
Sample Efficient Graph-Based Optimization with Noisy Observations

Tan Nguyen

ACEMS, School of Mathematical Sciences
Queensland University of Technology

Ali Shameli

Stanford University

Yasin Abbasi-Yadkori

Adobe Research

Anup Rao

Adobe Research

Branislav Kveton

Adobe Research

Abstract

We study sample complexity of optimizing “hill-climbing friendly” functions defined on a graph under noisy observations. We define a notion of convexity, and we show that a variant of best-arm identification can find a near-optimal solution after a small number of queries that is independent of the size of the graph. For functions that have local minima and are nearly convex, we show a sample complexity for the classical simulated annealing under noisy observations. We show effectiveness of the greedy algorithm with restarts and the simulated annealing on problems of graph-based nearest neighbor classification as well as a web document re-ranking application.

1 Introduction

Stochastic optimization of a function defined on a large finite set frequently arises in many practical problems. Instances of this problem include finding the most attractive design for a webpage, the node with maximum influence in a social network, etc. There are a number of approaches to this problem. At one extreme, we can use global optimization methods such as simulated annealing, genetic algorithms, or cross-entropy methods (Rubinstein, 1997, Christian and Casella, 1999). Although limited theoretical performance guarantees are available, these methods perform well in a number of practical applications. In the other extreme, we can use the best-arm identification algorithms from bandit literature or hypothesis testing from statistics community (Mannor and Tsitsiklis, 2004, Even-Dar et al., 2006, Audibert and

Bubeck, 2010). We have stronger sample complexity results for this family of algorithms. The sample complexity states that, with high probability, the algorithm will return a near-optimal solution after a small number of observations, and this number typically grows polynomially with the size of the decision set. These methods however often perform poorly when the size of the decision set is large. In such problems, it is important to exploit the structure of the specific problems to speed up the optimization. If appropriate features are available and it can be assumed that the objective function is linear in these features, then algorithms designed for bandit linear optimization are applicable (Auer, 2002). More generally, kernel-based or neural network based bandit algorithms are available for problems with non-linear objective functions (Srinivas et al., 2010).

We are particularly interested in problems where the item similarity is captured by a graph structure. In general, and with no further conditions, we cannot hope to show non-trivial sample complexity rates. We observe that in many real-world applications, the objective function is easy and *hill-climbing friendly*, in the sense that from many nodes, there exist a monotonic path to the global minima. We make this property explicit by defining a notion of *convexity* for functions defined on graphs. Under this condition, we show that a hill-climbing procedure that uses a variant of best-arm identification as a sub-routine has a small sample complexity that is independent of the size of graph. In the presence of local minima, this greedy procedure might require many restarts, and might not be efficient in practice. Simulated annealing is commonly used in practice for such problems. We also define a notion of nearly convex functions that allows for existence of shallow local minima. We show that for nearly convex functions and using an appropriate estimation of function values, the classical simulated annealing procedure finds near optimal nodes after a small number of function evaluations.

While asymptotic convergence of simulated annealing is studied extensively in the literature, there are only few finite-

time convergence rates for this important algorithm. Sreenivas Pydi et al. (2018) show finite-time convergence rates for the algorithm, but they consider only deterministic functions and their rates scale with the size of the graph. These results, that are obtained under very general conditions, do not quite explain success of the simulated annealing algorithm in large-scale problems. In practice, simulated annealing finds a near-optimal solution with a much smaller number of function evaluations. Our bounds are in terms of the convexity of the function, and the rates do not scale with the size of the graph. Additionally, our results hold in the noisy setting. Bouttier and Gavra (2017) show convergence rates for simulated annealing applied to noisy functions. Their results, however, appear to have several gaps.

1.1 Related Work

A best-arm identification algorithm from bandit literature can find a near-optimal node, and the time and sample complexity is linear in the number of nodes (Mannor and Tsitsiklis, 2004, Even-Dar et al., 2006, Audibert and Bubeck, 2010, Jamieson and Nowak, 2014, Kaufmann et al., 2016). Such complexity is not acceptable when the size of the graph is very large. We are interested in designing and analyzing algorithms that find a near-optimal node, and the sample and computational complexity is sublinear in the number of nodes.

Bandit algorithms for graph functions are studied in Valko et al. (2014). The sample complexity of these algorithms scale with the smoothness of the function, but the computational complexity scales linearly with the number of nodes. We are interested in problems where the graph is very large, and we might have access to only local information. As we will see in the experiments, the SPECTRALBANDIT algorithm of Valko et al. (2014) is not applicable to problems with large graphs. Further, we study a different notion of regularity that is inspired by convexity in continuous optimization. Bandit problems for large action sets are also studied under a number of different notions of regularity (Bubeck et al., 2009, Abbasi-Yadkori et al., 2011, Carpentier and Valko, 2015, Grill et al., 2015).

1.2 Contributions

We cannot hope to achieve sublinear rate without further conditions on the function. We define a notion of strong convexity for functions defined on a graph. For strongly convex functions, we design an algorithm, called EXPLORE-DESCEND, to find the optimal node with guaranteed error probability. The EXPLORE-DESCEND algorithm uses a *best-arm identification* algorithm as a submodule. The best-arm identification problem is the problem of finding the optimal action given a sampling budget. For *nearly convex functions*, we show that the classical simulated annealing algorithm finds the global minima in a reasonable time.

For convex functions, the EXPLORE-DESCEND has lower sample complexity, but simulated annealing can handle non-convex functions.

We study the empirical performance of EXPLORE-DESCEND and simulated annealing in two applications. First, we consider the problem of content optimization in a digital marketing application. In each round, a user visits a webpage that is composed of a number of components. The learning agent chooses the contents of these components and shows the resulting page to the user. The user feedback, in the form of click/no click or conversion/no conversion, is recorded and used to update the decision making policy. The objective is to return a page configuration with near-optimal click-through rate. We would like to find such a configuration as quickly as possible with a small number of user interactions. In this problem, each page configuration is a node of the graph, and two nodes are connected if the corresponding page configurations differ in only one component.

