

Volumetric Spanners: an Efficient Exploration Basis for Learning

Elad Hazan
Technion - Israel Inst. of Tech.
ehazan@ie.technion.ac.il

Zohar Karnin
Yahoo Labs
zkarnin@ymail.com

Raghu Meka
Microsoft Research
meka@microsoft.com

Abstract

Numerous machine learning problems require an *exploration basis* - a mechanism to explore the action space. We define a novel geometric notion of exploration basis with low variance called volumetric spanners, and give efficient algorithms to construct such bases.

We show how efficient volumetric spanners give rise to an efficient and near-optimal regret algorithm for bandit linear optimization over general convex sets. Previously such results were known only for specific convex sets, or under special conditions such as the existence of an efficient self-concordant barrier for the underlying set.

1 Introduction

A fundamental difficulty in machine learning is environment exploration. A prominent example is the famed multi-armed bandit (MAB) problem, in which a decision maker iteratively chooses an action from a set of available actions and receives a payoff, without observing the payoff of all other actions she could have taken. The MAB problem displays an exploration-exploitation tradeoff, in which the decision maker trades exploring the action space vs. exploiting the knowledge already obtained to pick the best arm.

Another example in which environment exploration is crucial, or perhaps the main point, is active learning and experiment design. In these fields it is important to correctly identify the most informative queries so as to efficiently construct a solution.

Exploration is hardly summarized by picking an action uniformly at random¹. Indeed, sophisticated techniques from various areas of optimization, statistics and convex geometry have been applied to designing ever better exploration algorithms. To mention a few: [4] devise the notion of *barycentric spanners*, and use this construction to give the first low-regret algorithms for complex decision problems such as online routing. [1] use self-concordant barriers to build an efficient exploration strategy for convex sets in Euclidean space. [9] apply tools from convex geometry, namely the John ellipsoid to construct optimal-regret algorithms for bandit linear optimization, albeit not always efficiently.

In this paper we consider a generic approach to exploration, and quantify what efficient exploration with *low variance* requires in general. Given a set in Euclidean space, a low-variance exploration basis is a subset with the following property: given noisy estimates of a linear function over the basis, one can construct an estimate for the linear function over the entire set without increasing the variance of the estimates.

By definition, such low variance exploration bases are immediately applicable to noisy linear regression: given a low-variance exploration basis, it suffices to learn the function values only over the basis in order to interpolate the value of the underlying linear regressor over the entire decision set. This fact can be used for active learning as well as for the exploration component of a bandit linear optimization algorithm.

Henceforth we define a novel construction for a low variance exploration basis called **volumetric spanners** and give efficient algorithms to construct them. We further investigate the convex

¹although uniform sampling does at times work exceptionally well, e.g. [10]

geometry implications of our construction, and define the notion of a **minimal volumetric ellipsoid** of a convex body. We give structural theorems on the existence and properties of these ellipsoids, as well as constructive algorithms to compute them in several cases.

We complement our findings with an application to machine learning, in which we advance a well-studied open problem that has exploration as its core difficulty: an efficient and near-optimal regret algorithm for bandit linear optimization (BLO). We expect that volumetric spanners and volumetric ellipsoids can be useful elsewhere in experiment design and active learning. We briefly discuss the application to BLO next.

Bandit Linear Optimization Bandit linear optimization (BLO) is a fundamental problem in decision making under uncertainty that efficiently captures structured action sets. The canonical example is that of online routing in graphs: a decision maker iteratively chooses a path in a given graph from source to destination, the adversary chooses lengths of the edges of the graph, and the decision maker receives as feedback the length of the path she chose but no other information (see [4]). Her goal over many iterations is to attain an average travel time as short as that of the best fixed shortest path in the graph.

This decision problem is readily modeled in the “experts” framework, albeit with efficiency issues: the number of possible paths is potentially exponential in the graph representation. The BLO framework gives an efficient model for capturing such structured decision problems: iteratively a decision maker chooses a point in a convex set and receives as a payoff an adversarially chosen linear cost function. In the particular case of online routing, the decision set is taken to be the s-t-flow polytope, which captures the convex hull of all source-destination shortest paths in a given graph, and has a succinct representation with polynomially many constraints and low dimensionality. The linear cost function corresponds to a weight function on the graphs edges, where the length of a path is defined as the sum of weights of its edges.

The BLO framework captures many other structured problems efficiently, e.g., learning permutations, rankings and other examples (see [1]). As such, it has been the focus of much research in the past few years. The reader is referred to the recent survey of [8] for more details on algorithmic results for BLO. Let us remark that certain online bandit problems do not immediately fall into the convex BLO model that we address, such as combinatorial bandits studied in [10].

In this paper we contribute to the large literature on the BLO model by giving the first polynomial-time and near optimal-regret algorithm for BLO over general convex decision sets; see Section 6 for a formal statement . Previously efficient algorithms, with non-optimal-regret, were known over convex sets that admit an efficient self-concordant barrier [1], and optimal-regret algorithms were known over general sets [9] but these algorithms were not computationally efficient. Our result, based on volumetric spanners, is able to attain the best of both worlds.

1.1 Volumetric Ellipsoids and Spanners

We now describe the main convex geometric concepts we introduce and use for low variance exploration. To do so we first review some basic notions from convex geometry.

Let \mathbb{R}^d be the d -dimensional vector space over the reals. Given a set of vectors $S = \{v_1, \dots, v_t\} \subset \mathbb{R}^d$, we denote by $\mathcal{E}(S)$ the ellipsoid defined by S :

$$\mathcal{E}(S) = \left\{ \sum_{i \in S} \alpha_i v_i : \sum_i \alpha_i^2 \leq 1 \right\}.$$

By abuse of notation, we also say that $\mathcal{E}(S)$ is *supported* on the set S .

Ellipsoids play an important role in convex geometry and specific ellipsoids associated with a convex body have been used in previous works in machine learning for designing good exploration bases for convex sets $\mathcal{K} \subseteq \mathbb{R}^d$. For example, the notion of *barycentric spanners* which were introduced in the seminal work of [4] corresponds to looking at the ellipsoid of maximum volume supported by exactly d points from \mathcal{K}^2 . Barycentric Spanners have since been used as

²While the definition of [4] is not phrased as such, their analysis shows the existence of barycentric spanners by looking at the maximum volume ellipsoid.

an exploration basis in several works: In [13] for online bandit linear optimization, in [6] for a high probability counterpart of the online bandit linear optimization, in [17] for repeated decision making of approximable functions and in [12] for a stochastic version of bandit linear optimization. Another example is the work of [9] which looks at the minimum volume enclosing ellipsoid (MVEE) also known as the John ellipsoid (see Section 2 for more background on this fundamental object from convex geometry).

As will be clear soon, our definition of a *minimal volumetric ellipsoid* enjoys the best properties of the examples above enabling us to get more efficient algorithms. Similar to barycentric spanners, it is supported by a small (quasi-linear) set of points of \mathcal{K} . Simultaneously and unlike the barycentric counterpart, the volumetric ellipsoid contains the body \mathcal{K} , a property shared with the John ellipsoid.

Definition 1.1 (Volumetric Ellipsoids). *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a set in Euclidean space. For $S \subseteq \mathcal{K}$, we say that $\mathcal{E}(S)$ is a volumetric ellipsoid for \mathcal{K} if it contains \mathcal{K} . We say that $\mathcal{E}_{\mathcal{K}} = \mathcal{E}(S)$ is a minimal volumetric ellipsoid if it is a containing ellipsoid defined by a set of minimal cardinality*

$$\mathcal{E}_{\mathcal{K}} \in \min_{|S|} \{ \mathcal{E}(S) \text{ such that } S \subseteq \mathcal{K} \subseteq \mathcal{E}(S) \}.$$

We say that $|S|$ is the order of the minimal volumetric ellipsoid or of \mathcal{K} denoted $\mathbf{order}(\mathcal{K})$.

We discuss various geometric properties of volumetric ellipsoids later. For now, we focus on their utility in designing efficient exploration bases. To make this concrete and to simplify some terminology later on (and also to draw an analogy to barycentric spanners), we introduce the notion of *volumetric spanners*. Informally, these correspond to sets S that span all points in a given set with coefficients having Euclidean norm at most one. Formally:

Definition 1.2. *Let $\mathcal{K} \subseteq \mathbb{R}^d$ and let $S \subseteq \mathcal{K}$. We say that S is a volumetric spanner for \mathcal{K} if $\mathcal{K} \subseteq \mathcal{E}(S)$.*

Clearly, a set \mathcal{K} has a volumetric spanner of cardinality at most t if and only if $\mathbf{order}(\mathcal{K}) \leq t$.

Our goal in this work is to bound the order of arbitrary sets. A priori, it is not even clear if there is a universal bound (depending only on the dimension and not on the set) on the order S for compact sets \mathcal{K} . However, barycentric spanners and the John ellipsoid show that the order of any compact set in \mathcal{R}^d is at most $O(d^2)$. Our main structural result in convex geometry gives a nearly optimal linear bound on the order of all sets.

