Approximating Minimization Diagrams and Generalized Proximity Search[*]

Sariel Har-Peled[†]   Nirman Kumar[‡]

November 9, 2018

Abstract

We investigate the classes of functions whose minimization diagrams can be approximated efficiently in $\mathbb{R}^d$. We present a general framework and a data-structure that can be used to approximate the minimization diagram of such functions. The resulting data-structure has near linear size and can answer queries in logarithmic time. Applications include approximating the Voronoi diagram of (additively or multiplicatively) weighted points. Our technique also works for more general distance functions, such as metrics induced by convex bodies, and the nearest furthest-neighbor distance to a set of point sets. Interestingly, our framework works also for distance functions that do not comply with the triangle inequality. For many of these functions no near-linear size approximation was known before.

1. Introduction

Given a set of functions $\mathcal{F} = \left\{ f_i : \mathbb{R}^d \to \mathbb{R} \mid i = 1, \ldots, n \right\}$, their minimization diagram is the function $f_{\min}(\mathsf{q}) = \min\limits_{i=1,\ldots,n} f_i(\mathsf{q})$, for any $\mathsf{q} \in \mathbb{R}^d$. By viewing the graphs of these functions as manifolds in $\mathbb{R}^{d+1}$, the graph of the minimization diagram, also known as the *lower envelope* of $\mathcal{F}$, is the manifold that can be viewed from an observer at $-\infty$ on the x_{d+1} axis. Given a set of functions $\mathcal{F}$ as above, many problems in Computational Geometry can be viewed as computing the minimization diagram; that is, one preprocesses $\mathcal{F}$, and given a query point q, one needs to compute $f_{\min}(\mathsf{q})$ quickly. This typically requires $n^{O(d)}$ space if one is interested in logarithmic query time. If one is restricted to using linear space, then the query time deteriorates to $O\left(n^{1-O(1/d)}\right)$ [Mat92, Cha10]. There is substantial work on bounding the complexity of the lower envelope in various cases, how to compute it efficiently, and performing range search on them; see the book by Sharir and Agarwal [SA95].

Nearest neighbor. One natural problem that falls into this framework is the nearest neighbor (NN) search problem. Here, given a set P of n data points in a metric space $\mathcal{X}$, we need to preprocess P, such that given a query point $\mathsf{q} \in \mathcal{X}$, one can find (quickly) the point $\mathsf{n}_\mathsf{q} \in \mathsf{P}$ closest to q. Nearest neighbor search is a fundamental task used in numerous domains including machine learning, clustering, document retrieval, databases, statistics, and many others.

To see the connection to lower envelopes, consider a set of data points $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$ in $\mathbb{R}^d$. Next, consider the set of functions $\mathcal{F} = \{f_1, \ldots, f_n\}$, where $f_i(\mathsf{q}) = \|\mathsf{q} - \mathsf{p}_i\|$, for $i = 1, \ldots, n$. The graph of f_i is the set of points $\{(\mathsf{q}, f_i(\mathsf{q})) \mid \mathsf{q} \in \mathbb{R}^d\}$ (which is a cone in $\mathbb{R}^{d+1}$ with apex at $(\mathsf{p}_i, 0)$). Clearly the NN problem is to evaluate the minimization diagram of the functions at a query point q.

More generally, given a set of n functions, one can think of the minimization diagram defining a "distance function", by analogy with the above. The distance of a query point here is simply the "height" of the lower envelope at that point.

Exact nearest neighbor. The exact nearest neighbor problem has a naive linear time algorithm without any preprocessing. However, by doing some nontrivial preprocessing, one can achieve a sub-linear query time. In $\mathbb{R}^d$, this is facilitated by answering point location queries using a Voronoi diagram [dBCvKO08]. However, this approach is only suitable for low dimensions, as the complexity of the Voronoi diagram is $\Theta\left(n^{\lceil d/2 \rceil}\right)$ in the worst case. Specifically, Clarkson [Cla88] showed a data-structure with query time $O(\log n)$ time, and $O\left(n^{\lceil d/2 \rceil + \delta}\right)$ space, where $\delta > 0$ is a prespecified constant (the $O(\cdot)$ notation here hides constants that are exponential in the dimension). One can trade-off the space used and the query time [AM93]. Meiser [Mei93] provided a data-structure with query time $O\left(d^5 \log n\right)$ (which has polynomial dependency on the dimension), where the space used is $O\left(n^{d+\delta}\right)$. These solutions are impractical even for data-sets of moderate size if the dimension is larger than two.

Approximate nearest neighbor. In typical applications, however, it is usually sufficient to return an *approximate nearest neighbor* (**ANN**). Given an $\varepsilon > 0$, a $(1 + \varepsilon)$-ANN, to a query point q, is a point $y \in \mathsf{P}$, such that

$$\|\mathsf{q} - y\| \leq (1 + \varepsilon)\left\|\mathsf{q} - \mathsf{n}_\mathsf{q}\right\|,$$

where $\mathsf{n}_\mathsf{q} \in \mathsf{P}$ is the nearest neighbor to q in P. Considerable amount of work was done on this problem, see [Cla06] and references therein.

In high dimensional Euclidean space, Indyk and Motwani showed that ANN can be reduced to a small number of near neighbor queries [IM98, HIM12]. Next, using locality sensitive hashing they provide a data-structure that answers ANN queries in time (roughly) $\widetilde{O}\left(n^{1/(1+\varepsilon)}\right)$ and preprocessing time and space $\widetilde{O}\left(n^{1+1/(1+\varepsilon)}\right)$; here the $\widetilde{O}(\cdot)$ hides terms polynomial in $\log n$ and $1/\varepsilon$. This was improved to $\widetilde{O}\left(n^{1/(1+\varepsilon)^2}\right)$ query time, and preprocessing time and space $\widetilde{O}\left(n^{1+1/(1+\varepsilon)^2}\right)$ [AI08]. These bounds are near optimal [MNP06].

In low dimensions (i.e., $\mathbb{R}^d$ for small d), one can use linear space (independent of ε) and get ANN query time $O(\log n + 1/\varepsilon^{d-1})$ [AMN+98, Har11]. The trade-off for this logarithmic query time is of course an exponential dependence on d. Interestingly, for this data-structure, the approximation parameter ε is not prespecified during the construction; one needs to provide it only during the query. An alternative approach, is to use Approximate Voronoi Diagrams (AVD), introduced by Har-Peled [Har01], which is a partition of space into regions, of near-linear total complexity, typically with a representative point for each region that is an ANN for any point in the region. In particular, Har-Peled showed that there is such a decomposition of size $O\left((n/\varepsilon^d) \log^2 n\right)$, such that ANN queries can be answered in $O(\log n)$ time. Arya and Malamatos [AM02] showed how to build AVD's of linear complexity (i.e., $O(n/\varepsilon^d)$). Their construction uses Well-Separated Pair Decomposition [CK95]. Further trade-offs between query time and space for AVD's were studied by Arya *et al.* [AMM09].

Generalized distance functions: motivation. The algorithms for approximate nearest neighbor, extend to various metrics in $\mathbb{R}^d$, for example the well known ℓ_p metrics. In particular, previous constructions of AVD's extend to ℓ_p metrics [Har01, AM02] as well. However, these constructions fail even for a relatively simple and natural extension; specifically, multiplicative weighted Voronoi diagrams. Here, every site p, in the given point set P, has a weight ω_p, and the "distance" of a query point q to p is $f_\mathsf{p}(\mathsf{q}) = \omega_\mathsf{p} \|\mathsf{q} - \mathsf{p}\|$. The function f_p is the natural distance function induced by p. As with ordinary Voronoi diagrams, one can define the weighted Voronoi diagram as a partition of space into disjoint regions, one for each site p, such that in the region for p the function f_p is the one realizing the minimum among all the functions induced by the points of P. It is known that, even in the plane, multiplicative Voronoi diagrams can have quadratic complexity, and the minimizing distance function usually does not comply with the triangle inequality. Intuitively, such multiplicative Voronoi diagrams can be used to model facilities where the price of delivery to a client depends on the facility and the distance. Of course, this is only one possible distance function, and there are many other such functions that are of interest (e.g., multiplicative, additive, etc.).

When fast proximity and small space is not possible. Consider a set of segments in the plane, and we are interested in the nearest segment to a query point. Given n such segments and n such query points, this is an extension of Hopcroft's problem, which requires only to decide if there is any of the given points on any of the segments. There are lower bounds (in reasonable models) that show that Hopcroft's problem cannot be solved faster than $\Omega(n^{4/3})$ time [Eri96]. This implies that no multiplicative-error approximation for proximity search in this case is possible, if one insists on near linear preprocessing, and logarithmic query time.

When is fast ANN possible. So, consider a set of geometric objects where each one of them induces a natural distance function, measuring how far a point in space is from this object. Given such a collection of functions, the nearest neighbor for a query point is simply the function that defines the lower envelope "above" the query point (i.e., the object closest to the query point under its distance function). Clearly, this approach allows a generalization of the proximity search problem. In particular, the above question becomes, for what classes of functions, can the lower envelope be approximated up to $(1 + \varepsilon)$-multiplicative error, in logarithmic time? Here the preprocessing space used by the data structure should be near linear.

1.1. Our results

We characterize the conditions that are sufficient to approximate efficiently the minimization diagram of functions. Using this framework, one can quickly, approximately evaluate the lower envelope for large classes of functions that arise naturally from proximity problems. Our data-structure can be constructed in near linear time, uses near linear space, and answers proximity queries in logarithmic time (in constant dimension). Our framework is quite general and should be applicable to many distance functions, and in particular we present the following specific cases where the new data-structure can be used:

(A) **Multiplicative Voronoi diagrams.** Given a set of points P, where the ith point p_i has associated weight $w_i > 0$, for $i = 1, \ldots, n$, consider the functions $f_i(\mathsf{q}) = w_i \|\mathsf{q} - \mathsf{p}_i\|$. The minimization diagram for this set of functions, corresponds to the multiplicative weighted Voronoi diagram of the points. The approach of Arya and Malamatos [AM02] to construct AVD's using WSPD's fails for this problem, as that construction relies on the triangle inequality that the regular Euclidean distance posseses, which does not hold in this case.

3

We provide a near linear space AVD construction for this case. We are unaware of any previous results on AVD for multiplicatively weighted Voronoi diagrams.

(B) **Minkowski norms of fat convex bodies.** Given a bounded symmetric convex body C centered at the origin, it defines a natural metric; that is, for points u and v their distance, as induced by C, denoted by $\|\mathsf{u} - \mathsf{v}\|_C$, is the minimum x such that $xC + \mathsf{u}$ contains v. So, given a set of n data points $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$ and n centrally symmetric and bounded convex bodies $C_1, \ldots, C_n$, we define $f_i(\mathsf{q}) = \|\mathsf{p}_i - \mathsf{q}\|_{C_i}$, for $i = 1, \ldots, n$. Since each point induces a distance by a different convex body, this collection no longer defines a metric, and this makes the problem significantly more challenging. In particular, existing techniques for AVD and ANN cannot be readily applied. Intuitively, the fatness of the associated convex bodies turns out to be sufficient to approximate the associated distance function, see Section 5.2. The negative example for the case of segments presented above, indicates that this condition is also necessary.

(C) **Nearest furthest-neighbor.** Consider a situation where the given input is uncertain; specifically, for the ith point we are given a set of points $\mathsf{P}_i \subseteq \mathbb{R}^d$ where it might lie (the reader might consider the case where the ith point randomly chooses its location out of the points of P_i). There is a growing interest in how to handle such inputs, as real world measurements are fraught with uncertainty, see [DRS09, Agg09, AESZ12, AAH+13] and references therein. In particular, in the worst case, the distance of the query point q to the ith point, is the distance from q to the furthest-neighbor of q in P_i; that is, $\mathcal{F}_i(\mathsf{q}) = \max_{\mathsf{p} \in \mathsf{P}_i} \|\mathsf{q} - \mathsf{p}\|$. Thus, in the worst case, the nearest point to the query is $\mathcal{F}(\mathsf{q}) = \min_i \mathcal{F}_i(\mathsf{q})$. Using our framework we can approximate this function efficiently, using space $\widetilde{O}(n)$, and providing logarithmic query time. Note, that surprisingly, the space requirement is independent of the original input size, and only depends on the number of uncertain points.

Paper organization. In Section 2 we define our framework and prove some basic properties. Since we are trying to make our framework as inclusive as possible, its description is somewhat abstract. In Section 4, we describe the construction of the AVD and its associated data-structure. We describe in Section 5 some specific cases where the new AVD construction can be used. We conclude in Section 6.

2. Preliminaries

For the sake of simplicity of exposition, throughout the paper we assume that all the "action" takes place in the unit cube $[0, 1]^d$. Among other things this implies that all the queries are in this region. This can always be guaranteed by an appropriate scaling and translation of space. The scaling and translation, along with the conditions on functions in our framework, implies that outside the unit cube the approximation to the lower envelope can be obtained in constant time.

## 2.1. Informal description of the technique

Consider n points in the plane $\mathsf{p}_1, \ldots, \mathsf{p}_n$, where the "distance" from the ith point to a query q, is the minimum scaling of an ellipse $\mathcal{E}_i$ (centered at p_i), till it covers q, and let f_i denote this distance function. Assume that these ellipses are fat. Clearly each function f_i defines a deformed cone. Given a query point $\mathsf{q} \in \mathbb{R}^2$, we are interested in the first function graph being hit by a vertical ray shoot upward from $(\mathsf{q}, 0)$. In particular, let $f_{\min}(\mathsf{q}) = \min_{i=1,\ldots,n} f_i(\mathsf{q})$ be the minimization diagram of these functions.

As a first step to computing $f_{\min}(\mathsf{q})$, consider the decision version of this problem. Given a value r, we are interested in deciding if $f_{\min}(\mathsf{q}) \leq r$. That is, we want to decide if $\mathsf{q} \in \bigcup_i(\mathsf{p}_i + r\mathcal{E}_i)$. Of course, this is by itself a computationally expensive task, and as such we satisfy ourselves with an approximate decision to this procedure. Formally, we replace every ellipse by a collection of grid cells (of the right resolution), such that approximately it is enough to decide if the query point lies inside any of these grid cells – if it does, we know that $f_{\min}(\mathsf{q}) \leq (1 + \varepsilon)r$, otherwise $f_{\min}(\mathsf{q}) > r$. Of course, as depicted in the right, since the ellipses are of different sizes, the grid cells generated for each ellipse might belong to different resolutions, and might be of different sizes. Nevertheless, one can perform this point-location query among the marked grid squares quickly using a compressed quadtree.
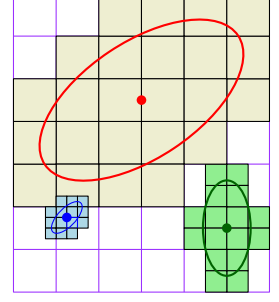
If we were interested only in the case where $f_{\min}(\mathsf{q})$ is guaranteed to be in some interval $[\alpha, \beta]$, then the problem would be easily solvable. Indeed, build a sequence of the above deciders $\mathcal{D}_1, \ldots, \mathcal{D}_m$, where $\mathcal{D}_i$ is for the distance $(1 + \varepsilon)^i\alpha$, and $m = \log_{1+\varepsilon}(\beta/\alpha)$. Clearly, doing a binary search over these deciders with the query point would resolve the distance query.

Sketchable. Unfortunately, in general, there is no such guarantee – which makes the problem significantly more challenging. Fortunately, for truly "large" distances a collection of such ellipses looks like a constant number of ellipse (at least in the approximate case). In the example of the figure above, for large enough distance, the ellipses looks like a single ellipse, as demonstrated in the figure on the right. Slight more formally, if $\bigcup_i(\mathsf{p}_i + r\mathcal{E}_i)$ is connected, then the set $\bigcup_i(\mathsf{p}_i + R\mathcal{E}_i)$ can be $(1 + \varepsilon)$-approximated by a constant number of these ellipses, if $R > \Omega(nr/\varepsilon)$. A family of functions having this property is sketchable. This suggests the problem is easy for very large distances.
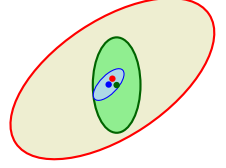
Critical values to search over. The above suggests that connectivity is the underlying property that enables us to simplify and replace a large set of ellipses, by a few ellipses, if we are looking at them from sufficiently far. This implies that the critical values when the level-set of the functions changes its connectivity are the values we should search over during the nearest neighbor search. Specifically, let r_i be the minimal r when the set $\bigcup_{k=1}^n(\mathsf{p}_k + r_i\mathcal{E}_k)$ has $n - i$ connected components, and let $r_1 \leq r_2 \leq \cdots \leq r_n$ be the resulting sequence. Using the above decision procedure, and a binary search, we can find the index j, such that $r_j f_{\min}(\mathsf{q}) \leq r_{j+1}$. Furthermore, the decision procedure for the distance r_j, reports which connected components of $\bigcup_{k=1}^n(\mathsf{p}_k + r_j\mathcal{E}_k)$ contains the query point q. Assume this connected components is formed by the first t functions; that is, $\bigcup_{k=1}^t(\mathsf{p}_k + r_j\mathcal{E}_k)$ is connected and contains q. There are two possibilities:

(A) If $f_{\min}(\mathsf{q}) \in \left[r_j, c_a(t/\varepsilon)r_j\right]$, then a binary search with the decision procedure would approximation $f_{\min}(\mathsf{q})$, where c_a is some constant.

(B) If $f_{\min}(\mathsf{q}) > (t/\varepsilon)r_j$ then this whole cluster of functions can be sketched and replaced by constant number of representative functions, and the nearest-neighbor search can now resolve directly by checking for each function in the sketch, what is the distance of the query point from it.

2.1.1. Challenges

There are several challenges in realizing the above scheme:

(A) We are interested in more general distance functions. To this end, we carefully formalize what conditions the underlying distance functions induced by each point has to fulfill so that our framework applies.

