

Chapter 7

Directional Width

Any unit vector

$$v \in S_{d-1} := \{x \in \mathbb{R}^d \mid \|x\| = 1\}$$

defines a *direction* in \mathbb{R}^d . In this chapter, we are interested in approximating the “extent” of a point set in a given direction. Actually, we will approximate the extent in all directions simultaneously. You are well-aware of this concept of extent; for example, you may read in a travel guide that the extent of Switzerland in north-south direction is roughly 220 km, and its extent in east-west direction is roughly 348 km. Here is the formal definition.

Definition 7.0.1 Let $K \subseteq \mathbb{R}^d$ be a compact set, $v \in S_{d-1}$. The width of K in direction v is defined as

$$w_v(K) = \max_{p \in K} p \cdot v - \min_{p \in K} p \cdot v.$$

Bilinearity of the scalar product easily yields (think about it!) the following alternative definition.

$$w_v(K) = \max_{p, q \in K} (p - q) \cdot v.$$

Because v is a unit vector, $(p - q) \cdot v$ is the (signed) length of $(p - q)$'s projection onto v . The directional width $w_v(K)$ is therefore the longest such projection you can get from two points in K . Another way of thinking about the directional width is as the minimum distance between any two hyperplanes with normal vector v that have K between them, see Figure 7.1.

In case of the Switzerland example, the north-south extent is the minimum distance between two latitudes that have the country in between them, and the east-west extent is the same with longitudes.

We have the following intuitive

Fact 7.0.2 Let $K \subseteq \mathbb{R}^d$ be a compact set. Then

$$w_v(K) = w_v(\text{conv}(K)), \quad \forall v \in S_{d-1}.$$

In the sequel, we restrict (as usual) to finite point sets.




Figure 7.1: The directional width

7.1 A constant-factor approximation

Let $P \subseteq \mathbb{R}^d$ be a finite point set. We start off by showing that the bounding cuboid that we have computed in Chapter 6 to approximate the volume of $\text{conv}(P)$ also approximates all directional widths.

Theorem 7.1.1 *Assume that $\mathbf{0} \in P \subseteq \mathbb{R}^d$, and that P is full-dimensional (not contained in any hyperplane). Let*

$$C = \left\{ \sum_{k=1}^d \lambda_k v_k \mid \ell_k \leq \lambda_k \leq u_k, k = 1, \dots, d \right\} \quad (7.1)$$

be the bounding cuboid of P computed by a call to the algorithm `MinCuboidApprox`(P, d). Then for all $v \in S_{d-1}$, the inequality

$$w_v(C) \leq 2^{d+1} w_v(P)$$

holds.

We would like to remark that this bound is pretty bad. In approximating the volume of the convex hull, we have argued that an exponential bound in d is not that disturbing, because the volume is exponential in the scale of the point set. Any directional width, however, only grows *linearly* with the scale, and in that situation, an exponential upper bound is quite embarrassing. It is not known whether the approximation factor in Theorem 7.1.1 can be improved to something polynomial in d . The trouble is that $\text{conv}(P)$ might be quite long and skinny, even though its bounding cuboid is “fat”.

In order to prove the theorem, we proceed as in the volume case. We first find a *lower* bound for $w_v(P)$, then an *upper* bound for $w_v(C)$. Comparing these bounds then yields the result.

Let us briefly recall how the algorithm `MinCuboidApprox`(P, d) works, and how it sets the vales v_k, u_k, ℓ_k that define the cuboid in the statement of Theorem 7.1.1.

In round k (running from d to 1), the algorithm first chooses a point q_k of maximum norm in its current point set $P^{(k)}$ (where $P^{(d)} = P$). It then sets v_k to $q_k / \|q_k\|$ so that we get a unit vector; u_k is simply $\|q_k\|$, while ℓ_k is chosen as $\min_{p \in P^{(k)}} p \cdot v_k$. It follows that $|\ell_k| \leq u_k$.