Theorem 1.1. *Any compact set $\mathcal{K} \subseteq \mathbb{R}^d$ admits a volumetric ellipsoid of order at most $12d$. Further, if $\mathcal{K} = \{v_1, \dots, v_n\}$ is a discrete set, then a volumetric ellipsoid for \mathcal{K} of order at most $12d$ can be constructed in time $O(n^{3.5} + dn^3 + d^5)$.*

We emphasize the last part of the above theorem giving an algorithm for finding volumetric spanners of small size; this could be useful in using our structural results for algorithmic purposes. We also give a different algorithmic construction for the discrete case (a set of n vectors) in Section 5. While being sub-optimal by logarithmic factors (gives an ellipsoid of order $O(d(\log d)(\log n))$) this alternate construction has the advantage of being simpler and more efficient to compute.

1.2 Approximate Volumetric Spanners

Theorem 1.1 shows the existence of good volumetric spanners and also gives an efficient algorithm for finding such a spanner in the discrete case, i.e. when \mathcal{K} is finite and given explicitly. However, the existence proof uses the John ellipsoid in a fundamental way and it is not known how to compute (even approximately) the John ellipsoid efficiently for the case of general convex bodies. For such computationally difficult cases, we introduce a natural relaxation of volumetric ellipsoids which can be computed efficiently for a bigger class of bodies and is similarly useful. The relaxation comes from requiring that the ellipsoid of small support contain all but an ε fraction of the points

³We note that our definition allows for multi-sets, meaning that S may contain the same vector more than once

in \mathcal{K} (under some distribution). In addition, we also require that the measure of points decays exponentially fast w.r.t their $\mathcal{E}(S)$ -norm (see precise definition in next section); this property gives us tighter control on the set of points not contained in the ellipsoid. When discussing a measure over the points of a body the most natural one is the uniform distribution over the body. However, it makes sense to consider other measures as well and our approximation results in fact hold for a wide class of distributions.

Definition 1.3. Let $S \subseteq \mathcal{R}^d$ be a set of vectors and let V be the matrix whose columns are the vectors of S . We define the semi-norm

$$\|x\|_{\mathcal{E}(S)} = \sqrt{x^\top (VV^\top)^{-1}x},$$

where $(VV^\top)^{-1}$ is the Moore-Penrose pseudo-inverse of VV^\top . Let \mathcal{K} be a convex set in \mathcal{R}^d and p a distribution over it. Let $\varepsilon > 0$. A (p, ε) -exp-volumetric spanner of \mathcal{K} is a set $S \subseteq \mathcal{K}$ such that for any $\theta > 1$

$$\Pr_{x \sim p} [\|x\|_{\mathcal{E}(S)} \geq \theta] \leq \varepsilon^{-\theta}.$$

We prove that spanners as above can be efficiently obtained for any log-concave distribution:

Theorem 1.2. Let \mathcal{K} be a convex set in \mathcal{R}^d and p a log-concave distribution over it. By sampling $O(d + \log^2(1/\varepsilon))$ i.i.d. points from p one obtains, w.p. at least $1 - \exp(-\sqrt{\max\{\log(1/\varepsilon), d\}})$, a (p, ε) -exp-volumetric spanner for \mathcal{K} .

1.3 Structure of the paper

In the next section we list the preliminaries and known results from measure concentration, convex geometry and online learning that we need. In Section 3 we show the existence of small size volumetric spanners. In sections 4 and 5 we give efficient constructions of volumetric spanners for continuous and discrete sets, respectively. We then proceed to describe the application of our geometric results to bandit linear optimization in Section 6.

2 Preliminaries

We now describe several preliminary results we need from convex geometry and linear algebra. We start with some notation:

- A matrix $A \in \mathbb{R}^{d \times d}$ is positive semi-definite (PSD) when for all $x \in \mathcal{R}^d$ it holds that $x^\top Ax \geq 0$. Alternatively, when all of its eigenvalues are non-negative. We say that $A \succeq B$ if $A - B$ is PSD.
- Given an ellipsoid $\mathcal{E}(S) = \{\sum_i \alpha_i v_i : \sum_i \alpha_i^2 \leq 1\}$, we shall use the notation $\|x\|_{\mathcal{E}(S)} \triangleq \sqrt{x^\top (VV^\top)^{-1}x}$ to denote the (Minkowski) semi-norm defined by the ellipsoid, where V is the matrix with the vectors v_i 's as columns.
- Throughout, we denote by I_d the $d \times d$ identity matrix.

We next describe properties of the John ellipsoid which plays an important role in our proofs.

2.1 The Fritz John Ellipsoid

Let $\mathcal{K} \subseteq \mathbb{R}^n$ be an arbitrary convex body. Then, the **John ellipsoid** of \mathcal{K} is the minimum volume ellipsoid containing \mathcal{K} . This ellipsoid is unique and its properties have been the subject of important study in convex geometry since the seminal work of [16] (see [5] and [15] for historic information).

Suppose that we have linearly transformed \mathcal{K} such that its minimum volume enclosing ellipsoid (MVEE) is the unit sphere; in convex geometric terms, \mathcal{K} is in *John's position*. The celebrated

decomposition theorem of [16] gives a characterization of when a body is in John's position and will play an important role in our constructions (the version here is from [5]).

Below we consider only symmetric convex sets, i.e. sets that admit the following property: if $x \in \mathcal{K}$ then also $-x \in \mathcal{K}$. The sets encountered in machine learning applications are most always symmetric, since estimating a linear function on a point x is equivalent to estimating it on its polar $-x$, and negating the outcome.

Theorem 2.1 ([5]). *Let $\mathcal{K} \in \mathcal{R}^d$ be a symmetric set such that its MVEE is the unit sphere. Then there exist $m \leq d(d+1)/2 - 1$ contact points of \mathcal{K} and the sphere u_1, \dots, u_m and non-negative weights c_1, \dots, c_m such that $\sum_i c_i u_i = 0$ and $\sum c_i u_i u_i^T = I_d$.*

The contact points of a convex body have been extensively studied in convex geometry and they also make for an appealing exploration basis in our context. Indeed, [9] use these contact points to attain an optimal-regret algorithm for BLO. Unfortunately we know of no efficient algorithm to compute, or even approximate, the John ellipsoid for a general convex set, thus the results of [9] do not yield efficient algorithms for BLO.

For our construction of volumetric spanners we need to compute the MVEE of a discrete symmetric set, which can be done efficiently. We make use of the following (folklore) result:

Theorem 2.2 (folklore, see e.g. [18, 11]). *Let $\mathcal{K} \subseteq \mathcal{R}^d$ be a set of n points. It is possible to compute an ε -approximate MVEE for \mathcal{K} (an enclosing ellipsoid of volume at most $(1 + \varepsilon)$ that of the MVEE) that is also supported on a subset of \mathcal{K} in time $O(n^{3.5} \log \frac{1}{\varepsilon})$.*

The run-time above is attainable via the ellipsoid method or path-following interior point methods (see references in theorem statement). An approximation algorithm rather than an exact one is necessary in a real-valued computation model, and the logarithmic dependence on the approximation guarantee is as good as one can hope for in general.

The above theorem allows us to efficiently compute a linear transformation such that the MVEE of \mathcal{K} is essentially the unit sphere. We can then use linear programming to compute an approximate decomposition like in John's theorem as follows.

Theorem 2.3. *Let $\{x_1, \dots, x_n\} = \mathcal{K} \subseteq \mathcal{R}^d$ be a set of n points and assume that:*

1. \mathcal{K} is symmetric.
2. The John Ellipsoid of \mathcal{K} is the unit sphere.

Then it is possible, in $O((\sqrt{n} + d)n^3)$ time, to compute non-negative weights c_1, \dots, c_n such that (1) $\sum_i c_i x_i = 0$ and (2) $\sum_{i=1}^n c_i x_i x_i^T = I_d$.

Proof. Denote the MVEE of \mathcal{K} by \mathcal{E} and let V be its corresponding $d \times d$ matrix, meaning V is such that $\|y\|_{\mathcal{E}}^2 = y^T V^{-1} y \leq 1$ for all $y \in \mathcal{K}$. By our assumptions $I_d = V$.

As \mathcal{K} is symmetric and its MVEE is the unit ball, according to Theorem 2.1, there exist $m \leq d(d+1)/2 - 1$ contact points u_1, \dots, u_m of \mathcal{K} with the unit ball and a vector $c' \in \mathcal{R}^m$ such that $c' \geq 0$, $\sum c'_i = d$ and $\sum c'_i u_i u_i^T = I_d$. It follows that the following LP has a feasible solution: Find $c \in \mathcal{R}^n$ such that $c \geq 0$, $\sum c_i \leq d$ and $\sum c_i u_i u_i^T = I_d$. The described LP has $O(n + d^2)$ constraints and n variables. It can thus be solved in time $O(d + \sqrt{n})n^3$ via interior point methods. \square

We next state the results from probability theory that we need.

