et al. (2011) 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 efficient for solving Problem (33). The idea is then to use a stochastic (unbiased) estimate for the gradient ∇F . In the following remark, we provide a method to compute an unbiased estimate of the gradient using n samples from $\tilde{f}(S_i, \mathbf{z})$, where $\mathbf{z} \sim P$ and S_i 's, $i = 1, \dots, n$, are carefully chosen sets.

Remark 10 (*Constructing an Unbiased Estimator of the Gradient in Multilinear Extensions*) Recall that $f(S) = \mathbb{E}_{\mathbf{z} \sim P}[\tilde{f}(S, \mathbf{z})]$. In terms of the multilinear extensions, we obtain $F(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim P}[\tilde{F}(\mathbf{x}, \mathbf{z})]$, where F and \tilde{F} denote the multilinear extension of f and \tilde{f} , respectively. So $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$ is an unbiased estimator of $\nabla F(\mathbf{x})$ when $\mathbf{z} \sim P$. Note that finding the gradient of \tilde{F} may not be easy as it contains exponentially many terms. Instead, we can provide computationally cheap unbiased estimators for $\nabla \tilde{F}(\mathbf{x}, \mathbf{z})$. It can easily be shown that

$$\frac{\partial \tilde{F}}{\partial x_i} = \tilde{F}(\mathbf{x}, \mathbf{z}; x_i \leftarrow 1) - \tilde{F}(\mathbf{x}, \mathbf{z}; x_i \leftarrow 0). \quad (35)$$

where for example by $(\mathbf{x}; x_i \leftarrow 1)$ we mean a vector which has value 1 on its i -th coordinate and is equal to \mathbf{x} elsewhere. To create an unbiased estimator for $\frac{\partial \tilde{F}}{\partial x_i}$ at a point \mathbf{x} with realization \mathbf{z} we can simply sample a set S by including each element in it independently with probability x_i and use $\tilde{f}(S \cup \{i\}, \mathbf{z}) - \tilde{f}(S \setminus \{i\}, \mathbf{z})$ as an unbiased estimator for the i -th partial derivative of \tilde{F} . We can sample one single set S and use the above trick for all the coordinates. This involves n function computations for \tilde{f} .

Indeed, the stochastic gradient ascent method proposed by Hassani et al. (2017) can be used to solve the multilinear extension problem in (34) 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 of Karimi et al. (2017) achieves a $(1 - 1/e)$ approximation solution only when each $\tilde{f}(\cdot, \mathbf{z})$ is a coverage function. Here, we show that SCG achieves the first $(1 - 1/e)$ tight approximation guarantee for the discrete stochastic submodular Problem (33). More precisely, we show that SCG finds a solution for (34), 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 11 *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 2 with iterates \mathbf{x}_t , and recall the definition of the multilinear extension function F in (34). 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 the following*

$$|\nabla_j F(\mathbf{x}_{t+1}) - \nabla_j F(\mathbf{x}_t)| \leq m_f \sqrt{r} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|, \quad (36)$$

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

Proof See Appendix I. ■

The result in Lemma 11 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 of the matroid.

Let us now explain how the variance of the stochastic gradients of F relates to the variance of the marginal values of f . Recall that the stochastic function \tilde{F} is a multilinear extension of the stochastic set function \tilde{f} , and it can be shown that

$$\nabla_j \tilde{F}(\mathbf{x}, \mathbf{z}) = \tilde{F}(\mathbf{x}, \mathbf{z}; x_j = 1) - \tilde{F}(\mathbf{x}, \mathbf{z}; x_j = 0). \quad (37)$$

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

$$\mathbb{E} \left[\|\nabla \tilde{F}(\mathbf{x}, \mathbf{z}) - \nabla F(\mathbf{x})\|^2 \right] \leq n \max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]. \quad (38)$$

Using the result of Lemma 11, the expression in (38), and a coordinate-wise analysis, the bounds in Theorem 7 can be improved and specified for the case of multilinear extension maximization problem as we show in the following theorem.

Theorem 12 Consider *Stochastic Continuous Greedy (SCG)* outlined in Algorithm 2. Recall the definition of the multilinear extension function F in (34) and the definitions of r and m_f in Lemma 11. Further, set the averaging parameter as $\rho_t = 4/(t + 8)^{2/3}$. If Assumption 4 holds and the function f is monotone and submodular, then the iterate \mathbf{x}_T generated by SCG satisfies the inequality

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - 1/e)OPT - \frac{15DK}{T^{1/3}} - \frac{m_f r D^2}{2T}, \quad (39)$$

where $K := \max\{3\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|, 2\sqrt{n}\sqrt{\max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]} + \sqrt{3r}m_f D\}$ and OPT is the optimal value of Problem (34).

Proof The proof is similar to the proof of Theorem 7. Check Appendix J. ■

The result of Theorem 12 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 (39) are independent of n , except K that is at most proportional to \sqrt{n} . As a result, we have the following guarantee for SCG in the case of multilinear functions.

Corollary 13 Consider *Stochastic Continuous Greedy (SCG)* outlined in Algorithm 2. Suppose the conditions in Theorem 12 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. As a consequence, maximizing a stochastic Submodular set function with SCG requires $\mathcal{O}(n^{5/2}/\epsilon^3)$ evaluations of the function \tilde{f} in order to reach a $(1 - 1/e)OPT - \epsilon$ solution.

Proof According to the result in Theorem 12, 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. Since each iteration requires access to an unbiased estimator of the gradient $\nabla F(\mathbf{x})$ and it can be computed by n samples from $\tilde{f}(S_i, \mathbf{z})$ (Remark 10), then the total number of calls to the function \tilde{f} to reach a $(1 - 1/e)OPT - \epsilon$ solution is of order $\mathcal{O}(n^{5/2}/\epsilon^3)$ for the SCG method. ■

The result in Corollary 13 shows that after at most $\mathcal{O}(n^{5/2}/\epsilon^3)$ function evaluations of the stochastic set function \tilde{f} the iterates generated by SCG achieves a continuous solution \mathbf{x}_t with an objective function value that satisfies $\mathbb{E}[F(\mathbf{x}_t)] \geq (1 - 1/e)OPT - \epsilon$ where OPT is the optimal objective function value of Problem (34). Further, by using a lossless rounding scheme we can obtain a discrete set S^\dagger such that $\mathbb{E}[f(S^\dagger)] \geq (1 - 1/e) \max_{S \in \mathcal{I}} f(S) - \epsilon$.

Remark 14 According to Corollary 13, to reach a $(1 - (1/e)OPT - \epsilon)$ solution the total number of calls to the function \tilde{f} for our proposed SCG method is $\mathcal{O}(n^{5/2}/\epsilon^3)$. The result in [Badanidiyuru and Vondrak, 2013] shows that for the deterministic case where we have access to f (not \tilde{f}), after $\mathcal{O}(n^2/\epsilon^4)$ calls to f , one can find $(1 - (1/e)OPT - \epsilon)$ solution. Hence, if we apply both methods to the deterministic submodular case, the overall complexity of our method is slightly worse than the one in [Badanidiyuru and Vondrak, 2013] in terms of dependency on n , while the dependency of our result on ϵ is better than the one in

[Badanidiyuru and Vondrak, 2013]. We would like to highlight that the proposed method and its complexity bound in [Badanidiyuru and Vondrak, 2013] are for the deterministic setting, and if one modifies the method in [Badanidiyuru and Vondrak, 2013] to be applicable to the considered stochastic setting, its query complexity $\mathcal{O}(n^2/\epsilon^4)$ will further increase.

Indeed, by following similar steps we can extend the result for NMSCG to the discrete submodular maximization problem when the objective function is non-monotone and stochastic. We formally prove this claim in the following theorem.

Theorem 15 *Consider Non-monotone Stochastic Continuous Greedy (NMSCG) outlined in Algorithm 3. Recall the definition of the multilinear extension function F in (34) and the definitions of r and m_f in Lemma 11. Further, set the averaging parameter as $\rho_t = 4/(t + 8)^{2/3}$. If Assumption 4 holds and the function f is non-monotone and submodular, then the iterate \mathbf{x}_T generated by SCG satisfies the inequality*

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1/e)OPT - \frac{15DK}{T^{1/3}} - \frac{m_f r D^2}{2T}, \quad (40)$$

where $K := \max\{3\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|, 2\sqrt{n}\sqrt{\max_{j \in [n]} \mathbb{E}[f(\{j\}, \mathbf{z})^2]} + \sqrt{3r}m_f D\}$ and OPT is the optimal value of Problem (34).

Proof The proof is similar to the proof of Theorem 9. Check Appendix K. ■

Corollary 16 *Consider Non-monotone Stochastic Continuous Greedy (NMSCG) outlined in Algorithm 3. Suppose the conditions in Theorem 15 are satisfied. Then, the sequence of iterates generated by NMSCG achieves a $(1/e)OPT - \epsilon$ solution after $\mathcal{O}(n^{3/2}/\epsilon^3)$ iterations. As a consequence, maximizing a stochastic Submodular set function with NMSCG requires $\mathcal{O}(n^{5/2}/\epsilon^3)$ evaluations of the function f in order to reach a $(1/e)OPT - \epsilon$ solution.*

5.1. Convergence Bounds Based on Curvature

For the continuous greedy method it has been shown that if the submodular function f has a curvature $c \in [0, 1]$ the algorithm reaches a $(1/c)(1 - e^{-c})$ approximation guarantee. In this section, we show that the same improvement can be established for SCG in the stochastic setting. To do so, we first formally define the curvature c of a monotone submodular function f as

$$c := 1 - \min_{S, j \notin S} \frac{f(S \cup \{j\}) - f(S)}{f(\{j\})}. \quad (41)$$

Indeed, smaller curvature value c leads to an easier submodular maximization problem, and, in this case, we should be able to achieve a tighter approximate solution. In the following theorem, we match this expectation and show that if $c < 1$ the bound in (27) can be improved.

Theorem 17 *Consider the proposed Stochastic Continuous Greedy (SCG) defined in (20)-(22). Further recall the definition of the function f curvature c in (41). If Assumption 4*

is satisfied and the function f is monotone and submodular, then the expected objective function value for the iterate \mathbf{x}_T generated by SCG satisfies the inequality

$$\mathbb{E}[F(\mathbf{x}_T)] \geq \frac{1}{c}(1 - e^{-c})OPT - \frac{15DK}{T^{1/3}} - \frac{m_f r D^2}{2T}, \quad (42)$$

where $K := \max\{3\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|, 2\sqrt{n}\sqrt{\max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]} + \sqrt{3r}m_f D\}$.

Proof Check Appendix L. ■

Corollary 18 *Consider Stochastic Continuous Greedy (SCG) outlined in Algorithm 2. Suppose the conditions in Theorem 17 are satisfied. Then, the sequence of iterates generated by SCG achieves a $((1 - e^{-c})/c)OPT - \epsilon$ solution after $\mathcal{O}(n^{3/2}/\epsilon^3)$ iterations. As a consequence, maximizing a stochastic submodular set function with SCG requires $\mathcal{O}(n^{5/2}/\epsilon^3)$ function evaluations.*

6. Numerical Experiments

In this section, we compare the performances of the proposed stochastic conditional gradient method with state-of-the-art algorithms in both convex and submodular settings.

6.1. Convex Setting

We first compare the proposed SFW algorithm and mini-batch FW for a stochastic quadratic program of the form (2). Then, we compare their performances in solving a matrix completion problem. In this section, by mini-batch FW we refer to a variant of FW that simply replaces gradients by a mini-batch of stochastic gradients.

