

Noisy Submodular Maximization via Adaptive Sampling with Applications to Crowdsourced Image Collection Summarization

Adish Singla
ETH Zurich
adish.singla@inf.ethz.ch

Sebastian Tschiatschek
ETH Zurich
sebastian.tschiatschek@inf.ethz.ch

Andreas Krause
ETH Zurich
krausea@ethz.ch

Abstract

We address the problem of maximizing an unknown submodular function that can only be accessed via noisy evaluations. Our work is motivated by the task of summarizing content, e.g., image collections, by leveraging users' feedback in form of clicks or ratings. For summarization tasks with the goal of maximizing coverage and diversity, submodular set functions are a natural choice. When the underlying submodular function is unknown, users' feedback can provide noisy evaluations of the function that we seek to maximize. We provide a generic algorithm – EXPGREEDY – for maximizing an unknown submodular function under cardinality constraints. This algorithm makes use of a novel exploration module – TOPX – that proposes good elements based on adaptively sampling noisy function evaluations. TOPX is able to accommodate different kinds of observation models such as value queries and pairwise comparisons. We provide PAC-style guarantees on the quality and sampling cost of the solution obtained by EXPGREEDY. We demonstrate the effectiveness of our approach in an interactive, crowdsourced image collection summarization application.

Introduction

Many applications involve the selection of a subset of items, e.g., summarization of content on the web. Typically, the task is to select a subset of items of limited cardinality with the goal of maximizing their utility. This utility is often measured via properties like diversity, information, relevance or coverage. Submodular set functions naturally capture the fore-mentioned notions of utility. Intuitively, submodular functions are set functions that satisfy a natural diminishing returns property. They have been widely used in diverse applications, including content summarization and recommendations, sensor placement, viral marketing, and numerous machine learning and computer vision tasks (Krause and Guestrin 2008; Bilmes 2015).

Summarization of image collections. One of the motivating applications, which is the subject of our experimental evaluation, is summarization of image collections. Given a collection of images, say, from *Venice* on Flickr, we would

like to select a small subset of these images that summarizes the theme *Cathedrals in Venice*, cf., Figure 1. We cast this summarization task as the problem of maximizing a submodular set function f under cardinality constraints.

Submodular maximization. Usually, it is assumed that f is known, i.e., f can be evaluated exactly by querying an oracle. In this case, a greedy algorithm is typically used for maximization. The greedy algorithm builds up a set of items by picking those of highest marginal utility in every iteration, given the items selected that far. Despite its greedy nature, this algorithm provides the best constant factor approximation to the optimal solution computable in polynomial time (Nemhauser, Wolsey, and Fisher 1978).

Maximization under noise. However, in many realistic applications, the function f is *not known* and can only be evaluated up to (additive) noise. For instance, for the image summarization task, (repeatedly) querying users' feedback in form of clicks or ratings on the individual images or image-sets can provide such noisy evaluations. There are other settings for which marginal gains are hard to compute exactly, e.g., computing marginal gains of nodes in viral marketing applications (Kempe, Kleinberg, and Tardos 2003) or conditional information gains in feature selection tasks (Krause and Guestrin 2005). In such cases, one can apply a naive *uniform* sampling approach to estimate all marginal gains up to some error ϵ and apply the standard greedy algorithm. While simple, this uniform sampling approach could have high sample complexity, rendering it impractical for real-world applications. In this work, we propose to use *adaptive* sampling strategies to reduce sample complexity while maintaining high quality of the solutions.

Our approach for adaptive sampling. The first key insight for efficient adaptive sampling is that we need to estimate the marginal gains only up to the necessary confidence to decide for the best item to select next. However, if the difference in the marginal gains of the best item and the second best item is small, this approach also suffers from high sample complexity. To overcome this, we exploit our second key insight that instead of focusing on selecting the single best item, it is sufficient to select an item from a small subset of *any* size $l \in \{1, \dots, k\}$ of high quality items, where k is the cardinality constraint. Based on these insights, we propose a novel exploration module – TOPX – that proposes good elements by adaptively sampling noisy function evaluations.

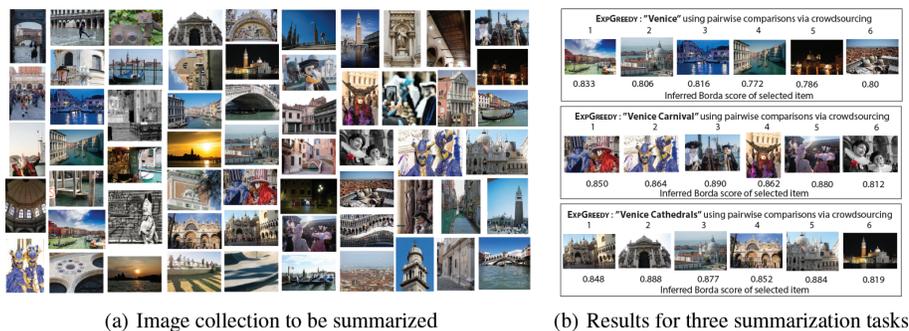


Figure 1: (a) Image collection \mathcal{V} to be summarized; (b) Summaries obtained using pairwise comparisons via crowdsourcing for the themes (i) Venice, (ii) Venice Carnival, and (iii) Venice Cathedrals. Images are numbered from 1 to 6 in the order of selection.

