New Constructions of Weak Epsilon-Nets

Jiří Matoušek* Uli Wagner†

Abstract

A finite set $N \subset \boldsymbol{R}^d$ is a *weak ε-net* for an n-point set $X \subset \boldsymbol{R}^d$ (with respect to convex sets) if N intersects every convex set K with $|K \cap X| \geq \varepsilon n$. We give an alternative, and arguably simpler, proof of the fact, first shown by Chazelle et al. [8], that every point set X in $\boldsymbol{R}^d$ admits a weak ε-net of cardinality $O(\varepsilon^{-d}\,\mathrm{polylog}(1/\varepsilon))$. Moreover, for a number of special point sets (e.g., for points on the moment curve), our method gives substantially better bounds. The construction yields an algorithm to construct such weak ε-nets in time $O(n \ln(1/\varepsilon))$.

1 Introduction

Weak ε-nets with respect to convex sets as defined in the abstract were introduced by Haussler and Welzl [10] and later found many applications in discrete geometry, most notably in the spectacular proof of the Hadwiger-Debrunner (p, q)-conjecture by Alon and Kleitman [3].

For $0 < \varepsilon < 1$ and $X \subset \boldsymbol{R}^d$, let $f(X, \varepsilon)$ denote the minimum cardinality of a weak ε-net for X, and let

$$f(d, \varepsilon) = \sup\{f(X, \varepsilon) : X \subset \boldsymbol{R}^d \text{ finite}\}.$$

Alon et al. [2] proved that $f(d, \varepsilon)$ is finite for every $d \geq 1$ and $\varepsilon > 0$. They established the bounds $f(2, \varepsilon) = O(\varepsilon^{-2})$ and $f(d, \varepsilon) \leq C_d \varepsilon^{-(d+1-\delta_d)}$, where C_d depends only on d and δ_d is a positive number that tends to 0 (exponentially fast) as $d \to \infty$. Chazelle et al. [8] improved the bound for all fixed dimensions $d \geq 3$ to $O(\varepsilon^{-d} \ln(1/\varepsilon)^{b(d)})$, with a suitable constant $b(d)$.

Our main result is an alternative, and arguably simpler, proof of the upper bound $f(d, \varepsilon) = O(\varepsilon^{-d}\,\mathrm{polylog}(1/\varepsilon))$ for fixed dimension d (with constants and the exponent of the logarithm depending on d). Our proof is based on the following partition theorem [12]: For a finite point set $X \subset \boldsymbol{R}^d$ and an integer parameter r, $2 \leq r < |X|$, there is a partition of X into $\Theta(r)$ parts of roughly equal size such that no hyperplane crosses the convex hull of more than $O(r^{1-1/d})$ parts.

Besides proving the upper bound mentioned above for general point sets in arbitrary dimension, Chazelle et al. [8] also construct weak ε-nets of size $O(\varepsilon^{-1}\,\mathrm{polylog}(1/\varepsilon))$ for planar point sets in convex position. Our approach can be seen as a generalization of that construction.

For special classes of point sets, such as point sets on the moment curve, or more generally, point sets on a k-dimensional algebraic variety of bounded degree, partitions with smaller hyperplane crossing number are available, and in such cases, we get correspondingly smaller weak ε-nets.

*Department of Applied Mathematics and Institute for Theoretical Computer Science (ITI), Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic, matousek@kam.mff.cuni.cz

†Institut für Theoretische Informatik, ETH Zürich, CH-8092 Zürich, Switzerland, uli@inf.ethz.ch

Partitions with small hyperplane crossing number can also be constructed efficiently, so our construction can be turned into an algorithm for computing weak ε-nets. This will be discussed in Section 5. The time required for computing a weak ε-net of size $O(\varepsilon^{-d} \operatorname{polylog}(1/\varepsilon))$ for an n-point set X in $\boldsymbol{R}^d$ is $O(n \ln(1/\varepsilon))$ (with constants depending on d). We note that Chazelle et al. [7] gave an algorithm with running time $n(1/\varepsilon)^{O(1)}$.

An earlier version of this paper [15] also included a proof of the fact that for points uniformly distributed on the $(d-1)$-dimensional sphere, there are weak ε-nets of size $O(\varepsilon^{-1} \ln(1/\varepsilon)^2)$ (with a constant of proportionality depending on d). We are grateful to János Pach for pointing out to us that this had already been proved by Bradford and Capoyleas [6]. (For the case $d = 2$, Chazelle et al. [8] showed an $O(\varepsilon^{-1})$ bound.)

2 Toolbox

Here we list several results from discrete geometry that we will use. Proofs and references can be found, for instance, in [13].

Centerpoints. Let X be a finite set of points in $\boldsymbol{R}^d$. A point $q \in \boldsymbol{R}^d$ (not necessarily in X) is called a *centerpoint* for X if every halfspace containing q contains at least $|X|/(d+1)$ points of X.

