

Analysis of Incomplete Data and an Intrinsic-Dimension Helly Theorem

Jie Gao · Michael Langberg · Leonard J. Schulman

Received: 12 April 2007 / Revised: 12 June 2008 / Accepted: 7 July 2008 /

Published online: 23 September 2008

© Springer Science+Business Media, LLC 2008

Abstract The analysis of incomplete data is a long-standing challenge in practical statistics. When, as is typical, data objects are represented by points in \mathbb{R}^d , incomplete data objects correspond to affine subspaces (lines or Δ -flats). With this motivation we study the problem of finding the *minimum intersection radius* $r(\mathcal{L})$ of a set of lines or Δ -flats \mathcal{L} : the least r such that there is a ball of radius r intersecting every flat in \mathcal{L} . Known algorithms for finding the minimum enclosing ball for a point set (or clustering by several balls) do not easily extend to higher-dimensional flats, primarily because “distances” between flats do not satisfy the triangle inequality. In this paper we show how to restore geometry (i.e., a substitute for the triangle inequality) to the problem, through a new analog of Helly’s theorem. This “intrinsic-dimension” Helly theorem states: for any family \mathcal{L} of Δ -dimensional convex sets in a Hilbert space, there exist $\Delta + 2$ sets $\mathcal{L}' \subseteq \mathcal{L}$ such that $r(\mathcal{L}) \leq 2r(\mathcal{L}')$. Based upon this we present an algorithm that computes a $(1 + \varepsilon)$ -core set $\mathcal{L}' \subseteq \mathcal{L}$, $|\mathcal{L}'| = O(\Delta^4/\varepsilon)$, such that the

Communicated by Pankaj Agarwal.

An extended abstract appeared in ACM–SIAM Symposium on Discrete Algorithms, 2006.

Work was done when J. Gao was with Center for the Mathematics of Information, California Institute of Technology.

Work was done when M. Langberg was a postdoctoral scholar at the California Institute of Technology. Research supported in part by NSF grant CCF-0346991.

Research of L.J. Schulman supported in part by an NSF ITR and the Okawa Foundation.

J. Gao

Department of Computer Science, Stony Brook University, Stony Brook, NY, 11794, USA
e-mail: jgao@cs.sunysb.edu

M. Langberg (✉)

Computer Science Division, The Open University of Israel, Raanana 43107, Israel
e-mail: mikel@openu.ac.il

L.J. Schulman

Department of Computer Science, California Institute of Technology, Pasadena, CA 91125, USA
e-mail: schulman@caltech.edu

ball centered at a point c with radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects every element of \mathcal{L} . The running time of the algorithm is $O(n^{\Delta+1}d \text{ poly}(\Delta/\varepsilon))$. For the case of lines or line segments ($\Delta = 1$), the (expected) running time of the algorithm can be improved to $O(nd \text{ poly}(1/\varepsilon))$. We note that the size of the core set depends only on the dimension of the input objects and is independent of the input size n and the dimension d of the ambient space.

Keywords Clustering · k -center · Core set · Incomplete data · Helly theorem · Approximation · Inference

1 Introduction

One of the great challenges in computational theory is the extraction of patterns from massive and high-dimensional data sets. A common difficulty associated with such data sets is that entries are incomplete—a few questions are left blank on a questionnaire; weather records for a region omit the figures for one weather station for a short period because of a malfunction; stock exchange data is absent for one stock on one day because of a trading suspension; and so forth. How should we process the partial data? Statisticians approach this problem in a variety of ways: deleting incomplete entries; filling in incomplete entries based on the most similar complete entry (“hot deck imputation”); filling in incomplete entries with the sample mean (“mean substitution”); or using a learning algorithm or criterion (EM, max likelihood) to infer a missing entry [18]. All of these are attempts to cope concurrently with two difficulties: (1) The statistical relationship between the present and missing data is usually not known. This precludes a universal answer to which approach is most statistically sound. (2) There is a combinatorial explosion inherent in trying out all candidate assignments to the missing values. The present paper offers a new approach to the problem of incomplete data: an approach rooted in the geometry of the data set.

From the computational point of view, a data item with d representative features is typically represented by a point in \mathbb{R}^d , each dimension corresponding to a feature; frequently one obtains good results by approximating the similarity of two items by their Euclidean distance, after choosing a good scaling of the axes. The most elementary form of data analysis for such a data set is to find the smallest ball that approximates the data set, whether in terms of the sum of distances to the center of the ball, maximum distance to the center, etc. An immediate generalization is to find a small number of balls (a “clustering” of the data) which between them cover the points (different interpretations of “cover” lead to the well-known k -median problem, k -center problem, etc.). There is copious work on these problems in the machine learning and algorithms literature.

A data item that is lacking information about one or more features corresponds to a line or a flat in \mathbb{R}^d , whose dimension is the number of missing features. There is no difficulty in assigning a distance between two such flats; it is simply the distance between the nearest points on the flats. So we can seek to cluster the flats so as to minimize some objective function. Right away there is a major difficulty: “distances” between flats do not satisfy the triangle inequality. The problem is not that the triangle inequality is slightly violated, but that *no* relaxation of it holds. No matter how far

apart lines a and c are, there is always a line b that intersects both. This problem defeats many existing algorithmic approaches for “clustering”-type tasks, and for good reason—the geometry seems, in a genuine sense, to be absent. What does it mean to cluster lines, if “ a resembles b ” and “ b resembles c ” imply nothing at all about a resembling c ?

In this paper we initiate work on data analysis for convex sets of low dimension inside an ambient space of possibly high dimension. Specifically, we assume that each input object is a convex subset of a flat of dimension Δ within a Hilbert space. For the existence theorems, this space may be infinite-dimensional, while for the algorithmic statements, we take it to be \mathbb{R}^d for a value of d that we shall consider to be much higher than Δ . Our measure of the similarity amongst a collection of convex sets \mathcal{L} is the *minimum intersection radius* $r(\mathcal{L})$, the least r such that there is a ball of radius r intersecting every flat in \mathcal{L} . The center of the optimum ball (selected arbitrarily in the degenerate case that it is not unique) is termed the *minimum intersection center* and denoted $c(\mathcal{L})$. Intuitively this center is the best explanation of the incomplete input data; the minimum intersection radius captures how well this center fits the data, in the sense that every incomplete data item can be *completed* to a point within that distance of the center. Significantly, therefore, our model makes a functional prediction for reconstructing missing data: among all points of the flat, use that which is closest to $c(\mathcal{L})$. Thus in addition to the role of our model in learning aggregate properties of the data set, it also provides an inference mechanism about the missing features of individual records.

Our approach can also be described under the shape fitting framework [5], where one asks for a shape, a point in our case, that best fits the input set, lines or flats, under some criterion.

The core of our contribution is to show a way to restore geometry to the problem of analyzing incomplete data, in spite of the failure of the triangle inequality. This restoration goes through a variant of Helly’s theorem. Suppose that we blow up each line or flat ℓ to a cylinder or a slab that encloses all the points within distance r from ℓ . Helly’s theorem says that if every $d + 1$ of these have a common intersection, then all of them have a common intersection. In other words, Helly’s theorem restores geometry because if every subcollection of $d + 1$ out of the n “data flats” are within distance r of some “explanation point”, then all of the n lines are within distance r of an explanation point.

As it stands, however, this chain of reasoning is too weak. The dimension of the ambient space, d , is typically of the order of hundreds or thousands, much larger than the maximum dimension Δ of the individual data items. We redress this gap by developing a version of Helly’s theorem that takes into account the low dimension of the sets involved. Beginning with the case of lines ($\Delta = 1$), we show that if every 3 of the n “data lines” are within distance r of some “explanation point,” then all of the n lines are within distance $2r$ of some explanation point. Notice that we are now free of the “extrinsic” dimension of the ambient space and depend only on the intrinsic dimensionality, $\Delta = 1$, of the data sets. This result can be extended to any Δ -dimensional convex objects \mathcal{L} , for any $0 \leq \Delta \leq d$, as follows: if every subset of $\Delta + 2$ convex objects of dimension at most Δ in a Hilbert space are within distance r of some point, then all of the objects are within distance $2r$ of some point. This result

is optimal in the sense that there exist configurations in which any $\Delta + 1$ convex objects of dimension at most Δ in \mathbb{R}^d have a minimum intersection radius that is strictly smaller than $1/2$ of that of \mathcal{L} . We call this result the intrinsic-dimension Helly theorem:

Theorem 1.1 (Intrinsic-dimension Helly theorem) *For any n convex sets of dimension at most Δ in a Hilbert space, $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_n\}$, there exist $\Delta + 2$ sets $\mathcal{L}' \subset \mathcal{L}$ such that $r(\mathcal{L}) \leq 2r(\mathcal{L}')$.*

Note that when $r = 0$ (i.e., when the sets of \mathcal{L} intersect), Theorem 1.1 directly generalizes Helly's theorem (except that it is weaker by 1 in the case that the ambient space is of finite dimension d , and $\Delta = d$). The implications of the theorem to the analysis of incomplete data are immediate: given a collection \mathcal{L} of n convex sets of dimension at most Δ , a 2-approximation of the minimum intersection radius of \mathcal{L} results from enumerating all subsets \mathcal{L}' of \mathcal{L} of size $\Delta + 2$ and determining the largest $r(\mathcal{L}')$. Actually as will be seen later, the implicit $O(n^{\Delta+2}d \text{poly}(\Delta))$ -runtime can be replaced by an algorithm with expected running time $n(d \text{poly}(\Delta) + 2^{\text{poly}(\Delta)})$ and a suitable center for \mathcal{L} (not generally equal to $c(\mathcal{L}')$) is identified as part of the same process. To obtain such a running time, we assume throughout that each Δ -dimensional convex set can be represented by $\text{poly}(\Delta)$ linear constraints. Here, and throughout this work, each appearance of the term “poly(x)” refers to a function of value at most αx^β for some universal constants α and β .

Next, we provide a method to achieve an approximation ratio of $1 + \varepsilon$ (for any $\varepsilon > 0$) for the minimum intersection radius. A subset $\mathcal{L}' \subseteq \mathcal{L}$ is said to be an α -core set, with respect to $r(\mathcal{L})$, if the minimum intersection radius $r(\mathcal{L}')$ of \mathcal{L}' approximates $r(\mathcal{L})$ within a multiplicative factor of α . Theorem 1.1 says that when \mathcal{L} is a set of lines in d -dimensional Euclidian space, one can find a 2-core set \mathcal{L}' of size 3; and in general, if \mathcal{L} consists of Δ -dimensional convex sets, there exists a 2-core set \mathcal{L}' of size $\Delta + 2$. For general values of $\alpha = 1 + \varepsilon$ and for \mathcal{L} consisting of Δ -dimensional convex sets in \mathbb{R}^d , we show that for any $\varepsilon > 0$, there exists a $(1 + \varepsilon)$ -core set of size $O(\Delta^4/\varepsilon)$.

Such a core set can be found in time $O(n^{\Delta+1}d \text{poly}(\Delta/\varepsilon))$. For the case of lines ($\Delta = 1$), the running time of the algorithm can be improved to $O(nd \text{poly}(1/\varepsilon))$. Notice that the size of the $(1 + \varepsilon)$ -core set only depends on ε and Δ , and is independent of the total input n or the dimension of the ambient space. After a preliminary version of this paper was made public, Sarel Har-Peled [10] noticed that one may trade the size of the $(1 + \varepsilon)$ -core set with a better running time. Namely, it is possible to find a $(1 + \varepsilon)$ -core set of size $\Delta^{O(1/\varepsilon)}$ in expected running time $n(d\Delta^{O(1/\varepsilon)} + 2^{\Delta^{O(1/\varepsilon)}})$.