Our contributions. Our main contributions are:

- We provide a greedy algorithm EXPGREEDY for submodular maximization of a function via noisy evaluations. The core part of this algorithm, the exploration module TOPX, is invoked at every iteration and implements a novel adaptive sampling scheme for efficiently selecting a small set of items with high marginal utilities.
- Our theoretical analysis and experimental evaluation provide insights of how to trade-off the quality of subsets selected by EXPGREEDY against the number of evaluation queries performed by TOPX.
- We demonstrate the applicability of our algorithms in a real-world application of crowdsourcing the summarization of image collections via eliciting crowd preferences based on (noisy) pairwise comparisons, cf., Figure 1.

Related Work

Submodular function maximization (offline). Submodular set functions $f(S)$ arise in many applications and, therefore, their optimization has been studied extensively. For example, a celebrated result of Nemhauser, Wolsey, and Fisher (1978) shows that non-negative monotone submodular functions under cardinality constraints can be maximized up to a constant factor of $(1 - 1/e)$ by a simple greedy algorithm. Submodular maximization has furthermore been studied for a variety of different constraints on S , e.g., matroid constraints or graph constraints (Krause et al. 2006; Singh, Krause, and Kaiser 2009), and in different settings, e.g., distributed optimization (Mirzasoleiman et al. 2013). When the function f can only be evaluated up to (additive) noise, a naive *uniform* sampling approach has been employed to estimate all the marginal gains (Kempe, Kleinberg, and Tardos 2003; Krause and Guestrin 2005), an approach that could have high sampling complexity.

Learning submodular functions. One could approach the maximization of an unknown submodular function by first learning the function of interest and subsequently optimizing it. Tschitschek et al. (2014) present an approach for learning linear mixtures of known submodular component functions for image collection summarization. Our work is complimentary to that approach wherein we directly target the subset selection problem without learning the underlying function. In general, learning submodular functions from

data is a difficult task — Balcan and Harvey (2011) provide several negative results in a PAC-style setting.

Submodular function maximization (online). Our work is also related to online submodular maximization with (opaque) bandit feedback. Streeter and Golovin (2008) present approaches for maximizing a sequence of submodular functions in an online setting. Their adversarial setting forces them to use conservative algorithms with slow convergence. Yue and Guestrin (2011) study a more restricted setting, where the objective is an (unknown) linear combination of known submodular functions, under stochastic noise. While related in spirit, these approaches aim to minimize cumulative regret. In contrast, we aim to identify a single good solution performing as few queries as possible — the above mentioned results do not apply to our setting.

Best identification (pure exploration bandits). In exploratory bandits, the learner first explores a set of actions under time / budget constraints and then exploits the gathered information by choosing the estimated best action (top^1 identification problem) (Even-Dar, Mannor, and Mansour 2006; Bubeck, Munos, and Stoltz 2009). Beyond best individual actions, Zhou, Chen, and Li (2014) design an (ϵ, δ) -PAC algorithm for the top^m identification problem where the goal is to return a subset of size m whose aggregate utility is within ϵ compared to the aggregate utility of the m best actions. Chen et al. (2014) generalize the problem by considering combinatorial constraints on the subsets that can be selected, e.g., subsets must be size m , represent matchings, etc. They present general learning algorithms for all decision classes that admit offline maximization oracles. Dueling bandits are variants of the bandit problem where feedback is limited to relative preferences between pairs of actions. The best-identification problem is studied in this weaker information model for various notions of ranking models (e.g., *Borda* winner), cf., Busa-Fekete and Hüllermeier (2014). However, in contrast to our work, the reward functions considered by Chen et al. (2014) and other existing algorithms are modular, thus limiting the applicability of these algorithms. Our work is also related to contemporary work by Hassidim and Singer (2015), who treat submodular maximization under noise. While their algorithms apply to persistent noise, their technique is computationally demanding, and does not enable one to use noisy preference queries.

Problem Statement

Utility model. Let $\mathcal{V} = \{1, 2, \dots, N\}$ be a set of N items. We assume a utility function $f: 2^{\mathcal{V}} \rightarrow \mathbb{R}$ over subsets of \mathcal{V} . Given a set of items $S \subseteq \mathcal{V}$, the utility of this set is $f(S)$. Furthermore, we assume that f is *non-negative, monotone* and *submodular*. Monotone set functions satisfy $f(S) \leq f(S')$ for all $S \subseteq S' \subseteq \mathcal{V}$; and submodular functions satisfy the following diminishing returns condition: for all $S \subseteq S' \subseteq \mathcal{V} \setminus \{a\}$, it holds that $f(S \cup \{a\}) - f(S) \geq f(S' \cup \{a\}) - f(S')$. These conditions are satisfied by many realistic, complex utility functions (Krause and Guestrin 2011; Krause and Golovin 2012). Concretely, in our image collection summarization example, \mathcal{V} is a collection of images, and f is a function that assigns every summary $S \subseteq \mathcal{V}$ a score, preferring relevant and diverse summaries.