Center Point Theorem. *For every finite point set $X \in \boldsymbol{R}^d$, there exists a centerpoint.*

The Löwner-John Ellipsoid. We will make use of the fact that every d-dimensional convex body K contains an ellipsoid of volume $\Omega(\operatorname{vol} K)$; this is a consequence of the following:

John's Lemma. *Let $K \subset \boldsymbol{R}^d$ be a compact convex body with nonempty interior. Then there exists an ellipsoid E such that*

$$E \subseteq K \subseteq dE,$$

where dE is the ellipsoid that arises from E by expanding it from its center by a factor of d.

VC-Dimension and ("Strong") ε-Nets. Let X be an arbitrary set and $\mathcal{F}$ a family of subsets of X. A set $A \subseteq X$ is *shattered* by $\mathcal{F}$ if every subset of A can be obtained as the intersection of A with some set $S \in \mathcal{F}$. The *VC-dimension* of $\mathcal{F}$ is defined as

$$\text{VC-dim}(\mathcal{F}) = \sup\{|A| : A \subseteq X \text{ is shattered by } \mathcal{F}\}.$$

Now suppose that X is equipped with a probability measure μ and that $\mathcal{F}$ is a system of μ-measurable sets, and let $\varepsilon > 0$. A subset $N \subseteq X$ is called an ε-net for $\mathcal{F}$ with respect to μ if $N \cap S \neq \emptyset$ for every $S \in \mathcal{F}$ with $\mu(S) \geq \varepsilon$.

Epsilon-Net Theorem. *There is a constant C such that for all X, μ and $\mathcal{F}$ as above and all ε, $0 < \varepsilon \leq \frac{1}{2}$, there exists an ε-net for $\mathcal{F}$ w.r.t. μ of size at most $CD\varepsilon^{-1} \ln(1/\varepsilon)$, where $D :=$ VC-dim$(\mathcal{F})$.*

Let us briefly compare this with the definition of weak ε-nets for convex sets. The measure considered in that definition is the normalized counting measure on a finite set in $\boldsymbol{R}^d$ (but the definition could equally well be phrased with an arbitrary probability measure). The main difference is that the VC-dimension of the system of all convex sets in $\boldsymbol{R}^d$ has *infinite* VC-dimension, and so the epsilon-net theorem does not give anything.

3 A General Construction via Partitions

We begin with an auxiliary statement concerning centerpoints.

Lemma 1. *Let $X \subset \boldsymbol{R}^d$ be a finite point set. Then there are subsets $T_1, T_2, \ldots, T_d \subseteq X$, $|T_j| \leq d$, such that $K := \bigcap_{j=1}^{d} \operatorname{conv}(T_j) \neq \emptyset$ and the lexicographic minimum of K is a centerpoint of X.*

Proof. Let $\mathcal{Y}$ be the system of all $Y \subseteq X$ such $|Y| > \frac{d}{d+1}|X|$ and there is an open halfspace γ with $Y = X \cap \gamma$. Then the intersection $C := \bigcap_{Y \in \mathcal{Y}} \operatorname{conv}(Y)$ is nonempty, and its points are centerpoints of X; see, e.g., the proof of the centerpoint theorem in [13]. Let q be the lexicographic minimum of C. By a standard argument using Helly's theorem (see, e.g., [13], Lemma 8.1.2), there are sets $Y_1, Y_2, \ldots, Y_t \in \mathcal{Y}$, $t \leq d$, such that q is also the lexicographic minimum of $\bigcap_{j=1}^{t} \operatorname{conv}(Y_j)$. Assuming that no Y_j can be omitted without violating this property, q has to lie on the boundary of $\operatorname{conv}(Y_j)$ for each j. Then by Carathéodory's theorem, for each j we can choose an at most d-point $T_j \subseteq Y_j$ with $q \in \operatorname{conv}(T_j)$. Then $T_1, \ldots, T_t$ have the property required in the lemma.

In order to simplify notation in the sequel, the lemma speaks about exactly d sets T_j. If the above proof yields fewer sets, we thus repeat some of them the appropriate number of times. □

The crucial ingredient of our construction are partitions with small hyperplane crossing numbers. In the following, $\boldsymbol{R}_+$ will denote the set of positive real numbers, and we will use the notation "$\dot\cup$" for the union of pairwise disjoint sets.

Definition 2. Let $\kappa : \boldsymbol{R}_+ \to \boldsymbol{R}_+$ be a nondecreasing sublinear function. We say that a finite point set $Y \subset \boldsymbol{R}^d$ is *κ-partitionable* if, for every r, $2 \leq r < |Y|$, there exists a partition $Y = Y_1 \dot\cup \ldots \dot\cup Y_t$, $t \leq r$, such that $|Y|/r \leq |Y_i| \leq 2|Y|/r$ for all i and the *hyperplane crossing number* of the partition is at most $\kappa(r)$, i.e. no hyperplane h crosses more than $\kappa(r)$ of the sets Y_i.

Here, a hyperplane h is said to *cross* a point set Z if it intersects $\operatorname{conv}(Z)$ but does not contain it.

Further, a point set $X \subset \boldsymbol{R}^d$ is *hereditarily* κ-partitionable if every subset $Y \subseteq X$ is κ-partitionable.

Here is our main theorem.