(B) The above scheme requires (roughly) quadratic space to be realized. To reduce the space to near linear, we need be more aggressive about replacing clusters of points/functions by sketches. To this end, we replace our global scheme by a recursive scheme that starts with the "median" critical value, and fork the search at this value using the decision procedure. Now, when continuing the search above this value, we replace every cluster (at this resolution) by its sketch.

(C) Computing this "median" value directly is too expensive. Instead we randomly select a function, we compute the connectivity radius of this single distance function with the remaining functions. With good probability this value turns out to be good.

(D) We need to be very careful to avoid accumulation in the error as we replace clusters by sketches.

2.2. Notations and basic definitions

Given $\mathsf{q} \in \mathbb{R}^d$ and $\mathsf{P} \subseteq \mathbb{R}^d$ a non-empty closed set, the ***distance*** of q to P is $\mathsf{d}(\mathsf{q}, \mathsf{P}) = \min_{x \in \mathsf{P}} \|\mathsf{q} - x\|$.

For a number $\ell > 0$, the ***grid*** of side-length ℓ, denoted by G_ℓ, is the natural tiling of $\mathbb{R}^d$, with cubes of side-length ℓ (i.e. with a vertex at the origin). A cube $\square$ is ***canonical*** if it belongs to G_ℓ, ℓ is a power of 2, and $\square \subseteq [0,1]^d$. Informally, a canonical cube (or cell) is a region that might correspond to a cell in a quadtree having the unit cube as the root region.

Definition 2.1. *To approximate a set $X \subseteq [0,1]^d$, up to distance r, consider the set $\mathsf{G}_{\approx r}(X)$ of all the canonical grid cells of G_ℓ that have a non-empty intersection with X, where $\ell = 2^{\lfloor \log_2(r/\sqrt{d}) \rfloor}$. Let $\cup \mathsf{G}_{\approx r}(X) = \bigcup_{\square \in \mathsf{G}_{\approx r}(X)} \square$, denote the union of cubes of $\mathsf{G}_{\approx r}(X)$.*

Observe that $X \subseteq \cup \mathsf{G}_{\approx r}(X) \subseteq X \oplus \mathsf{B}(0, r)$, where $\oplus$ denotes the Minkowski sum, and $\mathsf{B}(0, r)$ is the ball of radius r centered at the origin.

Definition 2.2. *For $\ell \geq 0$ and a function $f : \mathbb{R}^d \to \mathbb{R}$, the ℓ **sublevel set** of f is the set $f_{\preceq \ell} = \left\{ \mathsf{p} \in \mathbb{R}^d \,\middle|\, f(\mathsf{p}) \leq \ell \right\}$. For a set of functions $\mathcal{F}$, let $\mathcal{F}_{\preceq \ell} = \bigcup_{f \in \mathcal{F}} f_{\preceq \ell}$.*

Definition 2.3. *Given a function f and $\mathsf{q} \in \mathbb{R}^d$ their distance$_f$ is $\mathsf{d}(\mathsf{q}, f) = f(\mathsf{q})$. Given two functions f and g, their **distance$_f$** $\mathsf{d}(f, g)$ is the minimum $l \geq 0$ such that $f_{\preceq l} \cap g_{\preceq l} \neq \emptyset$. Similarly, for two sets of function, $\mathcal{F}$ and $\mathcal{G}$, their **distance$_f$** is*

$$\mathsf{d}(\mathcal{F}, \mathcal{G}) = \min_{f \in \mathcal{F}, g \in \mathcal{G}} \mathsf{d}(f, g).$$

Example 2.4. To decipher these somewhat cryptic definitions, the reader might want to consider the standard settings of regular Voronoi diagrams. Here, we have a set P of n points. The ith point $\mathsf{p}_i \in \mathsf{P}$ induces the natural function $f_i(\mathsf{q}) = \|\mathsf{q} - \mathsf{p}_i\|$. We have:

(A) The graph of f_i in $\mathbb{R}^{d+1}$ is a cone "opening upwards" with an apex at $(\mathsf{p}_i, 0)$.

(B) The ℓ sublevel set of f_i (i.e., $(f_i)_{\preceq \ell}$) is a ball of radius ℓ centered at p_i.

(C) The distance$_f$ of q from f_i is the Euclidean distance between q and p_i.

(D) Consider two subsets of points $X, Y \subseteq \mathsf{P}$ and let $\mathcal{F}_X$ and $\mathcal{F}_Y$ be the corresponding sets of functions. The distance$_f$ $\ell = \mathsf{d}(\mathcal{F}_X, \mathcal{F}_Y)$ is the minimum radius of balls centered at points of X and Y, such that there are two balls from the two sets that intersect; that is, ℓ is half

6

the minimum distance between a point of X and a point of Y. In particular, if the union of balls of radius ℓ centered at X is connected i.e. $(\mathcal{F}_X)_{\preceq \ell}$ is connected, and similarly for Y, then $(\mathcal{F}_X \cup \mathcal{F}_Y)_{\preceq \ell}$ is connected. This is the critical value where two connected components of the sublevel set merge.

The distance$_f$ function behaves to some extent like a distance function: (i) $\mathrm{d}(f,g)$ always exists, and (ii) (symmetry) $\mathrm{d}(f,g) = \mathrm{d}(g,f)$, Also, we have $f_{\preceq \mathrm{d}(f,g)} \neq \emptyset$. We extend the above definition to sets of functions. Note that the triangle inequality does not hold for $\mathrm{d}(\cdot,\cdot)$.

Observation 2.5. *Suppose that f and g are two functions such that $\mathrm{d}(f,g) > 0$ and $\mathsf{q} \in \mathbb{R}^d$. Then, $\max(\mathrm{d}(\mathsf{q},f), \mathrm{d}(\mathsf{q},g)) \geq \mathrm{d}(f,g)$.*

Definition 2.6. *Let $B_1, B_2, \ldots, B_m$ be n connected, nonempty sets in $\mathbb{R}^d$. This collection of sets is **connected** if $\cup_i B_i$ is connected.*

2.2.1. Sketches

A key idea underlying our approach is that is that any set of functions of interest should look like a single (or a small number of functions) from "far" enough. Indeed, given a set of points $\mathsf{P} \subseteq \mathbb{R}^d$, they look like a single point (as far as distance), if the distance from $\mathcal{CH}(\mathsf{P})$ is at least $2\mathsf{diam}(\mathsf{P})/\varepsilon$.

Definition 2.7 (cl($\mathcal{F}$)). *Given a set of functions $\mathcal{G}$, if $\mathcal{G}$ contains a single function then the **connectivity level** $\mathsf{cl}(\mathcal{G})$ is 0; otherwise, it is the minimum $\ell \geq 0$, such that the collection of sets $f_{\preceq \ell}$ for $f \in \mathcal{G}$ is connected, see Definition 2.6.*

Remark 2.8. *It follows from Definition 2.7 that at level $\ell = \mathsf{cl}(\mathcal{G})$, each of the sets $f_{\preceq \ell}$ for $f \in \mathcal{G}$ are nonempty and connected and further their union $\mathcal{G}_{\preceq \ell}$ is also connected. This can be relaxed to require that the intersection graph of the sets $f_{\preceq \ell}$ for $f \in \mathcal{G}$ is connected (this also implies they are nonempty). Notice that, if at level ℓ, the sublevel sets are connected, then the relaxed definition is equivalent to Definition 2.7. However, the relaxed definition introduces more technical baggage, and for all the interesting applications we have, the sublevel sets $f_{\preceq y}$ are connected at all levels y they are nonempty. Therefore, in the interest of brevity, and to keep the presentation simple, we mandate that the sublevel sets be connected at ℓ. In fact, it would not harm to assume that the sublevel sets are connected whenever nonempty.*

Definition 2.9. *Given a set of functions $\mathcal{G}$ and $\delta \geq 0, y_0 \geq 0$, a (δ, y_0)-**sketch** for $\mathcal{G}$ is a (hopefully small) subset $\mathcal{H} \subseteq \mathcal{G}$, such that $\mathcal{G}_{\preceq y} \subseteq \mathcal{H}_{\preceq (1+\delta)y}$, for all $y \geq y_0$.*

It is easy to see that for any $\mathcal{G}, \delta \geq 0, y_0 \geq 0$, if $\mathcal{H} \subseteq \mathcal{G}$ is a (δ, y_0)-sketch, then for any $\delta' \geq \delta, y_0' \geq y_0, \mathcal{H}' \supseteq \mathcal{H}$ it is true that $\mathcal{H}'$ is a (δ', y_0')-sketch for $\mathcal{G}$. Trivially, for any $\delta \geq 0, y_0 \geq 0$, it is true that $\mathcal{H} = \mathcal{G}$ is a (δ, y_0)-sketch.

2.3. Conditions on the functions

We require that the set of functions under consideration satisfy the following conditions.

(P1) **Compactness.** For any $y \geq 0$ and $i = 1, \ldots, n$, the set $(f_i)_{\preceq y}$ is compact.

(P2) **Bounded growth.** For any $f \in \mathcal{F}$, there is a function $\lambda_f : \mathbb{R}^+ \to \mathbb{R}^+$, called the ***growth function***, such that for any $y \geq 0$ and $\varepsilon > 0$, if $f_{\preceq y} \neq \emptyset$, then $\lambda_f(y) \geq \mathsf{diam}(f_{\preceq y})/\zeta$, where ζ is an absolute constant, the ***growth constant***, depending only on the family of functions

and not on n and such that if $\mathsf{q} \in \mathbb{R}^d$ with $\mathsf{d}(\mathsf{q}, f_{\preceq y}) \leq \varepsilon \lambda_f(y)$, then $f(\mathsf{q}) \leq (1+\varepsilon)y$. This is equivalent to $f_{\preceq y} \oplus \mathsf{B}(0, \varepsilon \lambda_f(y)) \subseteq f_{\preceq (1+\varepsilon)y}$, where $\mathsf{B}(\mathsf{u}, r)$ is the ball of radius r centered at u.

(P3) **Existence of a sketch.** Given $\delta > 0$ and a subset $\mathcal{G} \subseteq \mathcal{F}$, there is a $\mathcal{H} \subseteq \mathcal{G}$ with $|\mathcal{H}| = O\left(1/\delta^{\mathsf{c}_{\mathrm{sk}}}\right)$ and $y_0 = O\big(\mathsf{cl}(\mathcal{G})\big(|\mathcal{G}|/\delta\big)^{\mathsf{c}_{\mathrm{sk}}}\big)$ such that, $\mathcal{H}$ is an (δ, y_0)-sketch, where c_{sk} is some positive integer constant that depends on the given family of functions.

We also require some straightforward properties from the computation model:

(C1) $\forall \mathsf{q} \in \mathbb{R}^d$ and $1 \leq i \leq n$, the value $f_i(\mathsf{q}) = \mathsf{d}(\mathsf{q}, f_i)$ is computable in $O(1)$ time.

(C2) For any $y \geq 0, r > 0$ and i, the set of grid cells approximating the sublevel set $(f_i)_{\preceq r}$ of f_i, that is $(f_i)_{\preceq y, \approx r} = \mathsf{G}_{\approx r}\big((f_i)_{\preceq y}\big)$ (see Definition 2.1), is computable in linear time in its size.

(C3) For any $f_i, f_j \in \mathcal{F}, 1 \leq i, j \leq n$ the distance$_f$ $\mathsf{d}(f_i, f_j)$ is computable in $O(1)$ time.

We also assume that the growth function $\lambda_{(f_i)}(y)$ from Condition (P2) be in fact computable easily i.e. in $O(1)$ time.

Remark 2.10. *We will use Condition (C2) for a given y and i only for r at least $\Omega\big(\varepsilon \lambda_{(f_i)}(y)\big)$ i.e. we will use a grid on the sublevel set at a low enough resolution typically ε times its growth function value at that point, which by Condition (C2) is also $\Omega\Big(\varepsilon \mathsf{diam}\big((f_i)_{\preceq y}\big)\Big)$. As such the number of grid cells in the grid used is $O(1/\varepsilon^d)$.*

2.3.1. Properties

The following are basic properties that the functions under consideration have. Since these properties are straightforward but their proof is somewhat tedious, we delegate their proof to Appendix B.

In the following, let $\mathcal{F}$ be a set of functions that satisfy the conditions above.

(L1) For any $f \in \mathcal{F}$, either $f_{\preceq 0} = \emptyset$ or $f_{\preceq 0}$ consists of a single point. (See Lemma B.1$_{\mathrm{p31}}$.)

(L2) If $\mathsf{cl}(\mathcal{G}) = 0$ for any non-empty subset $\mathcal{G}$ then $|\mathcal{G}| = 1$. (See Definition 2.7 and Observation B.2$_{\mathrm{p32}}$.)

(L3) Let $f \in \mathcal{G}$ and $y \geq 0$. For any $\mathsf{u}, \mathsf{v} \in f_{\preceq y}$, we have $\mathsf{uv} \subseteq \mathcal{G}_{\preceq (1+\zeta/2)y}$, where uv denotes the segment joining u to v. (See Lemma B.3$_{\mathrm{p32}}$.)

(L4) Let $A_1, \ldots, A_m \subseteq \mathbb{R}^d$ be compact connected sets, uv be a segment such that $\mathsf{uv} \cap A_i \neq \emptyset$, for $i = 1, \ldots, k$ and $\mathsf{uv} \subseteq \bigcup_{i=1}^k A_i$. Then, the sets $A_1, \ldots, A_k$ are connected. (See Lemma B.4$_{\mathrm{p32}}$.)

(L5) For any $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{F}$, $\delta \geq 0$ and $y \geq 0$, such that $\mathcal{H}$ is a (δ, y)-sketch for $\mathcal{G}$, we have that, $\mathsf{cl}(\mathcal{H}) \leq (1+\delta)(1+\zeta/2) \max(y, \mathsf{cl}(\mathcal{G}))$. (See Lemma B.5$_{\mathrm{p32}}$.)

(L6) Let $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{F}$, such that $\mathcal{H}$ is a (δ, y_0)-sketch for $\mathcal{G}$ for some $\delta \geq 0$ and $y_0 \geq 0$. Let q be a point such that $\mathsf{d}(\mathsf{q}, \mathcal{G}) \geq y_0$. Then we have that $\mathsf{d}(\mathsf{q}, \mathcal{H}) \leq (1+\delta)\mathsf{d}(\mathsf{q}, \mathcal{G})$. (See Lemma B.6$_{\mathrm{p33}}$.)

2.3.2. Computing the connectivity level

We implicitly assume that the above relevant quantities can be computed efficiently. For example given some $\delta > 0$, and y_0 as per the bound in condition (P3), a (δ, y_0)-sketch can be computed in time $O(|\mathcal{G}|/\delta^{\mathsf{c}_{\mathrm{sk}}})$ time. We also assume that $\mathsf{cl}(\mathcal{G})$ can be computed efficiently without resorting to the "brute force" method. The brute force method computes the individual distance$_f$ of the functions and then computes a MST on the graph defined by vertices as the functions and edge lengths as their distance$_f$. Then $\mathsf{cl}(\mathcal{G})$ is the longest edge of this MST.

3. Summary of results

Our main result is the following, the details of which are delegated to Section 4.

Theorem 3.1. *Let $\mathcal{F}$ be a set of n functions in $\mathbb{R}^d$ that complies with our assumptions, see Section 2.3, and has sketch constant $c_{sk} \geq d$. Then, one can build a data-structure to answer ANN for this set of functions, with the following properties:*
- *(A) The query time is $O(\log n + 1/\varepsilon^{c_{sk}})$.*
- *(B) The preprocessing time is $O\big(n\varepsilon^{-2c_{sk}} \log^{2c_{sk}+1} n\big)$.*
- *(C) The space used is $O\big(n\varepsilon^{-d-1-c_{sk}} \log^2 n\big)$.*

One can transform the data-structure into an AVD, and in the process improve the query time (the space requirement slightly deteriorates). See Section 4 for details.

Corollary 3.2. *Let $\mathcal{F}$ be a set of n functions in $\mathbb{R}^d$ that complies with our assumptions, see Section 2.3, and has sketch constant $c_{sk} \geq d$. Then, one can build a data-structure to answer ANN for this set of functions, with the following properties:*
- *(A) The improved query time is $O(\log n)$.*
- *(B) The preprocessing time is $O\big(n/\varepsilon^{O(1)} \log^{2c_{sk}+1} n\big)$.*
- *(C) The space used is $S = O\big(n/\varepsilon^{O(1)} \log^2 n\big)$.*

In particular, we can compute an AVD of complexity $O(S)$ for the given functions. That is, one can compute a space decomposition, such that every region has a single function associated with it, and for any point in this region, this function is the $(1+\varepsilon)$-ANN among the functions of $\mathcal{F}$. Here, a region is either a cube, or the set difference of two cubes.

3.1. Distance functions for which the framework applies

3.1.1. Multiplicative distance functions with additive offsets

We are given n points in $\mathbb{R}^d$, where the point p_i has weight $w_i > 0$, and an **_offset_** $\alpha_i \geq 0$ associated with it, for $i = 1, \ldots, n$. The **_multiplicative distance with offset_** induced by the ith point is $f_i(q) = w_i \|q - p_i\| + \alpha_i$. In Section 5.1 we prove that these distance functions comply with the conditions of Section 2.3, and in particular we get the following result.

Theorem 3.3. *Consider a set P of n points in $\mathbb{R}^d$, where the ith point p_i has additive weight $\alpha_i \geq 0$ and multiplicative weight $w_i > 0$. The ith point induces the additive/multiplicative distance function $f_i(q) = w_i \|q - p_i\| + \alpha_i$. Then one can compute a $(1+\varepsilon)$-AVD for these distance functions, with near linear space complexity, and logarithmic query time. See Theorem 3.1_{p9} for the exact bounds.*

3.1.2. Scaling distance

Somewhat imprecisely, a connected body O centered at a point ρ is α-rounded fat if it is α-fat (that is, there is radius r such that $\mathsf{ball}(\rho, r) \subseteq O \subseteq \mathsf{ball}(\rho, \alpha r)$), and from any point p on the boundary of O the "cone" $\mathcal{CH}(\mathsf{ball}(\rho, r) \cup p)$ is contained inside O (i.e., every boundary point sees a large fraction of the "center" of the object). We also assume that the boundary of each object O has constant complexity.

