

Geometric Computations on Indecisive and Uncertain Points

Allan Jørgensen* Maarten Löffler† Jeff M. Phillips‡

Abstract

We study computing geometric problems on *uncertain* points. An uncertain point is a point that does not have a fixed location, but rather is described by a probability distribution. When these probability distributions are restricted to a finite number of locations, the points are called *indecisive* points. In particular, we focus on geometric shape-fitting problems and on building compact distributions to describe how the solutions to these problems vary with respect to the uncertainty in the points. Our main results are: (1) a simple and efficient randomized approximation algorithm for calculating the distribution of any statistic on uncertain data sets; (2) a polynomial, deterministic and exact algorithm for computing the distribution of answers for any LP-type problem on an indecisive point set; and (3) the development of *shape inclusion probability* (SIP) functions which captures the ambient distribution of shapes fit to uncertain or indecisive point sets and are admissible to the two algorithmic constructions.

1 Introduction

In gathering data there is a trade-off between quantity and accuracy. The drop in the price of hard drives and other storage costs has shifted this balance towards gathering enormous quantities of data, yet with noticeable and sometimes intentionally tolerated increased error rates. However, often as a benefit from the large data sets, models are developed to describe the pattern of the data error.

Let us take as an example Light Detection and Ranging (LIDAR) data gathered for Geographic Information Systems (GIS) [40], specifically height values at millions of locations on a terrain. Each data point (x, y, z) has an x -value (longitude), a y -value (latitude), and a z -value (height). This data set is gathered by a small plane flying over a terrain with a laser aimed at the ground measuring the distance from the plane to the ground. Error can occur due to inaccurate estimation of the plane's altitude and position or artifacts on the ground distorting the laser's distance reading. But these errors are well-studied and can be modeled by replacing each data point with a probability distribution of its actual position. Greatly simplifying, we could represent each data point as a 3-variate normal distribution centered at its recorded value; in practice, more detailed uncertainty models are built.

Similarly, large data sets are gathered and maintained for many other applications. In robotic mapping [55, 22] error models are provided for data points gathered by laser range finders and other sources. In data mining [1, 6] original data (such as published medical data) are often perturbed by a known model to preserve anonymity. In spatial databases [28, 53, 15] large data sets may be summarized as probability distributions to store them more compactly. Data sets gathered by crawling the web have many false positives, and allow for models of these rates. Sensor networks [20] stream in large data sets collected by cheap and thus inaccurate sensors. In protein structure determination [51] every atom's position is imprecise due to inaccuracies in reconstruction techniques and the inherent flexibility in the protein. In summary, there are many large data sets with modeled errors and this uncertainty should be dealt with explicitly.

*formerly: MADALGO, Department of Computer Science, University of Aarhus, Denmark. jallan@madalgo.au.dk

†formerly: Computer Science Department, University of California, Irvine, USA. mloffler@uci.edu

‡School of Computing, University of Utah, USA. jeffp@cs.utah.edu



Fig. 1: (a) An example input consisting of $n = 3$ sets of $k = 6$ points each. (b) One of the 6^3 possible samples of $n = 3$ points.

1.1 The Input: Geometric Error Models

The input for a typical computational geometry problem is a set P of n points in \mathbb{R}^2 , or more generally \mathbb{R}^d . In this paper we consider extensions of this model where each point is also given a model of its uncertainty. This model describes for each point a distribution or bounds on the point's location, if it exists at all.

- Most generally, we describe these data as *uncertain points* $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$. Here each point's location is described by a probability distribution μ_i (for instance by a Gaussian distribution). This general model can be seen to encompass the forthcoming models, but is often not worked with directly because of the computational difficulties arisen from its generality. For instance in tracking uncertain objects a particle filter uses a discrete set of locations to model uncertainty [46] while a Kalman filter restricts the uncertainty model to a Gaussian distribution [33].
- A more restrictive model we also study in this paper are *indecisive points* where each point can take one of a finite number of locations. To simplify the model (purely for making results easier to state) we let each point have exactly k possible locations, forming the domain of a probability distribution. That is each uncertain point P_i is at one of $\{p_{i,1}, p_{i,2}, \dots, p_{i,k}\}$. Unless further specified, each location is equally likely with probability $1/k$, but we can also assign each location a weight $w_{i,j}$ as the probability that P_i is at $p_{i,j}$ where $\sum_{j=1}^k w_{i,j} = 1$ for all i .

Indecisive points appear naturally in many applications. They play an important role in databases [19, 7, 18, 16, 54, 2, 17], machine learning [10], and sensor networks [59] where a limited number of probes from a certain data set are gathered, each potentially representing the true location of a data point. Alternatively, data points may be obtained using imprecise measurements or are the result of inexact earlier computations. However, the results with detailed algorithmic analysis generally focus on one-dimensional data; furthermore, they often only return the expected value or the most likely answer instead of calculating a full distribution.

- An *imprecise point* is one where its location is not known precisely, but it is restricted to a range. In one-dimension these ranges are modeled as uncertainty intervals, but in 2 or higher dimensions they become geometric regions. An early model to quantify imprecision in geometric data, motivated by finite precision of coordinates, is ε -*geometry*, introduced by Guibas et. al. [26], where each point was only known to be somewhere within an ε -radius ball of its guessed location. The simplicity of this model has provided many uses in geometry. Guibas et. al. [27] define *strongly convex* polygons: polygons that are guaranteed to stay convex, even when the vertices are perturbed by ε . Bandyopadhyay and Snoeyink [8] compute the set of all potential simplices in \mathbb{R}^2 and \mathbb{R}^3 that could belong to the Delaunay triangulation. Held and Mitchell [31] and Löffler and Snoeyink [41] study the problem of preprocessing a set of imprecise points under this model, so that when the true points are specified later some computation can be done faster.

A more involved model for imprecision can be obtained by not specifying a single ε for all the points, but allowing a different radius for each point, or even other shapes of imprecision regions. This allows for

modeling imprecision that comes from different sources, independent imprecision in different dimensions of the input, etc. This extra freedom in modeling comes at the price of more involved algorithmic solutions, but still many results are available. Nagai and Tokura [47] compute the union and intersection of all possible convex hulls to obtain bounds on any possible solution, as does Ostrovsky-Berman and Joskowicz [48] in a setting allowing some dependence between points. Van Kreveld and Löffler [37] study the problem of computing the smallest and largest possible values of several geometric extent measures, such as the diameter or the radius of the smallest enclosing ball, where the points are restricted to lie in given regions in the plane. Kruger [38] extends some of these results to higher dimensions.

Although imprecise points do not traditionally have an associated probability distribution associated to them, we argue that they can still be considered a special case of our uncertain points, since we can impose e.g. a uniform distribution on the regions, and then ask question about the smallest or largest non-zero probability values of some function, which would correspond to bounds in the classical model.

- A *stochastic point* p has a fixed location, but which only exists with a probability ρ . These points arise naturally in many database scenarios [7, 17] where gathered data has many false positives. Recently in geometry Kamousi, Chan, and Suri [35, 34] considered geometric problems on stochastic points and geometric graphs with stochastic edges. These stochastic data sets can be interpreted as uncertain point sets as well by allowing the probability distribution governing uncertain points to have a certain probability of not existing, or rather the integral of the distribution is ρ instead of always 1.

1.2 The Output: Distributional Representations

This paper studies how to compute distributions of statistics over uncertain data. These distributions can take several forms. In the simplest case, a distribution of a single value has a one-dimensional domain. The technical definition yields a simpler exposition when the distribution is represented as a cumulative density function, which we refer to as a *quantization*. This notion can be extended to a multi-dimensional cumulative density function (a *k-variate quantization*) as we measure multiple variables simultaneously. Finally, we also describe distributions over shapes defined on uncertain points (e.g. minimum enclosing ball). As the domains of these shape distributions are a bit abstract and difficult to work with, we convey this information as a *shape inclusion probability* or *SIP*; for any point in the domain of the input point sets we describe the probability the point is contained in the shape.

This model of uncertain data has been studied in the database community but for different types of problems on usually one-dimensional data, such as indexing [2, 54, 32], ranking [17], nearest neighbors [14] and creating histograms [16].