Quadratic Programming. Consider a positive definite matrix $\mathbf{A} \in \mathbb{S}_{++}^n$, a vector $\mathbf{b} \in \mathbb{R}^n$, a random variable $\mathbf{z} \in \mathbb{R}^n$, and the random diagonal matrix $\text{diag}(\mathbf{z}) \in \mathbb{R}^{n \times n}$ defined by \mathbf{z} . The function F is defined as

$$F(\mathbf{x}) = \mathbb{E}[\tilde{F}(\mathbf{x}, \mathbf{z})] = \mathbb{E}\left[\frac{1}{2}\mathbf{x}^T(\mathbf{A} + \text{diag}(\mathbf{z}))\mathbf{x} + (\mathbf{b} + \mathbf{z})^T\mathbf{x}\right]. \quad (43)$$

We assume that each element of \mathbf{z} is sampled from a normal distribution $\mathcal{N}(0, \sigma)$. Therefore, the objective function can be simplified to $F(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{A}\mathbf{x} + \mathbf{b}^T\mathbf{x}$. Further, we assume that the set \mathcal{C} is defined as $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid l \leq x_i \leq u\}$. Here, we assume that the distribution is unknown to the algorithm and at each iteration we only have access to the stochastic gradients $\nabla\tilde{F}(\mathbf{x}, \mathbf{z}) = (\mathbf{A} + \text{diag}(\mathbf{z}))\mathbf{x} + \mathbf{b} + \mathbf{z}$.

In our experiments, we set the dimension of the problem to $n = 5$ and the lower bound and upper bounds for the set \mathcal{C} to $l = 10$ and $u = 100$. We construct \mathbf{A} and \mathbf{b} in such a way that the optimal solution of the unconstrained set, namely $-\mathbf{A}^{-1}\mathbf{b}$, does not belong to the set \mathcal{C} .

Figure 1 demonstrates the suboptimality gap $F(\mathbf{x}_T) - F(\mathbf{x}^*)$ for the iterates generated by the proposed SFW method (with batch size $b = 1$) as well as the naive stochastic

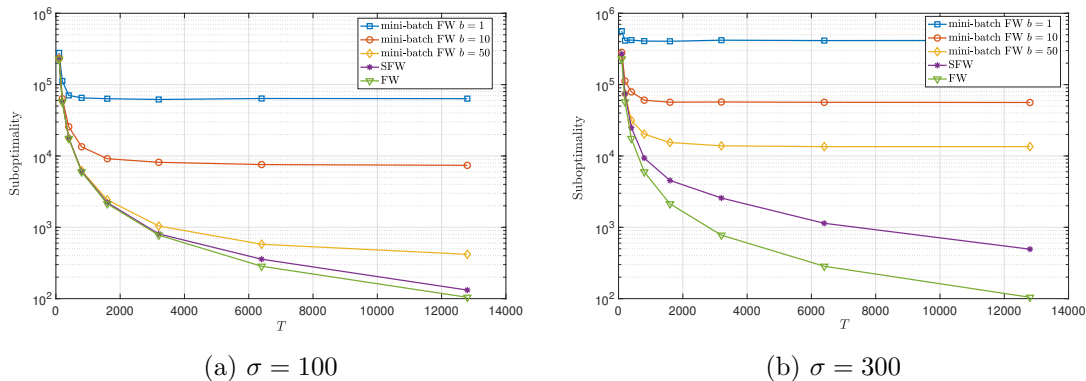


Figure 1: Comparison of the performances of (deterministic) FW, mini-batch FW ($b = 1$, $b = 10$, $b = 50$), and the proposed SFW method for the stochastic quadratic convex program defined in (43). The left and right plots correspond to the cases that the variance is $\sigma = 100$ and $\sigma = 300$, respectively. In both settings, the proposed SFW method that only uses a single stochastic gradient ($b = 1$) performs close to the full-batch (deterministic) FW algorithm, and it outperforms the mini-batch FW method with even larger batch sizes ($b = \{1, 10, 50\}$). This comparison is in terms of number of iterations. If we compare the algorithms in terms of number of evaluated stochastic gradients the gap between SFW and mini-batch FW algorithms with batch sizes larger than $b = 1$ will be even more significant as in this experiment SFW uses only a single stochastic gradient per iteration.

implementation of FW with batch sizes $b = \{1, 10, 50\}$ for the cases that the total number of iterations is $T = \{100, 200, 400, 800, 1600, 3200, 6400, 12800\}$. We further illustrates the performance of the (deterministic) FW as a benchmark. Indeed, to perform the update of FW we use the exact gradient $\mathbf{A}\mathbf{x} + \mathbf{b}$ at each iteration. Note that the optimal solution \mathbf{x}^* and the optimal objective function value $F(\mathbf{x}^*)$ are pre-computed by solving the quadratic program $\min_{\mathbf{x} \in \mathcal{C}} \frac{1}{2}\mathbf{x}^T \mathbf{A}\mathbf{x} + \mathbf{b}^T \mathbf{x}$. The left and right plots correspond to the cases that $\sigma = 100$ and $\sigma = 300$, respectively. In the left plot, which corresponds to the case that $\sigma = 100$, we observe that our proposed SFW method performs similar to the FW algorithm, while it uses only a single noisy stochastic gradient per iteration. The vanilla mini-batch FW method with a single stochastic gradient evaluation ($b = 1$) performs poorly. By increasing the size of the batch, the performance of the mini-batch FW improves, but it still underperforms SFW. In the right plot, which corresponds to the case with a larger variance, naturally the gap between the deterministic FW method and the stochastic algorithms becomes more significant. In this case, we observe that mini-batch FW even with large batch size of $b = 100$ is significantly worse than the proposed SFW method that only uses a single stochastic gradient per iteration. It is worth mentioning that increasing the batch size in mini-batch FW accelerates convergence and improves convergence accuracy, however, the suboptimality saturates at some point. In contrast, SFW converges to the optimal objective function at a sublinear rate in both small and large variance cases, matching our theory.

Matrix Completion. In this experiment, we study the performance of our proposed SFW algorithm in solving a matrix completion problem which is a canonical application of conditional gradient (Frank-Wolfe) type methods. We focus on a special case of matrix completion in which matrices are assumed to be symmetric. In particular, consider a symmetric matrix $\mathbf{C} \in \mathbb{S}^n$, where we only have access to a subset of its indices indicated by \mathcal{O} . Note that as \mathbf{C} is symmetric, if (i, j) is observed, i.e., $(i, j) \in \mathcal{O}$, then the pair (j, i) is also known, i.e., $(j, i) \in \mathcal{O}$. Our goal is to find a symmetric positive semidefinite matrix \mathbf{X} such that its elements in the set \mathcal{O} are close to the ones for \mathbf{C} , while its nuclear rank is smaller than a threshold. In other words, we focus on the optimization problem

$$\begin{aligned} \min f(\mathbf{X}) &:= \frac{1}{2} \sum_{(i,j) \in \mathcal{O}} \|\mathbf{X}_{ij} - \mathbf{C}_{ij}\|^2 \\ \text{s. t. } &\|\mathbf{X}\| \succeq \mathbf{0}, \quad \|\mathbf{X}\|_* \leq \alpha. \end{aligned} \tag{44}$$

In our simulations, we set the dimension to $n = 200$. We form the observation matrix as $\mathbf{C} = \hat{\mathbf{X}} + \mathbf{E}$. Here, $\hat{\mathbf{X}}$ is defined as $\hat{\mathbf{X}} = \mathbf{W}\mathbf{W}^T$ where $\mathbf{W} \in \mathbb{R}^{n \times r}$ has independent normal distributed entries, and \mathbf{E} is defined as $\mathbf{E} = \frac{1}{10}(\mathbf{L} + \mathbf{L}^T)$ where $\mathbf{L} \in \mathbb{R}^{n \times n}$ has independent normal distributed entries. In our experiments, we set the rank to $r = 10$ and the bound on the nuclear norm to $\alpha = \|\hat{\mathbf{X}}\|_*$, where $\|\hat{\mathbf{X}}\|_*$ is the nuclear norm of the matrix $\hat{\mathbf{X}}$. For settings that $\hat{\mathbf{X}}$ is not known in advance, one might use different choices of α and pick the one that performs better. We further define the set of observed entries \mathcal{O} by sampling the elements of the upper triangular part of \mathbf{C} uniformly at random with probability 0.8. Therefore, the size of the set \mathcal{O} is around $0.8 \times 200^2 = 32,000$. In the realization that we use the set \mathcal{O} has 31,884 elements.

To solve (44), by using FW method, we need to solve the subproblem (Hazan et al., 2016, Chapter 7)

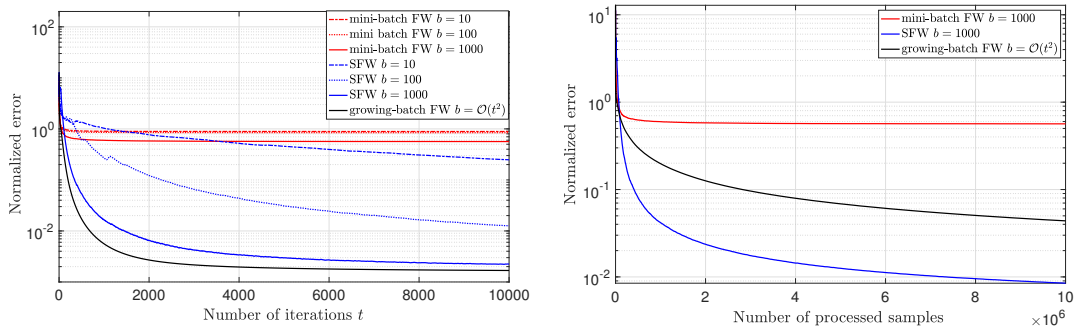
$$\begin{aligned} \min \text{tr}(\nabla f(\mathbf{X}_t)^T \mathbf{V}) \\ \text{s. t. } &\|\mathbf{V}\| \succeq \mathbf{0}, \quad \|\mathbf{V}\|_* \leq \alpha. \end{aligned} \tag{45}$$

where the gradient $\nabla f(\mathbf{X}_k) \in \mathbb{R}^{n \times n}$ is defined as $\nabla f(\mathbf{X}_t)_{i,j} = \mathbf{X}_{ij} - \mathbf{C}_{ij}$ if $(i, j) \in \mathcal{O}$, and $\nabla f(\mathbf{X}_t)_{i,j} = 0$, otherwise. It can be shown that the solution to the subproblem (45) is given by

$$\mathbf{V}_t = \begin{cases} \alpha \mathbf{v}_n \mathbf{v}_n^T & \text{if } \lambda_n \geq 0, \\ \mathbf{0} & \text{if } \lambda_n < 0, \end{cases} \tag{46}$$

where λ_n is the smallest eigenvalue of the gradient $\nabla f(\mathbf{X}_t)$ and \mathbf{v}_n is its corresponding eigenvector.

Indeed, evaluation of the gradient $\nabla f(\mathbf{X}_t)$ requires access to all observed elements in the set \mathcal{O} which can be computationally costly. In such cases, one may use a subset of the set \mathcal{O} as an unbiased estimate of the gradient. In our experiments, we consider (i) the mini-batch FW method that uses b elements of \mathcal{O} to compute a stochastic approximation of $\nabla f(\mathbf{X}_t)$, (ii) the growing mini-batch FW method proposed by Hazan and Luo (2016) which uses a batch size of $b = \mathcal{O}(t^2)$ at step t , and (iii) the proposed SFW method that uses the average of stochastic gradients over time as suggested in (3).