2.2 Distributions and Measure Concentration

For a set \mathcal{K} , let $x \sim \mathcal{K}$ denote a uniformly random vector from \mathcal{K} .

Definition 2.1. *A distribution over \mathcal{R}^d is log-concave if its probability distribution function (pdf) p is such that for all $x, y \in \mathcal{R}^d$ and $\lambda \in [0, 1]$,*

$$p(\lambda x + (1 - \lambda)y) \geq p(x)^\lambda p(y)^{1-\lambda}$$

Two log-concave distributions of interest to us are (1) the uniform distribution over a convex body and (2) a distribution over a convex body where $p(x) \propto \exp(L^\top x)$, where L is some vector in \mathcal{R}^d . The following result shows that given oracle access to the pdf of a log-concave distribution we can sample from it efficiently. An oracle to a pdf accepts as input a point $x \in \mathcal{R}^d$ and returns the value $p(x)$.

Lemma 2.1 ([19], Theorems 2.1 and 2.2). *Let p be a log-concave distribution over \mathbb{R}^d and let $\delta > 0$. Then, given oracle access to p , i.e. and oracle computing its pdf for any point in \mathbb{R}^d , there is an algorithm which approximately samples from p such that:*

1. *The total variation distance between the produced distribution and the distribution defined by p is no more than δ . That is, the difference between the probabilities of any event in the produced and actual distribution is bounded by δ .*
2. *The algorithm requires a pre-processing time of $\tilde{O}(d^5)$.*
3. *After pre-processing, each sample can be produced in time $\tilde{O}(d^4/\delta^4)$, or amortized time of $\tilde{O}(d^3/\delta^4)$ if more than d samples are needed.*

Definition 2.2 (Isotropic position). *A random variable x is said to be in isotropic position (or isotropic) if*

$$\mathbf{E}[x] = 0, \quad \mathbf{E}[xx^\top] = I_d.$$

A set $\mathcal{K} \subseteq \mathbb{R}^d$ is said to be in isotropic position if $x \sim \mathcal{K}$ is isotropic. Similarly, a distribution p is isotropic if $x \sim p$ is isotropic.

Henceforth we use several results regarding the concentration of log-concave isotropic random vectors. We use slight modification where the center of the distribution is not necessarily in the origin. For completeness we present the proof of the modified theorems in Appendix A

Theorem 2.4 (Theorem 4.1 in [2]). *Let p be a log-concave distribution over \mathcal{R}^d in isotropic position. There is a constant C such that for all $t, \delta > 0$, the following holds for $n = \frac{Ct^4 d \log^2(t/\delta)}{\delta^2}$. For independent random vectors $x_1, \dots, x_n \sim p$, with probability at least $1 - \exp(-t\sqrt{d})$,*

$$\left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^\top - I_d \right\| \leq \delta.$$

Corollary 2.1. *Let p be a log-concave distribution over \mathcal{R}^d and $x \sim p$. Assume that x is such that $\mathbf{E}[xx^\top] = I_d$. Then, there is a constant C such that for all $t \geq 1, \delta > 0$, the following holds for $n = \frac{Ct^4 d \log^2(t/\delta)}{\delta^2}$. For independent random vectors $x_1, \dots, x_n \sim p$, with probability at least $1 - \exp(-t\sqrt{d})$,*

$$\left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^\top - I_d \right\| \leq \delta.$$

Theorem 2.5 (Theorem 1.1 in [14]). *There exist constants c, C such that the following holds. Let p be a log-concave distribution over \mathcal{R}^d in isotropic position and let $x \sim p$. Then, for all $\theta \geq 0$,*

$$\Pr \left[\left| \|x\| - \sqrt{d} \right| > \theta\sqrt{d} \right] \leq C \exp(-c\sqrt{d} \cdot \min(\theta^3, \theta)).$$

Corollary 2.2. *Let p be a log-concave distribution over \mathcal{R}^n and let $x \sim p$. Assume that $\mathbf{E}[xx^\top] = I_d$. Then for some universal C, c it holds for any $\theta \geq 3$ that*

$$\Pr \left[\|x\| > \theta\sqrt{d} \right] \leq C \exp \left(-c\theta\sqrt{d} \right)$$

The following theorem provides a concentration bound for random vectors originating from an arbitrary distribution.

Theorem 2.6 ([20]). *Let X be a vector-valued random variable over \mathbb{R}^d with $\mathbf{E}[XX^\top] = \Sigma$ and $\|\Sigma^{-1/2}X\|^2 \leq R$. Then, for independent samples X_1, \dots, X_M from X , and $M \geq CR \log(R/\epsilon)/\epsilon^2$ the following holds with probability at least $1/2$:*

$$\left\| \frac{1}{M} \sum_{i=1}^M X_i - \Sigma \right\| \leq \epsilon \|\Sigma\|.$$

Finally, we also make use of barycentric spanners in our application to BLO and we briefly describe them next.

2.3 Barycentric Spanners

Definition 2.3 ([4]). *A barycentric spanner of $\mathcal{K} \subseteq \mathcal{R}^d$ is a set of d points $S = \{u_1, \dots, u_d\} \subseteq \mathcal{K}$ such that any point in \mathcal{K} may be expressed as a linear combination of the elements of S using coefficients in $[-1, 1]$. For $C > 1$, S is a C -approximate barycentric spanner of \mathcal{K} if any point in \mathcal{K} may be expressed as a linear combination of the elements of S using coefficients in $[-C, C]$*

In [4] it is shown that any compact set has a barycentric spanner. Moreover, they show that given an oracle with the ability to solve linear optimization problems over \mathcal{K} , an approximate barycentric spanner can be efficiently obtained. In the following sections we will use this constructive result.

Theorem 2.7 (Proposition 2.5 in [4]). *Let \mathcal{K} be a compact set in \mathcal{R}^d that is not contained in any proper linear subspace. Given an oracle for optimizing linear functions over \mathcal{K} , for any $C > 1$, it is possible to compute a C -approximate barycentric spanner for \mathcal{K} , using $O(d^2 \log_C(d))$ calls to the optimization oracle.*

3 Existence of Volumetric Ellipsoids and Spanners

In this section we show the existence of low order volumetric ellipsoids proving our main structural result, Theorem 1.1. Before we do so, we first state a few simple properties of volumetric ellipsoids (recall the definition of **order** from Definition 1.1):

- The definition of **order** is linear invariant: for any invertible linear transformation $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $K \subseteq \mathbb{R}^d$, $\mathbf{order}(K) = \mathbf{order}(TK)$.

Proof. Let $S \subseteq K$ be such that $K \subseteq \mathcal{E}(S)$. Then, clearly $TK \subseteq \mathcal{E}(TS)$. Thus, $\mathbf{order}(TK) \leq \mathbf{order}(K)$. The same argument applied to T^{-1} and TK shows that $\mathbf{order}(K) \leq \mathbf{order}(TK)$. \square

- The minimal volumetric ellipsoid is not unique in general; see example in figure 1. Further, it is in general different from the John ellipsoid.
- For non-degenerate bodies \mathcal{K} , their order is naturally lower bounded by d , and there are examples in which it is strictly more than d (e.g., figure 1).

To prove Theorem 1.1 we shall use John's decomposition theorem, Theorem 2.1.

Proof. of Theorem 1.1

Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a compact set. Without loss of generality assume that \mathcal{K} is symmetric and contains the origin; we can do so as in the following we only look at outer products of the form vv^\top for vectors $v \in \mathcal{K}$. Further, as $\mathbf{order}(\mathcal{K})$ is invariant under linear transformations, we can also assume that \mathcal{K} has been linearly transformed to be in John's position.

Now, by Theorem 2.1, there are $m = O(d^2)$ points $S = \{u_1, \dots, u_m\}$ on the enclosing unit ball that intersect \mathcal{K} (on the boundary of \mathcal{K}), and non-negative weights c_1, \dots, c_m such that $\sum_{i \in S} c_i u_i u_i^\top = I_d$. This implies by taking trace on both sides that $\sum_i c_i = d$.




Figure 1: In \mathbb{R}^2 the order of the volumetric ellipsoid of the equilateral triangle centered at the origin is at least 3. If the vertices are $[0, 1], [-\frac{\sqrt{3}}{2}, -\frac{1}{2}], [\frac{\sqrt{3}}{2}, -\frac{1}{2}]$, then the eigenpoles of the ellipsoid of the bottom two vertices are $[0, \frac{2}{3}], [2, 0]$. The second figure shows one possibility for a volumetric ellipsoid by adding $\frac{3}{4}$ of the first vertex to the previous ellipsoid. This shows the ellipsoid to be non-unique, as it can be rotated three ways.

Our idea is to start with the vectors u_1, \dots, u_m as a starting point for a volumetric spanner. However, this set has $O(d^2)$ points which is larger than what we want. We now prune these contact points via the sparsification methods of [7]. We use the following lemma that is a corollary of their Theorem 3.1.

