Proceedings of Machine Learning Research vol 125:1–19, 2020 33rd Annual Conference on Learning Theory

Adaptive Submodular Maximization under Stochastic Item Costs

Srinivasan Parthasarathy SPARTHA@US.IBM.COM
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

Constrained maximization of non-decreasing utility functions with submodularity-like properties is at the core of several AI and ML applications including viral marketing, pool-based active learning, adaptive feature selection, and sensor deployment. In this work, we develop adaptive policies for maximizing such functions when both the utility function and item costs may be stochastic.

First, we study maximization of an adaptive weak submodular function which is submodular in an approximate and probabilistic sense, subject to a stochastic fractional knapsack constraint, which requires total expected item cost to be within a given capacity. We present the β-GREEDY policy for this problem; our approximation guarantee for it recovers many known greedy maximization guarantees as special cases. Next, we study maximization of an adaptive submodular function, which is submodular in a probabilistic sense, subject to a stochastic knapsack constraint, which requires the total item cost to be within a given capacity with probability one. We present the MIX policy for this problem; our approximation guarantee for it is the first known approximation guarantee for maximizing a non-linear utility function subject to a stochastic knapsack constraint. Using alternative parameterizations of MIX, we also derive the first known approximation guarantees for maximizing an adaptive submodular function subject to a deterministic knapsack constraint.

Our guarantees are powered by an innovative differential analysis technique which models the β-GREEDY policy using a continuous-time stochastic reward process of a particle whose reward is related to the optimal utility through a differential inequality. The solution to this inequality yields our β-GREEDY guarantee. We combine differential analysis with a variety of other ideas to derive our MIX guarantees.

1. Introduction

Constrained maximization of non-decreasing utility functions with submodularity-like properties is at the core of several AI and ML applications. Submodularity is a property of diminishing returns which states that the increment in utility obtained by adding an item to an itemset decreases as this itemset is expanded through the addition of other items. In applications such as viral marketing, active learning, and sensor deployment, the utility function is stochastic and is submodular in a probabilistic sense. Golovin and Krause (2011) introduced the family of *adaptive submodular* functions for modeling such settings.[1] In other applications such as sparse regression and dictionary selection, the utility function is deterministic and submodular in an approximate sense. Das and Kempe (2018) introduced the concept of *submodularity ratio* for characterizing such *weak submodular* functions. Fujii and Sakaue (2019) generalized this concept for utility functions that are stochastic and are submodular in a sense that is both approximate and probabilistic. They introduced the concept of *adaptive submodularity ratio* for characterizing such *adaptive weak submodular* functions.

1. As is standard in literature, in this paper, an 'adaptive function' denotes a function that is stochastic.

We study maximization of these functions under stochastic item costs. Concretely, we have a ground set of items E; each $e \in E$ is endowed with a random state $S(e)$ as well as a random cost $C(e)$ where S and C are distributed according to known priors $\mathcal{P}$ and $\mathcal{Q}$ respectively.[2] The realizations of $S(e)$ and $C(e)$ can be observed once item e is selected. An *adaptive policy* selects items based on states and costs of items selected in previous iterations, in order to maximize the expected utility of selected itemset subject to a constraint on the total cost. In the stochastic fractional knapsack problem, the total expected cost conditional on any realization of item states needs to be within a given capacity B. We study this problem for the adaptive weak submodular function family which includes all other function families mentioned above as special cases. In the stochastic knapsack problem, total cost is required to be within the given capacity B with probability (w.p.) 1. Deterministic knapsack is a special case of stochastic knapsack with deterministic item costs. We study stochastic and deterministic knapsack problems for the adaptive submodular function family which includes submodular functions as a special case.

Our formulation encapsulates a wide range of practical applications, a few of which are illustrated next. **a) Viral marketing:** Consider a marketing campaign where we adaptively select users in a social network and 'influence' them by offering them a free trial. Users who are influenced may also influence others through a stochastic network diffusion process. Our objective, namely, the expected number of people in the network who are eventually influenced can be modeled as an adaptive monotone submodular function (Golovin and Krause, 2011; Sun et al., 2018). The cost of selecting a user could represent the duration of free trial required in order for the user to begin influencing neighbors; this may be stochastic and known only distributionally in advance. **b) Pool-based active learning:** Given a set of unlabeled examples, consider the problem of adaptively selecting a small subset for labeling s.t. their labels imply the labels of all other examples. Our objective, namely, the expected total 'shrinkage' in version space (the set of classifiers consistent with labeled examples), can be modeled as an adaptive monotone submodular function (Golovin and Krause, 2011; Chen and Krause, 2013; Cuong et al., 2013, 2014). Cost could represent effort required for labeling an example; this may be stochastic and known only distributionally in advance. **c) Adaptive feature selection for sparse regression:** A learner has access to a response vector and a prior distribution over a set of features, where uncertainty in features is due to measurement noise. The learner can adaptively query features in order to observe them accurately. Our objectively, namely, the squared multiple correlation or the R^2 fit, can be modeled as an adaptive weak submodular function when the measurement noise is independent across features (Fujii and Sakaue, 2019). The cost of querying a feature may be stochastic and known only distributionally in advance.

1.1. Our Contributions

1) Adaptive Weak Submodular Function + Stochastic Fractional Knapsack: We present the β-GREEDY policy for stochastic fractional knapsack. When the utility function f is adaptive weak submodular, we prove that β-GREEDY achieves an approximation guarantee of $1 - e^{-\alpha\beta}$. Here, $\alpha \in [0, 1]$ is the adaptive submodularity ratio which measures the degree to which f is submodular and $\beta \in [0, 1]$ is the approximation factor of the greedy step, which ensures that β-GREEDY is robust to approximate greedy selection and achieves a provably-good guarantee even when access to various quantities (such as priors $\mathcal{P}$ and $\mathcal{Q}$) is only approximate.

2. For any given item $e \in E$, $C(e)$ is independent of S and other item costs.

Significance: Our β-GREEDY guarantee recovers many known guarantees for greedy maximization as special cases. Previous guarantees for maximization of weak and adaptive weak submodular functions (Das and Kempe, 2018; Khanna et al., 2017; Bian et al., 2017; Fujii and Sakaue, 2019) were derived under a cardinality constraint which is the case of all item costs equal to 1. Das and Kempe (2018) proved an approximation guarantee of $1-e^{-\alpha}$ for the GREEDY policy with exact greedy selection ($\beta = 1$) for cardinality constrained maximization of a weak submodular function with submodularity ratio α. Fujii and Sakaue (2019) generalized this guarantee (with $\beta = 1$) for adaptive weak submodular functions. Golovin and Krause (2011) proved an approximation guarantee of $1 - e^{-\beta}$ for β-GREEDY under a fractional knapsack constraint with deterministic item costs and an adaptive submodular function; the latter is equivalent to an adaptive weak submodular function with $\alpha = 1$ (Fujii and Sakaue, 2019). These guarantees generalize the classic $1 - 1/e$ approximation guarantee of Nemhauser et al. (1978) for cardinality constrained maximization of submodular functions. Our β-GREEDY guarantee generalizes (Das and Kempe, 2018; Fujii and Sakaue, 2019; Golovin and Krause, 2011; Nemhauser et al., 1978) by simultaneously handling 1) item costs that are stochastic, 2) adaptive weak submodular utility function characterized by adaptive submodularity ratio $\alpha \in (0, 1]$, and 3) approximate greedy selection characterized by $\beta \in (0, 1]$.

2) Adaptive Submodular Function + Stochastic Knapsack: We present the MIX policy for stochastic knapsack which is a probabilistic mix of two greedy policies. When the utility function f is adaptive submodular, we prove that MIX achieves an approximation guarantee of $\frac{1}{6}\left(1 - e^{-\beta/4}\right)$.

Significance: Previous guarantees for stochastic knapsack problem were derived under the assumption that the utility function is *linear* (Dean et al., 2008; Bhalgat et al., 2011; Bhalgat, 2011; Halman et al., 2014; Gupta et al., 2011; Li and Yuan, 2013; Ma, 2014). Under the linearity assumption, utility $f(A)$ of an itemset A is the sum of item-wise utilities $\sum_{e\in A} f(e)$ where the utility $f(e)$ of item e is unrelated to the composition of itemset A. This assumption breaks down when the utility function is submodular or adaptive submodular. Our work presents the first guarantees for stochastic knapsack under non-linear utility functions.

