

Sparse Approximation via Generating Point Sets*

Avrim Blum[†]

Sariel Har-Peled[‡]

Benjamin Raichel[§]

October 25, 2017

For a set P of n points in the unit ball $\mathbf{b} \subseteq \mathbb{R}^d$, consider the problem of finding a small subset $T \subseteq P$ such that its convex-hull ε -approximates the convex-hull of the original set. Specifically, the Hausdorff distance between the convex hull of T and the convex hull of P should be at most ε . We present an efficient algorithm to compute such an ε' -approximation of size k_{alg} , where ε' is a function of ε , and k_{alg} is a function of the minimum size k_{opt} of such an ε -approximation. Surprisingly, there is no dependence on the dimension d in either of the bounds. Furthermore, every point of P can be ε -approximated by a convex-combination of points of T that is $O(1/\varepsilon^2)$ -sparse.

Our result can be viewed as a method for sparse, convex autoencoding: approximately representing the data in a compact way using sparse combinations of a small subset T of the original data. The new algorithm can be kernelized, and it preserves sparsity in the original input.

1. Introduction

Sparse approximation and coresets. Let P be a set of n points (observations) in the unit ball $\mathbf{b} \subseteq \mathbb{R}^d$, and let \mathcal{C}_P denote the convex-hull of P . Consider the problem of finding a small ε -coreset $T \subseteq P$ for projection width; that is, given any line ℓ in \mathbb{R}^d , consider the projections of \mathcal{C}_T and \mathcal{C}_P onto the line ℓ – these are two intervals $I_T \subseteq I_P$, and we require that $I_P \subseteq (1 + \varepsilon)I_T$. Such coresets have size $O(1/\varepsilon^{(d-1)/2})$, and lead to numerous efficient approximation algorithms in low-dimensions, see [AHV05]. In particular, such an ε -coreset guarantees that the Hausdorff distance between \mathcal{C}_T and \mathcal{C}_P is at most ε .

While such coresets can have size $\Omega(1/\varepsilon^{(d-1)/2})$ in the worst case, data may have structure allowing much smaller coresets to exist even in high dimensional spaces. For example, consider a dataset P in which all points are ε -close to one of k different lines. Then taking the extreme dataset points associated with each line results in $2k$ points, such that every $p \in P$ is 2ε -close to the convex hull of those points. More generally, the union of any two datasets which have good approximations of sizes k and k' , respectively, has one of size at most $k+k'$. Thus, it is natural to ask whether one can approximate the smallest such coreset, in terms of both its size and approximation quality.

*A preliminary version of this paper appeared in SODA 16 [BHR16]. The full version of the paper is also available on the arxiv [BHR15].

[†]Department of Computer Science, Carnegie Mellon University; avrim@cs.cmu.edu. Work was conducted while on sabbatical at the University of Illinois. Work on this paper was partially supported by NSF award CCF-1415460, and CCF-1525971.

[‡]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; <http://sarielhp.org/>. Work on this paper was partially supported by a NSF AF awards CCF-1421231, and CCF-1217462.

[§]Department of Computer Science; University of Texas at Dallas; Richardson, TX 75080, USA; benjamin.raichel@utdallas.edu; <http://utdallas.edu/~benjamin.raichel>. Work on this paper was partially supported by the University of Illinois Graduate College Dissertation Completion Fellowship.

Technique	$\varepsilon' \equiv d_H(\mathcal{C}_P, \mathcal{C}_T)$	$k_{\text{alg}} \equiv T $	Result
ε -nets	$\leq \varepsilon$	$O(dk_{\text{opt}} \log k_{\text{opt}})$	Lemma 3.2
Greedy set cover	$\leq (1 + \delta)\varepsilon$	$O((k_{\text{opt}}/(\varepsilon\delta)^2) \log n)$	Lemma 3.3
Greedy clustering	$\leq 8\varepsilon^{1/3} + \varepsilon$	$O(k_{\text{opt}}/\varepsilon^{2/3})$	Theorem 4.6 No dependency on d or n

Figure 1.1: Summary of our results: Given a set P contained in the unit ball of \mathbb{R}^d , such that there is a subset $P_{\text{opt}} \subseteq P$ of size k_{opt} , and $d_H(\mathcal{C}_P, \mathcal{C}_{P_{\text{opt}}}) \leq \varepsilon$, the above results compute an approximate set $T \subseteq P$. Note, that any point in P has an $O(1/\varepsilon^2)$ -sparse $(\varepsilon + \varepsilon')$ -approximation using T , because of the underlying sparsity – see **Lemma 2.6**.

The problem in matrix form. Given a collection P of n points (observations) in the unit ball $\mathbf{b} \subseteq \mathbb{R}^d$, viewed as column vectors, find a $d \times k$ matrix M such that each $p \in P$ can be approximately reconstructed as a *sparse, convex combination* of the columns of M . That is, for each $p \in P$ there exists a sparse non-negative vector x whose entries sum to one such that $p \approx Mx$. This problem is trivial if we allow $k = n$: simply make each data point $p \in P$ into a column of M , allowing the i th data point to be perfectly reconstructed using $x = e_i$, where e_i is the i th vector in the standard basis. The goal is to do so using $k \ll n$, so that M and the x 's can be viewed as an (approximate) compressed representation of the p 's.

Input assumption. We are given a set P of n points in \mathbb{R}^d all with norm at most one. Suppose that there *exists* a $d \times k_{\text{opt}}$ matrix M , such that

- (A) each column of M is a convex combination of the observations p , and
- (B) each $p \in P$ can be ε -approximately reconstructed as a convex combination of the columns of M : that is, for each $p \in P$ there exists a non-negative vector x whose entries sum to one such that $\|p - Mx\| \leq \varepsilon$.

Stated geometrically, the assumption is that the input P is contained in the unit ball \mathbf{b} (centered at the origin), and there exists a set $P_{\text{opt}} \subseteq \mathcal{C}_P$, of size k_{opt} , such that for any point $p \in P$, we have that p is ε -close to $\mathcal{C}_{P_{\text{opt}}}$, where $\mathcal{C}_{P_{\text{opt}}}$ denotes the convex-hull of P_{opt} . Formally, being ε -close means that the distance of p to the set $\mathcal{C}_{P_{\text{opt}}}$ is at most ε .

Our results. We present efficient algorithms for computing a $d \times k_{\text{alg}}$ matrix M' , consisting of k_{alg} points of P , such that each $p \in P$ can be ε' -approximately reconstructed as a sparse convex combination of the columns of M' , where k_{alg} and ε' are not too large, see **Figure 1.1** for details. Here, *sparse* means that only relatively few of the columns of M' would be used to represent (approximately) each point of *data*.

Stated in geometric terms, the algorithm computes a set T of k_{alg} points (these will be points from P) such that every point in P is ε' -close to the convex hull of T and moreover can be approximately reconstructed using a sparse convex combination of T .

The reader may notice that sparsity is not mentioned in the assumption about P_{opt} ($\equiv M$) and yet appears in the conclusion about T ($\equiv M'$). This is because convex combinations have the property that sparsity can be achieved almost for free, at the expense of a small amount of reconstruction error (see **Lemma 2.6**). This is to some extent the same reason that a large margin separator can be represented

using a small number of support vectors.