Lemma 3.1 ([7]). *Let u_1, \dots, u_m be unit vectors and let $p \in \Delta(m)$ be a distribution over $[m]$ such that $d \sum_{i=1}^m p_i u_i u_i^\top = I_d$. Then, there exists a (possibly multi) set $S \subseteq \{u_1, \dots, u_m\}$ such that $\sum_{v \in S} v v^\top \succeq I_d$ and $|S| \leq 12d$. Moreover, such a set can be computed in time $O(md^3)$.*

Proof of Lemma 3.1. We first state Theorem 3.1 of [7].

Theorem 3.1. *Let v_1, \dots, v_m be vectors in \mathcal{R}^d and let $c > 1$. Assume that $\sum_i v_i v_i^\top = I_d$. There exist scalars $s_i \geq 0$ for $i \in [m]$ with $|\{i | s_i > 0\}| \leq cd$ such that*

$$I_d \preceq \sum_i s_i v_i v_i^\top \preceq \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}} I_d$$

Furthermore, these scalars can be found in time $O(cd^3m)$.

Let $v_i = \sqrt{p_i} u_i$. We fix c as some constant whose value will be determined later. Clearly for these vectors v_i it holds that $d \sum_i v_i v_i^\top = I_d$. It follows from the above theorem that there exist some scalars $s_i \geq 0$ for which at most cd are non-zeros and

$$I_d \preceq d \sum_i s_i v_i v_i^\top \preceq \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}} I_d \quad (1)$$

Our set S will be composed by taking each u_i $\lceil ds_i p_i \rceil$ many times. Plugging in equation (1) shows that indeed

$$\sum_{w \in S} w w^\top \succeq I_d$$

and it remains to bound the size of S i.e., $\sum \lceil ds_i p_i \rceil$. By taking the trace of the expression and dividing by d we get that

$$\sum_i s_i \text{Trace}(v_i v_i^\top) \preceq \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}}$$

Plugging in the expressions for v_i along with u_i being unit vectors (hence $\text{Trace}(u_i u_i^\top) = 1$) lead to

$$\sum_i s_i p_i \preceq \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}}$$

It follows that

$$\sum \lceil ds_i p_i \rceil \leq d \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}} + |\{i | s_i \geq 0\}| \leq d \left(c + \frac{c+1+2\sqrt{c}}{c+1-2\sqrt{c}} \right)$$

By optimizing c we get $\sum \lceil ds_i p_i \rceil \leq 12d$, and the lemma is proved. \square

Lemma 3.1 proves the existence of a spanner of size at most $12d$ for any compact set in \mathcal{R}^d . Furthermore, the running time to find the set is $O(md^3) = O(d^5)$, using the bound of $m = O(d^2)$. This along with the running time for obtaining the contact points leads to a total running time of $O(n^{3.5} + dn^3 + d^5)$. \square

4 Approximate Volumetric Spanners

In this section we provide a construction for (p, ε) -exp-volumetric spanner (as in Definition 1.3), proving Theorem 1.2. We start by providing a more technical definition of a spanner. Note that unlike previous definitions, the following is not impervious to linear operators and will only be used to aid our construction.

Definition 4.1. A β -relative-spanner is a discrete subset $S \subseteq \mathcal{K}$ such that for all $x \in \mathcal{K}$, $\|x\|_{\mathcal{E}(S)}^2 \leq \beta \|x\|^2$.

A first step is a spectral characterization of relative spanners:

Lemma 4.1. Let $S = \{v_1, \dots, v_T\} \subseteq \mathcal{K}$ span \mathcal{K} and be such that

$$W = \sum_{i=1}^T v_i v_i^\top \succeq \frac{1}{\beta} I_d$$

Then S is a β -relative-spanner.

Proof. Let $V \in \mathbb{R}^{d \times T}$ be a matrix whose columns are the vectors of S . As $VV^\top = W \succeq \frac{1}{\beta} I_d$ we have that

$$\beta I_d \succeq (VV^\top)^{-1}$$

It follows that

$$\|x\|_{\mathcal{E}(S)} = x^\top (VV^\top)^{-1} x \leq \beta \|x\|^2$$

as required. \square

Algorithm 1

- 1: Input: An oracle to $x \sim p$, where p is a distribution over \mathcal{K} , $T \in \mathbb{N}$.
 - 2: **for** $t = 1$ to T **do**
 - 3: Choose $u_t \sim p$
 - 4: **end for**
 - 5: **return** $S = \{u_1, \dots, u_T\}$.
-

Proof of Theorem 1.2. We analyze Algorithm 1, previously defined within Theorem 1.2 assuming the vectors are sampled exactly according to the log-concave distribution. The result involving an approximate sample, which is necessary for implementing the algorithm in the general case, is an immediate application of Lemma 2.1 and Corollary 2.1.

Our analysis of the algorithm is for $T = C(d + \log^2(1/\varepsilon))$ samples, where C is some sufficiently large constant. Assume first that $\mathbf{E}_{x \sim p}[xx^\top] = I_d$. Let $W = \sum_{i=1}^T u_i u_i^\top$. Then, for $C > 0$ large enough, by Corollary 2.1, $\|\frac{1}{T}W - I_d\| \leq 1/2$ w.p. at least $1 - \exp(-\sqrt{d})$. Therefore, S spans \mathbb{R}^d and

$$\frac{1}{T}W \succeq I_d - \frac{1}{2}I_d = \frac{1}{2}I_d$$

Thus according to Lemma 4.1, S is a $(2/T)$ -relative spanner. Consider the case where $\Sigma = \mathbf{E}_{x \sim p}[xx^\top]$ is not necessarily the identity. By the above analysis we get that $\Sigma^{-1/2}S = \{\Sigma^{-1/2}u_1, \dots, \Sigma^{-1/2}u_T\}$ form a $(2/T)$ -relative spanner for $\Sigma^{-1/2}\mathcal{K}$. This is since the r.v defined as $\Sigma^{-1/2}x$ where $x \sim p$ is log-concave. The latter along with corollary 2.2 implies that for any $\theta \geq 1$,

$$\Pr_{x \sim p} \left[\|\Sigma^{-1/2}x\| \geq 3\theta\sqrt{d} \right] \leq c_1 \exp\left(-c_2\theta\sqrt{d}\right) \quad (2)$$

for some universal constants $c_1, c_2 > 0$. It follows that for our set S and any $\theta \geq 1$,

$$\begin{aligned}
\Pr_{x \sim p} [\|x\|_{\mathcal{E}(S)} > \theta] &= \Pr_{x \sim p} [\|\Sigma^{-1/2}x\|_{\mathcal{E}(\Sigma^{-1/2}S)} > \theta] && \|x\|_{\mathcal{E}(S)} = \|\Sigma^{-1/2}x\|_{\mathcal{E}(\Sigma^{-1/2}S)} \\
&\leq \Pr_{x \sim p} [\|\Sigma^{-1/2}x\| > \theta\sqrt{T/2}] && \Sigma^{-1/2}S \text{ is a } 2/T\text{-relative-spanner} \\
&= \Pr_{x \sim p} \left[\|\Sigma^{-1/2}x\| > 3\theta\sqrt{d}\sqrt{\frac{C}{18}} \cdot \sqrt{1 + \frac{\log^2(1/\varepsilon)}{d}} \right] && T = C(d + \log^2(1/\varepsilon)) \\
&\leq c_1 \exp \left(-c_2\theta\sqrt{d}\sqrt{\frac{C}{18}} \cdot \sqrt{1 + \frac{\log^2(1/\varepsilon)}{d}} \right) && \text{Equation (2), } C \geq 18 \\
&\leq \exp \left(-\theta\sqrt{d + \log^2(1/\varepsilon)} \right) && C \text{ sufficiently large} \\
&\leq \varepsilon^{-\theta}
\end{aligned}$$

□

In our application of volumetric spanners to BLO, we also need the following relaxation of volumetric spanners where we allow ourselves the flexibility to scale the ellipsoid:

Definition 4.2. *A ρ -ratio-volumetric spanner S of \mathcal{K} is a subset $S \subseteq \mathcal{K}$ such that for all $x \in \mathcal{K}$, $\|x\|_{\mathcal{E}(S)} \leq \rho$.*

One example for such an approximate spanner with $\rho = \sqrt{d}$ is a barycentric spanner (Definition 2.3). In fact, it is easy to see that a C -approximate barycentric spanner is a $C\sqrt{d}$ -ratio-volumetric spanner. The following is immediate from Theorem 2.7.

Corollary 4.1. *Let \mathcal{K} be a compact set in \mathcal{R}^d that is not contained in any proper linear subspace. Given an oracle for optimizing linear functions over \mathcal{K} , for any $C > 1$, it is possible to compute a $C\sqrt{d}$ -ratio-volumetric spanner S of \mathcal{K} of cardinality $|S| = d$, using $O(d^2 \log_C(d))$ calls to the optimization oracle.*

5 Fast Volumetric Spanners for Discrete Sets

In this section we describe a different algorithm that constructs volumetric spanners for discrete sets. The order of the spanners we construct here is suboptimal (in particular, there is a dependence on the size of the set \mathcal{K} which we didn't have before). However, the algorithm is particularly simple and efficient to implement (takes time linear in the size of the set).