Observation model. In classical submodular optimization, f is assumed to be known, i.e., f can be evaluated exactly by an *oracle*. In contrast, we only assume that noisy evaluations of f can be obtained. For instance, in our summarization example, one way to evaluate f is to query users to rate a summary, or to elicit user preferences via pairwise comparisons of different summaries. In the following, we formally describe these two types of queries in more detail:

(1) *Value queries.* In this variant, we query the value for $f(a|S) = f(\{a\} \cup S) - f(S)$ for some $S \subseteq \mathcal{V}$ and $a \in \mathcal{V}$. We model the noisy evaluation or response to this query by a random variable $X_{a|S}$ with unknown sub-Gaussian distribution. We assume that $X_{a|S}$ has mean $f(a|S)$ and that repeated queries for $f(a|S)$ return samples drawn i.i.d. from the unknown distribution.

(2) *Preference queries.* Let $a, b \in \mathcal{V}$ be two items and $S \subseteq \mathcal{V}$. The preference query aims at determining whether an item a is preferred over item b in the context of S (i.e., item a has larger marginal utility than another item b). We model the noisy response of this pairwise comparison by the random variable $X_{a>b|S}$ that takes values in $\{0, 1\}$. We assume that $X_{a>b|S}$ follows some unknown distribution and satisfies the following two properties: (i) $X_{a>b|S}$ has mean larger than 0.5 iff $f(a|S) > f(b|S)$; (ii) the mapping from utilities to probabilities is monotone in the sense, that if given some set S , and given that the gaps in utilities satisfy $f(a|S) - f(b|S) \geq f(a'|S) - f(b'|S)$ for items $a, b, a', b' \in \mathcal{V}$, then the mean of $X_{a>b|S}$ is greater or equal to the mean of $X_{a'>b'|S}$. For instance, the distribution induced by the commonly used Bradley-Terry-Luce preference model (Bradley and Terry 1952; Luce 1959) satisfies these conditions. We again assume that repeated queries return samples drawn i.i.d. from the unknown distribution.

Value queries are natural, if f is approximated via stochastic simulations (e.g., as in Kempe, Kleinberg, and Tardos (2003)). On the other hand, preference queries may be a more natural way to learn what is relevant / interesting to users compared to asking them to assign numerical scores, which are difficult to calibrate.

Objective. Our goal is to select a set of the items $S \subseteq \mathcal{V}$ with $|S| \leq k$ that maximizes the utility $f(S)$. The optimal

solution to this problem is given by

$$S^{opt} = \arg \max_{S \subseteq \mathcal{V}, |S| \leq k} f(S). \quad (1)$$

Note that obtaining optimal solutions to problem (1) is intractable (Feige 1998). However, a greedy optimization scheme based on the marginal utilities of the items can provide a solution S^{greedy} such that $f(S^{greedy}) \geq (1 - \frac{1}{e}) \cdot f(S^{opt})$, i.e., a solution that is within a constant factor of the optimal solution can be efficiently determined.

In our setting, we can only evaluate the unknown utility function f via noisy queries and thus cannot hope to achieve the same guarantees. The key idea is that in a stochastic setting, our algorithms can make repeated queries and aggregate noisy evaluations to obtain sufficiently accurate estimates of the marginal gains of items. We study our proposed algorithms in a PAC setting, i.e., we aim to design algorithms that, given positive constants (ϵ, δ) , determine a set S that is ϵ -competitive relative to a reference solution with probability of at least $1 - \delta$. One natural baseline is a constant factor approximation to the optimal solution, i.e., we aim to determine a set S such that with probability at least $1 - \delta$,

$$f(S) \geq (1 - \frac{1}{e}) \cdot f(S^{opt}) - \epsilon. \quad (2)$$

Our objective is to achieve the desired (ϵ, δ) -PAC guarantee while minimizing sample complexity (i.e., the number of evaluation queries performed).

Submodular Maximization Under Noise

We now present our algorithm EXPGREEDY for maximizing submodular functions under noise. Intuitively, it aims to mimic the greedy algorithms in noise-free settings, ensuring that it selects a good element in each iteration. Since we cannot evaluate the marginal gains exactly, we must experiment with different items, and use statistical inference to select items of high value. This experimentation, the core part of our algorithm, is implemented via a novel exploration module called TOPX.

The algorithm EXPGREEDY, cf., Algorithm 1, iteratively builds up a set $S \subseteq \mathcal{V}$ by invoking TOPX($\epsilon', \delta', k', S$) at every iteration to select the next item. TOPX returns candidate items that could potentially be included in S to maximize its utility. The simplest adaptive sampling strategy that TOPX could implement is to estimate the marginal gains up to the necessary confidence to decide for the next best item to select (*top*¹ identification problem). However, if the difference in the marginal gains of the best and the second best item is small, this approach could suffer from high sample complexity. Extending ideas from the randomized greedy algorithm (Buchbinder et al. 2014), we show in Theorem 1 that in every iteration, instead of focusing on *top*¹, it is sufficient for EXPGREEDY to select an item from a small subset of items with high utility. This corresponds to the following two conditions on TOPX:

1. TOPX returns a subset $A \neq \emptyset$ of size at most k' .
2. With probability at least $1 - \delta'$, the items in A satisfy

$$\frac{1}{|A|} \sum_{a \in A} f(a|S) \geq \max_{\substack{B \subseteq \mathcal{V} \\ |B|=|A|}} \left[\frac{1}{|B|} \sum_{b \in B} f(b|S) \right] - \epsilon'. \quad (3)$$

Algorithm 1: EXPGREEDY

- 1 **Input:** Ground set \mathcal{V} ; No. of items to pick k ; $\epsilon, \delta > 0$;
 - 2 **Output:** Set of items $S \subseteq \mathcal{V}$: $|S| \leq k$, such that S is ϵ -competitive with probability at least $(1 - \delta)$;
 - 3 **Initialize:** $S = \emptyset$
foreach $j = 1, \dots, k$ **do**
 - 4 $A = \text{TOPX}(\epsilon', \delta', k', S)$
 - 5 Sample s uniformly at random from A
 - 6 $S = S \cup \{s\}$
 - 7 **return** S
-

At any iteration, satisfying these two conditions is equivalent to solving a top^l identification problem for $l = |A|$ with PAC-parameters (ϵ', δ') . TOPX essentially returns a set of size l containing items of largest marginal utilities given S . Let us consider two special cases, (i) $l = 1$ and (ii) $l = k$ for all $j \in \{1, \dots, k\}$ iterations. Then, in the noise free setting, EXPGREEDY for case (i) mimics the classical greedy algorithm (Nemhauser, Wolsey, and Fisher 1978) and for case (ii) the randomized greedy algorithm (Buchbinder et al. 2014), respectively. As discussed in the next section, the sample complexity of these two cases can be very high. The key insight we use in EXPGREEDY is that if we could efficiently solve the top^l problem for *any* size $l \in \{1, \dots, k\}$ (i.e., $|A|$ is neither necessarily 1 or k), then the solution to the submodular maximization problem is guaranteed to be of high quality. This is summarized in the following theorem:

Theorem 1. *Let $\epsilon > 0, \delta \in (0, 1)$. Using $\epsilon' = \frac{\epsilon}{k}, \delta' = \frac{\delta}{k}$ and $k' = k$ for invoking TOPX, Algorithm 1 returns a set S that satisfies $\mathbb{E}[f(S)] \geq (1 - \frac{1}{e}) \cdot f(S^{opt}) - \epsilon$ with probability at least $1 - \delta$. For the case that $k' = 1$, the guarantee is $f(S) \geq (1 - \frac{1}{e}) \cdot f(S^{opt}) - \epsilon$ with probability at least $1 - \delta$.*

The proof is provided in Appendix A of the extended version of paper (Singla, Tschischek, and Krause 2016).

As it turns out, solutions of the greedy algorithm in the noiseless setting S^{greedy} often have utility larger than $(1 - \frac{1}{e}) \cdot f(S^{opt})$. Therefore, we also seek algorithms that with probability at least $1 - \delta$ identify solutions satisfying

$$f(S) \geq f(S^{greedy}) - \epsilon. \quad (4)$$

This can be achieved according to the following theorem:

Theorem 2. *Let $\delta \in (0, 1)$. Using $\epsilon' = 0, \delta' = \frac{\delta}{k}$ and $k' = 1$ for invoking TOPX, Algorithm 1 returns a set S that satisfies $f(S) = f(S^{greedy})$ with probability at least $1 - \delta$.*

If we set $k' = 1$ and $\epsilon = 0$, then condition (3) is actually equivalent to requiring that TOPX, with high probability, identifies the element with largest marginal utility. The proof follows by application of the union bound. This theorem ensures that if we can construct a corresponding exploration module, we can successfully compete with the greedy algorithm that has access to f . However, this can be prohibitively expensive in terms of the required number of queries performed by TOPX, cf., Appendix B of the extended version of this paper (Singla, Tschischek, and Krause 2016).

Exploration Module TOPX

In this section, we describe the design of our exploration module TOPX used by EXPGREEDY.

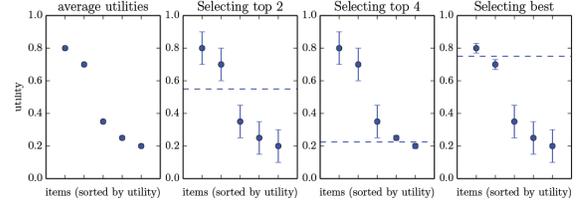


Figure 2: An example to illustrate the idea of top^l item selection with $N = 5$ items and parameter $k' = 4$. While both the top^1 identification ($l = 1$) and the top^4 identification ($l = k'$) have high sample complexity, the top^2 identification ($l = 2 \in \{1, \dots, k'\}$) is relatively easy.

TOPX with Value Queries

We begin with the observation that for fixed $l \in \{1, \dots, k'\}$ (for instance $l = 1$ or $l = k'$), an exploration module TOPX satisfying condition (3) can be implemented by solving a top^l identification problem. This can be seen as follows.