Related work. In comparison with the recent provable algorithms for autoencoding of Arora *et al.* [AGM14], our result does not require any distributional assumptions on the x 's or p 's, e.g., that the $p \in P$ were produced by choosing x from a particular distribution and then computing Mx and adding random noise. It also does not require that the columns of M be incoherent (nearly orthogonal). However, we do require that the columns of M be convex combinations of the points $p \in P$ and that they can approximately reconstruct the $p \in P$ via convex combinations, so our results are incomparable to those of Arora *et al.* [AGM14]. Work on related encoding or dictionary learning problems in the full rank case has been done by Spielman *et al.* [SWW12], and efficient algorithms for finding minimal and sparse Boolean representations under anchor-set assumptions were given by Balcan *et al.* [BBV15].

1.1. The results in detail

Our results are summarized in Figure 1.1.

- (A) **Sparse nearest-neighbor in high dimensions.** For a set of points P in the unit ball $\mathbf{b} \subseteq \mathbb{R}^d$ and any point of $p \in \mathcal{C}_P$, one can find a point $p' \in \mathcal{C}_P$ that is the convex combination of $O(1/\varepsilon^2)$ points of P , such that $\|p - p'\| \leq \varepsilon$. This is of course well known by now [Cla10], and we describe (for the sake of completeness) the surprisingly simple iterative algorithm (which is similar to the Perceptron algorithm) to compute such a representation in Section 2.2. This sparse representation is sometimes referred to as an approximate Carathéodory theorem [Bar15], and it also follows from the analysis of the Perceptron algorithm [Nov62] – see Remark 2.7.
- (B) **Geometric hitting set.** Our problem can be interpreted as (a somewhat convoluted) geometric hitting set problem. In particular, one can apply the Clarkson [Cla93] polytope approximation algorithm to this problem, thus yielding an $O(d \log k_{\text{opt}})$ approximation. For the sake of completeness, we describe this in detail in Section 3.1. (Since d might be large, this approximation is somewhat less attractive.)
- (C) **The greedy approach.** A natural approach is to try and solve the problem using the greedy algorithm. Here, this requires some work, and the resulting algorithm is a combination of the algorithm from (A) with greedy set cover for the ranges defined in (B). We initialize an instance of the algorithm from (A) for each point $p \in P$ whose job is to either find a hyperplane through p separating it from $P \setminus \{p\}$ by a large margin or else to approximate p as a combination of a few support-vectors in $P \setminus \{p\}$. At each step, we find the point $p' \in P$ that causes as many of these algorithms to perform an update as possible, and add it into our set T . The key issue is to prove that the procedure halts after a limited number of steps. This algorithm is described in Section 3.2.
- (D) **Using greedy clustering.** The second algorithm, and our main contribution, is more similar in spirit to the Gonzalez algorithm for k -center clustering: Repeatedly find the point $p \in P$ that is farthest from the convex hull of the points of T and then add it into T if this distance is greater than some threshold (a similar idea was used for subspace approximation [HV04, Lemma 5.2]). The key issue here is to prove that some measure of significant progress is made each time a new point is added. Somewhat surprisingly, after $O(k_{\text{opt}}/\varepsilon^{2/3})$ iterations, the resulting set is an $O(\varepsilon^{1/3})$ -approximation to the original set of points. Note, that unlike the other results mentioned above, there is no dependence on the dimension or the input size.

An additional property of all the above algorithms is that the points T found will be actual dataset points and the algorithms only require dot-product access to the data. This means that the algorithms

can be kernelized. Additionally, much as with CUR decompositions of matrices, since the points T are data points, they will preserve sparsity if the dataset P was sparse.

2. Preliminaries

For a set $X \subseteq \mathbb{R}^d$, \mathcal{C}_X denotes the *convex hull* of X . For two sets $P, P' \subseteq \mathbb{R}^d$, we denote by $d(P, P') = \min_{p \in P} \min_{p' \in P'} \|p - p'\|$ the *distance* between P and P' . For a point $q \in \mathbb{R}^d$, its distance to the set P is $d(q, P) = d(\{q\}, P)$, and its *projection* or *nearest neighbor* in P is the point $\text{nn}(q, P) = \arg \min_{p \in P} \|q - p\|$.

2.1. Sparse convex-approximation: Problem statement and background

For a set Y in \mathbb{R}^d , its *one sided Hausdorff distance* from X is $d(Y \rightarrow X) = \max_{y \in Y} d(y, X)$.

Definition 2.1. Consider two sets $P_{\text{in}}, P_{\text{out}} \subseteq \mathbb{R}^d$. A set $U \subseteq \mathcal{C}_{P_{\text{out}}}$ is a **δ -approximation** to P_{in} from P_{out} if $d(\mathcal{C}_{P_{\text{in}}} \rightarrow \mathcal{C}_U) \leq \delta$. In words, every point of $\mathcal{C}_{P_{\text{in}}}$ is within distance δ from a point of \mathcal{C}_U . In the **discrete δ -approximation** version, we require that $U \subseteq P_{\text{out}}$. We use $\text{opt}(P_{\text{in}}, P_{\text{out}}, \delta)$ to denote any minimum cardinality discrete δ -approximation to P_{in} from P_{out} , and $k_{\text{opt}}(P_{\text{in}}, P_{\text{out}}, \delta) = |\text{opt}(P_{\text{in}}, P_{\text{out}}, \delta)|$ to denote its size. We drop the phrase “from P_{out} ” when it is clear from the context.

Problem 2.2. Given sets $P_{\text{in}}, P_{\text{out}} \subseteq \mathbb{R}^d$, compute (or approximate) $\text{opt}(P_{\text{in}}, P_{\text{out}}, \delta)$.

For the majority of the paper we focus on the natural special case when $P = P_{\text{in}} = P_{\text{out}}$. The **Hausdorff distance** between sets X and Y is defined as $d_H(X, Y) = \max(d(Y \rightarrow X), d(X \rightarrow Y))$.

Lemma 2.3. (i) Let C be a convex-set in \mathbb{R}^d , then the function $f(p) = d(p, C)$ is convex, where $p \in \mathbb{R}^d$.

(ii) A convex-function f , over a convex bounded domain $D \subseteq \mathbb{R}^d$, attains its maximum in a boundary point of D .

(iii) For bounded point sets $U, P \subseteq \mathbb{R}^d$, such that $U \subseteq \mathcal{C}_P$, we have $d_H(\mathcal{C}_U, \mathcal{C}_P) = d(P \rightarrow \mathcal{C}_U)$.

Proof: This is all well known, and we include the proof for the sake of completeness.

(i) Consider any two points p, y in \mathbb{R}^d , and let $p' = \text{nn}(p, C)$ and $y' = \text{nn}(y, C)$. For any $t \in [0, 1]$, we have by convexity that $z = tp + (1 - t)y \in py$ and $z' = tp' + (1 - t)y' \in C$. Therefore, by the triangle inequality, we have

$$\begin{aligned} f(z) &= f(tp + (1 - t)y) \leq \|z - z'\| = \|(tp + (1 - t)y) - (tp' + (1 - t)y')\| \\ &= \|t(p - p') + (1 - t)(y - y')\| \leq \|t(p - p')\| + \|(1 - t)(y - y')\| \\ &= t\|p - p'\| + (1 - t)\|y - y'\| = tf(p) + (1 - t)f(y). \end{aligned}$$

(ii) If p is the interior of D then there are extremal points p_1, \dots, p_d of D , and constants $\alpha_1, \dots, \alpha_d \in [0, 1]$, such that $\sum_i \alpha_i = 1$ and $p = \sum_i \alpha_i p_i$. As such, by convexity, we have $f(p) = f(\sum_i \alpha_i p_i) \leq \sum_i \alpha_i f(p_i) \leq \max_i f(p_i)$.

