

A PTAS for Computing the Supremum of Gaussian Processes

Raghu Meka*

Abstract

We give a polynomial time approximation scheme (PTAS) for computing the supremum of a Gaussian process. That is, given a finite set of vectors $V \subseteq \mathbb{R}^d$, we compute a $(1 + \varepsilon)$ -factor approximation to $\mathbb{E}_{X \leftarrow \mathcal{N}^d} [\sup_{v \in V} |\langle v, X \rangle|]$ deterministically in time $\text{poly}(d) \cdot |V|^{O_\varepsilon(1)}$. Previously, only a constant factor deterministic polynomial time approximation algorithm was known due to the work of Ding, Lee and Peres [DLP11]. This answers an open question of Lee [Lee10] and Ding [Din11].

The study of supremum of Gaussian processes is of considerable importance in probability with applications in functional analysis, convex geometry, and in light of the recent breakthrough work of Ding, Lee and Peres [DLP11], to random walks on finite graphs. As such our result could be of use elsewhere. In particular, combining with the recent work of Ding [Din11], our result yields a PTAS for computing the cover time of bounded degree graphs. Previously, such algorithms were known only for trees.

Along the way, we also give an explicit oblivious linear estimator for semi-norms in Gaussian space with optimal query complexity.

1 Introduction

The study of supremum of Gaussian processes is a major area of study in probability and functional analysis as epitomized by the celebrated *majorizing measures* theorem of Fernique and Talagrand (see [LT91], [Tal05] and references therein). There is by now a rich body of work on obtaining tight estimates and characterizations of the supremum of Gaussian processes with several applications in analysis [Tal05], convex geometry [Pis99] and more. Recently, in a striking result, Ding, Lee and Peres [DLP11] used the theory to resolve the Winkler-Zuckerman *blanket time* conjectures [WZ96], indicating the usefulness of Gaussian processes even for the study of combinatorial problems over discrete domains.

Ding, Lee and Peres [DLP11] used the powerful *Dynkin isomorphism theory* and majorizing measures theory to establish a structural connection between the cover time (and blanket time) of a graph G and the supremum of a Gaussian process associated with the Gaussian Free Field on G . They then use this connection to resolve the Winkler-Zuckerman blanket time conjectures and to obtain the first deterministic polynomial time constant factor approximation algorithm for computing the cover time of graphs. This latter result resolves an old open question of Aldous and Fill (1994).

Besides showing the relevance of the study of Gaussian processes to discrete combinatorial questions, the work of Ding, Lee and Peres gives evidence that studying Gaussian processes could even be an important algorithmic tool; an aspect seldom investigated in the rich literature on

*Institute for Advanced Study, Princeton. Email: raghu@ias.edu

Gaussian processes in probability and functional analysis. Here we address the corresponding computational question directly, which given the importance of Gaussian processes in probability, could be of use elsewhere. In this context, the following question was asked by Lee [Lee10] and Ding [Din11].

Question 1.1. For every $\varepsilon > 0$, is there a deterministic polynomial time algorithm that given a set of vectors $v_1, \dots, v_m \in \mathbb{R}^d$, computes a $(1 + \varepsilon)$ -factor approximation to $\mathbb{E}_{X \leftarrow \mathcal{N}^d}[\sup_i |\langle v_i, X \rangle|]$ ¹.

We remark that Lee [Lee10] and [Din11] actually ask for an approximation to $\mathbb{E}_{X \leftarrow \mathcal{N}^d}[\sup_i \langle v_i, X \rangle]$ (and not the supremum of the absolute value). However, this formulation results in a somewhat artificial asymmetry and for most interesting cases these two are essentially equivalent: if $\mathbb{E}_{X \leftarrow \mathcal{N}^d}[\sup_i \langle v_i, X \rangle] = \omega(\max_i \|v_i\|_2)$, then $\mathbb{E}_{X \leftarrow \mathcal{N}^d}[\sup_i |\langle v_i, X \rangle|] = (1 + o(1)) \mathbb{E}_{X \leftarrow \mathcal{N}^d}[\sup_i \langle v_i, X \rangle]$ ². We shall overlook this distinction from now on.

There is a simple randomized algorithm for the problem: sample a few Gaussian vectors and output the median supremum value for the sampled vectors. This however, requires $O(d \log d / \varepsilon^2)$ random bits. Using Talagrand’s majorizing measures theorem, Ding, Lee and Peres give a deterministic polynomial time $O(1)$ -factor approximation algorithm for the problem. This approach is inherently limited to not yield a PTAS as the majorizing measures characterization is bound to lose a universal constant factor. Here we give a PTAS for the problem thus resolving the above question.

Theorem 1.2. *For every $\varepsilon > 0$, there is a deterministic algorithm that given a set of vectors $v_1, \dots, v_m \in \mathbb{R}^d$, computes a $(1 + \varepsilon)$ -factor approximation to $\mathbb{E}_{x \leftarrow \mathcal{N}^d}[\sup_i |\langle v_i, x \rangle|]$ in time $\text{poly}(d) \cdot m^{\tilde{O}(1/\varepsilon^2)}$.*

Our approach is comparatively simpler than the work of Ding, Lee and Peres, using some classical *comparison inequalities* in convex geometry.

