APPROXIMATING THE RADII OF POINT SETS[*]

KASTURI VARADARAJAN[†], S. VENKATESH[‡], YINYU YE[§], AND JIAWEI ZHANG[¶]

Abstract. We consider the problem of computing the outer-radii of point sets. In this problem, we are given integers n, d, and k, where $k \leq d$, and a set P of n points in $\Re^d$. The goal is to compute the *outer k-radius* of P, denoted by $\mathcal{R}_k(P)$, which is the minimum over all $(d-k)$-dimensional flats F of $\max_{p \in P} d(p, F)$, where $d(p, F)$ is the Euclidean distance between the point p and flat F. Computing the radii of point sets is a fundamental problem in computational convexity with many significant applications. The problem admits a polynomial time algorithm when the dimension d is constant [U. Faigle, W. Kern, and M. Streng, *Math. Program.*, 73 (1996), pp. 1–5]. Here we are interested in the general case in which the dimension d is not fixed and can be as large as n, where the problem becomes NP-hard even for $k = 1$. It is known that $R_k(P)$ can be approximated in polynomial time by a factor of $(1 + \varepsilon)$ for any $\varepsilon > 0$ when $d - k$ is a fixed constant [M. Bădoiu, S. Har-Peled, and P. Indyk, in *Proceedings of the ACM Symposium on the Theory of Computing*, 2002; S. Har-Peled and K. Varadarajan, in *Proceedings of the ACM Symposium on Computing Geometry*, 2002]. A polynomial time algorithm that guarantees a factor of $O(\sqrt{\log n})$ approximation for $R_1(P)$, the width of the point set P, is implied by the results of Nemirovski, Roos, and Terlaky [*Math. Program.*, 86 (1999), pp. 463–473] and Nesterov [*Handbook of Semidefinite Programming Theory, Algorithms*, Kluwer Academic Publishers, Norwell, MA, 2000]. In this paper, we show that $R_k(P)$ can be approximated by a ratio of $O(\sqrt{\log n})$ for any $1 \leq k \leq d$, thus matching the previously best known ratio for approximating the special case $R_1(P)$, the width of point set P. Our algorithm is based on semidefinite programming relaxation with a new mixed deterministic and randomized rounding procedure. We also prove an inapproximability result that gives evidence that our approximation algorithm is doing well for a large range of k. We show that there exists a constant $\delta > 0$ such that the following holds for any $0 < \varepsilon < 1$: there is no polynomial time algorithm that approximates $R_k(P)$ within $(\log n)^\delta$ for all k such that $k \leq d - d^\varepsilon$ unless NP $\subseteq$ DTIME $[2^{(\log m)^{O(1)}}]$. Our inapproximability result for $R_k(P)$ extends a previously known hardness result of Brieden [*Discrete Comput. Geom.*, 28 (2002), pp. 201–209] and is proved by modifying Brieden's construction using basic ideas from probabilistically checkable proofs (PCP) theory.

Key words. approximation algorithms, semidefinite programming, computational convexity

AMS subject classifications. 68W20, 68W25, 68W40

DOI. 10.1137/050627472

1. Introduction. Computing the outer k-radius of a point set is a fundamental problem in computational convexity with applications in global optimization, data

mining, statistics, and clustering, and it has received considerable attention in the computational geometry literature [20, 21, 22]. In this problem, we are given integers n, d, and k, where $k \leq d$, and a set P of n points in $\Re^d$. A *flat* or *affine subspace* F in $\Re^d$ is specified by a point $q \in \Re^d$ and a linear subspace H; it is defined as $F = \{q + h | h \in H\}$. The dimension of the flat F is defined to be the dimension of the linear subspace H. For any flat F, let $R(P, F) = \max_{p \in P} d(p, F)$ denote the *radius* of the flat F with respect to P, where $d(p, F)$ is the Euclidean distance between the point p and the flat F. The goal is to compute the *outer k-radius* of P, denoted by $R_k(P)$, which is the minimum of $R(P, F)$ over all $(d - k)$-dimensional flats F. A $(d - k)$-flat is simply a flat of dimension $d - k$. Roughly speaking, the outer k-radius $R_k(P)$ measures how well the point set P can be approximated by an affine subspace of dimension $d - k$. A few special cases of $R_k(P)$ which have received particular attention include $R_1(P)$, half of the *width* of P; $R_d(P)$, the radius of the minimum enclosing ball of P; and $R_{d-1}(P)$, the radius of the minimum enclosing cylinder of P.

When the dimension d is a fixed constant, $R_k(P)$ can be computed exactly in polynomial time [15]. It is also known that $R_k(P)$ can be approximated by a factor of $(1+\varepsilon)$ for any $\varepsilon > 0$ in $O(n + f_d(\frac{1}{\varepsilon}))$ time [2, 6], where f_d is a polynomial for every fixed d. In this paper, we are interested in the general scenario when the dimensions k and d are not fixed and d can be as large as n.

When the dimensions k and d are part of the input, the complexity of computing/approximating $R_k(P)$ depends on the parameter $d - k$. It is well known that the problem is polynomial time solvable when $d - k = 0$, i.e., the minimum enclosing ball of a set of points can be computed in polynomial time (Gritzmann and Klee [20]). Megiddo [25] shows that the problem of determining whether there is a line that intersects a set of balls is NP-hard. In his reduction, the balls have the same radius, which implies that computing the radius $R_{d-1}(P)$ of the min-enclosing cylinder of a set of points P is NP-hard. Bǎdoiu, Har-Peled, and Indyk [7] show that $R_{d-1}(P)$ can be approximated in polynomial time by a factor of $(1 + \varepsilon)$ for any $\varepsilon > 0$. Har-Peled and Varadarajan [22, 23] generalize the result and show that $R_k(P)$ can be approximated by a factor of $(1 + \varepsilon)$ for any $\varepsilon > 0$ when $d - k$ is constant.[1]

More hardness results are known when $d - k$ becomes large or when k becomes small. Bodlaender et al. [10] show that the problem is NP-hard when $k = 1$. This is true even for the case $n = d + 1$ [20]. Gritzmann and Klee [20] also show that it is NP-hard to compute $R_k(P)$ if $k \leq c \cdot d$ for any fixed $0 < c < 1$. These negative results are further improved by Brieden, Gritzmann, and Klee [11] and Brieden [14], the latter of which has shown that it is NP-hard to approximate $R_1(P)$, the width of a point set, to within *any* constant factor.