3) Adaptive Submodular Function + Deterministic Knapsack: The MIX policy is parameterized by two hyper-parameters. We present two alternative parameterizations of MIX than the one used in **2)** and derive two new guarantees for adaptive submodular maximization under deterministic knapsack. The first of these is a $\frac{1}{2}\left(1 - e^{-\beta/2}\right)$ approximation guarantee, which (for the special case of deterministic knapsack) improves over the $\frac{1}{6}\left(1 - e^{-\beta/4}\right)$ approximation guarantee in **2)**. The second is a $\frac{\beta}{1+\beta}\left(1 - e^{-\beta}\right)$ approximation guarantee which improves over the first when $\beta \geq 0.377$.

Significance: A variety of guarantees exist for maximizing a submodular function under a deterministic knapsack constraint (Khuller et al., 1999; Sviridenko, 2004; Leskovec et al., 2007; Sviridenko et al., 2015; Yoshida, 2019; Calinescu et al., 2011; Cuong and Xu, 2016; Kawase et al., 2019). We develop the first guarantees for maximizing an adaptive submodular function under a deterministic knapsack constraint. Our $\frac{\beta}{1+\beta}\left(1 - e^{-\beta}\right)$ guarantee generalizes a $\frac{1}{2}(1-1/e)$ guarantee due to Leskovec et al. (2007) for deterministic submodular functions (with $\beta = 1$).

4) Differential Analysis: One of the central contributions of this paper is an innovative differential analysis technique which lies at the heart of all our approximation guarantees. In this technique, we model the execution of β-GREEDY through a continuous-time stochastic reward process involving a particle which traverses the positive x-axis and continuously collects reward at a stochastic rate. We relate the particle's reward to β-GREEDY and optimal utilities through a differential inequality. This inequality can be solved through standard Laplace transform machinery for differential equations resulting in our β-GREEDY guarantee. We combine differential analysis with other ideas to derive our MIX guarantees. In particular, we use differential analysis in conjunction with the idea of truncated costs introduced by Dean et al. (2008) and the idea of policy emulation to derive our MIX guarantee for stochastic knapsack.

1.2. Applications

β-GREEDY and MIX policies enable provably-good sequential decision making in a wide variety of applications. Examples of applications where β-GREEDY guarantee holds include all of the following. Examples of applications where MIX guarantees hold include applications below that are characterized by adaptive submodularity or submodularity.

Bipartite influence maximization and adaptive feature selection (Fujii and Sakaue, 2019) are some applications characterized by adaptive weak submodularity. Sparse regression, dictionary selection (Das and Kempe, 2018), Bayesian A-optimality in experimental design, and maximization of LPs with combinatorial constraints (Bian et al., 2017) are some applications characterized by weak submodularity. Viral marketing, sensor deployment, active learning (Asadpour et al., 2008; Golovin and Krause, 2011; Bai et al., 2006; Chen and Krause, 2013; Cuong et al., 2013, 2014; Fujii and Kashima, 2016; Golovin et al., 2010; González-Banos, 2001; Krause and Golovin, 2014; Krause et al., 2008; Sun et al., 2018), active detection (Chen et al., 2014), maximizing value of information (Krause and Guestrin, 2007; Chen et al., 2015), preference elicitation, face detection (Gabillon et al., 2014), active planning (Hollinger et al., 2013), touch based localization Javdani et al. (2013), database query discovery (Shen et al., 2014), community sensing (Singla and Krause, 2013), and sequential Bayesian search Wen et al. (2013) are some applications characterized by adaptive submodularity. Influence maximization in social networks (Kempe et al., 2003; Leskovec et al., 2007), information gain maximization in graphical models (Krause and Guestrin, 2005), multi-document summarization (Lin and Bilmes, 2011), and image collection summarization (Tschiatschek et al., 2014) are some applications characterized by submodularity.

2. Preliminaries

Adaptive Stochastic Optimization: Let E denote the ground set of items. Each $e \in E$ is endowed with a random state $S(e) \in \mathcal{O}$. The state-relation S which maps each item to its state is distributed according to a known prior $\mathcal{P}$. Utility function $F : 2^E \times \mathcal{O}^E \to \mathbb{R}$ maps each (itemset, state-relation) pair to a utility value. For a given $A \subseteq E$, $F(A, S)$ is a random variable (r.v.) which is fixed when realization of S is given. Item states *need not* be mutually independent. Each $e \in E$ is endowed with a random cost $C(e) \in \mathbb{R}_{>0}$. The cost-vector C which maps each item to its cost is distributed according to a known prior $\mathcal{Q}$. S and C are independent and item costs are mutually independent. For ease of exposition, we will assume $\mathcal{P}$ and $\mathcal{Q}$ are discrete distributions.

A policy is *singly adaptive* if it selects items based on states of items selected in previous iterations. A policy is *doubly adaptive* if it selects items based on both states and costs of items selected in previous iterations. A *truncated* policy π_t is derived from a doubly adaptive policy π_d and is parameterized by a budget; π_t executes π_d until its total cost exceeds the budget at which point π_t discards the last item selected and terminates immediately. Let $\mathcal{SP}$, $\mathcal{DP}$ and $\mathcal{TP}$ denote the set of singly adaptive, doubly adaptive and truncated policies respectively.

Problem Formulation: For an r.v. X, let $supp(X)$ denote the support of X. Let $A(\pi)$ denote the set of items selected (and not discarded) by policy π. Define the stochastic function $f : 2^E \to \mathbb{R}_{\geq 0}$ s.t. $\forall A \subseteq E$, $f(A) = F(A, S)$. For $A \subseteq E$, let $C(A) \triangleq \sum_{e \in A} C(e)$ denote the total cost of

A. Define the r.v.s $f(\pi) \triangleq f(A(\pi))$ and $C(\pi) \triangleq C(A(\pi))$. The integral and fractional stochastic knapsack problems are stated formally in ($\mathcal{I}$) and ($\mathcal{F}$) respectively.

$$\max_{\pi \in \mathcal{TP}} \ \mathbf{E}[f(\pi)] \quad \text{subject to} \quad \Pr[C(\pi) \leq B] = 1 \tag{$\mathcal{I}$}$$

$$\max_{\pi \in \mathcal{DP}} \ \mathbf{E}[f(\pi)] \quad \text{subject to} \quad \mathbf{E}[C(\pi)|S = \phi] \leq B, \ \ \forall \phi \in supp(S) \tag{$\mathcal{F}$}$$

Note that ($\mathcal{I}$) is defined over truncated policies. This is standard practice in stochastic knapsack literature since an optimal doubly adaptive policy which is feasible for ($\mathcal{I}$) can be arbitrarily worse than an optimal truncated policy which is feasible for ($\mathcal{I}$). Also note the use of conditional expectation in ($\mathcal{F}$). This serves two purposes. From a theoretical perspective, our analysis of ($\mathcal{F}$) acts as a key building block for our analysis of ($\mathcal{I}$). From a practical perspective, it enforces a useful fairness condition in budget allocation which states that budget for certain states cannot be increased at the expense of others in order to maximize expected utility.

Adaptive Functions: For $E_1 \subseteq E$ and $e \in E \setminus E_1$, let the r.v. $\Delta f_{E_1}(e) \triangleq f(E_1 \cup \{e\}) - f(E_1)$ denote the marginal utility of e w.r.t. E_1. Analogously, for $E_2 \subseteq E$, define $\Delta f_{E_1}(E_2) \triangleq f(E_1 \cup E_2) - f(E_1)$. A subrealization φ is a relation which maps each item in a subset of E to a state; this subset is denoted by $dom(\varphi)$. '$\varphi \subseteq S$' denotes the event that φ is consistent with S, i.e., $\forall e \in dom(\varphi) : S(e) = \varphi(e)$. Define the set of subrealizations $\mathcal{S} \triangleq \{\varphi | \Pr[\varphi \subseteq S] > 0\}$. Given a subrealization φ and $e \in E \setminus dom(\varphi)$, let the conditional r.v. $\Delta f(e|\varphi) \triangleq [\Delta f_{dom(\varphi)}(e)|\varphi \subseteq S]$ denote the marginal utility of e w.r.t. φ. Analogously, for $E_2 \subseteq E$, define $\Delta f(E_1|\varphi) \triangleq [\Delta f_{dom(\varphi)}(E_2)|\varphi \subseteq S]$. For subrealizations φ_1, φ_2, we write $\varphi_1 \subseteq \varphi_2$ to denote the following: '$dom(\varphi_1) \subseteq dom(\varphi_2)$ and $\forall e \in dom(\varphi_1) : \varphi_1(e) = \varphi_2(e)$'.