We explain our result on estimating semi-norms with respect to Gaussian measures mentioned in the abstract in Section 2.2.

We next discuss some applications of our result to computing cover times of graphs as implied by the works of Ding, Lee and Peres [DLP11] and Ding [Din11].

1.1 Application to Computing Cover Times of Graphs

The study of random walks on graphs is an important area of research in probability, algorithm design, statistical physics and more. As this is not the main topic of our work, we avoid giving formal definitions and refer the readers to [AF], [Lov93] for background information.

Given a graph G on n -vertices, the cover time, $\tau_{\text{cov}}(G)$, of G is defined as the expected time a random walk on G takes to visit all the vertices in G when starting from the worst possible vertex in G . Cover time is a fundamental parameter of graphs and is extensively studied. Algorithmically, there is a simple randomized algorithm for approximating the cover time - simulate a few trials of the random walk on G for $\text{poly}(n)$ steps and output the median cover time. However, without randomness the problem becomes significantly harder. This was one of the motivations of the work

¹Throughout, \mathcal{N} denotes the univariate Gaussian distribution with mean 0 and variance 1 and for a distribution \mathcal{D} , $X \leftarrow \mathcal{D}$ denotes a random variable with distribution \mathcal{D} . By a α -factor approximation to a quantity X we mean a number p such that $p \leq X \leq \alpha p$.

²This follows from standard concentration bounds for supremum of Gaussian processes; we do not elaborate on it here as we ignore this issue.

of Ding, Lee and Peres [DLP11] who gave the first deterministic constant factor approximation algorithm for the problem, improving on an earlier work of Kahn, Kim, Lovász and Vu [KKLV00] who obtained a deterministic $O((\log \log n)^2)$ -factor approximation algorithm. For the simpler case of trees, Feige and Zeitouni [FZ09] gave a FPTAS.

Ding, Lee and Peres also conjectured that the cover time of a graph G (satisfying a certain reasonable technical condition) is asymptotically equivalent to the supremum of an explicitly defined Gaussian process—the Gaussian Free Field on G . However, this conjecture though quite interesting on its own, is not enough to give a PTAS for cover time; one still needs a PTAS for computing the supremum of the relevant Gaussian process. Our main result provides this missing piece, thus removing one of the obstacles in their posited strategy to obtain a PTAS for computing the cover time of graphs. Recently, Ding [Din11] showed the main conjecture of Ding, Lee and Peres to be true for bounded-degree graphs and trees. Thus, combining his result (see Theorem 1.1 in [Din11]) with Theorem 1.2 we get a PTAS for computing cover time on bounded degree graphs with $\tau_{hit}(G) = o(\tau_{cov}(G))$ ³. As mentioned earlier, previously, such algorithms were only known for trees [FZ09].

2 Outline of Algorithm

The high level idea of our PTAS is as follows. Fix the set of vectors $V = \{v_1, \dots, v_m\} \subseteq \mathbb{R}^d$ and $\varepsilon > 0$. Without loss of generality suppose that $\max_{v \in V} \|v\|_2 = 1$. We first reduce the dimension of V by projecting V onto a space of dimension of $O((\log m)/\varepsilon^2)$ á la the classical Johnson-Lindenstrauss lemma (JLL). We then give an algorithm that runs in time polynomial in the number of vectors but exponential in the underlying dimension. Our analysis relies on two elegant comparison inequalities in convex geometry—Slepian’s lemma [Sle62] for the first step and Kanter’s lemma [Kan77] for the second step. We discuss these modular steps below.

2.1 Dimension Reduction

We project the set of vectors $V \subseteq \mathbb{R}^d$ to \mathbb{R}^k for $k = O((\log m)/\varepsilon^2)$ to preserve all pairwise (Euclidean) distances within a $(1 + \varepsilon)$ -factor as in the Johnson-Lindenstrauss lemma (JLL). We then show that the expected supremum of the *projected* Gaussian process is within a $(1 + \varepsilon)$ factor of the original value. The intuition is that, the supremum of a Gaussian process, though a global property, can be controlled by pairwise correlations between the variables. To quantify this, we use Slepian’s lemma, that helps us relate the supremum of two Gaussian processes by comparing pairwise correlations. Finally, observe that using known derandomizations of JLL, the dimension reduction can be done deterministically in time $\text{poly}(d, m, 1/\varepsilon)$ [EIO02].

Thus, to obtain a PTAS it would be enough to have a deterministic algorithm to approximate the supremum of a Gaussian process in time exponential in the dimension $k = O((\log m)/\varepsilon^2)$. Unfortunately, a naive argument by discretizing the Gaussian measure in \mathbb{R}^k leads to a run-time of at least $k^{O(k)}$; which gives a $m^{O((\log \log m)/\varepsilon^2)}$ algorithm. This question was recently addressed by Dadush and Vempala [DV12], who needed a similar sub-routine for their work on computing *M-Ellipsoids* of convex sets and give a deterministic algorithm with a run-time of $(\log k)^{O(k)}$. Combining their algorithm with the dimension reduction step gives a deterministic $m^{O((\log \log \log m)/\varepsilon^2)}$

³The hitting time $\tau_{hit}(G)$ is defined as the maximum over all pairs of vertices $u, v \in G$ of the expected time for a random walk starting at u to reach v . See the discussion in [Din11] for why this is a reasonable condition.

time algorithm for approximating the supremum. We next get rid of this $\omega(1)$ dependence in the exponent.