Theorem 3. *Suppose a finite point set $X \subset \boldsymbol{R}^d$ is hereditarily $O(r^{1-1/c})$-partitionable for some constant c, $1 \leq c \leq d$. That is, X is hereditarily κ-partitionable with $\kappa(r) \leq br^{1-1/c}$ for some constant b (for convenience in later calculations, we assume $b \geq 1$).*

Then, for every $\varepsilon > 0$, there is a weak ε-net of size

$$O\left(\varepsilon^{-c} \ln(1/\varepsilon)^a\right)$$

for X, where a depends only on d, c, and b; it can be shown that $a = \max\{cd^2(\ln 2 + \ln(d+1) + \ln b), 2\}$ suffices.

Proof. We first briefly outline the main idea for the proof of Theorem 3: Let r and k be two parameters, to be specified later, which satisfy $1 \leq k < r/2 \leq |X|/4$. Let $X = X_1 \dot\cup \ldots \dot\cup X_t$, $t \leq r$, be a partition for X as introduced in Definition 2 above, i.e., such that $|X|/r \leq |X_i| \leq 2|X|/r$ for all i and that no hyperplane crosses more than $\kappa(r)$ of the parts X_i. For a subset

$A \subseteq X$ with $|A| \geq \varepsilon|X|$, we will distinguish the following two cases: If A intersects "few" of the X_i's, then there must be some i such that the density of $A \cap X_i$ in X_i is at least some suitable ε' which is significantly larger than ε. We take care of this case by inductively constructing weak ε'-nets for all the X_i's (this is where we use the fact that suitable partitions exist for all subsets of X). On the other hand, in the case that A intersects "many" of the parts X_i, we will use the fact that the partition has low hyperplane crossing number to show that $\operatorname{conv}(A)$ contains a point q from a certain set N' of points defined by suitable constant-size subsets of X.

We now present the proof proper, which is subdivided into three parts:

1. (Recursive construction of the set N)
   We may assume that $\varepsilon \leq \varepsilon_0$ for some suitable constant ε_0 depending on b, c, d (otherwise, we construct a weak ε_0-net of size C/ε_0, $C = C(\varepsilon_0, d)$, by the method in [2], say). We may also assume that $|X| \geq 2/\varepsilon_0$, else we may take $N := X$ as a weak ε-net.

   Next, we construct the set N' mentioned above: We pick a *transversal* for the X_i's, i.e., a subset $T \subseteq X$ that contains exactly one point p_i from each X_i. Let $T_1, T_2, \ldots, T_d$ be subsets of T consisting of at most d points each. Whenever $\bigcap_{j=1}^{d} \operatorname{conv}(T_j) \neq \emptyset$, we define $q = q(T_1, \ldots, T_d)$ as the lexicographic minimum of $\bigcap_{j=1}^{d} \operatorname{conv}(T_j)$. We take N' to consist of all such points:

$$N' := \{q(T_1, \ldots, T_d) : T_j \subseteq T,\ |T_j| \leq d \text{ for } 1 \leq j \leq d, \text{ and } \bigcap_{j=1}^{d} \operatorname{conv}(T_j) \neq \emptyset\}. \quad (1)$$

   We observe that since $|T| \leq r$, we have $|N'| \leq r^{d^2}$. Further, for $1 \leq i \leq t$, let N_i be an inductively constructed weak $\frac{r\varepsilon}{2k}$-net for the set X_i. We define

$$N := N' \cup \bigcup_{i=1}^{t} N_i. \quad (2)$$

2. (Proof of the weak ε-net property for N)
   Suppose that $A \subseteq X$, $|A| \geq \varepsilon|X|$. We distinguish two cases.

   (i) If A intersects at most k of the sets X_i, then for some i, $|A \cap X_i| \geq \frac{\varepsilon}{k}|X| \geq \frac{r\varepsilon}{2k}|X_i|$. Thus, $\operatorname{conv}(A)$ intersects N_i.

   (ii) Let A intersect more than k of the sets X_i; say $A \cap X_i \neq \emptyset$ for $1 \leq i \leq k+1$ (see Fig. 1). Let $p_1, \ldots, p_{k+1}$ be the corresponding points from the transversal T chosen above, i.e. $X_i \cap T = \{p_i\}$. By Lemma 1, N' contains a point q that is a centerpoint of $\{p_1, \ldots, p_{k+1}\}$.
   We claim that $q \in \operatorname{conv}(A)$. Otherwise, q can be strictly separated from A by a hyperplane h, say $q \in h^-$ and $A \subset h^+$, where h^+ and h^- denote the open halfspaces bounded by h. Then, by the centerpoint property, at least $\frac{k+1}{d+1}$ of the points p_i lie in h^-, say $p_i \in h^-$ for $1 \leq i \leq \lceil \frac{k+1}{d+1} \rceil$. It follows that for these indices i, X_i contains a point from h^- as well as a point from h^+ (since $\emptyset \neq A \cap X_i \subset h^+$).