On the positive side, the algorithms of Nemirovski, Roos, and Terlaky [26] and Nesterov [27] imply that $R_1(P)$, or equivalently the width of the point set P, can be approximated within a factor of $O(\sqrt{\log n})$. Another algorithm for approximating the width of a point set is given by Brieden et al. [12, 13], and their algorithm has a performance guarantee $\sqrt{d/\log d}$ that is measured in the dimension d. Their algorithm in fact works for any convex body given in terms of appropriate "oracles"; the number of calls to the oracle is polynomial in the dimension d. They also show that this is the best possible result in the oracle model even if randomization is allowed. (Their algorithm actually gives a $\sqrt{d/\log n}$ approximation with poly(n) calls to the oracle, where n is the number of points in the set.) It is not clear if their algorithm can be extended to compute $R_k(P)$.

---

[1] Note that R_k^{opt} in [23] is the same as R_{d-k} in this paper

The problem of efficiently computing the low-rank approximation of matrices has received considerable attention recently; see [1, 5, 17] and the references cited in these papers. This problem corresponds to computing the best $(d-k)$-dimensional subspace that fits a point set, where the quality of a subspace is the *sum* of the square of the distance of each point from the flat. The problem is therefore related to the one we study in this paper, where the quality of a flat is the maximum over the point-flat distances. However, the low-rank approximation problem can be solved in polynomial time for any $1 \leq k \leq d$.

Our results and an overview. We show that $R_k(P)$ can be approximated in polynomial time by a factor of $O(\sqrt{\log n})$ for all $1 \leq k \leq d$, thereby generalizing the result of Nemirovski, Roos, and Terlaky [26] to all values of k. Our algorithm is based on a semidefinite programming (SDP) relaxation with a mixed deterministic and randomized rounding procedure, in contrast to all other purely randomized rounding procedures used for semidefinite programming approximation.

Generally speaking, the problem of computing $R_k(P)$ can be formulated as a quadratic minimization problem. SDP problems (where the unknowns are represented by positive semidefinite matrices) have recently been developed for approximating such problems; see, for example, Goemans and Williamson [18]. In the case of $k = 1$, computing $R_1(P)$ corresponds to a SDP problem plus an additional requirement that the rank of the unknown matrix equals 1. Removing the rank requirement, the SDP problem becomes a relaxation of the original problem and can be solved within any given accuracy $\epsilon > 0$ in time polynomial in $\ln \frac{1}{\epsilon}$ and the dimension of the data specifying the problem. Once obtaining an optimal solution, say X, of the SDP relaxation, one would like to generate a rank-1 matrix, say $\hat{X} = yy^T$, from X, where y is a column vector and serves as a solution to the original problem. Such rank reduction is called "rounding." Many rounding procedures are proposed, and almost all of them are randomized; see, for example, [9].

One particular procedure has been proposed by Nemirovski, Roos, and Terlaky [26] which can be used for approximating $R_1(P)$. Their procedure is a simple randomized rounding that can be described as follows: an optimal solution X of the SDP relaxation, whose rank could be as large as d, can be represented as (for example, by eigenvector decomposition)

$$X = \lambda_1 v_1 v_1^T + \lambda_2 v_2 v_2^T + \cdots + \lambda_d v_d v_d^T.$$

Then one can generate a single vector y by taking a random linear combination of the vectors $\sqrt{\lambda_1} v_1, \sqrt{\lambda_2} v_2, \ldots, \sqrt{\lambda_d} v_d$, where the coefficients of the combination take values of -1 or 1 uniformly and independently.

For the case $k \geq 2$, the SDP relaxation that we describe is best viewed as a direct relaxation of the problem of computing $R_k(P)$, rather than one that is obtained via a quadratic program formulation of $R_k(P)$. We then need to generate k rank-1 matrices from X, the optimal solution of the SDP relaxation, such that

$$\hat{X} = \sum_{i=1}^{k} y_i y_i^T,$$

where y_is are orthogonal to each other. Our rounding procedure works as follows: having obtained an optimal solution for the SDP relaxation with

$$X = \lambda_1 v_1 v_1^T + \lambda_2 v_2 v_2^T + \cdots + \lambda_d v_d v_d^T,$$

we deterministically partition the vectors $v_1, v_2, \ldots, v_d$ into k groups where group j may contain n_j vectors and each group can be seen as a single semidefinite matrix with rank n_j. We then generate one vector from each group using the randomized rounding procedure similar to that of Nemirovski, Roos, and Terlaky [26]. The k vectors generated by this rounding procedure will automatically satisfy the condition that any pair of them must be orthogonal to each other. We then manage to show that the quality of these vectors yields an approximation ratio of no more than $O(\sqrt{\log n})$.

We also prove an inapproximability result that gives evidence that our approximation algorithm is close to the best possible for a large range of k. We show that there exists a constant $\delta > 0$ such that the following holds for any $0 < \varepsilon < 1$: there is no polynomial time algorithm that approximates $R_k(P)$ within $(\log n)^\delta$ for all k such that $k \leq d - d^\varepsilon$ unless NP $\subseteq \tilde{P}$. $\tilde{P}$ denotes the complexity class DTIME$[2^{(\log m)^{O(1)}}]$, which is sometimes referred to as deterministic quasi-polynomial time. That is, $\tilde{P}$ contains the set of all problems for which there is an algorithm that runs in time $2^{(\log m)^{O(1)}}$ on inputs of size m.