Given input parameter S , for each $a \in \mathcal{V} \setminus S$ define its value $v_a = f(a|S)$, and define $v_a = 0$ for each $a \in S$. Then, the value of any set A (in the context of S) is $v(A) := \sum_{a \in A} v_a$, which is a modular (additive) set function. Thus, in this setting (fixed l), meeting condition (3) requires identifying a set A maximizing a *modular* set function under cardinality constraints from noisy queries. This corresponds to the top^l best-arm identification problem. In particular, Chen et al. (2014) proposed an algorithm – CLUCB – for identifying a best subset for modular functions under combinatorial constraints. At a high level, CLUCB maintains upper and lower confidence bounds on the item values, and adaptively samples noisy function evaluations until the pessimistic estimate (lower confidence bound) for the current best top^l subset exceeds the optimistic estimate (upper confidence bound) of any other subset of size l . The worst case sample complexity of this problem is characterized by the gap between the values of l -th item and $(l+1)$ -th item (with items indexed in descending order according to the values v_a).

As mentioned, the sample complexity of the top^l problem can vary by orders of magnitude for different values of l , cf., Figure 2. Unfortunately, we do not know the value of l with the lowest sample complexity in advance. However, we can modify CLUCB to jointly estimate the marginal gains and solve the top^l problem with lowest sample complexity.

The proposed algorithm implementing this idea is presented in Algorithm 2. It maintains confidence intervals for marginal item values, with the confidence radius of an item i computed as $rad_t(i) = R\sqrt{2 \log(\frac{4Nt^3}{\delta'}) / T_t(i)}$, where the corresponding random variables $X_{i|S}$ are assumed to have an R -sub-Gaussian tail and $T_t(i)$ is the number of observations made for item i by time t . If the exact value of R is not known, it can be upper bounded by the range (as long as the variables have bounded centered range). Upon termination, the algorithm returns a set A that satisfies condition (3). Extending results from Chen et al. (2014), we can bound the sample complexity of our algorithm as follows. Consider the gaps $\Delta_l = f(\pi(l)|S) - f(\pi(l+1)|S)$, where $\pi: \mathcal{V} \rightarrow \{1, \dots, N\}$ is a permutation of the items such that

Algorithm 2: TOPX

1 **Input:** Ground set \mathcal{V} ; Set S ; integer k' ; $\epsilon', \delta' > 0$;
2 **Output:** Set of items $A \subseteq \mathcal{V}$: $|A| \leq k'$, such that A satisfies (3) with probability at least $1 - \delta'$;
3 **Initialize:** For all $i = 1, \dots, |\mathcal{V}|$: observe $X_{i|S}$ and set v_i to that value, set $T_1(i) = 1$;
 for $t = 1, \dots, \mathbf{do}$
4 Compute confidence radius $rad_t(i)$ for all i ;
5 Initialize list of items to query $Q = \square$;
 for $l = 1, \dots, k'$ **do**
6 $M_t = \arg \max_{B \subseteq \mathcal{V}, |B|=l} \sum_{i \in B} v_i$;
7 **foreach** $i = 1, \dots, |\mathcal{V}|$ **do**
8 If $i \in M_t$ set $\tilde{v}_i = v_i - rad_t(i)$, otherwise
 set $\tilde{v}_i = \bar{v}_i + rad_t(i)$
9 $\tilde{M}_t = \arg \max_{B \subseteq \mathcal{V}, |B|=l} \sum_{i \in B} \tilde{v}_i$;
10 **if** $[\sum_{i \in \tilde{M}_t} \tilde{v}_i - \sum_{i \in M_t} \tilde{v}_i] \leq l \cdot \epsilon'$ **then**
11 Set $A = M_t$ and **return** A
12 Set $q = \arg \max_{i \in (M \setminus \tilde{M}) \cup (\tilde{M} \setminus M)} rad_t(i)$;
13 Update list of items to query:
 $Q = Q.append(q)$
 foreach $q \in Q$ **do**
14 Query and observe output $X_{q|S}$;
15 Update empirical means v_{t+1} using the output;
16 Update observation counts $T_{t+1}(q) = T_t(q) + 1$
 and $T_{t+1}(j) = T_t(j)$ for all $j \neq q$;

$f(\pi(1)|S) \geq f(\pi(2)|S) \geq \dots \geq f(\pi(N)|S)$. For every fixed l , the sample complexity of identifying a set of top l items is characterized by this gap Δ_l . The reason is that if Δ_l is small, many samples are needed to ensure that the confidence bounds rad_t are small enough to distinguish the top l elements from the runner-up. Our key insight is that we can be adaptive to the largest Δ_l , for $l \in \{1, \dots, k'\}$. That is, as long as there is *some* value of l with large Δ_l , we will be able to enjoy low sample complexity (cf., Figure 2):

Theorem 3. *Given $\epsilon' > 0, \delta' \in (0, 1), S \subseteq \mathcal{V}$ and k' , Algorithm 2 returns a set $A \subseteq \mathcal{V}, |A| \leq k'$ that with probability at least $1 - \delta'$ satisfies condition (3) using at most*

$$T \leq \mathcal{O} \left(k' \min_{l=1, \dots, k'} \left[R^2 H_{(l, \epsilon')} \log \left(\frac{R^2}{\delta'} H_{(l, \epsilon')} \right) \right] \right)$$

samples, where $H_{(l, \epsilon')} = N \min \{ \frac{4}{\Delta_l^2}, \frac{1}{\epsilon'^2} \}$.