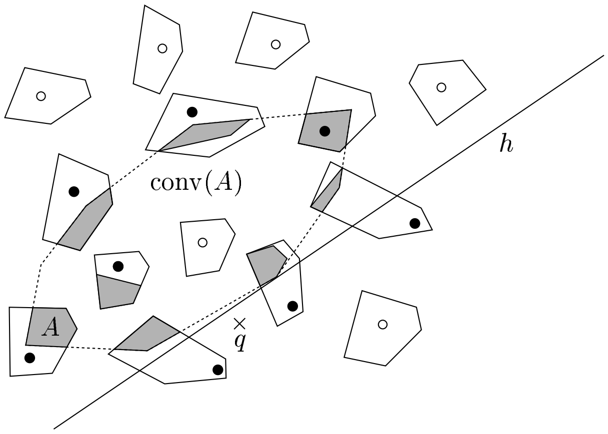

Figure 1: Checking the weak ε-net property.

Therefore, h crosses at least $\lceil \frac{k+1}{d+1} \rceil$ of the sets X_i. This leads to a contradiction if we set

$$k := \lfloor (d+1)\kappa(r) \rfloor, \tag{3}$$

which of course we do.

3. (Estimating the size of N)
We have to prove that for the exponent a as in the statement of Theorem 3,

$$|N| \leq C\varepsilon^{-c} \ln(1/\varepsilon)^a \tag{4}$$

for some suitable constant C. We proceed by induction on $|X|$. We first choose a sufficiently small constant ε_0; the subsequent calculations will show that a suitable value is

$$\ln \ln(1/\varepsilon_0) = 2cd^2(\ln 2 + \ln(d+1) + \ln(b)). \tag{5}$$

As remarked in Step 1 above, there is a constant C, depending on ε_0 and d, such that for $|X| < 2/\varepsilon_0$ or $\varepsilon > \varepsilon_0$, we even have $|N| < C/\varepsilon$.

Thus, we may assume $\varepsilon \leq \varepsilon_0$ and $|X| \geq 2/\varepsilon$. It remains to handle the inductive step and to specify the partition parameter r. Note that by (3),

$$\frac{k}{r} \leq \frac{(d+1)br^{1-\frac{1}{c}}}{r} = \frac{(d+1)b}{r^{1/c}}. \tag{6}$$

We set

$$r := \varepsilon^{-c/d^2} \ln(1/\varepsilon)^{1/d^2}. \tag{7}$$

(Observe that by our assumptions on $|X|$ and ε and because of $b \geq 1$, Equations (3) and (7) produce admissible values of k and r, i.e. $1 \leq k < r/2 \leq |X|/4$.)

Now, inductively, we have

$$\begin{aligned}
|N| &= |N'| + \sum_{i=1}^{t} |N_i| \\
&\leq r^{d^2} + rC(2k/r)^c \varepsilon^{-c} \Big[\ln(1/\varepsilon) + \ln(2k/r)\Big]^a \\
&\leq \varepsilon^{-c} \ln(1/\varepsilon) \\
&\quad + C\varepsilon^{-c} \ln(1/\varepsilon)^a 2^c (d+1)^c b^c \left(1 + w(\varepsilon) - \frac{1}{d^2}\right)^a,
\end{aligned}$$

where $w(\varepsilon) := \frac{\ln 2 + \ln(d+1) + \ln b - \frac{1}{cd^2} \ln\ln(1/\varepsilon)}{\ln(1/\varepsilon)}$. It follows that $|N| \leq C\varepsilon^{-c} \ln(1/\varepsilon)^a$, as desired, provided that

$$2^c (d+1)^c b^c (1 + w(\varepsilon) - \tfrac{1}{d^2})^a + \tfrac{1}{C \ln(1/\varepsilon)^{a-1}} \leq 1.$$

By choice of ε_0, we have for all $\varepsilon \leq \varepsilon_0$ that $w(\varepsilon) \leq -x(\varepsilon)$, where $x(\varepsilon) := \frac{\ln\ln(1/\varepsilon)}{2cd^2 \ln(1/\varepsilon)}$. Moreover, we may assume that $C \geq 1$, say, and we have $a \geq 2$, hence $\frac{1}{C \ln(1/\varepsilon)^{a-1}} \leq \frac{1}{\ln(1/\varepsilon)} =: y(\varepsilon)$. So it suffices to show that

$$2^c (d+1)^c b^c \left(1 - \frac{1}{d^2} - x(\varepsilon)\right)^a \leq 1 - y(\varepsilon).$$

We take logarithms and use the facts that $1 + t \leq e^t$ for all real t (for the left-hand side), and that $1 - y > e^{-2y}$ for small $y > 0$ (for the right-hand side). Thus, we see that it is enough to show

$$c(\ln 2 + \ln(d+1) + \ln b) - \frac{a}{d^2} - a \cdot x(\varepsilon) + 2y(\varepsilon) \leq 0.$$

But the last two terms together are at most zero, because $a \geq 2$ and $\ln\ln(1/\varepsilon) \geq 2cd^2$ for $\varepsilon \leq \varepsilon_0$, and the first two terms together are at most zero by choice of a. This completes the proof of Theorem 3.