To the best of our knowledge, this is the first work to address core sets for collections \mathcal{L} that consist of Δ -dimensional convex sets. We summarize the core set result by the following two theorems:

Theorem 1.2 ($(1 + \varepsilon)$ -core set for line segments) *Let $\varepsilon > 0$ be sufficiently small. Let \mathcal{L} be a set of lines or line segments $\{\ell_1, \dots, \ell_n\}$ in \mathbb{R}^d . There exists a subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $O(1/\varepsilon)$ such that $r(\mathcal{L}')(1 + \varepsilon) \geq r(\mathcal{L})$. The set \mathcal{L}' and a center c , such that the ball centered at c of radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects all lines or line segments in \mathcal{L} , can be found in expected time $O(nd \text{poly}(1/\varepsilon))$.*

Theorem 1.3 ($(1 + \varepsilon)$ -Core sets) *Let $\varepsilon > 0$ be sufficiently small. Let \mathcal{L} be a set of convex sets of dimension $\leq \Delta$, $\mathcal{L} = \{\ell_1, \dots, \ell_n\}$, in \mathbb{R}^d . There exists a subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $O(\Delta^4/\varepsilon)$ such that $r(\mathcal{L}')(1 + \varepsilon) \geq r(\mathcal{L})$. The set \mathcal{L}' and a center c , such that the ball centered at c of radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects all sets in \mathcal{L} , can be found in time $O(n^{\Delta+1}d \text{ poly}(\Delta/\varepsilon))$.*

Note that the algorithm of Theorem 1.2 is randomized and the running time is stated in *expectation*, while the algorithm of Theorem 1.3 is deterministic. This follows from the use of ‘LP-type’ programming (e.g., [19]) in the preliminary stages of the algorithm of Theorem 1.2. LP-type programming is not used in the algorithm of Theorem 1.3 (as it will not improve the stated running time).

As described above, the main focus of this paper is the near-optimal representation of a set of incomplete data entries (Δ -flats) by a single ball of minimal radius. Naturally, this is only the first step toward a more comprehensive theory that should provide algorithms for *clustering* incomplete data entries by providing several balls of small radius, at least one of which intersects each of the Δ -flats. It is easy to see that minimizing this radius is NP-hard, from the NP-hardness of the k -center problem for points. Our work already implies an initial result in this area: using our core set method, there is a straightforward $\tilde{O}(n^k)$ -time algorithm to obtain a 2-approximate k -clustering of n lines. Using a more elaborate analysis, for $k = 2$ and 3, the running time above can be reduced to $\tilde{O}(n)$ [9] (here, \tilde{O} hides polylogarithmic factors in n and linear factors in d).

1.1 Related Work

Clustering and shape fitting problems on points have been actively studied in recent years. One of the powerful techniques is to devise a core set, i.e., a small subset of representative points S' of S such that the optimization problems on S' is a good approximation to the optimal solution on S [5]. Precisely, a subset S' is a $(1 + \varepsilon)$ -core set of S if $(1 + \varepsilon)\mu(S') \geq \mu(S)$, where μ is a monotonic measure function. Agarwal et al. provided a framework for computing a $(1 + \varepsilon)$ -core set for a set of points S in \mathbb{R}^d with respect to many measure functions that depend on the extent of the point set, such as diameter, width, radius of the minimum enclosing ball, and volume of the smallest enclosing box [4]. The basic idea is to find a subset of points of size $O(1/\varepsilon^{(d-1)/2})$ whose convex hull approximates the convex hull of S . For some of the problems such as the minimum enclosing ball or ellipsoid, there is an incremental algorithm that computes a $(1 + \varepsilon)$ -core set of size that depends only on ε [6, 7, 15, 17]. Thus one can apply brute-force algorithms on the small core set S' and obtain efficient approximation algorithms for the optimization problems on S . Indeed, many geometric optimization problems such as minimum enclosing ball, k -clustering, and various shape fitting problems can be solved efficiently by using a small core set [2, 7, 11–13, 15, 17]. However, to the best of our knowledge, no work has been done on devising a core set for lines or flats with respect to a natural quality measure.

The study of core sets for points can not be directly applied to core sets for lines or flats. For a set of lines or flats, there is no natural definition of “convex hull.” Our core set algorithms for lines or flats are more related with the incremental core set

algorithm for points S in \mathbb{R}^d with respect to the radius of the minimum enclosing ball [7], which is described as follows. The algorithm starts with S' being a pair of furthest away points and computes the minimum enclosing ball of S' . If all the points are included in the minimum enclosing ball enlarged by a factor of $(1 + \varepsilon)$, then S' is a core set. Otherwise, a point outside the enlarged ball is added to S' . It can be shown that for each step, the radius of the minimum enclosing ball of S' is increased by a factor of $1 + O(\varepsilon^2)$. After $O(1/\varepsilon^2)$ steps, the algorithm terminates. However, there is a major difficulty to apply this algorithm for a set of lines \mathcal{L} : there is a situation where by adding each extra line, the minimum intersection radius of the current subset \mathcal{L}' stays the same but the minimum intersection radius of \mathcal{L}' , $r(\mathcal{L}')$, is still far away from the real value $r(\mathcal{L})$. A substantial part of this paper is devoted to showing that a carefully selected set of two lines (or $\Delta + 1$, more generally, Δ -flats) can improve the minimum intersection radius substantially.

We also note that there has been work on “clustering points with lines” [1, 3, 12], where one finds a set of lines \mathcal{L} such that the set of cylinders with radius r and axis as the lines of \mathcal{L} covers all the input points S . The problem we study in this paper can be phrased as “clustering lines with a point.” There does not exist an obvious connection between these two problems as a natural duality does not exist.

1.2 Organization

The remainder of the paper is organized as follows. We start with a few preliminaries in Sect. 2. In Sect. 3 we present the proof of Theorem 1.1. In Sect. 4 we present the proof of Theorem 1.2. The proof of Theorem 1.3 is very similar to that of Theorem 1.2 and is sketched in Sect. 5. To ease our presentation, the proof of many technical claims appear in the Appendix.

2 Preliminaries, Definitions, and Notation

We denote by $B_r(c)$ the ball centered at a center c with radius r in \mathbb{R}^d . We denote by $d(\cdot, \cdot)$ the Euclidean distance function. The distance between two points p, q is also written as $|pq|$. A ball B in \mathbb{R}^d is said to “cover” a set ℓ if $B \cap \ell \neq \emptyset$. Throughout this work, we consider the approximation parameter $\varepsilon > 0$ to be sufficiently small (as to satisfy several inequalities in our proofs).

Definition 2.1 A Δ -flat in a Hilbert space is a Δ -dimensional affine subspace. The dimension of a convex set in a Hilbert space is the least dimension of any flat containing it.

Definition 2.2 The *minimum intersection ball* $B(\mathcal{L})$ of a collection of convex sets \mathcal{L} in a Hilbert space is defined to be (one of) the minimum radius balls that covers all the sets in \mathcal{L} . The center of the minimum intersection ball is called the *minimum intersection center*, denoted as $c(\mathcal{L})$. The radius of the minimum intersection ball is called the *minimum intersection radius*, denoted as $r(\mathcal{L})$. Namely, $B(\mathcal{L}) = B_{r(\mathcal{L})}(c(\mathcal{L}))$.

Definition 2.3 (($1 + \varepsilon$)-Core set) Let \mathcal{L} be a set of convex sets in a Hilbert space. A subset \mathcal{L}' of \mathcal{L} is said to be a $(1 + \varepsilon)$ -core set w.r.t. the minimum intersection radius of \mathcal{L} if $r(\mathcal{L}) \leq (1 + \varepsilon)r(\mathcal{L}')$. A subset \mathcal{L}' of \mathcal{L} is said to be a ‘strong’ $(1 + \varepsilon)$ -core set w.r.t. the minimum intersection radius of \mathcal{L} if $r(\mathcal{L}) \leq (1 + \varepsilon)r(\mathcal{L}')$ and the minimum intersecting center $c(\mathcal{L}')$ of \mathcal{L}' is an approximation to the minimum intersecting center of \mathcal{L} , namely, $B_{r(\mathcal{L}')(1+\varepsilon)}(c(\mathcal{L}'))$ intersect all elements of \mathcal{L} . All core sets constructed in this work will be strong core sets. We thus omit the adjective ‘strong’ in the discussions and statements throughout this work.

We begin by noting that the minimum intersection radius and center of \mathcal{L} can be found in polynomial time up to an absolute error δ using convex programming. The proof of the following lemma appears in Appendix A.

Lemma 2.4 Let \mathcal{L} be a set of n convex sets with dimension at most Δ in \mathbb{R}^d , where each convex set can be represented by $\text{poly}(\Delta)$ convex constraints. Then $c(\mathcal{L})$ and $r(\mathcal{L})$ can be computed to an absolute precision $\delta > 0$ in time $O(\sqrt{n}(d^3 + d^2n) \text{poly}(\Delta) \log(n/\delta))$.

3 Intrinsic-Dimension Helly Theorem

We now prove Theorem 1.1 stated in the Introduction and restated here for convenience:

Theorem 1.1. For any n convex sets of dimension at most Δ in a Hilbert space, $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_n\}$, there exist $\Delta + 2$ sets $\mathcal{L}' \subset \mathcal{L}$ such that $r(\mathcal{L}) \leq 2r(\mathcal{L}')$.

Proof For each $(\Delta + 2)$ -tuple $\mathbf{i} = \{i_1, \dots, i_{\Delta+2}\}$ in $\{1, \dots, n\}$, let $B_{\mathbf{i}}$ be a minimum intersection ball of the subset $\mathcal{L}_{\mathbf{i}} = \{\ell_{i_1}, \dots, \ell_{i_{\Delta+2}}\}$, let $c_{\mathbf{i}}$ be the center of $B_{\mathbf{i}}$, and let $r_{\mathbf{i}}$ be the radius of $B_{\mathbf{i}}$. Let the largest radius among the $r_{\mathbf{i}}$ ’s be r . Now we claim that we can find a ball of radius $2r$ that intersects all the sets in \mathcal{L} . Consider the set ℓ_1 and denote by I_j the set of points on ℓ_1 with distance no more than $2r$ from ℓ_j . That is, $I_j = \{p \in \ell_1 \mid d(p, \ell_j) \leq 2r\}$. I_j is convex, since I_j is the intersection of two convex objects, the set ℓ_1 and the set of points of distance $2r$ from ℓ_j . Notice that I_j is a convex set of dimension Δ in ℓ_1 .

Consider any $(\Delta + 1)$ -tuple $\mathbf{i}' = \{i_1, \dots, i_{\Delta+1}\}$ in $\{2, \dots, n\}$. Now we claim that the corresponding $\Delta + 1$ sets $\{I_{i_1}, \dots, I_{i_{\Delta+1}}\}$ have nonempty intersection. Let \mathbf{i} be the $(\Delta + 2)$ -tuple obtained from \mathbf{i}' by adding an additional index of value 1, namely $\mathbf{i} = \{1, i_1, \dots, i_{\Delta+1}\}$. By the above discussion we have that $r_{\mathbf{i}} \leq r$. Let $c_{\mathbf{i}}$ be the center of the minimum intersection ball $B_{\mathbf{i}}$ of $\mathcal{L}_{\mathbf{i}}$. Let c'_1 be the point on the set ℓ_1 that is closest to $c_{\mathbf{i}}$. As the point $c_{\mathbf{i}}$ is within distance r from all sets in $\mathcal{L}_{\mathbf{i}}$, we have that the point c'_1 is within distance $2r$ from all the sets in $\mathcal{L}_{\mathbf{i}}$. This implies that c'_1 is in I_j for all $j \in \mathbf{i}'$. See Fig. 1.