A proof sketch is given in Appendix C of the extended version of this paper (Singla, Tschitschek, and Krause 2016).

TOPX with Preference Queries

We now show how noisy preference queries can be used. As introduced previously, we assume that there exists an underlying preference model (unknown to the algorithm) that induces probabilities $P_{i>j|S}$ for item i to be preferred over item j given the values $f(i|S)$ and $f(j|S)$. In this work, we focus on identifying the *Borda winner*, i.e., the item i maximizing the Borda score $P(i|S)$, formally given as $\frac{1}{(N-1)} \cdot \sum_{j \in \mathcal{V} \setminus \{i\}} P_{i>j|S}$. The Borda score measures the probability that item i is preferred to another item chosen uniformly at random. Furthermore, in our model where

we assume that an increasing gap in the utilities leads to monotonic increase in the induced probabilities, it holds that the top l items in terms of marginal gains are the top l items in terms of Borda scores. We now make use of a result called *Borda reduction* (Jamieson et al. 2015; Busa-Fekete and Hüllermeier 2014), a technique that allows us to reduce preference queries to value queries, and to consequently invoke Algorithm 2 with small modifications.

Defining values v_i via Borda score. For each item $i \in \mathcal{V}$, Algorithm 2 (step 3,15) tracks and updates mean estimates of the values v_i . For preference queries, these values are replaced with the Borda scores. The sample complexity in Theorem 3 is then given in terms of these Borda scores. For instance, for the Bradley-Terry-Luce preference model (Bradley and Terry 1952; Luce 1959), the Borda score for item i is given by $\frac{1}{(N-1)} \cdot \sum_{j \in \mathcal{V} \setminus \{i\}} \frac{1}{1 + \exp(-\beta(f(i|S) - f(j|S)))}$. Here, β captures the problem difficulty: $\beta \rightarrow \infty$ corresponds to the case of noise-free responses, and $\beta \rightarrow 0$ corresponds to uniformly random binary responses. The effect of β is further illustrated in the synthetic experiments.

Incorporating noisy responses. Observing the value for item i in the preference query model corresponds to pairing i with an item in the set $\mathcal{V} \setminus \{i\}$ selected uniformly at random. The observed preference response provides an unbiased estimate of the Borda score for item i . More generally, we can pick a small set Z_i of fixed size τ selected uniformly at random with replacement from $\mathcal{V} \setminus \{i\}$, and compare i against each member of Z_i . Then, the observed Borda score for item i is calculated as $\frac{1}{\tau} \cdot \sum_{j \in Z_i} X_{i>j|S}$. Here, τ is a parameter of the algorithm. One can observe that the cost of one preference query is τ times the cost of one value query. The effect of τ will be further illustrated in the experiments.

The subtle point of terminating Algorithm 2 (step 10) is discussed in Appendix D of the extended version of this paper (Singla, Tschitschek, and Krause 2016).

Experimental Evaluation

We now report on the results of our synthetic experiments.

Experimental Setup

Utility function f and set of items \mathcal{V} . In the synthetic experiments we assume that there is an underlying submodular utility function f which we aim to maximize. The exploration module TOPX performs value or preference queries, and receives noisy responses based on model parameters and the marginal gains of the items for this function. For our experiments, we constructed a realistic probabilistic coverage utility function following the ideas from El-Arini et al. (2009) over a ground set of $N = 60$ items. Details of this construction are not important for the results stated below and can be found in Appendix E of the extended version of this paper (Singla, Tschitschek, and Krause 2016).

Benchmarks and Metrics. We compare several variants of our algorithm, referred to by EXPGREEDY_θ , where θ refers to the parameters used to invoke TOPX. In particular, $\theta = O$ means $k' = 1$ and $\epsilon' = \epsilon/k$ (i.e., competing with $(1 - \frac{1}{\epsilon})$ of $f(S^{opt})$); $\theta = G$ means $k' = 1$ but $\epsilon' = 0$