2.2 Oblivious Linear Estimators for Semi-Norms

We in fact, solve a more general problem by constructing an optimal *linear estimator* for semi-norms in Gaussian space. Let $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}_+$ be a semi-norm, i.e., φ is homogeneous and satisfies triangle inequality. For normalization purposes, we assume that $1 \leq \mathbb{E}_{x \leftarrow \mathcal{N}^k}[\varphi(x)]$ and that the Lipschitz constant of φ is at most $k^{O(1)}$. Note that the supremum function $\varphi_V(x) = \sup_{v \in V} |\langle v, x \rangle|$ satisfies these conditions. Our goal will be to compute a $(1 + \varepsilon)$ -factor approximation to $\mathbb{E}_{x \leftarrow \mathcal{N}^k}[\varphi(x)]$ in time $2^{O_\varepsilon(k)}$.

Theorem 2.1. *For every $\varepsilon > 0$, there exists a deterministic algorithm running in time $(1/\varepsilon)^{O(k)}$ and space $\text{poly}(k, 1/\varepsilon)$ that computes a $(1 + \varepsilon)$ -factor approximation to $\mathbb{E}_{X \leftarrow \mathcal{N}^k}[\varphi(X)]$ using only oracle access to φ .*

Our algorithm has the additional property of being an *oblivious linear estimator*: the set of query points does not depend on φ and the output is a positive weighted sum of the evaluations of φ on the query points. Further, the construction is essentially optimal as any such oblivious estimator needs to make at least $(1/\varepsilon)^{\Omega(k)}$ queries (see [Section 7](#)). In comparison, the previous best bound of Dadush and Vempala [[DV12](#)] needed $(\log k)^{O(k)}$ queries. We also remark that the query points of our algorithm are essentially the same as that of Dadush and Vempala, however our analysis is quite different and leads to better parameters.

As in the analysis of the dimension reduction step, our analysis of the oblivious estimator relies on a comparison inequality—Kantor’s lemma—that allows us to “lift” a simple estimator for the univariate case to the multi-dimensional case.

We first construct a symmetric distribution μ on \mathbb{R} that has a simple *piecewise flat graph* and *sandwiches* the one-dimensional Gaussian distribution in the following sense. Let ν be a “shrinking” of μ defined to be the probability density function (pdf) of $(1 - \varepsilon)x$ for $x \leftarrow \mu$. Then, for every symmetric interval $I \subseteq \mathbb{R}$, $\mu(I) \leq \mathcal{N}(I) \leq \nu(I)$.

Kantor’s lemma [[Kan77](#)] says that for pdf’s μ, ν as above that are in addition *unimodal*, the above relation carries over to the product distributions μ^k, ν^k : for every symmetric convex set $K \subseteq \mathbb{R}^k$, $\mu^k(K) \leq \mathcal{N}^k(K) \leq \nu^k(K)$. This last inequality immediately implies that semi-norms cannot *distinguish* between μ^k and \mathcal{N}^k : for any semi-norm φ , $\mathbb{E}_{\mu^k}[\varphi(x)] = (1 \pm \varepsilon) \mathbb{E}_{\mathcal{N}^k}[\varphi(x)]$. We then suitably prune the distribution μ^k to have small support and prove [Theorem 4.1](#).

Our main result, [Theorem 1.2](#), follows by first reducing the dimension as in the previous section and applying [Theorem 4.1](#) to the semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}_+$, $\varphi(x) = \sup_i |\langle u_i, x \rangle|$ for the projected vectors $\{u_1, \dots, u_m\}$.

3 Dimension Reduction

The use of JLL type random projections for estimating the supremum comes from the following comparison inequality for Gaussian processes. We call a collection of real-valued random variables $\{X_t\}_{t \in T}$ a Gaussian process if every finite linear combination of the variables has a normal distribution with mean zero. For a reference to Slepian’s lemma we refer the reader to Corollary 3.14 and the following discussion in [[LT91](#)].

Theorem 3.1 (Slepian’s Lemma [Sle62]). *Let $\{X_t\}_{t \in T}$ and $\{Y_t\}_{t \in T}$ be two Gaussian processes such that for every $s, t \in T$, $\mathbb{E}[(X_s - X_t)^2] \leq \mathbb{E}[(Y_s - Y_t)^2]$. Then, $\mathbb{E}[\sup_t X_t] \leq \mathbb{E}[\sup_t Y_t]$.*

We also need a derandomized version of the Johnson-Lindenstrauss Lemma.