(a) Comparison in terms of number of iterations (b) Comparison in terms of number of samples

Figure 2: Comparison of the normalized error $\sum_{(i,j) \in \mathcal{O}} \|\mathbf{X}_{ij} - \mathbf{C}_{ij}\|^2 / \sum_{(i,j) \in \mathcal{O}} \|\mathbf{C}_{ij}\|^2$ of the mini-batch FW algorithm, the growing mini-batch FW method proposed in (Hazan and Luo, 2016), and our proposed SFW method for solving the matrix completion problem defined in (44). For all the choices of batch sizes considered here ($b = \{10, 100, 1000\}$), the proposed SFW method outperforms the corresponding version of the mini-batch FW algorithm. As we observe, increasing the size of the mini-batch b improves the convergence speed of both algorithms, but the mini-batch FW algorithm with $b = 1000$ still underperforms our proposed method with $b = 10$. In terms of number of iterations, the growing-batch FW method outperforms our proposed SFW method by using growing batches of size $b = \mathcal{O}(t^2)$, e.g., at iteration $t = 10^4$ it uses 10^8 samples; however, in terms of number of samples processed, SFW has the best performance as illustrated in the right plot.

Figure 2 illustrates the convergence paths of mini-batch FW and SFW for batch sizes of $b = \{10, 100, 1000\}$ as well as the convergence path of growing-batch FW in terms of both number of iterations and number of samples processed. Here, the normalized error is defined as

$$\text{Normalized error} := \frac{\sum_{(i,j) \in \mathcal{O}} \|\mathbf{X}_{ij} - \mathbf{C}_{ij}\|^2}{\sum_{(i,j) \in \mathcal{O}} \|\mathbf{C}_{ij}\|^2}.$$

Note that the stepsize for all the three algorithms is $\gamma_t = \frac{1}{t+1}$ and the averaging parameter for our proposed SFW algorithm is $\rho_t = \frac{1}{(t+1)^{2/3}}$. As we observe in Figure 2a, even for a large batch size of $b = 1000$, the mini-batch FW algorithm cannot obtain a normalized error better than 0.55 after 10,000 iterations. On the other hand, SFW with a small batch size of $b = 10$ achieves an error of 0.25 after 10,000 iterations. Indeed, by increasing the batch size for SFW its performance becomes better. In particular, SFW with $b = 1000$ achieves a normalized error of 2.3×10^{-3} after 10,000 iterations. We would like to highlight that even for the case that we set $b = 1000$, we use less than 3.2% of the observed elements – The number of observed elements is 31,884. However, the best performance in terms of number of iterations belongs to the growing-batch FW method that uses a batch size of $b = \mathcal{O}(t^2)$. This is not surprising as the convergence rate of the growing-batch FW algorithm is $\mathcal{O}(1/t)$, while the convergence rate of our proposed SFW is $\mathcal{O}(1/t^{1/3})$. However, the number of processed samples at each iteration by our method is




Figure 3: Comparison of the normalized error $\sum_{(i,j) \in \mathcal{O}} \|\mathbf{X}_{ij} - \mathbf{C}_{ij}\|^2 / \sum_{(i,j) \in \mathcal{O}} \|\mathbf{C}_{ij}\|^2$ of the proposed SFW method with different choices of the averaging parameter ρ_t for solving the matrix completion problem defined in (44). In this case we set the other parameters to $b = 100$ and $\gamma_t = 1/(t + 1)$. As suggested by our theory, the best performance belongs to the case that $\rho_t = \mathcal{O}(t^{-2/3})$.

much smaller than the one for growing-batch FW when t becomes large. To be more precise, after T iterations our method uses Tb samples, where b is a constant much smaller than T , while the growing-batch FW uses $\mathcal{O}(T^3)$ samples. Therefore, to have a better comparison between these algorithms we also compare their normalized errors versus the number of processed samples which is a more accurate measure for comparing the sample complexity of these algorithms. As we observe in Figure 2b, the proposed SFW method outperforms both mini-batch FW and growing-batch FW algorithms when we compare their normalized errors versus number of samples used. Note that in theory, both growing-batch FW and SFW may require processing $\mathcal{O}(1/\epsilon^3)$ samples to reach a suboptimality gap of $f(\mathbf{x}_t) - f^* \leq \epsilon$, but in practice we observe that the proposed SFW method outperforms growing-batch FW.

We proceed to study the effect of the averaging parameter ρ_t on the convergence of SFW. Our theoretical bound suggests that the best convergence guarantee is achieved when $\rho_t = \mathcal{O}(t^{-2/3})$. In this experiment we aim to check if this choice is reasonable relative to other possible sublinear rates. To do so, we compare the convergence paths of SFW with four different choices $\rho_t = \mathcal{O}(t^{-1/3})$, $\rho_t = \mathcal{O}(t^{-1/2})$, $\rho_t = \mathcal{O}(t^{-2/3})$, and $\rho_t = \mathcal{O}(t^{-1})$. As it can be observed in Figure 3, the best performance among these four choices belongs to $\rho_t = \mathcal{O}(t^{-2/3})$ used in our theoretical results. We would like to highlight that this experiment does not prove that $\rho_t = \mathcal{O}(t^{-2/3})$ is the optimal choice.

6.2. Submodular Setting

For the submodular setting, 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 in (33). 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 over the set of users. 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):

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

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 $O(n)$ time (the linear program step reduces to finding the largest k elements of a vector of length n). Therefore, this 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 very large.

In our experiments, we consider 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 of F are given in Remark 10.
- (ii) Stochastic Gradient Ascent (SG/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., 2017b; 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., 2017b) 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 users are picked and the function f is estimated by the average over 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

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

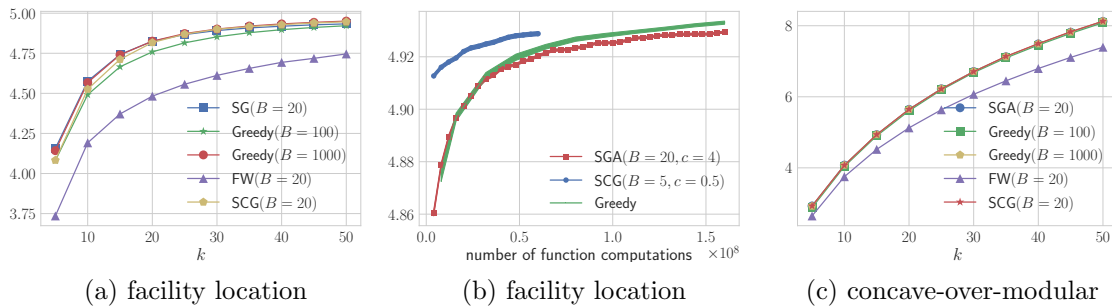


Figure 4: Comparison of the performances of SG/SGA, Greedy, FW, and SCG in a movie recommendation application. In all plots, y-axis represents the objective function value $f(S)$. Fig. 4a illustrates the performance of the algorithms in terms of the facility-location objective value w.r.t. the cardinality constraint size k after $T = 2000$ iterations. Fig. 4b 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. 4c demonstrates the concave-over-modular objective function value vs. the size of the cardinality constraint k after running the algorithms for $T = 2000$ iterations.

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

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

Figure 4 depicts the performance of different algorithms for the two proposed objective functions. As Figures 4a and 4c 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 4b shows the performance of the algorithms with respect to the number of times the simple functions $f(\cdot, i)$ are evaluated. Note that the total number of (simple) function evaluations for SGA and SCG 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.

7. Conclusion

In this paper, we developed stochastic conditional gradient methods for solving convex minimization and submodular maximization problems. The main idea of the proposed methods in both domains was using a momentum term in the stochastic gradient approximation step to reduce the noise of the stochastic approximation. In particular, in the convex setting, we proposed a stochastic variant of the Frank-Wolfe algorithm called **Stochastic Frank-Wolfe** that achieves an ϵ -suboptimal objective function value after at

most $\mathcal{O}(1/\epsilon^3)$ iterations, if the expected objective function is smooth. The main advantage of the **Stochastic Frank-Wolfe** method (SFW) comparing to other stochastic conditional gradient methods for stochastic convex minimization is the required condition on the batch size. In particular, the batch size for SFW can be as small as 1, while the state-of-the-art stochastic conditional gradient methods require a growing batch size.

In the submodular setting, 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 further extended our result to the non-monotone case and introduced **Non-monotone Stochastic Continuous Greedy** that obtains a $(1/e)$ approximation guarantee. 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. Likewise, we provided the first $1/e$ approximation guarantee for maximizing a *non-monotone* 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 the stochastic setting.

Acknowledgments

This work was partially 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 research of H. Hassani is supported by NSF HDR TRIPODS award 1934876, NSF award CPS-1837253, NSF award CIF-1910056, and NSF CAREER award CIF-1943064, and Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) under award FA9550-20-1-0111. A. Karbasi is partially supported by NSF CAREER (IIS-1845032), ONR YIA (N00014-19-1-2406), and AFOSR YIP (FA9550-18-1-0160).

Appendix

A. Proof of Lemma 1

Use the definition $\mathbf{d}_t := (1 - \rho_t)\mathbf{d}_{t-1} + \rho_t\nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)$ to write $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2$ as

$$\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 = \|\nabla F(\mathbf{x}_t) - (1 - \rho_t)\mathbf{d}_{t-1} - \rho_t\nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2. \quad (47)$$

Add and subtract the term $(1 - \rho_t)\nabla F(\mathbf{x}_{t-1})$ to the right hand side of (47), regroup the terms and expand the squared term to obtain

$$\begin{aligned}
 & \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \\
 &= \|\nabla F(\mathbf{x}_t) - (1 - \rho_t)\nabla F(\mathbf{x}_{t-1}) + (1 - \rho_t)\nabla F(\mathbf{x}_{t-1}) - (1 - \rho_t)\mathbf{d}_{t-1} - \rho_t\nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 \\
 &= \|\rho_t(\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)) + (1 - \rho_t)(\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})) + (1 - \rho_t)(\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1})\|^2 \\
 &= \rho_t^2\|\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 + (1 - \rho_t)^2\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2 \\
 &\quad + (1 - \rho_t)^2\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 \\
 &\quad + 2\rho_t(1 - \rho_t)(\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t))^T(\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})) \\
 &\quad + 2\rho_t(1 - \rho_t)(\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t))^T(\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}) \\
 &\quad + 2(1 - \rho_t)^2(\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1}))^T(\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}). \tag{48}
 \end{aligned}$$

Compute the conditional expectation $\mathbb{E}[\cdot | \mathcal{F}_t]$ for both sides of (48), and use the fact that $\nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)$ is an unbiased estimator of the gradient $\nabla F(\mathbf{x}_t)$, i.e., $\mathbb{E}[\nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t) | \mathcal{F}_t] = \nabla F(\mathbf{x}_t)$, to obtain

$$\begin{aligned}
 & \mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 | \mathcal{F}_t] \\
 &= \rho_t^2\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 | \mathcal{F}_t] + (1 - \rho_t)^2\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2 \\
 &\quad + (1 - \rho_t)^2\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 + 2(1 - \rho_t)^2(\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1}))^T(\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}). \tag{49}
 \end{aligned}$$

Use the condition in Assumption 3 to replace $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \nabla\tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 | \mathcal{F}_t]$ by its upper bound σ^2 . Further, using Young's inequality to substitute the inner product $2\langle \nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1}), \nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1} \rangle$ by the upper bounded $\beta_t\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 + (1/\beta_t)\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2$ where $\beta_t > 0$ is a free parameter. Applying these substitutions into (49) yields

$$\begin{aligned}
 \mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 | \mathcal{F}_t] &\leq \rho_t^2\sigma^2 + (1 - \rho_t)^2(1 + \beta_t^{-1})\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2 \\
 &\quad + (1 - \rho_t)^2(1 + \beta_t)\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2. \tag{50}
 \end{aligned}$$

According to Assumption 2, the norm $\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|$ is bounded above by $L\|\mathbf{x}_t - \mathbf{x}_{t-1}\|$. In addition, the condition in Assumption 4 implies that $L\|\mathbf{x}_t - \mathbf{x}_{t-1}\| = L\frac{1}{T}\|\mathbf{v}_t - \mathbf{x}_t\| \leq \frac{1}{T}LD$. Therefore, we can replace $\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|$ in (50) by its upper bound $\frac{1}{T}LD$ and write

$$\begin{aligned}
 & \mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 | \mathcal{F}_t] \\
 &\leq \rho_t^2\sigma^2 + \gamma_t^2(1 - \rho_t)^2(1 + \beta_t^{-1})L^2D^2 + (1 - \rho_t)^2(1 + \beta_t)\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2. \tag{51}
 \end{aligned}$$

Since we assume that $\rho_t \leq 1$ we can replace all the terms $(1 - \rho_t)^2$ in (51) by $(1 - \rho_t)$. Applying this substitution into (51) and setting $\beta := \rho_t/2$ lead to the inequality

$$\begin{aligned}
 \mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 | \mathcal{F}_t] &\leq \rho_t^2\sigma^2 + \gamma_t^2(1 - \rho_t)(1 + \frac{2}{\rho_t})L^2D^2 \\
 &\quad + (1 - \rho_t)(1 + \frac{\rho_t}{2})\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2. \tag{52}
 \end{aligned}$$

Now using the inequalities $(1 - \rho_t)(1 + (2/\rho_t)) \leq (2/\rho_t)$ and $(1 - \rho_t)(1 + (\rho_t/2)) \leq (1 - \rho_t/2)$ we obtain

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \mid \mathcal{F}_t] \leq \rho_t^2 \sigma^2 + \frac{2L^2 D^2 \gamma_t^2}{\rho_t} + \left(1 - \frac{\rho_t}{2}\right) \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2, \quad (53)$$

and the claim in (9) follows.