Since, for any $(\Delta + 1)$ -tuple $\mathbf{i}' = \{i_1, \dots, i_{\Delta+1}\}$, the corresponding $\Delta + 1$ convex sets I_j have nonempty intersection and these sets are embedded in the Δ -dimensional set ℓ_1 , by Helly’s Theorem [14] all the sets I_j , $2 \leq j \leq n$ have a nonempty intersection. Now let o be a point in $\bigcap_j I_j$, $2 \leq j \leq n$. The ball centered at o with radius $2r$ intersects all n sets of \mathcal{L} . □

Fig. 1 Proof of the reduction from the minimum intersection radius of n sets to that of $\Delta + 2$ sets





Fig. 2 A simplex Ω_2 of dimension 2 with three sets ℓ_1, ℓ_2, ℓ_3 as the 3 faces that have pairwise intersection but no common intersection



The case of $r = 0$ is closest in form to the original Helly theorem:

Corollary 3.1 For any n convex sets of dimension at most Δ in a Hilbert space, $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_n\}$, if every $\Delta + 2$ sets in \mathcal{L} intersect, then all sets in \mathcal{L} intersect.

The Helly number $\Delta + 2$ in Theorem 1.1 cannot be reduced as shown by the following construction.

Theorem 3.2 (Optimality) For any Δ , there exists a set of n convex sets $\mathcal{L} = \{\ell_1, \dots, \ell_n\}$ such that for every subset $\mathcal{L}' \subseteq \mathcal{L}$ of size less than $\Delta + 2$, it holds that $r(\mathcal{L}) > 2r(\mathcal{L}')$.

Proof For any Δ , consider the $(\Delta + 1)$ -dimensional simplex, namely the set

$$\Omega_{\Delta+1} = \left\{ (p_1, \dots, p_{\Delta+2}) \mid \sum_{i=1}^{\Delta+2} p_i = 1; \forall i \ p_i \in \mathbb{R}, p_i \geq 0 \right\}.$$

We show in Fig. 2 a case where $\Delta = 1$. Our sets ℓ_i will be subsets of $\Omega_{\Delta+1}$ of dimension Δ . $|\mathcal{L}| = \Delta + 2$. For $i = 1$ to $\Delta + 2$, define $\ell_i = \{(p_1, \dots, p_{\Delta+2}) \in \Omega_{\Delta+1} \mid p_i = 0\}$ (i.e., each ℓ_i is a Δ -dimensional facet of the simplex). It is not hard to verify that every subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $\Delta + 1$ has a nonempty intersection—if $\ell_i \notin \mathcal{L}'$, then the unit vector with ‘1’ in the i th coordinate is in \mathcal{L}' . Thus $r(\mathcal{L}') = 0$. However, the sets in \mathcal{L} do not have a common intersection. $r(\mathcal{L}) > 0$.

Since this construction lies in $\mathbb{R}^{\Delta+1}$, the only choice of parameters for which Theorem 1.1 is not tight is the obvious case of $\Delta = d$. □

Theorem 1.1 indicates a straightforward algorithm to find a 2-approximation to the minimum intersection radius by simply taking the maximum radius of all $(\Delta + 2)$ -tuples of sets that include ℓ_1 . In Appendix B, we show that one can improve the running time of this naive algorithm (and actually give an alternative proof

to Theorem 1.1) using the methodology of ‘LP-type’ programming (e.g., [19]) to $n(d \text{ poly}(\Delta) + 2^{\text{poly}(\Delta)})$.

Lemma 3.3 *Let \mathcal{L} be a set of n convex sets with dimension at most Δ in \mathbb{R}^d . Let $\ell \in \mathcal{L}$. The minimum radius ball $B_r(c)$ that covers the sets \mathcal{L} with center c on ℓ , and a set $\mathcal{L}' = \{\ell_1, \dots, \ell_{\Delta+2}\} \subseteq \mathcal{L}$ for which $r(\mathcal{L}) \leq 2r(\mathcal{L}')$ can be found in expected time $n(d \text{ poly}(\Delta) + 2^{\text{poly}(\Delta)})$.*

4 (1 + ε)-Core Set for Lines

The proof of Theorem 1.2 we present shortly strongly builds upon the notion of a $(1 + \varepsilon)$ -approximate intersection center, with $\varepsilon > 0$. In what follows we define $(1 + \varepsilon)$ -approximate intersection centers and state Theorem 4.4 which addresses a certain property of these centers. We then prove Theorem 1.2 based on Theorem 4.4. In this section we will assume that \mathcal{L} is a set of n lines. The general case in which \mathcal{L} also includes line segments is proven in a very similar manner and is given in detail in Appendix D.

Definition 4.1 A $(1 + \varepsilon)$ -approximate intersection ball of a set \mathcal{L} in \mathbb{R}^d is a ball of radius $(1 + \varepsilon)r(\mathcal{L})$ that intersects all sets in \mathcal{L} . The center of a $(1 + \varepsilon)$ -approximate intersection ball is called a $(1 + \varepsilon)$ -approximate intersection center. We denote by $C_\varepsilon(\mathcal{L})$ the set of $(1 + \varepsilon)$ -approximate intersection centers of \mathcal{L} .

Definition 4.2 Let ℓ be a line in \mathbb{R}^d . A cylinder of radius r with axis ℓ in \mathbb{R}^d is defined as the set of points in \mathbb{R}^d which are of distance at most r from ℓ .

Denote by $C_i(\varepsilon)$ the cylinder of radius $(1 + \varepsilon)r$ with ℓ_i as the axis, $\ell_i \in \mathcal{L}$, $r = r(\mathcal{L})$. Then the set of $(1 + \varepsilon)$ -approximate intersection centers is exactly the common intersection of all the cylinders $C_i(\varepsilon)$. Thus we have:

Observation 4.3 *For any set \mathcal{L} , $C_\varepsilon(\mathcal{L})$ is convex.*

Theorem 4.4 *$C_\varepsilon(\mathcal{L})$ is included in a cylinder of radius $25\sqrt{\varepsilon}r(\mathcal{L})$ with axis parallel to one of the lines $\ell_i \in \mathcal{L}$. Moreover this axis passes through $c(\mathcal{L})$.*

The proof of Theorem 4.4 is based on Lemmas 4.5 and 4.6. The proofs of the Lemmas and Theorem 4.4 are nontrivial and rather technical; they appear in Appendix C.

Lemma 4.5 *Suppose that c is a minimum intersection center of a set of lines \mathcal{L} and p is a $(1 + \varepsilon)$ -approximate intersection center with $|cp| \geq \alpha\sqrt{\varepsilon}r(\mathcal{L})$. Then there exists a line $\ell \in \mathcal{L}$ such that the angle between ℓ and the line segment cp is bounded (by above) by $\arcsin(\sqrt{2 + \varepsilon}/\alpha)$.*

Lemma 4.6 *$C_\varepsilon(\mathcal{L})$ does not include a 2-dimensional disk with radius $\frac{5}{2}\sqrt{\varepsilon}r(\mathcal{L})$.*

With Theorem 4.4, we now prove Theorem 1.2. The idea is similar with the construction of a $(1 + \varepsilon)$ -core set P' for a set of points P in \mathbb{R}^d such that the radius of the minimum enclosing ball of P is bounded by $(1 + \varepsilon)$ times that of P' [7]. The basic idea in [7] is to add a point not covered by the minimum enclosing ball of the current core set such that the minimum radius is increased substantially. However, a direct application of this idea does not work for the case of lines. One can find a scenario where adding a line cannot improve the minimum intersection radius. We will show that a careful selection of two lines can always increase the minimum intersection radius by a substantial factor.

Recall that in the section we assume that \mathcal{L} consists of lines (and not line segments). We thus rephrase Theorem 1.2 stated in the Introduction. As mentioned, extending our result to the case of line segments is analyzed in Appendix D.

Theorem 1.2 (Restated for Lines) *Let $\varepsilon > 0$ be sufficiently small. Let \mathcal{L} be a set of lines $\{\ell_1, \dots, \ell_n\}$ in \mathbb{R}^d . There exist a subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $O(1/\varepsilon)$ such that $r(\mathcal{L}')(1 + \varepsilon) \geq r(\mathcal{L})$. The set \mathcal{L}' and a center c , such that the ball centered at c of radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects all lines or line segments in \mathcal{L} , can be found in expected time $O(nd \text{ poly}(1/\varepsilon))$.*

Proof We first show how to find a core set of size $O(1/\varepsilon^2)$. We then use ideas of [16] to improve the size to $O(1/\varepsilon)$.

Step I: Find a core set of size $O(1/\varepsilon^2)$ Let $\varepsilon > 0$ be a small constant. In what follows we present an algorithm for finding \mathcal{L}' . Our algorithm is greedy and strongly builds upon Theorem 4.4. Roughly speaking, the main idea of our algorithm is as follows. We start out by picking a subset of \mathcal{L}' of size 3 according to Lemma 3.3. For these lines, it holds that $\alpha r(\mathcal{L}') \geq r(\mathcal{L})$ where $\alpha = 2$. The set \mathcal{L}' is a good starting point, but we still need to reduce the value α above to $(1 + \varepsilon)$. We do this in a series of steps. In each step, a line or two are added to \mathcal{L}' , and α reduces by a factor of $(1 - \varepsilon^2/200^2)$. Hence, after $O(1/\varepsilon^2)$ such steps we are in a situation in which $(1 + \varepsilon)r(\mathcal{L}') \geq r(\mathcal{L})$, and we are done. The second part of the theorem (regarding efficiency matters) will follow from the detailed description of the algorithm.

We will assume w.l.o.g. that we start with a subset \mathcal{L}' of size 3 for which $r(\mathcal{L}') \geq 2$. If this is not the case, we can map our instance \mathcal{L} to an equivalent *blown-up* instance for which the inequality specified above will hold (specifically, each line defined by a point $o \in \mathbb{R}^d$ and a direction $b \in \mathbb{R}^d$ will be mapped to a line in direction b passing through $c \cdot o$ for a sufficiently large blowup factor $c > 0$). Here we assume the (standard) computational model in which real numbers may be manipulated freely.

In what follows we focus on one iteration of the algorithm. Let $\delta = \varepsilon^3$. Let \mathcal{L}' be the subset defined by the algorithm so far. Using Lemma 2.4 with precision δ , let c be an approximation to $c(\mathcal{L}')$ and r be an approximation to $r(\mathcal{L}')$. Namely, $d(c, c(\mathcal{L}')) \leq \delta$ and $|r - r(\mathcal{L}')| \leq \delta$. Notice that applying Lemma 2.4 with precision δ can be done in time $d \text{ poly}(1/\varepsilon)$ (here we use the fact that the total dimension of the lines involved in the computation is independent of d). Also, as we assume that $r(\mathcal{L}') \geq 2$, it holds that $1 + \delta \geq \frac{r(\mathcal{L}')}{r} \geq 1 - \delta$.

To describe one iteration of our algorithm we will need to compute the axis of the cylinder containing certain approximate intersection centers of \mathcal{L}' . Namely, we

consider the set of $(1 + 2\varepsilon^2/200^2)$ -approximate intersection centers of \mathcal{L}' which (by Theorem 4.4 and the properties of c and r) are included in a cylinder of radius $\frac{25\sqrt{2}\varepsilon}{200}(r + \delta) + \delta \leq \frac{\varepsilon}{2}r(\mathcal{L}')$ passing through c . Let ℓ' be the axis of this cylinder. Recall from Theorem 4.4 that ℓ' is parallel to one of the lines in \mathcal{L}' . In what follows we assume that ℓ' is known. This assumption can be done w.l.o.g. as we may run the procedure described below several times, each time setting ℓ' to be parallel to a different line in \mathcal{L}' . As we soon show, we halt our procedure, or proceed to add lines to \mathcal{L}' in an additional iteration only if certain conditions hold. These conditions are sure to be encountered once we have chosen the correct line ℓ' .