If $P^{(k)} \neq \{\mathbf{0}\}$, the point set gets projected onto the hyperplane through the origin orthogonal to v_k , and the algorithm recursively proceeds with the resulting point set $P^{(k-1)}$ in round $k - 1$. Our assumption that P is full-dimensional guarantees that the algorithm actually reaches round 1, and that the vectors q_1, \dots, q_d (or their scaled versions v_1, \dots, v_d) form an orthogonal (or orthonormal) basis of \mathbb{R}^d .

Note that $u_k - \ell_k$ is exactly the width of $P^{(k)}$ in direction v_d , in formulas:

$$w_{v_k}(P^{(k)}) = u_k - \ell_k.$$

We also need to recall that the point set $P^{(k)}$ can be expressed in terms of the original set P via

$$P^{(k)} = \{p^{(k)} \mid p \in P, \}, \quad p^{(k)} = p - \sum_{i=k+1}^d (p \cdot v_i) v_i. \quad (7.2)$$

A lower bound for $w_v(P)$. As in the Chapter 6, let $p_k \in P$ be the preimage of $q_k \in P^{(k)}$, meaning that $q_k = p_k^{(k)}$.

Because the directional width cannot increase when we remove points, we have that

$$w_v(P) \geq w_v(\{\mathbf{0}, p_k\}) = |p_k \cdot v|, \quad k = 1, \dots, d,$$

and this implies

$$w_v(P) \geq \max_{k=1}^d |p_k \cdot v|. \quad (7.3)$$

This is our lower bound.

An upper bound for $w_v(C)$. Since we have $q_k = \|q_k\| v_k = u_k v_k$ and $|\ell_k| \leq u_k$ for all k , the cuboid (7.1) is contained in the cuboid

$$C' = \left\{ \sum_{k=1}^d \lambda_k q_k \mid -1 \leq \lambda_k \leq 1 \right\}, \quad (7.4)$$

so when we upper-bound the directional widths of the latter, we also have valid bounds for the original cuboid C .

We now want to express C' in terms of the p_k , in order to be able to compare this with our lower bound. Slightly rewriting the formula for $p^{(k)}$ in (7.2), for $p = p_k$ (observe that $q_k = p_k^{(k)}$) yields

$$p_k = q_k + \sum_{i=k+1}^d \frac{p_k \cdot v_i}{u_i} q_i, \quad k = 1, \dots, d. \quad (7.5)$$

This shows that the p_k are linear combinations of the q_i , where i ranges from k to d ; the equations can compactly be written in matrix form as

$$\begin{pmatrix} p_1^T \\ p_2^T \\ \vdots \\ p_{d-1}^T \\ p_d^T \end{pmatrix} = \begin{pmatrix} 1 & (p_1 \cdot v_2)/u_2 & \cdots & (p_1 \cdot v_{d-1})/u_{d-1} & (p_1 \cdot v_d)/u_d \\ 0 & 1 & \cdots & (p_2 \cdot v_{d-1})/u_{d-1} & (p_2 \cdot v_d)/u_d \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & (p_{d-1} \cdot v_d)/u_d \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \begin{pmatrix} q_1^T \\ q_2^T \\ \vdots \\ q_{d-1}^T \\ q_d^T \end{pmatrix}.$$

Let us denote the transformation matrix in the latter equation by A . Since A is upper-triangular with the diagonal entries being equal to 1, A is invertible, and we get

$$\begin{pmatrix} q_1^T \\ q_2^T \\ \vdots \\ q_{d-1}^T \\ q_d^T \end{pmatrix} = A^{-1} \begin{pmatrix} p_1^T \\ p_2^T \\ \vdots \\ p_{d-1}^T \\ p_d^T \end{pmatrix}.$$