B. Proof of Lemma 2

Based on the L -smoothness of the expected function F we show that $F(\mathbf{x}_{t+1})$ is bounded above by

$$F(\mathbf{x}_{t+1}) \leq F(\mathbf{x}_t) + \nabla F(\mathbf{x}_t)^T (\mathbf{x}_{t+1} - \mathbf{x}_t) + \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2. \quad (54)$$

Replace the terms $\mathbf{x}_{t+1} - \mathbf{x}_t$ in (54) by $\gamma_{t+1}(\mathbf{v}_t - \mathbf{x}_t)$ and add and subtract the term $\gamma_{t+1} \mathbf{d}_t^T (\mathbf{v}_t - \mathbf{x}_t)$ to the right hand side of the resulted inequality to obtain

$$F(\mathbf{x}_{t+1}) \leq F(\mathbf{x}_t) + \gamma_{t+1} (\nabla F(\mathbf{x}_t) - \mathbf{d}_t)^T (\mathbf{v}_t - \mathbf{x}_t) + \gamma_{t+1} \mathbf{d}_t^T (\mathbf{v}_t - \mathbf{x}_t) + \frac{L\gamma_{t+1}^2}{2} \|\mathbf{v}_t - \mathbf{x}_t\|^2. \quad (55)$$

Since $\langle \mathbf{x}^*, \mathbf{d}_t \rangle \geq \min_{v \in \mathcal{C}} \langle \mathbf{v}, \mathbf{d}_t \rangle = \langle \mathbf{v}_t, \mathbf{d}_t \rangle$, we can replace the inner product $\langle \mathbf{v}, \mathbf{d}_t \rangle$ in (55) by its upper bound $\langle \mathbf{x}^*, \mathbf{d}_t \rangle$. Applying this substitution leads to

$$F(\mathbf{x}_{t+1}) \leq F(\mathbf{x}_t) + \gamma_{t+1} (\nabla F(\mathbf{x}_t) - \mathbf{d}_t)^T (\mathbf{v}_t - \mathbf{x}_t) + \gamma_{t+1} \mathbf{d}_t^T (\mathbf{x}^* - \mathbf{x}_t) + \frac{L\gamma_{t+1}^2}{2} \|\mathbf{v}_t - \mathbf{x}_t\|^2. \quad (56)$$

Add and subtract $\gamma_{t+1} \nabla F(\mathbf{x}_t)^T (\mathbf{x}^* - \mathbf{x}_t)$ to the right hand side of (56) and regroup the terms to obtain

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\leq F(\mathbf{x}_t) + \gamma_{t+1} (\nabla F(\mathbf{x}_t) - \mathbf{d}_t)^T (\mathbf{v}_t - \mathbf{x}^*) + \gamma_{t+1} \nabla F(\mathbf{x}_t)^T (\mathbf{x}^* - \mathbf{x}_t) \\ &\quad + \frac{LD\gamma_{t+1}^2}{2} \|\mathbf{v}_t - \mathbf{x}_t\|^2. \end{aligned} \quad (57)$$

Using the Cauchy-Schwarz inequality, we can show that the inner product $(\nabla F(\mathbf{x}_t) - \mathbf{d}_t)^T (\mathbf{v}_t - \mathbf{x}^*)$ is bounded above by $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| \|\mathbf{v}_t - \mathbf{x}^*\|$. Moreover, the inner product $\nabla F(\mathbf{x}_t)^T (\mathbf{x}^* - \mathbf{x}_t)$ is upper bounded by $F(\mathbf{x}^*) - F(\mathbf{x}_t)$ due to the convexity of the function F . Applying these substitutions into (57) implies that

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\leq F(\mathbf{x}_t) + \gamma_{t+1} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| \|\mathbf{v}_t - \mathbf{x}^*\| - \gamma_{t+1} (F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \frac{L\gamma_{t+1}^2}{2} \|\mathbf{v}_t - \mathbf{x}^*\|^2 \\ &\leq F(\mathbf{x}_t) + \gamma_{t+1} D \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \gamma_{t+1} (F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \frac{LD^2\gamma_{t+1}^2}{2}, \end{aligned} \quad (58)$$

where the second inequality holds since $\|\mathbf{v}_t - \mathbf{x}^*\| \leq D$ according to Assumption 1. Finally, subtract the optimal objective function value $F(\mathbf{x}^*)$ from both sides of (58) and regroup the terms to obtain

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*) \leq (1 - \gamma_{t+1})(F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \gamma_{t+1} D \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| + \frac{LD^2\gamma_{t+1}^2}{2}, \quad (59)$$

and the claim in (10) follows.

C. Proof of Theorem 3

Computing the expectation of both sides of (10) with respect to \mathcal{F}_0 yields

$$\mathbb{E} [F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)] \leq (1 - \gamma_{t+1})\mathbb{E} [(F(\mathbf{x}_t) - F(\mathbf{x}^*))] + \gamma_{t+1}D\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|] + \frac{LD^2\gamma_{t+1}^2}{2} \quad (60)$$

Using Jensen's inequality, the expression $\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|]$ can be substituted by its upper bound $\sqrt{\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]}$ to obtain

$$\begin{aligned} & \mathbb{E} [F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)] \\ & \leq (1 - \gamma_{t+1})\mathbb{E} [(F(\mathbf{x}_t) - F(\mathbf{x}^*))] + \gamma_{t+1}D\sqrt{\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]} + \frac{LD^2\gamma_{t+1}^2}{2}. \end{aligned} \quad (61)$$

Next, we study the rate that the sequence of expected gradient errors $\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ converges to zero. To do so, we first prove the following lemma which is an extension of Lemma 8 in (Mokhtari and Ribeiro, 2015).

Lemma 19 *Consider the scalars $b \geq 0$ and $c > 1$. Let ϕ_t be a sequence of real numbers satisfying*

$$\phi_t \leq \left(1 - \frac{c}{(t + t_0)^\alpha}\right) \phi_{t-1} + \frac{b}{(t + t_0)^{2\alpha}}, \quad (62)$$

for some $0 \leq \alpha \leq 1$ and $t_0 \geq 0$. Then, the sequence ϕ_t converges to zero at the following rate

$$\phi_t \leq \frac{Q}{(t + t_0 + 1)^\alpha}, \quad (63)$$

where $Q := \max\{\phi_0(t_0 + 1)^\alpha, b/(c - 1)\}$.

Proof We prove the claim in (63) by induction. First, note that $Q \geq \phi_0(t_0 + 1)^\alpha$ and therefore $\phi_0 \leq Q/(t_0 + 1)^\alpha$ and the base step of the induction holds true. Now assume that the condition in (63) holds for $t = k$, i.e.,

$$\phi_k \leq \frac{Q}{(k + 1 + t_0)^\alpha}. \quad (64)$$

The goal is to show that (63) also holds for $t = k + 1$. To do so, first set $t = k + 1$ in the expression in (62) to obtain

$$\phi_{k+1} \leq \left(1 - \frac{c}{(k + 1 + t_0)^\alpha}\right) \phi_k + \frac{b}{(k + 1 + t_0)^{2\alpha}}. \quad (65)$$

According to the definition of Q , we know that $b \leq Q(c - 1)$. Moreover, based on the induction hypothesis it holds that $\phi_k \leq \frac{Q}{(k+1+t_0)^\alpha}$. Using these inequalities and the expression in (65) we can write

$$\phi_{k+1} \leq \left(1 - \frac{c}{(k + 1 + t_0)^\alpha}\right) \frac{Q}{(k + 1 + t_0)^\alpha} + \frac{Q(c - 1)}{(k + 1 + t_0)^{2\alpha}}. \quad (66)$$

Pulling out $\frac{Q}{(k+1+t_0)^\alpha}$ as a common factor and simplifying and reordering terms it follows that (66) is equivalent to

$$\phi_{k+1} \leq Q \left(\frac{(k+1+t_0)^\alpha - 1}{(k+1+t_0)^{2\alpha}} \right). \quad (67)$$

Based on the inequality

$$((k+1+t_0)^\alpha - 1)((k+1+t_0)^\alpha + 1) < (k+1+t_0)^{2\alpha}, \quad (68)$$

the result in (67) implies that

$$\phi_{k+1} \leq \left(\frac{Q}{(k+1+t_0)^\alpha + 1} \right). \quad (69)$$

Since $(k+1+t_0)^\alpha + 1 \geq (k+1+t_0+1)^\alpha = (k+t_0+2)^\alpha$, the result in (69) implies that

$$\phi_{k+1} \leq \left(\frac{Q}{(k+2+t_0)^\alpha} \right), \quad (70)$$

and the induction step is complete. Therefore, the result in (63) holds for all $t \geq 0$. \blacksquare

Now using the result in Lemma 19 we can characterize the convergence of the sequence of expected errors $\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ to zero. To be more precise, compute the expectation of both sides of the result in (9) with respect to \mathcal{F}_0 and set $\gamma_t = 2/(t+8)$ and $\rho_t = 4/(t+8)^{2/3}$ to obtain

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \leq \left(1 - \frac{2}{(t+8)^{2/3}} \right) \mathbb{E} [\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2] + \frac{16\sigma^2 + 2L^2D^2}{(t+8)^{4/3}}. \quad (71)$$

According to the result in Lemma 19, the inequality in (71) implies that

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \leq \frac{Q}{(t+9)^{2/3}}, \quad (72)$$

where $Q = \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 2L^2D^2\}$. This result is achieved by setting $\phi_t = \mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$, $\alpha = 2/3$, $b = 16\sigma^2 + 2L^2D^2$, $c = 2$, and $t_0 = 8$ in Lemma 19.