To prove the lower bound result, we start with a two-prover protocol for 3SAT in which the verifier has very low error probability. Such a protocol is obtained as a consequence of the probabilistically checkable proofs (PCP) theorem of Arora et al. [3], and Arora and Safra [4] and the parallel repetition theorem of Raz [28]. The construction of Brieden [14] then implies a reduction from Max-3SAT to width computation such that the ratio of the width of point sets that correspond to satisfiable instances to those that correspond to unsatisfiable instances is large. This separation gives us the inapproximability result for the width. This result can then be extended to an inapproximability result for $R_k(P)$ for a large range of k.

The remainder of this paper is organized as follows: in section 2, we present our algorithm for approximating the outer k-radius $R_k(P)$ of a point set P. In section 3, we describe our inapproximability results. We make some concluding remarks in section 4.

2. Approximating the radius. We now present the quadratic program formulation of the outer k-radius problem and its SDP relaxation. It will be helpful to first introduce some notation that will be used later. The trace of a given square matrix A, denoted by $\text{Tr}(A)$, is the sum of the entries on the main diagonal of A. We use I to denote the identity matrix whose dimension will be clear in the context. The inner product of two vectors p and q is denoted by $\langle p, q \rangle$. The 2-norm of a vector x, denoted by $\|x\|$, is defined by $\sqrt{\langle x, x \rangle}$. For a matrix X, we use the notation $X \succeq 0$ to mean that X is a positive semidefinite matrix. For simplicity, we assume that P is symmetric in the sense that if $p \in P$, then $-p \in P$. This is without loss of generality for the following reason: we may, by performing a translation if necessary, assume that $0 \in P$. Denote the set $\{-p | p \in P\}$ by $-P$, and let $Q = P \cup -P$. It is clear that $R_k(P) \leq R_k(Q) \leq 2R_k(P)$. Therefore, if we found a good approximation for $R_k(Q)$, then it must also be a good approximation for $R_k(P)$.

Since P is a symmetric point set, the best $(d-k)$-flat for P contains the origin so that it is a subspace. Thus, the square of $R_k(P)$ can be defined by the optimal value of the following quadratic minimization problem:

(1)
$$
\begin{aligned}
R_k(P)^2 := \quad & \text{Minimize} && \alpha \\
& \text{Subject to} && \sum_{i=1}^{k} \langle p, x_i \rangle^2 \leq \alpha \;\; \forall p \in P, \\
& && \|x_i\|^2 = 1, \;\; i = 1, \ldots, k, \\
& && \langle x_i, x_j \rangle = 0 \;\; \forall i \neq j.
\end{aligned}
$$

Assume that $x_1, x_2, \ldots, x_k \in \Re^d$ is the optimal solution of (1). Then one can easily verify that the matrix $X = x_1 x_1^T + x_2 x_2^T + \cdots + x_k x_k^T$ is a feasible solution for the following semidefinite program:

(2)
$$
\begin{aligned}
\alpha_k^* := \quad & \text{Minimize} \quad \alpha \\
& \text{Subject to} \quad \text{Tr}(pp^T X) \; (= p^T X p) \leq \alpha \; \forall p \in P, \\
& \qquad\qquad \text{Tr}(X) = k, \\
& \qquad\qquad I - X \succeq 0, \; X \succeq 0.
\end{aligned}
$$

It follows that $\alpha_k^* \leq R_k(P)^2$. The following lemma follows from the above observations.

LEMMA 1. *There exists an integer $r \geq k$ such that we can compute, in polynomial time, r nonnegative reals $\lambda_1, \lambda_2, \ldots, \lambda_r$ and r orthogonal unit vectors $v_1, v_2, \ldots, v_r$ such that*
1. $\sum_{i=1}^r \lambda_i = k$.
2. $\max_{1 \leq i \leq r} \lambda_i \leq 1$.
3. $\sum_{i=1}^r \lambda_i \langle p, v_i \rangle^2 \leq R_k(P)^2$ *for any $p \in P$.*

Proof. We solve the semidefinite program (2) and let X^* be an optimal solution of (2). We claim that the rank of X^*, say r, is at least k. This follows from the fact that $\text{Tr}(X^*) = k$ and $I - X^* \succeq 0$. In other words, $\text{Tr}(X^*) = k$ implies that the sum of the eigenvalues of X^* is equal to k, and $I - X^* \succeq 0$ implies that all the eigenvalues are less than or equal to 1. Therefore, X^* has at least k nonzero eigenvalues, which implies that the rank of X^* is at least k. Let $\lambda_1, \lambda_2, \ldots, \lambda_r$ be the r nonnegative eigenvalues and $v_1, v_2, \ldots, v_r$ be the corresponding eigenvectors (see [26, p. 466] for details on computing the eigenvalues and eigenvectors in polynomial time). Then we have $\sum_{i=1}^r \lambda_i = k$ and $\max_{1 \leq i \leq r} \lambda_i \leq 1$. Furthermore, for any $p \in P$,

$$
\sum_{i=1}^r \lambda_i \langle p, v_i \rangle^2 = \text{Tr}(pp^T \sum_{i=1}^r \lambda_i v_i v_i^T) = \text{Tr}(pp^T X^*) \leq \alpha_k^* \leq R_k(P)^2. \qquad \square
$$

2.1. Deterministic first rounding. In this section, we prove a lemma concerning how to deterministically group the eigenvalues and their eigenvectors. The proof of the lemma is elementary, but it plays an important role for proving our main result.

LEMMA 2. *The index set $\{1, 2, \ldots, r\}$ can be partitioned into k sets $I_1, I_2, \ldots, I_k$ such that, for any $i : 1 \leq i \leq k$, $\sum_{j \in I_i} \lambda_j \geq \frac{1}{2}$.*

Proof. Recall that $\sum_{j=1}^r \lambda_j = k$ and $0 \leq \lambda_j \leq 1$ for all j. Without loss of generality, we can assume that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r$. Our partitioning algorithm is the same as the longest-processing-time heuristic algorithm for the parallel machine scheduling problem. The algorithm works as follows:
1. For $i = 1, 2, \ldots, k$, set $I_i = \emptyset$, and let $L_i = 0$. Let $I = \{1, 2, \ldots, r\}$.
2. While $I \neq \emptyset$,
    choose j from I with the smallest index;
    choose set i with the smallest value L_i.
    Let $I_i := I_i \cup \{j\}$, $L_i := L_i + \lambda_j$, and $I := I - \{j\}$.