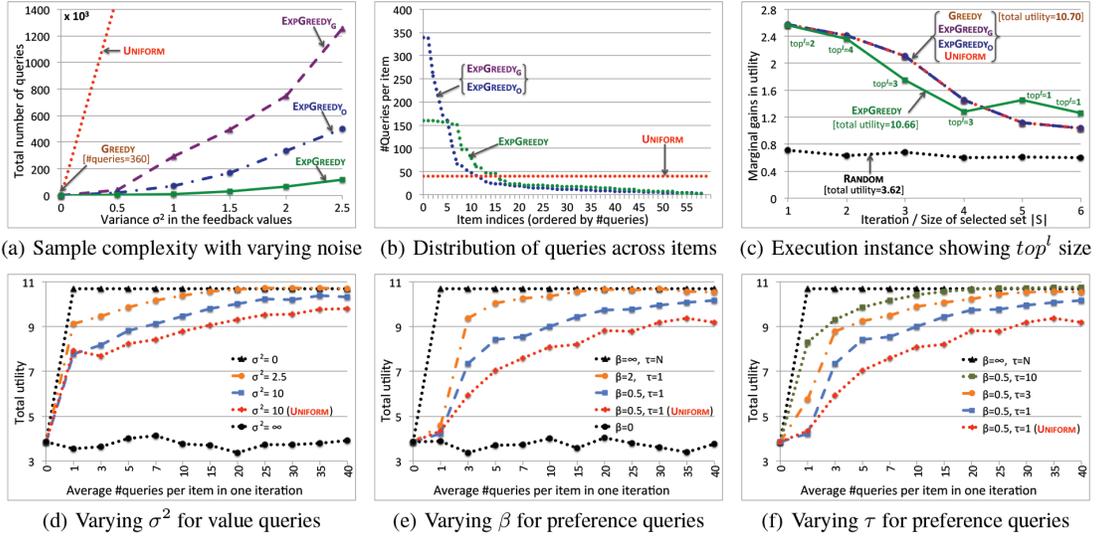


Figure 3: Experimental results using synthetic function f and simulated query responses: (a)–(d) results are for value queries and (e)–(f) results are for preference queries. (a) EXPGREEDY dramatically reduces the sample complexity compared to UNIFORM and the other adaptive baselines. (b) EXPGREEDY adaptively allocates queries to identify largest gap Δ_l . (c) An execution instance showing the marginal gains and sizes of top^l solutions returned by TOPX. (d)–(f) EXPGREEDY is robust to noise and outperforms UNIFORM baseline.

(i.e., competing with $f(S^{greedy})$); omitting θ means $k' = k$ and $\epsilon = \epsilon/k$. As benchmarks, we compare the performance with the deterministic greedy algorithm GREEDY (with access to noise-free evaluations), as well as with random selection RANDOM. As a natural and competitive baseline, we compare our algorithms against UNIFORM (Kempe, Kleinberg, and Tardos 2003; Krause and Guestrin 2005) — replacing our TOPX module by a naive exploration module that uniformly samples all the items for best item identification. For all the experiments, we used PAC parameters ($\epsilon = 0.1, \delta = 0.05$) and a cardinality constraint of $k = 6$.

Results

Sample Complexity. In Figure 3(a), we consider value queries, and compare the number of queries performed by different algorithms under varying noise-levels until convergence to the solution with the desired guarantees. For variance σ^2 , we generated the query responses by sampling uniformly from the interval $[\mu - \sigma^2, \mu + \sigma^2]$, where μ is the expected value of that query. For reference, the query cost of GREEDY with access to the unknown function f is marked (which equals $N \cdot k$). The sample complexity differs by orders of magnitude, i.e., the number of queries performed by EXPGREEDY grows much slower than that of EXPGREEDY_G and EXPGREEDY_O. The sample complexity of UNIFORM is worse by further orders of magnitude compared to any of the variants of our algorithm.

Varying σ^2 for value queries and (β, τ) for preference queries. Next, we investigate the quality of obtained solutions for a limited budget on the total number of queries that can be performed. Although convergence may be slow, one may still get good solutions early in many cases. In Figures 3(d)–3(f) we vary the available average budget per item per iteration on the x-axis — the total budget available in

terms of queries that can be performed is equivalent to $N \cdot k$ times the average budget shown on x-axis. In particular, we compare the quality of the solutions obtained by EXPGREEDY for different σ^2 in Figure 3(d) and for different (β, τ) in Figure 3(e)–3(f), averaged over 50 runs. The parameter β controls the noise-level in responses to the preference queries, whereas τ is the algorithm’s parameter indicating how many pairwise comparisons are done in one query to reduce variance. For reference, we show the extreme case of $\sigma^2 = \infty$ and $\beta = 0$ which is equivalent to RANDOM; and the case of $\sigma^2 = 0$ and $(\beta = \infty, \tau = N)$ which is equivalent to GREEDY. In general, for higher σ in value queries, and for lower β or smaller τ in preference queries, more budget must be spent to achieve solutions with high utility.

Comparison with UNIFORM exploration. In Figure 3(d) and Figure 3(e)–3(f), we also report the quality of solutions obtained by UNIFORM for the case of $\sigma^2 = 10$ and $(\beta = 0.5, \tau = 1)$ respectively. As we can see in the results, in order to achieve a desired value of total utility, UNIFORM may require up to 3 times more budget in comparison to that required by EXPGREEDY. Figure 3(b) compares how different algorithms allocate budget across items (i.e., the distribution of queries), in one of the iterations. We observe that EXPGREEDY does more exploration across different items compared to EXPGREEDY_G and EXPGREEDY_O. However, the exploration is heavily skewed in comparison to UNIFORM because of the adaptive sampling by TOPX. Figure 3(c) shows the marginal gain in utility of the considered algorithms at different iterations for a particular execution instance (no averaging over multiple executions of the algorithms was performed). For EXPGREEDY, the size of top^l solutions returned by TOPX in every iteration is indicated in the Figure, demonstrating that TOPX adaptively allocates queries to efficiently identify the largest gap Δ_l .

Image Collection Summarization

We now present results on a crowdsourced image collection summarization application, performed on Amazon’s Mechanical Turk platform. As our image set \mathcal{V} , we retrieved 60 images in some way related to the city of Venice from Flickr, cf., Figure 1(a). A total of over 100 distinct workers participated per summarization task.