Theorem 3.2 ([EIO02]). *For every $\varepsilon > 0$, there exists a deterministic $(dm^2(\log m + 1/\varepsilon)^{O(1)})$ -time algorithm that given vectors $v_1, \dots, v_m \in \mathbb{R}^d$ computes a linear mapping $A : \mathbb{R}^d \rightarrow \mathbb{R}^k$ for $k = O((\log m)/\varepsilon^2)$ such that for every $i, j \in [m]$, $\|v_i - v_j\|_2 \leq \|A(v_i) - A(v_j)\|_2 \leq (1 + \varepsilon)\|v_i - v_j\|_2$.*

Combining the above two theorems immediately implies the following.

Lemma 3.3. *For every $\varepsilon > 0$, there exists a deterministic $(dm^2(\log m + 1/\varepsilon)^{O(1)})$ -time algorithm that given vectors $v_1, \dots, v_m \in \mathbb{R}^d$ computes a linear mapping $A : \mathbb{R}^d \rightarrow \mathbb{R}^k$ for $k = O((\log m)/\varepsilon^2)$ such that*

$$\mathbb{E}_{x \leftarrow \mathcal{N}^d} [\sup_i |\langle v_i, x \rangle|] \leq \mathbb{E}_{y \leftarrow \mathcal{N}^k} [\sup_i |\langle A(v_i), y \rangle|] \leq (1 + \varepsilon) \mathbb{E}_{x \leftarrow \mathcal{N}^d} [\sup_i |\langle v_i, x \rangle|]. \quad (3.1)$$

Proof. Let $V = \{v_1, \dots, v_m\} \cup \{-v_1, \dots, -v_m\}$ and let $\{X_v\}_{v \in V}$ be the Gaussian process where the joint distribution is given by $X_v \equiv \langle v, x \rangle$ for $x \leftarrow \mathcal{N}^d$. Then, $\mathbb{E}_{x \leftarrow \mathcal{N}^d} [\sup_i |\langle v_i, x \rangle|] = \mathbb{E}[\sup_v X_v]$.

Let $A : \mathbb{R}^d \rightarrow \mathbb{R}^k$ be the linear mapping as given by Theorem 3.2 applied to V . Let $\{Y_v\}_{v \in V}$ be the “projected” Gaussian process with joint distribution given by $Y_v \equiv \langle A(v), y \rangle$ for $y \leftarrow \mathcal{N}^k$. Then, $\mathbb{E}_{y \leftarrow \mathcal{N}^k} [\sup_i |\langle v_i, y \rangle|] = \mathbb{E}[\sup_v Y_v]$.

Finally, observe that for any $u, v \in V$,

$$\mathbb{E}[(X_u - X_v)^2] = \|u - v\|_2^2 \leq \|A(u) - A(v)\|_2^2 = \mathbb{E}[(Y_u - Y_v)^2] \leq (1 + \varepsilon)^2 \mathbb{E}[(X_u - X_v)^2].$$

Combining the above inequality with Slepian’s lemma Lemma 3.1 applied to the pairs of processes $(\{X_v\}_{v \in V}, \{Y_v\}_{v \in V})$ and $(\{Y_v\}_{v \in V}, \{(1 + \varepsilon)X_v\}_{v \in V})$ it follows that

$$\mathbb{E}[\sup_v X_v] \leq \mathbb{E}[\sup_v Y_v] \leq \mathbb{E}[\sup_v (1 + \varepsilon)X_v] = (1 + \varepsilon) \mathbb{E}[\sup_v X_v].$$

The lemma now follows. □

4 Oblivious Estimators for Semi-Norms in Gaussian Space

In the previous section we reduced the problem of computing the supremum of a d -dimensional Gaussian process to that of a Gaussian process in $k = O((\log m)/\varepsilon^2)$ -dimensions. Thus, it suffices to have an algorithm for approximating the supremum of Gaussian processes in time exponential in the dimension. We will give such an algorithm that works more generally for all semi-norms.

Let $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}_+$ be a semi-norm. That is, φ satisfies the triangle inequality and is homogeneous. For normalization purposes we assume that $1 \leq \mathbb{E}_{\mathcal{N}^k}[\varphi(X)]$ and the Lipschitz constant of φ is at most $k^{O(1)}$.

Theorem 4.1. *For every $\varepsilon > 0$, there exists a set $S \subseteq \mathbb{R}^k$ with $|S| = (1/\varepsilon)^{O(k)}$ and a function $p : \mathbb{R}^k \rightarrow \mathbb{R}_+$ computable in $\text{poly}(k, 1/\varepsilon)$ time such that the following holds. For every semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}_+$,*

$$(1 - \varepsilon) \left(\sum_{x \in S} p(x) \varphi(x) \right) \leq \mathbb{E}_{X \leftarrow \mathcal{N}^k} [\varphi(X)] \leq (1 + \varepsilon) \left(\sum_{x \in S} p(x) \varphi(x) \right).$$

Moreover, successive elements of S can be enumerated in $\text{poly}(k, 1/\varepsilon)$ time and $O(k \log(1/\varepsilon))$ space.

[Theorem 2.1](#) follows immediately from the above.