Now we proceed by replacing the term $\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]$ in (61) by its upper bound in (72) and γ_{t+1} by $2/(t+9)$ to write

$$\mathbb{E} [F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)] \leq \left(1 - \frac{2}{t+9} \right) \mathbb{E} [(F(\mathbf{x}_t) - F(\mathbf{x}^*))] + \frac{2D\sqrt{Q}}{(t+9)^{4/3}} + \frac{2LD^2}{(t+9)^2}. \quad (73)$$

Note that we can write $(t+9)^2 = (t+9)^{4/3}(t+9)^{2/3} \geq (t+9)^{4/3}9^{2/3} \geq 4(t+9)^{4/3}$. Therefore,

$$\mathbb{E} [F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)] \leq \left(1 - \frac{2}{t+9} \right) \mathbb{E} [(F(\mathbf{x}_t) - F(\mathbf{x}^*))] + \frac{2D\sqrt{Q} + LD^2/2}{(t+9)^{4/3}}. \quad (74)$$

Now we proceed to prove by induction for $t \geq 0$ that

$$\mathbb{E}[F(\mathbf{x}_t) - F(\mathbf{x}^*)] \leq \frac{Q'}{(t+9)^{1/3}} \quad (75)$$

where $Q' = \max\{9^{1/3}(F(\mathbf{x}_0) - F(\mathbf{x}^*)), 2D\sqrt{Q} + LD^2/2\}$. To do so, first note that $Q' \geq 9^{1/3}(F(\mathbf{x}_0) - F(\mathbf{x}^*))$, and, therefore, $F(\mathbf{x}_0) - F(\mathbf{x}^*) \leq Q'/9^{1/3}$. This leads to the base of induction for $t = 0$. Now assume that the inequality (75) holds for $t = k$, i.e., $\mathbb{E}[F(\mathbf{x}_k) - F(\mathbf{x}^*)] \leq \frac{Q'}{(k+9)^{1/3}}$ and we aim to show that it also holds for $t = k + 1$.

To do so first set $t = k$ in (74) and replace $\mathbb{E}[F(\mathbf{x}_k) - F(\mathbf{x}^*)]$ by its upper bounds $\frac{Q'}{(k+9)^{1/3}}$ (as guaranteed by the hypothesis of the induction) to obtain

$$\mathbb{E}[F(\mathbf{x}_{k+1}) - F(\mathbf{x}^*)] \leq \left(1 - \frac{2}{k+9}\right) \frac{Q'}{(k+9)^{1/3}} + \frac{2D\sqrt{Q} + LD^2/2}{(k+9)^{4/3}}. \quad (76)$$

Now as in the proof of Lemma 19, replace $2D\sqrt{Q} + LD^2/2$ by Q' and simplify the terms to reach the inequality

$$\mathbb{E}[F(\mathbf{x}_{k+1}) - F(\mathbf{x}^*)] \leq Q' \left(\frac{k+8}{(k+9)^{4/3}} \right) \leq \frac{Q'}{(k+10)^{1/3}}, \quad (77)$$

and the induction is complete. Therefore, the inequality in (75) holds for all $t \geq 0$.

D. Proof of Theorem 4

To prove the claim in (13) we first show that the sum $\sum_{t=1}^{\infty} \rho_t \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2$ is finite almost surely. To do so, we construct a supermartingale using the result in Lemma 1. Let's define the stochastic process ζ_t as

$$\zeta_t := \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 + 2L^2D^2 \sum_{u=t+1}^{\infty} \frac{\gamma_u^2}{\rho_u} + \sigma^2 \sum_{u=t+1}^{\infty} \rho_u^2. \quad (78)$$

Note that ζ_t is well defined because the sums on the the right hand side of (78) are finite according to the hypotheses of Theorem 4. Further, define the stochastic process ξ_t as

$$\xi_t := \frac{\rho_{t+1}}{2} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2. \quad (79)$$

Considering the definitions of the sequences ζ_t and ξ_t and expression (9) in Lemma 1 we can write

$$\mathbb{E}[\zeta_t | \mathbf{x}_t] \leq \zeta_{t-1} - \xi_{t-1}. \quad (80)$$

Since the sequences ζ_t and ξ_t are nonnegative it follows from (80) that they satisfy the conditions of the supermartingale convergence theorem; see e.g. (Theorem E7.4 in (Solo and Kong, 1994)). Therefore, we can conclude that: (i) The sequence ζ_t converges almost surely to a limit. (ii) The sum $\sum_{t=0}^{\infty} \xi_t < \infty$ is almost surely finite. Hence, the second result implies that

$$\sum_{t=1}^{\infty} \rho_t \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 < \infty, \quad \text{a.s.} \quad (81)$$

Based on expression (10) in Lemma 2, we know that the suboptimality $F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*)$ is upper bounded by

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*) \leq (1 - \gamma_{t+1})(F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \gamma_{t+1}D\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| + \frac{LD^2\gamma_{t+1}^2}{2}. \quad (82)$$

Further use Young's inequality to replace $\gamma_{t+1}\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|$ by the sum $\beta_{t+1}\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 + \gamma_{t+1}^2/\beta_{t+1}$ where β_{t+1} is a free positive parameter. Set $\beta_{t+1} = \rho_{t+1}$ to obtain $\gamma_{t+1}\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| \leq \rho_{t+1}\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 + \gamma_{t+1}^2/\rho_{t+1}$. Applying this substitution into (82) implies that

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}^*) \leq (1 - \gamma_{t+1})(F(\mathbf{x}_t) - F(\mathbf{x}^*)) + \rho_{t+1}D\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 + \frac{D\gamma_{t+1}^2}{\rho_{t+1}} + \frac{LD^2\gamma_{t+1}^2}{2}. \quad (83)$$

To conclude the almost sure convergence of the sequence $F(\mathbf{x}_t) - F(\mathbf{x}^*)$ to zero from the expression in (82) we first state the following Lemma from (Bertsekas and Tsitsiklis, 1996).

Lemma 20 *Let $\{X_t\}$, $\{Y_t\}$, and $\{Z_t\}$ be three sequences of numbers such that $Y_t \geq 0$ for all $t \geq 0$. Suppose that*

$$X_{t+1} \leq X_t - Y_t + Z_t, \quad \text{for } t = 0, 1, 2, \dots \quad (84)$$

and $\sum_{t=0}^{\infty} Z_t < \infty$. Then, either $X_t \rightarrow -\infty$ or else $\{X_t\}$ converges to a finite value and $\sum_{t=0}^{\infty} Y_t < \infty$.

From now on we focus on realizations that support $\sum_{t=1}^{\infty} \rho_t \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 < \infty$ which have probability 1, according to the result in (81).

Consider the outcome of Lemma 20 with the identifications $X_t = F(\mathbf{x}_t) - F(\mathbf{x}^*)$, $Y_t = \gamma_{t+1}(F(\mathbf{x}_t) - F(\mathbf{x}^*))$, and $Z_t = \rho_{t+1}D\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 + \frac{D\gamma_{t+1}^2}{\rho_{t+1}} + \frac{LD^2\gamma_{t+1}^2}{2}$. Since the sequence $X_t = F(\mathbf{x}_t) - F(\mathbf{x}^*)$ is always non-negative, the first outcome of Lemma 20 is impossible and therefore we obtain that $F(\mathbf{x}_t) - F(\mathbf{x}^*)$ converges to a finite limit and

$$\sum_{t=0}^{\infty} \gamma_{t+1}F(\mathbf{x}_t) - F(\mathbf{x}^*) < \infty. \quad (85)$$

Recall that both of these results hold almost surely, since they are valid for the realization that $\sum_{t=1}^{\infty} \rho_t \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 < \infty$, which occur with probability 1 as shown in (81). The result in (85) implies that $\liminf_{t \rightarrow \infty} F(\mathbf{x}_t) - F(\mathbf{x}^*) = 0$ almost surely. Moreover, we know that the sequence $\{F(\mathbf{x}_t) - F(\mathbf{x}^*)\}$ almost surely converges to a finite limit. Combining these two observation we obtain that the finite limit is zero, and, therefore, $\lim_{t \rightarrow \infty} F(\mathbf{x}_t) - F(\mathbf{x}^*) = 0$ almost surely. Hence, the claim in (13) follows.

E. Proof of Lemma 6

By following the steps from (47) to (49) in the proof of Lemma 1 we can show that

$$\begin{aligned} & \mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \mid \mathcal{F}_t] \\ &= \rho_t^2 \mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 \mid \mathcal{F}_t \right] + (1 - \rho_t)^2 \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 \\ & \quad + (1 - \rho_t)^2 \|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2 + 2(1 - \rho_t)^2 \langle \nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1}), \nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1} \rangle. \end{aligned} \quad (86)$$

The term $\mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \nabla \tilde{F}(\mathbf{x}_t, \mathbf{z}_t)\|^2 \mid \mathcal{F}_t \right]$ can be bounded above by σ^2 according to Assumption 6. Based on Assumptions 4 and 5, we can also show that the squared norm $\|\nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1})\|^2$ is upper bounded by $L^2 D^2 / T^2$. Moreover, the inner product $2\langle \nabla F(\mathbf{x}_t) - \nabla F(\mathbf{x}_{t-1}), \nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1} \rangle$ can be upper bounded by $\beta_t \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 + (1/\beta_t)L^2 D^2 / T^2$ using Young's inequality (i.e., $2\langle \mathbf{a}, \mathbf{b} \rangle \leq \beta \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 / \beta$ for any $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ and $\beta > 0$) and the conditions in Assumptions 4 and 5, where $\beta_t > 0$ is a free scalar. Applying these substitutions into (49) leads to

$$\begin{aligned} & \mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \mid \mathcal{F}_t \right] \\ & \leq \rho_t^2 \sigma^2 + (1 - \rho_t)^2 \left(1 + \frac{1}{\beta_t}\right) \frac{L^2 D^2}{T^2} + (1 - \rho_t)^2 (1 + \beta_t) \|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2. \end{aligned} \quad (87)$$

Now by following the steps from (50) to (53) and computing the expected value with respect to \mathcal{F}_0 we obtain

$$\mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \right] \leq \left(1 - \frac{\rho_t}{2}\right) \mathbb{E} \left[\|\nabla F(\mathbf{x}_{t-1}) - \mathbf{d}_{t-1}\|^2 \right] + \rho_t^2 \sigma^2 + \frac{2L^2 D^2}{\rho_t T^2}. \quad (88)$$

Define $\phi_t := \mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \right]$ and set $\rho_t = \frac{4}{(t+8)^{2/3}}$ to obtain

$$\phi_t \leq \left(1 - \frac{2}{(t+8)^{2/3}}\right) \phi_{t-1} + \frac{16\sigma^2}{(t+8)^{4/3}} + \frac{L^2 D^2 (t+8)^{2/3}}{2T^2}. \quad (89)$$

Now use the conditions $8 \leq T$ and $t \leq T$ to replace $1/T$ in (89) by its upper bound $2/(t+8)$. Applying this substitution leads to

$$\phi_t \leq \left(1 - \frac{2}{(t+8)^{2/3}}\right) \phi_{t-1} + \frac{16\sigma^2 + 2L^2 D^2}{(t+8)^{4/3}}. \quad (90)$$

Now using the result in Lemma 19, we obtain that

$$\phi_t \leq \frac{Q}{(t+9)^{2/3}}, \quad (91)$$

where $Q := \max\{5\phi_0, 16\sigma^2 + 2L^2 D^2\}$. Replacing ϕ_t by its definition $\mathbb{E} \left[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2 \right]$ yields (26).