It is clear that when the algorithm stops, the sets $I_1, I_2, \ldots, I_k$ are a partition of $\{1, 2, \ldots, r\}$. Now we prove the lemma by contradiction. Assume that there exists some t such that $\sum_{j \in I_t} \lambda_j < \frac{1}{2}$.

We now claim that, for all i, $\sum_{j \in I_i} \lambda_j \leq 1$. Otherwise, suppose $\sum_{j \in I_{t'}} \lambda_j > 1$ for some t'. Note that $\lambda_j \leq 1$, for every j, and thus there are at least two eigenvalues

assigned to $I_{t'}$. Denote the last element within $I_{t'}$ by s'. It follows that $\sum_{j \in I_{t'}} \lambda_j - \lambda_{s'} = \sum_{j \in I_{t'} \setminus \{s'\}} \lambda_j \leq \sum_{j \in I_t} \lambda_j$ since, otherwise, we would have not assigned $\lambda_{s'}$ to $I_{t'}$ in the algorithm. However, since $\sum_{j \in I_t} \lambda_j < \frac{1}{2}$, we must have $\sum_{j \in I_{t'}} \lambda_j - \lambda_{s'} = \sum_{j \in I_{t'} \setminus \{s'\}} \lambda_j < \frac{1}{2}$. Thus, $\lambda_{s'} > \sum_{j \in I_{t'}} \lambda_j - \frac{1}{2} > \frac{1}{2}$. This is impossible since $\lambda_{s'}$ is the last eigenvalue assigned to $I_{t'}$, which implies $\lambda_{s'} \leq \lambda_j$ for every $j \in I_{t'}$, and we have already proved that there must exist an l such that $s' \neq l \in I_{t'}$ and $\lambda_l \leq \sum_{j \in I_{t'} \setminus \{s'\}} \lambda_j < \frac{1}{2}$. Therefore, $\sum_{j \in I_i} \lambda_j \leq 1$ for all i, and in particular $\sum_{j \in I_t} \lambda_j < \frac{1}{2}$. It follows that $\sum_{i=1}^{k} \sum_{j \in I_i} \lambda_j < k$. However, we know that since $I_1, I_2, \ldots, I_k$ is a partition of the index set $\{1, 2, \ldots, r\}$, $\sum_{i=1}^{k} \sum_{j \in I_i} \lambda_j = \sum_{j=1}^{r} \lambda_j = k$. This results in a contradiction. Therefore, such a t does not exist, and the proof is completed. □

Notice that the running time of the partitioning algorithm is bounded by $O(r \cdot k)$.[2]

2.2. Randomized second rounding. Let us now assume that we have found $I_1, I_2, \ldots, I_k$. Then our next randomized rounding procedure works as follows:

1. Generate an r-dimensional random vector ϕ such that each entry of ϕ takes value, independently, -1 or 1 with probability $\frac{1}{2}$ each way.
2. For $i = 1, 2, \ldots, k$, let

$$x_i = \frac{\sum_{j \in I_i} \phi_j \sqrt{\lambda_j} \cdot v_j}{\sqrt{\sum_{j \in I_i} \lambda_j}}.$$

The following lemmas show that $x_1, x_2, \ldots, x_k$ form a feasible solution for the original problem. In other words, they are k orthogonal unit vectors.

LEMMA 3. *For* $i = 1, 2, \ldots, k$, $\|x_i\| = 1$.

Proof. Recall that $\langle v_l, v_j \rangle = 0$ for any $l \neq j$ and $\|v_j\| = 1$. By definition,

$$\|x_i\|^2 = \left\langle \frac{\sum_{j \in I_i} \phi_j \sqrt{\lambda_j} v_j}{\sqrt{\sum_{j \in I_i} \lambda_j}}, \frac{\sum_{j \in I_i} \phi_j \sqrt{\lambda_j} v_j}{\sqrt{\sum_{j \in I_i} \lambda_j}} \right\rangle$$

$$= \frac{1}{\sum_{j \in I_i} \lambda_j} \sum_{j \in I_i} \langle \phi_j \sqrt{\lambda_j} v_j, \phi_j \sqrt{\lambda_j} v_j \rangle$$

$$= \frac{1}{\sum_{j \in I_i} \lambda_j} \sum_{j \in I_i} (\phi_j)^2 \lambda_j \|v_j\|^2$$

$$= 1. \quad \square$$

LEMMA 4. *If* $s \neq t$, *then* $\langle x_s, x_t \rangle = 0$.

Proof. Since for any $j \in I_s$ and $l \in I_t$, $\langle v_j, v_l \rangle = 0$,

$$\langle x_s, x_t \rangle = \left\langle \frac{\sum_{j \in I_s} \phi_j \sqrt{\lambda_j} v_j}{\sqrt{\sum_{j \in I_s} \lambda_j}}, \frac{\sum_{j \in I_t} \phi_j \sqrt{\lambda_j} v_j}{\sqrt{\sum_{j \in I_t} \lambda_j}} \right\rangle$$

$$= \frac{1}{\sqrt{\sum_{j \in I_s} \lambda_j \cdot \sum_{j \in I_t} \lambda_j}} \left\langle \sum_{j \in I_s} \phi_j \sqrt{\lambda_j} v_j, \sum_{j \in I_t} \phi_j \sqrt{\lambda_j} v_j \right\rangle$$

$$= 0. \quad \square$$

---

[2] An alternative way of partitioning the eigenvalues is the following: first, put the eigenvalues that are greater than or equal to $1/2$ into distinct subsets. If the number of such eigenvalues, say l, is not less than k, then we are done. Otherwise, arbitrarily put the remaining eigenvalues into $k - l$ subsets such that the sum of eigenvalues in each subset is greater than or equal to $1/2$. This method was suggested by an anonymous referee of a preliminary version of this paper.