Proof of [Theorem 2.1](#). Follows by enumerating over the set S and computing $\sum_{x \in S} p(x)\varphi(x)$ by querying φ on the points in S . \square

We now prove [Theorem 4.1](#). Here and henceforth, let γ denote the pdf of the standard univariate Gaussian distribution. Fix $\varepsilon > 0$ and let $\delta > 0$ be a parameter to be chosen later. Let $\mu \equiv \mu_{\varepsilon, \delta}$ be the pdf which is a piecewise-flat approximator to γ obtained by spreading the mass γ gives to an interval $I = [i\delta, (i+1)\delta)$ evenly over I . Formally, $\mu(z) = \mu(-z)$ and for $z > 0$, $z \in [i\delta, (i+1)\delta)$,

$$\mu(z) = \frac{\gamma([i\delta, (i+1)\delta))}{\delta}. \quad (4.1)$$

Clearly, μ defines a symmetric distribution on \mathbb{R} . We will show that for $\delta \ll \varepsilon$ sufficiently small, semi-norms cannot *distinguish* the product distribution μ^k from \mathcal{N}^k :

Lemma 4.2. *Let $\delta = (2\varepsilon)^{3/2}$. Then, for every semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$,*

$$(1 - \varepsilon) \mathbb{E}_{X \leftarrow \mu^k} [\varphi(X)] \leq \mathbb{E}_{Z \leftarrow \mathcal{N}^k} [\varphi(Z)] \leq \mathbb{E}_{X \leftarrow \mu^k} [\varphi(X)].$$

We first prove [Theorem 4.1](#) assuming the above lemma, whose proof is deferred to the next section.

Proof of [Theorem 4.1](#). Let $\hat{\mu}$ be the symmetric distribution supported on $\delta(\mathbb{Z} + 1/2)$ with pdf defined by

$$\hat{\mu}(\delta(i + 1/2)) = \mu([i\delta, (i+1)\delta)),$$

for $i \geq 0$. Further, let $X \leftarrow \mu^k$, $\hat{X} \leftarrow \hat{\mu}^k$, $Z \leftarrow \mathcal{N}^k$.

We claim that $\mathbb{E}[\varphi(\hat{X})] = (1 \pm \varepsilon) \mathbb{E}[\varphi(Z)]$. Let Y be uniformly distributed on $[-\delta, \delta]^k$ and observe that random variable $X \equiv \hat{X} + Y$ in law. Therefore,

$$\begin{aligned} \mathbb{E}[\varphi(X)] &= \mathbb{E}[\varphi(\hat{X} + Y)] = \mathbb{E}[\varphi(\hat{X})] \pm \mathbb{E}[\varphi(Y)] = \mathbb{E}[\varphi(\hat{X})] \pm \delta \mathbb{E}[\varphi(Y/\delta)] \\ &= \mathbb{E}[\varphi(\hat{X})] \pm \delta \mathbb{E}_{Z' \in_u [-1, 1]^k} [\varphi(Z')] = \mathbb{E}[\varphi(\hat{X})] \pm \delta \mathbb{E}[\varphi(Z)] \quad (\text{Lemma 5.7}). \end{aligned} \quad (4.2)$$

Thus, by [Lemma 4.2](#),

$$\mathbb{E}[\varphi(\hat{X})] = (1 \pm O(\varepsilon)) \mathbb{E}[\varphi(Z)] \quad (4.3)$$

We next prune $\hat{\mu}^k$ to reduce its support. Define $p : \mathbb{R}^k \rightarrow \mathbb{R}_+$ by $p(x) = \hat{\mu}^k(x)$. Clearly, $p(x)$ being a product distribution is computable in $\text{poly}(k, 1/\varepsilon)$ time.

Let $S = (\delta(\mathbb{Z} + 1/2))^k \cap B_2(3\sqrt{k})$, where $B_2(r) \subseteq \mathbb{R}^k$ denotes the Euclidean ball of radius r . As φ has Lipschitz constant bounded by $k^{O(1)}$, a simple calculation shows that throwing away all points in the support of \hat{X} outside S does not change $\mathbb{E}[\varphi(\hat{X})]$ much. It is easy to check that for $x \notin S$, $p(x) \leq \exp(-\|x\|_2^2/4)/(2\pi)^{k/2}$. Therefore,

$$\begin{aligned} \mathbb{E}[\varphi(\hat{X})] &= \sum_x p(x)\varphi(x) = \sum_{x \in S} p(x)\varphi(x) + \sum_{x \notin S} p(x)\varphi(x) \\ &= \sum_{x \in S} p(x)\varphi(x) \pm \sum_{x \notin S} \frac{\exp(-\|x\|_2^2/4)}{(2\pi)^{k/2}} \cdot (k^{O(1)}\|x\|_2) = \sum_{x \in S} p(x)\varphi(x) \pm o(1). \end{aligned} \quad (4.4)$$

From Equation 4.3 and the above equation we get (recall that $\mathbb{E}[\varphi(Z)] \geq 1$)

$$\mathbb{E}[\varphi(Z)] = (1 \pm O(\varepsilon)) \left(\sum_{x \in S} p(x) \varphi(x) \right),$$

which is what we want to show.