(iii) By (i), the function $d(p, \mathcal{C}_U)$ is convex. By (ii), its maximum over \mathcal{C}_P is attained at a point of P . We thus have that

$$\begin{aligned} d_H(\mathcal{C}_U, \mathcal{C}_P) &= \max(d(\mathcal{C}_U \rightarrow \mathcal{C}_P), d(\mathcal{C}_P \rightarrow \mathcal{C}_U)) = d(\mathcal{C}_P \rightarrow \mathcal{C}_U) = \max_{p \in \mathcal{C}_P} d(p, \mathcal{C}_U) = \max_{p \in P} d(p, \mathcal{C}_U) \\ &= d(P \rightarrow \mathcal{C}_U). \end{aligned} \quad \blacksquare$$

Definition 2.4. Consider any set $P \subseteq \mathbb{R}^d$. A set $U \subseteq \mathcal{C}_P$ is a **δ -approximation** to P if $d_H(\mathcal{C}_U, \mathcal{C}_P) \leq \delta$. By the above lemma, this is equivalent to every point of P being in distance at most δ from a point of \mathcal{C}_U . In the **discrete δ -approximation** version, we require that $U \subseteq P$. Let $\text{opt}(P, \delta)$ be any minimum cardinality δ -approximation to P , and let $k_{\text{opt}}(P, \delta) = |\text{opt}(P, \delta)|$ denote its size.

Problem 2.5. Given a set $P \subseteq \mathbb{R}^d$ and value δ , compute (or approximate) $\text{opt}(P, \delta)$.

Example. Consider a unit radius sphere $\mathbb{S}^{(d-1)}$ in \mathbb{R}^d centered at the origin, and let P be a δ' -packing on $\mathbb{S}^{(d-1)}$ (i.e., every point in $\mathbb{S}^{(d-1)}$ is at distance at most δ' from a point of P , and any two points of P are at distance at least δ' from each other). It is easy to verify that such a δ' -packing has size $\Theta(1/(\delta')^{d-1})$. Furthermore, for any $\delta > 0$, and an appropriate absolute constant c (independent of the dimension or δ), setting $\delta' = c\sqrt{\delta}$, we have the property that for any point $p \in P$, $d(p, \mathcal{C}_{P \setminus \{p\}}) > \delta$. That is, any δ -approximation to P requires $\Omega(1/\delta^{(d-1)/2})$ points.

On the other hand, let $P_{\text{out}} = \{\pm de_i \mid i = 1, \dots, d\}$, where e_i denotes the i th orthonormal vector, having zero in all coordinates except for the i th coordinate where it is 1. Clearly, $\mathbb{S}^{(d-1)} \subseteq \mathcal{C}_{P_{\text{out}}}$, and as such $k_{\text{opt}}(P, P_{\text{out}}, \delta) \leq |P_{\text{out}}| = 2d$, with equality for $\delta = 0$.

Throughout this paper we require that P_{out} be contained in the unit ball, disallowing this latter type of “trivial” solution, and furthermore having the property that a successful approximation also yields a sparse solution essentially for free, as shown next in [Lemma 2.6](#).

2.2. Computing the approximate distance to the convex hull

The following is well known, and is included for the sake of completeness, see [\[HKMR15\]](#). It also follows readily from the Preceptron algorithm (see [Remark 2.7](#) below).

Lemma 2.6. Let $P \subseteq \mathbb{R}^d$ be a point set, $\varepsilon > 0$ be a parameter, and let $q \in \mathbb{R}^d$ be a given query point. Then, one can compute, in $O(|P|d/\varepsilon^2)$ time, a point $t \in \mathcal{C}_P$, such that $\|q - t\| \leq d(q, \mathcal{C}_P) + \varepsilon\Delta$, where $\Delta = \text{diam}(P)$. Furthermore, t is a convex combination of $O(1/\varepsilon^2)$ points of P .


Proof: The algorithm is iterative, computing a sequence of points t_0, \dots, t_i inside \mathcal{C}_P that approach q . Initially, $p_0 = t_0$ is the closest point of P to q . In the i th iteration, the algorithm computes the vector $v_i = q - t_{i-1}$, and the point $p_i \in P$ that is extremal in the direction of v_i . Now, the algorithm sets t_i to be the closest point to q on the segment $s_i = t_{i-1}p_i$, and continues to the next iteration, for $M = O(1/\varepsilon^2)$ iterations. The algorithm returns the point t_M as the desired answer.

By induction, the point $t_i \in \mathcal{C}_{\{p_0, \dots, p_i\}}$. Furthermore, observe that the distance of the points t_0, t_1, \dots from q is monotonically decreasing. In particular, for all $i > 0$, t_i must fall in the middle of the segment s_i , as otherwise, p_i would be closer to q than p_0 , a contradiction to the definition of p_0 .

Project the point p_i to the segment $t_{i-1}q$, and let y_i be the projected point. Observe that $\|q - y_i\|$ is a lower bound on $d(q, \mathcal{C}_P)$. Therefore, if $\|y_i - t_{i-1}\| \leq \varepsilon\Delta$ then we are done, as $\|q - t_{i-1}\| \leq \|t_{i-1} - y_i\| + \|y_i - q\| \leq \varepsilon\Delta + d(q, \mathcal{C}_P)$. (In particular, one can use this as alternative stopping condition for the algorithm, instead of counting iterations.)

So, let α be the angle $\angle p_i t_{i-1} q$. Observe that as $t_{i-1}p_i \subseteq \mathcal{C}_P$, it follows that $\|t_{i-1} - p_i\| \leq \text{diam}(P) = \Delta$. Furthermore, $\cos \alpha = \frac{\|y_i - t_{i-1}\|}{\|t_{i-1} - p_i\|} > \frac{\varepsilon\Delta}{\Delta} = \varepsilon$, since $\|y_i - t_{i-1}\| > \varepsilon\Delta$. Hence, $\sin \alpha = \sqrt{1 - \cos^2 \alpha} \leq \sqrt{1 - \varepsilon^2} \leq 1 - \varepsilon^2/2$. Let $\ell_{i-1} = \|q - t_{i-1}\|$. We have that

$$\ell_i = \|q - t_i\| = \|q - t_{i-1}\| \sin \alpha \leq (1 - \varepsilon^2/2)\ell_{i-1}.$$



Analyzing the number of iterations required by the algorithm is somewhat tedious. If $\ell_0 = \|q - t_0\| \geq (4/\varepsilon^2)\Delta$ then the algorithm would be done in one iteration as otherwise $\ell_1 \leq \ell_0 - 2\Delta$, which is impossible. In particular, after $4/\varepsilon^2$ iterations the distance ℓ_i shrinks by a factor of two, and as such, after $O((1/\varepsilon^2) \log(1/\varepsilon))$ iterations the algorithm is done.

One can do somewhat better. By the above, we can assume that $d(q, P) = O(\Delta/\varepsilon^2)$. Now, set $\varepsilon_j = 1/2^{2+j}$. By the above, after $n_0 = O((1/\varepsilon_0^2) \log(1/\varepsilon_0)) = O(1)$ iterations, $\ell_{n_0} \leq d(q, \mathcal{C}_P) + \text{diam}(P)/4$. For $j \geq 1$, let $n_j = 4/(\varepsilon_j)^2$, and observe that, after $\nu_j = n_j + \sum_{k=0}^{j-1} n_k$ iterations, we have that

$$\ell_{\nu_j} \leq (d(q, \mathcal{C}_P) + \varepsilon_{j-1}\Delta)/2 \leq d(q, \mathcal{C}_P) + \varepsilon_j\Delta.$$

In particular, stopping as soon as $\varepsilon_j \leq \varepsilon$, we have the desired guarantee, and the number of iterations needed is $M = O(1) + \sum_{j=0}^{\lceil \lg 1/\varepsilon \rceil} 4/\varepsilon_j^2 = O(1/\varepsilon^2)$. ■

In our use of [Lemma 2.6](#), P and q will always be contained in the unit ball, so we can remove the Δ term in the bound if we wish since $\Delta \leq 2$.