For such an object, its **_scaling distance_** to a point q from O is the minimum t, such that $q \in tO$ (where the scaling is done around its center ρ). Given n α-rounded fat objects, it is natural to ask for the Voronoi diagram induced by their scaling distance.

These natural distance functions induced by such a set of objects complies with the framework of Section 2.3, see Section 5.2 for details. As such, we get the following result.

Theorem 3.4. *Consider a set $\mathcal{O}$ of α-rounded fat objects in $\mathbb{R}^d$, for some constant α. Then one can compute a $(1 + \varepsilon)$-AVD for the scaling distance functions induced by $\mathcal{O}$, with near linear space complexity, and logarithmic query time. See Theorem 3.1_{p9} and Corollary 3.2 for the exact bounds.*

3.1.3. Nearest furthest-neighbor

For a set of points $\mathsf{S} \subseteq \mathbb{R}^d$ and a point q, the ***furthest-neighbor distance*** of q from S, is $\mathcal{F}_{\mathsf{S}}(\mathsf{q}) = \max_{\mathsf{s} \in \mathsf{S}} \|\mathsf{q} - \mathsf{s}\|$; that is, it is the furthest one might have to travel from q to arrive to a point of S. For example, S might be the set of locations of facilities, where it is known that one of them is always open, and one is interested in the worst case distance a client has to travel to reach an open facility. The function $\mathcal{F}_{\mathsf{S}}(\cdot)$ is known as the ***furthest-neighbor Voronoi*** diagram, and while its worst case combinatorial complexity is similar to the regular Voronoi diagram, it can be approximated using constant size representation (in low dimensions), see [Har99].

Given n sets of points $\mathsf{P}_1, \ldots, \mathsf{P}_n$ in $\mathbb{R}^d$, we are interested in the distance function $\mathcal{F}(\mathsf{q}) = \min_i \mathcal{F}_i(\mathsf{q})$, where $\mathcal{F}_i(\mathsf{q}) = \mathcal{F}_{\mathsf{P}_i}(\mathsf{q})$. This quantity arises naturally when one tries to model uncertainty [AAH$^+$13]; indeed, let P_i be the set of possible locations of the ith point (i.e., the location of the ith point is chosen randomly, somehow, from the set P_i). Thus, $\mathcal{F}_i(\mathsf{q})$ is the worst case distance to the ith point, and $\mathcal{F}(\mathsf{q})$ is the worst-case nearest neighbor distance to the random point-set generated by picking the ith point from P_i, for $i = 1, \ldots, n$. We refer to $\mathcal{F}(\cdot)$ as the ***nearest furthest-neighbor*** distance, and we are interested in approximating it.

We prove in Section 5.3 that the distance functions $\mathcal{F}_1, \ldots, \mathcal{F}_n$ comply with the conditions of the framework, and we get the following result.

Theorem 3.5. *Given n point sets $\mathsf{P}_1, \ldots, \mathsf{P}_n$ in $\mathbb{R}^d$ with a total of m points, and a parameter $\varepsilon > 0$, one can preprocess the points into an AVD, of size $\widetilde{O}(n)$, for the nearest furthest-neighbor distance defined by these point sets. One can now answer $(1 + \varepsilon)$-approximate NN queries for this distance in $O(\log n)$ time. (Note, that the space and query time used, depend only on n, and not on the input size.)*

4. Constructing the **AVD**

The input is a set $\mathcal{F}$ of n functions satisfying the conditions of Section 2.3, and a number $0 < \varepsilon \leq 1$. We preprocess $\mathcal{F}$, such that given a query point q one can compute a $f \in \mathcal{F}$, where $\mathrm{d}(\mathsf{q}, \mathcal{F}) \leq \mathrm{d}(\mathsf{q}, f) \leq (1 + \varepsilon)\mathrm{d}(\mathsf{q}, \mathcal{F})$.

4.1. Building blocks

4.1.1. Near neighbor

Given a set of functions $\mathcal{G}$, a real number $\alpha \geq 0$, and a parameter $\varepsilon > 0$, a ***near-neighbor*** data-structure $\mathcal{D}_{near} = \mathcal{D}_{\mathrm{nr}}(\mathcal{G}, \varepsilon, \alpha)$ can decide (approximately) if a point has distance$_f$ larger or smaller than α. Formally, for a query point q a near-neighbor query answers yes if $\mathrm{d}(\mathsf{q}, \mathcal{G}) \leq \alpha$, and no if $\mathrm{d}(\mathsf{q}, \mathcal{G}) > (1 + \varepsilon)\alpha$. It can return either answer if $\mathrm{d}(\mathsf{q}, \mathcal{G}) \in \left(\alpha, (1 + \varepsilon)\alpha\right]$. If it returns yes, then it also returns a function $f \in \mathcal{G}$ such that $\mathrm{d}(\mathsf{q}, f) \leq (1 + \varepsilon)\alpha$. The query time of this data-structure is denoted by $T_{\leq}(m)$, where $m = |\mathcal{G}|$.

Lemma 4.1. *Given a set of m functions $\mathcal{G} \subseteq \mathcal{F}$, $\alpha > 0$ and $\varepsilon > 0$. One can construct a data structure (which is a compressed quadtree), of size $O\left(m/\varepsilon^d\right)$, in $O\left(m\varepsilon^{-d} \log(m/\varepsilon)\right)$ time, such that*

given any query point $\mathsf{q} \in \mathbb{R}^d$ *one can answer a* $(1 + \varepsilon)$*-approximate near-neighbor query for the distance* α*, in time* $T_\leq(m) = O(\log(m/\varepsilon))$.

Proof: For each $f \in \mathcal{G}$ consider the canonical grid set $\mathsf{G}_{\approx r_f}(f_{\preceq\alpha})$, where $r_f = \varepsilon\lambda_f(\alpha) \geq \varepsilon\mathsf{diam}(f_{\preceq\alpha})/\zeta$, where $\lambda_f(\cdot), \zeta$ are the growth function and the growth constant, see (P2). The sublevel set of interest is $\mathcal{G}_{\preceq\alpha}$ and its approximation is $\mathcal{C} = \bigcup_{f \in \mathcal{G}} \mathsf{G}_{\approx r_f}(f_{\preceq\alpha})$, as the bounded growth condition (P2) implies that $f_{\preceq\alpha} \subseteq \mathsf{G}_{\approx r_f}(f_{\preceq\alpha}) \subseteq f_{\preceq(1+\varepsilon)\alpha}$. The set of canonical cubes $\mathcal{C}$ can be stored in a compressed quadtree $\mathcal{T}$, and given a query point we can decide if a point is covered by some cube of $\mathcal{C}$ by performing a point location query in $\mathcal{T}$.

By Remark 2.10, $\left|\mathsf{G}_{\approx r_f}(f_{\preceq\alpha})\right| = O(\varepsilon^{-d})$. As such, the total number of canonical cubes in $\mathcal{C}$ is $O(m/\varepsilon^d)$, and the compressed quadtree for storing them can be computed in $O(m\varepsilon^{-d}\log(m/\varepsilon))$ time [Har11].

We mark a cell of the resulting quadtree by the function whose sublevel set it arose from (ties can be resolved arbitrarily). During query, if q is found in one of the cells we return yes and the function associated with the cell, otherwise we return no.

If we have that $\mathsf{d}(\mathsf{q}, \mathcal{G}) \leq \alpha$, then the query point q will be found in one of the marked cells, since they cover $\mathcal{G}_{\preceq\alpha}$. As such, the query will return yes. Moreover, if the query does return a yes, then it belongs to a cube of $\mathcal{C}$ that is completely covered by $\mathcal{G}_{\preceq(1+\varepsilon)\alpha}$, as desired. ∎

4.1.2. Interval data structure

Given a set of functions $\mathcal{G}$, real numbers $0 \leq \alpha \leq \beta$, and $\varepsilon > 0$, the ***interval data structure*** returns for a query point q, one of the following:

(A) If $\mathsf{d}(\mathsf{q}, \mathcal{G}) \in \left[\alpha, \beta\right]$, then it returns a function $g \in \mathcal{G}$ such that $\mathsf{d}(\mathsf{q}, g) \leq (1+\varepsilon)\mathsf{d}(\mathsf{q}, \mathcal{G})$. It might also return such a function for values outside this interval.

(B) "$\mathsf{d}(\mathsf{q}, \mathcal{G}) < \alpha$". In this case it returns a function $g \in \mathcal{G}$ such that $\mathsf{d}(\mathsf{q}, g) < \alpha$.

(C) "$\mathsf{d}(\mathsf{q}, \mathcal{G}) > \beta$".

The time to perform an interval query is denoted by $T_r(m, \alpha, \beta)$.

Lemma 4.2. *Given a set of* m *functions* $\mathcal{G}$*, an interval* $[\alpha, \beta]$ *and an approximation parameter* $\tau > 0$*, one can construct an interval data structure of size* $O\bigl(m\tau^{-d-1}\log(4\beta/\alpha)\bigr)$*, in time* $O\bigl(m\tau^{-d-1}\log(4\beta/\alpha)\log(m/\tau)\bigr)$*, such that given a query point* q *one can answer* $(1+\tau)$*-approximate nearest neighbor query for the distances in the interval* $[\alpha, \beta]$*, in time* $T_r(m, \alpha, \beta, f) = O\left(\log\dfrac{m\log(4\beta/\alpha)}{\tau}\right)$.

Proof: Using Lemma 4.1, build a $(1 + \tau/4)$-near neighbor data-structure D_i for $\mathcal{G}$, for distance $r_i = (\alpha/2)(1+\tau/4)^i$, for $i = 0, \ldots, L = \left\lceil\log_{1+\tau/4}(4\beta/\alpha)\right\rceil = O\bigl(\tau^{-1}\log(4\beta/\alpha)\bigr)$. Clearly, an interval query can be answered in three stages:

(A) Perform a point-location query in D_0. If the answer is yes then $\mathsf{d}(\mathsf{q}, \mathcal{G}) < \alpha$. We can also return a function $g \in \mathcal{G}$ with $\mathsf{d}(\mathsf{q}, g) < \alpha$.

(B) Similarly, perform a point-location query in D_L. If the answer is no then $\mathsf{d}(\mathsf{q}, \mathcal{G}) > \beta$ and we are done.

(C) It must be that $\mathsf{d}(\mathsf{q}, \mathcal{G}) \in \left[r_i, r_{i+1}\right]$ for some i. Find this i by performing a binary search on the data-structures $\mathsf{D}_0, \ldots, \mathsf{D}_L$, for the first i such that D_i returns no, but D_{i+1} returns yes. Clearly, D_{i+1} provides us with the desired $(1 + \tau/4)^2$-ANN to the query point.

To get the improved query time, observe that we can overlay these compressed quadtrees $\mathsf{D}_0, \ldots, \mathsf{D}_L$ into a single quadtree. For every leaf (or compressed node) of this quadtree we compute the original node with the lowest value covering this node. Clearly, finding the desired distance can now be resolved by a single point-location query in this overlay of quadtrees. The total

size of these quadtrees is $S = O\big(L\big(m/\tau^d\big)\big)$, and the total time to compute these quadtrees is $T_1 = O\big(L\big(m/\tau^d\big)\log(m/\tau)\big)$, and the time to compute their overlay is $O(S\log L)$. The time to perform a point-location query in the overlayed quadtree is $O(\log S)$. ∎

Lemma 4.2 readily implies that if somehow a priori we know the nearest neighbor distance$_f$ lies in an interval of values of polynomial spread, then we would get the desired data-structure by just using Lemma 4.2. To overcome this unbounded spread problem, we would first argue that, under our assumptions, there are only linear number of intervals where interesting things happen to the distance$_f$ function.

4.1.3. Connected components of the sublevel sets

Given a finite set X and a partition of it into disjoint sets $X = X_1 \cup \cdots \cup X_k$, let this partition be denoted by $\langle X_1, \ldots, X_k \rangle_X$. For $1 \le i \le k$, each X_i is a **part** of the partition.

Definition 4.3. *For two partitions $P_A = \langle A_1, \ldots, A_k \rangle_X$ and $P_B = \langle B_1, \ldots, B_l \rangle_X$ of the same set X, P_B is a **refinement** of P_A, denoted by $P_B \sqsubseteq P_A$, if for any B_i there exists a set A_{j_i}, such that $B_i \subseteq A_{j_i}$. In the other direction, P_A is a **coarsening** of P_B.*

Observation 4.4. *Given partitions Π, Ξ of a finite set X, if $\Pi \sqsubseteq \Xi$ then $|\Xi| \le |\Pi|$.*

Definition 4.5. *Given partitions $\Pi = \langle X_1, \ldots, X_k \rangle_X \sqsubseteq \Xi = \big\langle X_1', \ldots, X_{k'}' \big\rangle_X$, let $\phi(\Pi, \Xi, i)$ be the function that return the set of indices of sets in Π whose union is $X_i' \in \Xi$.*

Observation 4.6. *Given partitions $\Pi \sqsubseteq \Xi$ of a set X with n elements. The partition function $\phi(\Pi, \Xi, \cdot)$ can be computed in $O(n)$ time. For any $1 \le i \le |\Xi|$, the set $\phi(\Pi, \Xi, i)$ can be returned in $O\big(|\phi(\Pi, \Xi, i)|\big)$ time, and its size can be returned in $O(1)$ time.*

Definition 4.7. *For $\mathcal{G} \subseteq \mathcal{F}$ and $\ell > 0$, consider the intersection graph of the sets $f_{\preceq \ell}$, for all $f \in \mathcal{G}$. Each connected component is a **cluster** of $\mathcal{G}$ at level ℓ. And the partition of $\mathcal{G}$ by these clusters, denoted by $\mathsf{C}(\mathcal{G}, \ell)$, is the ℓ-**clustering** of $\mathcal{G}$.*

The values ℓ at which the ℓ-clustering of $\mathcal{F}$ changes are, intuitively, the critical values when the sublevel set of $\mathcal{F}$ changes and which influence the AVD. These values are critical in trying to decompose the nearest neighbor search on $\mathcal{F}$ into a search on smaller sets.

Observation 4.8. *If $0 \le a \le b$ then $\mathsf{C}(\mathcal{G}, a) \sqsubseteq \mathsf{C}(\mathcal{G}, b)$.*

The following lemma testifies that we can approximate the ℓ-clustering quickly, for any number ℓ.

Lemma 4.9. *Given $\mathcal{G} \subseteq \mathcal{F}$, $\ell \ge 0$, and $\varepsilon > 0$, one can compute, in $O\big(\frac{m}{\varepsilon^d}\log(m/\varepsilon)\big)$ time, a partition $\Psi = \Psi_\varepsilon(\mathcal{G}, \ell)$, such that $\mathsf{C}(\mathcal{G}, \ell) \sqsubseteq \Psi \sqsubseteq \mathsf{C}(\mathcal{G}, (1 + \varepsilon)\ell)$, where $m = |\mathcal{G}|$.*

Proof: For each $f \in \mathcal{G}$, tile the sublevel sets $(f)_{\preceq \ell}$ by canonical cubes of small enough diameter, such that bounded growth condition (P2) assures all cubes are inside $(f)_{\preceq(1+\varepsilon)\ell}$. To this end, for $f \in \mathcal{G}$, set $r_f = \varepsilon\lambda_f(\ell) \ge \varepsilon\mathsf{diam}(f_{\preceq\ell})/\zeta$, and compute the set $\mathcal{C}_f = \bigcup_{f \in \mathcal{G}} \Big(\mathsf{G}_{\approx r_f}\big((f)_{\preceq\ell}\big)\Big)$, see Definition 2.1. It is easy to verify that we have that

$$(\mathcal{G})_{\preceq\ell} \subseteq \cup\mathcal{C}_f \subseteq (\mathcal{G})_{\preceq(1+\varepsilon)\ell}. \tag{1}$$

By assumption, we have that $|\mathcal{C}_f| = O(1/\varepsilon^d)$, and the total number of canonical cubes in all the sets $\mathcal{C}_f$ for $f \in \mathcal{G}$ is $O(m/\varepsilon^d)$. We throw all these canonical cubes into a compressed quadtree, this takes $O\big((m/\varepsilon^d)\log(m/\varepsilon)\big)$ time. Here, every node of the compressed quadtree is marked if it belongs to some of these sets, and if so, to which of the sets. Two sets $\cup\mathcal{C}_f$ and $\cup\mathcal{C}_g$ intersect, if and only if there are two canonical cubes, in these two sets, such that they overlap; that is, one of them is a sub-cube of the other. Initialize a union-find data-structure, and traverse the compressed quadtree using **DFS**, keeping track of the current connected component, and performing a union operation whenever encountering a marked node (i.e., all the canonical nodes associated with it, are unionized into the current connected component). Finally, we perform a union operation for all the cells in $\mathcal{C}_f$, for all $f \in \mathcal{G}$. Clearly, this results in the desired connected components of the intersection graph of $\cup\mathcal{C}_f$ (note, that we consider two sets as intersecting only if their interiors intersect). Translating each such connected set of canonical cubes back to the functions that gave rise to them, results in the desired partition. ∎

Remark 4.10. *The partition Ψ computed by Lemma 4.9 is monotone, that is, for $\ell \leq \ell'$ and $\varepsilon \leq \varepsilon'$, we have $\Psi_\varepsilon(\mathcal{G}, \ell) \sqsubseteq \Psi_{\varepsilon'}(\mathcal{G}, \ell')$. Moreover, for each cluster $C \in \Psi_\varepsilon(\mathcal{G}, \ell)$, we have that $\mathsf{cl}(C) \leq (1 + \varepsilon)\ell$.*

4.1.4. Computing a splitting distance

Definition 4.11. *Given a partition $\Psi = \Psi_\varepsilon(\mathcal{G}, \ell)$ of $\mathcal{G}$, with $m = |\Psi|$ clusters, a distance x is a* **splitting distance** *if $m/4 \leq |\Psi_1(\mathcal{G}, x/4)|$ and $|\Psi_1(\mathcal{G}, x)| \leq (7/8)m$.*

Lemma 4.12. *Given a partition $\Psi = \Psi_\varepsilon(\mathcal{G}, \ell)$ of $\mathcal{G}$, one can compute a splitting distance for it, in expected $O(n(\log n + t))$ time, where $n = |\mathcal{G}|$ and t is the maximum cluster size in Ψ.*

Proof: For each cluster $C \in \Psi$, let r_C be its distance$_f$ from all the functions in $\mathcal{G} \setminus C$; that is $r_C = \min_{f \in C} \min_{g \in \mathcal{G} \setminus C} \mathrm{d}(f, g)$. Note that $r_C \geq \ell$. Now, let $r_1 \leq r_2 \leq \cdots \leq r_m$ be these distance$_f$ distances for the m clusters of Ψ. We randomly pick a cluster $C \in \Psi$ and compute $\ell' = r_C$ for it by brute force – computing the distance$_f$ of each function of C with the functions of $\mathcal{G} \setminus C$.