It follows that

$$\left(\sum_{k=1}^d \lambda_k q_k \right)^T = (\lambda_1, \dots, \lambda_d) \begin{pmatrix} q_1^T \\ q_2^T \\ \vdots \\ q_{d-1}^T \\ q_d^T \end{pmatrix} = (\lambda_1, \dots, \lambda_d) A^{-1} \begin{pmatrix} p_1^T \\ p_2^T \\ \vdots \\ p_{d-1}^T \\ p_d^T \end{pmatrix} = \left(\sum_{k=1}^d \mu_k p_k \right)^T,$$

with

$$(\mu_1, \dots, \mu_d) = (\lambda_1, \dots, \lambda_d) A^{-1}.$$

Plugging this into (7.4) yields

$$C' = \left\{ \sum_{k=1}^d \mu_k p_k \mid (\mu_1, \dots, \mu_d) = (\lambda_1, \dots, \lambda_d) A^{-1}, -1 \leq \lambda_k \leq 1 \right\}. \quad (7.6)$$

It is important to understand that nothing fancy is going on here; the matrix A^{-1} is simply a basis transformation matrix: vectors whose coordinates $\lambda_1, \dots, \lambda_d$ are given with respect to the basis (q_1, \dots, q_d) are expressed using coordinates μ_1, \dots, μ_d with respect to the basis (p_1, \dots, p_d) .

In general, any invertible matrix may arise as a basis transformation matrix; but in our case, the matrix has a special structure, due to the way the p_k are constructed in the algorithm.

Observation 7.1.2 *All off-diagonal entries of A are at most one in absolute value.*

Proof. Consider any off-diagonal entry

$$\frac{p_k \cdot v_i}{u_i}, \quad i > k.$$

The fact that q_i was chosen in round i shows that

$$|p_k^{(i)} \cdot v_i| \leq \|p_k^{(i)}\| \leq \|q_i\| = u_i, \quad (7.7)$$

with $p_k^{(i)} \in P^{(i)}$ the iterated projection of p_k . On the other hand, computing $p_k^{(i)}$ according to (7.2) yields

$$p_k^{(i)} = p_k - \sum_{j=i+1}^d (p_k \cdot v_j) v_j,$$

and using pairwise orthogonality of the v_j gives

$$p_k^{(i)} \cdot v_i = p_k \cdot v_i.$$

Together with (7.7), we therefore get

$$\left| \frac{p_k \cdot v_i}{u_i} \right| = \left| \frac{p_k^{(i)} \cdot v_i}{u_i} \right| \leq \frac{u_i}{u_i} = 1.$$

□

This observation now implies bounds on the absolute values of the μ_k in (7.6). In an exercise, we ask you to show that in A^{-1} , the sum of absolute values in column k is at most 2^{k-1} . Because μ_k is a weighted sum of these values, with weights in $[-1, 1]$, we also get that $|\mu_k| \leq 2^{k-1}$ for all μ_k appearing in (7.6). This can be used to prove

Lemma 7.1.3 *The width of C' in any direction v satisfies*

$$w_v(C') \leq 2(2^d - 1) \max_{k=1}^d |p_k \cdot v|.$$

Together with $w_v(C) \leq w_v(C')$ and the lower bound (7.3), Theorem 7.1.1 follows. It remains to show the lemma.

Proof. Since C' is centrally symmetric ($x \in C'$ implies $-x \in C'$), we have

$$w_v(C') = 2 \max_{x \in C'} x \cdot v,$$

and it suffices to compute the latter maximum. For $x \in C'$ and μ_1, \dots, μ_d such that $x = \sum_{k=1}^d \mu_k p_k$ we get

$$x \cdot v \leq |x \cdot v| = \left| \sum_{k=1}^d \mu_k (p_k \cdot v) \right| \leq \sum_{k=1}^d |\mu_k| |p_k \cdot v| \leq \sum_{k=1}^d 2^{k-1} \max_{k=1}^d |p_k \cdot v| = (2^d - 1) \max_{k=1}^d |p_k \cdot v|.$$

Because this holds for all x , we have proved a bound for $\max_{x \in C'} x \cdot v$, and the lemma follows. □

7.2 Core sets for directional width

Based on the constant-factor approximation from the previous section, we can now construct *core sets*. Let us state the theorem right away.