F. Proof of Theorem 7

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

$$\begin{aligned} F(\mathbf{x}_{t+1}) & \geq F(\mathbf{x}_t) + \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle - \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ & = F(\mathbf{x}_t) + \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2, \end{aligned} \quad (92)$$

where the equality follows from the update in (22). Since \mathbf{v}_t is in the set \mathcal{C} , it follows from Assumption 4 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 (92) to obtain

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{v}_t, \mathbf{d}_t \rangle + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2} \\ &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{x}^*, \mathbf{d}_t \rangle + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2}. \end{aligned} \quad (93)$$

Note that the second inequality in (93) holds since based on (21) we can write

$$\langle \mathbf{x}^*, \mathbf{d}_t \rangle \leq \max_{\mathbf{v} \in \mathcal{C}} \{ \langle \mathbf{v}, \mathbf{d}_t \rangle \} = \langle \mathbf{v}_t, \mathbf{d}_t \rangle. \quad (94)$$

Now add and subtract the inner product $\langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle / T$ to the RHS of (93) to get

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle + \frac{1}{T} \langle \mathbf{v}_t - \mathbf{x}^*, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2}. \quad (95)$$

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

$$\langle \mathbf{v}_t - \mathbf{x}^*, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle \geq -\frac{\beta_t}{2} \|\mathbf{v}_t - \mathbf{x}^*\|^2 - \frac{1}{2\beta_t} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2, \quad (96)$$

for any $\beta_t > 0$. By applying these substitutions into (95) we obtain

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{1}{T} (F(\mathbf{x}^*) - F(\mathbf{x}_t)) - \frac{LD^2}{2T^2} - \frac{1}{2T} \left(\beta_t \|\mathbf{v}_t - \mathbf{x}^*\|^2 + \frac{\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2}{\beta_t} \right). \quad (97)$$

Replace $\|\mathbf{v}_t - \mathbf{x}^*\|^2$ by its upper bound $4D^2$ and compute the expected value of (97) to write

$$\begin{aligned} &\mathbb{E}[F(\mathbf{x}_{t+1})] \\ &\geq \mathbb{E}[F(\mathbf{x}_t)] + \frac{1}{T} \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_t)] - \frac{1}{2T} \left[4\beta_t D^2 + \frac{\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]}{\beta_t} \right] - \frac{LD^2}{2T^2}. \end{aligned} \quad (98)$$

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 (26). Further, set $\beta_t = (Q^{1/2})/(2D(t+9)^{1/3})$ and regroup the resulted expression to obtain

$$\mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_{t+1})] \leq \left(1 - \frac{1}{T}\right) \mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_t)] + \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \frac{LD^2}{2T^2}. \quad (99)$$

By applying the inequality in (99) recursively for $t = 0, \dots, T-1$ we obtain

$$\mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_T)] \leq \left(1 - \frac{1}{T}\right)^T (F(\mathbf{x}^*) - F(\mathbf{x}_0)) + \sum_{t=0}^{T-1} \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \sum_{t=0}^{T-1} \frac{LD^2}{2T^2}. \quad (100)$$

Note that we can write

$$\begin{aligned}
\sum_{t=0}^{T-1} \frac{1}{(t+9)^{1/3}} &\leq \frac{1}{9^{1/3}} + \int_{t=0}^{T-1} \frac{1}{(t+9)^{1/3}} dt \\
&= \frac{1}{9^{1/3}} + \frac{3}{2}(t+9)^{2/3} \Big|_{t=T-1} - \frac{3}{2}(t+9)^{2/3} \Big|_{t=0} \\
&\leq \frac{3}{2}(T+8)^{2/3} \\
&\leq \frac{15}{2}T^{2/3}
\end{aligned} \tag{101}$$

where the last inequality holds since $(T+8)^{2/3} \leq 5T^{2/3}$ for any $T \geq 1$. By simplifying the terms on the right hand side (100) and using the inequality in (101) we can write

$$\mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_T)] \leq \frac{1}{e}(F(\mathbf{x}^*) - F(\mathbf{x}_0)) + \frac{15DQ^{1/2}}{T^{1/3}} + \frac{LD^2}{2T}. \tag{102}$$

Here, we use the fact that $F(\mathbf{x}_0) \geq 0$, and hence the expression in (102) can be simplified to

$$\mathbb{E}[F(\mathbf{x}_T)] \geq (1 - 1/e)F(\mathbf{x}^*) - \frac{15DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \tag{103}$$

and the claim in (27) follows.

G. Proof of Theorem 8

Following the steps of the proof of Theorem 7 we can derive the inequality

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{1}{T}\langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle + \frac{1}{T}\langle \mathbf{v}_t - \mathbf{x}^*, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{LD^2}{2T^2}. \tag{104}$$

Using the definition of weak DR-submodularity and monotonicity of F we can show that $\langle \mathbf{x}^*, \nabla F(\mathbf{x}_t) \rangle \geq \gamma(F(\mathbf{x}^*) - F(\mathbf{x}_t))$. Further, based on Young's inequality 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 $\beta_t > 0$. Applying these substitutions into (104) leads to

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{\gamma}{T}(F(\mathbf{x}^*) - F(\mathbf{x}_t)) - \frac{LD^2}{2T^2} - \frac{1}{2T} \left(\beta_t \|\mathbf{v}_t - \mathbf{x}^*\|^2 + \frac{\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2}{\beta_t} \right). \tag{105}$$

Substitute $\|\mathbf{v}_t - \mathbf{x}^*\|^2$ by its upper bound $4D^2$ and compute the expected value of (105) to write

$$\begin{aligned}
&\mathbb{E}[F(\mathbf{x}_{t+1})] \\
&\geq \mathbb{E}[F(\mathbf{x}_t)] + \frac{\gamma}{T}\mathbb{E}[F(\mathbf{x}^*) - F(\mathbf{x}_t)] - \frac{1}{2T} \left[4\beta_t D^2 + \frac{\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2]}{\beta_t} \right] - \frac{LD^2}{2T^2}. \tag{106}
\end{aligned}$$

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 (26). Further, set $\beta_t = (Q^{1/2})/(2D(t+9)^{1/3})$ and regroup the resulted expression to obtain

$$\mathbb{E} [F(\mathbf{x}^*) - F(\mathbf{x}_{t+1})] \leq \left(1 - \frac{\gamma}{T}\right) \mathbb{E} [F(\mathbf{x}^*) - F(\mathbf{x}_t)] + \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \frac{LD^2}{2T^2}. \quad (107)$$

By applying the inequality in (107) recursively for $t = 0, \dots, T-1$ we obtain

$$\mathbb{E} [F(\mathbf{x}^*) - F(\mathbf{x}_T)] \leq \left(1 - \frac{\gamma}{T}\right)^T (F(\mathbf{x}^*) - F(\mathbf{x}_0)) + \sum_{t=0}^{T-1} \frac{2DQ^{1/2}}{(t+9)^{1/3}T} + \sum_{t=0}^{T-1} \frac{LD^2}{2T^2}. \quad (108)$$

Simplify the terms on the right hand side (108) and use (101) to obtain

$$\mathbb{E} [F(\mathbf{x}^*) - F(\mathbf{x}_T)] \leq e^{-\gamma}(F(\mathbf{x}^*) - F(\mathbf{x}_0)) + \frac{15DQ^{1/2}}{T^{1/3}} + \frac{LD^2}{2T}. \quad (109)$$

Here, we use the fact that $F(\mathbf{x}_0) \geq 0$, and hence the expression in (109) can be simplified to

$$\mathbb{E} [F(\mathbf{x}_T)] \geq (1 - e^{-\gamma})F(\mathbf{x}^*) - \frac{15DQ^{1/2}}{T^{1/3}} - \frac{LD^2}{2T}, \quad (110)$$

and the claim in (29) follows.

H. Proof of Theorem 9

Using the update of the NMSCG method we can write

$$\begin{aligned} \mathbf{x}_{i,t+1} &= \mathbf{x}_{i,t} + \frac{1}{T} \mathbf{v}_{i,t} \\ &\leq \mathbf{x}_{i,t} + \frac{1}{T} (\bar{\mathbf{u}}_i - \mathbf{x}_{i,t}) \\ &\leq \left(1 - \frac{1}{T}\right) \mathbf{x}_{i,t} + \frac{1}{T} \bar{\mathbf{u}}_i \\ &= \left(1 - \frac{1}{T}\right)^t \mathbf{x}_{i,0} + \frac{1}{T} \bar{\mathbf{u}}_i \sum_{j=0}^{t-1} \left(1 - \frac{1}{T}\right)^j \\ &= \bar{\mathbf{u}}_i \left(1 - \left(1 - \frac{1}{T}\right)^{t+1}\right) \end{aligned} \quad (111)$$

where the first inequality follows by the condition $\mathbf{v}_{i,t} \leq \bar{\mathbf{u}}_i - \mathbf{x}_{i,t}$. The result in (111) implies that $\mathbf{x}_t \leq \bar{\mathbf{u}}(1 - (1 - 1/T)^t)$. Therefore, according to Lemma 3 in (Bian et al., 2017a) which is a generalized version of Lemma 7 in (Chekuri et al., 2015) it follows that

$$F(\mathbf{x}_t \vee \mathbf{x}^*) \geq (1 - (1/T))^t F(\mathbf{x}^*). \quad (112)$$

Using this result and the Taylor's expansion of the objective function F we can write

$$\begin{aligned} F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) &\geq \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle - \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2, \end{aligned} \quad (113)$$

where the equality follows by the update of NMSCG. Add and subtract $\langle \mathbf{d}_t, \mathbf{v}_t \rangle$ to the right hand side of (113) to obtain

$$\begin{aligned} F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) &\geq \frac{1}{T} \langle \mathbf{d}_t, \mathbf{v}_t \rangle + \frac{1}{T} \langle \nabla F(\mathbf{x}_t) - \mathbf{d}_t, \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2 \\ &\geq \frac{1}{T} \langle \mathbf{d}_t, \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle + \frac{1}{T} \langle \nabla F(\mathbf{x}_t) - \mathbf{d}_t, \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2, \end{aligned} \quad (114)$$

where the second inequality is valid since $\{\mathbf{d}_t^T \mathbf{v}_t\} \geq \{\mathbf{d}_t^T \mathbf{y}\}$ for all $\mathbf{y} \leq \bar{\mathbf{u}} - \mathbf{x}_t$ and we know that $\mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \leq \bar{\mathbf{u}} - \mathbf{x}_t$. Now add and subtract $(1/T) \langle \nabla F(\mathbf{x}_t), \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle$ to the right hand side of (114) to obtain

$$\begin{aligned} &F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \\ &\geq \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle + \frac{1}{T} \langle \mathbf{d}_t - \nabla F(\mathbf{x}_t), \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle \\ &\quad + \frac{1}{T} \langle \nabla F(\mathbf{x}_t) - \mathbf{d}_t, \mathbf{v}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2 \\ &= \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle + \frac{1}{T} \langle \nabla F(\mathbf{x}_t) - \mathbf{d}_t, \mathbf{v}_t + \mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^* \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2 \\ &\geq \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{x}_t \vee \mathbf{x}^* - \mathbf{x}_t \rangle - \frac{2D}{T} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \frac{L}{2T^2} \|\mathbf{v}_t\|^2. \end{aligned} \quad (115)$$

The last inequality holds since the inner product $\langle \nabla F(\mathbf{x}_t) - \mathbf{d}_t, \mathbf{v}_t + \mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^* \rangle$ can be upper bounded by $\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| \|\mathbf{v}_t + \mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^*\|$ using the Cauchy-Schwartz inequality and further we can upper bound the norm $\|\mathbf{v}_t + \mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^*\|$ by $2D$ since $\|\mathbf{v}_t + \mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^*\| \leq \|\mathbf{v}_t\| + \|\mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^*\|$ and both $\mathbf{x}_t - \mathbf{x}_t \vee \mathbf{x}^*$ and \mathbf{v}_t belong to the set \mathcal{C} . This holds due to the assumption that \mathcal{C} is down-closed.