If the ball centered at c of radius $(1 + \varepsilon)(1 - \delta)r \leq (1 + \varepsilon)r(\mathcal{L}')$ intersects \mathcal{L} , halt and output the set \mathcal{L}' . Otherwise, using Lemma 2.4 with precision δ , for each $\ell \in \mathcal{L} \setminus \mathcal{L}'$, we compute an estimate r_ℓ to $r(\mathcal{L}' \cup \{\ell\})$. Namely, $|r_\ell - r(\mathcal{L}' \cup \{\ell\})| \leq \delta$ and thus $1 - \delta \leq \frac{r(\mathcal{L}' \cup \{\ell\})}{r_\ell} \leq 1 + \delta$. If there exists a line $\ell \in \mathcal{L}$ such that $r_\ell \geq (1 + \varepsilon^2/200^2)\frac{1+\delta}{1-\delta}r$ (implying that $r(\mathcal{L}' \cup \{\ell\}) \geq (1 + \varepsilon^2/200^2)r(\mathcal{L}')$), add ℓ to \mathcal{L}' and proceed in an additional iteration of the algorithm.

We are now in a situation that for every line $\ell \in \mathcal{L}$, the radius of the minimum intersection ball of $\mathcal{L}' \cup \{\ell\}$ is very close to the radius of the minimum intersection ball of \mathcal{L}' , namely, $r_\ell < (1 + \varepsilon^2/200^2)\frac{1+\delta}{1-\delta}r$. This implies that $r(\mathcal{L}' \cup \{\ell\}) < (\frac{1+\delta}{1-\delta})^2(1 + \varepsilon^2/200^2)r(\mathcal{L}') \leq (1 + 2\varepsilon^2/200^2)r(\mathcal{L}')$. We now conclude that the center of the minimum intersection ball of $\mathcal{L}' \cup \{\ell\}$ is in the set of $(1 + 2\varepsilon^2/200^2)$ -approximate intersection centers of \mathcal{L}' , which in turn, is in the $\frac{\varepsilon}{2}r(\mathcal{L}')$ cylinder around ℓ' . In this case, we use the axis ℓ' to find a pair of lines that when added to \mathcal{L}' will increase $r(\mathcal{L}')$ substantially.

For each line $\ell_i \in \mathcal{L} \setminus \mathcal{L}'$, we now compute a certain interval I_i on ℓ' . Namely, we define I_i to be the set of points x on ℓ' such that the ball of radius $(1 - \delta)(1 + \varepsilon)r$ centered at x intersects the sets in $\mathcal{L}' \cup \{\ell_i\}$. It is not hard to verify that for each line ℓ_i , this interval I_i is not empty. Indeed, consider the minimum intersection center c_i^* of $\mathcal{L}' \cup \{\ell_i\}$. As mentioned above, c_i^* is in the $\frac{\varepsilon}{2}r(\mathcal{L}')$ cylinder around ℓ' . Consider the projection of c_i^* onto the line ℓ' . Denote this projection by c'_i . It now follows from the triangle inequality that the ball of radius $(1 + \varepsilon/2 + 2\varepsilon^2/200^2)r(\mathcal{L}') \leq \frac{1-\delta}{1+\delta}(1 + \varepsilon)r(\mathcal{L}') \leq (1 - \delta)(1 + \varepsilon)r$ centered at c'_i covers $\mathcal{L}' \cup \{\ell_i\}$, which implies that $c'_i \in I_i$.

If for all pairs of lines ℓ_i and ℓ_j the corresponding intervals intersect, then by Helly’s theorem, there is a point c' in the intersection of all the intervals. This implies that the ball of radius $(1 - \delta)(1 + \varepsilon)r \leq (1 + \varepsilon)r(\mathcal{L}')$ centered at c' covers all the lines, and we may halt the algorithm and output the set \mathcal{L}' .

Finally, if there are two lines ℓ_i and ℓ_j with corresponding intervals that do not intersect, then we claim that $r(\mathcal{L}' \cup \{\ell_i, \ell_j\}) \geq (1 + 2\varepsilon^2/200^2)r(\mathcal{L}')$ and we may add both ℓ_i and ℓ_j to \mathcal{L}' to increase the radius of \mathcal{L}' substantially and proceed in an additional iteration of the algorithm. Assume for contradiction that $r(\mathcal{L}' \cup \{\ell_i, \ell_j\}) < (1 + 2\varepsilon^2/200^2)r(\mathcal{L}')$ and let c^* be the minimum intersection center of $\mathcal{L}' \cup \{\ell_i, \ell_j\}$. As the ball of radius $(1 + 2\varepsilon^2/200^2)r(\mathcal{L}')$ centered at c^* also covers \mathcal{L}' , it follows that the distance of c^* from ℓ' is at most $\frac{\varepsilon}{2}r(\mathcal{L}')$. As before, consider the projection of c^* onto the line ℓ' . Denote this projection by c' . It now follows that the ball of radius $(1 + \frac{\varepsilon}{2} + 2\varepsilon^2/200^2)r(\mathcal{L}') \leq \frac{1-\delta}{1+\delta}(1 + \varepsilon)r(\mathcal{L}') \leq (1 - \delta)(1 + \varepsilon)r$ centered at c' covers $\mathcal{L}' \cup \{\ell_i, \ell_j\}$. This implies that the point c' is in the intervals I_i and I_j corresponding to ℓ_i and ℓ_j , which is a contradiction.

To summarize, in each iteration we either terminate with a set \mathcal{L}' for which $r(\mathcal{L}) \leq r(\mathcal{L}')(1 + \varepsilon)$ or we add 1 or 2 additional lines to \mathcal{L}' and increase the radius of \mathcal{L}' by a factor of at least $1 + \varepsilon^2/200^2$. After $O(1/\varepsilon^2)$ iterations we have a $(1 + \varepsilon)$ -core set. This concludes the proof for an $O(1/\varepsilon^2)$ sized core set.

Step II: Find a core set of size $O(1/\varepsilon)$ For a core set of size $O(1/\varepsilon)$, recall the iterative structure of the algorithm above. Our starting point was a 2 approximate core set of size 3, and our goal was a $(1 + \varepsilon)$ approximate core set. This was obtained by a series of steps in which an α approximate core set was enhanced to obtain an $\alpha(1 - \varepsilon^2/c)$ approximate core set (for some constant c). In all these steps, the value of ε remained fixed. We now show the benefit of changing the value of ε throughout the algorithm.

For an integer k , consider obtaining a $(1 + 2^{-k})$ approximate core set by performing $\simeq k$ rounds of the algorithm above. Namely, starting with a sufficiently small constant $\varepsilon = 2^{-k_0}$ (resulting from the proof of Step I above), in our first round we turn the 2 approximate core set of size 3 into a $(1 + 2^{-k_0})$ approximate core set. This first round will involve $O(2^{2k_0})$ iterations of the algorithm in Step I described above. In the second round, setting $\varepsilon = 2^{-k_0-1}$, we turn the $(1 + 2^{-k_0})$ approximate core set into a $(1 + 2^{-k_0-1})$ approximate core set. This second round will involve $O(2^{k_0+1})$ iterations of the algorithm in Step I. In general, in round i , setting $\varepsilon = 2^{-k_0-i}$, we turn the $(1 + 2^{-k_0-(i-1)})$ approximate core set into a $(1 + 2^{-k_0-i})$ approximate core set. The i th round will involve $O(2^{k_0+i})$ iterations of the algorithm in Step I. This process is continued until we obtain a $(1 + 2^{-k})$ approximate core set.

It follows that the first round increases the size of the core set by $O(2^{2k_0})$ and the remaining rounds increase the size of the core set by $c'2^{k_0+i}$, for some constant c' . Thus the total size of the resulting core set is $O(2^k)$ once $k_0 \leq \sqrt{k}$. We conclude our assertion for sufficiently small values of ε . \square

5 $(1 + \varepsilon)$ -Core Set for Convex Sets of Dimension $\leq \Delta$

The algorithm in the previous section can be extended to the general case of convex sets with dimension at most Δ , resulting in the proof of Theorem 1.3 stated in the Introduction and restated here for convenience.

Theorem 1.3. *Let $\varepsilon > 0$ be sufficiently small. Let \mathcal{L} be a set of convex sets of dimension $\leq \Delta$, $\mathcal{L} = \{\ell_1, \dots, \ell_n\}$, in \mathbb{R}^d . There exist a subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $O(\Delta^4/\varepsilon)$ such that $r(\mathcal{L}')(1 + \varepsilon) \geq r(\mathcal{L})$. The set \mathcal{L}' and a center c , such that the ball centered at c of radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects all sets in \mathcal{L} , can be found in time $O(n^{\Delta+1}d \text{ poly}(\Delta/\varepsilon))$.*

As in the previous section, the proof of Theorem 1.3 uses the notion of a $(1 + \varepsilon)$ -approximate intersection center. In what follows we state Theorem 5.2 which addresses $(1 + \varepsilon)$ -approximate intersection centers for Δ -dimensional convex sets. Then the proof of Theorem 1.3 is given. A detailed proof of Theorem 5.2 and several related Lemmas are given in Appendix D.

Definition 5.1 Let ℓ be a convex set with dimension $\leq \Delta$ in \mathbb{R}^d . A slab in \mathbb{R}^d of width r with axis ℓ is defined as the set of points in \mathbb{R}^d which are of distance at most r from ℓ . In other words, a slab in \mathbb{R}^d of width r with axis ℓ is defined to be the Minkowski sum of ℓ with a ball of radius r .

Theorem 5.2 *The set of $(1 + \varepsilon)$ -approximate intersection centers of a collection \mathcal{L} of convex sets with dimension at most Δ , $C_\varepsilon(\mathcal{L})$, is included in a Δ -slab of width $\beta\sqrt{\varepsilon(\Delta + 2)^3}r(\mathcal{L})$ for some constant β .*

Proof of Theorem 1.3 The basic idea is the same as in Theorem 1.2. We present the proof for a $(1 + \varepsilon)$ approximate core set of size $O(1/\varepsilon^2)$. To obtain a core set of size $O(1/\varepsilon)$, we use the ideas of [16] described in the proof of Theorem 1.2.

Our proof for a $(1 + \varepsilon)$ approximate core set of size $O(1/\varepsilon^2)$ is based on Theorem 1.2 with one major difference. As in Theorem 1.2, we show that an iterative algorithm, which starts from a set \mathcal{L}' of size $\Delta + 2$ and wisely adds $\Delta + 1$ new elements of \mathcal{L} in each iteration, will result in a core set of the desired size. However, in contrast to the algorithm presented in Theorem 1.2, our proof of which $\Delta + 1$ elements of \mathcal{L} to be added in each iteration is not constructive. We thus suggest an exhaustive search. Details follow.

We now turn to the proof for the existence of a small size core set. We start with $\Delta + 2$ sets $\mathcal{L}' \subseteq \mathcal{L}$ according to Theorem 1.1 such that $r(\mathcal{L}) \leq \alpha \cdot r(\mathcal{L}')$, $\alpha = 2$. Let $c = c(\mathcal{L}')$ be the minimum intersection center of \mathcal{L}' and ℓ' be the axis of the slab that contains the collection of $(1 + \frac{\varepsilon^2}{4\beta^2(\Delta+2)^3})$ -approximate intersection centers of \mathcal{L}' for some constant β in Theorem 5.2. Define I_i to be a subset of ℓ' such that a point $p \in I_i$ has distance at most $(1 + \varepsilon)r(\mathcal{L}')$ away from the sets $\mathcal{L}' \cup \{\ell_i\}$, $\ell_i \in \mathcal{L} \setminus \mathcal{L}'$.

If among all I_i , every $\Delta + 1$ of them have a nonempty intersection, then $\bigcap_i I_i \neq \emptyset$ by Helly’s theorem. Thus the ball with radius $(1 + \varepsilon)r(\mathcal{L}')$ centered at a point $c' \in \bigcap_i I_i$ intersects with every set in \mathcal{L} , and we are done (\mathcal{L}' is the core set).