1.3 Contributions

For each type of distributional representation of output we study, the goal is a function from some domain to a range of $[0, 1]$. For the general case of uncertain points, we provide simple and efficient, randomized approximation algorithms that results in a function that *everywhere* has error at most ε . Each variation of the algorithm runs in $O((1/\varepsilon^2)(\nu + \log(1/\delta))T)$ time and produces an output of size $O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ where ν describes the complexity of the output shape (i.e. VC-dimension), δ is the probability of failure, and T is the time it takes to compute the geometric question on certain points. These results are quite practical as experimental results demonstrate that the constant for the big-Oh notation is approximately 0.5. Furthermore, for one-dimensional output distributions (quantizations) the size can be reduced to $1/\varepsilon$, and for k -dimensional distributions to $O((k/\varepsilon) \log^4(1/\varepsilon))$. We also extend these approaches to allow for geometric approximations based on α -kernels [4, 3].

For the case of indecisive points, we provide deterministic and exact, polynomial-time algorithms for all LP-type problems with constant combinatorial dimension (e.g. minimum enclosing ball). We also provide evidence that for problems outside this domain, deterministic exact algorithms are not available, in particular showing that diameter is $\#P$ -hard despite having an output distribution of polynomial size. Finally, we consider deterministic algorithms for uncertain point sets with continuous distributions describing the location

of each point. We describe a non-trivial range space on these input distributions from which an ε -sample creates a set of indecisive points, from which this algorithm can be performed to deterministically create an approximation to the output distribution.

2 Preliminaries

This section provides formal definitions for existing approximation schemes related to our work as well as the for our output distributions.

2.1 Approximation Schemes: ε -Samples and α -Kernels

This work allows for three types of approximations. The most natural in this setting is controlled by a parameter ε which denotes the error tolerance for probability. That is an ε -approximation for any function with range in $[0, 1]$ measuring probability can return a value off by at most an additive ε . The second type of error is a parameter δ which denotes the chance of failure of a randomized algorithm. That is a δ -approximate randomized algorithm will be correct with probability at least $1 - \delta$. Finally, in specific contexts we allow a geometric error parameter α . In our context, an α -approximate geometric algorithm can allow a relative α -error in the width of any object. This is explained more formally below. It should be noted that these three types of error cannot be combined into a single term, and each needs to be considered separately. However, the ε and δ parameters have a well-defined trade-off.

ε -Samples. For a set P let \mathcal{A} be a set of subsets of P . In our context usually P will be a point set and the subsets in \mathcal{A} could be induced by containment in a shape from some family of geometric shapes. For example, \mathcal{J}_+ describes one-sided intervals of the form $(-\infty, t)$. The pair (P, \mathcal{A}) is called a *range space*. We say that $Q \subset P$ is an ε -sample of (P, \mathcal{A}) if

$$\forall R \in \mathcal{A} \left| \frac{\phi(R \cap Q)}{\phi(Q)} - \frac{\phi(R \cap P)}{\phi(P)} \right| \leq \varepsilon,$$

where $|\cdot|$ takes the absolute value and $\phi(\cdot)$ returns the measure of a point set. In the discrete case $\phi(Q)$ returns the cardinality of Q . We say \mathcal{A} *shatters* a set S if every subset of S is equal to $R \cap S$ for some $R \in \mathcal{A}$. The cardinality of the largest discrete set $S \subseteq P$ that \mathcal{A} can shatter is the *VC-dimension* of (P, \mathcal{A}) .

When (P, \mathcal{A}) has constant VC-dimension ν , we can create an ε -sample Q of (P, \mathcal{A}) , with probability $1 - \delta$, by uniformly sampling $O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ points from P [57, 39]. There exist deterministic techniques to create ε -samples [43, 13] of size $O(\nu(1/\varepsilon^2) \log(1/\varepsilon))$ in time $O(\nu^{3\nu} n ((1/\varepsilon^2) \log(\nu/\varepsilon))^\nu)$. A recent result of Bansal [9] (also see this simplification [42]) can slightly improve this bound to $O(1/\varepsilon^{2-1/2\nu})$, following an older existence proof [45], in time polynomial in n and $1/\varepsilon$. When P is a point set in \mathbb{R}^d and the family of ranges \mathcal{R}_d is determined by inclusion in axis-aligned boxes, then an ε -sample for (P, \mathcal{R}_d) of size $O((d/\varepsilon) \log^{2d}(1/\varepsilon))$ can be constructed in $O((n/\varepsilon^3) \log^{6d}(1/\varepsilon))$ time [49].

For a range space (P, \mathcal{A}) the *dual range space* is defined (\mathcal{A}, P^*) where P^* is all subsets $\mathcal{A}_p \subseteq \mathcal{A}$ defined for an element $p \in P$ such that $\mathcal{A}_p = \{A \in \mathcal{A} \mid p \in A\}$. If (P, \mathcal{A}) has VC-dimension ν , then (\mathcal{A}, P^*) has VC-dimension $\leq 2^{\nu+1}$. Thus, if the VC-dimension of (\mathcal{A}, P^*) is constant, then the VC-dimension of (P, \mathcal{A}) is also constant [44].

When we have a distribution $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$, such that $\int_{x \in \mathbb{R}^d} \mu(x) dx = 1$, we can think of this as the set P of all points in \mathbb{R}^d , where the weight w of a point $p \in \mathbb{R}^d$ is $\mu(p)$. To simplify notation, we write (μ, \mathcal{A}) as a range space where the ground set is this set $P = \mathbb{R}^d$ weighted by the distribution μ .

α -Kernels. Given a point set $P \in \mathbb{R}^d$ of size n and a direction $u \in \mathbb{S}^{d-1}$, let $P[u] = \arg \max_{p \in P} \langle p, u \rangle$, where $\langle \cdot, \cdot \rangle$ is the inner product operator. Let $\omega(P, u) = \langle P[u] - P[-u], u \rangle$ describe the width of P in direction u . We say that $K \subseteq P$ is an α -kernel of P if for all $u \in \mathbb{S}^{d-1}$

$$\omega(P, u) - \omega(K, u) \leq \alpha \cdot \omega(P, u).$$




Fig. 2: (a) The true form of a monotonically increasing function from $\mathbb{R} \rightarrow \mathbb{R}$. (b) The ε -quantization R as a point set in \mathbb{R} . (c) The inferred curve h_R in \mathbb{R}^2 .

α -kernels of size $O(1/\alpha^{(d-1)/2})$ [4] can be calculated in time $O(n + 1/\alpha^{d-3/2})$ [12, 58]. Computing many extent related problems such as diameter and smallest enclosing ball on K approximates the problem on P [4, 3, 12].

2.2 Problem Statement

Let $\mu_i : \mathbb{R}^d \rightarrow \mathbb{R}^+$ describe the probability distribution of an uncertain point P_i where the integral $\int_{q \in \mathbb{R}^d} \mu_i(q) dq = 1$. We say that a set Q of n points is a *support* from \mathcal{P} if it contains exactly one point from each set P_i , that is, if $Q = \{q_1, q_2, \dots, q_n\}$ with $q_i \in P_i$. In this case we also write $Q \in \mathcal{P}$. Let $\mu_{\mathcal{P}} : \mathbb{R}^d \times \mathbb{R}^d \times \dots \times \mathbb{R}^d \rightarrow \mathbb{R}^+$ describe the distribution of supports $Q = \{q_1, q_2, \dots, q_n\}$ under the joint probability over each $q_i \in P_i$. For brevity we write the space $\mathbb{R}^d \times \dots \times \mathbb{R}^d$ as \mathbb{R}^{dn} . For this paper we will assume $\mu_{\mathcal{P}}(q_1, q_2, \dots, q_n) = \prod_{i=1}^n \mu_{P_i}(q_i)$, so the distribution for each point is independent, although this restriction can be easily removed for all randomized algorithms.

Quantizations and their approximations. Let $f : \mathbb{R}^{dn} \rightarrow \mathbb{R}^k$ be a function on a fixed point set. Examples include the radius of the minimum enclosing ball where $k = 1$ and the width of the minimum enclosing axis-aligned rectangle along the x -axis and y -axis where $k = 2$. Define the “dominates” binary operator \preceq so that $(p_1, \dots, p_k) \preceq (v_1, \dots, v_k)$ is true if for every coordinate $p_i \leq v_i$. Let $\mathbb{X}_f(v) = \{Q \in \mathbb{R}^{dn} \mid f(Q) \preceq v\}$. For a query value v define, $F_{\mu_{\mathcal{P}}}(v) = \int_{Q \in \mathbb{X}_f(v)} \mu_{\mathcal{P}}(Q) dQ$. Then $F_{\mu_{\mathcal{P}}}$ is the cumulative density function of the distribution of possible values that f can take[□]. We call $F_{\mu_{\mathcal{P}}}$ a *quantization* of f over $\mu_{\mathcal{P}}$.