Definition 1 (Golovin and Krause (2011)) *Function f is adaptive submodular if it satisfies (1), (2), and (3).*

$$\mathbf{E}[f(\emptyset)] = 0 \tag{1}$$

$$\forall \varphi \in \mathcal{S}, \forall e \in E \setminus dom(\varphi) : \mathbf{E}[\Delta f(e|\varphi)] \geq 0 \tag{2}$$

$$\forall \varphi_2 \in \mathcal{S}, \forall \varphi_1 \subseteq \varphi_2, \forall e \in E \setminus dom(\varphi_2) : \mathbf{E}[\Delta f(e|\varphi_1)] \geq \mathbf{E}[\Delta f(e|\varphi_2)] \tag{3}$$

Let the conditional r.v. $\Delta f(\pi|\varphi) \triangleq \Delta f(A(\pi)|\varphi)$ denote the marginal utility of $A(\pi)$ w.r.t. φ.

Definition 2 (Fujii and Sakaue (2019)) *Function f is adaptive weak submodular with adaptive submodularity ratio $\alpha \in [0, 1]$ if it satisfies (1), (2) and (4).*

$$\forall \varphi \in \mathcal{S}, \forall \pi \in \mathcal{SP} : \ \alpha \mathbf{E}[\Delta f(\pi|\varphi)] \leq \sum_{e \in E \setminus dom(\varphi)} \Pr[e \in A(\pi)|\varphi \subseteq S] \mathbf{E}[\Delta f(e|\varphi)] \tag{4}$$

Fujii and Sakaue (2019) showed that (4) holds with $\alpha = 1$ if and only if f is adaptive submodular. Adaptive functions are stochastic generalizations of their deterministic counterparts. In particular, we can model deterministic utility and cost functions by letting $\mathcal{P}$ and $\mathcal{Q}$ be degenerate distributions over support of size one.

3. β-GREEDY Policy

The β-GREEDY policy is illustrated in Algorithm 1. It is a singly adaptive policy whose candidate e_φ during the current iteration is the item which maximizes the ratio of expected marginal utility over expected cost, conditional on the current subrealization φ. This candidate selection step which is implemented by the $\arg\max$ function in Algorithm 1 can be β-approximate. In order to simplify our analysis, we require the $\arg\max$ function to be deterministic: i.e., its output e_φ is unique given its input φ. If its residual capacity (current value of b) is insufficient, it uses randomized rounding to either select e_φ or terminate. The rounding probability is s.t. the expected cost of β-GREEDY equals its budget (initial value of b), conditional on any realization of S. Initializing $b \leftarrow B$ ensures that β-GREEDY is a feasible policy for instance $\mathcal{F}$. Theorem 3 states our approximation guarantee for β-GREEDY which we prove in Section 3.1.

Algorithm 1 β-GREEDY policy

Initialize: $b \leftarrow$ budget, $\varphi \leftarrow \emptyset$, *stop* $\leftarrow$ *false*
repeat
  $e_\varphi \leftarrow \arg\max_{e \in E \setminus dom(\varphi)} \mathbf{E}[\Delta f(e|\varphi)]/\mathbf{E}[C(e)]$
  if $b < \mathbf{E}[C(e_\varphi)]$ **then** *stop* $\leftarrow$ *true* w.p. $(1 - b/\mathbf{E}[C(e)])$
  if $\neg$*stop* **then** select e_φ; update $\varphi \leftarrow \varphi \cup \{(e_\varphi, S(e_\varphi))\}$ and $b \leftarrow b - \mathbf{E}[C(e)]$
until $dom(\varphi) = E$ *or* $b \leq 0$ *or stop*

3.1. Analysis of β-GREEDY

Theorem 3 *Let f be an adaptive weak submodular function satisfying (1), (2), and (4). Let π^* denote an optimal doubly adaptive policy for stochastic fractional knapsack instance ($\mathcal{F}$). Let σ denote the β-GREEDY policy initialized with budget $b \leftarrow B$. Then, $\mathbf{E}[f(\sigma)] \geq \left(1 - e^{-\alpha\beta}\right) \mathbf{E}[f(\pi^*)]$ where the expectations are over the distributions of S, C and internal random choices of σ and π^*.*

Proof outline: Our proof strategy for Theorem 3 is as follows. We will introduce a continuous time stochastic process which models the execution of σ using a particle which accumululates reward over time at a stochastic rate. We relate the rate function of the particle to the expected utility of σ during various stages of its execution in Lemma 5. We build on this lemma to relate the rate function to the expected utility of π^* using a differential inequality in Lemma 7. The solution to this differential inequality proves Theorem 3. With the exception of two supporting lemmas (4 and 6) whose proofs are in Appendix A, the complete proof of Theorem 3 is presented below.

Complete proof: Let σ^∞ denote the β-GREEDY policy initialized with budget $b \leftarrow \mathbf{E}[C(E)]$. This ensures that σ^∞ will terminate only after selecting all items in E. We model the execution of σ^∞ using a continuous-time stochastic reward process ν, which is also referred to as the rate function. This process is best imagined through a particle that starts at the origin at time $t = 0$, and moves along the positive x-axis at unit speed. Let $C(\varphi) \triangleq C(dom(\varphi))$. For $x \in [0, \mathbf{E}[C(E)])$, define:

$$\theta(x) \triangleq \arg\max_{\{\varphi | \varphi \text{ is the subrealization at the start of some iteration of } \sigma^\infty \text{ s.t. } \mathbf{E}[C(\varphi)] \leq x\}} \mathbf{E}[C(\varphi)] \tag{5}$$

$$\nu(x) \triangleq \Delta f(e_{\theta(x)} | \theta(x)) / \mathbf{E}[C(e_{\theta(x)})] \tag{6}$$

$\theta(x)$ is the last subrealization observed by σ^∞ before its expected cost exceeded x. Recall from Algorithm 1 that $e_{\theta(x)}$ is the item selected by σ^∞ during the iteration at the start of which its subrealization is $\theta(x)$. $\nu(x)$ is the *rate* at which the particle accumulates reward when it reaches position x. Both $\theta(x)$ and $\nu(x)$ are r.v.s and the latter is fixed when the former is given. Suppose φ is the subrealization of σ^∞ at the start of an iteration. (5) and (6) imply $\forall x \in [\mathbf{E}[C(\varphi)], \mathbf{E}[C(\varphi)] + \mathbf{E}[C(e_\varphi)])$, $\theta(x) = \varphi$ and $\mathbf{E}[\nu(x)|\theta(x) = \varphi] = \mathbf{E}[\Delta f(e_\varphi|\varphi)]/\mathbf{E}[C(e_\varphi)]$. Lemma 4 follows from the fact that $\arg\max$ function in Algorithm 1 is deterministic. Its proof is in Appendix A.

Lemma 4 *$\forall x \in [0, \mathbf{E}[C(E)]), \forall \varphi \in supp(\theta(x))$, events '$\theta(x) = \varphi$' and '$\varphi \subseteq S$' are equivalent.*

Lemma 5 *The rate function ν, the function θ, and the utility of σ are related as follows.*

$$\mathbf{E}[f(\sigma)] = \int_0^B \mathbf{E}[\nu(x)]dx \tag{7}$$

$$\forall x \in [0, \mathbf{E}[C(E)]): \ \mathbf{E}[f(\theta(x))] \leq \int_0^x \mathbf{E}[\nu(t)]dt \tag{8}$$

Proof Given $S = \phi$, the sequence of subrealizations seen by σ^∞ at the start of its iterations is fixed. Let this be $\varphi_1 = \emptyset, \varphi_2, \ldots, \varphi_{|E|}$. Let k be s.t. $\theta(B) = \varphi_k$. During iterations $i \in \{1, \ldots k-1\}$, from the definition of ν in (6), the expected increment in σ's utility is as follows:

$$\mathbf{E}[\Delta f_{\varphi_i}(e_{\varphi_i})|S = \phi] = \int_{\mathbf{E}[C(\varphi_i)]}^{\mathbf{E}[C(\varphi_{i+1})]} [\nu(x)|S = \phi]dx \tag{9}$$

During iteration k, due to randomized rounding, the expected increment in σ's utility is as follows:

$$\mathbf{E}[f(\sigma) - f(\varphi_k)|S = \phi] = \frac{(B - \mathbf{E}[C(\varphi_k)])\mathbf{E}[\Delta f_{\varphi_k}(e_{\varphi_k})|S = \phi]}{\mathbf{E}[C(e_{\varphi_k})]} = \int_{\mathbf{E}[C(\varphi_k)]}^B [\nu(x)|S = \phi]dx \tag{10}$$