Let i be the rank of $\ell' = r_C$ among $r_1, \ldots, r_m$. With probability $1/2$, we have that $m/4 \leq i \leq (3/4)m$. If so we have that:

- (A) All the clusters that correspond to $r_i, \ldots, r_m$ are singletons in the partition $\Psi_1(\mathcal{G}, \ell/4)$, as the distance$_f$ of each one of these clusters is larger than ℓ'. We conclude that $|\Psi_1(\mathcal{G}, \ell'/4)| \geq m/4$.
- (B) All the clusters of Ψ that correspond to $r_1, \ldots, r_i$ are contained inside a larger cluster of $\Psi_1(\mathcal{G}, \ell')$ (i.e., they were merged with some other cluster). But then, the number of clusters in $\Psi_1(\mathcal{G}, \ell')$ is at most $(7/8)m$. Indeed, put an edge between such a cluster, to the cluster realizing the smallest distance$_f$ with it. This graph has at least $e \geq m/4$ edges, and it is easy to see that each component of size at least 2 in the underlying undirected graph has the same number of edges as vertices. As such the number of singleton components is at most $m - e$ while the number of components of size at least 2 is at most $e/2$. It follows that the total number of components is at most $m - e/2 \leq 7m/8$. Since each such component corresponds to a cluster in $\Psi_1(\mathcal{G}, \ell')$ the claim is proved.

Now, compute $\Psi_1(\mathcal{G}, \ell')$ and $\Psi_1(\mathcal{G}, \ell'/4)$ using Lemma 4.9. With probability at least half they have the desired sizes, and we are done. Otherwise, we repeat the process. In each iteration we spend $O(n(\log n + t))$ time, and the probability for success is half. As such, in expectation the number of rounds needed is constant. ∎

13

```
Search( 𝒢, ϒ, q )
   // 𝒢:  set of functions
   // ϒ = Ψ₁(𝒢, ℓ) for some value ℓ
   // Invariant:   d(q, 𝒢) > ℓN
   if |ϒ| = 1 then
        return d(q, 𝒢) = min_{f∈𝒢} d(q, f)          (*)
   x ← compute a splitting distance of ϒ, see Lemma 4.12

   // Perform an interval approximate nearest
   //     neighbor query on the interval [x/8N, x8N]
   //     for the set 𝒢, see Lemma 4.2.
   if d(q, 𝒢) ∈ [x/8N, x8N²] or (1 + ε/4)-ANN found then
        return nearest function found by the
                  (1 + ε/4)-approximate interval query.
   if d(q, 𝒢) < x/8N then
        f ← 2-approximate near neighbor query on 𝒢
                  and distance x/8, see Lemma 4.1.
        Find cluster C ∈ Ψ₁(𝒢, x/4), such that f ∈ C,
                  see Lemma 4.9.
        return Search( C, ϒ[C], q )
   if d(q, 𝒢) > x8N² then
        return Search( 𝒢, Ψ₁(𝒢, xN), q )          (**)
```

Figure 4.1: Search algorithm: We are given a query point q, and an approximation parameter $\varepsilon > 0$. The quantity N is a parameter to be specified shortly. Initially, we call this procedure on the set of functions $\mathcal{F}$ with Υ being the partition of $\mathcal{F}$ into singletons (i.e., $\ell = 0$). Here, $\Upsilon[C]$ denotes the partition of C induced by the partition Υ.

4.2. The search procedure

4.2.1. An initial "naive" implementation

The search procedure is presented in Figure 4.1.

Lemma 4.13. **Search**$(\mathcal{G}, \Upsilon, q)$ *returns a function* $f \in \mathcal{G}$, *such that* $d(q, f) \leq (1 + \varepsilon) d(q, \mathcal{G})$. *The depth of the recursion of* **Search** *is* $h = O(\log n)$, *where* $n = |\mathcal{G}|$.

Proof: The proof is by induction on the size of Υ. If $|\Upsilon| = 1$, then the function realizing $d(q, \mathcal{G})$ is returned, and the claim is true.

Let x be the computed splitting distance of Υ. Next, the procedure perform an $(1 + \varepsilon/4)$-approximate interval nearest-neighbor query for q on the range $[x/8N, x8N]$. If this computed the approximate nearest neighbor then we are done.

Otherwise, it must be that either $d(q, \mathcal{G}) < x/8N$ or $d(q, \mathcal{G}) > 8Nx$, and significantly, we know which of the two options it is:

(A) If $d(q, \mathcal{G}) < x/8N$ then doing an approximate near-neighbor query on $\mathcal{G}$ and distance $x/8$, returns a function $f \in \mathcal{G}$ such that $d(q, f) \leq x/4$. Clearly, the nearest neighbor to q must be in the cluster containing f in the partition $\Psi_1(\mathcal{G}, x/4)$, and **Search** recurses on this cluster. Now, by induction, the returned ANN is correct.

14

Since x is a splitting distance of Υ, see Definition 4.11, we have $|\Upsilon|/4 \leq |\Psi_1(\mathcal{G}, x/4)|$ and $\Upsilon \sqsubseteq \Psi_1(\mathcal{G}, x/4)$. As such, since C is one of the clusters of $\Psi_1(\mathcal{G}, x/4)$, the induced partition of C by Υ (i.e., $\Upsilon[C]$), can have at most $(1 - 1/4)|\Upsilon|$ clusters.

(B) Otherwise, we have $d(q, \mathcal{G}) > x \cdot 8N$. Since x is a splitting distance, we have that $|\Psi_1(\mathcal{G}, x)| \leq (7/8)|\Upsilon|$, see Definition 4.11. We recurse on $\mathcal{G}$, and a partition that has fewer clusters, and by induction, the returned answer is correct.

In each step of the recursion, the partition shrunk by at least a fraction of $7/8$. As such, after a logarithmic number of recursive calls, the procedure is done. ∎

4.2.2. But where is the beef? Modifying Search to provide fast query time

The reader might wonder how we are going to get an efficient search algorithm out of **Search**, as the case that Υ is a single cluster, still requires us to perform a scan on all the functions in this cluster and compute their distance$_f$ from the query point q. Note however, we have the invariant that the distance of interest is polynomially larger than the connectivity level of each of the clusters of Υ. In particular, precomputing for all the sets of functions such that (*) might be called on, their $\varepsilon/8$-sketches, and answering the query by computing the distance on the sketches, reduces the query time to $O(1/\varepsilon^{c_{sk}} + \log^2 n)$ (assuming that we precomputed all the data-structures used by the query process). Indeed, an interval query takes $O(\log n)$ time, and there $O(\log n)$ such queries. The final query on the sketch takes time proportional to the sketch size which is $O(1/\varepsilon^{c_{sk}})$.

As such, the major challenge is not making the query process fast, but rather building the search structure quickly, and arguing that it requires little space.

4.2.3. Sketching a sketch

To improve the efficiency of the preprocessing for **Search**, we are going to use sketches more aggressively. Specifically, for each of the clusters of Υ, we can compute their δ-sketches, for $\delta = \varepsilon/(8h) = O(\varepsilon/\log n)$, see Lemma 4.13. From this point on, when we manipulate this cluster, we do it on its sketch. To make this work set $N = n^{4c_{sk}}$, see (P3)$_{p8}$ and Lemma B.5.

The only place in the algorithm where we need to compute the sketches, is in (**) in Figure 4.1. Specifically, we compute $\Psi_1(\mathcal{G}, xN)$, and for each new cluster $C \in \Psi_1(\mathcal{G}, xN)$, we combine all the sketches of the clusters $D \in \Upsilon$ such that $D \subseteq C$ into a single set of functions. We then compute a δ-sketch for this set, and this sketch is this cluster from this point on. In particular, the recursive calls to **Search** would send the sketches of the clusters, and not the clusters themselves. Conceptually, the recursive call would also pass the minimum distance where the sketches are active – it is easy to verify that we use these sketches only at distances that are far away and are thus allowable (i.e., the sketches represent the functions they correspond to, well in these distances).

Importantly, whenever we compute such a new set, we do so for a distance that is bigger by a polynomial factor (i.e., N) than the values used to create the sketches of the clusters being merged. Indeed, observe that $x > \ell$ and as such xN is N times bigger than ℓ (an upper bound on the value used to compute the input sketches).

As such, all these sketches are valid, and can be used at this distance (or any larger distance). Of course, the quality of the sketch deteriorates. In particular, since the depth of recursion is h, the worst quality of any of the sketches created in this process is at most $(1 + \delta)^h \leq 1 + \varepsilon/4$.

Significantly, before using such a sketch, we would shrink it by computing a $\varepsilon/8$-sketch of it. This would reduce the sketch size to $O(1/\varepsilon^{c_{sk}})$. Note, however, that this still does not help us as far as recursion - we must pass the larger δ-sketches in the recursive call of (**).

This completes the description of the search procedure. It is still unclear how to precompute all the data-structures required during the search. To do that, we need to better understand what the search process does.

4.3. The connectivity tree, and the preprocessing

Given a set of functions $\mathcal{F}$, create a tree tracking the connected components of the MST of the functions. Formally, initially we start with n singletons (which are the leafs of the tree) that are labeled with the value zero, and we store them in a set $\mathcal{F}$ of active nodes. Now, we compute for each pair of sets of functions $X, Y \in \mathcal{F}$ the distance$_f$ $\mathrm{d}(X, Y)$, and let X', Y' be the pair realizing the minimum of this quantity. Merge the two sets into a new set $Z = X' \cup Y'$, create a new node for this set having the node for X' and Y' as children, and set its label to be $\mathrm{d}(X', Y')$. Finally, remove X' and Y' from $\mathcal{F}$ and insert Z into it. Repeat till there is a single element in $\mathcal{F}$. Clearly, the result is a tree that tracks the connected components of the MST.

To make the presentation consistent, let $\mathrm{d}_{\approx}(X, Y)$ be the minimum x such that $\Psi_1(X \cup Y, x)$ is connected. Computing $\mathrm{d}_{\approx}(X, Y)$ can be done by computing $\mathrm{d}_{\approx}(f, g)$ for each pair of functions separately. This in turn, can be done by first computing $\alpha = \mathrm{d}(f, g)$ and observing that r is between $\alpha/4$ and α. In particular, r must be a power of two, so there are only 3 candidate values to consider, and which is the right one can be decided using Lemma 4.9.

So, in the above, we use $\mathrm{d}_{\approx}(\cdot, \cdot)$ instead of $\mathrm{d}(\cdot, \cdot)$, and let $\mathcal{H}$ be the resulting tree. For a value ℓ, let $\mathrm{L}_{\mathcal{H}}(\ell)$ be the set of nodes such that their label is smaller than ℓ, but their parent label is larger than ℓ. It is easy to verify that $\mathrm{L}_{\mathcal{H}}(\ell)$ corresponds to $\Psi = \Psi_1(\mathcal{F}, \ell)$; indeed, every cluster $C \in \Psi$ corresponds to a node $u \in \mathrm{L}_{\mathcal{H}}(\ell)$, such that the set of functions stored in the leaves of the subtree of u, denoted by $\mathrm{F}(u)$ is C. The following can be easily proved by induction.

Lemma 4.14. *Consider a recursive call* **Search**$(\mathcal{G}, \Upsilon, \mathsf{q})$ *made during the search algorithm execution. Then* $\mathcal{G} = \mathrm{F}(u)$, *and* $\Upsilon = \left\{ \mathrm{F}(v) \,\middle|\, v \in \mathrm{L}_{\mathcal{H}}(\ell) \text{ and } v \text{ is in the subtree of } u \right\}$.
That is, a recursive call of **Search** *corresponds to a subtree of* $\mathcal{H}$.

Of course, not all possible subtrees are candidates to be such a recursive call. In particular, **Search** can now be interpreted as working on a subtree T of $\mathcal{H}$, as follows:
  (A) If T is a single node u, then find the closet function to $\mathrm{F}(u)$. Using the sketch this can be done quickly.
  (B) Otherwise, computes a distance x, such that the number of nodes in the level $\mathrm{L}_T(x)$ is roughly half the number of leaves of T.
  (C) Using interval data-structure determine if the distance$_f$ $\mathrm{d}(\mathsf{q}, \mathrm{F}(T))$ is in the range $[x/8N, x8N^2]$. If so, we found the desired ANN.
  (D) If $\mathrm{d}(\mathsf{q}, \mathrm{F}(T)) > x8N^2$ then continue recursively on portion of T above $\mathrm{L}_T(x)$.
  (E) If $\mathrm{d}(\mathsf{q}, \mathrm{F}(T)) < x/8N$ then we know the node $u \in \mathrm{L}_T(x)$ such that the ANN belongs to $\mathrm{F}(u)$. Continue the search recursively on the subtree of T rooted at u.

That is, **Search** breaks T into subtrees, and continues the search recursively on one of the subtrees. Significantly, every such subtree has constant fraction of the size of T, and every edge of T belongs to a single such subtree.

The preprocessing now works by precomputing all the data-structures required by **Search**. Of course, the most natural approach would be to precompute $\mathcal{H}$, and build the search tree by simulating the above recursion on $\mathcal{H}$. Fortunately, this is not necessary, we simulate running **Search**, and investigate all the different recursive calls. We thus only use the above $\mathcal{H}$ in analyzing the preprocessing running time. See Figure 4.1$_{\mathrm{p14}}$.

In particular, given a subtree T with m edges, the corresponding partition Υ would have at most m sets. Each such set would have a δ-sketch, and we compute a $\varepsilon/8$-sketch for each one of these sketches. Namely, the input size here is $M = O(m/\delta^{c_{sk}})$. Computing the $\varepsilon/8$-sketches for each one of these sketches reduces the total number of functions to $M' = O(m/\varepsilon^{c_{sk}})$, and takes $U_1 = O(M/\varepsilon^{c_{sk}}) = O\big(m(\varepsilon\delta)^{-c_{sk}}\big)$ time, see Section 2.3.2. Computing the splitting distance, using Lemma 4.12, takes $U_2 = O(M'\log M' + 1/\varepsilon^{c_{sk}}) = O(m\varepsilon^{-c_{sk}}\log m)$ time. Computing the interval data-structure Lemma 4.2 takes $U_3 = O\big(M'\varepsilon^{-d-1}\log n\log M'\big)$ time, and requires $S_1 = O\big(M'\varepsilon^{-d-1}\log n\big)$ space. This breaks T into edge disjoint subtrees $T_1, \ldots, T_t$, and we compute the search data-structure for each one of them separately (each one of these subtrees is smaller by a constant fraction of the original tree). Finally, we need to compute the δ-sketches for the clusters sent to the appropriate recursive calls, and this takes $U_4 = O(M/\delta^{c_{sk}})$, by Section 2.3.2.

Every edge of the tree T gets charged for the amount of work spent in building the top level data-structure. That is, the top level amortized work each edge of T has to pay is

$$O\Big((U_1 + U_2 + U_3 + U_4)/m\Big)$$
$$= O\Big((\varepsilon\delta)^{-c_{sk}} + \varepsilon^{-c_{sk}}\log m + \varepsilon^{-d-1-c_{sk}}\log^2 n + \delta^{-2c_{sk}}\Big)$$
$$= O\big(\varepsilon^{-2c_{sk}}\log^{2c_{sk}} n\big),$$

assuming $c_{sk} \geq 2$. Since an edge of T gets charged at most $O(\log n)$ times by this recursive construction, we conclude that the total preprocessing time is $O\big(n\varepsilon^{-2c_{sk}}\log^{2c_{sk}+1} n\big)$.

As for the space, we have by the same argumentation, that each edge requires $O(\log n \cdot (S_1/m)) = \big(\varepsilon^{-d-1-c_{sk}}\log^2 n\big)$. As such, the overall space used by the data-structure is $\big(n\varepsilon^{-d-1-c_{sk}}\log^2 n\big)$. As for the query time, it boils down to $O(\log n)$ interval queries, and then scanning one $O(\varepsilon)$-sketch. As such, this takes $O\big(\log^2 n + 1/\varepsilon^{c_{sk}}\big)$ time.

4.4. The result

Restatement of Theorem 3.1$_{\mathbf{p9}}$. Let $\mathcal{F}$ be a set of n functions in $\mathbb{R}^d$ that complies with our assumptions, see Section 2.3, and has sketch constant $c_{sk} \geq d$. Then, one can build a data-structure to answer ANN for this set of functions, with the following properties:

    (A) The query time is $O(\log n + 1/\varepsilon^{c_{sk}})$.
    (B) The preprocessing time is $O\big(n\varepsilon^{-2c_{sk}}\log^{2c_{sk}+1} n\big)$.
    (C) The space used is $O\big(n\varepsilon^{-d-1-c_{sk}}\log^2 n\big)$.