The workers were queried for pairwise preferences as follows. They were told that our goal is to summarize images from Venice, motivated by the application of selecting a small set of pictures to send to friends after returning from a trip to Venice. The detailed instructions can be found in Appendix F of the extended version of this paper (Singla, Tschitschek, and Krause 2016). We ran three instances of the algorithm for three distinct summarization tasks for the themes (i) Venice, (ii) Venice Carnival, and (iii) Venice Cathedrals. The workers were told the particular theme for summarization. Additionally, the set of images already selected and two proposal images a and b were shown. Then they were asked which of the two images would improve the summary more if added to the already selected images.

For running EXPGREEDY, we used an average budget of 25 queries per item per iteration and $\tau = 3$, cf., Figure 3(f). Note that we did not make use of the function f constructed for the synthetic experiments at all, i.e., the function we maximized was not known to us. The results of this experiment are shown in Figure 1(b), demonstrating that our methodology works in real-word settings, produces high quality summaries and captures the semantics of the task.

Conclusions

We considered the problem of cardinality constrained submodular function maximization under noise, i.e., the function to be optimized can only be evaluated via noisy queries. We proposed algorithms based on novel adaptive sampling strategies to achieve high quality solutions with low sample complexity. Our theoretical analysis and experimental evaluation provide insights into the trade-offs between solution quality and sample complexity. Furthermore, we demonstrated the practical applicability of our approach on a crowdsourced image collection summarization application.

Acknowledgments. We would like to thank Besmira Nushi for helpful discussions. This research is supported in part by SNSF grant 200021_137971 and the Nano-Tera.ch program as part of the Opensense II project.

References

Balcan, M., and Harvey, N. J. A. 2011. Learning submodular functions. In *STOC*.

Bilmes, J. 2015. Submodularity in machine learning applications. AAAI, Tutorial.

Bradley, R. A., and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* 39(3/4):324–345.

Bubeck, S.; Munos, R.; and Stoltz, G. 2009. Pure exploration in multi-armed bandits problems. In *ALT*, 23–37.

Buchbinder, N.; Feldman, M.; Naor, J.; and Schwartz, R. 2014. Submodular maximization with cardinality constraints. In *SODA*, 1433–1452.

Busa-Fekete, R., and Hüllermeier, E. 2014. A survey of preference-based online learning with bandit algorithms. In *Algorithmic Learning Theory*, 18–39. Springer.

Chen, S.; Lin, T.; King, I.; Lyu, M. R.; and Chen, W. 2014. Combinatorial pure exploration of multi-armed bandits. In *NIPS*, 379–387.

El-Arini, K.; Veda, G.; Shahaf, D.; and Guestrin, C. 2009. Turning down the noise in the blogosphere. In *KDD*.

Even-Dar, E.; Mannor, S.; and Mansour, Y. 2006. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *JMLR* 7:1079–1105.

Feige, U. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM* 45:314–318.

Hassidim, A., and Singer, Y. 2015. Submodular optimization under noise. Working paper.

Jamieson, K. G.; Katariya, S.; Deshpande, A.; and Nowak, R. D. 2015. Sparse dueling bandits. In *AISTATS*.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *KDD*.

Krause, A., and Golovin, D. 2012. Submodular function maximization. *Tractability: Practical Approaches to Hard Problems* 3:19.

Krause, A., and Guestrin, C. 2005. Near-optimal nonmyopic value of information in graphical models. In *UAI*.

Krause, A., and Guestrin, C. 2008. Beyond convexity: Submodularity in machine learning. ICML, Tutorial.

Krause, A., and Guestrin, C. 2011. Submodularity and its applications in optimized information gathering. *TIST* 2(4):32.

Krause, A.; Guestrin, C.; Gupta, A.; and Kleinberg, J. 2006. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *IPSN*, 2–10.

Luce, R. D. 1959. *Individual Choice Behavior: A theoretical analysis*. Wiley.

Mirzasoleiman, B.; Karbasi, A.; Sarkar, R.; and Krause, A. 2013. Distributed submodular maximization: Identifying representative elements in massive data. In *NIPS*.

Nemhauser, G.; Wolsey, L.; and Fisher, M. 1978. An analysis of the approximations for maximizing submodular set functions. *Math. Prog.* 14:265–294.

Singh, A.; Krause, A.; and Kaiser, W. J. 2009. Nonmyopic adaptive informative path planning for multiple robots. In *IJCAI*, 1843–1850.

Singla, A.; Tschitschek, S.; and Krause, A. 2016. Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization (extended version). <http://arxiv.org/abs/1511.07211>.

Streeter, M., and Golovin, D. 2008. An online algorithm for maximizing submodular functions. In *NIPS*, 1577–1584.

Tschitschek, S.; Iyer, R.; Wei, H.; and Bilmes, J. 2014. Learning mixtures of submodular functions for image collection summarization. In *NIPS*.

Yue, Y., and Guestrin, C. 2011. Linear submodular bandits and their application to diversified retrieval. In *NIPS*.

Zhou, Y.; Chen, X.; and Li, J. 2014. Optimal PAC multiple arm identification with applications to crowdsourcing. In *ICML*.