Ideally, we would return the function $F_{\mu_{\mathcal{P}}}$ so we could quickly answer any query exactly, however, for the most general case we consider, it is not clear how to calculate $F_{\mu_{\mathcal{P}}}(v)$ exactly for even a single query value v . Rather, we introduce a data structure, which we call an ε -quantization, to answer any such query approximately and efficiently, illustrated in Figure 2 for $k = 1$. An ε -quantization is a point set $R \subset \mathbb{R}^k$ which induces a function h_R where $h_R(v)$ describes the fraction of points in R that v dominates. Let $R_v = \{r \in R \mid r \preceq v\}$. Then $h_R(v) = |R_v|/|R|$. For an isotonic (monotonically increasing in each coordinate) function $F_{\mu_{\mathcal{P}}}$ and any value v , an ε -quantization, R , guarantees that $|h_R(v) - F_{\mu_{\mathcal{P}}}(v)| \leq \varepsilon$. More generally (and, for brevity, usually only when $k > 1$), we say R is a k -variate ε -quantization. An example of a 2-variate ε -quantization is shown in Figure 3. The space required to store the data structure for R is dependent only on ε and k , not on $|P|$ or $\mu_{\mathcal{P}}$.

$(\varepsilon, \delta, \alpha)$ -Kernels. Rather than compute a new data structure for each measure we are interested in, we can also compute a single data structure (a coreset) that allows us to answer many types of questions. For an isotonic function $F_{\mu_{\mathcal{P}}} : \mathbb{R}^+ \rightarrow [0, 1]$, an (ε, α) -quantization data structure M describes a function $h_M : \mathbb{R}^+ \rightarrow [0, 1]$ so for any $x \in \mathbb{R}^+$, there is an $x' \in \mathbb{R}^+$ such that (1) $|x - x'| \leq \alpha x$ and (2) $|h_M(x) - F_{\mu_{\mathcal{P}}}(x')| \leq \varepsilon$. An $(\varepsilon, \delta, \alpha)$ -kernel is a data structure that can produce an (ε, α) -quantization, with probability at least $1 - \delta$, for $F_{\mu_{\mathcal{P}}}$ where f measures the width in any direction and whose size depends only on ε , α , and δ . The notion of (ε, α) -quantizations is generalizes to a k -variate version, as do $(\varepsilon, \delta, \alpha)$ -kernels.

[□] For a function f and a distribution of point sets $\mu_{\mathcal{P}}$, we will always represent the cumulative density function of f over $\mu_{\mathcal{P}}$ by $F_{\mu_{\mathcal{P}}}$.




Fig. 3: (a) The true form of a xy -monotone 2-variate function. (b) The ε -quantization R as a point set in \mathbb{R}^2 . (c) The inferred surface h_R in \mathbb{R}^3 . (d) Overlay of the two images.

Shape inclusion probabilities. A *summarizing shape* of a point set $P \subset \mathbb{R}^d$ is a Lebesgue-measurable subset of \mathbb{R}^d that is determined by P . I.e. given a class of shapes \mathcal{S} , the summarizing shape $S(P) \in \mathcal{S}$ is the shape that optimizes some aspect with respect to P . Examples include the smallest enclosing ball and the minimum-volume axis-aligned bounding box. For a family \mathcal{S} we can study the *shape inclusion probability function* $s_{\mu_P} : \mathbb{R}^d \rightarrow [0, 1]$ (or *sip* function), where $s_{\mu_P}(q)$ describes the probability that a query point $q \in \mathbb{R}^d$ is included in the summarizing shape². For the more general types of uncertain points, there does not seem to be a closed form for many of these functions. In these cases we can calculate an ε -*sip* function $\hat{s} : \mathbb{R}^d \rightarrow [0, 1]$ such that $\forall_{q \in \mathbb{R}^d} |s_{\mu_P}(q) - \hat{s}(q)| \leq \varepsilon$. The space required to store an ε -*sip* function depends only on ε and the complexity of the summarizing shape.

3 Randomized Algorithm for ε -Quantizations

We develop several algorithms with the following basic structure (as outlined in Algorithm 3.1): (1) sample one point from each distribution to get a random point set; (2) construct the summarizing shape of the random point set; (3) repeat the first two steps $O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ times and calculate a summary data structure. This algorithm only assumes that we can draw a random point from μ_p for each $p \in P$ in constant time; if the time depends on some other parameters, the time complexity of the algorithms can be easily adjusted.

Algorithm 3.1 Approximate μ_P w.r.t. a family of shapes \mathcal{S} or function f_S

- 1: **for** $i = 1$ **to** $m = O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ **do**
 - 2: **for all** $p_j \in P$ **do**
 - 3: Sample q_j from μ_{p_j} .
 - 4: Set $V_i = f_S(\{q_1, q_2, \dots, q_n\})$.
 - 5: Reduce or Simplify the set $\mathcal{V} = \{V_i\}_{i=1}^m$.
-

Algorithm for ε -quantizations. For a function f on a point set P of size n , it takes $T_f(n)$ time to evaluate $f(P)$. We construct an approximation to F_{μ_P} as follows. First draw a sample point q_j from each μ_{p_j} for $p_j \in P$, then evaluate $V_i = f(\{q_1, \dots, q_n\})$. The fraction of trials of this process that produces a value dominated by v is the estimate of $F_{\mu_P}(v)$. In the univariate case we can reduce the size of \mathcal{V} by returning $2/\varepsilon$ evenly spaced points according to the sorted order.

Theorem 3.1. For a distribution μ_P of n points, with success probability at least $1 - \delta$, there exists an ε -quantization of size $O(1/\varepsilon)$ for F_{μ_P} , and it can be constructed in $O(T_f(n)(1/\varepsilon^2) \log(1/\delta))$ time.

² For technical reasons, if there are (degenerately) multiple optimal summarizing shapes, we say each is equally likely to be the summarizing shape of the point set.

Proof. Because $F_{\mu_P} : \mathbb{R} \rightarrow [0, 1]$ is an isotonic function, there exists another function $g : \mathbb{R} \rightarrow \mathbb{R}^+$ such that $F_{\mu_P}(t) = \int_{x=-\infty}^t g(x) dx$ where $\int_{x \in \mathbb{R}} g(x) dx = 1$. Thus g is a probability distribution of the values of f given inputs drawn from μ_P . This implies that an ε -sample of (g, \mathcal{J}_+) is an ε -quantization of F_{μ_P} , since both estimate within ε the fraction of points in any range of the form $(-\infty, x)$. This last fact can also be seen through a result by Dvoretzky, Kiefer, and Wolfowitz [21].

By drawing a random sample q_i from each μ_{p_i} for $p_i \in P$, we are drawing a random point set Q from μ_P . Thus $f(Q)$ is a random sample from g . Hence, using the standard randomized construction for ε -samples, $O((1/\varepsilon^2) \log(1/\delta))$ such samples will generate an $(\varepsilon/2)$ -sample for g , and hence an $(\varepsilon/2)$ -quantization for F_{μ_P} , with probability at least $1 - \delta$.

Since in an $(\varepsilon/2)$ -quantization R every value $h_R(v)$ is different from $F_{\mu_P}(v)$ by at most $\varepsilon/2$, then we can take an $(\varepsilon/2)$ -quantization of the function described by $h_R(\cdot)$ and still have an ε -quantization of F_{μ_P} . Thus, we can reduce this to an ε -quantization of size $O(1/\varepsilon)$ by taking a subset of $2/\varepsilon$ points spaced evenly according to their sorted order. \square

Multivariate ε -quantizations. We can construct k -variate ε -quantizations similarly using the same basic procedure as in Algorithm 3.1. The output V_i of f is now k -variate and thus results in a k -dimensional point. As a result, the reduction of the final size of the point set requires more advanced procedures.