Proof: The query time stated above is $O\big(\log^2 n + 1/\varepsilon^{c_{sk}}\big)$. To get the improved query time, we observe that **Search**, performs a sequence of point-location queries in a sequence of interval near neighbor data-structures (i.e., compressed quadtrees), and then it scans a set of functions of size $O(1/\varepsilon^{c_{sk}})$ to find the ANN. We take all these quadtrees spread through our data-structure, and assign them priority, where a quadtree $\mathcal{Q}_1$ has higher priority than a compressed quadtree $\mathcal{T}_2$ if $\mathcal{Q}_1$ is queried after $\mathcal{Q}_2$ for any search query. This defines an acyclic ordering on these compressed quadtrees. Overlaying all these compressed quadtrees together, one needs to return for the query point, the leaf of the highest priority quadtree that contains the query point. This can be easily done by scanning the compressed quadtree, and for every leaf computing the highest priority leaf that contains it (observe, that here we are overlaying only the nodes in the compressed quadtrees that are marked by some sublevel set – nodes that are empty are ignored).

A tedious but straightforward induction implies that doing a point-location query in the resulting quadtree is equivalent to running the search procedure as described above. Once we found

the leaf that contains the query point, we scan the sketch associated with this cell, and return the computed nearest-neighbor. ∎

Restatement of Corollary 3.2$_{\textbf{p9}}$. Let $\mathcal{F}$ be a set of n functions in $\mathbb{R}^d$ that complies with our assumptions, see Section 2.3, and has sketch constant $\mathsf{c_{sk}} \geq d$. Then, one can build a data-structure to answer ANN for this set of functions, with the following properties:

    (A) The improved query time is $O(\log n)$.
    (B) The preprocessing time is $O\big(n/\varepsilon^{O(1)} \log^{2\mathsf{c_{sk}}+1} n\big)$.
    (C) The space used is $S = O\big(n/\varepsilon^{O(1)} \log^2 n\big)$.

In particular, we can compute an AVD of complexity $O(S)$ for the given functions. That is, one can compute a space decomposition, such that every region has a single function associated with it, and for any point in this region, this function is the $(1+\varepsilon)$-ANN among the functions of $\mathcal{F}$. Here, a region is either a cube, or the set difference of two cubes.

Proof: We build the data-structure of Theorem 3.1, except that instead of linearly scanning the sketch during the query time, we preprocess each such sketch for an exact point-location query; that is, we compute the lower envelope of the sketch and preprocess it for vertical ray shooting [AE98]. This would require $O\big(1/\varepsilon^{O(1)}\big)$ space for each such sketch, and the linear scanning that takes $O(1/\varepsilon^{O(1)})$ time, now is replaced by a point-location query that takes $O\big(\log 1/\varepsilon^{O(1)}\big) = O(\log 1/\varepsilon) = O(\log n)$, as desired.

As for the second part, observe that every leaf of the compressed quadtree is the set difference of two canonical grid cells. The lower envelope of the functions associated with such a leaf, induce a partition of this leaf into regions with total complexity $O\big(1/\varepsilon^{O(1)}\big)$. ∎

5. Applications

We present some concrete classes of functions that satisfy our framework, and for which we construct AVD's efficiently.

5.1. Multiplicative distance functions with additive offsets

As a warm-up we present the simpler case of additively offset multiplicative distance functions. The results of this section are almost subsumed by more general results in Section 5.2. Here the sublevel sets look like expanding balls but there is a time lag before the balls even come into existence i.e. sublevel sets are empty up-to a certain level, this corresponds to the additive offsets. In Section 5.2 the sublevel sets are more general fat bodies but there is no additive offset. The results in the present section essentially give an AVD construction of approximate weighted Voronoi diagrams. More formally, we are given a set of points $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$. For $i = 1, \ldots, n$, the point p_i has weight $w_i > 0$, and a constant $\alpha_i \geq 0$ associated with it. We define $f_i(\mathsf{q}) = w_i \|\mathsf{q} - \mathsf{p}_i\| + \alpha_i$. Let $\mathcal{F} = \{f_1, \ldots, f_n\}$. We have, $(f_i)_{\preceq y} = \emptyset$ for $y < \alpha_i$ and $(f_i)_{\preceq y} = \mathsf{B}\big(\mathsf{p}_i, \frac{y-\alpha_i}{w_i}\big)$ for $y \geq \alpha_i$. Checking conditions (C1) and (C2) is trivial. As for (C3) we have the following easy lemma,

Lemma 5.1. *For any $1 \leq i, j \leq n$ we have*

$$\mathbb{d}(f_i, f_j) = \max\left(\max(\alpha_i, \alpha_j), \|\mathsf{p}_i - \mathsf{p}_j\| \frac{w_i w_j}{w_i + w_j} + \frac{\alpha_i w_j + \alpha_j w_i}{w_i + w_j}\right).$$

Proof: The ith distance function is $f_i(\mathsf{q}) = w_i \|\mathsf{q} - \mathsf{p}_i\| + \alpha_i$. As such, for $y < \max(\alpha_i, \alpha_j)$ either $(f_i)_{\preceq y} = \emptyset$, or $(f_j)_{\preceq y} = \emptyset$ and $(f_i)_{\preceq y} \cap (f_j)_{\preceq y} = \emptyset$. For $y \geq \max(\alpha_i, \alpha_j)$, we have

$$f_i(\mathsf{q}) \leq y \quad \implies \quad w_i \|\mathsf{q} - \mathsf{p}_i\| + \alpha_i \leq y \quad \implies \quad \|\mathsf{q} - \mathsf{p}_i\| \leq \frac{y - \alpha_i}{w_i},$$

which implies that $\mathsf{q} \in \mathsf{B}\left(\mathsf{p}_i, \frac{y - \alpha_i}{w_i}\right)$; that is, we have $(f_i)_{\preceq y} = \mathsf{B}\left(\mathsf{p}_i, \frac{y - \alpha_i}{w_i}\right)$ and $(f_j)_{\preceq y} = \mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right)$.

Now, if $\mathsf{p}_i = \mathsf{p}_j$ then the distance$_f$ distance between the two functions is the minimal value such that their sublevel sets are not empty, and this is $\max(\alpha_i, \alpha_j)$. In particular, the given expression

$$\alpha = \max\left(\max(\alpha_i, \alpha_j), \|\mathsf{p}_i - \mathsf{p}_j\| \frac{w_i w_j}{w_i + w_j} + \frac{\alpha_i w_j + \alpha_j w_i}{w_i + w_j}\right)$$

evaluates to $\max(\alpha_i, \alpha_j)$, as desired.

If $\mathsf{p}_i \neq \mathsf{p}_j$ the sublevel sets intersect for the first time when the balls $\mathsf{B}\left(\mathsf{p}_i, \frac{y - \alpha_i}{w_i}\right)$ and $\mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right)$ touch at a point that belongs to the segment $\mathsf{p}_i\mathsf{p}_j$. Clearly then we have

$$\|\mathsf{p}_i - \mathsf{p}_j\| = \frac{y - \alpha_i}{w_i} + \frac{y - \alpha_j}{w_j} \quad \implies \quad w_i w_j \|\mathsf{p}_i - \mathsf{p}_j\| = w_j(y - \alpha_i) + w_i(y - \alpha_j)$$

$$\implies \quad (w_j + w_j) y = w_i w_j \|\mathsf{p}_i - \mathsf{p}_j\| + w_j \alpha_i + w_i \alpha_j$$

$$\implies \quad y = \|\mathsf{p}_i - \mathsf{p}_j\| \frac{w_i w_j}{w_i + w_j} + \frac{\alpha_i w_j + \alpha_j w_i}{w_i + w_j}$$

Lemma 5.2. *Given $1 \leq i, j \leq n$ such that $w_i \leq w_j$. Suppose $y \geq \max(\alpha_i, \alpha_j)$. Then, $(f_j)_{\preceq y} \subseteq (f_i)_{\preceq(1+\delta)y}$ if and only if $y \geq \frac{\|\mathsf{p}_i - \mathsf{p}_j\| + \alpha_i/w_i - \alpha_j/w_j}{(1+\delta)/w_i - 1/w_j}$.*

Proof: For $y \geq \max(\alpha_i, \alpha_j)$ we have that $(f_i)_{\preceq y} = \mathsf{B}\left(\mathsf{p}_i, \frac{y - \alpha_i}{w_i}\right)$ and $(f_j)_{\preceq y} = \mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right)$. If $\mathsf{p}_i = \mathsf{p}_j$ then for any y such that $\frac{(1+\delta)y - \alpha_i}{w_i} \geq \frac{y - \alpha_j}{w_j}$ we will have that $(f_j)_{\preceq y} \subseteq (f_i)_{\preceq(1+\delta)y}$. Clearly this condition is also necessary. It is easy to verify that this is equivalent to the desired expression.

Consider the case $\mathsf{p}_i \neq \mathsf{p}_j$. Sufficiency: Notice that for any $\mathsf{u} \in \mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right)$ we have $\|\mathsf{u} - \mathsf{p}_i\| \leq \|\mathsf{p}_i - \mathsf{p}_j\| + \|\mathsf{p}_j - \mathsf{u}\| \leq \|\mathsf{p}_i - \mathsf{p}_j\| + \frac{y - \alpha_j}{w_j}$ by the triangle inequality. Therefore, if $\frac{(1+\delta)y - \alpha_i}{w_i} \geq \|\mathsf{p}_i - \mathsf{p}_j\| + \frac{y - \alpha_j}{w_j}$, then $\mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right) \subseteq \mathsf{B}\left(\mathsf{p}_i, \frac{(1+\delta)y - \alpha_i}{w_i}\right)$. This is exactly the stated condition. Indeed, by rearrangement,

$$y((1 + \delta)/w_i - 1/w_j) \geq \|\mathsf{p}_i - \mathsf{p}_j\| + \alpha_i/w_i - \alpha_j/w_j.$$

Necessity: Notice that $\mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right)$ has a boundary point at distance $\frac{y - \alpha_j}{w_j}$ from p_j on the directed line from p_i to p_j on the other side of p_j as p_i, while $\mathsf{B}\left(\mathsf{p}_i, \frac{(1+\delta)y - \alpha_i}{w_i}\right)$ has the intercept of $\frac{(1+\delta)y - \alpha_i}{w_i} - \|\mathsf{p}_i - \mathsf{p}_j\|$. For the condition to hold it must be true that $\frac{(1+\delta)y - \alpha_i}{w_i} - \|\mathsf{p}_i - \mathsf{p}_j\| \geq \frac{y - \alpha_j}{w_j}$, which is also the stated condition. ∎

It is easy to see that compactness (P1) and bounded growth (P2) hold for the set of functions $\mathcal{F}$ (for (P2) we can take the growth function $\lambda_{(f_i)}(y) = (y - \alpha_i)/w_i$ for $y \geq \alpha_i$ and the growth constant ζ to be 2). The following lemma proves the sketch property (P3).

19

Lemma 5.3. *For any $\mathcal{G} \subseteq \mathcal{F}$ and $\delta > 0$ there is a (δ, y_0)-sketch $\mathcal{H} \subseteq \mathcal{G}$ with $|\mathcal{H}| = 1$ and $y_0 = 3\mathsf{cl}(\mathcal{G}) |\mathcal{G}| /\delta$.*

Proof: If $|\mathcal{G}| = 1$ we can let $\mathcal{H} = \mathcal{G}$ and the result is easily seen to be true. Otherwise, let $l = \mathsf{cl}(\mathcal{G})$ for brevity. Observe that $l \geq \max\limits_{i: f_i \in \mathcal{G}} \alpha_i$, as otherwise some $(f_i)_{\preceq l} = \emptyset$ and cannot be part of a connected collection of sets. Let $|\mathcal{G}| = m \geq 2$ and let $\mathcal{G} = \{f_1, \ldots, f_m\}$, and assume that we have $w_1 \leq w_i, 1 \leq i \leq m$. We let $\mathcal{H} = \{f_1\}$, the function with the minimum associated weight. We are restricted to the range $l \geq \alpha_i, 1 \leq i \leq m$ so, $(f_i)_{\preceq l}$ is the ball $\mathsf{B}\left(\mathsf{p}_i, \frac{l - \alpha_i}{w_i}\right)$ for each $1 \leq i \leq m$. Since $\mathcal{G}_{\preceq l}$ is connected, it must be true that for $2 \leq j \leq m$ there exist a sequence of distinct indices, $1 = i_1, i_2, \ldots, i_{k-1}, i_k = j$ such that $\mathsf{B}\left(\mathsf{p}_{i_r}, \frac{l - \alpha_{i_r}}{w_{i_r}}\right) \cap \mathsf{B}\left(\mathsf{p}_{i_{r+1}}, \frac{l - \alpha_{i_{r+1}}}{w_{i_{r+1}}}\right) \neq \emptyset$ for $1 \leq r \leq k - 1$. By Lemma 5.1 we can write that,

$$l \geq \frac{\left\| \mathsf{p}_{i_r} - \mathsf{p}_{i_{r+1}} \right\| + \frac{\alpha_{i_r}}{w_{i_r}} + \frac{\alpha_{i_{r+1}}}{w_{i_{r+1}}}}{\frac{1}{w_{i_r}} + \frac{1}{w_{i_{r+1}}}}.$$

Rearranging,

$$\left\| \mathsf{p}_{i_r} - \mathsf{p}_{i_{r+1}} \right\| \leq l \left(\frac{1}{w_{i_r}} + \frac{1}{w_{i_{r+1}}} \right) - \left(\frac{\alpha_{i_r}}{w_{i_r}} + \frac{\alpha_{i_{r+1}}}{w_{i_{r+1}}} \right)$$

$$\leq \frac{2l}{w_1},$$

as $w_1 \leq w_i$, for $1 \leq i \leq m$. It follows by the triangle inequality and the above, that $\left\| \mathsf{p}_{i_1} - \mathsf{p}_{i_m} \right\| \leq \sum_{r=1}^{m-1} \left\| \mathsf{p}_{i_r} - \mathsf{p}_{i_{r+1}} \right\| \leq \frac{2(m-1)l}{w_1} \leq \frac{2ml}{w_1}$. Thus we have,

$$\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| \leq \frac{2ml}{w_1}, \tag{2}$$

for $j = 1, \ldots, m$. Let $y_0 = \frac{3l|\mathcal{G}|}{\delta} = \frac{3ml}{\delta}$. Then, for $y \geq y_0$ we have that,

$$y \geq \frac{3ml}{\delta} = \frac{\frac{2ml}{w_1} + \frac{ml}{w_1}}{\frac{\delta}{w_1}} \geq \frac{\frac{2ml}{w_1} + \frac{l}{w_1}}{\frac{\delta}{w_1}},$$

for $m \geq 2$. Using Eq. (2) and the above, we have for $y \geq y_0$ since $l \geq \alpha_1$,

$$y \geq \frac{\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| + \frac{l}{w_1}}{\frac{\delta}{w_1}} \geq \frac{\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| + \frac{\alpha_1}{w_1}}{\frac{\delta}{w_1}}.$$

It follows that for $y \geq y_0$,

$$y \geq \frac{\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| + \frac{\alpha_1}{w_1}}{\frac{\delta}{w_1}} \geq \frac{\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| + \frac{\alpha_1}{w_1} - \frac{\alpha_j}{w_j}}{\frac{\delta}{w_1} + \left(\frac{1}{w_1} - \frac{1}{w_j} \right)}$$

$$= \frac{\left\| \mathsf{p}_1 - \mathsf{p}_j \right\| + \frac{\alpha_1}{w_1} - \frac{\alpha_j}{w_j}}{\frac{(1+\delta)}{w_1} - \frac{1}{w_j}},$$

as $w_1 \leq w_j$ for $1 \leq j \leq m$. Thus, by Lemma 5.2, $\mathsf{B}\left(\mathsf{p}_j, \frac{y - \alpha_j}{w_j}\right) \subseteq \mathsf{B}\left(\mathsf{p}_1, \frac{(1+\delta)y - \alpha_1}{w_1}\right)$, for $y \geq y_0$ and therefore by definition, $\mathcal{H}$ is a (δ, y_0)-sketch for $\mathcal{G}$. ■
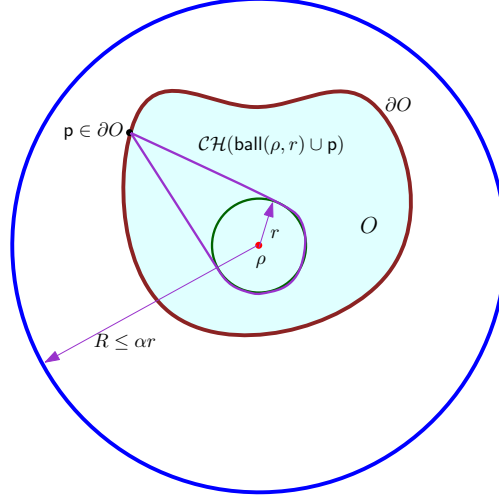
Figure 5.1: Being α-rounded fat.

We thus get the following result.

Restatement of Theorem 3.3$_{p9}$. *Consider a set* P *of* n *points in* $\mathbb{R}^d$, *where the* i*th point* p_i *has additive weight* $\alpha_i \geq 0$ *and multiplicative weight* $w_i > 0$. *The* i*th point induces the additive/multiplicative distance function* $f_i(\mathsf{q}) = w_i \|\mathsf{q} - \mathsf{p}_i\| + \alpha_i$. *Then one can compute a* $(1+\varepsilon)$-*AVD for these distance functions, with near linear space complexity, and logarithmic query time. See Theorem 3.1$_{p9}$ for the exact bounds.*

5.2. Scaling distance – generalized polytope distances

Let $O \subseteq \mathbb{R}^d$ be a compact set homeomorphic to $\mathsf{B}(0, 1)$ and containing a "center" point ρ in its interior. Then O is **_star shaped_** if for any point $\mathsf{v} \in O$ the entire segment $\rho\mathsf{v}$ is also in O. Naturally, any convex body O with any center $\rho \in O$ is star shaped. The t-**_scaling_** of O with a center ρ is the set $tO = \left\{ t(\mathsf{v} - \rho) + \rho \,\middle|\, \mathsf{v} \in O \right\}$.