For later use we remark that setting $r := \varepsilon^{-\beta}$ works equally well for any sufficiently small constant $\beta > 0$; we always get a bound $O(\varepsilon^{-c}(\ln(1/\varepsilon))^a)$ with a suitable a depending on d, c, b, β. □

4 Applications of the General Construction

We now derive some consequences of Theorem 3. Since any point set in $\boldsymbol{R}^d$ is (hereditarily) $O(r^{1-1/d})$-partitionable [12], we have re-proved:

Theorem 4 (Chazelle et al. [8]). *For every finite point set $X \subset \boldsymbol{R}^d$ and every $\varepsilon \in (0,1)$, there exists a weak ε-net of size $O(\varepsilon^{-d} \ln(1/\varepsilon)^a)$, for a suitable constant $a = a(d)$.*

Points on the moment curve. Let X be a subset of the moment curve $\gamma := \{(t, t^2, \ldots, t^d) : t \in \boldsymbol{R}\}$. Such point sets X are κ-partitionable with $\kappa(r) = d$: Given $X = \{p_1, \ldots, p_n\}$, where the points are numbered according to their order along the curve γ, we partition the points into "intervals" of the appropriate length. That is, let $s = \lceil n/r \rceil$ and $q := \lfloor n/s \rfloor$, and define

$X_i := \{p_{(i-1)s+1}, p_{(i-1)s+2}, \ldots, p_{is}\}$ for $1 \le i \le q-1$, and $X_q := \{p_{(q-1)s+1}, \ldots, p_n\}$. If a hyperplane h crosses an interval X_i, it intersects the moment curve γ within that interval. Therefore, at most d intervals can be crossed, since no hyperplane has more that d points of intersection with γ. Thus, in the notation of Theorem 3, we have $c = 1$ and $b = d$, and we obtain:

Proposition 5. *Every finite subset of the moment curve in $\boldsymbol{R}^d$ admits a weak ε-net of size $O(\varepsilon^{-1}\ln(1/\varepsilon)^{a(d)})$, with $a(d) \le d^2(\ln 2 + \ln(d+1) + \ln d)$).*

Points on an algebraic variety or on the boundary of a convex set. The following lemma summarizes some improved partitioning results for special point sets:

Lemma 6.

(i) *Let V be k-dimensional algebraic variety in $\boldsymbol{R}^d$, $1 \le k \le d-1$, of degree bounded by a constant D. Then any finite $X \subset V$ is $O(r^{1-1/c}(\ln r)^{1+1/c})$-partitionable for $c = \lfloor (d+k)/2 \rfloor$ (with the constant depending on k, d, and D).*

(ii) *Let V be the relative boundary of a $(k+1)$-dimensional convex set in $\boldsymbol{R}^d$. Then any finite $X \subset V$ is $O(r^{1-1/c}(\ln r)^{1+1/c})$-partitionable for $c = \lfloor (d+k)/2 \rfloor$ (with the constant depending on k and d).*

(iii) *Let V be k-dimensional algebraic variety in $\boldsymbol{R}^d$, $1 \le k \le d-1$, of degree bounded by a constant D. Then any finite $X \subset V$ is $O(r^{1-1/c}(\ln r)^{1+1/c})$-partitionable for $c = \max(k, 2k-4+\eta)$, where $\eta > 0$ is an arbitrarily small constant and the implicit constant depends on k, d, D, η.*

Part (i) is based on a zone theorem by Aronov et al. [4], and it is explicitly mentioned in [1] (in the proof of Theorem 6.3, as a consequence of a partition theorem formulated in an abstract setting). Part (ii) follows by exactly the same argument from another zone theorem of Aronov et al. [4], where one has the relative boundary of $(k+1)$-dimensional convex set instead of the variety V, with the same bound on the complexity of the zone. Finally, part (iii) follows from known results on decompositions of arrangements of semialgebraic sets [9], [11] by the technique described in [1], but it includes a simple observation which may be new in this context and of independent interest, and so we outline the proof in Section 6.

We thus obtain:

Theorem 7.

(i) *Let V be a k-dimensional algebraic variety in $\boldsymbol{R}^d$ of degree bounded by a constant D. Then for every finite $X \subset V$ and every $\varepsilon \in (0,1)$, there exists a weak ε-net of cardinality $O(\varepsilon^{-c}\ln(1/\varepsilon)^a)$ for*

$$c = \min(\lfloor (d+k)/2 \rfloor, \max(k, 2k-4+\eta)),$$

with an arbitrarily small $\eta > 0$ and with a and the implicit constant depending on d, k, η, D.

(ii) *Let V be the relative boundary of a $(k+1)$-dimensional convex set in $\boldsymbol{R}^d$. Then for every finite $X \subset V$ and every $\varepsilon \in (0,1)$, a weak ε-net of cardinality $O(\varepsilon^{-c}\ln(1/\varepsilon)^a)$ exists for $c = \lfloor (d+k)/2 \rfloor$, with a and the implicit constant depending on d and k.*

It is fair to remark that Theorem 3 does not apply directly in this case, since we have some extra logarithmic factors in the bounds on the crossing numbers of the partitions. However, the calculations in the proof of Theorem 3 go through almost unchanged, with a suitable larger exponent a.

5 The Algorithmic Side

Whenever partitions with small hyperplane crossing number can be computed efficiently, our proof of Theorem 3 immediately yields an algorithm for computing weak ε-nets.