Theorem 7.2.1 *Let $P \subseteq \mathbb{R}^d$, $|P| = n$ and choose $\varepsilon > 0$. Furthermore, let M be such that*

$$w_v(C) \leq Mw_v(P), \quad \forall v \in S_{d-1},$$

where C is the bounding cuboid of P obtained from `MinCuboidApprox`(P, d).

(i) *There exists a subset $S \subseteq P$ of size at most*

$$2 \left(\left\lceil \frac{2M}{\varepsilon} \right\rceil \right)^{d-1},$$

with the property that for all $v \in S_{d-1}$, $w_v(S) \geq (1 - \varepsilon)w_v(P)$.

(ii) *A set S as in (i)—an ε -core set of P with respect to directional widths—can be computed in time*

$$O(d^2n),$$

assuming that M is known.

With Theorem 7.1.1, we may choose $M = 2^{d+1}$. Note that even though the resulting bounds are horrible with respect to d , they are *independent* of n . This is what defines a core set: a subset whose size only depends on the dimension and the parameter ε , with the property that it approximates the quantity under consideration (in our case all directional widths) arbitrarily well as $\varepsilon \rightarrow 0$. We have seen core sets before: In Chapter 4 we have shown the existence of a subset of P whose size only depends on ε (and *not* even on d) such that its smallest enclosing ball is a $(1 + \varepsilon)$ -approximation for the smallest enclosing ball of P .

For directional widths, the size of a core set depends on d , and it is known that this must be the case. The bound we prove here is not best possible, though.

Proof. (i) We assume that $2M/\varepsilon$ is integer, otherwise we round up to the next larger integer (this explains the ceiling in the statement of the theorem). We tile the bounding cuboid C of P obtained from `MinCuboidApprox`(P, d) with

$$\left(\frac{2M}{\varepsilon} \right)^d$$

scaled-down copies of C which we call *cells* (the scaling factor is $\varepsilon/2M$). A *column* in the resulting tiling is any set of $2M/\varepsilon$ cells that have the same projection onto the hyperplane orthogonal to the vector v_d of the cuboid C . There are $(2M/\varepsilon)^{d-1}$ columns, see Figure 7.2 (left).




Figure 7.2: Constructing coresets for directional width

Within each column, we now find the highest and the lowest cell which contains points from P , see Figure 7.2 (middle).

Then we let S be a set of points obtained by choosing *exactly one* point from each such highest and lowest cell, see Figure 7.2 (right).

It is clear by construction that S has the required size, so we only need to show that it indeed approximates all directional widths according to the prescribed accuracy.

For this, let \mathcal{C} be the union of all cells that contain points of S . Since every point $p \in P$ has a cell from \mathcal{C} above it, and one below (both may coincide with the cell containing p), we know that $p \in \text{conv}(\mathcal{C})$ for all $p \in P$, and this means

$$w_v(P) \leq w_v(\text{conv}(\mathcal{C})) \stackrel{\text{Fact 7.0.2}}{=} w_v(\mathcal{C}), \quad \forall v \in S_{d-1}. \quad (7.8)$$

Now, because any cell c contributing to \mathcal{C} contains a point of S , the width of \mathcal{C} in any direction v cannot be larger than the width of S in direction v , plus *twice* the width of a single cell c in direction v , see Figure 7.3.

By definition of M , we have

$$w_v(c) = \frac{\varepsilon}{2M} w_v(\mathcal{C}) \leq \frac{\varepsilon}{2} w_v(P),$$

and this gives

$$\begin{aligned} w_v(S) &\geq w_v(\mathcal{C}) - 2w_v(c) \\ &\stackrel{(7.8)}{\geq} w_v(P) - \varepsilon w_v(P) \\ &= (1 - \varepsilon)w_v(P), \end{aligned}$$

as desired.