Given a star shaped object O with a center ρ, the **_scaling distance_** of a point q from O is the minimum t, such that $\mathsf{p} \in tO$, and let $F_O(\mathsf{q})$ denote this distance function. Note that, for any $y \geq 0$, the sublevel set $(F_O)_{\preceq y}$ is the y-scaling of O, that is $(F_O)_{\preceq y} = yO$.

Note, that for a point $\mathsf{p} \in \mathbb{R}^d$, if we take $O = \mathsf{B}(\mathsf{p}, 1)$ with center p, then $F_O(\mathsf{q}) = \|\mathsf{p} - \mathsf{q}\|$. That is, this distance notion is a strict extension of the Euclidean distance.

Henceforward, for this section, we assume that an object O contains the origin in its interior and the origin is the designated center, unless otherwise stated.

Definition 5.4. *Let* $O \subseteq \mathbb{R}^d$ *be a star shaped object centered at* ρ. *We say that* O *is* α-*fat if there is a number* r *such that,* $\mathsf{ball}(\rho, r) \subseteq O \subseteq \mathsf{ball}(\rho, \alpha r)$.

Definition 5.5. *Let* O *be a star shaped object centered at* ρ. *We say that* O *is* α-*rounded fat if there is a radius* r *such that, (i)* $\mathsf{ball}(\rho, r) \subseteq O \subseteq \mathsf{ball}(\rho, \alpha r)$ *and, (ii) For every point* p *in the boundary of* O, *the cone* $\mathcal{CH}(\mathsf{ball}(\rho, r) \cup \mathsf{p})$, *lies within* O, *see Figure 5.1.*

By definition any α-rounded fat object is also α-fat. However, it is not true that a α-fat object is necessarily α'-rounded fat for any α', that is even allowed to depend on α. The following useful result is easy to see, also see Figure 5.2 for an illustration.
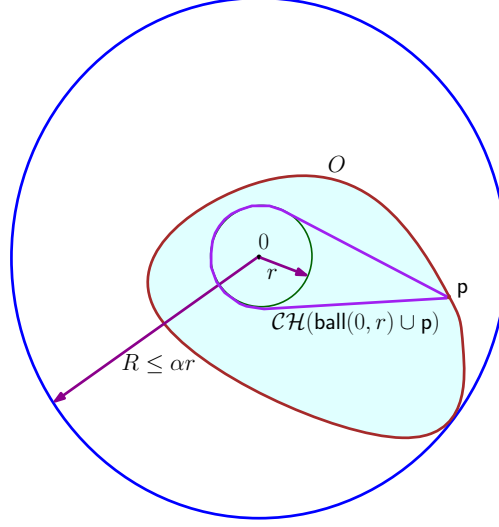
Figure 5.2: A α-fat convex body is α-rounded fat.

Lemma 5.6. *Let O be a α-fat object. If O is convex then O is also α-rounded fat.*

Given a set $\mathcal{O} = \{O_1, O_2, \ldots, O_n\}$ of n star shaped objects, consider the set $\mathcal{F}$ of n scaling distance functions, where the ith function, for $i = 1, \ldots, n$ is $f_i = F_{O_i}$. We assume that the boundary of each object O_i has constant complexity.

We next argue that $\mathcal{F}$ complies with the framework of Section 2.3. Using standard techniques, we can compute the quantities required in conditions (C1)–(C3)$_{\text{p8}}$ including the diameter of the sublevel set $\mathsf{diam}(yO_i) = y\mathsf{diam}(O_i)$. Also, trivially we have that condition (P1)$_{\text{p7}}$ is satisfied as the sublevel sets are dilations of the O_i and are thus compact by definition. The next few lemmas establish that both bounded growth (P2) and the sketch property (P3) are also true, if the objects are also α-rounded fat for some constant α.

Lemma 5.7. *Given $\alpha > 0$, suppose O is a star shaped object that is α-rounded fat. Then for any $c \geq 2\alpha$ and any $y \geq 0, \varepsilon > 0$ we have that $yO \oplus \mathsf{B}(0, (\varepsilon/c)\mathsf{diam}(yO)) \subseteq (1 + \varepsilon)yO$; that is,*

$$(F_O)_{\preceq y} \oplus \mathsf{B}\left(0, (\varepsilon/c)\mathsf{diam}\left((F_O)_{\preceq y}\right)\right) \subseteq (F_O)_{\preceq(1+\varepsilon)y}.$$

Proof: Since $(F_O)_{\preceq y} = yO$ we show that $yO \oplus \mathsf{B}(0, (\varepsilon/c)\mathsf{diam}(yO)) \subseteq (1 + \varepsilon)yO$. Let r be the radius guaranteed by Definition 5.5 for O. Clearly $\mathsf{diam}(yO) = y\mathsf{diam}(O) \leq 2y\alpha r$. We show that for every $\mathsf{p} \in \partial yO$ we have that $\mathsf{p} + \mathsf{B}(0, (\varepsilon/c)\mathsf{diam}(yO)) \subseteq (1+\varepsilon)yO$ [1]. Let $\mathsf{p} \in \partial yO$. It is sufficient to show that, $\mathsf{B}(\mathsf{p}, (2\varepsilon y\alpha r/c)) \subseteq (1 + \varepsilon)yO$. Clearly $\mathsf{p}' = (1 + \varepsilon)\mathsf{p} \in \partial(1 + \varepsilon)yO$. Since the cone, $\mathcal{CH}(\mathsf{B}(\rho, (1 + \varepsilon)yr) \cup \mathsf{p}')$ is in $(1 + \varepsilon)yO$, it is clear that the ball of radius,

$$x = \|\mathsf{p}' - \mathsf{p}\| \frac{\|\mathsf{p}'\|}{(1 + \varepsilon)yr} = \|\mathsf{p}' - \mathsf{p}\| \frac{\|\mathsf{p}\|}{yr}$$

is completely within $(1 + \varepsilon)yO$, see Figure 5.3. Now, $\|\mathsf{p} - \mathsf{p}'\| = \varepsilon \|\mathsf{p}\| \geq \varepsilon yr$, and $\|\mathsf{p}\| \geq yr$. It follows that $x \geq \varepsilon yr$. If we choose $c \geq 2\alpha$, the claimed result is easily seen to hold. ∎

---

[1] Topological arguments can show that for objects homeomorphic to balls, if this is true for boundary points p, then for any $\mathsf{p} \in O$, $\mathsf{p} + \mathsf{B}(0, (\varepsilon/c)\mathsf{diam}(yO)) \subseteq (1 + \varepsilon)yO$. We omit the technical argument here.
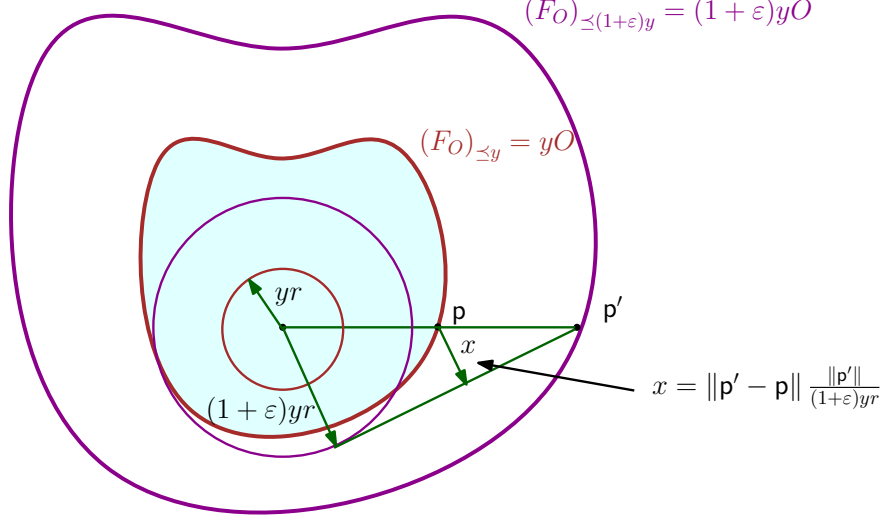
Figure 5.3: The $(1 + \varepsilon)$ expansion of yO contains $\mathsf{B}(\mathsf{p}, x)$.

By the above lemma we can take the growth function $\lambda_{F_{O_i}}(y) = \mathsf{diam}\left((F_{O_i})_{\preceq y}\right) = y\mathsf{diam}(O_i)$ and the growth constant, see $(\mathrm{P2})_{\mathrm{p7}}$, for the set of functions F_{O_i} to be $\zeta = c = 2\alpha$. If the object O is α-fat but not α'-rounded fat for any constant $\alpha' > 0$ then it may be that its scaling distance function grows arbitrarily quickly and thus fails to comply with our framework, see Figure 5.4. It is not hard to see that Lemma 5.7 implies that bounded growth (P2) is satisfied for all the functions $f_1, \ldots, f_n$ when the objects under consideration $O_1, \ldots, O_n$ are α-rounded fat. To show that condition (P3) is satisfied, is slightly harder.

Lemma 5.8. *Let $\mathcal{O}$ be a set of n star shaped objects $O_1, O_2, \ldots, O_n$. Let $\alpha \geq 1$ be any constant. Suppose that $O_1, \ldots, O_n$ are α-rounded fat. Then, for any $\delta > 0$, there is a subset $\mathcal{I} \subseteq \{1, 2, \ldots, n\}$ with $|\mathcal{I}| = O(\delta^{-d})$, such that for all $y \geq 0$, we have*

$$\bigcup_{i \in [n]} yO_i \subseteq \bigcup_{i \in \mathcal{I}} (1 + \delta)yO_i.$$

Moreover, for every $i \in \mathcal{I}$ we have that $\mathsf{diam}(O_i) = \Omega(\max_i \mathsf{diam}(O_i))$.

Proof: Recall our convention, that here all the bodies are centered at the origin. Clearly it is sufficient to show this for $y = 1$. Let r_i for $i = 1, \ldots, n$ be the radius of the ball satisfying the conditions of Definition 5.5 for object O_i. Assume that $r_1 \geq r_2 \cdots \geq r_n$. If $\alpha r_j \leq r_1$ for some j, then $O_j, \ldots, O_n$ are contained in O_1, we can add 1 to the set $\mathcal{I}$. From now we assume that for each $1 \leq j \leq n$, $\alpha r_j > r_1$ i.e. we can ignore the sufficiently small objects. Our index set $\mathcal{I}$ is a subset of these prefix indices for which, $\alpha r_i \geq r_1$. As such,

$$\mathsf{diam}(O_i) \geq 2r_i \geq \frac{2}{\alpha}r_1 \geq \frac{4}{\alpha^2}(2\alpha r_1) \geq \frac{4}{\alpha^2} \max_{1 \leq i \leq n} \mathsf{diam}(O_i),$$

for any $i \in \mathcal{I}$. It is easy to see that, $\bigcup_{i \in [n]} O_i \subseteq \bigcup_{i \in [n]} \mathsf{B}(0, \alpha r_i) \subseteq \mathsf{B}(0, \alpha r_1)$. We tile the ball $\mathsf{B}(0, \alpha r_1)$ with cubes of diameter at most $\delta \alpha r_1/c'$ where c' is a constant that we determine shortly. Notice that the number of such cubes is $O(\delta^{-d})$. Let $\mathcal{C}$ denote the set of these cubes. If $\mathsf{c} \cap O_i \neq \emptyset$ for some $1 \leq i \leq n$ and $\mathsf{c} \in \mathcal{C}$ then we add c to a set $\mathcal{A}$ and i to our index set $\mathcal{I}$. Notice that we choose at most one object among all objects that might intersect c. Now, $\bigcup_{i \in [n]} O_i \subseteq \bigcup_{\mathsf{c} \in \mathcal{A}} \mathsf{c}$, as $\bigcup_{\mathsf{c} \in \mathcal{A}} \mathsf{c}$
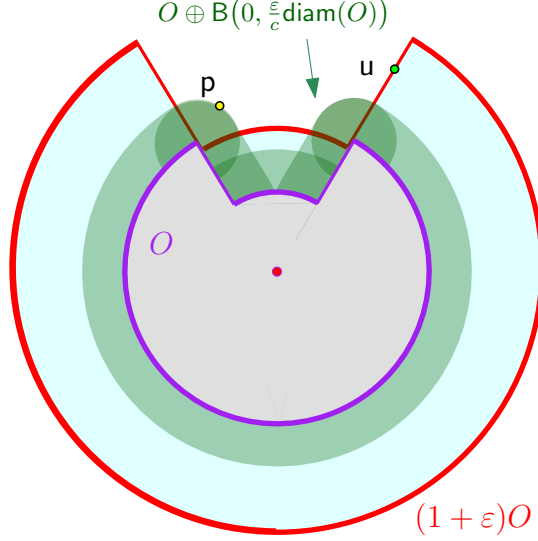
Figure 5.4: The object O is α-fat but not α'-rounded fat. In particular, the point p is in $O \oplus \mathsf{B}(0, (\varepsilon/c)\mathsf{diam}(O))$ but not in $(1+\varepsilon)O$. In particular, the scaling distance function is discontinuous at u.

covers $\mathsf{B}(0, \alpha r_1)$. Observe, $|\mathcal{I}| \le |\mathcal{A}| \le |\mathcal{C}| = O(\delta^{-d})$. We show that it is possible to choose c' so large that, $\bigcup_{\mathsf{c}\in\mathcal{A}} \mathsf{c} \subseteq \bigcup_{i\in\mathcal{I}} (1+\delta)O_i$. Since $\mathsf{c}\cap O_i \ne \emptyset$ and $\mathsf{diam}(\mathsf{c}) \le \delta\alpha r_1/c'$, $\mathsf{c} \subseteq O_i \oplus \mathsf{B}(0, \delta\alpha r_1/c')$. We choose c' large enough so that $\delta\alpha r_1/c' \le \delta\mathsf{diam}(O_i)/c$ where $c = 2\alpha$ is the constant from Lemma 5.7. Then we will have by Lemma 5.7,

$$
\begin{aligned}
\mathsf{c} &\subseteq O_i \oplus \mathsf{B}\big(0, \delta\alpha r_1/c'\big) \subseteq O_i \oplus \mathsf{B}(0, \delta\mathsf{diam}(O_i)/c) \\
&\subseteq (1+\delta)O_i,
\end{aligned}
$$

proving the claim. Now,

$$
\delta\alpha r_1/c' \le \delta\alpha^2 r_i/c' \le \frac{\delta\alpha^2 \mathsf{diam}(O_i)}{2c'} \le \frac{\delta\mathsf{diam}(O_i)}{c},
$$

if $c' = c\alpha^2/2 = \alpha^3$. ∎

Lemma 5.9. *Let $\alpha \ge 1$ be any constant. Let O be a star shaped object that is α-rounded fat, and let $\delta > 0$. Let $\mathsf{u} \in \mathbb{R}^d$ with $\|\mathsf{u}\| \le \delta\mathsf{diam}(O)/c$ where $c = 2\alpha$. Then we have that $O + \mathsf{u} \subseteq (1+\delta)O$.*

Proof: We have, $O + \mathsf{u} \subseteq O \oplus \mathsf{B}(0, \|\mathsf{u}\|) \subseteq O \oplus \mathsf{B}(0, \delta\mathsf{diam}(O)/c)$ as $\|\mathsf{u}\| \le \delta\mathsf{diam}(O)/c$. The result follows by appealing to Lemma 5.7. ∎

Lemma 5.10. *For $i = 1, \ldots, n$, let O_i be a star shaped object in $\mathbb{R}^d$ centered at a point p_i. Let $\mathcal{O} = \{O_1, \ldots, O_n\}$, $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$, and $\mathcal{F} = \{f_i \mid 1 \le i \le n\}$, where $f_i = F_{O_i}$, for $i = 1, \ldots, n$. For $i = 1, \ldots, n$, let r_i denote the radius of the ball for O_i from Definition 5.5, and let $r = \max_i r_i$. Then, $\mathsf{cl}(\mathcal{F}) \ge \mathsf{diam}(\mathsf{P})/(2n\alpha r)$.*

Proof: The claim is trivially true if $\mathsf{diam}(\mathsf{P}) = 0$, i.e. all the points p_i are the same. Let $l = \mathsf{cl}(\mathcal{F})$, for brevity. As we have $(f_i)_{\preceq l} = lO_i$, where the scaling is done around its center p_i, it follows that the sets lO_i for $i = 1, \ldots, n$ are connected. Since $lO_i \subseteq \mathsf{B}(\mathsf{p}_i, l\alpha r_i) \subseteq \mathsf{B}(\mathsf{p}_i, l\alpha r)$ it is easy to see that the

balls $B(p_i, l\alpha r)$ for $i = 1, \ldots, n$ are also connected. Let $u, v \in P$ be such that $\|u - v\| = \text{diam}(P)$. There is a sequence of distinct $i_1, \ldots, i_k \in \{1, \ldots, n\}$ such that $u = p_{i_1}, v = p_{i_k}$ and we have $B(p_{i_r}, l\alpha r) \cap B(p_{i_{r+1}}, l\alpha r) \neq \emptyset$ for $1 \leq r \leq k-1$. It follows that $\|p_{i_r} - p_{i_{r+1}}\| \leq 2l\alpha r, 1 \leq r \leq k-1$. By the triangle inequality,

$$\text{diam}(P) = \|u - v\| = \|p_{i_1} - p_{i_k}\| \leq \sum_{r=1}^{k-1} \|p_{i_r} - p_{i_{r+1}}\|$$

$$\leq \sum_{r=1}^{k-1} 2l\alpha r = 2(k-1)l\alpha r \leq 2nl\alpha r,$$

thus proving the claim. ∎

We can now show that condition $(P3)_{p8}$ holds for the F_{O_i}.