Now replace $\|\mathbf{v}_t\|^2$ by its upper bound D^2 and use the concavity of the function F in positive directions to write

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \geq \frac{1}{T} (F(\mathbf{x}_t \vee \mathbf{x}^*) - F(\mathbf{x}_t)) - \frac{2D}{T} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \frac{LD^2}{2T^2}. \quad (116)$$

Now use the expression in (112) to obtain

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \geq \frac{1}{T} \left[\left(1 - \frac{1}{T}\right)^t F(\mathbf{x}^*) - F(\mathbf{x}_t) \right] - \frac{2D}{T} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \frac{LD^2}{2T^2}. \quad (117)$$

Computing the expectation of both sides and replacing $\mathbb{E}[\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|]$ by its upper bound $\sqrt{Q}/((t+9)^{1/3})$ according to the result in (26) lead to

$$\mathbb{E}[F(\mathbf{x}_{t+1})] \geq \left(1 - \frac{1}{T}\right) \mathbb{E}[F(\mathbf{x}_t)] + \frac{1}{T} \left(1 - \frac{1}{T}\right)^t F(\mathbf{x}^*) - \frac{2D\sqrt{Q}}{T(t+9)^{1/3}} - \frac{LD^2}{2T^2}. \quad (118)$$

By applying the inequality in (118) recursively for $t = 0, \dots, T - 1$ we obtain

$$\begin{aligned}
 & \mathbb{E}[F(\mathbf{x}_T)] \\
 & \geq \left(1 - \frac{1}{T}\right)^T F(\mathbf{x}_0) + \sum_{i=0}^{T-1} \left(\frac{1}{T} \left(1 - \frac{1}{T}\right)^i F(\mathbf{x}^*) - \frac{2D\sqrt{Q}}{T(i+9)^{1/3}} - \frac{LD^2}{2T^2}\right) \left(1 - \frac{1}{T}\right)^{T-1-i} \\
 & = \left(1 - \frac{1}{T}\right)^T F(\mathbf{x}_0) + \left(1 - \frac{1}{T}\right)^{T-1} F(\mathbf{x}^*) - \sum_{i=0}^{T-1} \left(\frac{2D\sqrt{Q}}{T(i+9)^{1/3}} + \frac{LD^2}{2T^2}\right) \left(1 - \frac{1}{T}\right)^{T-1-i} \\
 & \geq \left(1 - \frac{1}{T}\right)^T F(\mathbf{x}_0) + \left(1 - \frac{1}{T}\right)^{T-1} F(\mathbf{x}^*) - \sum_{i=0}^{T-1} \left(\frac{2D\sqrt{Q}}{T(i+9)^{1/3}} + \frac{LD^2}{2T^2}\right), \tag{119}
 \end{aligned}$$

where the second inequality holds since $(1 - 1/T)$ is strictly smaller than 1. Replacing the sum $\sum_{i=0}^{T-1} \frac{1}{(i+9)^{1/3}}$ in (119) by its upper bound in (101) leads to

$$\mathbb{E}[F(\mathbf{x}_T)] \geq \left(1 - \frac{1}{T}\right)^T F(\mathbf{x}_0) + \left(1 - \frac{1}{T}\right)^{T-1} F(\mathbf{x}^*) - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{LD^2}{2T} \tag{120}$$

Use the fact that $F(\mathbf{x}_0) \geq 0$ and the inequality $(1 - \frac{1}{T})^{T-1} \geq e^{-1}$ to obtain

$$\mathbb{E}[F(\mathbf{x}_T)] \geq e^{-1} F(\mathbf{x}^*) - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{LD^2}{2T}, \tag{121}$$

and the claim in (32) follows.

I. Proof of Lemma 11

Based on the mean value theorem, we can write

$$\nabla F(\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t) - \nabla F(\mathbf{x}_t) = \frac{1}{T}\mathbf{H}(\tilde{\mathbf{x}}_t)\mathbf{v}_t, \tag{122}$$

where $\tilde{\mathbf{x}}_t$ is a convex combination of \mathbf{x}_t and $\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t$ and $\mathbf{H}(\tilde{\mathbf{x}}_t) := \nabla^2 F(\tilde{\mathbf{x}}_t)$. This expression shows that the difference between the coordinates of the vectors $\nabla F(\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t)$ and $\nabla F(\mathbf{x}_t)$ can be written as

$$\nabla_j F(\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t) - \nabla_j F(\mathbf{x}_t) = \frac{1}{T} \sum_{i=1}^n H_{j,i}(\tilde{\mathbf{x}}_t) v_{i,t}, \tag{123}$$

where $v_{i,t}$ is the i -th element of the vector \mathbf{v}_t and $H_{j,i}$ denotes the component in the j -th row and i -th column of the matrix \mathbf{H} . Hence, the norm of the difference $|\nabla_j F(\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t) - \nabla_j F(\mathbf{x}_t)|$ is bounded above by

$$|\nabla_j F(\mathbf{x}_t + \frac{1}{T}\mathbf{v}_t) - \nabla_j F(\mathbf{x}_t)| \leq \frac{1}{T} \left| \sum_{i=1}^n H_{j,i}(\tilde{\mathbf{x}}_t) v_{i,t} \right|. \tag{124}$$

Note here that the elements of the matrix $\mathbf{H}(\tilde{\mathbf{x}}_t)$ are less than the maximum marginal value (i.e. $\max_{i,j} |H_{i,j}(\tilde{\mathbf{x}}_t)| \leq \max_{i \in \{1, \dots, n\}} f(i) \triangleq m_f$). We thus get

$$|\nabla_j F(\mathbf{x}_t + \frac{1}{T} \mathbf{v}_t) - \nabla_j F(\mathbf{x}_t)| \leq \frac{m_f}{T} \sum_{i=1}^n |v_{i,t}|. \quad (125)$$

Note that at each round t of the algorithm, we have to pick a vector $\mathbf{v}_t \in \mathcal{C}$ s.t. the inner product $\langle \mathbf{v}_t, \mathbf{d}_t \rangle$ is maximized. Hence, without loss of generality we can assume that the vector \mathbf{v}_t is one of the extreme points of \mathcal{C} , i.e. it is of the form $\mathbf{1}_I$ for some $I \in \mathcal{I}$ (note that we can easily force integer vectors). Therefore by noticing that \mathbf{v}_t is an integer vector with at most r ones, we have

$$|\nabla_j F(\mathbf{x}_t + \frac{1}{T} \mathbf{v}_t) - \nabla_j F(\mathbf{x}_t)| \leq \frac{m_f \sqrt{r}}{T} \sqrt{\sum_{i=1}^n |v_{i,t}|^2}, \quad (126)$$

which yields the claim in (36).

J. Proof of Theorem 12

According to the Taylor's expansion of the function F near the point \mathbf{x}_t we can write

$$\begin{aligned} F(\mathbf{x}_{t+1}) &= F(\mathbf{x}_t) + \langle \nabla F(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{1}{2} \langle \mathbf{x}_{t+1} - \mathbf{x}_t, \mathbf{H}(\tilde{\mathbf{x}}_t) (\mathbf{x}_{t+1} - \mathbf{x}_t) \rangle \\ &= F(\mathbf{x}_t) + \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle + \frac{1}{2T^2} \langle \mathbf{v}_t, \mathbf{H}(\tilde{\mathbf{x}}_t) \mathbf{v}_t \rangle, \end{aligned} \quad (127)$$

where $\tilde{\mathbf{x}}_t$ is a convex combination of \mathbf{x}_t and $\mathbf{x}_t + \frac{1}{T} \mathbf{v}_t$ and $\mathbf{H}(\tilde{\mathbf{x}}_t) := \nabla^2 F(\tilde{\mathbf{x}}_t)$. Note that based on the inequality $\max_{i,j} |H_{i,j}(\tilde{\mathbf{x}}_t)| \leq \max_{i \in \{1, \dots, n\}} f(i) \triangleq m_f$, we can lower bound H_{ij} by $-m_f$. Therefore,

$$\begin{aligned} \langle \mathbf{v}_t, \mathbf{H}(\tilde{\mathbf{x}}_t) \mathbf{v}_t \rangle &= \sum_{j=1}^n \sum_{i=1}^n v_{i,t} v_{j,t} H_{ij}(\tilde{\mathbf{x}}_t) \\ &\geq -m_f \sum_{j=1}^n \sum_{i=1}^n v_{i,t} v_{j,t} = -m_f \left(\sum_{i=1}^n v_{i,t} \right)^2 = -m_f r \|\mathbf{v}_t\|^2, \end{aligned} \quad (128)$$

where the last equality is because \mathbf{v}_t is a vector with r ones and $n - r$ zeros (see the explanation in the proof of Lemma 11). Replace the expression $\langle \mathbf{v}_t, \mathbf{H}(\tilde{\mathbf{x}}_t) \mathbf{v}_t \rangle$ in (127) by its lower bound in (128) to obtain

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle - \frac{m_f r}{2T^2} \|\mathbf{v}_t\|^2. \quad (129)$$

In the following lemma we derive a variant of the result in Lemma 6 for the multilinear extension setting.

Lemma 21 Consider *Stochastic Continuous Greedy (SCG)* outlined in Algorithm 2, and recall the definitions of the function F in (34), the rank r , and $m_f \triangleq \max_{i \in \{1, \dots, n\}} f(i)$. If we set $\rho_t = \frac{4}{(t+8)^{2/3}}$, then for $t = 0, \dots, T$ and for $j = 1, \dots, n$ it holds

$$\mathbb{E} [\|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|^2] \leq \frac{Q}{(t+9)^{2/3}}, \quad (130)$$

where $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 3m_f^2 r D^2\}$.

Proof The proof is similar to the proof of Lemma 6. The main difference is to write the analysis for the j -th coordinate and replace L by $m_f \sqrt{r}$ as shown in Lemma 11. Then using the proof techniques in Lemma 6 the claim in Lemma 21 follows. \blacksquare

The rest of the proof is identical to the proof of Theorem 7, by following the steps from (92) to (103) and considering the bound in (130) we obtain

$$\mathbb{E} [F(\mathbf{x}_T)] \geq (1 - 1/e)F(\mathbf{x}^*) - \frac{2DQ^{1/2}}{T^{1/3}} - \frac{m_f r D^2}{2T}, \quad (131)$$

where $Q := \max\{5\|\nabla F(\mathbf{x}_0) - \mathbf{d}_0\|^2, 16\sigma^2 + 3m_f^2 r D^2\}$. Therefore, the claim in Theorem 12 follows by replacing σ^2 by $n \max_{j \in [n]} \mathbb{E}[\tilde{f}(\{j\}, \mathbf{z})^2]$ as shown in (38).

K. Proof of Theorem 15

Following the steps of the proof of Theorem 12 from (127) to (129) we obtain

$$F(\mathbf{x}_{t+1}) \geq F(\mathbf{x}_t) + \frac{1}{T} \langle \nabla F(\mathbf{x}_t), \mathbf{v}_t \rangle - \frac{m_f r}{2T^2} \|\mathbf{v}_t\|^2. \quad (132)$$

Then, by following the steps from (112) to (117) we obtain

$$F(\mathbf{x}_{t+1}) - F(\mathbf{x}_t) \geq \frac{1}{T} \left[\left(1 - \frac{1}{T}\right)^t F(\mathbf{x}^*) - F(\mathbf{x}_t) \right] - \frac{2D}{T} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \frac{m_f r D^2}{2T^2}. \quad (133)$$

Note that the result in Lemma 21 holds for both monotone and non-monotone functions F . Therefore, by computing the expected value of both sides of (117) we obtain Then, by following the steps from (112) to (117) we obtain

$$\mathbb{E} [F(\mathbf{x}_{t+1})] - \mathbb{E} [F(\mathbf{x}_t)] \geq \frac{1}{T} \left[\left(1 - \frac{1}{T}\right)^t F(\mathbf{x}^*) - \mathbb{E} [F(\mathbf{x}_t)] \right] - \frac{2D}{T} \frac{\sqrt{Q}}{(t+9)^{1/3}} - \frac{m_f r D^2}{2T^2}. \quad (134)$$

Now by following the steps from (118) to (121) the claim in (40) follows.