Remark 2.7. [Lemma 2.6](#) is known, and a variant of it follows readily from a result (from 1962) on the convergence of the Perceptron algorithm [[Nov62](#)]. Indeed, consider a set $P \subseteq \mathbb{R}^d$ and a query point $q \in \mathbb{R}^d$. Assume that $q \in \mathcal{C}_P$, and furthermore that q is the origin (translating space if needed to ensure this). Run the Perceptron algorithm learning a linear classifier that passes through the origin and classifies P as positive examples. Stop the algorithm after $M = 1/\varepsilon^2$ classification mistakes (since $q \in \mathcal{C}_P$, there will always be a mistake in P). Let p_1, \dots, p_M be the sequence of points on which mistakes were made and let $w = p_1 + \dots + p_M$ be the resulting hypothesis vector. By the analysis of [[Nov62](#)], we have $\|w\| \leq \text{diam}(P)\sqrt{M}$. This implies that the point $p' = w/M$, which is a convex combination of the points p_1, \dots, p_M , has length—and therefore distance from q —at most $\varepsilon \text{diam}(P)$.

Thus, we conclude that for any point $p \in \mathcal{C}_P$, and any $\varepsilon \in (0, 1)$, there is a point $p' \in \mathcal{C}_U$, which is a convex combination of $O(1/\varepsilon^2)$ points of P , such that $\|p - p'\| \leq \varepsilon \text{diam}(P)$. This is sometimes referred to as approximate Carathéodory theorem [[Bar15](#)].

We described the alternative algorithm (in the proof of [Lemma 2.6](#)) because it is more direct and slightly simpler in this case.


3. Approximations via hitting set algorithms

Here we look at two hitting set type algorithms for [Problem 2.2](#). An (α, β) -**approximation** of $\text{opt}(P_{\text{in}}, P_{\text{out}}, \varepsilon)$ is a set $U \subseteq P_{\text{out}}$ such that $d(\mathcal{C}_{P_{\text{in}}} \rightarrow \mathcal{C}_U) \leq \alpha$ and $|U| \leq \beta k_{\text{opt}}(P_{\text{in}}, P_{\text{out}}, \varepsilon)$, see [Definition 2.1](#).

As a warm-up exercise, we first present an $(\varepsilon, O(d \log k_{\text{opt}}))$ -approximation using approximation algorithms for hitting sets for set systems with bounded VC dimension. Then, we build on that to get a greedy algorithm providing a $((1 + \delta)\varepsilon, O((\varepsilon\delta)^{-2} \log n))$ -approximation.

3.1. Approximation via VC dimension

Definition 3.1. For a set $P \subseteq \mathbb{R}^d$ and a direction vector v , let p be the point of P extreme in the direction of v , and let h' be the hyperplane with normal v and tangent to \mathcal{C}_P at p . For a parameter ε , let h be the hyperplane formed by translating h' distance ε in the direction $-v$. The ε -**shadow** of h' (or v), is the halfspace $h^+(P, \varepsilon, v)$ bounded by h that contains p in its interior. In words, the ε -shadow of v is the outer supporting halfspace for P with a normal in the direction of v , translated in by distance ε .



Lemma 3.2. *Given sets P_{in} and P_{out} in \mathbb{R}^d with a total of n points, and a parameter ε , one can compute a $(\varepsilon, O(d \log k_{\text{opt}}))$ -approximation to the optimal discrete set $\text{opt}(P_{\text{in}}, P_{\text{out}}, \varepsilon)$ in polynomial time.*

Proof: For a direction v , consider the hyperplane h' tangent to $\mathcal{C}_{P_{\text{in}}}$ at an extremal point $p_v \in P_{\text{in}}$ in the direction of v , and its ε -shadow $h^+ = h^+(P_{\text{in}}, \varepsilon, v)$, see [Figure 3.1](#).

Clearly, any discrete ε -approximation $U \subseteq P_{\text{out}}$ to P_{in} , must contain at least one point of $P_{\text{out}} \cap h^+$, as otherwise the approximation fails for the point p_v (in particular, if such a halfspace has no point in P_{out} then there is no approximation). Now, consider the set system

$$\mathcal{S} = \left(P_{\text{out}}, \{P_{\text{out}} \cap h^+(P_{\text{in}}, \varepsilon, v) \mid v \text{ any unit vector}\} \right).$$

This set system has VC dimension at most $d + 1$, and in particular, for such a set system one can compute a $O(d \log k_{\text{opt}})$ approximation to its minimum size hitting set, which is the desired approximation in this case, see [[Har11](#), Section 6.3]. We describe the algorithm below, but first we verify that this indeed yields the desired approximation.

Consider a hitting set $U \subseteq P_{\text{out}}$ of \mathcal{S} . Let p be any point in $\mathcal{C}_{P_{\text{in}}}$, and let p' be the closest point to p in \mathcal{C}_U . If $\|p - p'\| \leq \varepsilon$, then we are done. Otherwise, consider the vector $v = p - p'$. Let z denote the hyperplane whose normal is v and which passes through the point p' , and let z^+ denote the open halfspace bounded by z and in the direction of v (i.e. containing p). As p' is the closest point to p in \mathcal{C}_U , z^+ has empty intersection with \mathcal{C}_U . Moreover, $h^+(P_{\text{in}}, \varepsilon, v) \subsetneq z^+$, as the bounding hyperplanes of both halfspaces have v as a normal, and the extreme point of $\mathcal{C}_{P_{\text{in}}}$ in the direction of v must be $> \varepsilon$ away from z (as p is at least this far in the direction of v). See [Figure 3.2](#). These two facts combined imply $h^+(P_{\text{in}}, \varepsilon, v) \cap \mathcal{C}_U = \emptyset$, a contradiction as $h^+(P_{\text{in}}, \varepsilon, v) \cap P_{\text{out}}$ is a set in \mathcal{S} that should have been hit.

As for the algorithm, Clarkson [[Cla93](#)] described how to compute this set via reweighting, but the following technique due to Long [[Lon01](#)] is easier to describe (we sketch it here for the sake of completeness). Consider the LP relaxation of the hitting set for this set system. Clearly, one can assign weights to points (between 0 and 1), such that the total weight of the points is at most k_{opt} , and for every range in \mathcal{S} the total weight of the points it covers is at least 1. Dividing this fractional solution by k_{opt} , we get a weighted set system, where every set has weight at least $\eta = 1/k_{\text{opt}}$, and total weight of the points is 1. That is, we can interpret these weights over the points as a measure, where all the sets of interests are η -heavy. A random sample of size $O((d/\eta) \log(1/\eta)) = O(k_{\text{opt}} d \log k_{\text{opt}})$ of P (according to the weights) is an η -net with constant probability [[HW87](#)], and stabs all the sets of \mathcal{S} , as desired. Should the random sample fail, one can sample again till success. \blacksquare

3.2. Approximation via a greedy algorithm

Lemma 3.3. *Let P_{in} and P_{out} be sets of points in \mathbb{R}^d contained in the unit ball, with a total of n points. For parameters $\varepsilon, \delta \in (0, 1)$, one can compute, in polynomial time, a $((1 + \delta)\varepsilon, O(\varepsilon^{-2} \delta^{-2} \log n))$ -approximation to the optimal discrete set $\text{opt}(P_{\text{in}}, P_{\text{out}}, \varepsilon)$.*

Proof: The algorithm is greedy – the basic idea is to restrict the set system of [Lemma 3.2](#) to the relevant active sets. Formally, let $U_0 = \{p_0\}$, where p_0 is some arbitrary point of P_{out} . For $i > 0$, in the i th




Figure 3.1: Circles and squares denote points of P_{in} and P_{out} , respectively.