If there are $\Delta + 1$ sets $\ell_1, \ell_2, \dots, \ell_{\Delta+1}$ such that their corresponding sets I_j , $j = 1, \dots, \Delta + 1$, do not have a common intersection, then it can be verified in the same way as in Theorem 1.2 that $r(\mathcal{L}' \cup \{\ell_1, \dots, \ell_{\Delta+1}\}) \geq (1 + \frac{\varepsilon^2}{4\beta^2(\Delta+2)^3})r(\mathcal{L}')$. Thus we add all the sets ℓ_j , $j = 1, \dots, \Delta + 1$, to \mathcal{L}' and go to the next iteration.

Thus, for each iteration, at most $\Delta + 1$ sets are added to \mathcal{L}' and the value α is decreased by a factor of $(1 - \frac{\varepsilon^2}{2\beta^2(\Delta+2)^3})$. After $O(\Delta^3/\varepsilon^2)$ steps, there are $O(\Delta^4/\varepsilon^2)$ sets in \mathcal{L}' such that $r(\mathcal{L}) \leq (1 + \varepsilon)r(\mathcal{L}')$.

This concludes our proof for the existence of a $(1 + \varepsilon)$ -core set for convex sets of dimension at most Δ . For an algorithm to compute this core set, notice that the axis ℓ' of the slab containing $C_\varepsilon(\mathcal{L}')$ is not known. Thus we will exhaustively try $\Delta + 1$ tuples of the remaining sets of \mathcal{L} at each iteration. We will terminate our algorithm once no $\Delta + 1$ tuple will increase the minimum intersecting radius significantly. The running time of the algorithm follows from the same arguments as in Theorem 1.2. \square

We would like to mention that the running time of the algorithm above for general Δ is greater than that implied by standard convex programming (Lemma 2.4). However, in our algorithm in addition to returning an approximate center c , we also

return the core set \mathcal{L}' whose existence should be viewed as the main contribution of this theorem. As suggested by Sariel Har-Peled [10], the following is an alternative algorithm with an improved running time at the price of a larger core set.

Theorem 5.3 *Let $\varepsilon > 0$ be sufficiently small. Let \mathcal{L} be a set of convex sets of dimension $\leq \Delta$, $\mathcal{L} = \{\ell_1, \dots, \ell_n\}$, in \mathbb{R}^d . There exist a subset $\mathcal{L}' \subseteq \mathcal{L}$ of size $\Delta^{O(1/\varepsilon)}$ such that $r(\mathcal{L}')(1 + \varepsilon) \geq r(\mathcal{L})$. The set \mathcal{L}' and a center c , such that the ball centered at c of radius $(1 + \varepsilon)r(\mathcal{L}')$ intersects all sets in \mathcal{L} , can be found in expected time $n(d\Delta^{O(1/\varepsilon)} + 2^{\Delta^{O(1/\varepsilon)}})$.*

Proof of Theorem 5.3 As in the previous proof, we start with a small core-set \mathcal{L}' of size $\Delta + 2$ such that $r(\mathcal{L}) \leq 2 \cdot r(\mathcal{L}')$. We proceed with a series of iterations. In each iteration both \mathcal{L}' and $r(\mathcal{L}')$ are going to grow substantially. The following algorithm will yield a core set of size $\Delta^{O(1/\varepsilon^2)}$, this may be reduced to $\Delta^{O(1/\varepsilon)}$ by applying the techniques of [16] described in Theorem 1.2.

Consider an intermediate iteration. Let \mathcal{L}' be the current core-set, $r' = r(\mathcal{L}')$ its minimum intersecting radius, and c' the center of a ball of radius r' that covers \mathcal{L}' . Consider the affine subspace H that includes all sets in \mathcal{L}' . As each set in \mathcal{L}' has dimension at most Δ , the total dimension h of H is at most $\Delta|\mathcal{L}'|$. It holds that $c' \in H$. Using a slightly refined analysis to that presented in Lemma 3.3, one can find the minimum radius ball $B_{c''}(r'')$ that intersects all sets in \mathcal{L} with center in H in time $n(d \text{poly}(h) + 2^{\text{poly}(\Delta)+O(h \log h)})$ (recall that each convex set in \mathcal{L} is defined by at most $\text{poly}(\Delta)$ constraints). Moreover, the algorithm used in Lemma 3.3 also yields a subset $\mathcal{L}'' \subseteq \mathcal{L}$ of size $h + 1$ such that the minimum radius ball that covers \mathcal{L}'' with center in H has radius r'' . If $r'' \leq (1 + \varepsilon)r'$, then we are done (\mathcal{L}' is the core-set we were looking for). Otherwise, we now claim that by taking $\mathcal{L}^* = \mathcal{L}' \cup \mathcal{L}''$ as our new core-set we get $r(\mathcal{L}^*) \geq r(\mathcal{L}')(1 + \varepsilon^2/3)$.

Let $B_{c^*}(r^*)$ be a minimum intersecting ball of \mathcal{L}^* . Notice that c^* is no longer necessarily on H . Let δ be the distance of c^* from H . We now consider two cases, one in which δ is large and the second in which δ is small. In both cases we will use the projection c_H^* of c^* to H .

Case 1: $\delta \geq r' \frac{2\varepsilon + \varepsilon^2}{2(1 + \varepsilon)}$. There exists a set ℓ of distance at least r' from c_H^* (this follows from the fact that H includes all sets in \mathcal{L}' and by the definition of r'). This implies that the distance between c^* and ℓ is at least $\sqrt{(r')^2 + \delta^2} \geq r'(1 + \varepsilon^2/3)$.

Case 2: $\delta < r' \frac{2\varepsilon + \varepsilon^2}{2(1 + \varepsilon)}$. Recall that $r'' > (1 + \varepsilon)r'$. Thus there exists a set ℓ in \mathcal{L}^* of distance at least $(1 + \varepsilon)r'$ from c_H^* . As the distance between c^* and c_H^* is δ , the set ℓ is of distance at least $r'(1 + \varepsilon) - \delta \geq r'(1 + \varepsilon^2/3)$ from c^* . Hence, we have shown that $r(\mathcal{L}^*) \geq r(\mathcal{L}')(1 + \varepsilon^2/3)$.

After $O(1/\varepsilon^2)$ such iterations we will finally obtain a $(1 + \varepsilon)$ core-set \mathcal{L}' of size $\Delta^{O(1/\varepsilon^2)}$. After modifying the algorithm to obtain a core set of size $\Delta^{O(1/\varepsilon)}$, the expected running time (which will be dominated by the executions of Lemma 3.3) sums up to $n(d\Delta^{O(1/\varepsilon)} + 2^{\Delta^{O(1/\varepsilon)}})$. □

Acknowledgements The authors would like to thank Sariel Har-Peled for suggesting to analyze the algorithm of Theorem 5.3, Piyush Kumar for pointing out the work of [16] and its application in the context of clustering lines, and Lin Xiao for several helpful discussions on convex programming.

Appendix A: Lemma 2.4

Lemma 2.4. *Let \mathcal{L} be a set of n convex sets with dimension at most Δ in \mathbb{R}^d , where each convex set can be represented by $\text{poly}(\Delta)$ convex constraints. Then $c(\mathcal{L})$ and $r(\mathcal{L})$ can be computed to an absolute precision $\delta > 0$ in time $O(\sqrt{n}(d^3 + d^2n) \text{poly}(\Delta) \log(n/\delta))$.*

Proof We can formulate this problem by convex programming. Assume that c is a point in \mathbb{R}^d and r is the intersection radius of c with respect to \mathcal{L} . Each convex set ℓ_i has dimension at most Δ and is represented as follows. ℓ_i is in a Δ -dimensional space B_i , which has origin $o^{(i)}$ and is spanned by Δ unit vectors $\vec{b}_j^{(i)}$, $j = 1, \dots, \Delta$. Therefore, each point in ℓ_i can be represented by $o^{(i)} + \sum_{j=1}^{\Delta} \lambda_j^{(i)} \vec{b}_j^{(i)}$, where $\lambda_j^{(i)}$ is a scalar. Inside B_i , the convex set ℓ_i is specified by m_i convex constraints $f_j^{(i)}(\lambda^{(i)}) \leq 0$, $j = 1, \dots, m_i$, where $\lambda^{(i)} = \{\lambda_j^{(i)}\}$ (here we assume that $m_i = \text{poly}(\Delta)$). We can find the minimum intersection radius and center of \mathcal{L} by solving the following optimization problem:

$$\begin{aligned} & \min r \\ & \text{s.t.} \quad \left\| c - o^{(i)} - \sum_{j=1}^{\Delta} \lambda_j^{(i)} \vec{b}_j^{(i)} \right\| \leq r, \quad 1 \leq i \leq n; \\ & \quad \quad f_j^{(i)}(\lambda^{(i)}) \leq 0, \quad 1 \leq j \leq m_i, 1 \leq i \leq n. \end{aligned}$$

This is a convex optimization problem with variables $\lambda_j^{(i)}$, r , and c . The total number of variables is $k = n\Delta + d + 1$. The total number of constraints is N , where $N = \sum_i m_i + n = O(n \text{poly}(\Delta))$. Thus one can find the solution up to precision $\delta > 0$ in time $O(\sqrt{N}(n\Delta + d)^3 \log(N/\delta))$ by using a generic interior point method [8].

Specifically, define x to be a column vector of all the variables $\lambda_j^{(i)}$, r , and c . Denote $f_0(x)$ as the objective function r and $f_i(x) \leq 0$ all the N constraints. Define $f(x)$ as $tf_0(x) + \phi(x)$, where $t > 0$ is a parameter in the barrier method and $\phi(x)$ is the barrier function, $\phi(x) = -\sum_{i=1}^N \log(-f_i(x))$. Basically the barrier method performs $O(\sqrt{N} \log(N/\delta))$ Newton iterations. Each Newton iteration solves the system $Hy = g$, where $H = \nabla^2 f(x)$, $g = \nabla f(x)$. H is the Hessian matrix of $f(x)$ and is symmetric positive definite. H has dimension k by k , where $k = n\Delta + d + 1$. The system can be solved via Cholesky factorization with $H = LL^T$, $y = -L^{-T}L^{-1}g$, where L is an upper triangular matrix. This can be done in time $O(k^3) = O((n\Delta + d)^3)$. Putting everything together, the interior point method takes time $O(\sqrt{N}(n\Delta + d)^3 \log(N/\delta))$ to find the center and radius up to precision δ .

A more careful analysis by exploring the sparsity of H shows a better bound on the running time. Specifically, the lower right submatrix in H with dimension $n\Delta$ by $n\Delta$ has n blocks, each of size Δ by Δ along the diagonal. Thus the Cholesky factorization of H takes $O(d^3 + nd^2\Delta + nd\Delta^2 + n\Delta^3) = O(d^3 + nd^2\Delta)$ operations. The running time of the algorithm can be improved to $O(\sqrt{N}(d^3 + d^2n\Delta) \log(N/\delta)) = O(\sqrt{n}(d^3 + d^2n) \text{poly}(\Delta) \log(n/\delta))$. □

Appendix B: Lemma 3.3

Lemma 3.3. *Let \mathcal{L} be a set of n convex sets with dimension at most Δ in \mathbb{R}^d . Let $\ell \in \mathcal{L}$. The minimum radius ball $B_r(c)$ that covers the sets \mathcal{L} with center c on ℓ , and a set $\mathcal{L}' = \{\ell_1, \dots, \ell_{\Delta+2}\} \subseteq \mathcal{L}$ for which $r(\mathcal{L}) \leq 2r(\mathcal{L}')$ can be found in expected time $n(d \text{ poly}(\Delta) + 2^{\text{poly}(\Delta)})$.*

Proof Let \mathcal{L} be a set of n convex sets of dimension Δ in \mathbb{R}^d . Let $\ell_0 \in \mathcal{L}$. In what follows we study the problem of finding the minimum radius ball covering \mathcal{L} with center on ℓ_0 . We show that this problem falls in the abstract framework of ‘LP-type’ problems and can be solved by the randomized algorithm of [19] in expected running time $n(d \text{ poly}(\Delta) + 2^{\text{poly}(\Delta)})$. The algorithm of [19] not only finds the minimum radius ball (say $B_r(c)$ centered at c of radius r) covering \mathcal{L} with center on ℓ_0 , it also returns a subset \mathcal{L}' of \mathcal{L} of size $\Delta + 1$ such that the minimum radius ball covering \mathcal{L}' with center on ℓ_0 is also $B_r(c)$. This implies that $2r(\mathcal{L}' \cup \ell_0) \geq r(\mathcal{L})$. Indeed, $r \geq r(\mathcal{L})$ and $r(\mathcal{L}' \cup \ell_0) \geq \frac{r}{2}$, otherwise by projecting onto ℓ_0 and using the triangle inequality one could find a ball of radius less than r that covers \mathcal{L}' with center on ℓ_0 .