Our second application is the problem of nearest neighbor search and classification. Given a set of points $\mathcal{V} = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$, a query point $y \in \mathbb{R}^d$, and a distance function $\ell : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, we are interested in solving $\operatorname{argmin}_{x \in \mathcal{V}} \ell(x, y)$. A trivial solution to this problem is to examine all points in \mathcal{V} and return the point with the smallest distance value. The computational complexity of this method is $O(n)$, and is not practical in large-scale problems. An approximate nearest neighbor method returns an approximate solution in a sublinear time. A class of approximate nearest neighbor methods that is particularly suitable for big-data applications is the graph-based nearest neighbor search Arya and Mount (1993). In a graph-based search, we construct a proximity graph in the training phase, and perform a hill-climbing search in the test phase. To improve the performance, we perform an additional smoothing procedure that replaces the value of a node by the average function values in a vicinity of the node. In practice, these average values are estimated by performing random walks, and hence the problem is a graph optimization problem with noisy observations. We show that the proposed graph optimization technique outperforms popular nearest neighbor search methods such as KDTree and LSH in two image classification problems. Interestingly, and compared to these methods, the computational complexity of the proposed technique appears to be less sensitive to the size of the training data. This property is particularly appealing in big data applications.

1.3 Notations

We use $[K]$ to denote the set $\{1, 2, \dots, K\}$. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with n nodes $\mathcal{V} = \{1, \dots, n\}$ and the set of edges \mathcal{E} . Let $f : \mathcal{V} \rightarrow [0, 1]$ be some unknown function defined on \mathcal{V} . Let $x^* = \operatorname{argmin}_{x \in \mathcal{V}} f(x)$ be the global minimizer of f .

The goal is to find a node x with small loss $f(x) - f(x^*)$. In this paper, we study the problem where the evaluation of f is noisy, such that we can only observe $g(x; \eta) = f(x) + \eta$, where η is a zero-mean R -sub-Gaussian random variable, meaning that for any $\lambda \in \mathbb{R}$, $\mathbf{E}(e^{\lambda\eta}) \leq e^{\lambda^2 R^2/2}$.

Let \mathcal{N}_x denote the set of neighbors of node $x \in \mathcal{V}$. Let $\overline{\mathcal{N}}_x = \mathcal{N}_x \cup \{x\}$. For simplicity, we assume all nodes have the same degree and we let $d = |\mathcal{N}_x|$ denote the number of neighbors. We sometimes write f_x to denote the function value $f(x)$. For two nodes $x, y \in \mathcal{V}$, we use $\mathcal{P}_{x,y}$ to denote all paths from x to y in the graph. We use \mathcal{P}_x to denote all paths starting from node x . We use $\ell(p)$ to denote the length of a path.

1.4 Convexity

The general discrete optimization problem is hard for an unrestricted function f and graph \mathcal{G} . As such, we study a restricted class of problems that allow for efficient algorithms. Let $\Delta_{x,z} = f_x - f_z$ be the amount of improvement if we move from node x to the neighbor $z \in \mathcal{N}_x$. We say path $p = \{x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_k\}$ is m -strongly convex if $\Delta_{x_{i-1}, x_i} - \Delta_{x_i, x_{i+1}} \geq m\Delta_{x_i, x_{i+1}} > 0$, for all $i = 1, \dots, k$. Sometimes we use Δ_{x_i} to denote $\Delta_{x_i, x_{i+1}}$ if the m -strongly convex path is clear from the context. We use $\overline{\Delta}_x$ to denote $\max_{z \in \mathcal{N}_x} \Delta_{x,z}$.

Definition 1 (Convex Functions) A function $f : \mathcal{V} \rightarrow \mathbb{R}$ defined on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is (strongly) convex if from any node $x \in \mathcal{V}$, there exists a (strongly) convex path to the global minima x^* .

For a nearly convex function, as defined next, the convexity condition is not satisfied at all points.

Definition 2 (Nearly Convex Functions) Let $\alpha > 0$ be a parameter. Let \mathcal{C} be the set of points such that $\overline{\Delta}_x \geq \alpha(f_x - f_{x^*})$. We say function $f : \mathcal{V} \rightarrow \mathbb{R}$ is (α, c, r) -nearly convex if for any $x \in \mathcal{V} \setminus \mathcal{C}$, there exists a $y \in \mathcal{C}$ and a $p \in \mathcal{P}_{xy}$ such that $\ell(p) \leq r$ and $\max_{z \in p} f(z) - f(x) \leq c$.

A path that satisfies the above conditions is called a low energy path. A convex function is also a $(\frac{m}{m+1}, 0, 0)$ -nearly convex function.

Lemma 1. We have that $f_x - f_{x^*} \leq \frac{m+1}{m} \Delta_x$.

Proof. Let $p = \{y = x \rightarrow y_1 \rightarrow \dots \rightarrow y_k\}$ be a strongly convex path. Then

$$\Delta_{y_k} \leq \frac{1}{1+m} \Delta_{y_{k-1}} \leq \dots \leq \frac{1}{(1+m)^k} \Delta_x.$$

We have that

$$f_x - f_{x_k} = \sum_{i=1}^k (f_{y_{i-1}} - f_{y_i}) = \sum_{i=1}^k \Delta_{y_{i-1}}$$

$$= \Delta_x + \sum_{i=1}^{k-1} \Delta_{y_i} \leq \sum_{i=0}^{k-1} \frac{\Delta_x}{(1+m)^i}.$$

We conclude that $f_x - f_{y_k} \leq \frac{m+1}{m} \Delta_x$, from which the statement follows. \square

Concave and nearly concave functions are similarly defined. The convexity condition allows for efficient algorithms. The intuition behind our analysis is that, by strong convexity, if node x is far from the global minima the improvement Δ_x is large. When degree d is relatively small, the local search methods have a sufficiently large probability of hitting a good direction (the strongly convex path). So either we are already close to the global minima, or we are far from the global minima and with a constant probability we make a large improvement.