Now we establish a bound on the performance of our algorithm. First, let us introduce Bernstein's theorem (see, for example, [26]), which is a form of the Chernoff bound.

LEMMA 5. *Let ϕ be a random vector whose entries are independent and either 1 or -1 with probability $\frac{1}{2}$ each way. Then, for any vector e and $\beta > 0$,*

$$\text{prob}\{\langle \phi, e\rangle^2 > \beta\|e\|^2\} < 2 \cdot \exp\left(-\frac{\beta}{2}\right).$$

Let $C_{ip} = \sum_{j \in I_i} \lambda_j \langle p, v_j\rangle^2$. Then we have

LEMMA 6. *For each $i = 1, 2, \ldots, k$ and each $p \in P$, we have*

$$\text{prob}\{\langle p, x_i\rangle^2 > 12\log(n) \cdot C_{ip}\} < \frac{2}{n^3}.$$

Proof. Given i and p, define an $|I_i|$-dimensional vector e such that its entries are $\sqrt{\lambda_j}\langle p, v_j\rangle$, $j \in I_i$, respectively. Furthermore, we define the $|I_i|$-dimensional vector $\phi|_{I_i}$ whose entries are those of ϕ with indices in I_i. First notice that

$$\|e\|^2 = \sum_{j \in I_i}(\sqrt{\lambda_j}\langle p, v_j\rangle)^2 = \sum_{j \in I_i}\lambda_j \cdot \langle p, v_j\rangle^2 = C_{ip}.$$

On the other hand, since $\sum_{j \in I_i}\lambda_j \geq \frac{1}{2}$,

$$\langle p, x_i\rangle^2 = \left\langle p, \frac{\sum_{j \in I_i}\sqrt{\lambda_j}v_j\phi_j}{\sqrt{\sum_{j \in I_i}\lambda_j}}\right\rangle^2$$

$$\leq 2\left\langle p, \sum_{j \in I_i}\sqrt{\lambda_j}v_j\phi_j\right\rangle^2$$

$$= 2\left(\sum_{j \in I_i}\sqrt{\lambda_j}\phi_j\langle p, v_j\rangle\right)^2$$

$$= 2\left\langle \phi|_{I_i}, e\right\rangle^2.$$

Thus

$$\text{prob}\{\langle p, x_i\rangle^2 > 12\log(n)C_{ip}\} \leq \text{prob}\{\langle \phi|_{I_i}, e\rangle^2 > 6\log(n)\|e\|^2\}.$$

Therefore, the conclusion of the lemma follows by using Lemma 5 and by letting $\beta = 6\log(n)$. □

THEOREM 1. *We can compute in polynomial time a $(d-k)$-flat such that, with probability at least $1 - \frac{2}{n}$, the distance between any point $p \in P$ and F is at most $\sqrt{12\log(n)} \cdot R_k(P)$.*

Proof. For given $i = 1, 2, \ldots, k$ and $p \in P$, consider the event

$$B_{ip} = \{\phi|\langle p, x_i\rangle^2 > 12\log(n) \cdot C_{ip}\}$$

and $B = \bigcup_{i,p} B_{ip}$. The probability that the event B happens is bounded by

$$\sum_{i,p}\text{prob}\{\langle p, x_i\rangle^2 > 12\log(n) \cdot C_{ip}\} < \frac{2kn}{n^3} \leq \frac{2}{n}.$$

If B does not happen, then for any i and p,

$$\langle p, x_i \rangle^2 \le 12 \log(n) \cdot C_{ip}.$$

Therefore, for each $p \in P$,

$$\sum_{i=1}^{k} \langle p, x_i \rangle^2 \le 12 \log(n) \sum_{i=1}^{k} C_{ip} \le 12 \log(n) \cdot R_k(P)^2.$$

The last inequality follows from Lemma 1. This completes the proof by taking F as the subspace which is orthogonal to the vectors $x_1, x_2, \ldots, x_k$. $\square$

3. The inapproximability results. We start with formal definitions of the problems that will be used in the sequence of reductions from 3SAT to computing the outer k-radius $R_k(P)$ of a set P of points. Our starting point will be the classic 3SAT problem, in which we are given a 3CNF formula and we want to know if there is an assignment to its variables that simultaneously satisfies all its clauses. The next problem we consider is the resricted quadratic programming problem as defined by Brieden [14].

DEFINITION 1 (ζ-restricted quadratic programming). *We are given nonnegative integers $\lambda, \tau, \kappa,$ and σ and nonnegative rational numbers $c_{p,q,a,b}$ for $p \in [\lambda]$, $q \in [\tau]$, $a \in [\kappa]$, and $b \in [\sigma]$. (For a nonnegative integer n, $[n]$ denotes the set $\{1, 2, \ldots, n\}$.) Our goal is to maximize*

$$f(x) = \sum_{p,q,a,b} c_{p,q,a,b} x_{p,a} y_{q,b}$$

over the polytope $P \subseteq \Re^{\lambda\kappa + \tau\sigma}$ described by

$$\sum_{a \in [\kappa]} x_{p,a} = 1 \text{ for } p \in [\lambda],$$

$$\sum_{b \in [\sigma]} y_{q,b} = 1 \text{ for } q \in [\tau],$$

$$0 \le x_{p,a} \le 1 \text{ for } p \in [\lambda], a \in [\kappa],$$

$$0 \le y_{q,b} \le 1 \text{ for } q \in [\tau], b \in [\sigma].$$

We denote instances in which $\kappa, \sigma \le \zeta$ and $\lambda, \tau \le \Delta$ by ζ-restricted $QP[\Delta]$.

DEFINITION 2 (symmetric full-dimensional norm maximization). *We are given a string (n, m, A), where n and m are natural numbers and A is a rational $m \times n$ matrix. Our goal is to maximize*

$$f(x) = ||x||_2$$

over all vectors x that belong to the polytope $P = \{x| -1 \le Ax \le 1\}$. We denote an instance in which the number of rows of A is at most m and the number of columns is at most n by $NM[m, n]$.