We now sketch the proof that the problem at hand is an LP-type problem. Throughout our proof we use the notation in the paper [19] freely. Our problem is defined by a couple (\mathcal{L}, w) where w is the function on subsets \mathcal{L}' of \mathcal{L} which returns the minimum radius ball covering the sets in \mathcal{L}' with center on ℓ_0 . For various technical reasons, we alter w as to satisfy *basis regularity* (as described in [19]). Namely, $w(\mathcal{L}') = (r, c, \min(\Delta + 1, |\mathcal{L}'|))$ where $B_r(c)$ is the ball that covers \mathcal{L}' with center on ℓ_0 , which is lexicographically minimal (in the set $(\mathbb{R}, \mathbb{R}^d, \{1, \dots, \Delta + 1\})$). We use the notation \leq to denote lexicographic order ‘less than or equal’ of points of $w(\cdot)$. If no such minimum exists, we define $w(\mathcal{L}')$ to be the special character $-\infty$. To use the framework outlined in [19], it suffices to prove the following claims.

Claim 1 *For every $\mathcal{L}' \subseteq \mathcal{L}$, there exists a subset $\mathcal{L}'' \subseteq \mathcal{L}'$ of size at most $\Delta + 1$ such that $w(\mathcal{L}') = w(\mathcal{L}'')$, i.e., the combinatorial dimension of (\mathcal{L}, w) is $\Delta + 1$.*

Proof For $\mathcal{L}' = \{\ell_1, \dots, \ell_s\}$, $s > \Delta + 1$, let $w(\mathcal{L}') = (r, c, \Delta + 1)$. Let S_i^+ be the set of points p in ℓ_0 that are either (a) within distance greater than r from ℓ_i or (b) within distance r from ℓ_i and have lexicographic order greater than or equal to c (i.e., $p \geq c$). The set S_i^+ is exactly the centers p in ℓ_0 that are *no better* than the center c when considering only ℓ_i . Let S_i^- be the set $\ell_0 \setminus S_i^+$. It is not hard to verify that S_i^- is convex. As $w(\mathcal{L}') = (r, c, \Delta + 1)$, it follows that $\bigcap_{i=1}^s S_i^- = \emptyset$. By Helly’s theorem, there exist a set $\mathcal{L}'' \subset \mathcal{L}'$ of size $\Delta + 1$ such that $\bigcap_{\ell_i \in \mathcal{L}''} S_i^- = \emptyset$. This implies that $w(\mathcal{L}'') = w(\mathcal{L}')$. \square

Claim 2 (a) *Monotonicity:* For every $\mathcal{L}'' \subseteq \mathcal{L}' \subseteq \mathcal{L}$, it is the case that $w(\mathcal{L}'') \leq w(\mathcal{L}')$. (b) *Locality:* For every $\mathcal{L}'' \subseteq \mathcal{L}' \subseteq \mathcal{L}$ with $-\infty < w(\mathcal{L}'') = w(\mathcal{L}')$ and any $\ell \in \mathcal{L}$, if $w(\mathcal{L}') < w(\mathcal{L}' \cup \{\ell\})$, then $w(\mathcal{L}'') < w(\mathcal{L}'' \cup \{\ell\})$.

Proof Both (a) and (b) above follow from the definition on w . Namely, for monotonicity, the minimum radius ball in the solution $w(\mathcal{L}')$ intersect all the lines

in \mathcal{L}'' . Thus the lexicographically minimum ball for \mathcal{L}'' is smaller. For locality, the minimum radius ball in $w(\mathcal{L}')$ does not cover ℓ —otherwise $w(\mathcal{L}' \cup \{\ell\}) \leq w(\mathcal{L}')$, which contradicts with the assumption. Thus we must have $w(\mathcal{L}'') < w(\mathcal{L}'' \cup \{\ell\})$. If else, $w(\mathcal{L}'') \geq w(\mathcal{L}'' \cup \{\ell\})$, then the minimum radius ball in $w(\mathcal{L}'') = w(\mathcal{L}')$ covers ℓ . \square

The above two claims support the existence of *basis*. A basis \mathcal{L}'' is a subset of \mathcal{L} such that for any $\ell \in \mathcal{L}''$, it is the case that $w(\mathcal{L}'' \setminus \{\ell\}) < w(\mathcal{L}'')$. A basis \mathcal{L}'' for a subset \mathcal{L}' is a basis $\mathcal{L}'' \subseteq \mathcal{L}'$ such that $w(\mathcal{L}') = w(\mathcal{L}'')$.

Claim 3 *The following operations can be computed in expected time $d \text{poly}(\Delta) + 2^{\text{poly}(\Delta)}$: (a) Violation test: Given a basis \mathcal{L}'' and a set ℓ , decide if $w(\mathcal{L}'') < w(\mathcal{L}'' \cap \{\ell\})$ or not. (b) Basis Computation: Given a basis \mathcal{L}'' and a set ℓ , compute a basis for $\mathcal{L}'' \cap \{\ell\}$.*

Proof In what follows we sketch an algorithm which given \mathcal{L} computes $w(\mathcal{L})$. This will suffice to show that operations (a) and (b) above can be solved as asserted. We use the fact that given a set of n convex sets in \mathbb{R}^d (each defined by $\text{poly}(d)$ linear constraints) one can find their minimum intersecting ball in expected time $n2^{O(\sqrt{d \log d})}$ (see, for example, [19]).

Let \mathcal{L} be a set of n convex sets in \mathbb{R}^d of dimension at most Δ each. Assume that ℓ_0 is a Δ -flat. Recall that $w(\mathcal{L})$ corresponds to the minimum intersecting ball of \mathcal{L} with center on ℓ_0 . We first tie this problem with the problem of finding the minimum intersecting ball of a certain set of convex sets, without the restriction that the center is on ℓ_0 . For any $\ell \in \mathcal{L}$, define ℓ^- to be the *mirror image* of ℓ with respect to ℓ_0 . Namely, for a point $p \in \mathbb{R}^d$, let c be the point on ℓ_0 such that $d(p, \ell_0) = d(p, c)$. The mirror image of p with respect to ℓ_0 is the point p^- for which $d(p, \ell_0) = d(p^-, \ell_0)$, and the line through p, p^- passes through c . Now, $\ell^- = \{p^- | p \in \ell\}$, and $\mathcal{L}^- = \{\ell^- | \ell \in \mathcal{L}\}$.

Consider the set $\mathcal{L} \cup \mathcal{L}^-$. We now claim that the radius of the minimum intersecting ball of $\mathcal{L} \cup \mathcal{L}^-$ is equal to the radius of the minimum intersecting ball of \mathcal{L} with center on ℓ_0 . Indeed, let $B_r(c)$ be the minimum intersecting ball of $\mathcal{L} \cup \mathcal{L}^-$. Consider the projection of c on ℓ_0 , denoted as c_0 . It is not hard to verify that $B_r(c_0)$ is an intersecting ball of $\mathcal{L} \cup \mathcal{L}^-$. Moreover, any ball with center on ℓ_0 that covers \mathcal{L} will also cover \mathcal{L}^- .

Thus, to find the radius of the minimal intersecting ball of \mathcal{L}' with center on ℓ_0 , it suffices to compute the radius of the minimum intersecting ball of $\mathcal{L} \cup \mathcal{L}^-$. This is addressed above when ℓ_0 is a flat. If ℓ_0 is a convex polyhedron defined by the intersection of $\text{poly}(\Delta)$ half spaces, then the center of the minimal intersecting ball of \mathcal{L}' with center in ℓ_0 either lies in the interior of ℓ_0 or in the interior of one of its facets, or on one of the vertices of ℓ_0 . To find if the minimum intersecting ball of \mathcal{L}' with center in ℓ_0 falls in the interior of a certain facet, we use the procedure above when \mathcal{L}^- is defined with respect to the (subspace including the) facet under consideration. This implies that finding the minimal intersecting ball of \mathcal{L}' with center in ℓ_0 can be done by solving $2^{\text{poly}(\Delta)}$ occurrences of the general minimum intersecting ball problem, which all in all will take expected time $n2^{O(\text{poly}(\Delta) + \sqrt{d \log d})}$.

Given a set system \mathcal{L} , it is left to show how to obtain the asserted running time of $d \text{poly}(\Delta) + 2^{\text{poly}(\Delta)}$ to compute $w(\mathcal{L})$. The input in both the violation test and the

basis computation is a set system \mathcal{L} of size $O(\Delta)$. As each set in \mathcal{L} is of dimension at most Δ , the total set system lies in $O(\Delta^2)$ dimensions. Thus in time $d \text{poly}(\Delta) + 2^{\text{poly}(\Delta)}$ one can compute the radius of the minimum radius ball that intersects all sets in \mathcal{L} and is centered in ℓ_0 (the additional factor of $d \text{poly}(\Delta)$ appears as the dimension of the sets must be reduced from d to $\text{poly}(\Delta)$). However, we are also interested in the *minimum* (w.r.t. the lexicographic ordering) center c on ℓ_0 such that $B_r(c)$ intersects \mathcal{L} . Let $\ell_0^r \subseteq \ell_0$ be the convex region on ℓ_0 consisting of points c which are of distance r from every set in \mathcal{L} . Using several iterations of convex programming (similar to that presented in the proof of Lemma 2.4), the point in ℓ_0^r which is minimal w.r.t. the lexicographic ordering can be found in time $d \text{poly}(\Delta)$. Finally, in order to find a basis for \mathcal{L} , we can try removing sets from \mathcal{L} and take the minimal subset which has the same w value. \square

This concludes the proof of Lemma 3.3. \square

Appendix C: Lemmas 4.5, 4.6 and Theorem 4.4

Lemma 4.5. *Suppose that c is a minimum intersection center of a set of lines \mathcal{L} and p is a $(1 + \varepsilon)$ -approximate intersection center with $|cp| \geq \alpha \sqrt{\varepsilon} r(\mathcal{L})$. Then there exists a line $\ell \in \mathcal{L}$ such that the angle between ℓ and the line segment cp is bounded (by above) by $\arcsin(\sqrt{2 + \varepsilon}/\alpha)$.*


Proof Consider a point t on cp that moves infinitesimally away from c towards p . We argue that there must be a line $\ell \in \mathcal{L}$ such that $d(c, \ell) = r(\mathcal{L})$ and the distance from a point t on cp to ℓ is nondecreasing when t moves infinitesimally from c to p on cp . Otherwise, the distances from t to all the lines in \mathcal{L} are strictly less than $r(\mathcal{L})$ when t moves infinitesimally away from c . This contradicts with the fact that $r(\mathcal{L})$ is the minimum intersection radius.

Now suppose that for the line $\ell \in \mathcal{L}$, the distance from t to ℓ stays the same when t moves infinitesimally away from c . Then it must be that cp is parallel with ℓ . Therefore the angle between ℓ and cp is zero. The claim is true.