2 Approximating the Global Minimizer

We analyze two algorithms for the graph-based optimization. The first algorithm is a local search procedure where in each round, the learner attempts to identify a good neighbor and move there. This algorithm is analyzed under the strong convexity condition. The second algorithm is the well-known simulated annealing procedure with an exponential transition function. We provide high probability error bounds for the greedy method, while the sample complexity of the more complicated simulated annealing is analyzed in expectation.

2.1 The Local Search Algorithm and Best-Arm Identification

The greedy approach is an intuitive approach to the graph optimization problem: We start from a random node, and at each node and given a sampling budget, we explore its neighbors to find the best neighbor. The problem that we solve in each node can be viewed as a fixed-budget best-arm identification problem Audibert and Bubeck (2010). In a fixed-budget setting, the learner plays actions for a number of rounds in some fashion and returns an approximately optimal arm at the end. An example of an algorithm designed for the fixed-budget setting is the SUCCESSIVEREJECT algorithm of Audibert and Bubeck (2010).

Before describing the best-arm identification algorithm, we introduce some notation. Let K be an integer, $[K]$ be the set of arms in a bandit problem, μ_i be the mean value of arm i , and $\hat{\mu}_{i,t}$ be the empirical mean of arm i after t observations. Let $\mu_* = \max_{i \in [K]} \mu_i$ be the value of the optimal arm, and let Δ_i be the *optimality gap* of arm i , i.e. $\Delta_i = \mu_* - \mu_i$. Without loss of generality assume that $\Delta_1 \leq \Delta_2 \leq \dots \leq \Delta_K$. Let B be the budget. Define

$$H = \max_{i \in [K]} i \Delta_i^{-2},$$

$$\overline{\log}(K) = \frac{1}{2} + \sum_{i=2}^K \frac{1}{i},$$

$$B_k = \left\lceil \frac{1}{\overline{\log}(K)} \frac{B-K}{K+1-k} \right\rceil, B_0 = 0.$$

The SUCCESSIVEREJECT algorithm is a procedure with $K-1$ rounds where in each round one arm is eliminated. Let A_k be the remaining arms at the beginning of round $k \in [K]$. In round k , each arm $i \in A_k$ is selected for $B_k - B_{k-1}$ rounds. At the end of this round, the arm with the smallest empirical mean is removed.

Theorem 1 (Audibert and Bubeck (2010)). *The probability of error of the SUCCESSIVEREJECT algorithm with a budget of B is smaller than*

$$e_B = \frac{K(K-1)}{2} \exp\left(-\frac{B-K}{\overline{\log}(K)H}\right).$$

In the EXPLORE-DESCEND procedure, we use the above bandit algorithm to find the best neighbor in each round. More specifically, for a node $x \in \mathcal{V}$, we solve a best-arm identification problem with action values $\mu_i = f(x) - f(z_i)$, where z_i is the i th neighbor of node x . We then move to the chosen neighbor and repeat the process until the budget is consumed. We call this approach ‘‘Explore and Descend’’ and it is detailed in Algorithm 1.

Algorithm 1: Explore and Descend

Input : graph \mathcal{G} , starting node x_0 , budget T

Output : x such that $x = x^*$ with high probability

- 1 Per round budget $\{T_1, \dots, T_S\}$ such that $\sum_{s=1}^S T_s = T$;
 - 2 **for** $s = 0, \dots, S-1$ **do**
 - 3 $x_{s+1} \leftarrow \text{DescentOracle}(\mathcal{G}, x_s, T_s)$;
 - 4 **return** x_S
-

The algorithm depends on the subroutine $\text{DescentOracle}(\mathcal{G}, x_{s-1}, T_s)$. This subroutine is an implementation of the best-arm bandit algorithms described earlier with the decision set $\overline{\mathcal{N}}_x$ and budget T_s .

Let $p = \{y_0, y_1, \dots, y_S\}$ be a path such that $y_0 = x_0$ and $y_s = \arg\min_{y \in \overline{\mathcal{N}}_{y_{s-1}}} f(y)$. Assume S is large enough so that $y_S = x^*$ and w.l.o.g assume that $y_{S-1} \neq x^*$. Let $\Delta_1 = \arg\min_{y \in \mathcal{V} \setminus \{x^*\}} f(y) - f(x^*)$.

Lemma 2. *We have that $\Delta_{y_s, y_{s+1}} \geq (1+m)^{S-1-s} \Delta_1$.*

Proof. We prove by induction. The induction base holds for $s = S-1$ by definition of Δ_1 . Next, assume the statement holds for s . Let $\{y_s, z, \dots, z', y_S\}$ be a strongly convexity path. Thus,

$$\Delta_{y_s, z} \geq (1+m)^{S-1-s} \Delta_{z', y_S} \geq (1+m)^{S-1-s} \Delta_1.$$

On the other hand and by definition, $\Delta_{y_s, y_{s+1}} \geq \Delta_{y_s, z}$, which gives the desired result. \square

If we use the SUCCESSIVEREJECT algorithm as the subroutine in Algorithm 1, then the probability that the path p is not followed is smaller than

$$e'_T = \frac{d(d-1)}{2} \sum_{s=0}^{S-1} \exp\left(-\frac{T_s - d}{\overline{\log}(d)H_s}\right). \quad (1)$$