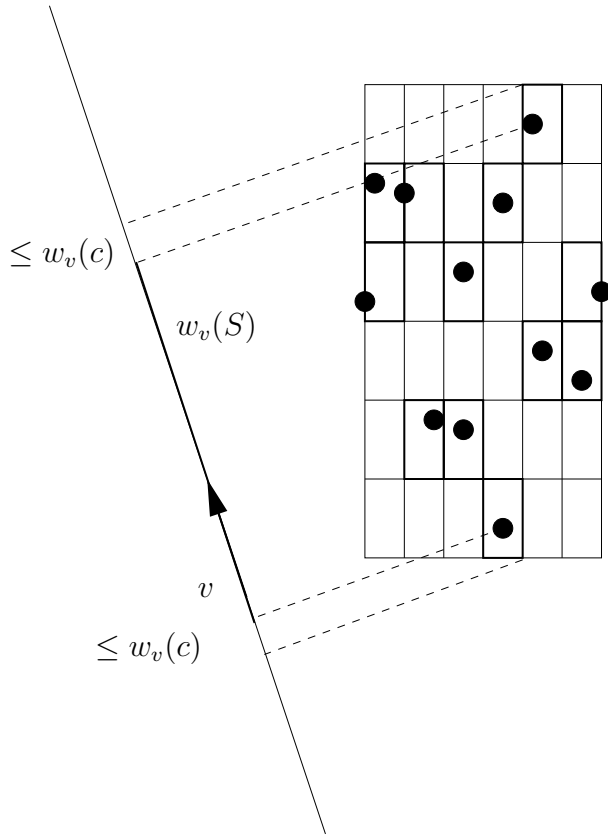


Figure 7.3: $w_v(\mathcal{C}) \leq w_v(S) + 2w_v(c)$

The construction for S we just gave can be turned into an algorithm of the required runtime. First compute in time $O(d^2n)$ the bounding cuboid of C , according to Theorem 6.1.1.

For any point $p \in P$, we have to compute a cell $c \subseteq C$ containing p .¹ Let us identify each cell with a unique d -tuple (a_1, \dots, a_d) of integers from $\{0, \dots, 2M/\varepsilon - 1\}$, where (a_1, \dots, a_d) encodes the cell

$$\{x \in C \mid a_i \leq \frac{2M(x \cdot v_i - \ell_i)}{\varepsilon(u_i - \ell_i)} \leq a_{i+1}\}.$$

We also say that the cell has *index* (a_1, \dots, a_d) . This may look complicated, but recall that by definition of u_i and ℓ_i , we have

$$\ell_i \leq x \cdot v_i \leq u_i, \quad \forall x \in C,$$

so the expression

$$\frac{2M(x \cdot v_i - \ell_i)}{\varepsilon(u_i - \ell_i)}$$

¹Cells overlap along their boundaries, so a point may be in several cells.

runs between 0 and $2M/\varepsilon$. The possible values of a_i therefore subdivide this range into $2M/\varepsilon$ equal-sized intervals, just what we are doing with the cells.

In order to find the index of a cell containing p , we simply compute

$$a_i := \left\lfloor \frac{2M(p \cdot v_i - \ell_i)}{\varepsilon(u_i - \ell_i)} \right\rfloor, \quad i = 1, \dots, d,$$

where we decrease this value to $2M/\varepsilon - 1$ in case the computation yields $2M/\varepsilon$. Assuming that the “floor” function can be evaluated in constant time, the index for p can be computed in $O(d^2n)$ time, $O(dn)$ for each scalar product. (Actually, we have already computed these scalar products during the algorithm `MinCuboidApprox`.)