Figure 3.2

iteration, consider the current convex set $C_{i-1} = \mathcal{C}_{U_{i-1}}$. For a point $q \in P_{\text{in}} \setminus C_{i-1}$, let $\text{nn}(q, C_{i-1})$ be its nearest point in C_{i-1} , and let $v_i(q)$ be the direction of the vector $q - \text{nn}(q, C_{i-1})$. In particular, consider the ε -shadow halfspace $h^+ = h^+(P_{\text{in}}, \varepsilon, v_i(q))$, see [Definition 3.1](#), which should be hit by the desired hitting set¹.

Let $Z_i \subseteq P_{\text{in}}$ be the set of points of P_{in} that are *unhappy*; that is, they are in distance $\geq (1 + \delta)\varepsilon$ from $\mathcal{C}_{U_{i-1}}$. We restrict our attention to the set system of active halfspaces; that is,

$$\mathcal{S}_i = \left(P_{\text{out}}, \{ P_{\text{out}} \cap h^+(P_{\text{in}}, \varepsilon, v_i(q)) \mid q \in Z_i \} \right).$$

(As before, if $P_{\text{out}} \cap h^+$ is empty, then no approximation is possible, and the algorithm is done.) Now, as in the classical algorithm for hitting set (or set cover), pick the point p_i in P_{out} that hits the largest number of ranges in \mathcal{S}_i , and add it to U_{i-1} to form U_i .

A point $q \in Z_i$, is *hit* in the i th iteration if $p_i \in h^+(P, \varepsilon, v_i(q))$, see [Figure 3.3](#). The argument of [Lemma 2.6](#) (or [Remark 2.7](#)) implies that after a point $q \in P_{\text{in}}$ is hit $c/(\varepsilon^2\delta^2)$ times, its distance to the convex-hull of the current points is smaller than $(1 + \delta)\varepsilon$, and it is no longer unhappy, where c is some sufficiently large constant. Indeed, using the notation of the proof [Lemma 2.6](#), if a point $q \in Z_i$ is hit in the i th iteration by a point p_i , and $d(q, \mathcal{C}_{U_{i-1}}) \leq (1 + \delta)\varepsilon$ then we are done. Otherwise, let $t_{i-1} = \text{nn}(q, \mathcal{C}_{U_{i-1}})$, and let y_i be the projection of p_i to the segment qt_{i-1} , see [Figure 2.1](#). We have that $\|y_i - t_{i-1}\| \geq \|q - t_{i-1}\| - \|q - y_i\| \geq (1 + \delta)\varepsilon - \varepsilon \geq \varepsilon\delta$, since $\|q - y_i\| \leq \varepsilon$ (as p_i and y_i are both in the ε -shadow of q). Now, the analysis of [Lemma 2.6](#) applies (with $\varepsilon\delta$ instead of ε), implying that after $O(1/(\varepsilon\delta)^2)$ iterations, the distance of q from the current convex-hull would be smaller than $(1 + \delta)\varepsilon$.




Figure 3.3

So, let n_i be the number of unhappy points in the beginning of the i th iteration, and observe that at least n_i/k_{opt} points are being hit in the i th iteration. In particular, let $\kappa = 2 \lceil ck_{\text{opt}}/(\varepsilon^2\delta^2) \rceil$, and observe that in the iterations between $i - \kappa$ and i , we have that the number of points being hit is at least $\sum_{j=i-\kappa}^i n_j/k_{\text{opt}} \geq 2n_i c/(\varepsilon^2\delta^2)$. This implies that $n_{i-\kappa} \geq 2n_i$. Otherwise, $n_{i-\kappa} < 2n_i$, implying that in this range of iterations $> N = n_{i-\kappa}c/(\varepsilon^2\delta^2)$ hits happened, which is impossible, as $n_{i-\kappa}$ points can be hit at most N times before they are all happy.

As such, after κ iterations of the greedy algorithm, the number of unhappy points drops by a factor of two, and we conclude that after $O(k_{\text{opt}}(\varepsilon\delta)^{-2} \log n)$ total iterations, the algorithm is done. \blacksquare

4. Approximating the convex hull in high dimensions

Here we provide an efficient bi-criteria approximation algorithm for [Problem 2.5](#). That is, the algorithm computes a subset $U \subseteq \mathcal{C}_P$, such that (i) $d_H(\mathcal{C}_U, \mathcal{C}_P) \leq O(\varepsilon^{1/3})\text{diam}(P)$, and (ii) $|U| \leq O(k_{\text{opt}}(P, \varepsilon)/\varepsilon^{2/3})$. Significantly, the computed set U is actually a subset of P , implying that the algorithm simultaneously solves both the continuous and discrete variants of the problem.

To simplify the presentation, in the remainder of this section we assume $\Delta = \text{diam}(P) = O(1)$, and hence drop most appearances of Δ .

¹The hitting set computed by the algorithm is somewhat weaker, only hitting all the $(1 + \delta)\varepsilon$ -shadows.

4.1. The algorithm

Let $\delta = 8\varepsilon^{1/3}$. The algorithm is greedy, similar in spirit to the Gonzalez algorithm for k -center clustering [Gon85] and subspace approximation algorithms [HV04, Lemma 5.2]. The algorithm starts with an arbitrary point $t_0 \in P$. For $i > 0$, in the i th iteration, the algorithm computes the point t_i in P which is furthest away from $\mathcal{C}_{U_{i-1}}$, where $U_{i-1} = \{t_0, \dots, t_{i-1}\}$. For now assume these distance queries are done exactly – later on we describe how to use approximate queries (i.e., Lemma 2.6). Let $r_i = d(t_i, \mathcal{C}_{U_{i-1}})$. The algorithm stops as soon as $r_i \leq \delta$, and outputs U_{i-1} .

Observation 4.1. *In the above algorithm, for all $i > 0$, the point t_i is a vertex of \mathcal{C}_P (so long as exact distance queries are used). In particular, if the output has to be a subset of the convex hull vertices, one can choose t_0 to be the extreme vertex in any direction.*

4.2. Analysis

By the termination condition of the algorithm, when the algorithm stops every point in P is in distance at most $\delta = 8\varepsilon^{1/3}$ away from $\mathcal{C}_{U_{i-1}}$, as desired. As for the number of rounds until termination, we argue that in each round there exists some point $o \in P_{\text{opt}}$ which is far from $\mathcal{C}_{U_{i-1}}$ (as specified in Claim 4.2) and such that $d(o, U_i) \leq (1 - \Omega(\varepsilon^{2/3}))d(o, U_{i-1})$.