Here $H_s = \max_{i \in [d]} i \Delta_{i,s}^{-2}$ and $\Delta_{i,s}$ is the i th gap at node y_s . Notice that the error bound in Theorem 1 is vacuous when $T \leq K$. In contrast, Inequality 1 provides a meaningful error bound even in this small budget regime. If we use the loose upper bounds $H \leq K \Delta_1^{-2}$ and $H_s \leq d \Delta_{1,s}^{-2}$, then the error probabilities of the above methods are bounded as

$$e_T \leq \frac{K(K-1)}{2} \exp\left(-\frac{(T-K)\Delta_1^2}{K\overline{\log}(K)}\right),$$

$$e'_T \leq \frac{d(d-1)}{2} \sum_{s=0}^{S-1} \exp\left(-\frac{(T_s-d)\Delta_{1,s}^2}{d\overline{\log}(d)}\right).$$

Consider the large budget regime where $T = bK$ for some $b \geq 1$. Then we have

$$e_T \leq \frac{K(K-1)}{2} \exp\left(-\frac{(b-1)\Delta_1^2}{\overline{\log}(K)}\right).$$

Choose $T_s = a_s d$. By Lemma 2, $\Delta_{1,s} \geq (1+m)^{S-1-s} \Delta_1$. Thus,

$$e'_T \leq \frac{d(d-1)}{2} \sum_{s=0}^{S-1} \exp\left(-\frac{(a_s-1)(1+m)^{2(S-1-s)}\Delta_1^2}{\overline{\log}(d)}\right).$$

Choose a_s such that $(a_s-1)(1+m)^{2(S-1-s)} = (b-1)v$ for some $v \geq 1$. Then,

$$e'_T \leq \frac{d(d-1)S}{2} \exp\left(-\frac{(b-1)v\Delta_1^2}{\overline{\log}(d)}\right),$$

which is smaller than the upper bound on e_T as long as $d(d-1)S/2 \leq K(K-1)/2$. Because $\sum_{s=0}^{S-1} T_s \leq T$, we need to have

$$\begin{aligned} T &\geq d \sum_{s=0}^{S-1} a_s \\ &= dS + d \sum_{s=0}^{S-1} (a_s - 1) \\ &= dS + d(b-1)v \sum_{s=0}^{S-1} \frac{1}{(1+m)^{2s}} \\ &= dS + d(b-1)v \frac{1-g^S}{1-g}, \end{aligned}$$

where $g = 1/(1+m)^2$. In particular, the budget condition is satisfied if $T \geq dS + \frac{d(b-1)v}{1-g}$.

2.2 Nearly Convex Problems: Simulated Annealing

Given that we have access only to noisy observations, we consider the following Metropolis-Hastings Algorithm with exponential weights:

$$P(x \rightarrow y) = \begin{cases} \frac{1}{d} \min \left(1, e^{\gamma(\hat{f}_{x,s} - \hat{f}_{y,s})} \right) & \text{if } y \in \mathcal{N}_x \\ 1 - \sum_{z \in \mathcal{N}_x} P(x \rightarrow z) & \text{if } y = x \end{cases}$$

To simplify the analysis, we use a fixed time-independent inverse temperature, although in practice a time-varying inverse temperature might be preferable. We estimate each function value by $s = 2r\gamma^2 R^2$ samples. Next, we provide sample complexity bounds for the above procedure in expectation.

Theorem 2. *For a strongly convex function, and given a constant ε and with the choice of $\gamma = \frac{d}{c\alpha\varepsilon}$, after $t \geq \frac{\log(\alpha(f_{x^*} - f_{x_0})/(\varepsilon d))}{\log(d/(d-\alpha))}$ rounds, we have*

$$\mathbf{E}[f(x_{t+1}) - f(x^*)] \leq \varepsilon.$$

For a nearly convex function, let $\beta = 1 - \frac{\alpha e^{-c\gamma}}{d^{r+1}}$. With the choice of $\gamma = 1/c$ and after $t \geq \frac{r}{\log(1/\beta)} \log(F\alpha\gamma)$ rounds, we have that

$$\mathbf{E}(f(x_{t+r+1}) - f(x^*)) \leq \frac{3}{\alpha\gamma} d^{r+1} e^r.$$

Proof. We use $\{x_0 \rightarrow x_1 \rightarrow x_2 \rightarrow \dots\}$ to denote a path generated by the simulated annealing algorithm. Let $y(x_t)$ be the closest point in \mathcal{C} on a low energy path from x_t and let $r(x_t) = r_t$ be the distance to this point from x_t . Let \mathcal{P}'_{x_t} be the set of paths of length less than $r_t + 1$ and starting at x_t such that at least one node on the path is not the same as x_t . Consider the low energy sub-path starting at x_t :

$$p = \{z_1 = x_t \rightarrow z_2 \rightarrow \dots \rightarrow z_{r(x_t)} = y(x_t)\}.$$

Let $z(p)$ be the terminating node in a path p . Given that function f is (α, c, r) -nearly convex, and given that noise is R -sub-Gaussian, probability that this low-energy path is taken by the algorithm is

$$\begin{aligned} \mathbf{P}(p) &\leq \frac{1}{d^{r_t+1}} \mathbf{E} \left(e^{-\sum_{k=1}^{r_t} \gamma(\hat{f}(z_k) - \hat{f}(z_{k-1}))} \right) \\ &= \frac{1}{d^{r_t+1}} \mathbf{E} \left(\exp \left(-\sum_{k=1}^{r_t} \gamma(\hat{f}(z_k) - f(z_k)) \right. \right. \\ &\quad \left. \left. + \sum_{k=1}^{r_t} \gamma(\hat{f}(z_{k-1}) - f(z_{k-1})) \right. \right. \\ &\quad \left. \left. + \sum_{k=1}^{r_t} \gamma(f(z_{k-1}) - f(z_k)) \right) \right) \\ &\leq \frac{1}{d^{r_t+1}} e^{-\gamma(f(z(p)) - f(z_1)) + 2r_t\gamma^2 R^2/s}. \end{aligned}$$