Theorem 8. *Suppose that X is a finite point set in $\boldsymbol{R}^d$ such that, for every subset $Y \subseteq X$ and for every r, $2 \le r < |Y|^\alpha$ with some constant $\alpha > 0$, we can compute in time $O(|Y| \ln r)$ a partition $Y = Y_1 \dot\cup \ldots \dot\cup Y_t$ with $t \le r$, $|Y|/r \le |Y_i| \le 2|Y|/r$ for all i, and hyperplane crossing number $O(r^{1-1/c})$.*

Then, for every $\varepsilon \in (0,1)$, we can find a weak ε-net of size $O(\varepsilon^{-c} \ln(1/\varepsilon)^{O(1)})$ for X in time $O(|X| \ln(1/\varepsilon))$.

For every set X of n points in $\boldsymbol{R}^d$, the assumptions of Theorem 8 are satisfied with $c = d$, see [12]. So we obtain:

Corollary 9. *For every n-point set $X \subset \boldsymbol{R}^d$ and every $\varepsilon \in (0,1)$, a weak ε-net of size $O(\varepsilon^{-d}(\ln(1/\varepsilon))^{O(1)})$ can be computed in time $O(n \ln(1/\varepsilon))$.*

Proof of Theorem 8. We construct a weak ε-net as in the proof of Theorem 3, setting $r := \varepsilon^{-\beta}$ with $\beta := \min(c/d^2, \alpha)$. According to the proof of Theorem 3, we obtain a weak ε-net of size $O(\varepsilon^{-c}(\ln(1/\varepsilon))^a)$ with a suitable constant a.

It remains to estimate the running time. We first note that the computation of $q(T_1, T_2, \ldots, T_d)$ as in our proof of Theorem 3, although not simple, is a constant-time operation for d fixed.

Let $g(n, \varepsilon)$ be the maximal time required to compute a weak ε-net for a subset $Y \subset X$ of cardinality n, $\varepsilon \ge (1/n)^{1/c}$ (for smaller ε, we can simply take the set Y as a weak ε-net). We have the recurrence

$$g(n, \varepsilon) = O(n \ln r) + r^{d^2} + \sum_i g(n_i, r\varepsilon/k), \tag{8}$$

where $k = O(r^{1-1/c})$ and $\sum_i n_i = n$.

With our choice of $r = \varepsilon^{-\beta}$, and since $(1/\varepsilon)^c \le n$, the first two terms in the recurrence for g together are at most $An \ln(1/\varepsilon)$ for some constant A. We also have $\frac{k}{r} = O(r^{-1/c}) \le B\varepsilon^{\beta/c}$ for a constant B. Assuming inductively that $g(n', \varepsilon') \le Cn' \ln(1/\varepsilon')$ for all $n' < n$ and $\varepsilon' > \varepsilon$, with a suitable constant C, we have

$$\begin{aligned} g(n, \varepsilon) &\le An \ln(1/\varepsilon) + \sum_i Cn_i \ln\left(\frac{k}{r} \cdot \frac{1}{\varepsilon}\right) \\ &= An \ln(1/\varepsilon) + Cn\Big[\ln(1/\varepsilon) + \ln(k/r)\Big] \\ &\le n\Big[A \ln(1/\varepsilon) + C \ln(1/\varepsilon) + C \ln B - (C\beta/c) \ln(1/\varepsilon)\Big] \\ &\le Cn \ln(1/\varepsilon), \end{aligned}$$

assuming that C is chosen so large that $C\beta/c \ge 2A$, say, and that ε is so small that $(\beta/c) \ln(1/\varepsilon) \ge 2 \ln B$. This finishes the proof of Theorem 8. $\square$

6 Partitions for Points on a Variety

Here we outline the proof of Lemma 6(iii). In order to use a general partition theorem from [1], we first recall the abstract framework defined there.

A *range space with elementary cells* is a triple $(X, \Gamma, \mathcal{E})$, where X is a ground set and Γ and $\mathcal{E}$ are set systems on X. The sets in Γ are called *ranges*, while those in $\mathcal{E}$ are called *elementary cells*. A range $\gamma \in \Gamma$ *crosses* a set $S \subseteq X$ if $\gamma \cap S \neq \emptyset$ and $S \not\subseteq \gamma$. A collection $\Xi \subseteq \mathcal{E}$ of elementary cells is called an *elementary cell decomposition* for a set $Q \subseteq \Gamma$ of ranges if $\bigcup \Xi = X$ and no $\gamma \in Q$ crosses any $e \in \Xi$ (both Q and Ξ are usually finite).

For example, one can take $X = \boldsymbol{R}^d$, let Γ be the set of all (closed) halfspaces in $\boldsymbol{R}^d$, and let $\mathcal{E}$ consists of all relatively open simplices in $\boldsymbol{R}^d$ (of all dimensions from 0 to d, and also including unbounded simplices, i.e., intersections of $d+1$ halfspaces). We note that a simplex $e \in \mathcal{E}$ is crossed by a hyperplane h, according to the definition used in the previous sections, iff it is crossed by at least one of the two closed halfspaces bounded by h. If Q is a finite set of halfspaces, then a triangulation of the arrangement of the bounding hyperplanes of Q is an elementary cell decomposition.