We first prove an inapproximability result for the case $k = 1$ and later extend it to a large range of k using a simple reduction. The crux of the proof is the following lemma.

LEMMA 7. *There is a constant $c > 1$ such that for any sufficiently large integer parameter $t \ge 1$, there is a reduction T from 3SAT formulas of size m to computing R_1 for a point set of size $n = 2^{O(t2^{3t} \log m)}$ in $d = 2^{O(t \log m)}$ dimensions such that the following hold:*

1. If ψ is satisfiable, then $R_1(T(\psi)) \geq w$ for some w.
2. If ψ is unsatisfiable, then $R_1(T(\psi)) \leq w'$ for some w'.
3. $\frac{w}{w'} \geq c^t$.

This reduction, including the computation of w and w', runs in time $2^{O(t2^{3t} \log m)}$.

   Proof. The proof involves a sequence of three reductions.

   From 3SAT to quadratic programming. Bellare and Rogaway [8, section 4] give a reduction from 3SAT to quadratic programming via a two-prover protocol for 3SAT. We use their reduction, but in order to get the right parameters in the hardness of approximation result for quadratic programming, we need to replace the one-round two-prover protocol that they start off with by a different one that is described in Feige [16, section 2.2]. For completeness, we now describe this two-prover protocol for 3SAT.

   The two-prover protocol. Feige [16, Proposition 2.1.2] shows there exists a polynomial time reduction T from 3CNF formulas to 3CNF formulas such that each clause of $T(\psi)$ has exactly three literals (corresponding to three different variables) and each variable appears in exactly five clauses and furthermore the following hold:

1. If ψ is satisfiable, then $T(\psi)$ is satisfiable.
2. If ψ is not satisfiable, then $T(\psi)$ is at most $(1-\epsilon)$-satisfiable for some constant $0 < \epsilon < 1$. That is, any assignment satisfies at most a fraction $(1 - \varepsilon)$ of all clauses in $T(\psi)$.

   Without the requirement that each variable appears in exactly five clauses, such a reduction is known to be a consequence of the PCP theorem [3]. We now describe the steps taken by the verifier in the two-prover protocol.

   1. Convert ψ to $T(\psi)$.

   2. Choose t clauses uniformly at random (with replacement) from $T(\psi)$. Ask prover P_1 for an assignment to the variables in each clause chosen.

   3. From each chosen clause, choose one of the three variables in that clause uniformly at random. We get t *distinguished* variables, possibly with repetitions. Ask the prover P_2 for an assignment to each of these t variables.

   4. Accept if, for each chosen clause, it is satisfied by the assignment received from prover P_1 and the assignments made by the two provers to the distinguished variable from the clause are consistent. (Acceptance means that the verifier declares ψ to be satisfiable.) For example, suppose $t = 2$ and the verifier chose the clauses $(\neg x \vee y \vee z)$ and $(\neg y \vee z \vee w)$ and chose x and w as the respective distingushed variables. Suppose that P_1 returned the values x_1, y_1, and z_1 for the variables in the first clause and the values y_2, z_2, and w_2 for the variables in the second clause. Suppose that P_2 returned values x_3 and w_3 for the distinguished variables. Then the verifier accepts if both $(\neg x_1 \vee y_1 \vee z_1)$ and $(\neg y_2 \vee z_2 \vee w_2)$ evaluate to true, $x_1 = x_3$, and $w_2 = w_3$.

   The prover P_1 is any function that on seeing ψ and the identity of the t clauses in $T(\psi)$ returns $3t$ bits that the verifier interprets as an assignment to the $3t$ variables in these clauses. Similarly, the prover P_2 is any function that on seeing ψ and the identity of the t distinguished variables returns t bits that the verifier interprets as an assignment to the distinguished variables.

   If ψ is satisfiable, so is $T(\psi)$, and there exist provers (functions) that will cause the verifier to accept on every outcome of the random choices. This can be seen by picking a satisfying assignment to $T(\psi)$ and defining the two provers so that they answer according to this assignment. If ψ is unsatisfiable, what is the maximum probability, over all choices of provers (functions) P_1 and P_2, that the verifier accepts ψ? By Raz's parallel repetition theorem [28], this *error probability* is bounded above by s^t for some $s < 1$ (where the s depends on ε). We refer the reader to Feige [16,

section 2.2] for a discussion of this and to the paper by Håstad [24] for more details on the use of two-prover protocols in inapproximability results. Also, note that in this protocol the questions to the two provers are at most $O(t \log m)$ bits long, where m is the input size, since $O(\log m)$ bits suffice to identify a clause or variable in $T(\psi)$. The answers from the two-provers P_1 and P_2 are $3t$ and t bits long.

We now plug this two-prover protocol into the reduction of Bellare and Rogaway [8, section 4] from SAT to restricted quadratic programming via two-prover protocols. Their description assumes for simplicity that the question and answer lengths of the two-prover protocol are the same, but their reduction works even if these sizes are different. Using the fact that the answer length is at most $3t$, we obtain the following.[3]

LEMMA 8 (Bellare and Rogaway [8]). *There is a constant $f > 1$ such that, for any sufficiently large integer $t \geq 1$, there is a reduction T_1 that maps $3CNF$ formulas of size m to 2^{3t}-restricted $QP[2^{O(t \log m)}]$ such that the following hold:*

1. *If ψ is satisfiable, then $OPT(T_1(\psi)) = w_1$ for some w_1.*
2. *If ψ is unsatisfiable, then $OPT(T_1(\psi)) \leq w_2$ for some w_2.*
3. *$\frac{w_1}{w_2} \geq f^t$.*

Moreover, this reduction, including the computation of w_1 and w_2, runs in time $2^{O(t \log m)}$.

From quadratic programming to norm maximization. Brieden [14, Theorem 3.4] describes a set of interesting reductions that converts an instance of quadratic programming to an instance of the norm maximization problem. Using this reduction, we obtain the following.