So consider some round i , the current set U_{i-1} , and the point $t_i \in P$ furthest away from $\mathcal{C}_{U_{i-1}}$. Let t'_i be the closest point to t_i in $\mathcal{C}_{U_{i-1}}$, and let $r_i = \|t_i - t'_i\|$. Let h_i be the hyperplane orthogonal to the segment $t_i t'_i$ and lying ε distance below t_i in the direction of t'_i . Let h_i^+ denote the closed halfspace having h_i as its boundary, and that contains t_i , see Figure 4.1. If no points of P_{opt} are in h_i^+ then $d(t_i, \mathcal{C}_{P_{\text{opt}}}) > \varepsilon$, which is impossible. Therefore, there must be a point $o_i \in P_{\text{opt}} \cap h_i^+$. Let o'_i be the closest point to o_i in $\mathcal{C}_{U_{i-1}}$.




Figure 4.1

Claim 4.2. $r_i - \varepsilon \leq \|o_i - o'_i\| \leq r_i$.

Proof: Let h'_i be the translation of h_i so it passes through t'_i , see Figure 4.1. We have that $r_i - \varepsilon = d(h'_i, h_i) \leq \|o_i - o'_i\|$, as o_i lies in h_i^+ (i.e., above h_i) and all of $\mathcal{C}_{U_{i-1}}$ lies below h'_i .

For the second part, for any $p \in \mathbb{R}^d$, let $f_{i-1}(p)$ be the distance of p from $\mathcal{C}_{U_{i-1}}$. By Lemma 2.3 (iii), and since $o_i \in P_{\text{opt}} \subseteq \mathcal{C}_P$, it follows that $\|o_i - o'_i\| \leq \max_{p \in \mathcal{C}_P} f_{i-1}(p) = \|t_i - t'_i\| = r_i$. ■

Lemma 4.3. *If $r_i \geq 8\varepsilon^{1/3}$ then $d(o_i, \mathcal{C}_{U_i}) \leq (1 - \varepsilon^{2/3})d(o_i, \mathcal{C}_{U_{i-1}})$.*

Proof: In the following, all entities are defined in the context of the i th iteration, and we omit the subscript i denoting this to simplify the exposition. Assume, for the time being, that the angle $\angle t t' o'$ is a right angle and the segment $t' o'$ has length $\ell = 1$, see Figure 4.2. This is the worst case configuration in terms of the new convex-hull \mathcal{C}_{U_i} getting closer to o , as can be easily seen.

Let z be the intersection of h with the ray emanating from o' in the direction $t - t'$. Let z' be the closest point to z on $o't$, let $\tau = \|z - z'\|$, and let ρ be the radius of the ball formed by $\text{ball}(o', r) \cap h$. See [Figure 4.2](#).




Figure 4.2

Rather than bounding the distance of o to \mathcal{C}_{U_i} directly, instead we use bounds on ρ and τ . Observe that $o \in h^+ \cap \text{ball}(o', r) \subseteq \text{ball}(z, \rho)$, and as such,

$$\|o - z\| \leq \rho. \text{ Now, we have } \rho = \sqrt{r^2 - \|z - o'\|^2} = \sqrt{r^2 - (r - \varepsilon)^2} = \sqrt{2r\varepsilon - \varepsilon^2} \leq \sqrt{2r\varepsilon}.$$

Let $\alpha = \angle zo't$ and $\beta = \pi/2 - \alpha = \angle to't'$, and observe that $\sin \alpha = \cos \beta = \ell / \sqrt{\ell^2 + r^2}$, where $\ell = \|o' - t'\| = 1$. Now, we have

$$\frac{\tau}{r - \varepsilon} = \sin \alpha = \frac{\ell}{\sqrt{\ell^2 + r^2}} = \frac{1}{\sqrt{1 + r^2}} \leq \sqrt{1 - \frac{r^2}{2}} \leq 1 - \frac{r^2}{4} \quad (4.1)$$

since $\ell = 1$ and $r \leq 1$.





Figure 4.3: Note, that o is not necessarily in the two dimensional plane depicted by the figure. All other points are in this plane.

Sanity condition: Consider the line which is the intersection of the hyperplane h and the two dimensional plane spanned by t, t' and o' (this line is denoted by h in the figures). Let u be the point in distance ρ on this line from z , on the side further away from t . Let t'' be the intersection of h with to' . Next, let u' be the nearest point to u on the segment to' , see [Figure 4.3](#).

We want to argue the distance between o and \mathcal{C}_{U_i} , can be bounded in terms of the distance between u and u' , however to do so we need to guarantee that u' is in the interior of this segment to' . Setting $\ell' = \|z - t''\|$, this happens if

$$\begin{aligned} \|u' - t''\| < \|t'' - o'\| &\iff \|u' - t''\| = (\rho + \ell') \cos \beta = (\rho + \ell') \frac{\ell'}{\|t'' - o'\|} < \|t'' - o'\| \\ &\iff (\rho + \ell') \ell' < \|t'' - o'\|^2 = (\ell')^2 + (r - \varepsilon)^2. \end{aligned}$$

Thus, we have to prove that $\rho \ell' < (r - \varepsilon)^2$. As $\ell' < \ell = 1$, we have that this is implied if $\rho \leq \sqrt{2r\varepsilon} < (r - \varepsilon)^2$, and this inequality holds if $r \geq 8\varepsilon^{1/3}$.



Back to the proof: We next bound the distance of o from \mathcal{C}_{U_i} . Observe that by rotating o around the line $o't$ we can assume that o lies on the plane spanned by t, t', o' and its distance to the segment to' has not changed. Now, the set of points in distance r' from the segment $o't$ is a hippodrome, and this hippodrome covers a connected portion of ball(o', r). For $r' = \|u - u'\|$, by the above sanity condition, this hippodrome covers all the points of ball(o', r) that are above h . This implies that o maximizes its distance to \mathcal{C}_{U_i} if $o = u$.

So, let $\tau' = \|u - u'\|$. By the above sanity condition the segment to' and uu' meet at a right angle, and hence by similar triangles (see Figure 4.3), we have

$$\tau' = \frac{\ell' + \rho}{\ell'}\tau = \tau + \rho\frac{\tau}{\ell'} = \tau + \rho\sin\beta = \tau + \rho\frac{r}{\sqrt{\ell^2 + r^2}} = \tau + \rho\frac{r}{\sqrt{1 + r^2}} \leq \tau + \rho r.$$

This implies, by Eq. (4.1), that

$$\frac{d(o, \mathcal{C}_{U_i})}{d(o, \mathcal{C}_{U_{i-1}})} \leq \frac{\|u - u'\|}{\|z - o'\|} = \frac{\tau'}{r - \varepsilon} \leq \frac{\tau}{r - \varepsilon} + \frac{\rho r}{r - \varepsilon} \leq \frac{\tau}{r - \varepsilon} + 2\rho \leq 1 - \frac{r^2}{4} + 2\sqrt{r\varepsilon} \leq 1 - \varepsilon^{2/3},$$

if $r \geq 8\varepsilon^{1/3}$. ■

Lemma 4.4. *Let P be a set of n points in \mathbb{R}^d with diameter $\Delta = \text{diam}(P)$, and let $\varepsilon > 0$ be a parameter, then one can compute a set $U \subseteq P$, such that*

- (i) $d_H(\mathcal{C}_U, \mathcal{C}_P) \leq (8\varepsilon^{1/3} + \varepsilon)\Delta$,
- (ii) $m = |U| \leq O(k_{\text{opt}}/\varepsilon^{2/3})$, where $k_{\text{opt}} = k_{\text{opt}}(P, \varepsilon)$, and
- (iii) the running time is $O(nm^2d/\varepsilon^2)$.