We now reason about the complexity of S . First, by a simple covering argument $|S| < (1/\delta)^{O(k)}$:

$$|S| < \frac{\text{Vol}(B_2(3\sqrt{k}) + [-\delta, \delta]^k)}{\text{Vol}([- \delta, \delta]^k)} = (1/\delta)^{O(k)} = (1/\varepsilon)^{O(k)},$$

where for sets $A, B \subseteq \mathbb{R}^k$, $A+B$ denotes the Minkowski sum and Vol denotes Lebesgue volume. This size bound almost suffices to prove Theorem 4.1 except for the complexity of enumerating elements from S . Without loss of generality assume that $R = 3\sqrt{n}/\delta$ is an integer. Then, enumerating elements in S is equivalent to enumerating integer points in the n -dimensional ball of radius R . This can be accomplished by going through the set of lattice points in the natural lexicographic order, and takes $\text{poly}(k, 1/\varepsilon)$ time and $O(k \log(1/\varepsilon))$ space per point in S . \square

5 Proof of Lemma 4.2

Our starting point is the following definition that helps us *compare* multivariate distributions when we are only interested in volumes of convex sets. We shall follow the notation of [Bal03].

Definition 5.1. *Given two symmetric pdf's, f, g on \mathbb{R}^k , we say that f is less peaked than g ($f \preceq g$) if for every symmetric convex set $K \subseteq \mathbb{R}^k$, $f(K) \leq g(K)$.*

We also need the following elementary facts. The first follows from the unimodality of the Gaussian density and the second from partial integration.

Fact 5.2. *For any $\delta > 0$ and μ as defined by Equation 4.1, μ is less peaked than γ .*

Fact 5.3. *Let f, g be distributions on \mathbb{R}^k with $f \preceq g$. Then for any semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$, $\mathbb{E}_f[\varphi(x)] \geq \mathbb{E}_g[\varphi(x)]$.*

Proof. Observe that for any $t > 0$, $\{x : \varphi(x) \leq t\}$ is convex. Let random variables $X \leftarrow f$, $Y \leftarrow g$. Then, by partial integration, $\mathbb{E}[\varphi(X)] = \int_0^\infty \varphi'(t) \Pr[\varphi(X) > t] dt \geq \int_0^\infty \varphi'(t) \Pr[\varphi(Y) > t] dt = \mathbb{E}[\varphi(Y)]$. \square

The above statements give us a way to compare the expectations of μ and γ for one-dimensional convex functions. We would now like to do a similar comparison for the product distributions μ^k and γ^k . For this we use Kanter's lemma [Kan77], which says that the relation \preceq is preserved under tensoring if the individual distributions have the additional property of being *unimodal*.

Definition 5.4. *A distribution f on \mathbb{R}^n is unimodal if f can be written as an increasing limit of a sequence of distributions each of which is a finite positively weighted sum of uniform distributions on symmetric convex sets.*

Theorem 5.5 (Kanter's Lemma [Kan77]; cf. [Bal03]). *Let μ_1, μ_2 be symmetric distributions on \mathbb{R}^n with $\mu_1 \preceq \mu_2$ and let ν be a unimodal distribution on \mathbb{R}^m . Then, the product distributions $\mu_1 \times \nu$, $\mu_2 \times \nu$ on $\mathbb{R}^n \times \mathbb{R}^m$ satisfy $\mu_1 \times \nu \preceq \mu_2 \times \nu$.*

We next show that μ “sandwiches” γ in the following sense.

Lemma 5.6. *Let ν be the pdf of the random variable $y = (1-\varepsilon)x$ for $x \leftarrow \mu$. Then, for $\delta \leq (2\varepsilon)^{3/2}$, $\mu \preceq \gamma \preceq \nu$.*

Proof. As mentioned above, $\mu \preceq \gamma$. We next show that $\gamma \preceq \nu$. Intuitively, ν is obtained by spreading the mass that γ puts on an interval $I = [i\delta, (i+1)\delta)$ evenly on the *smaller* interval $(1-\varepsilon)I$. The net effect of this operation is to push the pdf of μ closer towards the origin and for δ sufficiently small the inward push from this “shrinking” wins over the outward push from going to μ .

Fix an interval $I = [-i\delta(1-\varepsilon) - \theta, i\delta(1-\varepsilon) + \theta]$ for $0 \leq \theta < \delta(1-\varepsilon)$. Then,

$$\nu(I) = \nu([-i\delta(1-\varepsilon), i\delta(1-\varepsilon)]) + 2\nu([i\delta(1-\varepsilon), i\delta(1-\varepsilon) + \theta]) \quad (5.1)$$

$$= \gamma([-i\delta, i\delta]) + \frac{2\theta \cdot \gamma([i\delta, (i+1)\delta])}{\delta(1-\varepsilon)}. \quad (5.2)$$

We now consider two cases.