LEMMA 9 (Brieden [14]). *For any $\lambda > 0$, there is a reduction T_2 from restricted quadratic programming to symmetric full-dimensional norm maximization that maps 2^{3t}-restricted $QP[2^{O(t \log m)}]$ into $NM[2^{O(t2^{3t} \log m)}, 2^{O(t \log m)}]$ with the following property: for any input L of QP to T_2,*

$$\frac{OPT(L)}{(1+\lambda)} \leq OPT(T_2(L)) \leq (1+\lambda)OPT(L).$$

Moreover, the reduction T_2 runs in time $2^{O(t2^{3t} \log m)}$.

From norm maximization to width computation. The reduction from symmetric full-dimensional norm maximization to width computation is simple [19] and is in fact used by Brieden [14]. Let $a_i \in \Re^n$ be the vector that corresponds to the ith row of matrix A which is input to the norm-maximation problem for $1 \leq i \leq m$. Thus the norm-maximization problem is

(3)
$$\gamma := \quad \text{Maximize} \quad ||x||_2$$
$$\text{Subject to} \quad \langle a_i, x \rangle^2 \leq 1 \text{ for } 1 \leq i \leq m.$$

The reduction T_3 simply constructs a set B of points by adding, for each $1 \leq i \leq m$, the points a_i and $-a_i$ to B. Since B is a symmetric point set, $R_1(B)^2$ is given by the program

(4)
$$\text{Minimize} \quad \alpha$$
$$\text{Subject to} \quad \langle a_i, x \rangle^2 \leq \alpha \text{ for } 1 \leq i \leq m,$$
$$||x||^2 = 1.$$

---

[3] In Bellare's and Rogaway's reduction, the connection between the parameters of the two-prover protocol for 3SAT and the parameters of the resulting ζ-restricted $QP[\Delta]$ instance is as follows: ζ is exponential in the answer length, Δ is exponential in the question length, and the "gap" f^t in Lemma 8 is the reciprocal of the error probability of the protocol.

It is easy to verify that $\gamma = 1/R_1(B)$.

The reduction T claimed in Lemma 7 is obtained by composing the reductions T_1, T_2, and T_3. In particular, choose λ such that $(1 + \lambda)^2 < f$, and let

$$c = \frac{1}{(1 + \lambda)^2} f > 1.$$

It can now be checked that Lemma 7 holds with this choice of c.   □

THEOREM 2.
  1. *There exists a constant $\delta > 0$ such that the following holds: there is no quasi-polynomial time algorithm that approximates $R_1(P)$ within $(\log n)^\delta$ unless NP $\subseteq \tilde{P}$.*
  2. *Fix any constant $b \geq 1$. Then there is no quasi-polynomial time algorithm that approximates $R_1(P)$ within $(\log d)^b$ unless NP $\subseteq \tilde{P}$.*

Proof. To prove part 1, we apply the reduction of Lemma 7 with $t = \log \log m$ to obtain an instance of computing R_1 for a set of $n = 2^{O(t2^{3t} \log m)}$ in $d = 2^{O(t \log m)}$ dimensions. Choose $\delta' < \frac{\log c}{5}$. Then

$$\frac{c^t}{(t2^{3t})^{\delta'}} \geq \frac{c^t}{(2^{4t})^{\delta'}} \geq \left(\frac{c}{2^{4\delta'}}\right)^t > (2^{\delta'})^t \geq (\log m)^{\delta'}.$$

Thus $c^t > (t2^{3t} \log m)^{\delta'}$. Since $n = 2^{O(t2^{3t} \log m)}$, we can choose $\delta < \delta'$ such that, for n large enough,

$$c^t > (\log n)^\delta.$$

To prove part 2, we apply the reduction of Lemma 7 with $t = \frac{2p \log \log m}{\log c}$ for some sufficiently large constant p. Then,

$$\frac{c^t}{t^p} \geq \frac{2^{2p \log \log m}}{t^p} \geq 2^{p \log \log m} \frac{2^{p \log \log m}}{t^p} = 2^{p \log \log m} \left(\frac{\log m}{t}\right)^p > 2^{p \log \log m}$$

since $\log m > t$ for sufficiently large m.

Thus, $c^t > t^p 2^{p \log \log m} = (t \log m)^p$. Since the dimension d is $2^{O(t \log m)}$, it follows that for every constant $b \geq 1$, we can choose p large enough such that

$$c^t > (\log d)^b.$$

Observe that the reduction runs in quasi-polynomial time for our choice of t in both cases and hence the theorem follows.   □

We now give the easy reduction from width to the outer k-radius that proves the main result of this section.

THEOREM 3.
  1. *There exists a constant $\delta > 0$ such that the following holds for any $0 < \varepsilon < 1$: there is no quasi-polynomial time algorithm that approximates $R_k(P)$ within $(\log n)^\delta$ for all k such that $k \leq d - d^\varepsilon$ unless NP $\subseteq \tilde{P}$.*
  2. *Fix any $\varepsilon > 0$. Fix any constant $c \geq 1$. Then there is no quasi-polynomial time algorithm that approximates $R_k(P)$ within $(\log d)^c$ for all k such that $k \leq d - d^\varepsilon$ unless NP $\subseteq \tilde{P}$.*

Proof. Let P be a set of n points in $\Re^d$. We map P to a set P' of n points in $\Re^{d+k-1}$ using the function that takes a point $(x_1, \ldots, x_d) \in \Re^d$ to the point $(x_1, \ldots, x_d, 0, \ldots, 0)$. It is easily checked that $R_1(P) = R_k(P')$. Theorem 3 follows

from this reduction and some simple calculations: observe that the reduction runs in polynomial time even if we set k to be $d^{1/\varepsilon} - d + 1$. With this choice, the target dimension $d' := d + k - 1$ equals $d^{1/\varepsilon}$. Thus $k = d^{1/\varepsilon} - d + 1 \geq d' - d'^\varepsilon$. Theorem 3(1) now follows by applying Theorem 2(1). For part (2), we apply Theorem 2(2) with $b = 2c$. Since

$$(\log d)^{2c} \geq (\varepsilon \log d')^{2c} = (\varepsilon^2 \log d')^c (\log d')^c \geq (\log d')^c$$

for sufficiently large d', Theorem 3(2) also follows.   □

4. Conclusions. Finding efficient rounding methods for SDP relaxation plays a key role in constructing better approximation algorithms for various hard optimization problems. All of them developed to date are randomized in nature. Therefore, the mixed deterministic and randomized rounding procedure developed in this paper may have its own independent value. We expect to see more applications of the procedure in approximating various computational geometry and space embedding problems.