Theorem 3.2. *Given a distribution μ_P of n points, with success probability at least $1 - \delta$, we can construct a k -variate ε -quantization for F_{μ_P} of size $O((k/\varepsilon^2)(k + \log(1/\delta)))$ and in time $O(T_f(n)(1/\varepsilon^2)(k + \log(1/\delta)))$.*

Proof. Let \mathcal{R}_+ describe the family of ranges where a range $A_p = \{q \in \mathbb{R}^k \mid q \preceq p\}$. In the k -variate case there exists a function $g : \mathbb{R}^k \rightarrow \mathbb{R}^+$ such that $F_{\mu_P}(v) = \int_{x \preceq v} g(x) dx$ where $\int_{x \in \mathbb{R}^k} g(x) dx = 1$. Thus g describes the probability distribution of the values of f , given inputs drawn randomly from μ_P . Hence a random point set Q from μ_P , evaluated as $f(Q)$, is still a random sample from the k -variate distribution described by g . Thus, with probability at least $1 - \delta$, a set of $O((1/\varepsilon^2)(k + \log(1/\delta)))$ such samples is an ε -sample of (g, \mathcal{R}_+) , which has VC-dimension k , and the samples are also a k -variate ε -quantization of F_{μ_P} . Again, this specific VC-dimension sampling result can also be achieved through a result of Kiefer and Wolfowitz [36]. \square

We can then reduce the size of the ε -quantization R to $O((k^2/\varepsilon) \log^{2k}(1/\varepsilon))$ in $O(|R|(k/\varepsilon^3) \log^{6k}(1/\varepsilon))$ time [49] or to $O((k^2/\varepsilon^2) \log(1/\varepsilon))$ in $O(|R|(k^{3k}/\varepsilon^{2k}) \cdot \log^k(k/\varepsilon))$ time [13], since the VC-dimension is k and each data point requires $O(k)$ storage.

Also on k -variate statistics, we can query the resulting k -dimensional distribution using other shapes with bounded VC-dimension ν , and if the sample size is $m = O((1/\varepsilon^2)(\nu + \log(1/\delta)))$, then all queries have at most ε -error with probability at least $1 - \delta$. In contrast to the two above results, this statement seems to require the VC-dimension view, as opposed to appealing to the Kiefer-Wolfowitz line of work [21, 36].

3.1 $(\varepsilon, \delta, \alpha)$ -Kernels

The above construction works for a fixed family of summarizing shapes. In this section, we show how to build a single data structure, an $(\varepsilon, \delta, \alpha)$ -kernel, for a distribution μ_P in \mathbb{R}^{dn} that can be used to construct (ε, α) -quantizations for several families of summarizing shapes. This added generality does come at an increased cost in construction. In particular, an $(\varepsilon, \delta, \alpha)$ -kernel of μ_P is a data structure such that in any query direction $u \in \mathbb{S}^{d-1}$, with probability at least $1 - \delta$, we can create an (ε, α) -quantization for the cumulative density function of $\omega(\cdot, u)$, the width in direction u .

We follow the randomized framework described above as follows. The desired $(\varepsilon, \delta, \alpha)$ -kernel \mathcal{K} consists of a set of $m = O((1/\varepsilon^2) \log(1/\delta))$ $(\alpha/2)$ -kernels, $\{K_1, K_2, \dots, K_m\}$, where each K_j is an $(\alpha/2)$ -kernel of a point set Q_j drawn randomly from μ_P . Given \mathcal{K} , with probability at least $1 - \delta$, we can create an (ε, α) -quantization for the cumulative density function of width over μ_P in any direction $u \in \mathbb{S}^{d-1}$. Specifically, let $M = \{\omega(K_j, u)\}_{j=1}^m$.

Lemma 3.3. *With probability at least $1 - \delta$, M is an (ε, α) -quantization for the cumulative density function of the width of μ_P in direction u .*

Proof. The width $\omega(Q_j, u)$ of a random point set Q_j drawn from μ_P is a random sample from the distribution over widths of μ_P in direction u . Thus, with probability at least $1 - \delta$, m such random samples would create an ε -quantization. Using the width of the α -kernels K_j instead of Q_j induces an error on each random sample of at most $2\alpha \cdot \omega(Q_j, u)$. Then for a query width w , say there are γm point sets Q_j that have width at most w and $\gamma' m$ α -kernels K_j with width at most w ; see Figure 4. Note that $\gamma' > \gamma$. Let $\hat{w} = w - 2\alpha w$. For each point set Q_j that has width greater than w but the corresponding α -kernel K_j has width at most w , it follows that K_j has width greater than \hat{w} . Thus the number of α -kernels K_j that have width at most \hat{w} is at most γm , and thus there is a width w' between w and \hat{w} such that the number of α -kernels at most w' is exactly γm . \square




Fig. 4: (ε, α) -quantization M (white circles) and ε -quantization R (black circles) given a query width w .

Since each K_j can be computed in $O(n + 1/\alpha^{d-3/2})$ time, we obtain:

Theorem 3.4. *We can construct an $(\varepsilon, \delta, \alpha)$ -kernel for μ_P on n points in \mathbb{R}^d of size $O((1/\alpha^{(d-1)/2})(1/\varepsilon^2) \cdot \log(1/\delta))$ in $O((n + 1/\alpha^{d-3/2}) \cdot (1/\varepsilon^2) \log(1/\delta))$ time.*

The notion of (ε, α) -quantizations and $(\varepsilon, \delta, \alpha)$ -kernels can be extended to k -dimensional queries or for a series of up to k queries which all have approximation guarantees with probability $1 - \delta$.

Other coresets. In a similar fashion, coresets of a point set distribution μ_P can be formed using coresets for other problems on discrete point sets. For instance, sample $m = O((1/\varepsilon^2) \log(1/\delta))$ points sets $\{P_1, \dots, P_m\}$ each from μ_P and then store α -samples $\{Q_1 \subseteq P_1, \dots, Q_m \subseteq P_m\}$ of each. When we use random sampling in the second set, then not all distributions μ_{p_i} need to be sampled for each P_j in the first round. This results in an $(\varepsilon, \delta, \alpha)$ -sample of μ_P , and can, for example, be used to construct (with probability $1 - \delta$) an (ε, α) -quantization for the fraction of points expected to fall in a query disk. Similar constructions can be done for other coresets, such as ε -nets [30], k -center [5], or smallest enclosing ball [11].

3.2 Measuring the Error

We have established asymptotic bounds of $m = O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ random samples for constructing ε -quantizations. Now we empirically demonstrate that the constant hidden by the big-O notation is approximately 0.5, indicating that these algorithms are indeed quite practical.

As a data set, we consider a set of $n = 50$ sample points in \mathbb{R}^3 chosen randomly from the boundary of a cylinder piece of length 10 and radius 1. We let each point represent the center of 3-variate Gaussian distribution with standard deviation 2 to represent the probability distribution of an uncertain point. This set of distributions describes an uncertain point set $\mu_P : \mathbb{R}^{3n} \rightarrow \mathbb{R}^+$.

We want to estimate three statistics on μ_P : **dwid**, the width of the points set in a direction that makes an angle of 75° with the cylinder axis; **diam**, the diameter of the point set; and **seb₂**, the radius of the smallest enclosing ball (using code from Bernd Gärtner [25]). We can create ε -quantizations with m samples from μ_P , where the value of m is from the set $\{16, 64, 256, 1024, 4096\}$.

We would like to evaluate the ε -quantizations versus the ground truth function F_{μ_P} ; however, it is not clear how to evaluate F_{μ_P} . Instead, we create another ε -quantization Q with $\eta = 100000$ samples from μ_P , and treat this as if it were the ground truth. To evaluate each sample ε -quantization R versus Q we find

the maximum deviation (i.e. $d_\infty(R, Q) = \max_{q \in \mathbb{R}} |h_R(q) - h_Q(q)|$) with h defined with respect to `diam` or `dwid`. This can be done by for each value $r \in R$ evaluating $|h_R(r) - h_Q(r)|$ and $|(h_R(r) - 1/|R|) - h_Q(r)|$ and returning the maximum of both values over all $r \in R$.

Given a fixed “ground truth” quantization Q we repeat this process for $\tau = 500$ trials of R , each returning a $d_\infty(R, Q)$ value. The set of these τ maximum deviations values results in another quantization S for each of `diam` and `dwid`, plotted in Figure 5. Intuitively, the maximum deviation quantization S describes the sample probability that $d_\infty(R, Q)$ will be less than some query value.




Fig. 5: Shows quantizations of $\tau = 500$ trials for $d_\infty(R, Q)$ where Q and R measure `dwid` and `diam`. The size of each R is $m = \{16, 64, 256, 1024, 4096\}$ (from right to left) and the “ground truth” quantization Q has size $\eta = 100000$. Smooth, thick curves are $1 - \delta = 1 - \exp(-2m\varepsilon^2 + 1)$ where $\varepsilon = d_\infty(R, Q)$.

Note that the maximum deviation quantizations S are similar for both statistics (and others we tried), and thus we can use these plots to estimate $1 - \delta$, the sample probability that $d_\infty(R, Q) \leq \varepsilon$, given a value m . We can fit this function as approximately $1 - \delta = 1 - \exp(-m\varepsilon^2/C + \nu)$ with $C = 0.5$ and $\nu = 1.0$. Thus solving for m in terms of ε , ν , and δ reveals: $m = C(1/\varepsilon^2)(\nu + \log(1/\delta))$. This indicates the big-O notation for the asymptotic bound of $O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ [39] for ε -samples only hides a constant of approximately 0.5.