Case 1: $i \geq (1-\varepsilon)/\varepsilon$ so that $i\delta(1-\varepsilon) + \theta \leq i\delta$. Then, from the above equation,

$$\nu(I) \geq \gamma([-i\delta, i\delta]) \geq \gamma([-i\delta(1-\varepsilon) - \theta, i\delta(1-\varepsilon) + \theta]) = \gamma(I).$$

Case 2: $i < (1-\varepsilon)/\varepsilon$. Let $\alpha = (i+1)\delta = \delta/\varepsilon$. Then, as $1 - x^2/2 \leq e^{-x^2/2} \leq 1$,

$$\gamma([i\delta, i\delta + \theta]) \leq \theta \cdot \gamma(0), \quad \gamma([i\delta, (i+1)\delta]) \geq \delta \cdot \gamma(0) \cdot (1 - \alpha^2/2).$$

Therefore,

$$\begin{aligned} \nu(I) &= \gamma(I) - 2\gamma([i\delta, i\delta(1-\varepsilon) + \theta]) + \frac{2\theta \cdot \gamma([i\delta, (i+1)\delta])}{\delta(1-\varepsilon)} \\ &\geq \gamma(I) - 2\gamma([i\delta, i\delta + \theta]) + \frac{2\theta \cdot \gamma([i\delta, (i+1)\delta])}{\delta(1-\varepsilon)} \\ &\geq \gamma(I) - 2\theta\gamma(0) + \frac{2\theta \cdot \delta \cdot \gamma(0) \cdot (1 - \alpha^2/2)}{\delta(1-\varepsilon)} \\ &= \gamma(I) + \frac{2\theta\gamma(0)}{1-\varepsilon} \cdot (\varepsilon - \alpha^2/2) \geq \gamma(I), \end{aligned}$$

for $\alpha^2 \leq 2\varepsilon$, i.e., if $\delta \leq (2\varepsilon)^{3/2}$. □

Lemma 4.2 follows easily from the above two claims.

Proof of Lemma 4.2. Clearly, μ, ν, γ are unimodal and product of unimodal distributions is unimodal. Thus, from the above lemma and iteratively applying Kanter’s lemma we get $\mu^k \preceq \gamma^k \preceq \nu^k$. Therefore, by Fact 5.3, for any semi-norm φ ,

$$\mathbb{E}_{\mu^k}[\varphi(X)] \geq \mathbb{E}_{\gamma^k}[\varphi(Y)] \geq \mathbb{E}_{\nu^k}[\varphi(X)] = \mathbb{E}_{\mu^k}[\varphi((1-\varepsilon)X)] = (1-\varepsilon) \mathbb{E}_{\mu^k}[\varphi(X)].$$

□

We now prove the auxiliary lemma we used in proof of Theorem 4.1.

Lemma 5.7. *Let ρ be the uniform distribution on $[-1, 1]$. Then, $\gamma \preceq \rho$ and for any semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$, $\mathbb{E}_{\rho^k}[\varphi(x)] \leq \mathbb{E}_{\gamma^k}[\varphi(x)]$.*

Proof. It is easy to check that $\gamma \preceq \rho$. Then, by Kanter’s lemma $\gamma^k \preceq \rho^k$ and the inequality follows from Fact 5.3. □

6 A PTAS for Supremum of Gaussian Processes

Our main theorem, [Theorem 1.2](#), follows immediately from [Lemma 3.3](#) and [Theorem 2.1](#) applied to the semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$ defined by $\varphi(x) = \sup_{i \leq m} |\langle A(v_i), x \rangle|$.

7 Lowerbound for Oblivious Estimators

We now show that [Theorem 4.1](#) is optimal: any oblivious linear estimator for semi-norms as in the theorem must make at least $(C/\varepsilon)^k$ queries for some constant $C > 0$.

Let $S \subseteq \mathbb{R}^k$ be the set of query points of an oblivious estimator. That is, there exists a function $f : \mathbb{R}_+^S \rightarrow \mathbb{R}_+$ such that for any semi-norm $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}_+$, $f((\varphi(x) : x \in S)) = (1 \pm \varepsilon) \mathbb{E}_{Y \leftarrow \mathcal{N}^k}[\varphi(Y)]$. We will assume that f is monotone in the following sense: $f(x_1, \dots, x_{|S|}) \leq f(y_1, \dots, y_{|S|})$ if $0 \leq x_i \leq y_i$ for all i . This is clearly true for any linear estimator (and also for the median estimator). Without loss of generality suppose that $\varepsilon < 1/4$.

The idea is to define a suitable semi-norm based on S : define $\varphi : \mathbb{R}^k \rightarrow \mathbb{R}$ by $\varphi(x) = \sup_{u \in S} |\langle u/\|u\|_2, x \rangle|$. It is easy to check that for any $v \in S$, $\|v\|_2 \leq \varphi(v)$. Therefore, the output of the oblivious estimator when querying the Euclidean norm is at most the output of the estimator when querying φ . In particular,

$$(1 - \varepsilon) \mathbb{E}_{Y \leftarrow \mathcal{N}^k}[\|Y\|_2] \leq f((\|x\|_2 : x \in S)) \leq f((\varphi(x) : x \in S)) \leq (1 + \varepsilon) \mathbb{E}_{Y \leftarrow \mathcal{N}^k}[\varphi(Y)]. \quad (7.1)$$

We will argue that the above is possible only if $|S| > (C/\varepsilon)^k$. Let \mathcal{S}^{k-1} denote the unit sphere in \mathbb{R}^k . For the remaining argument, we shall view $Y \leftarrow \mathcal{N}^k$ to be drawn as $Y = RX$, where $X \in \mathcal{S}^{k-1}$ is uniformly random on the sphere and $R \in \mathbb{R}$ is independent of X and has a Chi-squared distribution with k degrees of freedom. Let $S(\varepsilon) = \cup_{u \in S} \{y \in \mathcal{S}^{k-1} : |\langle u/\|u\|_2, y \rangle| \geq 1 - 4\varepsilon\}$.