Proof: Recall that in any round before the algorithm terminates $r_i > \delta\Delta = 8\varepsilon^{1/3}\Delta$. Let $P_{\text{opt}} = \text{opt}(P, \varepsilon)$ be any optimal approximating set of size k_{opt} . In the i th iteration of the algorithm, for some point $o_i \in P_{\text{opt}}$, its distance to the convex hull of \mathcal{C}_{U_i} shrinks by a factor of $1 - \varepsilon^{2/3}$, by Lemma 4.3. Conceptually, we charge round i to o_i . Now, note that by Claim 4.2, $d(o_i, \mathcal{C}_{U_{i-1}}) \geq r_i - \varepsilon\Delta > (\delta - \varepsilon)\Delta \geq \Delta\delta/2$. Therefore, once the distance of an optimal point o to \mathcal{C}_{U_i} falls below $\Delta\delta/2 = 8\varepsilon^{1/3}\Delta/2$, it cannot be charged again in any future iteration. The initial distance of o to \mathcal{C}_{U_0} is at most Δ . As such, by Lemma 4.3, an optimal point o can get charged at most k times, where k is the smallest positive integer such that $(1 - \varepsilon^{2/3})^k \Delta \leq 4\varepsilon^{1/3}\Delta$, which holds if $\exp(-k\varepsilon^{2/3}) \leq 4\varepsilon^{1/3}$. Namely, $k = O(\varepsilon^{-2/3} \log 1/\varepsilon)$.

Using the same idea of decreasing values of ε , as done in Lemma 2.6, one can improve this bound to $O(1/\varepsilon^{2/3})$. We omit the easy but tedious details. We conclude that the number of iterations performed by the algorithm is at most $m = O(k_{\text{opt}}/\varepsilon^{2/3})$.

So the distance of all the points of P_{opt} from \mathcal{C}_{U_m} is at most $\delta\Delta$. Now, consider any point $p \in \mathcal{C}_P$. Let $t = \text{nn}(p, \mathcal{C}_{P_{\text{opt}}})$, and observe that $\|p - t\| \leq \varepsilon\Delta$. Since $t \in \mathcal{C}_{P_{\text{opt}}}$, we have that t can be written as a convex combination $t = \sum_{i=1}^{\nu} \alpha_i o_i$, where $\alpha_1, \dots, \alpha_{\nu} \geq 0$, $\sum_i \alpha_i = 1$, and $o_1, \dots, o_{\nu} \in P_{\text{opt}}$. For $i = 1, \dots, \nu$, let $o'_i = \text{nn}(o_i, \mathcal{C}_{U_m})$, and note that $t' = \sum_i \alpha_i o'_i \in \mathcal{C}_{U_m}$. Now observe that for all i , $\|o_i - o'_i\| \leq \delta\Delta$. In particular, $(o_i - o'_i) \in \text{ball}(0, \delta\Delta)$, and hence $\sum_i \alpha_i (o_i - o'_i) \in \text{ball}(0, \delta\Delta)$. Therefore $d(p, \mathcal{C}_{U_m}) \leq \|p - t'\| \leq \|p - t\| + \|t - t'\| \leq \varepsilon\Delta + \|\sum_i \alpha_i (o_i - o'_i)\| \leq (\varepsilon + \delta)\Delta$. We conclude that $d_H(\mathcal{C}_{U_m}, \mathcal{C}_P) \leq (\varepsilon + \delta)\Delta$.

As for the running time, at each iteration, the algorithm computes the point in P furthest away from \mathcal{C}_{U_i} . The analysis above assumes these queries are done exactly, which is expensive. However, by

Lemma 2.6 one can use faster $\varepsilon\Delta$ -approximate queries. Specifically, in each iteration, for each point $p \in P$ use **Lemma 2.6** to compute an additive $\varepsilon\Delta$ -approximation to its distance to \mathcal{C}_{U_i} , and then select the point in P with the largest returned approximate distance. It is easy to verify this does not change the correctness of the algorithm. Specifically, the point t_i chosen in the i th round, may now be $\varepsilon\Delta$ closer to the current convex hull than the furthest point, and so in the analysis of **Lemma 4.3**, o_i may lie as much as $\varepsilon\Delta$ above t_i . In particular, the length of τ does not change, however now ρ is only bounded by $2\sqrt{r\varepsilon}$ instead of $\sqrt{2r_i\varepsilon}$, and this constant factor difference only slightly degrades the constant in front of $\varepsilon^{2/3}$ in the lemma statement. The other effect is that when the algorithm stops the distance to the convex hull is bounded by $(8\varepsilon^{1/3} + \varepsilon)\Delta$, and this is accounted for in the above theorem statement.

Now using **Lemma 2.6** directly, it takes $O(nmd/\varepsilon^2)$ time per round to find the $\varepsilon\Delta$ approximate furthest point, and therefore the total running time is $O(nm^2d/\varepsilon^2)$. \blacksquare

4.2.1. Improving the running time further

The running time of the algorithm of **Lemma 4.4** can be improved further, but it requires some care. Let $L_{i-1} = \text{span}(U_{i-1})$ denote the linear subspace spanned by the point set U_{i-1} , with the orthonormal basis v_1, \dots, v_{i-1} . For any point $p \in P$, let p'_{i-1} denote its orthogonal projection onto the subspace; that is, $p'_{i-1} = \text{nn}(p, L_{i-1}) = \sum_{j=1}^{i-1} \langle p, v_j \rangle v_j$, and let $\ell_{i-1}(p) = \|p - p'_{i-1}\| = d(p, L_{i-1})$. Observe, that for any point $t \in L_{i-1}$ and any point $p \in \mathbb{R}^d$, we have that $\|p - t\| = \sqrt{\|p - p'_{i-1}\|^2 + \|p'_{i-1} - t\|^2}$ by the Pythagorean theorem, where p'_{i-1} is the projection of p to L_{i-1} .

As such, for any point $p \in P$, in the beginning of the i th iteration, the algorithm has the projection and distance of p to L_{i-1} ; that is, $p'_{i-1} = (\langle p, v_1 \rangle, \dots, \langle p, v_{i-1} \rangle)$. and $\ell_{i-1}(p)$. The algorithm also initially computes for each point $p \in P$ its norm $\|p\|^2$. Therefore, given any point $t \in L_{i-1}$, its distance to a point $p \in P$ can be computed in $O(i)$ time (instead of $O(d)$). The algorithm also maintains, for every point $p \in P$, an approximate nearest neighbor $\text{nn}_{i-1}(p) \in \mathcal{C}_{U_{i-1}}$; that is,

$$d(p, \mathcal{C}_{U_{i-1}}) \leq \|p - \text{nn}_{i-1}(p)\| \leq d(p, \mathcal{C}_{U_{i-1}}) + \varepsilon\Delta,$$

where $\Delta = \text{diam}(P)$. Naturally, the algorithm also maintains the distance $d_{i-1}(p) = \|p - \text{nn}_{i-1}(p)\|$.