Algorithm 2

- 1: Input $\mathcal{K} = \{x_1, \dots, x_n\} \subseteq \mathbb{R}^d$.
 - 2: **if** $n < Cd \log d$ **then**
 - 3: **return** $S \leftarrow \mathcal{K}$
 - 4: **end if**
 - 5: Compute $\Sigma = \sum_i x_i x_i^\top$ and let $u_i = \Sigma^{-1/2}x_i$.
 - 6: For $i \in [n]$, let $p_i = 1/2n + \|u_i\|^2/2d$. Let S be a random set obtained by drawing $M = Cd \log d$ samples with replacement from $[n]$ according to the distribution p_1, \dots, p_n .
 - 7: Verify that for at least $n/2$ vectors from $\{x_1, \dots, x_n\}$, it holds that $\|x_i\|_{\mathcal{E}(S)} \leq 1$. If that is not the case discard S and repeat the above step.
 - 8: Apply the algorithm recursively on the data points for which $\|x_i\|_{\mathcal{E}(S)} > 1$.
-

Theorem 5.1. *Given a set of vectors $\mathcal{K} = \{x_1, \dots, x_n\} \in \mathbb{R}^d$, Algorithm 2 outputs a volumetric spanner of size $O((d \log d)(\log n))$ and has an expected running time of $O(nd^2)$.*

Proof. Consider a single iteration of the algorithm with input $v_1, \dots, v_n \in \mathbb{R}^d$. We claim that the random set S obtained in step 6 satisfies the following condition with constant probability:

$$\Pr_{x \in \mathcal{K}} [\|x\|_{\mathcal{E}(S)} \leq 1] \geq 1/2 \quad (3)$$

Suppose the above statement is true. Then, the lemma follows easily as it implies that for the next iteration there are fewer than $n/2$ vectors. Hence, after $(\log n)$ recursive calls we will have a volumetric spanner. The total size of the set will be $O((d \log d)(\log n))$. To see the time complexity, consider a single run of the algorithm. The most computationally intensive steps are computing Σ and $\Sigma^{-1/2}$ which take time $O(nd^2)$ and $O(d^3)$ respectively. We also need to compute $(\sum_{v \in S} vv^\top)^{-1}$ (to compute the $\mathcal{E}(S)$ norm) which takes time $O(d^3 \log d)$, and compute the $\mathcal{E}(S)$ norm of all the vectors which requires $O(nd^2)$. As $n = \Omega(d \log d)$, it follows that a single iteration runs of a total expected time of $O(nd^2)$. Since the size of n is split in half between iterations, the claim follows.

We now prove that Equation 3 holds with constant probability

$$x_j^\top \left(\sum_{v \in S} vv^\top \right)^{-1} x_j = u_j^\top \left(\sum_{v \in S'} vv^\top \right)^{-1} u_j. \quad (4)$$

where $S' = \{\Sigma^{-1/2}v | v \in S\}$ is the (linearly) shifted version of S . Therefore, it suffices to show that with sufficiently high probability, the right hand side of the above equation is bounded by 1 for at least $n/2$ indices $j \in [n]$.

Note that $p_i = 1/2n + \|u_i\|^2/2d$ form a probability distribution: $\sum_i p_i = 1/2 + (\sum_i \|u_i\|^2)/2d = 1$. Let $X \in \mathbb{R}^d$ be a random variable with $X = u_i/\sqrt{p_i}$ with probability p_i for $i \in [n]$. Then, $\mathbf{E}[XX^\top] = I_d$. Further, for any $i \in [n]$

$$\|u_i\|^2/p_i \leq 2d.$$

Therefore, by Theorem 2.6, if we take $M = Cd(\log d)$ samples X_1, \dots, X_M for C sufficiently large, then with probability of at least $1/2$, it holds that

$$\sum_{i=1}^M X_i X_i^\top \succeq (M/2)I_d.$$

Let $T \subseteq [n]$ be the multiset corresponding to the indices of the sampled vectors X_1, \dots, X_M . The above inequality implies that

$$\sum_{i \in T} \frac{1}{p_i} u_i u_i^\top \succeq (M/2)I_d.$$

Now,

$$\sum_{v \in S'} vv^\top \succeq (\min_i p_i) \sum_{v \in S'} \frac{1}{p_i} vv^\top \succeq (\min_i p_i)(M/2)I_d \succeq (M/4n)I_d.$$

Therefore,

$$\begin{aligned} \sum_{i=1}^n u_i^\top \left(\sum_{v \in S'} vv^\top \right)^{-1} u_i &= \sum_{i=1}^n \text{Tr} \left(\left(\sum_{v \in S'} vv^\top \right)^{-1} (u_i u_i^\top) \right) \\ &= \text{Tr} \left(\left(\sum_{v \in S'} vv^\top \right)^{-1} \left(\sum_{i=1}^n u_i u_i^\top \right) \right) \\ &= \text{Tr} \left(\left(\sum_{v \in S'} vv^\top \right)^{-1} \right) \\ &\leq \frac{4nd}{M} \leq \frac{4n}{C \log d} \leq \frac{n}{2 \log d}, \end{aligned}$$

for C sufficiently large. Therefore, by Markov's inequality and Equation 4, it follows that Equation 3 holds with high probability. The theorem now follows. \square

6 Bandit Linear Optimization

Recall the problem of Bandit Linear Optimization (BLO): iteratively at each time sequence t , the environment chooses a loss vector L_t that is not revealed to the player. The player chooses a vector $x_t \in \mathcal{K}$ where $\mathcal{K} \subseteq \mathbb{R}^d$ is convex, and once she commits to her choice, the loss $\ell_t = x_t^\top L_t$ is revealed. The objective is to minimize the loss and specifically, the regret, defined as the strategy's loss minus the loss of the best fixed strategy of choosing some $x^* \in \mathcal{K}$ for all t . We henceforth assume that the loss vectors L_t 's are chosen from the polar of \mathcal{K} , meaning from $\{L : |L^\top x| \leq 1 \ \forall x \in \mathcal{K}\}$. In particular this means that the losses are bounded in absolute value, although a different choice of assumption (i.e. ℓ_∞ bound on the losses) can yield different regret bounds, see discussion in [3].

The problem of BLO is a natural generalization of the classical Multi-Armed Bandit problem and extremely useful for efficiently modeling decision making under partial feedback for structured problems. As such the research literature is rich with algorithms and insights into this fundamental problem. For a brief historical survey please refer to earlier sections of this manuscript. In this section we focus on the first efficient and optimal-regret algorithm, and thus immediately jump to Algorithm 3. We make the following assumptions over the decision set \mathcal{K} :

1. The set \mathcal{K} is equipped with a membership oracle. This implies via [19] (Lemma 2.1) that there exists an efficient algorithm for sampling from a given log-concave distribution over \mathcal{K} . Via the discussion in previous sections, this also implies that we can construct approximate (both types of approximations, see Definitions 4.2 and 1.3) volumetric spanners efficiently over \mathcal{K} .
2. The losses are bounded in absolute values by 1. That is, the loss functions are always chosen (by an oblivious adversary) from a convex set \mathcal{Z} such that \mathcal{K} is contained in its polar, i.e. $\forall L \in \mathcal{Z}, x \in \mathcal{K}, |L^\top x| \leq 1$. This implies that the set \mathcal{K} admits for any $\varepsilon > 0$ an ε -net, w.r.t the norm defined by \mathcal{Z} , whose size we denote by $|\mathcal{K}|_\varepsilon \leq (\varepsilon/2)^{-d}$.

For Algorithm 3 we prove the following optimal regret bound:

Algorithm 3 GeometricHedge with Volumetric Spanners Exploration

- 1: \mathcal{K} , parameters γ, η , horizon T .
 - 2: $p_1(x)$ uniform distribution over \mathcal{K} .
 - 3: **for** $t = 1$ to T **do**
 - 4: Let S'_t be a $(p_t, \exp(-(4\sqrt{d} + \log(2T))))$ -exp-volumetric spanner of \mathcal{K} .
 - 5: Let S''_t be a $2\sqrt{d}$ -ratio-volumetric spanner of \mathcal{K}
 - 6: Set S_t as the union of S'_t, S''_t .
 - 7: $\hat{p}_t(x) = (1 - \gamma)p_t(x) + \frac{\gamma}{|S_t|} \mathbf{1}_{x \in S_t}$
 - 8: sample x_t according to \hat{p}_t
 - 9: observe loss $\ell_t \triangleq L_t^\top x_t$
 - 10: Let $C_t \triangleq \mathbf{E}_{x \sim \hat{p}_t} [xx^\top]$
 - 11: $\hat{L}_t \triangleq \ell_t C_t^{-1} x_t$
 - 12: $p_{t+1}(x) \propto p_t(x) e^{-\eta \hat{L}_t^\top x}$
 - 13: **end for**
-

Theorem 6.1. *Under the assumptions stated above, and let $s = \max_t |S_t|$, $\eta = \sqrt{\frac{\log |\mathcal{K}|_{1/T}}{dT}}$ and let $\gamma = s \sqrt{\frac{\log(|\mathcal{K}|_{1/T})}{dT}}$. Algorithm 3 given parameters γ, η suffers a regret bounded by*

$$O \left((s + d) \sqrt{\frac{T \log |\mathcal{K}|_{1/T}}{d}} \right)$$

We note that while the size $\log(|\mathcal{K}|_{1/T})$ can be bounded by $d \log(T)$, in certain scenarios such as s-t paths in graphs it is possible to obtain sharper upper bounds that immediately imply better regret via Theorem 6.1.