We also ran these experiments to k -variate quantizations by considering the width in k different directions. As expected, the quantizations for maximum deviation can be fit with an equation $1 - \delta = 1 - \exp(-m\varepsilon^2/C + k)$ with $C = 0.5$, so $m \leq C(1/\varepsilon^2)(k + \log 1/\delta)$. For $k > 2$, this bound for m becomes too conservative; even fewer samples were needed.

4 Deterministic Computations on Indecisive Point Sets

In this section, we take as input a set of n indecisive points, and describe deterministic exact algorithms for creating quantizations of classes of functions on this input. We characterize problems when these deterministic algorithms can or can not be made efficient.

4.1 Polynomial Time Algorithms

We are interested in the distribution of the value $f(Q)$ for each support $Q \in \mathcal{P}$. Since there are k^n possible supports, in general we cannot hope to do anything faster than that without making additional assumptions about f . Define $\tilde{f}(\mathcal{P}, r)$ as the fraction (measured by weight) of supports of \mathcal{P} for which f gives a value smaller than or equal to r . In this version, for simplicity, we assume general position and that k^n can be described by $O(1)$ words, (handled otherwise in Appendix C). First, we will let $f(Q)$ denote the radius of the smallest enclosing disk of Q in the plane, and show how to solve the decision problem in polynomial time in that case. We then show how to generalize the ideas to other classes of measures.

Smallest enclosing disk. Consider the problem where f measures the radius of the smallest enclosing disk of a support and let all weights be uniform so $w(q_{i,j}) = 1$ for all i and j . Evaluating $\tilde{f}(\mathcal{P}, r)$ in time polynomial in n and k is not completely trivial since there are k^n possible supports. However, we can make use of the



Fig. 6: (a) The smallest enclosing circle of a set of points is defined by two or three points on the boundary. (b) This circle contains one purple (dark) point, four blue (medium) points, and two yellow (light) points. Hence there are $1 \times 4 \times 2 = 8$ samples that have this basis.




Fig. 7: (a) Example input with $n = 3$ and $k = 3$. (b) One possible basis, consisting of 3 points. This basis has one support: the basis itself. (c) Another possible basis, consisting of 2 points. This basis has three supports. (d) The graph showing for each diameter d how many supports do not exceed that diameter. This corresponds to the cumulative distribution of the radius of the smallest enclosing disk of these points.

fact that each smallest enclosing disk is in fact defined by a set of at most 3 points that lie on the boundary of the disk. For each support $Q \in \mathcal{P}$ we define $B_Q \subseteq Q$ to be this set of at most 3 points, which we call the *basis* for Q . Bases have the property that $f(Q) = f(B_Q)$.

Now, to avoid having to test an exponential number of supports, we define a *potential basis* to be a set of at most 3 points in \mathcal{P} such that each point is from a different P_i . Clearly, there are at most $(nk)^3$ possible potential bases, and each support $Q \in \mathcal{P}$ has one as its basis. Now, we only need to count for each potential basis the number of supports it represents. Counting the number of samples that have a certain basis is easy for the smallest enclosing circle. Given a basis B , we count for each indecisive point P that does not contribute a point to B itself how many of its members lie inside the smallest enclosing circle of B , and then we multiply these numbers. Figure 7 illustrates the idea.

Now, for each potential basis B we have two values: the number of supports that have B as their basis, and the value $f(B)$. We can sort these $O((nk)^3)$ pairs on the value of f , and the result provides us with the required distribution. We spend $O(nk)$ time per potential basis for counting the points inside and $O(n)$ time for multiplying these values, so combined with $O((nk)^3)$ potential bases this gives $O((nk)^4)$ total time.

Theorem 4.1. *Let \mathcal{P} be a set of n sets of k points. In $O((nk)^4)$ time, we can compute a data structure of $O((nk)^3)$ size that can tell us in $O(\log(nk))$ time for any value r how many supports of $Q \in \mathcal{P}$ satisfy $f(Q) \leq r$.*

LP-type problems. The approach described above also works for measures $f : \mathcal{P} \rightarrow \mathbb{R}$ other than the smallest enclosing disk. In particular, it works for LP-type problems [52] that have constant combinatorial dimension. An *LP-type problem* provides a set of constraints H and a function $\omega : 2^H \rightarrow \mathbb{R}$ with the following two properties:

MONOTONICITY: For any $F \subseteq G \subseteq H$, $\omega(F) \leq \omega(G)$.
LOCALITY: For any $F \subseteq G \subseteq H$ with $\omega(F) = \omega(G)$
and an $h \in H$ such that $\omega(G \cup h) > \omega(G)$
implies that $\omega(F \cup h) > \omega(F)$.

A *basis* for an LP-type problem is a subset $B \subset H$ such that $\omega(B') < \omega(B)$ for all proper subsets B' of B . And we say that B is a basis for a subset $G \subseteq H$ if $B \subseteq G$, $\omega(B) = \omega(G)$ and B is a basis. A constraint $h \in H$ *violates* a basis B if $\omega(B \cup h) > \omega(B)$. The radius of the smallest enclosing ball is an LP-type problem (where the points are the constraints and $\omega(\cdot) = f(\cdot)$) as are linear programming and many other geometric problems. Let the maximum cardinality of any basis be the *combinatorial dimension* of a problem.

For our algorithm to run efficiently, we assume that our LP-type problem has available the following algorithmic primitive, which is often assumed for LP-type problems with constant combinatorial dimension [52]. For a subset $G \subset H$ where B is known to be the basis of G and a constraint $h \in H$, a *violation test* determines in $O(1)$ time if $\omega(B \cup h) > \omega(B)$; i.e., if h *violates* B . More specifically, given an efficient violation test, we can ensure a stronger algorithmic primitive. A *full violation test* is given a subset $G \subset H$ with known basis B and a constraint $h \in H$ and determines in $O(1)$ time if $\omega(B) < \omega(G \cup h)$. This follows because we can test in $O(1)$ time if $\omega(B) < \omega(B \cup h)$; MONOTONICITY implies that $\omega(B) < \omega(B \cup h)$ only if $\omega(B) < \omega(B \cup h) \leq \omega(G \cup h)$, and LOCALITY implies that $\omega(B) = \omega(B \cup h)$ only if $\omega(B) = \omega(G) = \omega(G \cup h)$. Thus we can test if h violates G by considering just B and h , but if either MONOTONICITY or LOCALITY fail for our problem we cannot.

We now adapt our algorithm to LP-type problems where elements of each P_i are potential constraints and the ranking function is f . When the combinatorial dimension is a constant β , we need to consider only $O((nk)^\beta)$ bases, which will describe all possible supports.

The full violation test implies that given a basis B , we can measure the sum of probabilities of all supports of \mathcal{P} that have B as their basis in $O(nk)$ time. For each indecisive point P such that $B \cap P = \emptyset$, we sum the probabilities of all elements of P that do not violate B . The product of these probabilities times the product of the probabilities of the elements in the basis, gives the probability of B being the true basis. See Algorithm 4.1 where the indicator function applied $1(f(B \cup \{p_j\}) = f(B))$ returns 1 if p_j does not violate B and 0 otherwise. It runs in $O((nk)^{\beta+1})$ time.

Algorithm 4.1 Construct Probability Distribution for $f(\mathcal{P})$.

```

1: for all potential bases  $B \subset Q \in \mathcal{P}$  do
2:   for  $i = 1$  to  $n$  do
3:     if there is a  $j$  such that  $p_{ij} \in B$  then
4:       Set  $w_i = w(p_{ij})$ .
5:     else
6:       Set  $w_i = \sum_{j=1}^k w(p_{ij})1(f(B \cup \{p_j\}) = f(B))$ .
7:   Store a point with value  $f(B)$  and weight  $(1/k^n) \prod_i w_i$ .
```

As with the special case of smallest enclosing disk, we can create a distribution over the values of f given an indecisive point set \mathcal{P} . For each basis B we calculate $\mu(B)$, the summed probability of all supports that have basis B , and $f(B)$. We can then sort these pairs according to the value as f again. For any query value r , we can retrieve $\hat{f}(\mathcal{P}, r)$ in $O(\log(nk))$ time and it takes $O(n)$ time to describe (because of its long length).