L. Proof of Theorem 17

According to the result in Lemma 3 of (Hassani et al., 2017), it can be shown that when the function F has a curvature c then

$$\max_{\mathbf{v}} \mathbf{v}^T \nabla F(\mathbf{x}_t) \geq F(\mathbf{x}^*) - cF(\mathbf{x}_t), \quad (135)$$

Using this inequality we can write

$$\begin{aligned} \mathbf{v}_t^T \mathbf{d}_t &= \max_{\mathbf{v}} \{\mathbf{v}^T \mathbf{d}_t\} \\ &= \max_{\mathbf{v}} \{\mathbf{v}^T \nabla F(\mathbf{x}_t) + \mathbf{v}^T (\nabla F(\mathbf{x}_t) - \mathbf{d}_t)\} \\ &\geq \max_{\mathbf{v}} \{\mathbf{v}^T \nabla F(\mathbf{x}_t)\} - \max_{\mathbf{v}} \{\mathbf{v}^T (\nabla F(\mathbf{x}_t) - \mathbf{d}_t)\} \\ &\geq F(\mathbf{x}^*) - cF(\mathbf{x}_t) - D \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\|. \end{aligned} \quad (136)$$

Therefore, using the result in (129) we can write

$$\begin{aligned} F(\mathbf{x}_{t+1}) &\geq F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) \rangle - \frac{m_f r}{2T^2} \|\mathbf{v}_t\|^2 \\ &= F(\mathbf{x}_t) + \frac{1}{T} \langle \mathbf{v}_t, \mathbf{d}_t \rangle + \frac{1}{T} \langle \mathbf{v}_t, \nabla F(\mathbf{x}_t) - \mathbf{d}_t \rangle - \frac{L}{2T^2} \|\mathbf{v}_t\|^2 \\ &\geq F(\mathbf{x}_t) + \frac{1}{T} (F(\mathbf{x}^*) - cF(\mathbf{x}_t)) - \frac{2D}{T} \|\nabla F(\mathbf{x}_t) - \mathbf{d}_t\| - \frac{m_f r D^2}{2T^2}. \end{aligned} \quad (137)$$

Compute the expected value of both sides and use the result of Lemma 21 to obtain

$$\mathbb{E}[F(\mathbf{x}_{t+1})] \geq \mathbb{E}[F(\mathbf{x}_t)] + \frac{1}{T} (F(\mathbf{x}^*) - c\mathbb{E}[F(\mathbf{x}_t)]) - \frac{2D\sqrt{Q}}{T(t+9)^{1/3}} - \frac{m_f r D^2}{2T^2}. \quad (138)$$

Subtract $F^* := F(\mathbf{x}^*)$ from both sides and regroup the terms to obtain

$$\mathbb{E}[F(\mathbf{x}_{t+1})] - F^* \geq (1 - \frac{c}{T})(\mathbb{E}[F(\mathbf{x}_t)] - F^*) + \frac{1-c}{T} F^* - \frac{2D\sqrt{Q}}{T(t+9)^{1/3}} - \frac{m_f r D^2}{2T^2}. \quad (139)$$

Now applying the expression in (139) for $t = 0, \dots, T-1$ recursively yields

$$\begin{aligned} &\mathbb{E}[F(\mathbf{x}_T)] - F^* \\ &\geq (1 - \frac{c}{T})^T (F(\mathbf{x}_0) - F^*) + \sum_{i=0}^{T-1} \left[\frac{1-c}{T} F^* - \frac{2D\sqrt{Q}}{T(i+9)^{1/3}} - \frac{m_f r D^2}{2T^2} \right] (1 - \frac{c}{T})^{T-1-i} \\ &\geq (1 - \frac{c}{T})^T (F(\mathbf{x}_0) - F^*) + \frac{1-c}{T} F^* \sum_{i=0}^{T-1} (1 - \frac{c}{T})^{T-1-i} - \sum_{i=0}^{T-1} \left[\frac{2D\sqrt{Q}}{T(i+9)^{1/3}} + \frac{m_f r D^2}{2T^2} \right] \\ &\geq (1 - \frac{c}{T})^T (F(\mathbf{x}_0) - F^*) + \frac{1-c}{c} \left(1 - (1 - \frac{c}{T})^T \right) F^* - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{m_f r D^2}{2T} \\ &\geq -(1 - \frac{c}{T})^T F^* + \frac{1-c}{c} \left(1 - (1 - \frac{c}{T})^T \right) F^* - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{m_f r D^2}{2T} \\ &= \left(\frac{1-c}{c} - \frac{1}{c} (1 - \frac{c}{T})^T \right) F^* - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{m_f r D^2}{2T}, \end{aligned} \quad (140)$$

where in the second inequality we use the fact that $(1 - c/T) \leq 1$, in the third inequality we use the result in (101), and in the fourth inequality we use the assumption that $F(\mathbf{x}_0) \geq 0$. Therefore, by regrouping the terms we obtain that

$$\begin{aligned} \mathbb{E}[F(\mathbf{x}_T)] &\geq \frac{1}{c} \left(1 - \left(1 - \frac{c}{T}\right)^T\right) F^* - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{m_f r D^2}{2T} \\ &\geq \frac{1}{c} (1 - e^{-c}) F^* - \frac{15D\sqrt{Q}}{T^{1/3}} - \frac{m_f r D^2}{2T}, \end{aligned} \quad (141)$$

and the claim in (42) follows.

References

- Francis Bach. Submodular functions: from discrete to continuous domains. *arXiv preprint arXiv:1511.00394*, 2015.
- Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-dynamic programming*, volume 3 of *Optimization and neural computation series*. Athena Scientific, 1996.
- Andrew An Bian, Kfir Yehuda Levy, Andreas Krause, and Joachim M. Buhmann. Non-monotone continuous dr-submodular maximization: Structure and algorithms. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 486–496, 2017a.
- Andrew An Bian, Baharan Mirzasoleiman, Joachim M. Buhmann, and Andreas Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, pages 111–120, 2017b.
- Jeffrey Bilmes and Wenruo Bai. Deep submodular functions. *arXiv preprint arXiv:1701.08939*, 2017.
- Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- Niv Buchbinder and Moran Feldman. Constrained submodular maximization via a non-symmetric technique. *arXiv preprint arXiv:1611.03253*, 2016.
- Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *SODA*, pages 1433–1452, 2014.
- Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. *SIAM Journal on Computing*, 44(5):1384–1402, 2015.
- Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

- C. Chekuri, T. S. Jayram, and J. Vondrak. 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.
- Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- Michael Collins, Amir Globerson, Terry Koo, Xavier Carreras, and Peter L Bartlett. Exponentiated gradient algorithms for conditional random fields and max-margin markov networks. *Journal of Machine Learning Research*, 9(Aug):1775–1822, 2008.
- Nikhil R. Devanur and Kamal Jain. Online matching with concave returns. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 137–144, 2012.
- Reza Eghbali and Maryam Fazel. Designing smoothing functions for improved worst-case competitive ratio in online optimization. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3279–3287, 2016.
- Alina Ene and Huy L. Nguyen. Constrained submodular maximization: Beyond $1/e$. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 248–257, 2016.
- Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.
- Uriel Feige, Vahab S Mirrokni, and Jan Vondrak. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 570–579, 2011.
- Moran Feldman, Christopher Harshaw, and Amin Karbasi. Greed is good: Near-optimal submodular maximization via greedy optimization. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017*, pages 758–784, 2017.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics (NRL)*, 3(1-2):95–110, 1956.
- Satoru Fujishige. *Submodular functions and optimization*, volume 58. Elsevier, 2005.
- Satoru Fujishige and Shiguelo Isotani. A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, 7(1):3–17, 2011.

- Shayan Oveis Gharan and Jan Vondrák. Submodular maximization by simulated annealing. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1098–1116. SIAM, 2011.
- Daniel Golovin, Andreas Krause, and Matthew Streeter. Online submodular maximization under a matroid constraint with application to learning assignments. *arXiv preprint arXiv:1407.1082*, 2014.
- Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. *arXiv preprint arXiv:1708.03949*, 2017.
- Simon S Haykin. *Adaptive filter theory*. Pearson Education India, 2008.
- Elad Hazan and Satyen Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*, pages 1843–1850, 2012.
- Elad Hazan and Haipeng Luo. Variance-reduced and projection-free stochastic optimization. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 1263–1271, 2016.
- Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 427–435, 2013.
- Mohammad Karimi, Mario Lucic, Hamed Hassani, and Andreas Krause. Stochastic submodular maximization: The case of coverage functions. In *Advances in Neural Information Processing Systems*, 2017.
- László Lovász. Submodular functions and convexity. In *Mathematical Programming The State of the Art*, pages 235–257. Springer, 1983.
- Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5), 2007.
- Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, and Amin Karbasi. Fast constrained submodular maximization: Personalized data summarization. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 1358–1367, 2016.
- Aryan Mokhtari and Alejandro Ribeiro. Global convergence of online limited memory BFGS. *Journal of Machine Learning Research*, 16(1):3151–3181, 2015.
- Aryan Mokhtari, Alec Koppel, Gesualdo Scutari, and Alejandro Ribeiro. Large-scale non-convex stochastic optimization by doubly stochastic successive convex approximation. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 4701–4705. IEEE, 2017.

- G. L. Nemhauser and L. A. Wolsey. Maximizing submodular set functions: formulations and analysis of algorithms. *Studies on Graphs and Discrete Programming, volume 11 of Annals of Discrete Mathematics*, 1981.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions–I. *Mathematical Programming*, 14(1): 265–294, 1978.
- Arkadi Nemirovski and D Yudin. On Cezari’s convergence of the steepest descent method for approximating saddle point of convex-concave functions. In *Soviet Math. Dokl*, volume 19, 1978.
- Arkadii Nemirovskii, David Borisovich Yudin, and Edgar Ronald Dawson. Problem complexity and method efficiency in optimization. 1983.
- S. J. Reddi, S. Sra, B. Póczos, and A. Smola. Stochastic frank-wolfe methods for nonconvex optimization. *arXiv preprint arXiv:1607.08254*, 2016.
- Alejandro Ribeiro. Ergodic stochastic optimization algorithms for wireless communication and networking. *IEEE Transactions on Signal Processing*, 58(12):6369–6386, 2010.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- Andrzej Ruszczyński. Feasible direction methods for stochastic programming problems. *Mathematical Programming*, 19(1):220–229, 1980.
- Andrzej Ruszczyński. A merit function approach to the subgradient method with averaging. *Optimisation Methods and Software*, 23(1):161–172, 2008.
- Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- Victor Solo and Xuan Kong. *Adaptive signal processing algorithms: stability and performance*. Prentice-Hall, Inc., 1994.
- Tasuku Soma, Naonori Kakimura, Kazuhiro Inaba, and Ken-ichi Kawarabayashi. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 351–359, 2014.
- Matthew Staib and Stefanie Jegelka. Robust budget allocation via continuous submodular functions. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 3230–3240, 2017.
- Serban Stan, Morteza Zadimoghaddam, Andreas Krause, and Amin Karbasi. Probabilistic submodular maximization in sub-linear time. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 3241–3250, 2017.

- Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1134–1148, 2015.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Jan Vondrák. Submodularity in combinatorial optimization. 2007.
- Jan 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, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 67–74, 2008.
- Jan Vondrák. Submodularity and curvature : The optimal algorithm (combinatorial optimization and discrete algorithms). *RIMS Kokyuroku Bessatsu*, B23:253–266, 2010.
- Laurence A Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.
- Yang Yang, Gesualdo Scutari, Daniel P. Palomar, and Marius Pesavento. A parallel decomposition method for nonconvex stochastic multi-agent optimization problems. *IEEE Transactions on Signal Processing*, 64(11):2949–2964, 2016.