Because p is a low energy path, we have $\mathbf{P}(p) \geq \frac{1}{d^{r_t+1}} e^{-cr_t\gamma}$. Let $c_t = f(y(x_t)) - f(x_t)$, and $P' = \sum_{p' \in \mathcal{P}'_x} \mathbf{P}(p')$. Notice that given x_t , c_t and $y(x_t)$ are deterministic. We write

$$\begin{aligned} \mathbf{E}(f(x_{t+r_t+1})|x_t) &\leq \mathbf{P}(p) (f(x_t) + c_t - \Delta_{y(x_t)}) \\ &\quad + \sum_{p' \in \mathcal{P}'_x} \mathbf{P}(p') f(z(p')) + (1 - \mathbf{P}(p) - P') f(x_t). \end{aligned}$$

The first term on the right side is related to the event that the state follows the low energy path to the state $y(x_t)$ for r_t rounds, and then goes to the best immediate neighbor at state $y(x_t)$. The second term is related to the event that the state is not the same as x_t after $r_t + 1$ rounds. Finally, the last term is related to the event that the state stays in x_t for the next $r_t + 1$ rounds. If $\Delta_{y(x_t)} \leq c_t$, then by Lemma 1, we already have

$$\begin{aligned} f(y(x_t)) - f(x^*) &\leq \frac{\Delta_{y(x_t)}}{\alpha} \leq \frac{c}{\alpha} \Rightarrow \\ f(x_t) - f(x^*) &\leq (r_t + 1/\alpha)c. \end{aligned}$$

Otherwise if $c_t \leq \Delta_{y(x_t)}$, we continue as follows:

$$\begin{aligned} \mathbf{E}(f(x_{t+r_t+1})|x_t) &\leq f(x_t) + \frac{1}{d^{r_t+1}} e^{-cr_t\gamma} (c_t - \Delta_{y(x_t)}) \\ &\quad + \sum_{p' \in \mathcal{P}'_x} \mathbf{P}(p') (f(z(p')) - f(x_t)) \\ &\leq f(x_t) + \frac{1}{d^{r_t+1}} e^{-cr_t\gamma} (c_t - \Delta_{y(x_t)}) \\ &\quad + \sum_{\substack{p' \in \mathcal{P}'_x, \\ f(z(p')) \geq f(x_t)}} \mathbf{P}(p') (f(z(p')) - f(x_t)) \\ &\leq \sum_{\substack{p' \in \mathcal{P}'_x, \\ f_{z(p')} \geq f_{x_t}}} \frac{e^{-\gamma(f(z(p')) - f(x_t)) + 2r_t\gamma^2 R^2/s}}{d^{r_t+1}} (f_{z(p')} - f_{x_t}) \\ &\quad + f(x_t) + \frac{1}{d^{r_t+1}} e^{-cr_t\gamma} (c_t - \Delta_{y(x_t)}) \\ &\leq f(x_t) + \frac{1}{d^{r_t+1}} e^{-cr_t\gamma} (c_t - \Delta_{y(x_t)}) + \frac{e^{-1+2r_t\gamma^2 R^2/s}}{\gamma}, \end{aligned}$$

where the last step follows from inequality $e^{-\gamma b} \leq 1/(\gamma e)$. By Lemma 1, $\Delta_{y(x_t)} \geq \alpha(f(y(x_t)) - f(x^*))$. Thus,

$$\begin{aligned} c_t - \Delta_{y(x_t)} &\leq f(y(x_t)) - f(x_t) - \alpha(f(y(x_t)) - f(x^*)) \\ &= -(f(x_t) - f(x^*)) - f(x^*) + \alpha f(x^*) \\ &\quad + (1 - \alpha)f(y(x_t)) \\ &= -(f(x_t) - f(x^*)) \\ &\quad + (1 - \alpha)(f(y(x_t)) - f(x^*)) \\ &\leq -\alpha(f(x_t) - f(x^*)) + (1 - \alpha)cr_t, \end{aligned}$$

where we used $f(y(x_t)) \leq f(x_t) + cr_t$ in the last step. Let $k = t + r_t + 1$. We have,

$$\mathbf{E}(f(x_k)|x_t) - f(x^*) \leq \left(1 - \frac{\alpha e^{-cr_t\gamma}}{d^{r_t+1}} \right) (f(x_t) - f(x^*))$$

$$\begin{aligned}
 & + \frac{cr_t e^{-cr_t \gamma} (1 - \alpha)}{d^{r_t+1}} + \frac{e^{-1+2r_t \gamma^2 R^2/s}}{\gamma} \\
 \leq & \left(1 - \frac{\alpha e^{-cr_t \gamma}}{d^{r_t+1}}\right) (f(x_t) - f(x^*)) \\
 & + \frac{c}{e(1 + \log d)} + \frac{e^{-1+2r_t \gamma^2 R^2/s}}{\gamma} \\
 \leq & \beta^{\lfloor \frac{t}{r} \rfloor} (f(x_1) - f(x^*)) \\
 & + \frac{1}{1 - \beta} \left(\frac{c}{e(1 + \log d)} + \frac{e^{-1+2r_t \gamma^2 R^2/s}}{\gamma} \right),
 \end{aligned}$$

where the second step holds by $r_t \leq r$, $\gamma c = 1$, and $\max_r c r e^{-\gamma c r} d^{-r} \leq (c/e)/(1 + \log d)$, and β is defined in the theorem statement. Using the fact that $s = 2r\gamma^2 R^2$ and $\gamma c = 1$, and a simple calculation shows that after $t \geq \frac{r}{\log(1/\beta)} \log(F\alpha\gamma)$ rounds, we have that

$$\mathbf{E}(f(x_{t+r+1}) - f(x^*)) \leq \frac{3d^{r+1}e^r}{\alpha\gamma}.$$

If $c = 0$, then we get that

$$\begin{aligned}
 \mathbf{E}(f(x_k)|x_t) - f(x^*) & \leq \left(1 - \frac{\alpha}{d^{r+1}}\right) (f(x_t) - f(x^*)) + \frac{1}{\gamma} \\
 & \leq \beta^{\lfloor \frac{t}{r} \rfloor} (f(x_1) - f(x^*)) + \frac{1}{(1 - \beta)\gamma}.
 \end{aligned}$$