We next create a hash table of size $O(n)$ for the columns; this simply means, for all n tuples (a_1, \dots, a_d) coming from the points, we insert the $d - 1$ -tuple (a_1, \dots, a_{d-1}) into the hash table, at the same time maintaining the highest and lowest cell in each column (the ones with largest and smallest a_d), along with some point in this cell. Assuming that the hash value for each tuple can be computed in at most $O(d^2)$ time, we stay within the required time bound. From the hash table, we finally extract the core set S . \square

7.3 The dual view: extent of hyperplanes

Let G be a set of nonvertical hyperplanes in \mathbb{R}^d , and let us assume that $g \in G$ is the set

$$g = \{x \in \mathbb{R}^d \mid x_d = \sum_{i=1}^{d-1} g_i x_i + g_d\}.$$

This is the notation of the introductory Chapter. For a point $u \in \mathbb{R}^{d-1}$ and $g \in G$, define

$$g(u) = \sum_{i=1}^{d-1} g_i u_i + g_d$$

to be g “evaluated” at u . The *extent* of G at u is defined as

$$e_u(G) = \max_{g \in G} g(u) - \min_{g \in G} g(u),$$

see Figure 7.4 for an illustration.

Definition 7.3.1 For G a set of hyperplanes in \mathbb{R}^d and $\varepsilon \geq 0$, $F \subseteq G$ is an ε -core set of G with respect to extents if

$$e_u(F) \geq (1 - \varepsilon)e_u(G), \quad \forall u \in \mathbb{R}^{d-1}.$$

This is similar in spirit to ε -core sets for directional widths. We want a subset that approximates *all* possible extents up to a relative error of ε . In fact, as the section title already indicates, this is nothing but the directional width setup in disguise.




Figure 7.4: The extent of a set of nonvertical hyperplanes G at a point u

Theorem 7.3.2 Let $P \subseteq \mathbb{R}^d$ be a finite point set and let $S \subseteq P$. Furthermore, let P^* be the set of nonvertical hyperplanes dual to P according to the definition in the introductory chapter. This means,

$$P^* = \{p^* \mid p \in P\}, \quad p^* = \{x \in \mathbb{R}^d \mid x_d = \sum_{i=1}^{d-1} 2p_i x_i - p_d\},$$

where p_1, \dots, p_d are the coordinates of p . Then the following two statements are equivalent.

(i) S is an ε -core set of P w.r.t. directional widths, meaning that

$$w_v(S) \geq (1 - \varepsilon)w_v(P), \quad \forall v \in S_{d-1}.$$

(ii) S^* is an ε -core set of P^* w.r.t. extents, meaning that

$$e_u(S^*) \geq (1 - \varepsilon)e_u(P^*), \quad \forall u \in \mathbb{R}^{d-1}.$$

In particular, ε -core sets w.r.t. extents exist: simply take G , dualize it to a point set G^* and find a core set F^* with $F \subseteq G$ using Theorem 7.2.1. Then F is the desired core set of G w.r.t. extents (recall that the dual of the dual is the original object; this is needed in the argument).

Proof. Let us first assume that S is an ε -core set of P . Now fix $u \in \mathbb{R}^{d-1}$. Then there is a vector $v \in S_{d-1}$ such that

$$(u_1, \dots, u_{d-1}, -1/2) = \lambda v$$

for suitable $\lambda > 0$. To get v , simply scale $(u_1, \dots, u_{d-1}, -1/2)$ to unit length, see Figure 7.5.

For $p \in P$, let us evaluate the hyperplane p^* at u . We get

$$p^*(u) = \sum_{i=1}^{d-1} 2p_i u_i - p_d = 2p \cdot (u_1, \dots, u_{d-1}, -1/2) = 2\lambda(p \cdot v). \quad (7.9)$$




Figure 7.5: Constructing $v \in S_{d-1}$ from $u \in \mathbb{R}^{d-1}$

This is easily seen to imply (how?) that

$$\begin{aligned} e_u(P^*) &= 2\lambda w_v(P), \\ e_u(S^*) &= 2\lambda w_v(S). \end{aligned}$$

Using (i)—the fact that S is a core set—we get

$$e_u(S^*) = 2|\lambda|w_v(S) \geq 2|\lambda|(1 - \varepsilon)w_v(P) = (1 - \varepsilon)e_u(P^*).$$

This argument works for all u , so S^* is the desired core set and (ii) follows.