In (9) and (10), note that the conditional r.v. $[\nu(x)|S = \phi]$ is a fixed quantity and therefore equals the conditional expectation $\mathbf{E}[\nu(x)|S = \phi]$. Summing these increments, we have:

$$\mathbf{E}[f(\sigma) - f(\emptyset)|S = \phi] = \int_0^B \mathbf{E}[\nu(x)|S = \phi]dx \tag{11}$$

Deconditioning (11) and combining with (1) yields (7). From (9), we also have:

$$\mathbf{E}[f(\theta(x)) - f(\emptyset)|\theta(x) = \varphi] = \int_0^{\mathbf{E}[C(\varphi)]} \mathbf{E}[\nu(t)|\theta(x) = \varphi]dt \leq$$

$$\int_0^{\mathbf{E}[C(\varphi)]} \mathbf{E}[\nu(t)|\theta(x) = \varphi]dt + \frac{(x - \mathbf{E}[C(\varphi)])\mathbf{E}[\Delta f(e_\varphi|\varphi)]}{\mathbf{E}[C(e_\varphi)]} = \int_0^x \mathbf{E}[\nu(t)|\theta(x) = \varphi]dt \tag{12}$$

The inequality in (12) is due to (2). Deconditioning (12) and combining with (1) yields (8). ■

Lemma 6 is a consequence of (2). Its proof is in Appendix A.

Lemma 6 $\forall \varphi \in \mathcal{S}, \forall \pi \in \mathcal{DP}, \forall x \in [0, \mathbf{E}[C(E)]): \ \mathbf{E}[f(\pi)] \leq \mathbf{E}[f(A(\pi) \cup dom(\theta(x)))]$

Lemma 7 *The rate function ν is related to the utility of π^* as follows.*

$$\forall x \in [0, B), \ \mathbf{E}[f(\pi^*)] \leq \int_0^x \mathbf{E}[\nu(t)]dt + \frac{B\mathbf{E}[\nu(x)]}{\alpha\beta} \tag{13}$$

Proof

$$\mathbf{E}[f(A(\pi^*) \cup dom(\theta(x)))|\theta(x) = \varphi] = \mathbf{E}[f(\varphi)] + \mathbf{E}[\Delta f_\varphi(A(\pi^*))]$$

$$\leq \mathbf{E}[f(\varphi)] + \frac{\sum_{e \in E \setminus dom(\varphi)} \Pr[e \in A(\pi^*)|\varphi \subseteq S]\mathbf{E}[\Delta f_\varphi(e)]}{\alpha} \tag{14}$$

$$\leq \mathbf{E}[f(\varphi)] + \frac{\mathbf{E}[\Delta f_\varphi(e_\varphi)]}{\alpha\beta\mathbf{E}[C(e_\varphi)]} \sum_{e \in E \setminus dom(\varphi)} \Pr[e \in A(\pi^*)|\varphi \subseteq S]\mathbf{E}[C(e)] \tag{15}$$

$$\leq \mathbf{E}[f(\theta(x))|\theta(x) = \varphi] + \frac{\mathbf{E}[\nu(x)|\theta(x) = \varphi]B}{\alpha\beta} \tag{16}$$

(14) follows from (4) and (15) follows from the fact that the $\arg\max$ function in Algorithm 1 is β-approximate. (16) follows Lemma 4 and from the fact that π^* is a feasible policy for ($\mathcal{F}$). Deconditioning (16) and combining it with (8) and Lemma 6 proves this lemma. ■

Proof [Theorem 3] Let $G : [0, B] \to \mathbb{R}$ be a continuously differentiable function which satisfies the ordinary differential equation (ODE) $\mathbf{E}[f(\pi^*)] = G(t) + \frac{B}{\alpha\beta}\frac{dG}{dt}(t)$ and the boundary condition $F(0) = 0$. The solution to this differential equation which can be derived through standard Laplace transform based techniques is $G(t) = \left(1 - e^{-\frac{\alpha\beta t}{B}}\right)\mathbf{E}[f(\pi^*)]$. Let the particle referred to until now be P_1. Let P_2 be another particle which also starts at the origin at time $t = 0$ and moves on the positive x-axis at unit speed. Thus, at every time instant, P_1 and P_2 have the same location. P_2 accumulates reward at point x at rate $\frac{dG}{dt}(t)\big|_{t=x}$. (13) and the above ODE imply that at any point $x \in [0, B)$, if the total expected reward of P_1 equals the total reward of P_2, then P_1's expected rate of reward is no less than that of P_2. Hence at any point $x \in [0, B)$, P_1 accumulates as much reward as P_2 in expectation: i.e.,

$$\forall x \in [0, B) : \int_0^x \mathbf{E}[\nu(t)]dt \geq G(x) \implies \int_0^B \mathbf{E}[\nu(t)]dt = \lim_{x \to B} \int_0^x \mathbf{E}[\nu(t)]dt \geq G(B) \tag{17}$$

Theorem 3 now follows from (7) and (17). ■

4. MIX Policy

The MIX policy is illustrated in Algorithm 2. It is a truncated policy which is parameterized by a probability value p_ℓ and a cost threshold λ. MIX probabilistically chooses between two other policies, namely LIGHT and SINGLETON and executes the chosen policy. For item e, define the *truncated cost* of e as the r.v. $K(e) \triangleq \min\{C(e), B\}$. Item e is called *light* if $\mathbf{E}[K(e)] \leq \lambda$ and *heavy* otherwise. LIGHT selects light items iteratively based on the ratio of expected marginal utility over expected truncated cost. If LIGHT reaches an iteration in which its total cost exceeds its budget (value of b), it discards its final selection and terminates. SINGLETON selects a single item which maximizes $\Pr[C(e) \leq B]\mathbf{E}[f(e)]$. If the item's cost exceeds B, SINGLETON discards it. Theorems 8, 13, and 14 describe how to choose the values of p_ℓ and λ in order to derive our guarantees. Setting $b \leftarrow B$ during the initialization step of LIGHT ensures that MIX is a feasible policy for instance $\mathcal{I}$.

Algorithm 2 MIX(p_ℓ, λ) policy

Execute LIGHT with probability p_ℓ and execute SINGLETON with probability $1 - p_\ell$

Policy LIGHT

  Initialize: $b \leftarrow$ budget, $\varphi \leftarrow \emptyset$, $stop \leftarrow false$, $L \leftarrow \{e|\mathbf{E}[K(e)] \le \lambda\}$

  repeat

    select $e_\varphi \leftarrow \arg\max_{e \in L \setminus dom(\varphi)} \mathbf{E}[\Delta f(e|\varphi)]/\mathbf{E}[K(e)]$

    if $C(\varphi) + C(e_\varphi) \le b$ **then** $\varphi \leftarrow \varphi \cup \{(e_\varphi, S(e_\varphi))\}$

    else discard e_φ; $stop \leftarrow true$

  until $dom(\varphi) = L$ *or stop*

Policy SINGLETON

  select $e^* \leftarrow \arg\max_{e \in E} \Pr[C(e) \le B]\mathbf{E}[f(e)]$

  if $C(e^*) > B$ **then** discard e^*

4.1. Analysis of MIX

Let $\pi^\dagger$ denote an optimal truncated policy for instance $\mathcal{I}$ and let τ denote the MIX policy. Let τ_ℓ and τ_s denote the LIGHT and SINGLETON policies respectively. Theorem 8 states the approximation guarantee of MIX for stochastic knapsack which we prove in Section 4.1.1.

4.1.1. STOCHASTIC KNAPSACK

Theorem 8 *Let f be an adaptive submodular function. Set $p_\ell = \frac{1}{2}$ and $\lambda = B/6$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau)] \ge \frac{1}{6}\left(1 - e^{-\beta/4}\right)\mathbf{E}[f(\pi^\dagger)]$.

Proof outline: We split the expected utility of $\pi^\dagger$ into two portions: the portion contributed by light items (R) and the remaining portion contributed by heavy items ($\bar{R}$). We relate the expected utility of τ_ℓ to R in Lemma 11; this lemma is built on several ideas, including the notion of policy emulation which emulates a truncated policy using a singly adaptive policy, the notion of composite policy which combines elements of greedy and optimal policies, an upper bound on the truncated cost of $\pi^\dagger$ and the doubly adaptive policy from which it is derived, and the ideas of differential analysis introduced in Section 3.1. We relate the expected utility of τ_s to $\bar{R}$ in Lemma 12 by arguing that $\pi^\dagger$ cannot choose too many heavy items in expectation. Combining these lemmas yields Theorem 8. With the exception of a supporting lemma (10) which is presented in Appendix A, the complete proof of Theorem 8 is presented below.