Lemma 5.11. *Consider the setting of Lemma 5.10. Given $\delta > 0$, there is a index set $\mathcal{I} \subseteq \{1, \ldots, n\}$ with $|\mathcal{I}| = O(\delta^{-d})$ and $y_0 = O(l \cdot n/\delta)$ such that the functions $\left\{ f_j \mid j \in \mathcal{I} \right\}$ form a (δ, y_0)-sketch, where $l = \text{cl}(\mathcal{F})$.*

Proof: We provide a sketch of the proof as details are easy but tedious. For each $1 \leq i, j \leq n$ we consider the set of objects $O_{ij} = O_i + p_j - p_i$, i.e. O_{ij} is O_i translated so that it is centered at p_j. By Lemma 5.8 there is an index set $\mathcal{I} \subseteq \{1, \ldots, n\}$ with $|\mathcal{I}| = O(\delta^{-d})$ such that for all y and any fixed j with $1 \leq j \leq n$ we have that $\bigcup_{i \in [n]} yO_{ij} \subseteq \bigcup_{i \in \mathcal{I}} (1 + \delta/4) yO_{ij}$. Let r_i denote the radius of the ball for O_i from Definition 5.5, and let $r = \max_i r_i$. By Lemma 5.10, we have that, $l \geq \text{diam}(P)/(2n\alpha r)$. Lemma 5.8 finds a $\mathcal{I}$ such that for all $i \in \mathcal{I}$, $r_i \geq \Omega(r)$. A translated copy $O_{ij} = O_i + p_j - p_i$ is a translation by a vector $u = p_j - p_i$. As $l \geq \text{diam}(P)/(2n\alpha r)$, there is a $y_0 = O(ln/\delta)$ such that $\|p_j - p_i\| \leq \delta\text{diam}(y_0 O_i)/4c$ for all $1 \leq i, j \leq n$, where $c = 2\alpha$. Thus using Lemma 5.9, $(1 + \delta/4)y_0 O_i + (p_j - p_i) \subseteq (1 + \delta/4)^2 y_0 O_i \subseteq (1 + \delta)y_0 O_i$. Clearly this also holds for any $y \geq y_0$. Thus for $y \geq y_0$ we have $(1 + \delta)yO_i$, covers $yO_i + (p_j - p_i)$ for $1 \leq i \leq n$. It is then easy to see that $\left\{ f_i \mid i \in \mathcal{I} \right\}$ is a (δ, y_0)-sketch. ∎

We conclude that for α-rounded fat objects, the scaling distance function they define falls under our framework. We thus get the following result.

Restatement of Theorem 3.4$_{p10}$. Consider a set $\mathcal{O}$ of α-rounded fat objects in $\mathbb{R}^d$, for some constant α. Then one can compute a $(1 + \varepsilon)$-AVD for the scaling distance functions induced by $\mathcal{O}$, with near linear space complexity, and logarithmic query time. See Theorem 3.1$_{p9}$ and Corollary 3.2 for the exact bounds.

Note, that the result in Theorem 3.4 covers any symmetric convex metric. Indeed, given a convex symmetric shape C centered at the origin, the distance it induces for any pair of points $p, u \in \mathbb{R}^d$, is the scaling distance of C centered p to u (or, by symmetry, the scaling distance of p from C centered at u). Under this distance $\mathbb{R}^d$ is a metric space, and of course, the triangle inequality holds. By an appropriate scaling of space, which does not affect the norm (except for scaling it) we can make C fat, and now Theorem 3.4 applies. Of course, Theorem 3.4 is considerably more general, allowing each of the points to induce a different scaling distance function, and the distance induced does not have to comply with the triangle inequality.

5.3. Nearest furthest-neighbor

For a set of points $\mathsf{S} \subseteq \mathbb{R}^d$ and a point q, the ***furthest-neighbor distance*** of q from Q, is $\mathcal{F}_\mathsf{S}(\mathsf{q}) = \max_{\mathsf{s} \in \mathsf{S}} \|\mathsf{q} - \mathsf{s}\|$; that is, it is the furthest one might have to travel from q to arrive to a point of S. For example, S might be the set of locations of facilities, where it is known that one of them is always open, and one is interested in the worst case distance a client has to travel to reach an open facility. The function $\mathcal{F}_\mathsf{S}(\cdot)$ is known as the ***furthest-neighbor Voronoi*** diagram, and while its worst case combinatorial complexity is similar to the regular Voronoi diagram, it can be approximated using a constant size representation (in low dimensions), see [Har99].

Given n sets of points $\mathsf{P}_1, \ldots, \mathsf{P}_n$ in $\mathbb{R}^d$, we are interested in the distance function $\mathcal{F}(\mathsf{q}) = \min_i \mathcal{F}_i(\mathsf{q})$, where $\mathcal{F}_i(\mathsf{q}) = \mathcal{F}_{\mathsf{P}_i}(\mathsf{q})$. This quantity arises natural when one tries to model uncertainty; indeed, let P_i be the set of possible locations of the ith point (i.e., the location of the ith point is chosen randomly, somehow, from the set P_i). Thus, $\mathcal{F}_i(\mathsf{q})$ is the worst case distance to the ith point, and $\mathcal{F}(\mathsf{q})$ is the worst-case nearest neighbor distance to the random point-set generated by picking the ith point from P_i, for $i = 1, \ldots, n$. We refer to $\mathcal{F}(\cdot)$ as the ***nearest furthest-neighbor*** distance, and we are interested in its approximation.

A naive solution to this problem would maintain a data structure for computing the furthest neighbor approximately for each of the P_i and then just compute the minimum of those distances. A data-structure to compute a $1 - \varepsilon$ approximation to the furthest neighbor takes $O(1/\varepsilon^d)$ space for $O(1/\varepsilon^d)$ query time, see [Har99] although this was probably known before. Thus the entire data structure would take up total space of $O(n/\varepsilon^d)$ with a query time of $O(n/\varepsilon^d)$. By using our general framework we can speed up the computation. We will show that $\mathcal{F}_i$, for $i = 1, \ldots, n$ satisfy the conditions (P1) – (P3) and (C1)–(C3). By Theorem 3.1 we can prepare a data-structure of size $O(n\mathrm{polylog}\,(n))$ that allows us a query time of $O(\log n)$ to find the desired nearest furthest-neighbor approximately. In order to facilitate the computations of distance$_f$s we also maintain data structures for $(1 - \varepsilon/4)$-approximate furthest neighbor search for each of the point sets P_i for $i = 1, 2, \ldots, n$ where ε is the approximation parameter for approximating the nearest furthest-neighbor, i.e. the approximation parameter for the problem we are trying to solve. Also, for $\mu = \varepsilon^2/144$ we also maintain μ-coresets for computing the minimum enclosing ball (**MEB**) approximately for each P_i for $i = 1, 2, \ldots, n$. Each such coreset has $O(1/\varepsilon^2)$ points, see [BHI02, BC03b, BC03a]. For each i with $1 \leq i \leq n$, the radius of the MEB of the coreset points is a $(1 + \mu)$-approximation to the radius of the MEB of P_i.

5.3.1. Satisfaction of the conditions

Observation 5.12. *We have that* $(\mathcal{F}_i)_{\preceq y} = \bigcap_{\mathsf{u} \in \mathsf{P}_i} \mathsf{B}(\mathsf{u}, y)$*, and* $\mathsf{diam}\left((\mathcal{F}_i)_{\preceq y}\right) \leq 2y$.

Given the above observation, it is easy to see that Condition (P1) is true, as $(\mathcal{F}_i)_{\preceq y}$ is a finite intersection of compact sets. The following Lemma shows that Condition (P2) is also true, by letting the growth function $\lambda_{(\mathcal{F}_i)}(y) = y$. Since $y \geq \mathsf{diam}\left((\mathcal{F}_i)_{\preceq y}\right)/2$ by Observation 5.12, it follows that we can choose the growth constant ζ to be 2.

Lemma 5.13. *For any i with $1 \leq i \leq n$, if $(\mathcal{F}_i)_{\preceq y} \neq \emptyset$, it is true that,*

$$(\mathcal{F}_i)_{\preceq y} \oplus \mathsf{B}(0, \varepsilon y) \subseteq (\mathcal{F}_i)_{\preceq(1+\varepsilon)y}.$$

Proof: Consider any point q in $(\mathcal{F}_i)_{\preceq y} \oplus \mathsf{B}(0, \varepsilon y)$. It is easy to see that $\mathsf{B}(\mathsf{q}, (1 + \varepsilon)y) \supseteq \mathsf{P}_i$ by the triangle inequality, and so $\mathsf{q} \in (\mathcal{F}_i)_{\preceq(1+\varepsilon)y}$. $\blacksquare$

Condition (P3) is implied by the following,

Lemma 5.14. *Let $\mathcal{G} \subseteq \{\mathcal{F}_1, \ldots, \mathcal{F}_n\}$ denote any set of functions. Then, given any $\delta > 0$, there is a subset $\mathcal{H} \subseteq \mathcal{G}$ with $|\mathcal{H}| = 1$ and a y_0 with $y_0 = O(\mathsf{cl}(\mathcal{G})\,|\mathcal{G}|\,/\delta)$ such that $\mathcal{H}$ is a (δ, y_0)-sketch for $\mathcal{G}$.*

Proof: Without loss of generality, let $\mathcal{G} = \{\mathcal{F}_1, \ldots, \mathcal{F}_m\}$ where $m = |\mathcal{G}|$. Let $z = \mathsf{cl}(\mathcal{G})$. Since $(\mathcal{F}_i)_{\preceq z}$ for $i = 1, \ldots, m$ are all connected, and by Observation 5.12, for each i with $1 \leq i \leq m$ we have that $\mathsf{diam}\left((\mathcal{F}_i)_{\preceq z}\right) \leq 2z$ it follows that for any two points $\mathsf{u} \in \mathsf{P}_j, \mathsf{v} \in \mathsf{P}_k$ with $1 \leq j, k \leq m$ there are points $\mathsf{u}' \in (\mathcal{F}_j)_{\preceq z}, \mathsf{v}' \in (\mathcal{F}_k)_{\preceq z}$, such that $\|\mathsf{u} - \mathsf{u}'\| \leq z$, $\|\mathsf{v} - \mathsf{v}'\| \leq z$ (by definition of the function $\mathcal{F}_j$ and $\mathcal{F}_k$ respectively) and $\|\mathsf{u}' - \mathsf{v}'\| \leq 2mz$, by the bound on the diameter of the sublevel sets and the condition of being connected, which is the same as the intersection graph of the sets being connected. It follows by the triangle inequality that $\|\mathsf{u} - \mathsf{v}\| \leq 2(m+1)z \leq 4mz$ i.e. $\mathsf{diam}(\mathsf{P}_1 \cup \cdots \cup \mathsf{P}_m) \leq 4mz$. Let $\mathcal{H}$ be a set containing an arbitrary function from $\mathcal{G}$ say, $\mathcal{H} = \{\mathcal{F}_1\}$. It is not too hard to see (or one can apply Lemma 5.2 with the case $\alpha_i = 0, w_i = 1$), that for every $i = 1, \ldots, m$, $\mathsf{B}(\mathsf{v}, y) \subseteq \mathsf{B}(\mathsf{u}, (1+\delta)y)$ for any points $\mathsf{u}, \mathsf{v} \in \mathsf{P}_1 \cup \cdots \cup \mathsf{P}_m$ if $y \geq y_0 = 4mz/\delta$, from the bound on the diameter of $\mathsf{P}_1 \cup \cdots \cup \mathsf{P}_m$. Thus, for any i with $1 \leq i \leq m$ we have that,

$$(\mathcal{F}_i)_{\preceq y} = \bigcap_{\mathsf{v} \in \mathsf{P}_i} \mathsf{B}(\mathsf{v}, y) \subseteq \mathsf{B}(\mathsf{u}, (1+\delta)y),$$

for all $\mathsf{u} \in \mathsf{P}_1$ and $y \geq y_0$. As such,

$$(\mathcal{F}_i)_{\preceq y} \subseteq \bigcap_{\mathsf{u} \in \mathsf{P}_1} \mathsf{B}(\mathsf{u}, (1+\delta)y) = (\mathcal{F}_1)_{\preceq (1+\delta)y},$$

and the result follows. $\blacksquare$

Remark 5.15. *For $\mathcal{G} = \{\mathcal{F}_1, \ldots, \mathcal{F}_m\}$, notice that we can compute the above set $\mathcal{H}$ in $O(1)$ time and that we can compute a polynomial approximation to $\mathsf{cl}(\mathcal{G})$ in $O(m)$ time, since we can compute the diameter $\mathsf{diam}(\mathsf{P}_1 \cup \cdots \cup \mathsf{P}_m)$ approximately in $O(m)$ time - we simply take an arbitrary point in P_1 and compute furthest distances approximately for each of P_i for $1 \leq i \leq m$ and take the maximum of these. We can use the $O(1)$ time query algorithm for furthest neighbor for this purpose.*

We now consider the computability conditions (C1)– (C3). To compute $\mathbb{d}(\mathsf{q}, \mathcal{F}_i)$ we use the data structure for approximate furthest neighbor queries to get a $(1 - \varepsilon/4)$-approximation to this number. We run the preprocessing algorithm, see Section 4, with approximation parameter $\varepsilon/4$. By Remark 2.10, we only tile the sublevel sets $(\mathcal{F}_i)_{\preceq y}$ with canonical cubes of size (rounded to a power of two) $\varepsilon\lambda_{(\mathcal{F}_i)}(y)/4 = \varepsilon y/4$. Notice that the minimum y such that $(\mathcal{F}_i)_{\preceq y}$ is non-empty is clearly the radius of the MEB of the point set P_i and for this value $(\mathcal{F}_i)_{\preceq y}$ just includes the center of the MEB. Let u_i and z_i denote the center and radius of the exact MEB, and u_i' and z_i' denote those computed by using the coreset. Since $\mathsf{z}_i' \leq (1 + \mu)\mathsf{z}_i$, it not too hard to see that $\|\mathsf{u}_i - \mathsf{u}_i'\| \leq 3\sqrt{\mu}\mathsf{z}_i = \varepsilon\mathsf{z}_i/4$. This is implied for example by Lemma A.1, presented in Appendix A, which may also be of independent interest (this assumes $\mu < 1$ which is indeed true). We are required to tile the sublevel set $(\mathcal{F}_i)_{\preceq y}$ for some $y \geq \mathsf{z}_i'/(1 + \mu)$ using cubes of size roughly $\varepsilon y/4$, but we use in fact cubes of size roughly $\varepsilon y/c$ for some large constant c. One can consider such cubes at increasing distance from the point u_i'. Choosing any point within a cube one evaluates approximately the furthest neighbor distance of P_i, and checks if it is at most $y(1 + O(\varepsilon))$. If so, one includes the cube. Since all such cubes will intersect the ball around u_i' of radius y, or a slight expansion of it, the number of such

cubes is still $O(1/\varepsilon^d)$. Now, the procedure in fact guarantees that all subcubes intersecting $(\mathcal{F}_i)_{\preceq y}$ are found, but in fact there may be some that do not intersect it. However, this is not a problem as such cells will still be inside $(\mathcal{F}_i)_{\preceq(1+\varepsilon/4)y}$ which is what is really required. To see that this works should be intuitively clear. We omit the straightforward, but tedious proof. This settles Condition (C2). Notice that the distance$_f$ between $\mathcal{F}_i$ and $\mathcal{F}_j$ is the radius of the minimum enclosing ball of the point set $\mathsf{P}_i \cup \mathsf{P}_j$. Using the $(1 + \mu)$-coresets that we have for the MEB of P_i and P_j we can compute a $(1 + 2\mu)$-coreset for the MEB of $\mathsf{P}_i \cup \mathsf{P}_j$ by simply merging those coresets. This allows us to approximately compute the distance$_f$.

Remark 5.16. *For the computability conditions (C1)–(C3) we only showed approximate results, that is the distance$_f$s were computed approximately. In fact, to be conservative, we used $\varepsilon/4$ as the approximation parameter in the construction algorithm and the furthest neighbor data structure. As a tedious but straightforward argument can show, the main lemmas Lemma 4.1 and Lemma 4.2 for near neighbor and interval range queries as well as the ones for computing the connectivity and splitting radius Lemma 4.9 and Lemma 4.12 can work under such approximate computations, with the same running times.*

We thus get the following result.

Restatement of Theorem 3.5$_{\mathbf{p10}}$. Given n point sets $\mathsf{P}_1, \ldots, \mathsf{P}_n$ in $\mathbb{R}^d$ with a total of m points, and a parameter $\varepsilon > 0$, one can preprocess the points into an AVD, of size $\widetilde{O}(n)$, for the nearest furthest-neighbor distance defined by these point sets. One can now answer $(1 + \varepsilon)$-approximate NN queries for this distance in $O(\log n)$ time. (Note, that the space and query time used, depend only on n, and not on the input size.)

Proof: We only need to show how get the improved space and query time. Observe that every one of the sets P_i can be replaced by a subset $\mathsf{S}_i \subseteq \mathsf{P}_i$, of size $O(1/\varepsilon^d \log(1/\varepsilon))$, such that for any point $\mathsf{q} \in \mathbb{R}^d$, we have that $\mathcal{F}_{\mathsf{S}_i}(\mathsf{q}) \leq \mathcal{F}_{\mathsf{P}_i}(\mathsf{q}) \leq (1 + \varepsilon/4)\mathcal{F}_{\mathsf{S}_i}(\mathsf{q})$. Such a subset can be computed in $O(|\mathsf{P}_i|)$ time, see [Har99][2]. We thus perform this transformation for each one of the uncertain point sets $\mathsf{P}_1, \ldots, \mathsf{P}_n$, which reduces the input size to $O(n/\varepsilon^d \log(1/\varepsilon))$. We now apply our main result to the distance functions induced by the reduced sets $\mathsf{S}_1, \ldots, \mathsf{S}_n$. ∎

6. Conclusions

In this paper, we investigated what classes of functions have minimization diagrams that can be approximated efficiently – where our emphasis was on distance functions. We defined a general framework and the requirements on the distance functions to fall under it. For this framework, we presented a new data-structure, with near linear space and preprocessing time. This data-structure can evaluate (approximately) the minimization diagram of a query point in logarithmic time. Surprisingly, one gets an AVD (approximate Voronoi diagram) of this complexity; that is, a decomposition of space with near linear complexity, such that for every region of this decomposition a single function serves as an ANN for all points in this region.