Our proof actually works in the other direction (ii) \Rightarrow (i) as well when we are able to construct for *any* given $v \in S_{d-1}$ a point $u \in \mathbb{R}^{d-1}$ such that the equation 7.9 holds for suitable $\lambda \neq 0$. Unfortunately, this cannot be done: vectors v such that $v_d = 0$ have no preimage u under the mapping illustrated in Figure 7.5.

But there is a simple trick. If \tilde{v} is such an exceptional vector, we consider a sequence $v^{(i)}, i \in \mathbb{N}$ of vectors such that $v_d^{(i)} \neq 0$ and $\lim_{i \rightarrow \infty} v^{(i)} = \tilde{v}$. For any $v^{(i)}$ we get a preimage $u^{(i)}$ such that

$$p^*(u^{(i)}) = 2\lambda^{(i)}(p \cdot v^{(i)}), \lambda^{(i)} \neq 0.$$

To get $u^{(i)}$, choose the unique intersection point of the line through 0 and $v^{(i)}$ with the hyperplane $x_d = -1/2$. This intersection point exists if $v_d^{(i)} \neq 0$, but $\lambda^{(i)}$ might be negative.

Still, if condition (ii) of the theorem holds— S^* is a core set for P^* —we can argue as before that

$$w_{v^{(i)}}(S) = \left| \frac{1}{2\lambda^{(i)}} \right| e_{u^{(i)}}(S^*) \geq \left| \frac{1}{2\lambda^{(i)}} \right| (1 - \varepsilon) e_{u^{(i)}}(P^*) = (1 - \varepsilon) w_{v^{(i)}}(P).$$

Since w_v is a continuous function, this inequality also holds in the limit, for \tilde{v} . \square

7.4 Linearization

Why did we develop the dual view of directional width, as explained in the previous section? So far it seems that we didn't get anything new. The new thing is this: the core set algorithm from Section 7.2 can also be used to find a coresets of a set of *polynomials* with respect to extents, see Figure 7.6.

Before we describe how this is done, let us consider an application.




Figure 7.6: The extent of a set G of polynomials over \mathbb{R}^{d-1} at a point u

7.4.1 Core sets for directional width of moving points

We have a set P of *moving points*. Formally, a moving point is a function $p : \mathbb{R} \rightarrow \mathbb{R}^d$ such that $p(t)$ is the position of p at time t . If each coordinate of $p(t)$ is a polynomial in t , we speak about *algebraic motion*.

For example, we may have $p(t) = a + tb$ for fixed $a, b \in \mathbb{R}^d$, in which case all the polynomials are linear, and we have *linear motion*.

Now consider the problem of finding an ε -core set for P with respect to directional widths. The width of P in direction v at time t is defined as

$$w_{v,t}(P) = \max_{p \in P} p(t) \cdot v - \min_{p \in P} p(t) \cdot v.$$

A set $S \subseteq P$ is an ε -core set if

$$w_{v,t}(S) \geq (1 - \varepsilon)w_{v,t}(P), \quad \forall v \in S_{d-1}, t \in \mathbb{R}. \quad (7.10)$$

In order to obtain a picture as in Figure 7.6 for this problem, let us define functions $f_p : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ by

$$f_p(v, t) = p(t) \cdot v, \quad p \in P.$$

In case of algebraic motion, the coordinates of $p(t)$ are polynomials in t ; the coordinates of v are (linear) polynomials in the variables v_1, \dots, v_d , so all f_p are polynomials. Let G be the set of all f_p .

Assume we can find an ε -core set $F \subseteq G$ w.r.t. extents for the family of polynomials G , meaning that

$$e_{v,t}(F) \geq (1 - \varepsilon)e_{v,t}(G), \quad \forall (v, t) \in \mathbb{R}^{d+1},$$

where we define the extent as

$$e_{v,t}(G) = \max_{f_p \in G} f_p(v, t) - \min_{f_p \in G} f_p(v, t) = w_{v,t}(P),$$

as we did it before for a set of hyperplanes. Then, by definition of the f_p , the set $S := \{p \mid f_p \in F\}$ satisfies (7.10) and is therefore the desired core set.