If the distance from t to ℓ is strictly increasing when t moves infinitesimally away from c , the distance $d(t, \ell)$ is monotonically increasing as t moves linearly from c to p . Now we bound the angle between cp and ℓ as follows. Note that the point p is a $(1 + \varepsilon)$ -approximate intersection center, $r(\mathcal{L}) < d(p, \ell) \leq (1 + \varepsilon)r(\mathcal{L})$. We denote by ℓ' the line that is parallel with ℓ and goes through center c . Let q be the point on ℓ for which pq is perpendicular to ℓ , $|pq| = d(p, \ell) \leq (1 + \varepsilon) \cdot r(\mathcal{L})$. Let q' be the point on ℓ' for which qq' is perpendicular to ℓ' , $|qq'| = d(c, \ell) = r(\mathcal{L})$. Finally, let s be the point on ℓ for which cs is perpendicular to ℓ , $|cs| = d(c, \ell) = r(\mathcal{L})$. See Fig. 3.

Let H be the hyperplane with normal cs that passes through c , and let S be the (closed) half space defined by H that does not include ℓ . The segment cp (not including c) is either entirely contained in S or entirely contained in the complement of S . We now claim the cp is in S . Consider the ball B of radius $r(\mathcal{L})$ around s . The point c is on the boundary of both B and S , and the segment cp does not intersect the interior of B . Otherwise there would be a point on cp of distance less than $r(\mathcal{L})$ to s . This can only happen if cp is in S .

Fig. 3 If the distance between an approximate center p to the minimum center c is more than $\alpha\sqrt{\varepsilon} \cdot r(\mathcal{L})$, then there exists a line whose angle to line cp is bounded by $\arcsin(\sqrt{2 + \varepsilon}/\alpha)$



If p is in S , the inner angle of the triangle $\Delta pq'q$ at vertex q' , $\angle pq'q \geq \pi/2$. Thus,

$$|pq'| \leq \sqrt{|pq|^2 - |qq'|^2} \leq r(\mathcal{L})\sqrt{(1 + \varepsilon)^2 - 1} = r(\mathcal{L})\sqrt{2\varepsilon + \varepsilon^2}.$$

On the other hand, ℓ' is perpendicular to both the line qq' and the line pq . Thus ℓ' is perpendicular to the plane defined by the triangle $pq'q$. This implies that $\angle pq'c = \pi/2$. Therefore,

$$\sin \angle pcq' = \frac{|pq'|}{|pc|} \leq \frac{\sqrt{2\varepsilon + \varepsilon^2} \cdot r(\mathcal{L})}{\alpha\sqrt{\varepsilon} \cdot r(\mathcal{L})} = \sqrt{2 + \varepsilon}/\alpha.$$

Thus the angle between ℓ and line cp is $\angle pcq' \leq \arcsin(\sqrt{2 + \varepsilon}/\alpha)$. □

Lemma 4.6. $C_\varepsilon(\mathcal{L})$ does not include a 2-dimensional disk with radius $\frac{5}{2}\sqrt{\varepsilon}r(\mathcal{L})$.

Proof Assume that there is a disk D with radius $\frac{5}{2}\sqrt{\varepsilon}r(\mathcal{L})$ inside $C_\varepsilon(\mathcal{L})$. Denote by u the center of D . Consider two lines through u that are perpendicular to each other. Suppose that the intersections of these two lines with the boundary of D are a, b, c, d , see Fig. 4. $|ua| = |ub| = |uc| = |ud| = 5\sqrt{\varepsilon}r(\mathcal{L})/2$.

Now we claim that for every line $\ell \in \mathcal{L}$, $d(u, \ell) < r(\mathcal{L})$. Assume that, otherwise, there is a line $\ell \in \mathcal{L}$ such that $d(u, \ell) \geq r(\mathcal{L})$. We observe that when we move a point t continuously from u to a or from u to c , at least in one case the distance $d(t, \ell)$ is monotonically increasing. Similarly for nodes b, d . We take such two vectors, say \vec{ua}, \vec{ub} . The vectors \vec{ua}, \vec{ub} are perpendicular to each other. Furthermore, a, b are both inside $C_\varepsilon(\mathcal{L})$, so $d(a, \ell) \leq (1 + \varepsilon)r(\mathcal{L})$, $d(b, \ell) \leq (1 + \varepsilon)r(\mathcal{L})$. Also $|ua| = |ub| = 5\sqrt{\varepsilon}r(\mathcal{L})/2$. By applying the arguments of Lemma 4.5, the angle between ℓ to the line ua (ub) is at most $\arcsin(2\sqrt{2 + \varepsilon}/5) < \pi/4$, a contradiction.

Thus, $d(u, \ell) < r(\mathcal{L})$ for any line $\ell \in \mathcal{L}$. We conclude that u is an intersection center of \mathcal{L} with radius less than $r(\mathcal{L})$. This gives a contradiction to the definition of $r(\mathcal{L})$. □

Theorem 4.4. $C_\varepsilon(\mathcal{L})$ is included in a cylinder of radius $25\sqrt{\varepsilon}r(\mathcal{L})$ with axis parallel to one of the lines $\ell_i \in \mathcal{L}$. Moreover this axis passes through $c(\mathcal{L})$.




Fig. 4 The set of $(1 + \varepsilon)$ -approximate centers $C_\varepsilon(\mathcal{L})$ does not include a 2-dimensional disk with radius $\frac{5}{2}\sqrt{\varepsilon} \cdot r(\mathcal{L})$





Fig. 5 If $|cp| > 20\sqrt{\varepsilon} \cdot r(\mathcal{L})$ and the cylinder with axis cp and radius $20\sqrt{\varepsilon} \cdot r(\mathcal{L})$ does not include all the approximate centers, we can find a square R inside Δvcp with side length $5\sqrt{\varepsilon} \cdot r(\mathcal{L})$

Proof We first study the case where the set of minimum intersection centers is included in a finite radius ball centered at the origin. Then the set of approximate intersection centers is also included in a finite radius ball. Let c be a minimum intersection center. Denote by p the approximate center such that $|cp|$ is maximum. If $|cp| \leq 20\sqrt{\varepsilon}r(\mathcal{L})$, then the approximate centers are inside a ball with radius $20\sqrt{\varepsilon}r(\mathcal{L})$ centered at c and thus included in any cylinder with axis through c and radius $20\sqrt{\varepsilon}r(\mathcal{L})$. The claim is true.

If $|cp| > 20\sqrt{\varepsilon}r(\mathcal{L})$, we claim that the cylinder with axis cp and radius $20\sqrt{\varepsilon}r(\mathcal{L})$ includes $C_\varepsilon(\mathcal{L})$. Assume that, otherwise, there must be an approximate center v with distance more than $20\sqrt{\varepsilon}r(\mathcal{L})$ away from the line cp . Denote by q the reflection point of p on the line cp , $|cq| = |cp|$. Since p is the furthest away approximate center from c , the projections of the other approximate centers on the line including cp fall inside the line segment pq . Now we consider the triangle Δvcp . By the convexity of $C_\varepsilon(\mathcal{L})$, all the points inside Δvcp are $(1 + \varepsilon)$ -approximate centers. We claim that there must be a 2-dimensional square R with side length $5\sqrt{\varepsilon}r(\mathcal{L})$ inside Δvcp . Take the middle point of the line segment cp , denoted as s . $|cs| > 10\sqrt{\varepsilon}r(\mathcal{L})$. Take the point s' on line segments cv or vp such that the projection of s' on the line cp is s . Now we argue that a square of side length $5\sqrt{\varepsilon}r(\mathcal{L})$ with s as one corner and a portion of line segment sc as one of the sides must be completely inside Δvcp . If the projection of v on the line cp lies on the segment cp , as shown in Fig. 5(i), then the length of ss' is at least $10\sqrt{\varepsilon}r(\mathcal{L})$. Thus $\Delta css'$ must have R completely in Δvcp . If the projection of v on the line cp lies on the segment cq , as shown in Fig. 5(ii), the length of ss' is at least $5\sqrt{\varepsilon}r(\mathcal{L})$. Again it is not hard to verify that the square R is completely inside Δvcp .

Thus we can find a square R inside Δvcp with side length $5\sqrt{\varepsilon}r(\mathcal{L})$, as shown in Fig. 5. This implies a contradiction by Lemma 4.6. Thus the cylinder with axis cp and radius $20\sqrt{\varepsilon}r(\mathcal{L})$ includes all the approximate centers.

Fig. 6 There is a cylinder with axis parallel with one of the lines in \mathcal{L} and radius $25\sqrt{\varepsilon}r(\mathcal{L})$ that includes the collection of $(1 + \varepsilon)$ -approximate intersection centers



For a line $\ell_i \in \mathcal{L}$, consider the line ℓ'_i that is parallel with ℓ_i and goes through c . We claim that one of the cylinders with axis ℓ'_i and radius $25\sqrt{\varepsilon}r(\mathcal{L})$ includes $C_\varepsilon(\mathcal{L})$. Recall that q is the reflection point of p on the line cp and the projections of the other approximate centers on the line cp fall inside line segment pq . Let α satisfy $|cp| > \alpha\sqrt{\varepsilon}r(\mathcal{L})$; by Lemma 4.5 there is a line $\ell \in \mathcal{L}$ such that the angle between ℓ and the line cp is $\arcsin(\sqrt{2 + \varepsilon}/\alpha)$. Take ℓ' to be the line through c which is parallel with ℓ . The distance from p to ℓ' is at most $r(\mathcal{L})\sqrt{2\varepsilon + \varepsilon^2}$. Thus the distance from any point on the line segment pq to line ℓ' is no more than $r(\mathcal{L})\sqrt{2\varepsilon + \varepsilon^2}$, by simple geometry. Take any approximate center t'' , assume that its projection to line cp is t and the projection of t on line ℓ' is t' . The distance from t'' to ℓ' is

$$\begin{aligned} d(t'', \ell') &\leq d(t'', t) + d(t, t') \\ &\leq 20\sqrt{\varepsilon}r(\mathcal{L}) + r(\mathcal{L})\sqrt{2\varepsilon + \varepsilon^2} \\ &\leq 25\sqrt{\varepsilon}r(\mathcal{L}). \end{aligned}$$

Thus the cylinder with axis ℓ' and radius $25\sqrt{\varepsilon}r(\mathcal{L})$ includes the collection of $(1 + \varepsilon)$ -approximate centers.

If the set of minimum intersection centers is not included in a finite radius ball, then all the lines are parallel, and the cylinder that includes all the approximate centers must have an axis parallel with the lines in \mathcal{L} (otherwise it cannot cover the minimum intersection centers). Thus if there is an approximate center v outside the cylinder with radius $25\sqrt{\varepsilon}r(\mathcal{L})$, by essentially the same analysis used for the case in which the set of minimum intersection centers is included in a finite radius ball, we can find a square with side length more than $5\sqrt{\varepsilon}r(\mathcal{L})$ completely inside $C_\varepsilon(\mathcal{L})$. By using Lemma 4.6 we have a contradiction. Thus the cylinder with axis that equals the line of the minimum intersection centers and radius $25\sqrt{\varepsilon}r(\mathcal{L})$ includes all the approximate centers. \square

Appendix D: Line Segments and Convex Sets of Dimension Δ

Lemma D.1 For a set of $\Delta + 1$ orthogonal vectors $\{v_1, v_2, \dots, v_{\Delta+1}\}$ and a Δ -flat ℓ in $\mathbb{R}^{\Delta+1}$, there must be a vector v_j such that the angle between v_j and ℓ is at least $\arcsin(1/\sqrt{\Delta + 1})$.