Theorem 4.2. *Given a set \mathcal{P} of n indecisive point sets of size k each, and given an LP-type problem $f : \mathcal{P} \rightarrow \mathbb{R}$ with combinatorial dimension β , we can create the distribution of f over \mathcal{P} in $O((nk)^{\beta+1})$ time. The size of the distribution is $O(n(nk)^\beta)$.*

If we assume general position of \mathcal{P} relative to f , then we can often slightly improve the runtime needed to calculate $\mu(B)$ using range searching data structures. However, to deal with degeneracies, we may need to spend $O(nk)$ time per basis, regardless.

If we are content with an approximation of the distribution rather than an exact representation, then it is often possible to drastically reduce the storage and runtime following techniques discussed in Section 3.

Measures that fit in this framework for points in \mathbb{R}^d include smallest enclosing axis-aligned rectangle (measured either by area or perimeter) ($\beta = 2d$), smallest enclosing ball in the L_1 , L_2 , or L_∞ metric ($\beta = d + 1$), directional width of a set of points ($\beta = 2$), and, after dualizing, linear programming ($\beta = d$). These approaches also carry over naturally to deterministically create polynomial-sized k -variate quantizations.

4.2 Hardness Results

In this section, we examine some extent measures that do not fit in the above framework. First, diameter does not satisfy the LOCALITY property, and hence we cannot efficiently perform the full violation test. We show that a decision variant of diameter is #P-Hard, even in the plane, and thus (under the assumption that #P \neq P), there is no polynomial time solution. This result holds despite the fact that diameter has a combinatorial dimension of 2, implying that the associated quantization has at most $O((nk)^2)$ steps. Second, the area of the convex hull does not have a constant combinatorial dimension, thus we can show the resulting distribution may have exponential size.

Diameter. The *diameter* of a set of points in the plane is the largest distance between any two points. We will show that the counting problem of computing $\tilde{f}(\mathcal{P}, r)$ is #P-hard when f denotes the diameter.

Problem 4.3. *PLANAR-DIAM:* Given a parameter d and a set $\mathcal{P} = \{P_1, \dots, P_n\}$ of n sets, each consisting of k points in the plane, how many supports $Q \subseteq \mathcal{P}$ have $f(Q) \leq d$?

We will now prove that Problem 4.3 is #P-hard. Our proof has three steps. We first show a special version of #2SAT has a polynomial reduction from Monotone #2SAT, which is #P-complete [56]. Then, given an instance of this special version of #2SAT, we construct a graph with weighted edges on which the diameter problem is equivalent to this #2SAT instance. Finally, we show the graph can be embedded as a straight-line graph in the plane as an instance of PLANAR-DIAM.

Let 3CLAUSE-#2SAT be the problem of counting the number of solutions to a 2SAT formula, where each variable occurs in at most three clauses, and each variable is either in exactly one clause or is negated in exactly one clause. Thus, each distinct literal appears in at most two clauses.

Lemma 4.4. *Monotone #2SAT has a polynomial reduction to 3CLAUSE-#2SAT.*

Proof. The Monotone #2SAT problem counts the number satisfying assignments to a #2SAT instance where each clause has at most two variables and no variables are negated. Let $X = \{(x, y_1), (x, y_2), \dots, (x, y_u)\}$ be the set of u clauses which contain variable x in a Monotone #2SAT instance. We replace x with u variables $\{z_1, z_2, \dots, z_u\}$ and we replace X with the following $2u$ clauses $\{(z_1, y_1), (z_2, y_2), \dots, (z_u, y_u)\}$ and $\{(z_1, \neg z_2), (z_2, \neg z_3), \dots, (z_{u-1}, \neg z_u), (z_u, \neg z_1)\}$. The first set of clauses preserves the relation with other original variables and the second set of clauses ensures that all of the new variables have the same value (i.e. TRUE or FALSE). This procedure is repeated for each original variable that is in more than 1 clause. \square

We convert this problem into a graph problem by, for each variable x_i , creating a set $P_i = \{p_i^+, p_i^-\}$ of two points. Let $S = \bigcup_i P_i$. Truth assignments of variables correspond to a support as follows. If x_i is set TRUE, then the support includes p_i^+ , otherwise the support includes p_i^- . We define a distance function f between points, so that the distance is greater than d (long) if the corresponding literals are in a clause, and less than d (short) otherwise. If we consider the graph formed by only long edges, we make two observations. First, the maximum degree is 2, since each literal is in at most two clauses. Second, there are no cycles since a literal is only in two clauses if in one clause the other variable is negated, and negated variables are in only one clause. These two properties imply we can use the following construction to show that the PLANAR-DIAM problem is as hard as counting Monotone #2SAT solutions, which is #P-complete.

Lemma 4.5. *An instance of PLANAR-DIAM reduced from 3CLAUSE-#2SAT can be embedded so $\mathcal{P} \subset \mathbb{R}^2$.*

Proof. Consider an instance of 3CLAUSE-#2SAT where there are n variables, and thus the corresponding graph has n sets $\{P_i\}_{i=1}^n$. We construct a sequence Γ of $n' \in [2n, 4n]$ points. It contains all points from \mathcal{P}




Fig. 8: Embedded points are solid, at center of circles of radius d . Dummy points hollow. Long edges are drawn between points at distance greater than d .

and a set of at most as many dummy points. First organize a sequence Γ' so if two points q and p have a long edge, then they are consecutive. Now for any pair of consecutive points in Γ' which do not have a long edge, insert a dummy point between them to form the sequence Γ . Also place a dummy point at the end of Γ .

We place all points on a circle C of diameter $d/\cos(\pi/n')$, see Figure 8. We first place all points on a semicircle of C according to the order of Γ , so each consecutive points are π/n' radians apart. Then for every other point (i.e. the points with an even index in the ordering Γ) we replace it with its antipodal point on C , so no two points are within $2\pi/n'$ radians of each other. Finally we remove all dummy points. This completes the embedding of \mathcal{P} , we now need to show that only points with long edges are further than d apart.

We can now argue that only vertices which were consecutive in Γ are further than d apart, the remainder are closer than d . Consider a vertex p and a circle C_p of radius d centered at p . Let p' be the antipodal point of p on C . C_p intersects C at two points, at $2\pi/n'$ radians in either direction from p' . Thus only points within $2\pi/n'$ radians of p' are further than a distance d from p . This set includes only those points which are adjacent to p in Γ , which can only include points which should have a long edge, by construction. \square

Combining Lemmas 4.4 and 4.5:

Theorem 4.6. *PLANAR-DIAM is #P-hard.*

Convex hull. Our LP-type framework also does not work for any properties of the convex hull (e.g. area or perimeter) because it does not have constant combinatorial dimension; a basis could have size n . In fact, the complexity of the distribution describing the convex hull may be $\Omega(k^n)$, since if all points in \mathcal{P} lie on or near a circle, then every support $Q \subseteq \mathcal{P}$ may be its own basis of size n , and have a different value $f(Q)$.

5 Deterministic Algorithms for Approximate Computations on Uncertain Points

In this section we show how to approximately answer questions about most representations of independent uncertain points; in particular, we handle representations that have almost all ($1 - \varepsilon$ fraction) of their mass with bounded support in \mathbb{R}^d and is described in a compact manner (see Appendix A). Specifically, in this section, we are given a set $\mathcal{P} = \{P_1, P_2, P_3, \dots, P_n\}$ of n independent random variables over the universe \mathbb{R}^d , together with a set $\mu_{\mathcal{P}} = \{\mu_1, \mu_2, \mu_3, \dots, \mu_n\}$ of n probability distributions that govern the variables, that is, $X_i \sim \mu_i$. Again, we call a set of points $Q = \{q_1, q_2, q_3, \dots, q_n\}$ a support of \mathcal{P} , and because of the independence we have probability $Pr[\mathcal{P} = Q] = \prod_i Pr[P_i = p_i]$.

The main strategy will be to replace each distribution μ_i by a discrete point set P_i , such that the uniform distribution over P_i is “not too far” from μ_i (P_i is not the most obvious ε -sample of μ_i). Then we apply the algorithms from Section 4 to the resulting set of point sets. Finally, we argue that the result is in fact an ε -quantization of the distribution we are interested in. Using results from Section 3 we can simplify the

output in order to decrease the space complexity for the data structure, without increasing the approximation factor too much.