A *faithful linearization* of dimension d for a set system (X, Γ) is an injective mapping $\varphi\colon X \to \boldsymbol{R}^d$ such that $\Gamma = \{\varphi^{-1}(H) : H \text{ a closed halfspace in } \boldsymbol{R}^d\}$. Thus, (X, Γ) possessing a faithful linearization means that X can be identified with a subset of some $\boldsymbol{R}^d$ and Γ then consists of all intersections of that subset with halfspaces.

For our purposes, the results in [1] can be summarized as follows (a combination of Theorem 5.1 and Lemma 3.1 from [1]):

Theorem 10 ([1]). *Let $(X, \Gamma, \mathcal{E})$ be a range space with elementary cells, such that (X, Γ) has a faithful linearization of some constant dimension d, and such that the VC-dimension of the set system $\{\Gamma_e : e \in \mathcal{E}\}$ is bounded by a constant d_1, where Γ_e denotes the set of all ranges $\gamma \in \Gamma$ crossing the elementary cell e. Let us assume that every finite $Q \subseteq \Gamma$ admits an elementary cell decomposition consisting of at most $Cm^c(\log m)^{c_1}$ elementary cells, where $m = |Q|$ and C, $c > 1$, and $c_1 \geq 0$ are constants.*

Then for every n-point set $P \subseteq X$ and every r, $1 < r < n$, there exists a partition $P = P_1 \dot\cup P_2 \dot\cup \cdots \dot\cup P_t$, such that $\lfloor n/r \rfloor \leq |P_i| < 2\lfloor n/r \rfloor$ (thus $t = \Theta(r)$), and no $\gamma \in \Gamma$ crosses more than $O(r^{1-1/c}(\log r)^{1+c_1/c})$ of the P_i.

We recall that a *semialgebraic set* is a subset of $\boldsymbol{R}^d$ definable by a formula that is a Boolean combination of finitely many polynomial inequalities in the variables $x_1, \ldots, x_d$ with real coefficients (the book by Bochnak, Coste, and Roy [5] provides an extensive reference).

The *description complexity* of a semialgebraic set can be defined, for our purposes, as $\max(d, D, m)$, where D is the maximum of the degrees of the polynomials in the defining inequalities and m is the number of inequalities. Tarski's well known result on quantifier elimination implies that subsets of $\boldsymbol{R}^d$ definable by a first-order formula over the reals involving quantifiers, and in particular, images and inverse images of semialgebraic sets under polynomial maps, are semialgebraic. A *Tarski cell* is a semialgebraic set of description complexity bounded by a constant (possibly depending on other parameters declared as constants).

We also need to recall results of Chazelle et al. [9], with an improvement in dimension 4 by Koltun [11].

Theorem 11. *Let $f_1, \ldots, f_m \in \boldsymbol{R}[x_1, \ldots, x_d]$ be polynomials of degree at most D, where D is a constant. Then $\boldsymbol{R}^d$ can be partitioned into at most $T_d(m)$ Tarski cells so that the sign* *of each f_j is constant on each of these cells*[1] *and with the following bounds on $T_d(m)$:*

- $T_2(m) = O(m^2)$,
- $T_3(m) = O(m^3\beta(m))$, *where β is an extremely slow-growing function, with $\beta(m) = o(\log\log\log m)$, for instance, and*
- $T_d(m) = O(m^{2d-4+\eta})$ *for every fixed $d \geq 4$ and every fixed $\eta > 0$ (this includes the improvement from* [11]*).*

In the situation of Lemma 6(iii), we have an algebraic variety V of dimension k in $\boldsymbol{R}^d$ of degree bounded by D. So V can be written as the set of common zeros of (finitely many) d-variate polynomials with real coefficients, each of degree at most D.

The following fact is well known (see, e.g., the discussion of stratification of semialgebraic sets in [5] Section 9.1) and has been used in many computational-geometric papers, especially in the case $k = d-1$:

Lemma 12. *For any V be as above, there are a partition $V = V_1 \cup V_2 \cup \cdots \cup V_M$, with M a constant depending on k, d, D, and k-dimensional linear subspaces $L_1, L_2, \ldots, L_M$, such that each V_i is a semialgebraic set of description complexity bounded by a constant (depending on k, d, D) and the orthogonal projection $\pi_i\colon V_i \to L_i$ is injective.*

Proof of Lemma 6(iii). It suffices to prove partitionability for each V_i as in Lemma 12 separately. Indeed, we can first divide the given finite set $X \subset V$ into at most M parts, each of them contained in some V_i, and then we apply a partition theorem for each V_i separately, with $r = r(i)$ adjusted so that the parts have the size required for the partition of the original X. The hyperplane crossing number of the resulting partition of X is at most M-times worse than the worst of the crossing numbers for the individual V_i.

So let us consider an n-point $P \subseteq V_i$ and a given parameter r. In order to apply Theorem 10, we let V_i be the ground set (playing the role of X in that theorem), and let Γ be the set of all intersections of V_i with (closed) halfspaces. The inclusion map $V_i \to \boldsymbol{R}^d$ is thus a faithful linearization. We let $\mathcal{E}$ consist of all semialgebraic subsets of V_i of description complexity at most C_1, for a sufficiently large constant C_1. As can be shown, for example, by a linearization argument (see, e.g., [13], Section 10.3), the VC-dimension assumption in Theorem 10 is satisfied. It remains to exhibit suitable elementary cell decompositions.