Complete proof: Construct a fractional stochastic knapsack instance $\mathcal{L}$ whose utility function is f, whose cost function is K, whose capacity is $2B$, and whose ground set of items is the subset of light items L in Algorithm 2: i.e., items e s.t. $\mathbf{E}[K(e)] \le \lambda = B/6$. Let $\bar{\sigma}$ denote the β-GREEDY policy applied to $\mathcal{L}$ and initialized with budget $b \leftarrow 2B$. Let $\bar{\sigma}^\infty$ denote the β-GREEDY policy applied to $\mathcal{L}$ and initialized with budget $b \leftarrow \mathbf{E}[K(L)]$. This ensures that $\bar{\sigma}^\infty$ will terminate only after selecting all items in L. Analogous to (5) and (6), $\forall x \in [0, \mathbf{E}[K(L)])$, define:

$$\bar{\theta}(x) \triangleq \underset{\{\varphi|\varphi \text{ is the subrealization at the start of some iteration of } \bar{\sigma}^\infty \text{ s.t. } \mathbf{E}[K(\varphi)] \le x\}}{\arg\max} \mathbf{E}[K(\varphi)] \tag{18}$$

$$\bar{\nu}(x) \triangleq \Delta f(e_{\bar{\theta}(x)}|\bar{\theta}(x))/\mathbf{E}[K(e_{\bar{\theta}(x)})] \tag{19}$$

Lemma 9 *Set $\lambda = B/6$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau_\ell)] \ge \frac{1}{3}\int_0^{B/2} \mathbf{E}[\bar{\nu}(x)]dx$.

Proof Conditioning on the event $S = \phi$, suppose the sequence of subrealizations seen by $\bar{\sigma}^\infty$ equals $\varphi_1, \ldots, \varphi_{|L|+1}$ and the sequence of items selected by it equals $e_{\varphi_1}, \ldots, e_{\varphi_{|L|}}$. Item e_i will be included in τ_ℓ if the total cost of items $e_1, \ldots, e_i$ does not exceed B. Hence, we have:

$$[f(\tau_\ell)|S = \phi] = \sum_{i=1}^{|L|} \left[\Delta f(e_{\varphi_i}|\varphi_i)\mathbf{1}_{C(\varphi_i)+C(e_{\varphi_i})\leq B}\middle|S = \phi\right]$$

$$= \sum_{i=1}^{|L|} \left[\bar{\nu}(\mathbf{E}[K(\varphi_i)])\mathbf{E}[K(e_{\varphi_i})]\mathbf{1}_{C(\varphi_i)+C(e_{\varphi_i})\leq B}\middle|S = \phi\right]$$

$$\geq \sum_{i|\mathbf{E}[K(\varphi_i)]\leq B/2} \left[\bar{\nu}(\mathbf{E}[K(\varphi_i)])\mathbf{E}[K(e_{\varphi_i})]\mathbf{1}_{C(\varphi_i)+C(e_{\varphi_i})\leq B}\middle|S = \phi\right] \implies$$

$$\mathbf{E}[f(\tau_\ell)|S = \phi] \geq \sum_{i|\mathbf{E}[K(\varphi_i)]\leq B/2} \mathbf{E}\left[\bar{\nu}(\mathbf{E}[K(\varphi_i)])\middle|S = \phi\right]\mathbf{E}[K(e_{\varphi_i})]\Pr[C(\varphi_{i+1}) \leq B] \tag{20}$$

$$\geq \sum_{i|\mathbf{E}[K(\varphi_i)]\leq B/2} \mathbf{E}\left[\bar{\nu}(\mathbf{E}[K(\varphi_i)])\middle|S = \phi\right]\mathbf{E}[K(e_{\varphi_i})](1 - \mathbf{E}[K(\varphi_{i+1})]/B) \tag{21}$$

$$\geq \sum_{i|\mathbf{E}[K(\varphi_i)]\leq B/2} \mathbf{E}\left[\bar{\nu}(\mathbf{E}[K(\varphi_i)])\middle|S = \phi\right]\mathbf{E}[K(e_{\varphi_i})]\left(1 - \frac{4B}{6B}\right) \tag{22}$$

$$\geq \frac{1}{3}\int_0^{B/2} \mathbf{E}[\bar{\nu}(x)|S = \phi]dx \implies \mathbf{E}[f(\tau_\ell)] \geq \frac{1}{3}\int_0^{B/2} \mathbf{E}[\bar{\nu}(x)]dx \tag{23}$$

(20) follows due to r.v.s S and C being independent. (21) follows from a modified version of Markov's inequality (24). Its derivation below is from Dean et al. (2008).

$$\Pr[C(E_1) \geq B] \leq \mathbf{E}[\min\{C(E_1)/B, 1\}] \leq \mathbf{E}[\sum_{e\in E_1} \min\{C(e)/B, 1\}] = \mathbf{E}[K(E_1)]/B$$

$$\implies \forall E_1 \subseteq E,\ \Pr[C(E_1) \leq B] \geq 1 - \mathbf{E}[K(E_1)]/B \tag{24}$$

(22) holds because $\mathbf{E}[K(\varphi_{i+1})] = \mathbf{E}[K(\varphi_i)] + \mathbf{E}[K(e_{\varphi_i})] \leq B/2 + B/6$. (23) follows from the first inequality in (17). Hence, the lemma is proved. ■

We will emulate $\pi^\dagger$ using a singly adaptive policy $\pi^\ddagger$ as follows: $\pi^\ddagger$ selects items using the same logic used by $\pi^\dagger$ and observes the state $S(e)$ of each selected item e; but instead of observing the cost $C(e)$, it samples a value from the distribution of $C(e)$ and uses this sampled value in its execution. By construction, for any possible value ϕ of S, the distribution of the conditional r.v. $[A(\pi^\ddagger)|S = \phi]$ is identical to that of $[A(\pi^\dagger)|S = \phi]$. Let ϑ_i be an r.v. which denotes the subrealization of $\pi^\ddagger$ at the start of iteration i. For a fixed $e \in E$, let $e_i \leftarrow e$ denote the event that $\pi^\ddagger$ selected item e during iteration i. Let i^* denote the maximum value of i s.t. the probability of $\pi^\ddagger$ terminating in iteration i is positive. We prove Lemma 10 in Appendix A. The idea behind this lemma is that $g(e, \pi^\ddagger)$ is the contribution of item e to the expected utility of $\pi^\ddagger$, R is the contribution due to light items, and $\bar{R}$ is the contribution due to items that are heavy.

Lemma 10 *Define* $g(e, \pi^\ddagger) \triangleq \sum_{i=1}^{i^*} \sum_{\vartheta\in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e|\vartheta_i = \vartheta]\mathbf{E}[\Delta f(e|\vartheta)]$. *Then,* $\mathbf{E}[f(\pi^\ddagger)] = R + \bar{R}$, *where* $R \triangleq \sum_{e\in L} g(e, \pi^\ddagger)$ *and* $\bar{R} \triangleq \sum_{e\in E\setminus L} g(e, \pi^\ddagger)$.

Lemma 11 *Set $\lambda = B/6$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau_\ell)] \geq \frac{1}{3}(1 - e^{-\beta/4})R$.

Proof Let π^+ denote the doubly adaptive policy from which $\pi^\dagger$ is derived. For an itemset E_1 and policy π, define $K(E_1) \triangleq \sum_{e \in E_1} K(e)$ and $K(\pi) \triangleq K(A(\pi))$. By definition of truncated costs and policies, we have $K(\pi^\dagger) \leq C(\pi^\dagger) \leq B$ w.p. 1. Since π^+ selects all items in $\pi^\dagger$ and utmost one more item, and since $K(e) \leq B$ for every item w.p. 1, (25) holds.

$$\begin{aligned}\mathbf{E}[K(\pi^\ddagger)|S = \phi] &= \sum_{e \in E} \Pr[e \in A(\pi^\ddagger)]\mathbf{E}[K(e)] = \sum_{e \in E} \Pr[e \in A(\pi^\dagger)]\mathbf{E}[K(e)] \\ &\leq \sum_{e \in E} \Pr[e \in A(\pi^+)]\mathbf{E}[K(e)] = \mathbf{E}[K(\pi^+)|S = \phi] \leq 2B \end{aligned} \tag{25}$$

Fix $x \in [0, \mathbf{E}[K(L)])$. We will construct two distinct composite policies $\bar{\theta} \oplus \pi^\ddagger$ and $\pi^\ddagger \oplus \bar{\theta}$ as follows. Both these policies have two stages. $\bar{\theta} \oplus \pi^\ddagger$, in its first stage, selects items in the same manner as $\bar{\sigma}^\infty$ until it reaches an iteration at the start of which its subrealization $\varphi \in supp(\theta(x))$. At this point, it shifts to the second stage. In this stage, it selects items in the same manner as $\pi^\ddagger$. Every $e \in A(\pi^\ddagger)$ will be a candidate during some iteration in the second stage; if this candidate is already selected by $\bar{\theta} \oplus \pi^\ddagger$ in the first stage, there is nothing more to be done in this iteration, and $\bar{\theta} \oplus \pi^\ddagger$ moves on to the next iteration. $\bar{\theta} \oplus \pi^\ddagger$ terminates as soon as all items in $dom(\bar{\theta}(x)) \cup A(\pi^\ddagger)$ are selected.