Now, the algorithm does the following in the i th iteration:

- (A) Computes, in $O(n)$ time, the point $p \in P$ that maximizes $d_{i-1}(p)$.
- (B) Let p'_{i-1} be the projection of p to L_{i-1} . Computes, in $O(d)$ time, the new vector for the basis of L_i ; that is $v_i = (p - p'_{i-1}) / \|p - p'_{i-1}\|$. Now v_1, \dots, v_i is an orthonormal basis of the linear space L_i .
- (C) For every point $p \in P$, update its projection p'_{i-1} into L_{i-1} into the projection of p into L_i , by computing $\langle p, v_i \rangle$. Also, update $\ell_i(p) = \sqrt{\ell_{i-1}(p)^2 - \langle p, v_i \rangle^2}$.
- (D) Let P' denote the projected points of P into L_i . For every $p \in P$, we need to update $\text{nn}_{i-1}(p)$ to $\text{nn}_i(p)$ (and the associated distance). To this end, the algorithm of **Lemma 2.6** is called on p'_i and U_i (all lying in the subspace L_i which is of dimension i). Importantly, the algorithm of **Lemma 2.6** is being warm-started with the point $\text{nn}_{i-1}(p)$. Let $\#_i(p)$ be the number of iterations performed inside the algorithm of **Lemma 2.6** to update the nearest-neighbor to p . Observe, that the running time for p is $O(\#_i(p)i^2)$, since $i = |U_i|$, the points lie in an i dimensional space, and as such, every iteration of the algorithm of **Lemma 2.6** takes $O(i^2)$ time.

Lemma 4.5. For $m = O(k_{\text{opt}}/\varepsilon^{2/3})$, the running time of the above algorithm is $O(nm(d + m/\varepsilon^2 + m^2))$.

Proof: The algorithm performs $m = O(k_{\text{opt}}/\varepsilon^{2/3})$ iterations, and this bound the dimension of the output subspace. Every iteration of the algorithm takes $O(nd)$ time, except for the last portion of updating the approximate nearest point for all the points of P (i.e., (D)). The key observation is that $\sum_i (\#_i(p) - 1) = O(1/\varepsilon^2)$, since if the algorithm of Lemma 2.6 runs $\alpha = \#_i(p) > 1$ iterations, then the distance of p to the convex-hull shrinks by a factor of $(1 - \varepsilon^2/2)^\alpha$. Arguing as in the proof of Lemma 2.6, this can happen $O(1/\varepsilon^2)$ times before p is in distance at most $\varepsilon\Delta$ from the convex-hull, and can no longer be updated. As such, for a single point $p \in P$, the operations in (D) takes overall $\sum_{i=1}^m O(i^2(\#_i(p) - 1)) = O(m^2(m + 1/\varepsilon^2))$ time. This implies the overall running time of the algorithm is $O(n(dm + m^2/\varepsilon^2 + m^3))$. ■

4.2.2. The result

Theorem 4.6. *Let P be a set of n points in \mathbb{R}^d with diameter $\Delta = \text{diam}(P)$, and let $\varepsilon > 0$ be a parameter, then one can compute a set $U \subseteq P$, such that*

- (i) $d_H(\mathcal{C}_U, \mathcal{C}_P) \leq (8\varepsilon^{1/3} + \varepsilon)\Delta$, and
- (ii) $|U| \leq O(k_{\text{opt}}/\varepsilon^{2/3})$, where $k_{\text{opt}} = k_{\text{opt}}(P, \varepsilon)$.

The running time of the algorithm is $O(nm(d + m/\varepsilon^2 + m^2))$. for $m = O(k_{\text{opt}}/\varepsilon^{2/3})$. (Here, the constants hidden in the O are independent of the dimension.)

Remark. (A) The constants hidden in the O notation used of Theorem 4.6 are independent of the dimension. In comparison to the other algorithms in this paper, the approximation quality is slightly worse. However, the advantage is a drastic improvement in the size of the approximation.

(B) The running time of the algorithm of Theorem 4.6 can be further improved, by keeping track for each point $p \in P$, and each point $t \in U_i$, the distance of t from the hyperplane (in L_i) that determines whether or not the approximate nearest neighbor to p needs to be recomputed. By careful implementation, this can be done in the i th iteration in $O(in)$ time (updating $O(in)$ such numbers in this iteration). This improves the running time to $O(nm(d + m/\varepsilon^2))$. Motivated by our laziness we omit the messy details.

Remark. Note that the algorithm is a simple iterative process, which is oblivious to the value of the diameter $\Delta = \text{diam}(P)$ and does not use it directly anywhere. Nevertheless, after $O(k_{\text{opt}}/\varepsilon^{2/3})$ iterations the solution is an $(8\varepsilon^{1/3} + \varepsilon)\Delta$ -approximation to the convex hull. In practice, one may not know the value of k_{opt} , and so this value cannot be used in a stopping condition. However, it is easy to get a 2-approximation Δ' , such that $\Delta \leq \Delta' \leq 2\Delta$, by a linear scan of the points. Then, one can use the check $d(t_i, \mathcal{C}_{U_i}) = d_H(\mathcal{C}_P, \mathcal{C}_{U_i}) \leq (8\varepsilon^{1/3} + \varepsilon)\Delta'/2$ as a stopping condition, where U_i is the current approximation.

References

- [AGM14] S. Arora, R. Ge, and A. Moitra. New algorithms for learning incoherent and overcomplete dictionaries. In *Proc. 27th Annual Conf. Learning Theory (COLT)*, volume 35, pages 779–806. JMLR.org, 2014.
- [AHV05] P. K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximation via coresets. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*, Math. Sci. Research Inst. Pub. Cambridge, New York, NY, USA, 2005.

- [Bar15] S. Barman. Approximating nash equilibria and dense bipartite subgraphs via an approximate version of caratheodory’s theorem. In *Proc. 47th Annu. ACM Sympos. Theory Comput.* (STOC), pages 361–369. ACM, 2015.
- [BBV15] M.-F. Balcan, A. Blum, and S. Vempala. Efficient representations for lifelong learning and autoencoding. In *Proc. 28th Annual Conf. Learning Theory* (COLT), volume 40, pages 191–210. JMLR.org, 2015.
- [BHR15] A. Blum, S. Har-Peled, and B. Raichel. Sparse Approximation via Generating Point Sets. *ArXiv e-prints*, July 2015.
- [BHR16] A. Blum, S. Har-Peled, and B. Raichel. Sparse approximation via generating point sets. In Robert Krauthgamer, editor, *Proc. 27th ACM-SIAM Sympos. Discrete Algs.* (SODA), pages 548–557. SIAM, 2016.
- [Cla93] K. L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. 3th Workshop Algorithms Data Struct.* (WADS), volume 709 of *Lect. Notes in Comp. Sci.*, pages 246–252. Springer-Verlag, 1993.
- [Cla10] K. L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Trans. Algo.*, 6(4):63:1–63:30, 2010.
- [Gon85] T. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.
- [Har11] S. Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., Boston, MA, USA, 2011.
- [HKMR15] S. Har-Peled, N. Kumar, D. Mount, and B. Raichel. Space exploration via proximity search. In *Proc. 31st Annu. Sympos. Comput. Geom.* (SoCG), volume 34 of *LIPICs*, pages 374–389, 2015.
- [HV04] S. Har-Peled and K. R. Varadarajan. High-dimensional shape fitting in linear time. *Discrete Comput. Geom.*, 32(2):269–288, 2004.
- [HW87] D. Haussler and E. Welzl. ϵ -nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [Lon01] P. M. Long. Using the pseudo-dimension to analyze approximation algorithms for integer programming. In *Proc. 7th Workshop Algorithms Data Struct.* (WADS), volume 2125 of *Lect. Notes in Comp. Sci.*, pages 26–37, 2001.
- [Nov62] A.B.J. Novikoff. On convergence proofs on perceptrons. In *Proc. Symp. Math. Theo. Automata*, volume 12, pages 615–622, 1962.
- [SWW12] D.A. Spielman, H. Wang, and J. Wright. Exact recovery of sparsely-used dictionaries. In *Proc. 25th Annual Conf. Learning Theory* (COLT), pages 37.1–37.18, 2012.