Corollary 6.1. *There exist an efficient algorithm for BLO for any convex set \mathcal{K} with regret of*

$$O\left(\sqrt{dT \log |\mathcal{K}|_{1/T}}\right) = O\left(d\sqrt{T \log(T)}\right)$$

Proof. The spanner in step 4 of the algorithm does not have to be explicitly constructed. According to Theorem 1.2, to obtain such a spanner it suffices to sample sufficiently many points from the distribution p_t , hence this portion of the exploration strategy is identical to the exploitation strategy.

According to Corollary 4.1, a $2\sqrt{d}$ -ratio-volumetric spanner of size d can be efficiently constructed, given a linear optimization oracle which in turn can be efficiently implemented by the membership oracle for \mathcal{K} . Hence, it follows that for the purpose of the analysis, $s = d$ and the bound follows. \square

To prove the theorem we follow the general methodology used in analyzing the performance of the geometric hedge algorithm. The major deviation from standard technique is the following sub-exponential tail bound, which we use to replace the the standard absolute bound for $|\hat{L}_t x|$. After giving its proof and a few auxiliary lemmas, we give the proof of the main theorem.

Lemma 6.1. *Let $x \sim p_t$, $x_t \sim \hat{p}_t$ and let \hat{L}_t be defined according to x_t . It holds, for any $\theta > 1$ that*

$$\Pr\left[|\hat{L}_t^\top x| > \frac{\theta s}{\gamma}\right] \leq \exp(-2\theta)/T$$

Proof.

$$\begin{aligned} \Pr\left[|\hat{L}_t^\top x| > \frac{\theta s}{\gamma}\right] &\leq \Pr\left[\|x\|_{\mathcal{E}(S_t)} \cdot \|x_t\|_{\mathcal{E}(S_t)} \geq \theta\right] && \text{Lemma 6.2} \\ &\leq \Pr\left[\|x\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta} \quad \vee \quad \|x_t\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta}\right] \\ &\leq \Pr\left[\|x\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta}\right] + \Pr\left[\|x_t\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta}\right] \\ &\leq 2 \Pr\left[\|x\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta}\right] \end{aligned}$$

To justify the last inequality notice that $x \sim p_t$ and $x_t \sim \hat{p}_t$ where \hat{p}_t is a convex sum of p_t and a distribution q_t for which $\Pr_{y \sim q_t}[\|y\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta} > 1] = 0$. Before we continue recall that we can assume that $\sqrt{\theta} \leq 2\sqrt{d}$, since S_t'' is a $2\sqrt{d}$ -ratio-volumetric spanner .

$$\begin{aligned} \Pr\left[|\hat{L}_t^\top x| > \frac{\theta s}{\gamma}\right] &\leq 2 \Pr\left[\|x\|_{\mathcal{E}(S_t)} \geq \sqrt{\theta}\right] \\ &\leq 2 \exp(-\sqrt{\theta}(4\sqrt{d} + \log 2T)) \quad \text{property of exp-volumetric spanner} \\ &\leq \frac{1}{T} \exp(-2\sqrt{\theta} \cdot 4d) \\ &\leq \frac{1}{T} \exp(-2\theta) \quad \text{since } \theta \leq 4d \end{aligned}$$

\square

Lemma 6.2. *For all $x \in \mathcal{K}$ it holds that $|\hat{L}_t^\top x| \leq \frac{|S_t| \|x\|_{\mathcal{E}(S_t)} \|x_t\|_{\mathcal{E}(S_t)}}{\gamma}$.*

Proof. Let $x \in \mathcal{K}$. Denote by V_t the matrix whose columns are the elements of S_t and recall that $\|y\|_{\mathcal{E}(S_t)}^2 = y^\top (V_t V_t^\top)^{-1} y$. Since $C_t \triangleq \mathbf{E}_{x \sim \hat{p}_t}[xx^\top]$, it holds that

$$C_t \succeq \frac{\gamma}{|S_t|} \sum_{v \in S_t} vv^\top = \frac{\gamma}{|S_t|} V_t V_t^\top$$

since both matrices are full rank, it holds that

$$C_t^{-1} \preceq \frac{|S_t|}{\gamma} (V_t V_t^\top)^{-1}$$

Notice that due to the Cauchy-Schwartz inequality,

$$|x^\top \hat{L}_t| = |\ell_t| \cdot |x^\top C_t^{-1} x_t| \leq |\ell_t| \cdot \|x^\top C_t^{-1/2}\| \cdot \|C_t^{-1/2} x_t\|$$

The matrix $C_t^{-1/2}$ is defined as C_t is positive definite. Now,

$$\|x^\top C_t^{-1/2}\|^2 = x^\top C_t^{-1} x \leq x^\top \frac{|S_t|}{\gamma} (V_t V_t^\top)^{-1} x = \frac{|S_t|}{\gamma} \|x\|_{\mathcal{E}(S_t)}^2$$

Since the analog can be said for $\|C_t^{-1/2} x_t\|$ (as $x_t \in \mathcal{K}$), it follows that

$$|x^\top \hat{L}_t| \leq |\ell_t| \frac{|S_t| \|x\|_{\mathcal{E}(S_t)} \|x_t\|_{\mathcal{E}(S_t)}}{\gamma} \leq \frac{|S_t| \|x\|_{\mathcal{E}(S_t)} \|x_t\|_{\mathcal{E}(S_t)}}{\gamma}$$

The last inequality is since we assume the rewards are in $[-1, 1]$. \square

Implementation for general convex bodies. In the case where the set \mathcal{K} is a general convex body, the analysis must include the fact that we can only approximately sample a log-concave distribution over \mathcal{K} . As the main focus of our work is to prove a polynomial solution we present only a simple analysis yielding a running time polynomial in the dimension d and horizon T . It is likely that a more thorough analysis can substantially reduce the running time.

Corollary 6.2. *In the general case where an approximate sampling is required, algorithm 3 can be implemented with a running time of $\tilde{O}(d^5 + d^3 T^6)$ per iteration.*

Proof. Fix an error parameter δ , and let us run Algorithm 3 with the approximate samplers p'_t guaranteed by Theorem 2.1. Then, in each use of the sampler we are replacing the true distribution we should be using p_t , with a distribution p'_t such that statistical distance between p_t, p'_t is at most δ . Let us now analyze the error incurred by this approximation by bounding the loss from the first round onwards. Suppose the algorithm ran with the approximate sampler in the first round but the exact sampler in each round afterwards. Then, as the statistical distance between the distributions is at most δ and the loss in each round is bounded by 1 and there are T rounds, the net difference in expected regret between using p_1 and p'_1 will be at most $\delta \cdot T$. Similarly, if we ran the algorithm with $p'_1, \dots, p'_{i-1}, p'_i, p_{i+1}, \dots, p_T$ as opposed to $p'_1, \dots, p'_{i-1}, p_i, p_{i+1}, \dots, p_T$ (we are changing the i 'th distribution from exact to approximate), the net difference in expected regret would be at most $\delta \cdot T$. Therefore, the total additional loss we may incur for using the approximate oracles is at most $T \cdot (\delta T) = \delta T^2$. Thus, if we take $\delta = \Delta/T^2$, where Δ is the regret bound from Theorem 6.1, we get a regret bound of 2Δ . The required value of δ is bounded by $T^{-1.5}$. Applying Theorem 2.1 leads to a running time of $\tilde{O}(d^5 + d^3 T^6)$ per iteration. \square

6.1 Proof of Theorem 6.1

We continue the analysis of the Geometric Hedge algorithm similarly to [13, 9], under certain assumptions over the exploration strategy. For convenience we will assume that the set of possible arms \mathcal{K} is finite. This assumption holds w.l.o.g since if \mathcal{K} is infinite, a $\sqrt{1/T}$ -net of it can be considered as described earlier (this will have no effect on the computational complexity of our algorithm, but a mere technical convenience in the proof below).

Before proving the theorem we will require three technical lemmas. In the first we show that \hat{L}_t is an unbiased estimator of L_t . In the second, we bound a proxy of its variance. In the third, we bound a proxy of the expected value of its exponent.