After $t \geq \frac{r}{\log(1/\beta)} \log(F\alpha\gamma)$ rounds, we have that

$$\mathbf{E}(f(x_{t+r+1}) - f(x^*)) \leq \frac{2d^{r+1}}{\alpha\gamma}.$$

For a strongly convex function, following a similar argument, we have that

$$\mathbf{E}(f(x_{t+1}) - f(x^*)) \leq \left(1 - \frac{\alpha}{d}\right) (f(x_t) - f(x^*)) + \frac{1}{\gamma e}.$$

Thus, given ε and with the choice of γ in the theorem statement, after $t \geq \frac{\log(\alpha(f_{x^*} - f_{x_0})/(\varepsilon d))}{\log(d/(d-\alpha))}$ rounds, we have $\mathbf{E}[f(x_{t+1}) - f(x^*)] \leq \varepsilon$. \square

For nearly convex problems, the error bound in this theorem is meaningful only if parameter r is small. In our experiments data, this value is $r = 1$ or $r = 2$.

3 Experiments

We implemented EXPLORE-DESCEND and SIMULATED ANNEALING algorithms and compare them with SPECTRALBANDIT of Valko et al. (2014) and SUCCESSIVEREJECT of Audibert and Bubeck (2010). The SUCCESSIVEREJECT algorithm works by considering all nodes of the graph as a big multi-arm bandit problem. The SPECTRALBANDIT uses the adjacency matrix to calculate the Laplacian and

Eigenvectors. Both of these algorithms, therefore, require global information of the graph, while our algorithms assume only local information: from one node one can only access its neighbors.

We note that the SPECTRALBANDIT algorithm is originally designed to minimize the cumulative regret, which partly explains the poor performance in a best-arm identification problem. In the fixed budget best-arm identification setting, we run SPECTRALBANDIT until the budget is consumed and output the most frequently pulled arm, which is, in our experiments, better than taking the arm with the best empirical mean. The algorithm has a high time complexity because of its reliance on matrix operations on matrices and vectors of dimension n . As SPECTRALBANDIT does not scale well with the size of the graph due to its matrix operations, we generated a small synthetic graph in order to evaluate it.

3.1 Applications in Graph-Structured Best-Arm Identification: Synthetic Data

First, we evaluated different algorithms on synthetic graphs, which are generated as follows. Each node of the graph is a point (x, y) , where $x, y \in \{-D, \dots, -1, 0, 1, \dots, D\}$ for some $D \in \mathbb{Z}^+$. Each node is connected to all of its eight immediate neighbors on the plain. Additionally, random edges are added in, such that the degree of each node is 15. The mean function value is $f(x, y) = 0.8(1 - \frac{x^2+y^2}{2D^2}) \in [0, 0.8]$. This mean is unknown to the algorithms, i.e. when the algorithm requests the value of $f(x, y)$, it is returned with a stochastic value: 1 with probability $p = f(x, y)$ and 0 with probability $1 - p$. It is easy to see that this graph is concave by Definition 1.

The results of the experiment is presented in Figure 1. The performance measure is the average sub-optimality gap, i.e. $f(x^*) - f(\hat{x})$, where \hat{x} is the solution returned by the algorithm. The average is taken over 1000 trials. Overall, our algorithms significantly outperform SPECTRALBANDIT and SUCCESSIVEREJECT both in term of time and sub-optimality gap, especially when the budget is smaller than the graph size, which is our intended setting. For SIMULATED ANNEALING, we use a fixed inverse temperature $\gamma = 250$. The result could be further improved by optimizing a schedule for this parameter. Interestingly, the number of pulls for each function evaluation also has significant impact on the performance, which can be seen from the plots for *Sim Annealing 1* (1 pull) and *Sim Annealing 5* (5 pulls) in Figure 1. As for EXPLORE-DESCEND, we simply allocate the budget equally for each node in the descending path, with the maximum path length set to 4 for $D = 10$ and 20 for $D = 100$. This algorithm is the fastest and also offers the best performance. Source code is available at <https://github.com/tan1889/graph-opt>

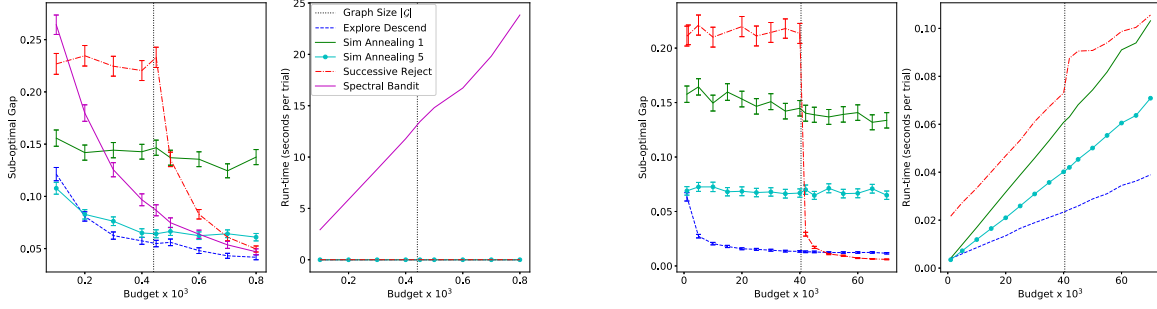


Figure 1: Comparing the performance of different algorithms on synthetic data. Two left figures: Small synthetic graph $D = 10$ (441 nodes). Two right figures: Large synthetic graph $D = 100$ (40401 nodes). Figures show average sub-optimality gaps over 1000 trials and run-time per trial of the algorithms. Same legend (on the second figure) for all figures.