It's a pretty strong core set: it approximates *all* directional widths of the moving point set at *any* time.

Here is the theorem that lets us reduce the extent problem for polynomials to the extent problem for hyperplanes in some higherdimensional space.

Theorem 7.4.1 *Let G be a set of n polynomials in $d - 1$ variables, and let k be the number of different nonconstant monomials over all members of G . Then the problem of finding an ε -core set for G w.r.t. extents can be reduced to the problem of finding an ε -core set of a set of n hyperplanes in \mathbb{R}^{k+1} , w.r.t. extents.*

Before we prove this, let us discuss what this means. Recall that a monomial is simply a product of variables. For example, the polynomial $x^2 + y^2 - 2xy$ is built from three monomials, x^2 , y^2 and xy . If we are dealing with linear polynomials

$$\sum_{i=1}^{d-1} g_i x_i + g_d,$$

(this is the setup of hyperplane extent), there are d monomials $x_1, \dots, x_{d-1}, 1$, but only $d - 1$ of them are nonconstant. In this case, we get $k = d$, and the statement of the theorem is trivial.

Let us next consider our motivating example of core sets for moving points. Assuming linear motion, the polynomials are of the form

$$f_p(v, t) = (a + bt) \cdot v = \sum_{i=1}^d a_i v_i + \sum_{i=1}^d b_i t v_i,$$

so we have a total of $2d$ monomials $v_1, \dots, v_d, t v_1, \dots, t v_d$ in all the f_p . In this case, the theorem holds with $k = 2d$.

Proof. Let us denote the k nonconstant monomials by $\varphi_1, \dots, \varphi_k$. Now we consider the mapping $\Phi : \mathbb{R}^{d-1} \rightarrow \mathbb{R}^k$, given by

$$\Phi(x_1, \dots, x_{d-1}) = (\varphi_1(x_1, \dots, x_{d-1}), \dots, \varphi_k(x_1, \dots, x_{d-1})).$$

Here, $\varphi_i(x_1, \dots, x_{d-1})$ is the evaluation of monomial φ_i at x_1, \dots, x_{d-1} . For example, evaluating x^2 at $(x, y) = (2, 3)$ gives 4.

Now let $f \in G$ be any polynomial, written in the form

$$f = \sum_{i=1}^k \beta_i \phi_i + \beta_{k+1}.$$

Here is the important observation: evaluating the polynomial f at $(x_1, \dots, x_{d-1}) \in \mathbb{R}^{d-1}$ yields the same result as evaluating the *hyperplane*

$$\{y \in \mathbb{R}^{k+1} \mid y_{k+1} = \sum_{i=1}^k \beta_i y_i + \beta_{k+1}\}$$

at $y = \Phi(x_1, \dots, x_{d-1})$.

Having a core set of the latter set of hyperplanes w.r.t. extents then yields a core set of the same quality for the original set of polynomials. The core set for the hyperplanes typically provides much more than we need: it works for the extents at *all* points $y \in \mathbb{R}^k$, while we only use it for points of the form $y = \Phi(x_1, \dots, x_{d-1})$. \square

This technique is called *linearization*. We have already seen an example for it in the chapter on VC-dimension: when you want to know whether a point $x \in \mathbb{R}^d$ is inside, on, or outside a certain ball in \mathbb{R}^d , the *lifting map* has allowed us to reduce this to the question whether the *lifted point*

$$\Phi(x) = \left(x_1, \dots, x_d, \sum_{i=1}^d x_i^2 \right)$$

is below, on, or above a certain hyperplane in \mathbb{R}^{d+1} . In this case, the linearization is more economical, because it has one coordinate for the *sum* of all monomials x_i^2 . Such savings are typical, and in that respect, the bound of Theorem 7.4.1 can be improved in many cases.

Bibliography