Lemma 6.3. *In each t , \hat{L}_t is an unbiased estimator of L_t*

Proof.

$$\hat{L}_t = \ell_t C_t^{-1} x_t = (L_t^\top x_t) C_t^{-1} x_t = C_t^{-1} (x_t x_t^\top) L_t$$

⁴Here, p_j 's are interpreted as the distribution given by the algorithm based on the distribution from previous round and p'_j is the approximate oracle for this distribution p_j .

Hence,

$$\mathbf{E}_{x_t \sim p_t} [\hat{L}_t] = C_t^{-1} \mathbf{E}_{x_t \sim p_t} [x_t x_t^\top] L_t = C_t^{-1} C_t L_t = L_t$$

□

Lemma 6.4. *Let $t \in [T]$, $x \sim p_t$ and $x_t \sim \hat{p}_t$. It holds that $\mathbf{E}[(\hat{L}_t^\top x)^2] \leq d/(1-\gamma) \leq 2d$*

Proof. For convenience, denote by q_t the uniform distribution over S_t - the exploration strategy at round t . First notice that for any $x \in \mathcal{K}$,

$$\begin{aligned} \mathbf{E}_{x_t \sim \hat{p}_t} [(\hat{L}_t^\top x)^2] &= x^\top \mathbf{E}_{x_t \sim \hat{p}_t} [\hat{L}_t \hat{L}_t^\top] x = x^\top \mathbf{E}_{x_t \sim \hat{p}_t} [\ell_t^2 C_t^{-1} x_t x_t^\top C_t^{-1}] x \\ &= \ell_t^2 x^\top C_t^{-1} \mathbf{E}_{x_t \sim \hat{p}_t} [x_t x_t^\top] C_t^{-1} x = \ell_t^2 x^\top C_t^{-1} x \\ &\leq x^\top C_t^{-1} x \end{aligned} \quad (5)$$

Next,

$$\mathbf{E}_{x \sim \hat{p}_t} [x^\top C_t^{-1} x] = \mathbf{E}_{x \sim \hat{p}_t} [C_t^{-1} \bullet x x^\top] = C_t^{-1} \bullet \mathbf{E}_{x \sim \hat{p}_t} [x x^\top] = C_t^{-1} \bullet C_t = \mathbf{Tr}(I_d) = d$$

Where we used linearity of expectation and denote $A \bullet B = \mathbf{Tr}(AB)$. Since C_t^{-1} is positive semi-definite,

$$(1-\gamma) \mathbf{E}_{x \sim \hat{p}_t} [x^\top C_t^{-1} x] \leq (1-\gamma) \mathbf{E}_{x \sim p_t} [x^\top C_t^{-1} x] + \gamma \mathbf{E}_{x \sim q_t} [x^\top C_t^{-1} x] = \mathbf{E}_{x \sim \hat{p}_t} [x^\top C_t^{-1} x] = d \quad (6)$$

The lemma follows from combining Equations 5 and 6. □

Lemma 6.5. *Denote by $\mathbf{1}_\phi$ the random variable taking a value of 1 if event ϕ occurred and 0 otherwise. Let $t \in [T]$, $x_t \sim \hat{p}_t$ and $x \sim p_t$. For \hat{L}_t defined by x_t it holds that*

$$\mathbf{E} \left[\exp(-\eta \hat{L}_t^\top x) \mathbf{1}_{-\eta \hat{L}_t^\top x > 1} \right] \leq \frac{2}{T}$$

Proof. Let f, F be the pdf and cdf of the random variable $Y = -\eta \hat{L}_t^\top x$ correspondingly. From Lemma 6.1 and the fact that $1/\eta = s/\gamma$ ($s = \max_t |S_t|$) we have that for any $\theta \geq 1$,

$$1 - F(\theta) \leq \frac{1}{T} e^{-2\theta}$$

and we'd like to prove that under this condition,

$$\mathbf{E}[e^Y \mathbf{1}_{Y > 1}] = \int_{\theta=1}^{\infty} e^\theta f(\theta) d\theta \leq \frac{2}{T}$$

which follows from the definition of the cdf and pdf:

$$\begin{aligned} \mathbf{E}[e^Y \mathbf{1}_{Y > 1}] &= \int_{\theta=1}^{\infty} e^\theta f(\theta) d\theta \\ &= \sum_{k=1}^{\infty} \int_{\theta=k}^{k+1} e^\theta f(\theta) d\theta \\ &\leq \sum_{k=1}^{\infty} e^{k+1} \int_{\theta=k}^{k+1} f(\theta) d\theta \\ &\leq \sum_{k=1}^{\infty} e^{k+1} (F(k+1) - F(k)) \\ &\leq \sum_{k=1}^{\infty} e^{k+1} (1 - F(k)) \\ &\leq \sum_{k=1}^{\infty} e^{k+1} \cdot \frac{1}{T} e^{-2k} \quad \text{Lemma 6.1} \\ &= \frac{e}{T} \sum_{k=1}^{\infty} e^{-k} = \frac{e}{T} \cdot \frac{e^{-1}}{1-e^{-1}} \leq \frac{2}{T} \end{aligned}$$

□

Proof of Theorem 6.1. For convenience we define within this proof for $x \in \mathcal{K}$, $\hat{\ell}_{1:t-1}(x) \triangleq \sum_{i=1}^{t-1} \hat{L}_i^\top x$ and let $\hat{\ell}_t(x) \triangleq \hat{L}_t^\top x$. Let $W_t = \sum_{x \in \mathcal{K}} \exp(-\eta \hat{\ell}_{1:t-1}(x))$. For all $t \in [T]$:

$$\begin{aligned} \mathbf{E} \left[\frac{W_{t+1}}{W_t} \right] &= \mathbf{E} \left[\sum_{x \in \mathcal{K}} \frac{\exp(-\eta \hat{\ell}_{1:t-1}(x)) \exp(-\eta \hat{\ell}_t(x))}{W_t} \right] \\ &= \mathbf{E}_{x_t \sim \hat{p}_t} \left[\sum_{x \in \mathcal{K}} p_t(x) \exp(-\eta \hat{\ell}_t(x)) \right] \\ &= \mathbf{E}_{x_t \sim \hat{p}_t, x \sim p_t} [\exp(-\eta \hat{\ell}_t(x))] \leq \\ &\leq 1 - \eta \mathbf{E}[\hat{L}_t^\top x] + \eta^2 \mathbf{E}[(\hat{L}_t^\top x)^2] + \mathbf{E} \left[\exp(-\eta \hat{L}_t^\top x) \mathbf{1}_{-\eta \hat{L}_t^\top x > 1} \right] \end{aligned}$$

using the inequality $\exp(y) \leq 1 + y + y^2 + \exp(y) \cdot \mathbf{1}_{y > 1}$

$$\leq 1 - \eta \mathbf{E}[\hat{L}_t^\top x] + \eta^2 \mathbf{E}[(\hat{L}_t^\top x)^2] + \frac{2}{T} \quad \text{Lemma 6.5}$$

Since \hat{L}_t is an unbiased estimator of L_t (Lemma 6.3) and according to Lemma 6.4, $\mathbf{E}[(\hat{L}_t^\top x)^2] \leq 2d$, we get:

$$\mathbf{E} \left[\frac{W_{t+1}}{W_t} \right] \leq 1 - \eta L_t^\top \mathbf{E}_{x \sim p_t} [x] + 2\eta^2 d + \frac{2}{T} \quad (7)$$

We now use Jensen's inequality:

$$\begin{aligned} \mathbf{E}[\log(W_T)] - \mathbf{E}[\log(W_1)] &= \mathbf{E}[\log(W_T/W_1)] \\ &= \sum_{t=1}^{T-1} \mathbf{E}[\log(W_{t+1}/W_t)] \\ &\leq \sum_{t=1}^{T-1} \log(\mathbf{E}[W_{t+1}/W_t]) \quad \text{Jensen} \\ &\leq \sum_{t=1}^{T-1} \log(1 - \eta L_t^\top \mathbf{E}_{x \sim p_t} [x] + 2\eta^2 d + \frac{2}{T}) \quad (7) \\ &\leq \sum_{t=1}^{T-1} -\eta L_t^\top \mathbf{E}_{x \sim p_t} [x] + 2\eta^2 d + \frac{2}{T} \quad \text{Due to } \ln(1+y) \leq y \text{ for all } y > -1 \\ &\leq 2 + 2\eta^2 Td - \eta \sum_t \mathbf{E}_{x \sim p_t} [L_t^\top x] \end{aligned}$$

Now, since $\log(W_1) = \log(|\mathcal{K}|)$ and $W_T \geq \exp(-\eta \hat{\ell}_{1:T}(x^*))$ for any $x^* \in \mathcal{K}$, by shifting sides of the above it holds for any $x^* \in \mathcal{K}$ that

$$\sum_t \mathbf{E}_{x \sim p_t} [L_t^\top x] - \sum_t L_t^\top x^* \leq \sum_t \mathbf{E}_{x \sim p_t} [L_t^\top x] + \mathbf{E}[\log W_T] \leq \frac{\log(|\mathcal{K}|) + 2}{\eta} + 2\eta Td$$