Now, by a standard volume argument, for any $y \in \mathcal{S}^{k-1}$, $\Pr_X[|\langle X, y \rangle| \geq 1 - 4\varepsilon] < (O(\varepsilon))^k$. Thus, by a union bound, $p = \Pr_X[X \in S(\varepsilon)] < |S| \cdot (O(\varepsilon))^k$. Further, for any $y \in \mathcal{S}^{k-1} \setminus S(\varepsilon)$, $\varphi(y) < 1 - 4\varepsilon$. Therefore,

$$\begin{aligned} \mathbb{E}_X[\varphi(X)] &= \Pr[X \notin S(\varepsilon)] \cdot \mathbb{E}[\varphi(X) | X \notin S(\varepsilon)] + \Pr[X \in S(\varepsilon)] \cdot \mathbb{E}[\varphi(X) | X \in S(\varepsilon)] \leq \\ &\quad (1 - p)(1 - 4\varepsilon) + p. \end{aligned}$$

Thus,

$$\mathbb{E}[\varphi(Y)] = \mathbb{E}[\varphi(RX)] = \mathbb{E}[R] \cdot \mathbb{E}[\varphi(X)] \leq \mathbb{E}[\|Y\|_2] \cdot ((1 - p)(1 - 4\varepsilon) + p). \quad (7.2)$$

Combining Equations [7.1](#) and [7.2](#), we get

$$1 - \varepsilon \leq (1 + \varepsilon) \cdot ((1 - p)(1 - 4\varepsilon) + p) < 1 - 3\varepsilon + 2p.$$

As $p < |S| \cdot (O(\varepsilon))^k$, the above leads to a contradiction unless $|S| > (C/\varepsilon)^k$ for some constant $C > 0$.

References

- [AF] D. Aldous and J. Fill, *Reversible markov chains and random walks on graphs*, In preparation, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.

- [Bal03] Keith Ball, *Convex geometry and functional analysis*, Handbook of the Geometry of Banach Spaces, vol. 1, Elsevier, 2003.
- [Din11] Jian Ding, *Asymptotics of cover times via gaussian free fields: bounded-degree graphs and general trees*, 2011, Preprint, available at <http://arxiv.org/abs/1103.4402>.
- [DLP11] Jian Ding, James R. Lee, and Yuval Peres, *Cover times, blanket times, and majorizing measures*, STOC, 2011, pp. 61–70.
- [DV12] Daniel Dadush and Santosh Vempala, *Deterministic construction of an approximate m-ellipsoid and its applications to derandomizing lattice algorithms*, SODA, 2012, pp. 1445–1456.
- [EIO02] Lars Engebretsen, Piotr Indyk, and Ryan O’Donnell, *Derandomized dimensionality reduction with applications*, SODA, 2002, pp. 705–712.
- [FZ09] U. Feige and O. Zeitouni, *Deterministic approximation for the cover time of trees*, 2009, Preprint, available at <http://arxiv.org/abs/0909.2005>.
- [Kan77] Marek Kanter, *Unimodality and dominance for symmetric random vectors*, Transactions of the American Mathematical Society **229** (1977), pp. 65–85.
- [KKLV00] Jeff Kahn, Jeong Han Kim, László Lovász, and Van H. Vu, *The cover time, the blanket time, and the matthews bound*, FOCS, 2000, pp. 467–475.
- [Lee10] James Lee, *Open question: Cover times and the gaussian free field*, 2010, Post in Blog tcsmath – “Open Question: Cover Times and the Gaussian Free Field” (Dec 9, 2010).
- [Lov93] László Lovász, *Random walks on graphs: A survey*, Combinatorics Paul Erdos is Eighty **2** (1993), no. 1, 1–46.
- [LT91] Michel Ledoux and Michel Talagrand, *Probability in banach spaces: Isoperimetry and processes*, Springer, Berlin, 1991.
- [Pis99] G. Pisier, *The volume of convex bodies and banach space geometry*, Cambridge tracts in mathematics, Cambridge University Press, 1999.
- [Sle62] D Slepian, *The one-sided barrier problem for gaussian noise*, Bell System Technical Journal (1962), 463–501.
- [Tal05] M. Talagrand, *The generic chaining: upper and lower bounds of stochastic processes*, Springer monographs in mathematics, Springer, 2005.
- [WZ96] Peter Winkler and David Zuckerman, *Multiple cover time*, Random Struct. Algorithms **9** (1996), no. 4, 403–411.