Let $Q \subseteq \Gamma$ be a set of m ranges (intersections of V_i with halfspaces). For $\gamma \in Q$, $\pi_i(\gamma) \subseteq L_i$ is a Tarski cell. Let $f_1, f_2, \ldots, f_{m_1} \in \boldsymbol{R}[x_1, \ldots, x_d]$ be all polynomials involved in the formulas defining the sets $\pi_i(\gamma)$ for $\gamma \in Q$; we have $m_1 = O(m)$.

We let $\boldsymbol{R}^k \equiv L_i = e_1 \cup e_2 \cup \cdots \cup e_t$ be a decomposition of L_i into at most $T_k(m_1)$ Tarski cells as in Theorem 11, and let $\tilde{e}_j = \pi_i^{-1}(e_j) \subseteq V_i$ be the inverse image of e_j. Each $\tilde{e}_j$ is a Tarski cell, and thus it belogs to $\mathcal{E}$ (for C_1 large enough). If some $\gamma \in Q$ crosses some $\tilde{e}_j$, then $\pi_i(\gamma)$ crosses e_j (as π_i is injective). It follows that there is a polynomial f_j involved in the formula defining $\pi_i(\gamma)$ whose sign on e_j is not constant, and this contradicts the construction of the e_j. Therefore, we have elementary cell decompositions in $(V_i, \Gamma, \mathcal{E})$ of size at most $T_k(O(m))$. The proof of Lemma 6(iii) is concluded by an application of Theorem 10. □

[1] Chazelle et al. formulated their results for polynomials with rational coefficients, but the construction and proof work with arbitrary real coefficients as well; this has been used many times in the literature.

7 Concluding Remarks

For ε fixed and sufficiently small and $d \to \infty$, the function $f(d, \varepsilon)$ is bounded from below by $e^{\Omega(\sqrt{d})}$ [14]. However, for *fixed* dimension d, while there are now several constructions that yield upper bounds of $O(\varepsilon^{-d} \operatorname{polylog}(1/\varepsilon))$ or slightly worse, no better lower bound than the obvious $f(d, \varepsilon) = \Omega(1/\varepsilon)$ seems to be known. There seems to be no convincing reason why $f(d, \varepsilon)$ should be substantially superlinear in $1/\varepsilon$.

Our construction rules out the popular points on the moment curve as candidate for a class of points sets that require weak ε-nets of substantially superlinear size.

References

[1] Pankaj K. Agarwal and Jiří Matoušek. On range searching with semialgebraic sets. *Discrete Comput. Geom.*, 11:393–418, 1994.

[2] Noga Alon, Imre Bárány, Zoltán Füredi, and Daniel Kleitman. Point selections and weak ε-nets for convex hulls. *Combin. Probab. Comput.*, 1(3):103–112, 1992.

[3] Noga Alon and Daniel J. Kleitman. Piercing convex sets and the Hadwiger-Debrunner (p, q)-problem. *Adv. Math.*, 96(1):103–112, 1992.

[4] Boris Aronov, Marco Pellegrini, and Micha Sharir. On the zone of a surface in a hyperplane arrangement. *Discrete Comput. Geom.*, 9(2):177–186, 1993.

[5] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry.* Springer-Verlag, New York, 1998.

[6] Phillip G. Bradford and Vasilis Capoyleas. Weak ε-nets for points on a hypersphere. *Discrete Comput. Geom.*, 18(1):83–91, 1997.

[7] Bernard Chazelle, Herbert Edelsbrunner, David Eppstein, Michelangelo Grigni, Leonidas Guibas, Micha Sharir, and Emo Welzl. Algorithms for weak epsilon-nets. Manuscript, 1995, available at `http://citeseer.nj.nec.com/24190.html`.

[8] Bernard Chazelle, Herbert Edelsbrunner, Michelangelo Grigni, Leonidas J. Guibas, Micha Sharir, and Emo Welzl. Improved bounds on weak ε-nets for convex sets. *Discrete Comput. Geom.*, 13:1–15, 1995.

[9] Bernard Chazelle, Herbert Edelsbrunner, Leonidas Guibas, and Micha Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. In *Proc. 16th Internat. Colloq. Automata Lang. Program.*, volume 372 of *Lecture Notes Comput. Sci.*, pages 179–192. Springer-Verlag, Berlin etc., 1989.

[10] David Haussler and Emo Welzl. ε-Nets and Simplex Range Queries. *Discrete Comput. Geom.*, 2:127–151, 1987.

[11] Vladlen Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, pages 56–65, 2001.

[12] Jiří Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8(3):315–334, 1992.

[13] Jiří Matoušek. *Lectures on Discrete Geometry.* Springer-Verlag, New York, 2002.

[14] Jiří Matoušek. A lower bound for weak ε-nets in high dimension. *Discrete Comput. Geom.*, 28:45–48, 2002.

[15] Jiří Matoušek and Uli Wagner. New constructions of weak epsilon-nets. In *Proc. 19th Annual ACM Symposium on Computational Geometry*, pages 129–135, 2003.