$\pi^\ddagger \oplus \bar{\theta}$, in its first stage, selects items in the same manner as $\pi^\ddagger$ until $\pi^\ddagger$ terminates at which point, it shifts to the second stage. In this stage, it selects items in the same manner as $\bar{\sigma}^\infty$ until it reaches an iteration at the start of which its subrealization φ (which contains item states from both stages 1 and 2) satisfies: $\exists \varphi' \in supp(\bar{\theta}(x))$ s.t. $\varphi' \subseteq \varphi$. Every $e \in \bar{\theta}(x)$ will be a candidate during some iteration in the second stage; if this candidate is already selected by $\pi^\ddagger \oplus \bar{\theta}$ in the first stage, there is nothing more to be done in this iteration, and $\pi^\ddagger \oplus \bar{\theta}$ moves on to the next iteration. $\pi^\ddagger \oplus \bar{\theta}$ terminates as soon as all items in $dom(\bar{\theta}(x)) \cup A(\pi^\ddagger)$ are selected. It should be clear from the description of these two composite policies that the itemsets selected by them are identical. We now have,

$$R = \mathbf{E}[f(\pi^\ddagger)] - \bar{R} \leq \mathbf{E}[f(\pi^\ddagger)] \leq \mathbf{E}[f(\pi^\ddagger \oplus \bar{\theta})] = \mathbf{E}[f(\bar{\theta} \oplus \pi^\ddagger)] \tag{26}$$

$$\leq \mathbf{E}[f(\bar{\theta})] + \frac{2B}{\beta}\mathbf{E}[\bar{\nu}(x)] \leq \int_0^x \mathbf{E}[\bar{\nu}(t)]dt + \frac{2B}{\beta}\mathbf{E}[\bar{\nu}(x)] \tag{27}$$

The second inequality in (26) follows from (3) and its formal proof is along the same lines as that of Lemma 6 in Appendix A. The first inequality in (27) follows from (25) along with the same reasoning as in the proof of Lemma (7). The second inequality follows from (8). We now have the differential inequality $\int_0^x \mathbf{E}[\bar{\nu}(t)]dt + \frac{2B\mathbf{E}[\bar{\nu}(x)]}{\beta} \geq R$. Using similar arguments as in the proof of Theorem 3 which resulted in (17), we have the following solution to this differential inequality: $\int_0^x \mathbf{E}[\bar{\nu}(t)]dt \geq \left(1 - e^{-\beta x/2B}\right) R$. Letting $x = B/2$ in this solution and combining it with Lemma 9 proves this lemma. ■

Lemma 12 *Set $\lambda = B/6$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau_s)] \geq \frac{1}{3}(1 - e^{-\beta/4})\bar{R}$.

Proof We have:

$$\sum_{e\in E\setminus L} \Pr[e \in A(\pi^+)]\mathbf{E}[K(e)] \leq 2B \implies \sum_{e\in E\setminus L} \Pr[e \in A(\pi^+)] \leq 12 \tag{28}$$

$$\left(\sum_{e\in E\setminus L} \Pr[e \in A(\pi^\dagger)]\mathbf{E}[f(e)]\right) \leq \sum_{e\in E\setminus L} \Pr[e \in A(\pi^+)] \Pr[C(e) \leq B]\mathbf{E}[f(e)] \tag{29}$$

$$\leq \sum_{e\in E\setminus L} \Pr[e \in A(\pi^+)] \left(\max_{e'\in E} \Pr[C(e') \leq B]\mathbf{E}[f(e')]\right) \leq 12\mathbf{E}[f(\tau_s)]/\beta \tag{30}$$

$$\implies \mathbf{E}[f(\tau_s)] \geq \frac{\beta}{12}\left(\sum_{e\in E\setminus L} \Pr[e \in A(\pi^\dagger)]\mathbf{E}[f(e)]\right)$$

$$\geq \frac{1}{3}\left(1-e^{-\beta/4}\right)\left(\sum_{e\in E\setminus L} \Pr[e \in A(\pi^\dagger)]\mathbf{E}[f(e)]\right) \geq \frac{1}{3}\left(1-e^{-\beta/4}\right)\bar{R} \tag{31}$$

(28) follows from the fact that for every item e that is heavy, $\mathbf{E}[K(e)] > \lambda = B/6$. (29) follows from the fact an item $e \in A(\pi^\dagger)$ only if the independent events $e \in A(\pi^+)$ and $C(e) \leq B$ occur. (30) follows from β-approximation of the $\arg\max$ function in SINGLETON. (31) follows from the fact that $x \geq 1 - e^{-x}$. This completes the proof. ∎

Proof [Theorem 8] Since $p_\ell = 1/2$, $\mathbf{E}[f(\tau)] = (\mathbf{E}[f(\tau_\ell)] + \mathbf{E}[f(\tau_s)])/2$. The theorem now follows by combining this equation with Lemmas 11 and 12. ∎

4.1.2. DETERMINISTIC KNAPSACK

Theorems 13 and 14 state two alternative approximation guarantees of MIX for deterministic knapsack using parameterizations that are different from the one used in Theorem 8. For deterministic knapsack, the guarantee in Theorem 13 is an improvement over the one in Theorem 8. The guarantee in Theorem 14 is an improvement over the one in Theorem 13 when $\beta \geq 0.377$.

Theorem 13 *Let f be an adaptive submodular function and let the cost function C be deterministic. Set $p_\ell = 1/2$ and $\lambda = B/2$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau)] \geq \frac{1}{2}\left(1-e^{-\beta/2}\right)\mathbf{E}[f(\pi^\dagger)]$.

Proof outline: The proof of this theorem follows along the same lines as that of the proof of Theorem 8, but with several simplifications due to the fact item costs are deterministic. The complete proof of Theorem 13 is in Appendix A.

Theorem 14 *Let f be an adaptive submodular function and let item costs be deterministic. Set $p_\ell = \frac{\beta}{1+\beta}$ and $\lambda = B$ in the* MIX *policy. Then,* $\mathbf{E}[f(\tau)] \geq \frac{\beta}{1+\beta}\left(1-e^{-\beta}\right)\mathbf{E}[f(\pi^\dagger)]$.

Proof outline: Taking advantage of the fact that item costs are deterministic, we first argue that $\pi^\dagger$ can be implemented as a singly adaptive policy which is feasible for $\mathcal{F}$. This ensures that β-GREEDY achieves a significant portion of the expected utility of $\pi^\dagger$. τ_ℓ selects all items selected by β-GREEDY with the potential exception of the last item; we argue that τ_s compensates for this loss to a significant extent, there by yielding Theorem 14. The complete proof of Theorem 14 is in Appendix A.

References

Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi. Stochastic submodular maximization. In *WINE*, volume 5385 of *Lecture Notes in Computer Science*, pages 477–489. Springer, 2008.

Xiaole Bai, Santosh Kumar, Dong Xuan, Ziqiu Yun, and Ten H. Lai. Deploying wireless sensors to achieve both coverage and connectivity. In *Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc'06, pages 131–142, New York, NY, USA, 2006. ACM. ISBN 1-59593-368-9. doi: 10.1145/1132905.1132921.

Anand Bhalgat. A $(2 + \epsilon)$-approximation algorithm for the stochastic knapsack problem. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.222.7341&rep=rep1&type=pdf, 2011.

Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. Improved approximation results for stochastic knapsack problems. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 1647–1665, Philadelphia, PA, USA, 2011. Society for Industrial and Applied Mathematics. URL http://dl.acm.org/citation.cfm?id=2133036.2133163.

Andrew An Bian, Joachim M. Buhmann, Andreas Krause, and Sebastian Tschiatschek. Guarantees for greedy maximization of non-submodular functions with applications. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pages 498–507. JMLR.org, 2017.

Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766, December 2011. ISSN 0097-5397. doi: 10.1137/080733991. URL http://dx.doi.org/10.1137/080733991.

Yuxin Chen and Andreas Krause. Near-optimal batch mode active learning and adaptive submodular optimization. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, pages I–160–I–168. JMLR.org, 2013.

Yuxin Chen, Hiroaki Shioi, Cesar Fuentes Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 55–63, Bejing, China, 22–24 Jun 2014. PMLR.

Yuxin Chen, Shervin Javdani, Amin Karbasi, J. Andrew Bagnell, Siddhartha Srinivasa, and Andreas Krause. Submodular surrogates for value of information. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 3511–3518. AAAI Press, 2015. ISBN 0-262-51129-0.

Nguyen Cuong and Huan Xu. Adaptive maximization of pointwise submodular functions with budget constraint. In *Advances in Neural Information Processing Systems 29*, pages 1244–1252, 2016.

Nguyen Viet Cuong, Wee Sun Lee, Nan Ye, Kian Ming A Chai, and Hai Leong Chieu. Active learning for probabilistic hypotheses using the maximum gibbs error criterion. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 1457–1465. 2013.

Nguyen Viet Cuong, Wee Sun Lee, and Nan Ye. Near-optimal adaptive pool-based active learning with general loss. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI'14, pages 122–131, Arlington, Virginia, United States, 2014. AUAI Press. ISBN 978-0-9749039-1-0.

Abhimanyu Das and David Kempe. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *J. Mach. Learn. Res.*, 19(1):74–107, January 2018. ISSN 1532-4435.

Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.

Kaito Fujii and Hisashi Kashima. Budgeted stream-based active learning via adaptive submodular maximization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, pages 514–522, USA, 2016. Curran Associates Inc. ISBN 978-1-5108-3881-9.

Kaito Fujii and Shinsaku Sakaue. Beyond adaptive submodularity: Approximation guarantees of greedy policy with adaptive submodularity ratio. 2019. arXiv:1904.10748 [cs.LG].

Victor Gabillon, Branislav Kveton, Zheng Wen, Brian Eriksson, and S. Muthukrishnan. Large-scale optimistic adaptive submodularity. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI'14, pages 1816–1823. AAAI Press, 2014.

Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Int. Res.*, 42(1):427–486, September 2011.

Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal bayesian active learning with noisy observations. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 766–774. Curran Associates, Inc., 2010.

H. González-Banos. A randomized art-gallery algorithm for sensor placement. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, SCG '01, pages 232–240, New York, NY, USA, 2001. ACM. ISBN 1-58113-357-X. doi: 10.1145/378583.378674.

Anupam Gupta, Ravishankar Krishnaswamy, Marco Molinaro, and R. Ravi. Approximation algorithms for correlated knapsacks and non-martingale bandits. In *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, FOCS '11, pages 827–836, Washington, DC, USA, 2011. IEEE Computer Society. ISBN 978-0-7695-4571-4. doi: 10.1109/FOCS.2011.48. URL https://doi.org/10.1109/FOCS.2011.48.

Nir Halman, Diego Klabjan, Chung Lun Li, James Orlin, and David Simchi-Levi. Fully polynomial time approximation schemes for stochastic dynamic programs. *SIAM Journal on Discrete Mathematics*, 28(4):1725–1796, 1 2014. ISSN 0895-4801. doi: 10.1137/130925153.

Geoffrey A Hollinger, Brendan Englot, Franz S Hover, Urbashi Mitra, and Gaurav S Sukhatme. Active planning for underwater inspection and the benefit of adaptivity. *The International Journal of Robotics Research*, 32(1):3–18, 2013.

Shervin Javdani, Matthew Klingensmith, J. Andrew Bagnell, Nancy S. Pollard, and Siddhartha S. Srinivasa. Efficient touch based localization through submodularity. In *ICRA*, pages 1828–1835. IEEE, 2013.

Yasushi Kawase, Hanna Sumita, and Takuro Fukunaga. Submodular maximization with uncertain knapsack capacity. *SIAM J. Discrete Math.*, 33(3):1121–1145, 2019. doi: 10.1137/18M1174428. URL https://doi.org/10.1137/18M1174428.

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, pages 137–146, New York, NY, USA, 2003. ACM. ISBN 1-58113-737-0.

Rajiv Khanna, Ethan R. Elenberg, Alexandros G. Dimakis, Sahand N. Negahban, and Joydeep Ghosh. Scalable greedy feature selection via weak submodularity. In *AISTATS*, volume 54 of *Proceedings of Machine Learning Research*, pages 1560–1568. PMLR, 2017.

Samir Khuller, Anna Moss, and Joseph (Seffi) Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, April 1999. ISSN 0020-0190.

Andreas Krause and Daniel Golovin. Submodular function maximization. In Lucas Bordeaux, Youssef Hamadi, and Pushmeet Kohli, editors, *Tractability : Practical Approaches to Hard Problems*, pages 71 – 104. Cambridge University Press, Cambridge, 2014.

Andreas Krause and Carlos Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, UAI'05, pages 324–331, Arlington, Virginia, United States, 2005. AUAI Press. ISBN 0-9749039-1-4. URL http://dl.acm.org/citation.cfm?id=3020336.3020377.

Andreas Krause and Carlos Guestrin. Near-optimal observation selection using submodular functions. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 2*, AAAI'07, pages 1650–1654, 2007. ISBN 978-1-57735-323-2.

Andreas Krause, Ajit Paul Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.

Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pages 420–429, New York, NY, USA, 2007. ACM.

Jian Li and Wen Yuan. Stochastic combinatorial optimization via poisson approximation. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 971–980, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488731. URL http://doi.acm.org/10.1145/2488608.2488731.

Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT '11, pages 510–520, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics. ISBN 978-1-932432-87-9. URL http://dl.acm.org/citation.cfm?id=2002472.2002537.

Will Ma. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1154–1163, Philadelphia, PA, USA, 2014. Society for Industrial and Applied Mathematics. ISBN 978-1-611973-38-9. URL http://dl.acm.org/citation.cfm?id=2634074.2634159.

George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294, 1978.

Yanyan Shen, Kaushik Chakrabarti, Surajit Chaudhuri, Bolin Ding, and Lev Novik. Discovering queries based on example tuples. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 493–504, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2376-5.

Adish Singla and Andreas Krause. Incentives for privacy tradeoff in community sensing. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013, November 7-9, 2013, Palm Springs, CA, USA*, 2013.

Lichao Sun, Weiran Huang, Philip S. Yu, and Wei Chen. Multi-round influence maximization. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery; Data Mining*, KDD '18, pages 2249–2258, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5552-0. doi: 10.1145/3219819.3220101.

Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.*, 32(1):41–43, January 2004.

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 '15, pages 1134–1148, Philadelphia, PA, USA, 2015. Society for Industrial and Applied Mathematics. URL http://dl.acm.org/citation.cfm?id=2722129.2722205.

Sebastian Tschiatschek, Rishabh Iyer, Haochen Wei, and Jeff Bilmes. Learning mixtures of submodular functions for image collection summarization. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'14, pages 1413–1421, Cambridge, MA, USA, 2014. MIT Press. URL http://dl.acm.org/citation.cfm?id=2968826.2968984.

Zheng Wen, Branislav Kveton, Brian Eriksson, and Sandilya Bhamidipati. Sequential bayesian search. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, pages II–226–II–234. JMLR.org, 2013.

Yuichi Yoshida. Maximizing a monotone submodular function with a bounded curvature under a knapsack constraint. *SIAM J. Discrete Math.*, 33(3):1452–1471, 2019. doi: 10.1137/16M1107644. URL https://doi.org/10.1137/16M1107644.

Appendix A. Proofs of Lemmas 4, 6, 10 and Theorems 13, 14

Proof [Lemma 4] φ is the subrealization of σ^∞ at the start of an iteration only if $\varphi \subseteq S$; i.e., $\theta(x) = \varphi \implies \varphi \subseteq S$. To prove the converse, consider any $\varphi \in supp(\theta(x))$. If φ is the subrealization of σ^∞ at the start of an iteration, then there is a sequence of subrealizations $seq(\varphi) \triangleq \langle \varphi_1 = \emptyset, \ldots, \varphi_{|dom(\varphi)|+1} = \varphi \rangle$ s.t. φ_i is the subrealization of σ^∞ at the start of iteration i. Since the candidate selection step ($\arg\max$ function in Algorithm 1) is deterministic, $seq(\varphi)$ must be unique.