General approach. Given a distribution $\mu_i : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ describing uncertain point P_i and a function f of bounded combinatorial dimension β defined on a support of \mathcal{P} , we can describe a straightforward range space $T_i = (\mu_i, \mathcal{A}_f)$, where \mathcal{A}_f is the set of ranges corresponding to the bases of f (e.g., when f measures the radius of the smallest enclosing ball, \mathcal{A}_f would be the set of all balls). More formally, \mathcal{A}_f is the set of subsets of \mathbb{R}^d defined as follows: for every set of β points which define a basis B for f , \mathcal{A}_f contains a range A that contains all points p such that $f(B) = f(B \cup \{p\})$. However, taking ε -samples from each T_i is *not* sufficient to create sets Q_i such that $\mathcal{Q} = \{Q_1, Q_2, \dots, Q_n\}$ so for all r we have $|\tilde{f}(\mathcal{P}, r) - \tilde{f}(\mathcal{Q}, r)| \leq \varepsilon$.

$\tilde{f}(\mathcal{P}, r)$ is a complicated joint probability depending on the n distributions and f , and the n straightforward ε -samples do not contain enough information to decompose this joint probability. The required ε -sample of each μ_i should model μ_i in relation to f and any instantiated point p_i representing μ_j for $i \neq j$. The following crucial definition allows for the range space to depend on any $n - 1$ points, including the possible locations of each uncertain point.

Let $\mathcal{A}_{f,n}$ describe a family of Lebesgue-measurable sets defined by $n - 1$ points $Z \subset \mathbb{R}^d$ and a value w . Specifically, $A(Z, w) \in \mathcal{A}_{f,n}$ is the set of points $\{p \in \mathbb{R}^d \mid f(Z \cup p) \leq w\}$. We describe examples of $\mathcal{A}_{f,n}$ in detail shortly, but first we state the key theorem using this definition. Its proof, delayed until after examples of $\mathcal{A}_{f,n}$, will make clear how $(\mu_i, \mathcal{A}_{f,n})$ exactly encapsulates the right guarantees to approximate $\hat{f}(\mathcal{P}, r)$, and thus why (μ_i, \mathcal{A}_f) does not.

Theorem 5.1. *Let $\mathcal{P} = \{P_1, \dots, P_n\}$ be a set of uncertain points where each $P_i \sim \mu_i$. For a function f , let Q_i be an ε' -sample of $(\mu_i, \mathcal{A}_{f,n})$ and let $\mathcal{Q} = \{Q_1, \dots, Q_n\}$. Then for any r , $|\hat{f}(\mathcal{P}, r) - \hat{f}(\mathcal{Q}, r)| \leq \varepsilon'n$.*

Smallest axis-aligned bounding box by perimeter. Given a set of points $P \subset \mathbb{R}^2$, let $f(P)$ represent the perimeter of the smallest axis-aligned box that contains P . Let each μ_i be a bivariate normal distribution with constant variance. Solving $f(P)$ is an LP-type problem with combinatorial dimension $\beta = 4$, and as such, we can describe the basis B of a set P as the points with minimum and maximum x - and y -coordinates. Given any additional point p , the perimeter of size ρ can only be increased to a value w by expanding the range of x -coordinates, y -coordinates, or both. As such, the region of \mathbb{R}^2 described by a range $A(P, w) \in \mathcal{A}_{f,n}$ is defined with respect to the bounding box of P from an edge increasing the x -width or y -width by $(w - \rho)/2$, or from a corner extending so the sum of the x and y deviation is $(w - \rho)/2$. See Figure 9(a).

Since any such shape defining a range $A(P, w) \in \mathcal{A}_{f,n}$ can be described as the intersection of $k = 4$ slabs along fixed axis (at 0° , 45° , 90° , and 135°), we can construct an (ε/n) -sample Q_i of $(\mu_i, \mathcal{A}_{f,n})$ of size $k = O((n/\varepsilon) \log^8(n/\varepsilon))$ in $O((n^6/\varepsilon^6) \log^{27}(n/\varepsilon))$ time [49]. From Theorem 5.1, it follows that for $\mathcal{Q} = \{Q_1, \dots, Q_n\}$ and any r we have $|\hat{f}(\mathcal{X}, r) - \hat{f}(\mathcal{Q}, r)| \leq \varepsilon$.

We can then apply Theorem 4.2 to build an ε -quantization of $f(\mathcal{X})$ in $O((nk)^5) = O(((n^2/\varepsilon) \log^8(n/\varepsilon))^5) = O((n^{10}/\varepsilon^5) \log^{40}(n/\varepsilon))$ time. The size can be reduced to $O(1/\varepsilon)$ within that time bound.

Corollary 5.2. *Let $\mathcal{P} = \{P_1, \dots, P_n\}$ be a set of indecisive points where each $P_i \sim \mu_i$ is bivariate normal with constant variance. Let f measure the perimeter of the smallest enclosing axis-aligned bounding box. We can create an ε -quantization of $f(\mathcal{P})$ in $O((n^{10}/\varepsilon^5) \log^{40}(n/\varepsilon))$ time of size $O(1/\varepsilon)$.*

Smallest enclosing disk. Given a set of points $P \subset \mathbb{R}^2$, let $f(P)$ represent the radius of the smallest enclosing disk of P . Let each μ_i be a bivariate normal distribution with constant variance. Solving $f(P)$ is an LP-type problem with combinatorial dimension $\beta = 3$, and the basis B of P generically consists of either 3 points which lie on the boundary of the smallest enclosing disk, or 2 points which are antipodal on the smallest enclosing disk. However, given an additional point $p \in \mathbb{R}^2$, the new basis B_p is either B or it is p along with 1 or 2 points which lie on the convex hull of P .




Fig. 9: (a) A shape from $\mathcal{A}_{f,n}$ for axis-aligned bounding box, measured by perimeter. (b) A shape from $\mathcal{A}_{f,n}$ for smallest enclosing ball using the L_2 metric in \mathbb{R}^2 . The curves are circular arcs of two different radii. (c) The same shape divided into wedges from $\mathcal{W}_{f,n}$.

We can start by examining all pairs of points $p_i, p_j \in P$ and the two disks of radius w whose boundary circles pass through them. If one such disk $D_{i,j}$ contains P , then $D_{i,j} \subset A(P, w) \in \mathcal{A}_{f,|P|+1}$. For this to hold, p_i and p_j must lie on the convex hull of P and no point that lies between them on the convex hull can contribute to such a disk. Thus there are $O(n)$ such disks. We also need to examine the disks created where p and one other point $p_i \in P$ are antipodal. The boundary of the union of all such disks which contain P is described by part of a circle of radius $2w$ centered at some $p_i \in P$. Again, for such a disk B_i to describe a part of the boundary of $A(P, w)$, the point p_i must lie on the convex hull of P . The circular arc defining this boundary will only connect two disks $D_{i,j}$ and $D_{k,i}$ because it will intersect with the boundary of B_j and B_k within these disks, respectively. An example of $A(P, w)$ is shown in Figure 9(b).

Unfortunately, the range space $(\mathbb{R}^2, \mathcal{A}_{f,n})$ has VC-dimension $O(n)$; it has $O(n)$ circular boundary arcs. So, creating an ε -sample of $T_i = (\mu_i, \mathcal{A}_{f,n})$ would take time exponential in n . However, we can decompose any range $A(P, w) \in \mathcal{A}_{f,n}$ into at most $2n$ “wedges.” We choose one point y inside the convex hull of P . For each circular arc on the boundary of $A(P, w)$ we create a wedge by coning that boundary arc to y . Let \mathcal{W}_f describe all wedge shaped ranges. Then $S = (\mathbb{R}^2, \mathcal{W}_f)$ has VC-dimension ν_S at most 9 since it is the intersection of 3 ranges (two halfspaces and one disk) that can each have VC-dimension 3. We can then create Q_i , an $(\varepsilon/2n^2)$ -sample of $S_i = (\mu_i, \mathcal{W}_f)$, of size $k = O((n^4/\varepsilon^2) \log(n/\varepsilon))$ in $O((n^2/\varepsilon)^{5+2 \cdot 9} \log^{2+9}(n/\varepsilon)) = O((n^{46}/\varepsilon^{23}) \log^{11}(n/\varepsilon))$ time, via Corollary A.2 (Appendix A). It follows that Q_i is an (ε/n) -sample of $T_i = (\mu_i, \mathcal{A}_{f,n})$, since any range $A(Z, w) \in \mathcal{A}_{f,n}$ can be decomposed into at most $2n$ wedges, each of which has counting error at most $\varepsilon/2n$, thus the total counting error is at most ε .

Invoking Theorem 5.1, it follows that $\mathcal{Q} = \{Q_1, \dots, Q_n\}$, for any r we have $|\hat{f}(\mathcal{P}, r) - \tilde{f}(\mathcal{Q}, r)| \leq \varepsilon$. We can then apply Theorem 4.2 to build an ε -quantization of $f(\mathcal{P})$ in $O((nk)^4) = O((n^{20}/\varepsilon^8) \log^4(n/\varepsilon))$ time. This is dominated by the time for creating the n (ε/n^2) -samples, even though we only need to build one and then translate and scale to the rest. Again, the size can be reduced to $O(1/\varepsilon)$ within that time bound.