We also showed some interesting classes of functions for which we get this AVD. For example, additive and multiplicative weighted distance functions. No previous results of this kind were known, and even in the plane, multiplicative Voronoi diagrams have quadratic complexity in the worst case (for which the AVD generated has near linear complexity for any constant dimension). The

---

[2]One computes an appropriate exponential grid, of size $O(1/\varepsilon^d \log(1/\varepsilon))$, and pick from each grid cell one representative point from the points stored inside this cell.

framework also works for Minkowski metrics of fat convex bodies, and nearest furthest-neighbor. However, our main result applies to even more general distance functions.

Several questions remain open for further research:

(A) Are the additional polylog factors in the space necessary? In particular, it seems unlikely that using WSPD's directly, as done by Arya and Malamatos [AM02], should work in the most general settings, so reducing the logarithmic dependency seems quite interesting. Specifically, can the Arya and Malamatos construction [AM02] be somehow adapted to this framework, possibly with some additional constraints on the functions, to get a linear space construction?

(B) On the applications side, are constant degree polynomials a good family amenable to our framework? Specifically, consider a polynomial $\tau(x)$ that is positive for all $x \geq 0$. Given a point u, we associate the distance function $f(\mathsf{q}) = \tau(\|\mathsf{q} - \mathsf{u}\|)$ with u. Given a set of such distance functions, under which conditions, can one build an AVD for these functions efficiently? (It is not hard to see that in the general case this is not possible, at least under our framework.)

References

[AAH+13] P. K. Agarwal, B. Aronov, S. Har-Peled, J. M. Phillips, K. Yi, and W. Zhang. Nearest neighbor searching under uncertainty ii. In *Proc. 32nd ACM Sympos. Principles Database Syst.*, page to appear, 2013.

[AE98] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*. Amer. Math. Soc., 1998.

[AESZ12] P. K. Agarwal, A. Efrat, S. Sankararaman, and W. Zhang. Nearest-neighbor searching under uncertainty. In *Proc. 31st ACM Sympos. Principles Database Syst.*, pages 225–236, 2012.

[Agg09] C. C. Aggarwal. *Managing and Mining Uncertain Data*. Springer, 2009.

[AI08] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1):117–122, 2008.

[AM93] P. K. Agarwal and J. Matoušek. Ray shooting and parametric search. *SIAM J. Comput.*, 22:540–570, 1993.

[AM02] S. Arya and T. Malamatos. Linear-size approximate Voronoi diagrams. In *Proc. 13th ACM-SIAM Sympos. Discrete Algs.*, pages 147–155, 2002.

[AMM09] S. Arya, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate nearest neighbor searching. *J. Assoc. Comput. Mach.*, 57(1):1–54, 2009.

[AMN+98] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J. Assoc. Comput. Mach.*, 45(6):891–923, 1998.

[BC03a] M. Bădoiu and K. L. Clarkson. Optimal coresets for balls. http://cm.bell-labs.com/who/clarkson/coresets2.pdf, 2003.

[BC03b] M. Bădoiu and K. Clarkson. Smaller coresets for balls. In *Proc. 14th ACM-SIAM Sympos. Discrete Algs.*, pages 801–802, 2003.

[BHI02]   M. Bădoiu, S. Har-Peled, and P. Indyk. Approximate clustering via coresets. In *Proc. 34th Annu. ACM Sympos. Theory Comput.*, pages 250–257, 2002.

[Cha10]   T. M. Chan. Optimal partition trees. In *Proc. 26th Annu. ACM Sympos. Comput. Geom.*, pages 1–10, 2010.

[CK95]   P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to *k*-nearest-neighbors and *n*-body potential fields. *J. Assoc. Comput. Mach.*, 42:67–90, 1995.

[Cla88]   K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17:830–847, 1988.

[Cla06]   K. L. Clarkson. Nearest-neighbor searching and metric space dimensions. In G. Shakhnarovich, T. Darrell, and P. Indyk, editors, *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, pages 15–59. MIT Press, 2006.

[dBCvKO08] M. de Berg, O. Cheong, M. van Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 3rd edition, 2008.

[DRS09]   N. N. Dalvi, C. Ré, and D. Suciu. Probabilistic databases: Diamonds in the dirt. *Commun. ACM*, 52(7):86–94, 2009.

[Eri96]   J. Erickson. New lower bounds for Hopcroft's problem. *Discrete Comput. Geom.*, 16:389–418, 1996.

[Har99]   S. Har-Peled. Constructing approximate shortest path maps in three dimensions. *SIAM J. Comput.*, 28(4):1182–1197, 1999.

[Har01]   S. Har-Peled. A replacement for Voronoi diagrams of near linear size. In *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 94–103, 2001.

[Har11]   S. Har-Peled. *Geometric Approximation Algorithms*. Amer. Math. Soc., 2011.

[HIM12]   S. Har-Peled, P. Indyk, and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. *Theory Comput.*, 8:321–350, 2012. Special issue in honor of Rajeev Motwani.

[IM98]   P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 604–613, 1998.

[Mat92]   J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992.

[Mei93]   S. Meiser. Point location in arrangements of hyperplanes. *Inform. Comput.*, 106:286–303, 1993.

[MNP06]   R. Motwani, A. Naor, and R. Panigrahi. Lower bounds on locality sensitive hashing. In *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, pages 154–157, 2006.

[SA95]   M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.

A. Bounding the size of intersection of balls of the same radius

Lemma A.1. *Let* u, z *be the center and radius of the MEB for a set of points* $P = \{p_1, \ldots, p_m\} \subseteq \mathbb{R}^d$. *Let* $\delta \geq 0$ *be any number. Let* $p \in \bigcap_{i=1}^m B(p_i, (1+\delta)z)$. *Then,*

$$\delta z \leq \|p - u\| \leq \sqrt{4\delta + 2\delta^2} z.$$

Proof: The first inequality follows from the triangle inequality. We use the fact that there are affinely independent points from P, on the surface of the MEB, such that u lies in their convex hull. Thus, assume without loss of generality that there are points $p_1, \ldots, p_k \in P$ that are affinely independent with $\|p_i - u\| = z$ and $\lambda_i \geq 0$ for $i = 1, 2, \ldots, k$ such that,

$$u = \sum_{i=1}^k \lambda_i p_i, \qquad \sum_{i=1}^k \lambda_i = 1.$$

We restrict our attention only to the points $p_1, \ldots, p_k$ since the region $\bigcap_{i=1}^k B(p_i, (1+\delta)z)$ contains the region $\bigcap_{i=1}^m B(p_i, (1+\delta)z)$. Consider an arbitrary point $p \in \bigcap_{i=1}^k B(p_i, (1+\delta)z)$. Let p' be the projection of p to the affine subspace spanned by $p_1, \ldots, p_k$. We first bound $\|p' - u\|$. It is easy to see that $p' - u$ satisfies $\langle p' - u, p_i \rangle \leq 0$ for some i with $1 \leq i \leq k$. Without loss of generality assume $i = 1$. It follows that,

$$\|p' - p_1\| \geq \sqrt{\|p' - u\|^2 + \|u - p_1\|^2} \geq \sqrt{z^2 + \|p' - u\|^2}.$$

On the other hand it must be the case that, $\|p' - p_1\| \leq \|p - p_1\| \leq (1+\delta)z$. As such, $(1+\delta)z \geq \sqrt{z^2 + \|p' - u\|^2}$, and we have that $\|p' - u\| \leq \sqrt{2\delta + \delta^2} z$. We also have that,

$$(1+\delta)^2 z^2 \geq \|p - p_1\|^2 = \|p - p'\|^2 + \|p' - p_1\|^2 \geq \|p - p'\|^2 + z^2,$$

implying that $\|p - p'\| \leq \sqrt{2\delta + \delta^2} z$. It follows by the Pythagorean theorem,

$$\|p - u\|^2 = \|p - p'\|^2 + \|p' - u\|^2 \leq 2(2\delta + \delta^2)z^2,$$

and thus $\|p - u\| \leq \sqrt{4\delta + 2\delta^2} z$. ∎

B. Basic properties of the functions

Lemma B.1. *Let* $\mathcal{F}$ *be a set of functions that satisfy the compactness (P1) and bounded growth (P2) conditions. Then, for any* $f \in \mathcal{F}$, *either* $f_{\leq 0} = \emptyset$ *or* $f_{\leq 0}$ *consists of a single point.*

Proof: If $f_{\leq 0}$ contains at least two points, then by compactness (P1) of $f_{\leq 0}$ there are two points $x, y \in f_{\leq 0}$ such that $\|x - y\| = \mathsf{diam}(f_{\leq 0}) > 0$. By the bounded growth (P2) it follows that

$$f_{\leq 0} \subseteq f_{\leq 0} \oplus B\left(0, \frac{\|x - y\|}{\zeta}\right) \subseteq f_{\leq 0} \oplus B(0, \lambda_f(0)) \subseteq f_{\leq 0},$$

using $\varepsilon = 1$ and the fact that $\lambda_f(0) \geq \mathsf{diam}(f_{\leq 0})/\zeta = \|x - y\|/\zeta$. Thus, $f_{\leq 0} \oplus B\left(0, \frac{\|x-y\|}{\zeta}\right) = f_{\leq 0}$. Clearly in $y \oplus B\left(0, \frac{\|x-y\|}{\zeta}\right)$ there is some y' such that $\|x - y'\| > \|x - y\|$ which contradicts that x and y is a diametrical pair in $f_{\leq 0}$. ∎

By the above lemma, we may assume that a symbolic perturbation guarantees that $\mathrm{d}(f,g) > 0$ for $f \neq g$. With this convention we have the following,

Observation B.2. *If* $\mathsf{cl}(\mathcal{G}) = 0$ *for any non-empty subset* $\mathcal{G}$ *then* $|\mathcal{G}| = 1$.

We also assume that the quantities $\mathrm{d}(f,g)$ are distinct for all distinct pairs of functions.

Lemma B.3. *Let* $f \in \mathcal{G}$ *and* $y \geq 0$. *Suppose* $\mathsf{u}, \mathsf{v} \in f_{\preceq y}$. *Then,* $\mathsf{uv} \subseteq \mathcal{G}_{\preceq (1+\zeta/2)y}$, *where* uv *denotes the segment joining* u *to* v.

Proof: If $\mathsf{u} = \mathsf{v}$, the claim is obvious. Using bounded growth (P2) with $\varepsilon = \zeta/2$, and the inequality $\lambda_f(y) \geq \mathsf{diam}(f_{\preceq y})/\zeta$, it follows that $f_{\preceq y} \oplus \mathsf{B}(0, \mathsf{diam}(f_{\preceq y})/2) \subseteq f_{\preceq (1+\zeta/2)y}$. Thus, $\mathsf{u} \oplus \mathsf{B}(0, \mathsf{diam}(f_{\preceq y})/2) \subseteq f_{\preceq (1+\zeta/2)y}$ as well as $\mathsf{B}(\mathsf{v}, \mathsf{diam}(f_{\preceq y})/2) \subseteq f_{\preceq (1+\zeta/2)y}$. Since $\|\mathsf{u} - \mathsf{v}\| \leq \mathsf{diam}(f_{\preceq y})$ it follows that the entire segment uv is in $f_{\preceq (1+\zeta/2)y}$. $\blacksquare$

Lemma B.4. *Let* $A_1, \ldots, A_m \subseteq \mathbb{R}^d$ *be compact connected sets. Let* uv *be any segment. Suppose that* $\mathsf{uv} \cap A_i \neq \emptyset$ *for all* $1 \leq i \leq k$ *and* $\mathsf{uv} \subseteq \bigcup_{i=1}^k A_i$. *Then, the sets* $A_i, 1 \leq i \leq k$, *are connected.*

Proof: It is sufficient to prove the claim for $A_i \subseteq \mathsf{uv}$, as the truth of the claim for compact sets $A_i \cap \mathsf{uv}$ implies the truth for A_i. Thus, assume $A_i \subseteq \mathsf{uv}$. Suppose the claim is false. Consider the intersection graph of the $A_i, 1 \leq i \leq k$. This graph has at least two components by assumption. Let $B_1, \ldots, B_l$ be the partition of $[1, k]$ that define these components i.e. for each $1 \leq i \leq l$, the sets $A_j, j \in B_i$ are connected, and $A_x \cap A_y = \emptyset$ for $1 \leq x, y \leq k$ if x, y belong to different B_i. Denote by $C_i = \bigcup_{j \in B_i} A_j$ for $1 \leq i \leq l$. Clearly each C_i is compact. By an easy compactness argument, there are distinct $1 \leq i_1, i_2 \leq l$ such that for points $\mathsf{s} \in C_{i_1}, \mathsf{t} \in C_{i_2}$, we have that $0 < \|\mathsf{s} - \mathsf{t}\| = \min_{1 \leq x \neq y \leq l, \mathsf{p} \in C_x, \mathsf{q} \in C_y} \|\mathsf{p} - \mathsf{q}\|$. However, this is impossible as s, t are distinct points on uv and the segment st is therefore covered by the $C_i, 1 \leq i \leq l$. It follows that a smaller distance between distinct C_i must be attainable. $\blacksquare$

Lemma B.5. *Suppose we are given* $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{F}$, $\delta \geq 0$ *and* $y \geq 0$, *and* $\mathcal{H}$ *is a* (δ, y)-*sketch for* $\mathcal{G}$. *Then,* $\mathsf{cl}(\mathcal{H}) \leq (1+\delta)(1 + \zeta/2) \max(y, \mathsf{cl}(\mathcal{G}))$.

Proof: Assume that $\mathcal{G} = \{f_1, \ldots, f_m\}$ and $\mathcal{H} = \{f_1, \ldots, f_k\}$ where $k \leq m$. If $m = 1$ then $k = 1$ and we have by definition $\mathsf{cl}(\mathcal{G}) = \mathsf{cl}(\mathcal{H}) = 0$ and the result clearly holds true. If $m > 1$, we need to show that $(f_i)_{\preceq y'}$, for $i = 1, \ldots, k$, are connected, where $y' = (1+\delta)(1 + \zeta/2)l$ and $l = \max(y, \mathsf{cl}(\mathcal{G}))$. Now by definition, $\mathcal{G}_{\preceq l}$ is a connected set. Consider any $1 \leq i \neq j \leq k$. Then there is a sequence of distinct indices $i = i_1, i_2, \ldots, i_s = j$ such that $(f_{i_r})_{\preceq l} \cap (f_{i_{r+1}})_{\preceq l} \neq \emptyset$ for $1 \leq r \leq s - 1$. Consider any such index say i_r such that $i_r > k$ i.e. $f_{i_r} \notin \mathcal{H}$. Since, $(f_{i_r})_{\preceq l} \cap (f_{i_{r-1}})_{\preceq l} \neq \emptyset$ and $(f_{i_r})_{\preceq l} \cap (f_{i_{r+1}})_{\preceq l} \neq \emptyset$ we can choose points $\mathsf{u} \in (f_{i_{r-1}})_{\preceq l} \cap (f_{i_r})_{\preceq l}$ and $\mathsf{v} \in (f_{i_r})_{\preceq l} \cap (f_{i_{r+1}})_{\preceq l}$. Now the entire segment $\mathsf{uv} \subseteq (f_{i_r})_{\preceq (1+\zeta/2)l}$ by Lemma B.3. Since $(1 + \zeta/2)l \geq y$ it follows by the sketch property (P3), that $\mathsf{uv} \subseteq (f_{i_r})_{\preceq (1+\zeta/2)l} \subseteq \mathcal{H}_{\preceq (1+\zeta/2)(1+\delta)l}$. By Lemma B.4 the sets in the minimal cover of uv by the sublevel sets $(f_i)_{\preceq (1+\zeta/2)(1+\delta)l}, 1 \leq i \leq k$, are connected. It follows that $(f_{i_r})_{\preceq (1+\zeta/2)l}$ can be replaced by a sub-collection of the $(f_i)_{\preceq (1+\zeta/2)(1+\delta)l}, 1 \leq i \leq k$ and the property of neighbor intersections is still valid in the chain. We replace each occurrence of the set $(f_{i_r})_{\preceq (1+\zeta/2)l}$ for $i_r > k$ by the corresponding chain. It is easy to see that the resulting chain connects up $(f_{i_1})_{\preceq (1+\zeta/2)(1+\delta)l}$ and $(f_{i_s})_{\preceq (1+\zeta/2)(1+\delta)l}$. Now, duplicate elements can be easily removed without affecting the neighbor intersection property of the chain. $\blacksquare$

The following testifies that a sketch approximates the distance $_f$ of a set of functions.

Lemma B.6. *Let $\mathcal{H} \subseteq \mathcal{G}$ be sets of functions, where $\mathcal{H}$ is a (δ, y_0)-sketch for $\mathcal{G}$ for some $\delta \geq 0$ and $y_0 \geq 0$. Let q be a point such that $\mathrm{d}(\mathsf{q}, \mathcal{G}) \geq y_0$. Then we have that $\mathrm{d}(\mathsf{q}, \mathcal{H}) \leq (1 + \delta)\mathrm{d}(\mathsf{q}, \mathcal{G})$.*

Proof: Let $l = \mathrm{d}(\mathsf{q}, \mathcal{G})$ and let $f \in \mathcal{G}$ be a witness that $\mathsf{q} \in f_{\preceq l}$. As $l \geq y_0$ we have that $f_{\preceq l} \subseteq \bigcup_{g \in \mathcal{H}} g_{\preceq(1+\delta)l}$ by the sketch property (Definition 2.9). As such there is some function $g \in \mathcal{H}$ such that $\mathsf{q} \in g_{\preceq(1+\delta)l}$. It follows that $\mathrm{d}(\mathsf{q}, g) \leq (1 + \delta)\mathrm{d}(\mathsf{q}, \mathcal{G})$. ∎