Proof Without loss of generality, we can assume that v_i 's are the unit vectors along the $\Delta + 1$ axes and ℓ passes through the origin. We take the unit normal vector v of ℓ . Then the angle between v_i and ℓ , denoted by θ_i , is $\pi/2 - \theta'_i$, where θ'_i is the angle

between vectors v_i and v . It holds that $\sum_{i=1}^{\Delta+1} \langle v, v_i \rangle^2 = 1$. Thus there exists i such that $\langle v, v_i \rangle \geq 1/\sqrt{\Delta+1}$. This implies that $\theta_i \geq \arcsin(1/\sqrt{\Delta+1})$. \square

Theorem 5.2. *The set of $(1 + \varepsilon)$ -approximate intersection centers of a collection \mathcal{L} of convex sets with dimension at most Δ , $C_\varepsilon(\mathcal{L})$, is included in a Δ -slab of width $\beta\sqrt{\varepsilon(\Delta+2)^3}r(\mathcal{L})$ for some constant β .*

Proof The proof is similar to the proof of Theorem 4.4 and is based on the analogs of Lemmas 4.5 and 4.6 stated as follows. We first start with the case of line segments.

Lemma D.2 *Suppose that c is a minimum intersection center of a set \mathcal{L} of lines or line segments and p is a $(1 + \varepsilon)$ -approximate intersection center, $|cp| \geq \alpha\sqrt{\varepsilon}r(\mathcal{L})$, with $\alpha > 2\sqrt{2 + \varepsilon}$, then there exists a line or line segment $\ell \in \mathcal{L}$ such that the angle between ℓ and the line segment cp is bounded (by above) by $\arcsin(\sqrt{2 + \varepsilon}/(\alpha - 2\sqrt{2 + \varepsilon}))$.*

Proof The proof is basically the same as that of Lemma 4.5. Using the notation of Lemma 4.5, we can essentially assume that the closest point s to c on ℓ is not one of the endpoints of the line segment ℓ . The same holds for the closest point q to p on ℓ . If otherwise, we can find a point \hat{c} on the line segment cp such that \hat{c} is the furthest point from c on cp such that \hat{c} 's closest point on ℓ is still s . Notice that $r(\mathcal{L}) \leq d(\hat{c}, \ell) \leq (1 + \varepsilon)r(\mathcal{L})$. Since the distance $d(t, \ell)$ is monotonically nondecreasing as t moves from c to p , the inner angle of $\Delta\hat{c}cs$ at c is at least $\pi/2$. Thus $|c\hat{c}| \leq r(\mathcal{L})\sqrt{(1 + \varepsilon)^2 - 1} = r(\mathcal{L})\sqrt{(2 + \varepsilon)\varepsilon}$. Similarly, one can find the corresponding point \hat{p} on $\hat{c}p$. $|p\hat{p}| \leq \sqrt{(2 + \varepsilon)\varepsilon}r(\mathcal{L})$. Thus $|\hat{c}\hat{p}| \geq (\alpha - 2\sqrt{2 + \varepsilon})\sqrt{\varepsilon}r(\mathcal{L})$. Now we follow the arguments in the case of lines. The angle between ℓ and cp is bounded by $\arcsin(\sqrt{2 + \varepsilon}/(\alpha - 2\sqrt{2 + \varepsilon}))$. \square

Lemma D.3 *Suppose that c is a minimum intersection center of a set \mathcal{L} of convex sets with dimension at most Δ and p is a $(1 + \varepsilon)$ -approximate intersection center, $|cp| \geq \alpha\sqrt{\varepsilon}r(\mathcal{L})$, $\alpha > 2\sqrt{2 + \varepsilon}$. Then there exists a convex set $\ell \in \mathcal{L}$ such that the angle between ℓ and the interval cp is bounded (by above) by $\arcsin(\sqrt{2 + \varepsilon}/(\alpha - 2\sqrt{2 + \varepsilon}))$.*

Proof The proof is very similar to the proof of Lemmas 4.5 and D.2. There is a convex set $\ell \in \mathcal{L}$ such that $d(c, \ell) = r(\mathcal{L})$, $d(p, \ell) \leq (1 + \varepsilon)r(\mathcal{L})$ and $d(t, \ell)$ is monotonically nondecreasing as t moves from c to p . Let s be the point on ℓ closest to c and q the point on ℓ closest to p . Now we know that $|cs| = r(\mathcal{L})$ and $|pq| \leq (1 + \varepsilon)r(\mathcal{L})$.

We note that the line segment sq is completely inside ℓ , due to the fact that ℓ is convex. Further, as a point t moves on cp from c to p , the distance between t and the line segment sq is nondecreasing. By Lemma D.2, the lines cp and sq have a small angle. Thus the angle between cp and ℓ is no more than $\arcsin(\sqrt{2 + \varepsilon}/(\alpha - 2\sqrt{2 + \varepsilon}))$. \square

Lemma D.4 *$C_\varepsilon(\mathcal{L})$ does not include a $(\Delta + 1)$ -dimensional ball with radius $\gamma\sqrt{\varepsilon}r(\mathcal{L})$ for $\gamma \geq (2 + \sqrt{\Delta + 1})\sqrt{2 + \varepsilon}$.*

Proof The proof is very similar to the proof of Lemma 4.6. Suppose that there is a $(\Delta + 1)$ -dimensional ball inside $C_\varepsilon(\mathcal{L})$. Denote by u the center of this ball. Again, we claim that $d(u, \ell) < r(\mathcal{L})$ for every $\ell \in \mathcal{L}$, which contradicts the definition of $r(\mathcal{L})$. Now suppose that there is a convex set $\ell \in \mathcal{L}$ such that $d(u, \ell) \geq r(\mathcal{L})$, we argue a contradiction. For any line j through u , when a point t moves away from u along j in one of the two directions, there must be a direction such that $d(t, \ell)$ is monotonically increasing. Thus we can find a set of $\Delta + 1$ orthogonal vectors $v_i \in C_\varepsilon(\mathcal{L})$ such that $d(t, \ell)$ is monotonically increasing when t moves from u along any of these vectors. Further the length of these vectors is at least $\gamma\sqrt{\varepsilon}r(\mathcal{L})$. By Lemma D.3, we can find $\Delta + 1$ orthogonal vectors of length $\gamma\sqrt{\varepsilon}r(\mathcal{L})$ centered at u such that the angle between each vector to the convex set ℓ is no more than $\arcsin\left(\frac{\sqrt{2+\varepsilon}}{\gamma-2\sqrt{2+\varepsilon}}\right) \leq \arcsin(1/\sqrt{\Delta+1})$. By considering any Δ -flat that contains ℓ , this contradicts Lemma D.1. \square

Finally we prove Theorem 5.2. We roughly follow the proof of Theorem 4.4. Let S be the set $\{p_0, p_1, \dots, p_{\Delta+1}\}$ of vertices that define a $(\Delta + 1)$ -dimensional simplex of maximum volume in $C_\varepsilon(\mathcal{L})$. For each vertex p_i of S , define the height r_i of p_i as the distance between p_i and the Δ -dimensional flat including $p_0, \dots, p_{i-1}, p_{i+1}, \dots, p_{\Delta+1}$. Let $r = \min_i r_i$ and assume that it is obtained at p_0 . On one hand, it holds that S (and thus $C_\varepsilon(\mathcal{L})$) includes¹ a $(\Delta + 1)$ -dimensional ball of radius $r/(\Delta + 2)$ centered at $p = \frac{1}{\Delta+2} \sum_{i=0}^{\Delta+1} p_i$. On the other, consider the Δ -dimensional flat V including the points $p_1, \dots, p_{\Delta+1}$. It must be the case that the Δ -slab with axis V and radius r contains $C_\varepsilon(\mathcal{L})$. Indeed, otherwise one could find a point $p'_0 \in C_\varepsilon(\mathcal{L})$ of distance greater than r from V implying that the simplex $S' = \{p'_0, p_1, \dots, p_{\Delta+1}\}$ is in $C_\varepsilon(\mathcal{L})$ and has volume greater than that of S (a contradiction). By Lemma D.4, this implies that $r \leq (\Delta + 2)\gamma\sqrt{\varepsilon}r(\mathcal{L})$ for γ as defined in Lemma D.4. By choosing an appropriate constant β we have $r \leq \beta\sqrt{(\Delta + 2)^3}r(\mathcal{L})$, as in the claim. \square

References

1. Agarwal, P.K., Procopiuc, C.M.: Approximation algorithms for projective clustering. In: SODA'00: Proceedings of the Eleventh Annual ACM–SIAM Symposium on Discrete Algorithms, Philadelphia, PA, USA, 2000, pp. 538–547. SIAM, Philadelphia (2000)
2. Agarwal, P.K., Procopiuc, C.M., Varadarajan, K.R.: Approximation algorithms for k -line center. In: ESA'02: Proceedings of the 10th Annual European Symposium on Algorithms, London, UK, 2002, pp. 54–63. Springer, Berlin (2002)
3. Agarwal, P.K., Procopiuc, C.M., Varadarajan, K.R.: A $(1 + \varepsilon)$ -approximation algorithm for 2-line-center. *Comput. Geom. Theory Appl.* **26**(2), 119–128 (2003)
4. Agarwal, P.K., Har-Peled, S., Varadarajan, K.R.: Approximating extent measures of points. *J. ACM* **51**(4), 606–635 (2004)
5. Agarwal, P., Har-Peled, S., Varadarajan, K.R.: Geometric approximation via coresets. In: *Current Trends in Combinatorial and Computational Geometry*. Cambridge University Press, Cambridge (2005)

¹Let V_i be the affine subspace defined by $p_0, \dots, p_{i-1}, p_{i+1}, p_{\Delta+1}$. The containment follows directly from the fact that $d(p_i, V_i) \geq r$ and thus $d(p, V_i) \geq \frac{1}{\Delta+2}r$. To prove the latter, assume w.l.o.g. (by rotation and translation) that V_i is a subspace and only the first $\Delta + 1$ coordinates of points in V_i may be nonzero.

6. Bádoiu, M., Clarkson, K.L.: Smaller core-sets for balls. In: SODA'03: Proceedings of the Fourteenth Annual ACM–SIAM Symposium on Discrete Algorithms, Philadelphia, PA, USA, 2003, pp. 801–802. SIAM, Philadelphia (2003)
7. Bádoiu, M., Har-Peled, S., Indyk, P.: Approximate clustering via core-sets. In: STOC'02: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, New York, NY, USA, 2002, pp. 250–257. Assoc. Comput. Mach., New York (2002)
8. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2003)
9. Gao, J., Langberg, M., Schulman, L.: Clustering lines: classification of incomplete data. Manuscript (2006)
10. Har-Peled, S.: Private communication
11. Har-Peled, S., Mazumdar, S.: On coresets for k -means and k -median clustering. In: STOC'04: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, New York, NY, USA, 2004, pp. 291–300. Assoc. Comput. Mach., New York (2004)
12. Har-Peled, S., Varadarajan, K.: Projective clustering in high dimensions using core-sets. In: SCG'02: Proceedings of the Eighteenth Annual Symposium on Computational Geometry, New York, NY, USA, 2002, pp. 312–318. Assoc. Comput. Mach., New York (2002)
13. Har-Peled, S., Wang, Y.: Shape fitting with outliers. *SIAM J. Comput.* **33**, 269–285 (2003)
14. Helly, E.: Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Dtsch. Math.-Ver.* **32**, 175–176 (1923)
15. Kumar, P., Yildirim, E.A.: Minimum volume enclosing ellipsoids and core sets. *J. Optim. Theory Appl.* **126**(1), 1–21 (2005)
16. Kumar, P., Mitchell, J.S.B., Yildirim, A.: Computing core-sets and approximate smallest enclosing hyperspheres in high dimensions. In: Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments, pp. 45–55 (2003)
17. Kumar, P., Mitchell, J.S.B., Yildirim, E.A.: Approximate minimum enclosing balls in high dimensions using core-sets. *J. Exp. Algorithmics* **8**, 1.1 (2003)
18. Little, R.J.A., Rubin, D.B.: Statistical Analysis with Missing Data. Wiley, New York (1986)
19. Matousek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. *Algorithmica* **16**(4/5), 498–516 (1996)