Figure 2: Comparing the performance of algorithms on web document reranking data. Left: Average sub-optimal gap over 1000 trials. Right: Run-time (s) per trial. Same legend for both figures.

3.2 Applications in Graph-Structured Best-Arm Identification: Web Document Reranking

To demonstrate the performance of our algorithm on real-world non-concave problems, we used data from Yandex Personalized Web Search Challenge to build a graph as follows. Each query forms a graph, whose nodes correspond to lists of 5 items (documents) for the given query. Two nodes are connected if they have 4 or more items in common at the same positions. The value of a node is a Bernoulli random variable with mean equal to the click-through rate of the associated list. The goal is to find the node with maximum click-through rate, i.e. the most relevant list. We chose the query that generated the largest possible graph (query no. 8107157) of 4514 nodes. As there were many small partitions in this graph, we took the largest partition as the input for our experiment. The resulting graph has 3992 nodes with degree varying from 1 to 171 (mean=35) and a maximum function value at 0.747.

For non-concave graphs, EXPLORE-DESCEND needs to

make multiple restarts. We set the number of restarts to $1 + budget/1000$ and allocate the *budget* equally between all restarts, then, for each restart, equally between each node in the path. All other parameters are the same as before.

The results of the experiment is presented in Figure 2. In the intended setting, our algorithms significantly outperform SUCCESSIVE REJECT. For very small budget, SIMULATED ANNEALING is better than EXPLORE-DESCEND, but this is reversed as the budget gets bigger. Additionally, for this graph we don't see the big advantage of *Sim Annealing 5* over *Sim Annealing 1* as was the case before. Outside of the intended setting, SUCCESSIVE REJECT quickly becomes the best algorithm when the budget gets larger than the graph size. Although, this algorithm requires global information about the graph, which may not be always feasible.

3.3 Applications in Graph-Based Nearest-Neighbor Classification

In this section, we use the proposed graph-based optimization methods in a graph-based nearest neighbor search problem. The graph-based nearest neighbor method takes a training set, constructs a proximity graph on this set, and when queried in the test phase, performs a hill-climbing search to find an approximate nearest neighbor. More details are given in Appendix A. This procedure is particularly well suited to big-data problems; In such problems, points will have close neighbors, and so the geodesic distance with respect to even a simple metric such as Euclidean metric should provide an appropriate metric. Further, and as we will show, computational complexity of the graph-based method in the test phase scales gracefully with the size of the training set, a property that is particularly important in big data applications. The intuition is that, in big-data problems, although a descend path might be longer, the objective function is generally more smooth and hence easier to minimize.

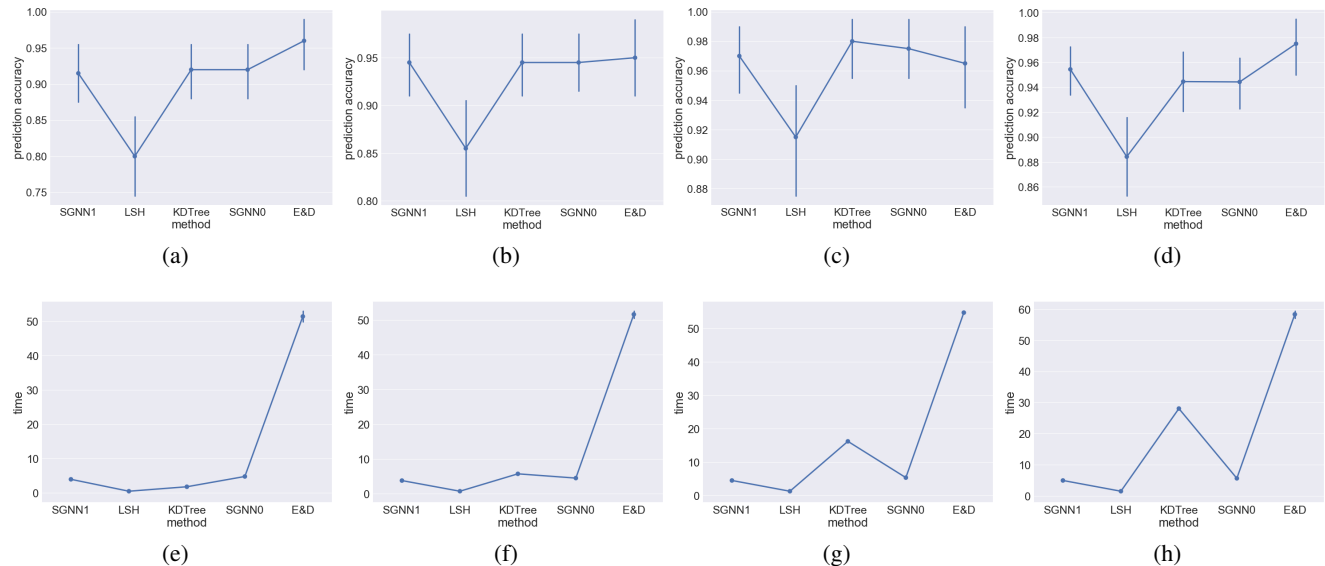


Figure 3: Prediction accuracy and running time of different methods on MNIST dataset as the size of training set increases. (a,e) Using 25% of training data. (b,f) Using 50% of training data. (c,g) Using 75% of training data. (d,h) Using 100% of training data.