REFERENCES

[1] D. ACHLIOPTAS AND F. MCSHERRY, *Fast computation of low rank matrix approximations*, in Proceedings of the ACM Symposium on the Theory of Computing, 2001.

[2] P. K. AGARWAL, S. HAR-PELED, AND K. R. VARADARAJAN, *Approximating extent measures of points*, J. ACM, 51 (2004), pp. 606–635.

[3] S. ARORA, C. LUND, R. MOTWANI, M. SUDAN, AND M. SZEGEDY, *Proof verification and hardness of approximation problems*, J. ACM, 45 (1998), pp. 501–555.

[4] S. ARORA AND S. SAFRA, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM, 45 (1998), pp. 70–122.

[5] Y. AZAR, A. FIAT, A. KARLIN, F. MCSHERRY, AND J. SAIA, *Spectral analysis of data*, in Proceedings of the ACM Symposium on the Theory of Computing, 2001.

[6] G. BAREQUET AND S. HAR-PELED, *Efficiently approximating the minimum-volume bounding box of a point set in three dimensions*, J. Algorithms, 38 (2001), pp. 91–109.

[7] M. BĂDOIU, S. HAR-PELED, AND P. INDYK, *Approximate clustering via core-sets*, in Proceedings of the ACM Symposium on the Theory of Computing, 2002.

[8] M. BELLARE AND P. ROGAWAY, *The complexity of approximating a nonlinear program*, Math. Program. B, 69 (1995), pp. 429–441.

[9] D. BERTSIMAS AND Y. YE, *Semidefinite Relaxations, Multivariate Normal Distributions, and Order Statistics*, Handbook Combin. Optim. 3, D.-Z. Du and P. M. Pardalos, eds., Kluwer Academic Publishers, Norwell, MA, 1998, pp. 1–19.

[10] H. L. BODLAENDER, P. GRITZMANN, V. KLEE, AND J. VAN LEEUWEN, *The computational complexity of norm maximization*, Combinatorica, 10 (1990), pp. 203–225.

[11] A. BRIEDEN, P. GRITZMANN, AND V. KLEE, *Inapproximability of some geometric and quadratic optimization problems*, in Approximation and Complexity in Numerical Optimization: Continuous and Discrete Problems, P. M. Pardalos, ed., Kluwer Academic Publishers, Norwell, MA, 2000, pp. 96–115.

[12] A. BRIEDEN, P. GRITZMANN, R. KANNAN, V. KLEE, L. LOVASZ, AND M. SIMONOVITS, *Deterministic and randomized polynomial-time approximation of radii*, Mathematika, 48 (2001), pp. 63–105.

[13] A. BRIEDEN, P. GRITZMANN, R. KANNAN, V. KLEE, L. LOVASZ, AND M. SIMONOVITS, *Approximation of diameters: Randomization doesn't help*, in Proceedings of the IEEE Symposium on the Foundations of Computer Science, 1998, pp. 244–251.

[14] A. BRIEDEN, *Geometric optimization problems likely not contained in APX*, Discrete Comput. Geom., 28 (2002), pp. 201–209.

[15] U. FAIGLE, W. KERN, AND M. STRENG, *Note on the computational complexity of j-radii of polytopes in R^n*, Math. Program., 73 (1996), pp. 1–5.

[16] U. FEIGE, *A threshold of $\ln n$ for approximating set cover*, J. ACM, 45 (1998), pp. 634–652.

[17] A. Frieze, R. Kannan, and S. Vempala, *Fast Monte-Carlo algorithms for finding low rank approximations*, J. ACM, 51 (2004), pp. 1025–1041.

[18] M. X. Goemans and D. P. Williamson, *Improved approximation algorithms for maximum cut and satisfiability problems using semi-definite programming*, J. ACM, 42 (1995), pp. 1115–1145.

[19] P. Gritzmann and V. Klee, *Inner and outer j-radii of convex bodies in finite-dimensional normed spaces*, Discrete Comput. Geom., 7 (1992), pp. 255–280.

[20] P. Gritzmann and V. Klee, *Computational complexity of inner and outer j-radii of polytopes in finite-dimensional normed spaces*, Math. Program., 59 (1993), pp. 162–213.

[21] P. Gritzmann and V. Klee, *On the complexity of some basic problems in computational convexity:* I. *Containment problems*, Discrete Math., 136 (1994), pp. 129–174.

[22] S. Har-Peled and K. Varadarajan, *Projective clustering in high dimensions using core-sets*, in Proceedings of the 18th Annual Symposium on Computational Geometry, ACM Press, 2002, pp. 312–318.

[23] S. Har-Peled and K. Varadarajan, *High-dimensional shape fitting in linear time*, Discrete Comput. Geom., 32 (2004), pp. 269–288.

[24] J. Håstad, *Some optimal inapproximability results*, J. ACM, 48 (2001), pp. 798–859.

[25] N. Megiddo, *On the complexity of some geometric problems in unbounded dimension*, J. Symbolic Comput., 10 (1990), pp. 327–334.

[26] A. Nemirovski, C. Roos, and T. Terlaky, *On maximization of quadratic forms over intersection of ellipsoids with common center*, Math. Program., 86 (1999), pp. 463–473.

[27] Yu. Nesterov, *Global quadratic optimization via conic relaxation*, in Handbook of Semidefinite Programming Theory, Algorithms, and Applications, H. Wolkowicz, R. Saigal, and L. Vandenberghe, eds., Kluwer Academic Publishers, Norwell, MA, 2000.

[28] R. Raz, *A parallel repetition theorem*, SIAM J. Comput., 27 (1998), pp. 763–803.