Suppose the events $\varphi \subseteq S$ and $\theta(x) = \varphi'$ occured for some $\varphi' \neq \varphi$. Consider the largest index k s.t. the first k subrealizations in $seq(\varphi)$ are equal to the corresponding k subrealizations in $seq(\varphi')$. Since the first subrealization in any such sequence is $\emptyset$, we have $k \geq 1$. Hence, we have $\varphi, \varphi' \subseteq S$, and $\exists k \geq 1$ s.t. the k^{th} subrealizations in $seq(\varphi)$ and $seq(\varphi')$ are equal, and the $k+1^{st}$ subrealizations differ. This can happen only if the output of $\arg\max$ function is not unique when φ_k is its input, which contradicts the fact that $\arg\max$ is deterministic. Hence, the lemma holds. ■

Proof [Lemma 6] Fix x. Construct a composite policy $\pi \oplus \theta$ with two stages as follows. The first stage of $\pi \oplus \theta$ is identical to π. Once π terminates, in shifts to the second stage. In this stage, $\pi \oplus \theta$ starts executing σ^∞ as if it is starting afresh: if a candidate e during a second stage iteration was already selected by $\pi \oplus \theta$ in its first stage, there is nothing more to be done in this iteration, and $\pi \oplus \theta$ moves on to the next iteration. Policy $\pi \oplus \theta$ terminates as soon as it reaches a second stage iteration where its subrealization φ (which includes all items selected in the two stages) satisfies: $\exists \varphi' \in supp(\theta(x))$ s.t $\varphi' \subseteq \varphi$. Clearly, the set of items selected by $\pi \oplus \theta$ in its first stage equals $A(\pi)$ and the set of all items selected by it equals $A(\pi) \cup dom(\theta(x))$. The expected first stage utility equals $\mathbf{E}[f(\pi)]$. During an iteration of the second stage, if $\pi \oplus \theta$ selects an item, the expected increment in utility provided by it is non-negative due to (2). Hence, the lemma is proved. ■

Proof [Lemma 10] Define $h(e) \triangleq \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \mathbf{1}_{e_i \leftarrow e} \mathbf{1}_{\vartheta_i = \vartheta} \Delta f_{dom(\vartheta)}(e)$. We now have, $f(\pi^\ddagger) = \sum_{e \in E} h(e)$. Taking expectations on both sides, we have:

$$\mathbf{E}[f(\pi^\ddagger)] = \sum_{e \in E} \mathbf{E}[h(e)]$$

$$= \sum_{e \in E} \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e | \vartheta_i = \vartheta] \mathbf{E}[\Delta f_{dom(\vartheta)}(e) | \vartheta_i = \vartheta \wedge e_i \leftarrow e]$$

$$= \sum_{e \in E} \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e | \vartheta_i = \vartheta] \mathbf{E}[\Delta f_{dom(\vartheta)}(e) | \vartheta \subseteq S \wedge \vartheta_i = \vartheta \wedge e_i \leftarrow e] \quad (32)$$

$$= \sum_{e \in E} \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e | \vartheta_i = \vartheta] \sum_{\phi \in supp(S) | \vartheta \subseteq \phi} \Pr[S = \phi | \vartheta \subseteq S \wedge \vartheta_i = \vartheta \wedge e_i \leftarrow e] \mathbf{E}[\Delta f_{dom(\vartheta)}(e) | \vartheta \subseteq S \wedge S = \phi \wedge \vartheta_i = \vartheta \wedge e_i \leftarrow e]$$

$$= \sum_{e \in E} \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e | \vartheta_i = \vartheta] \sum_{\phi \in supp(S) | \vartheta \subseteq \phi} \Pr[S = \phi | \vartheta \subseteq S] \mathbf{E}[\Delta f_{dom(\vartheta)}(e) | S = \phi] \quad (33)$$

$$= \sum_{e \in E} \sum_{i=1}^{i^*} \sum_{\vartheta \in supp(\vartheta_i)} \Pr[\vartheta_i = \vartheta] \Pr[e_i \leftarrow e | \vartheta_i = \vartheta] \mathbf{E}[\Delta f(e | \vartheta)] = \sum_{e \in E} g(e, \pi^\ddagger)$$

(32) follows from the fact that the event $\vartheta_i = \vartheta$ implies the event $\vartheta \subseteq S$. (33) follows from the facts that conditioned on $S = \phi$, the r.v. $\Delta f_{dom(\vartheta)}(e)$ is fully determined, and conditioned on $\vartheta \subseteq S$, the events $\vartheta_i = \vartheta \wedge e_i \leftarrow e$ which are fully determined by the internal random choices of $\pi^\ddagger$ along with the cost vector C is conditionally independent of the r.v. S, the state vector. ■

Proof [Theorem 13] The proof of this theorem follows along the same lines as Theorem 8 with a few key modifications. Since item costs are deterministic and $\mathcal{I}$ requires total cost $\leq B$ w.p. 1, we can drop all items e with $C(e) > B$ from the given instance and assume WLOG that the cost of every item is $\leq B$. This also implies that for every item, $C(e) = K(e)$. Further, since item costs are deterministic, during each iteration of $\pi^\dagger$, it can correctly determine if the inclusion of its candidate will cause the knapsack to overflow without having to select the item in the first place. In other words, $\pi^\dagger$ can be implemented as a singly adaptive policy whose total cost is $\leq B$ w.p. 1. Hence, $\pi^\dagger = \pi^+$, the doubly adaptive policy from which it is derived, and they are both feasible singly adaptive policies for instances $\mathcal{I}$ and $\mathcal{F}$ with $C(\pi^+) = K(\pi^+) \leq B$ w.p. 1.

Since $\lambda = B/2$, $\Pr[C(\varphi_{i+1}) \leq B] = 1$ in (20). Hence, we have $\mathbf{E}[f(\tau_\ell)] \geq \int_0^{B/2} \mathbf{E}[\bar{\nu}(x)]dx$ in (23). Combining this with the fact that $C(\pi^+) = K(\pi^+) \leq B$ w.p. 1 (i.e., replacing $2B$ with B in (25)), Lemma 11 can now be tightened as $\mathbf{E}[f(\tau_\ell)] \geq (1 - e^{-\beta/2})R$. We also have $\sum_{e \in E \setminus L} \Pr[e \in$

$A(\pi^+)] \leq 2$ in (28), which in turn yields $\mathbf{E}[f(\tau_s)] \geq \left(1 - e^{-\beta/2}\right) \bar{R}$ in (31). Putting these expressions together with the facts that $p_\ell = 1/2$ and $\mathbf{E}[f(\tau)] = (\mathbf{E}[f(\tau_\ell)] + \mathbf{E}[f(\tau_s)])/2$ yields the theorem. ■

Proof [Theorem 14] Recall that σ is the β-GREEDY policy applied to instance $\mathcal{F}$. Also recall from proof of Theorem 13 that $\pi^\dagger$ is a feasible singly adaptive policy for $\mathcal{F}$. Since f is adaptive submodular, $\alpha = 1$ and from Theorem 3, we have:

$$\mathbf{E}[f(\sigma)] \geq (1 - e^{-\beta})\mathbf{E}[f(\pi^\dagger)] \tag{34}$$

Since $\lambda = B$, the set of light items in τ is E. Hence, τ_ℓ selects items in the same sequence as σ with the exception of the final candidate e_r of σ. This item will always overflow the knapsack capacity and hence will be omitted by τ_ℓ but it may be included by σ through randomized rounding. Further, the expected utility of τ_s is at least β times the expected utility of any item in E. Hence:

$$\mathbf{E}[f(\sigma)] \leq \mathbf{E}[f(\theta(B))] + \mathbf{E}[\Delta f_{\theta(B)}(e_{\theta(B)})]$$

$$= \mathbf{E}[f(\tau_\ell)] + \sum_{\varphi} \Pr[\theta(B) = \varphi]\mathbf{E}[\Delta f_\varphi(e_\varphi)] \leq \mathbf{E}[f(\tau_\ell)] + \max_{e \in E} \mathbf{E}[f(e)] \tag{35}$$

$$\leq \mathbf{E}[f(\tau_\ell)] + \frac{\mathbf{E}[f(\tau_s)]}{\beta} = \frac{1+\beta}{\beta}(p_\ell \mathbf{E}[f(\tau_\ell)] + (1-p_\ell)\mathbf{E}[f(\tau_s)]) = \frac{1+\beta}{\beta}\mathbf{E}[f(\tau)] \tag{36}$$

(35) follows from adaptive submodularity of f. Combining (34) and (36) yields the theorem. ■