We apply the local search and simulated annealing methods along with an additional smoothing to the problem of nearest neighbor search. For SIMULATED ANNEALING, we call the resulting algorithm SGNN for Smoothed Graph-based Nearest Neighbor search. The pseudocode of the algorithm and more details are in Appendix A. The Explore-Descend is denoted by E&D in these experiments. We compared the proposed methods with the state-of-the-art nearest neighbor search methods (KDTree and LSH) in two image classification problems (MNIST and COIL-100). In an approximate nearest neighbor search problem, it is crucial to have sub-linear time complexity, and thus SPECTRALBANDIT and SUCCESSIVEREJECT are not applicable here.

Figure 3 (a-d) shows the accuracy of different methods on different portions of MNIST dataset. The graphs in these experiments are not concave, but ($\alpha = 0.3, c = 0, r = 2$)—nearly concave by Definition 2. The results for COIL-100 are shown in Appendix A. As the size of training set increases, the prediction accuracy of all methods improve. Figure 3 (e-h) shows that the test phase runtime of the SGNN method has a more modest growth for larger datasets. In contrast, KDTree becomes much slower for larger training datasets. The LSH method generally performs poorly, and it is hard to make it competitive with other methods. When using all training data, the SGNN method has roughly the same accuracy, but it has less than 20% of the test phase runtime of KDTree.

4 Conclusions and Future Work

We studied sample complexity of stochastic optimization of graph functions. We defined a notion of convexity for graph functions, and we showed that under the convexity condition, a greedy algorithm and the simulated annealing enjoy sample complexity bounds that are independent of the size of the graph. An interesting future work is the study of cumulative regret in this setting.

We showed effectiveness of the proposed techniques in a web document re-ranking problem as well as a graph-based nearest neighbor search problem. The computational complexity of the resulting nearest neighbor method scales gracefully with the size of the dataset, which is particularly appealing in big-data applications. Further quantification of this property remains for future work.

Acknowledgement

TN was supported by the Australian Research Council Centre of Excellence for Mathematical and Statistics Frontiers (ACEMS).

References

- Yasin Abbasi-Yadkori, David Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *NIPS*, 2011.
- Sunil Arya and David M. Mount. Algorithms for fast vector quantization. In *IEEE Data Compression Conference*, 1993.
- Jean-Yves Audibert and Sébastien Bubeck. Best Arm Identification in Multi-Armed Bandits. In *COLT - 23th Conference on Learning Theory - 2010*, page 13 p., Haifa, Israel, June 2010. URL <https://hal-enpc.archives-ouvertes.fr/hal-00654404>.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *JMLR*, (3):397–422, 2002.
- C. Bouttier and I. Gavra. Convergence rate of a simulated annealing algorithm with noisy observations. *ArXiv e-prints*, March 2017.
- M.R. Brito, E.L. Chavez, A.J. Quiroz, and J.E. Yukich. Connectivity of the mutual k-nearest-neighbor graph in clustering and outlier detection. *Statistics & Probability Letters*, 35(1):33–42, 1997.
- Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. Online optimization in ξ -armed bandits. In *NIPS*, 2009.
- Alexandra Carpentier and Michal Valko. Simple regret for infinitely many armed bandits. In *International Conference on Machine Learning*, 2015.
- Jie Chen, Haw ren Fang, and Yousef Saad. Fast approximate knn graph construction for high dimensional data via recursive lanczos bisection. *Journal of Machine Learning Research*, 10:1989–2012, 2009.
- P. R. Christian and G. Casella. *Monte Carlo Statistical Methods*. Springer New York, 1999.
- Michael Connor and Piyush Kumar. Fast construction of k-nearest neighbor graphs for point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):599–608, 2010.
- Wei Dong, Moses Charikar, and Kai Li. Efficient k-nearest neighbor graph construction for generic similarity measures. In *WWW*, 2011.
- David Eppstein, Michael S. Paterson, , and F. Frances Yao. On nearest-neighbor graphs. *Discrete & Computational Geometry*, 17(3):263–282, 1997.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *JMLR*, 7:1079–1105, 2006.
- Jean-Bastien Grill, Michal Valko, and Rémi Munos. Black-box optimization of noisy functions with unknown smoothness. In *NIPS*, 2015.
- Kiana Hajebi, Yasin Abbasi-Yadkori, Hossein Shahbazi, and Hong Zhang. Fast approximate nearest-neighbor search with k-nearest neighbor graph. In *IJCAI*, 2011.
- Kevin Jamieson and Robert Nowak. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *2014 48th Annual Conference on Information Sciences and Systems, CISS 2014*, pages 1–6, 2014.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *JMLR*, 17:1–42, 2016.
- Shie Mannor and John N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *JMLR*, 5:623–648, 2004.
- Gary L. Miller, Shang-Hua Teng, William Thurston, , and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *Journal of the ACM*, 44(1):1–29, 1997.
- Erion Plaku and Lydia E. Kavradi. Distributed computation of the knn graph for large high-dimensional point sets. *Journal of Parallel and Distributed Computing*, 67(3):346–359, 2007.
- R.Y. Rubinstein. Optimization of computer simulation models with rare events. *European Journal of Operations Research*, (99):89–112, 1997.
- M. Sreenivas Pydi, V. Jog, and P.-L. Loh. Graph-based ascent algorithms for function maximization. *ArXiv e-prints*, 2018.
- Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010.
- Michal Valko, Rémi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, 2014.
- Jing Wang, Jingdong Wang, Gang Zeng, Zhuowen Tu, Rui Gan, and Shipeng Li. Scalable k-nn graph construction for visual descriptors. In *CVPR*, 2012.