Finally, by noticing that

$$\sum_t \mathbf{E}_{x \sim \hat{p}_t} [L_t^\top x] - \sum_t \mathbf{E}_{x \sim p_t} [L_t^\top x] \leq \gamma T$$

we obtain a bound of

$$\mathbf{E}[\text{Regret}] = \mathbf{E} \left[\sum_t L_t^\top x_t \right] - \sum_t L_t^\top x^* = \sum_t \mathbf{E}_{x \sim \hat{p}_t} [L_t^\top x] - \text{Loss}(x^*) \leq \frac{\log(|\mathcal{K}|) + 2}{\eta} + 2\eta Td + \gamma T$$

on the expected regret. By plugging in the values of η, γ we get the bound of

$$O \left((s+d) \sqrt{\frac{T \log(|\mathcal{K}|)}{d}} \right)$$

as required. \square

Discussion: Notice that to obtain a (p, ε) -exp-volumetric spanner for a log-concave distribution p over a body \mathcal{K} we simply choose sufficiently many i.i.d samples from p . Since in the above algorithm p_t is always log-concave, it follows that S'_t consists of i.i.d samples from p_t , meaning that if we would not have required S''_t , the exploration and exploration strategies would be the same! Since we still require the set S''_t , there exists a need for a separate exploration strategy. Interestingly, the $2\sqrt{d}$ -ratio-volumetric spanner is obtained by taking a barycentric spanner, which is the exploration strategy of [13].

References

- [1] J.D. Abernethy, E. Hazan, and A. Rakhlin. Interior-point methods for full-information and bandit online learning. *IEEE Transactions on Information Theory*, 58(7):4164–4175, 2012.
- [2] Radoslaw Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of American Mathematical Society*, 23:535–561, 2010.
- [3] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Minimax policies for combinatorial prediction games. In Sham M. Kakade and Ulrike von Luxburg, editors, *COLT*, volume 19 of *JMLR Proceedings*, pages 107–132. JMLR.org, 2011.
- [4] Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *J. Comput. Syst. Sci.*, 74(1):97–114, 2008.
- [5] Keith Ball. An elementary introduction to modern convex geometry. In *Flavors of Geometry*, pages 1–58. Univ. Press, 1997.
- [6] Peter L. Bartlett, Varsha Dani, Thomas P. Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *COLT*, pages 335–342, 2008.
- [7] Joshua Batson, Daniel A Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.
- [8] S. Bubeck and N. Cesa-Bianchi. *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*, volume 5 of *Foundations and Trends in Machine Learning*. NOW, 2012.
- [9] Sébastien Bubeck, Nicolò Cesa-Bianchi, and Sham M. Kakade. Towards minimax policies for online linear optimization with bandit feedback. *Journal of Machine Learning Research - Proceedings Track*, 23:41.1–41.14, 2012.
- [10] Nicolò Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *J. Comput. Syst. Sci.*, 78(5):1404–1422, 2012.
- [11] S. Damla Ahipasaoglu, Peng Sun, and Michael J. Todd. Linear convergence of a modified frankwolfe algorithm for computing minimum-volume enclosing ellipsoids. *Optimization Methods and Software*, 23(1):5–19, 2008.
- [12] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, pages 355–366, 2008.
- [13] Varsha Dani, Sham M Kakade, and Thomas P Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, pages 345–352, 2007.
- [14] Olivier Gudon and Emanuel Milman. Interpolating thin-shell and sharp large-deviation estimates for isotropic log-concave measures. *Geometric and Functional Analysis*, 21(5):1043–1068, 2011.
- [15] Martin Henk. Löwner-John ellipsoids. *Documenta Mathematica*, pages 95–106, 2012.
- [16] F. John. Extremum Problems with Inequalities as Subsidiary Conditions. In K. O. Friedrichs, O. E. Neugebauer, and J. J. Stoker, editors, *Studies and Essays: Courant Anniversary Volume*, pages 187–204. Wiley-Interscience, New York, 1948.
- [17] Sham M Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009.
- [18] Leonid G Khachiyan. Rounding of polytopes in the real number model of computation. *Mathematics of Operations Research*, 21(2):307–320, 1996.

- [19] László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [20] M. Rudelson. Random vectors in the isotropic position. *Journal of Functional Analysis*, 164(1):60–72, 1999.

A Concentration bounds for non centered isotropic log concave distributions

We begin by proving an auxiliary lemma used in the proof of Corollary 2.1.

Lemma A.1. *Let $\delta > 0$, $t \geq 1$, let d be a positive integer and let $n = \frac{Ct^4 d \log^2(t/\delta)}{\delta^2}$ for some sufficiently large universal constant C . Let y_1, \dots, y_n be i.i.d d -dimensional vectors from an isotropic log-concave distribution. Then*

$$\Pr \left[\left\| \frac{1}{n} \sum y_i \right\| > \delta \right] \leq \exp(-t\sqrt{d})$$

Proof. For convenience let $S_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n y_i$. Since the y 's are independent, S_n is also log-concave distributed. Notice that $\mathbf{E}[S_n] = 0$ and $\mathbf{E}[S_n S_n^T] = \frac{1}{n} \sum \mathbf{E}[y_i y_i^T] = I_d$ hence S_n is isotropic. Now,

$$\Pr \left[\left\| \frac{1}{n} \sum y_i \right\| > \delta \right] = \Pr [\|S_n\| > \sqrt{n}\delta] = \Pr \left[\|S_n\| > \sqrt{d} \cdot \sqrt{Ct^4 \log^2(t/\delta)} \right] \leq \Pr \left[\|S_n\| > \sqrt{d} + \sqrt{d} \cdot \frac{1}{2} t\sqrt{C} \right].$$

The last inequality holds for $t \geq 1$ and $C \geq 4$. It now follows from Theorem 2.5 that

$$\Pr \left[\left\| \frac{1}{n} \sum y_i \right\| > \delta \right] \leq c_1 \exp(-c_2 t\sqrt{Cd})$$

where c_1, c_2 are some universal constants. Since $t\sqrt{d} \geq 1$, setting $C \geq \left(\frac{1+\log(c_1)}{c_2}\right)^2$ proves the claim. \square

Proof of Corollary 2.1. Let $a = \mathbf{E}[x]$ and let $\tilde{a} = \frac{1}{n} \sum x_i$. Notice that

$$\mathbf{E}[(x-a)(x-a)^T] = \mathbf{E}[xx^T] - \mathbf{E}[x]a^T - a\mathbf{E}[x] + aa^T = I_d - aa^T$$

is a PSD matrix hence $\|a\| \leq 1$. Consider the following equality.

$$\frac{1}{n} \sum_{i=1}^n (x_i - a)(x_i - a)^T = \frac{1}{n} \sum_{i=1}^n x_i x_i^T - a\tilde{a}^T - \tilde{a}a^T + aa^T$$

According to Lemma A.1, w.p. at least $1 - \exp(-t\sqrt{d})$,

$$\|\tilde{a} - a\| \leq \delta$$

in which case, since $\|a\| \leq 1$ and according to the triangle inequality,

$$\left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T - I_d \right\| \leq \left\| \frac{1}{n} \sum_{i=1}^n (x_i - a)(x_i - a)^T - (I_d - aa^T) \right\| + 2\delta$$

According to Theorem 2.4, w.p. at least $1 - \exp(-t\sqrt{d})$

$$\left\| \frac{1}{n} \sum_{i=1}^n (x_i - a)(x_i - a)^T - (I_d - aa^T) \right\| \leq \delta$$

and the corollary follows. \square

Proof of Corollary 2.2. Let $\mathbf{E}[x] = a$. Consider the r.v $y = x - a$. It holds that $\mathbf{E}[y] = 0$ and $\mathbf{E}[yy^T] = I_d - aa^T$. Notice that we can derive that

$$\|a\| \leq 1 \tag{8}$$

As $\mathbf{E}[yy^T]$ is a PSD matrix. Also, it is easy to verify that y is log-concave distributed. We now consider the r.v⁵ $z = (I_d - aa^t)^{-1/2}y$. It is easy to verify that the distribution of z is also log-concave and isotropic. It follows, from Theorem 2.5 that for any $\theta \geq 2$

$$\Pr \left[\|z\| > \theta\sqrt{d} \right] \leq \Pr \left[\|z\| - \sqrt{d} > \frac{1}{2}\theta\sqrt{d} \right] \leq C' \exp(-c\theta\sqrt{d})$$

By using equation 8 we get that for $\theta > 3$

$$\Pr \left[\|x\| > \theta\sqrt{d} \right] \leq \Pr \left[\|y\| > \theta\sqrt{d} - 1 \right] \leq \Pr \left[\|z\| > (\theta - 1/\sqrt{d})\sqrt{d} \right] \leq C' \exp(c' - c'\theta\sqrt{d}).$$

The last inequality holds since $\theta - 1/\sqrt{d} \geq 2$. □

⁵if $I_d - aa^T$ is not of full rank then y is in fact supported in an affine subspace of rank $d - 1$ and we can continue the analysis there.