Corollary 5.3. *Let $\mathcal{P} = \{P_1, \dots, P_n\}$ be a set of indecisive points where each $P_i \sim \mu_i$ is bivariate normal with constant variance. Let f measure the radius of the smallest enclosing disk. We can create an ε -quantization of $f(\mathcal{P})$ in $O((n^{46}/\varepsilon^{23}) \log^{11}(n/\varepsilon))$ time of size $O(1/\varepsilon)$.*

Now that we have seen two concrete examples, we prove Theorem 5.1. More examples can be found in Appendix B.

Proof of Theorem 5.1. When each P_i is drawn from a distribution μ_i , then we can write $\hat{f}(\mathcal{P}, r)$ as the probability that $f(\mathcal{P}) \leq r$ as follows. Let $1(\cdot)$ be the indicator function, i.e., it is 1 when the condition is true and 0 otherwise.

$$\hat{f}(\mathcal{P}, r) = \int_{p_1} \mu_1(p_1) \dots \int_{p_n} \mu_n(p_n) 1(f(\{p_1, p_2, \dots, p_n\}) \leq r) dp_n dp_{n-1} \dots dp_1$$

Consider the inner most integral

$$\int_{p_n} \mu_n(p_n) \mathbf{1}(f(\{p_1, p_2, \dots, p_n\}) \leq r) dp_n,$$

where $\{p_1, p_2, \dots, p_{n-1}\}$ are fixed. The indicator function $\mathbf{1}(\cdot) = 1$ is true when $f(\{p_1, p_2, \dots, p_{n-1}, p_n\}) \leq r$, and hence p_n is contained in a shape $A(\{p_1, \dots, p_{n-1}\}, r) \in \mathcal{A}_{f,n}$. Thus if we have an ε' -sample Q_n for $(\mu_n, \mathcal{A}_{f,n})$, then we can guarantee that

$$\int_{p_n} \mu_n(p_n) \mathbf{1}(f(\{p_1, p_2, \dots, p_n\}) \leq r) dp_n \leq \frac{1}{|Q_n|} \sum_{p_n \in Q_n} \mathbf{1}(f(\{p_1, p_2, \dots, p_{n-1}, p_n\}) \leq r) + \varepsilon'.$$

We can then move the ε' outside and change the order of the integrals to write:

$$\hat{f}(\mathcal{X}, r) \leq \frac{1}{|Q_n|} \sum_{p_n \in Q_n} \left(\int_{p_1} \mu_1(p_1) \dots \int_{p_{n-1}} \mu_{n-1}(p_{n-1}) \mathbf{1}(f(\{p_1, \dots, p_n\}) \leq r) dp_{n-1} \dots dp_1 \right) + \varepsilon'.$$

Repeating this procedure n times we get:

$$\hat{f}(\mathcal{X}, r) \leq \left(\prod_{i=1}^n \frac{1}{|Q_i|} \right) \sum_{p_1 \in Q_1} \dots \sum_{p_n \in Q_n} \mathbf{1}(f(\{p_1, \dots, p_n\}) \leq r) + \varepsilon' n = \tilde{f}(\mathcal{Q}, r) + \varepsilon' n,$$

where $\mathcal{Q} = \bigcup_i Q_i$.

Similarly we can achieve a symmetric lower bound for $\hat{f}(\mathcal{X}, r)$. □

6 Shape Inclusion Probabilities

So far, we have been concerned only with computing probability distributions of single-valued functions on a set of points. However, many geometric algorithms produce more than just a single value. In this section, we consider dealing with uncertainty when computing a two-dimensional shape (such as the convex hull, or MEB) of a set of points directly.

6.1 Randomized Algorithms

We can also use a variation of Algorithm 3.1 to construct ε -shape inclusion probability functions. For a point set $Q \subset \mathbb{R}^d$, let the summarizing shape $S_Q = \mathcal{S}(Q)$ be from some geometric family \mathcal{S} so $(\mathbb{R}^d, \mathcal{S})$ has bounded VC-dimension ν . We randomly sample m point sets $\mathcal{Q} = \{Q_1, \dots, Q_m\}$ each from μ_P and then find the summarizing shape $S_{Q_j} = \mathcal{S}(Q_j)$ (e.g. minimum enclosing ball) of each Q_j . Let this set of shapes be $S^{\mathcal{Q}}$. If there are multiple shapes from \mathcal{S} which are equally optimal, choose one of these shapes at random. For a set of shapes $S' \subseteq \mathcal{S}$, let $S'_p \subseteq S'$ be the subset of shapes that contain $p \in \mathbb{R}^d$. We store $S^{\mathcal{Q}}$ and evaluate a query point $p \in \mathbb{R}^d$ by counting what fraction of the shapes the point is contained in, specifically returning $|S'_p \cap S^{\mathcal{Q}}| / |S^{\mathcal{Q}}|$ in $O(\nu |S^{\mathcal{Q}}|) = O(\nu m)$ time. In some cases, this evaluation can be sped up with point location data structures.

To state the main theorem most cleanly, for a range space $(\mathbb{R}^d, \mathcal{S})$, denote its dual range space as (\mathcal{S}, P^*) where P^* is all subsets $S_p \subseteq \mathcal{S}$, for any point $p \in \mathbb{R}^d$, such that $S_p = \{S \in \mathcal{S} \mid p \in S\}$. Recall that the VC-dimension ν of (\mathcal{S}, P^*) is at most $2^{\nu'+1}$ where ν' is the VC-dimension of $(\mathbb{R}^d, \mathcal{S})$, but is typically much smaller.

Theorem 6.1. *Consider a family of summarizing shapes \mathcal{S} with dual range space (\mathcal{S}, P^*) with VC-dimension ν , and where it takes $T_{\mathcal{S}}(n)$ time to determine the summarizing shape $\mathcal{S}(Q)$ for any point set $Q \subset \mathbb{R}^d$ of size n . For a distribution μ_P of a point set of size n , with probability at least $1 - \delta$, we can construct an ε -sip function of size $O((\nu/\varepsilon^2)(\nu + \log(1/\delta)))$ and in time $O(T_{\mathcal{S}}(n)(1/\varepsilon^2) \log(1/\delta))$.*

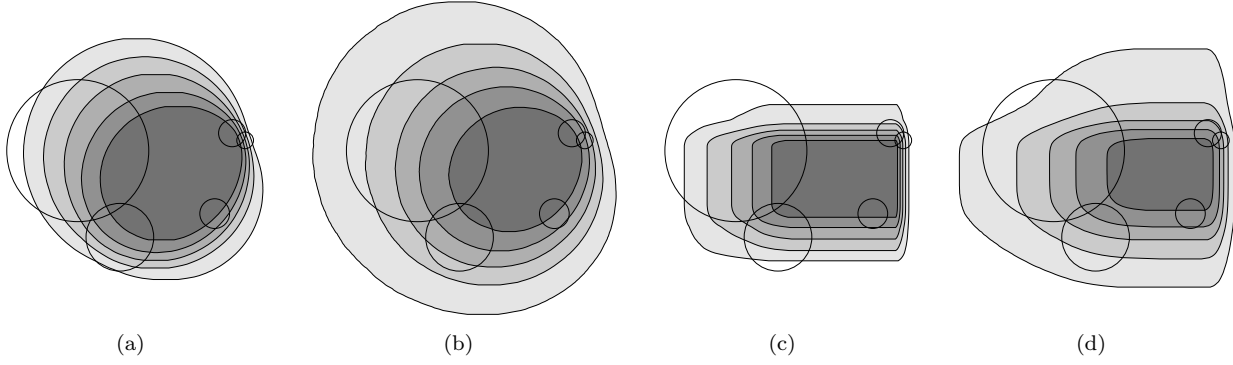


Fig. 10: The sip for the smallest enclosing ball (a,b) or smallest enclosing axis-aligned rectangle (c,d), for uniformly (a,c) or normally (b,d) distributed points. Isolines are drawn for $p \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$.

Proof. Using the above algorithm, sample $m = O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ point sets Q from μ_P and generate the m summarizing shapes S_Q . Each shape is a random sample from \mathcal{S} according to μ_P , and thus S^Ω is an ε -sample of (\mathcal{S}, P^*) .

Let $w_{\mu_P}(S)$, for $S \in \mathcal{S}$, be the probability that S is the summarizing shape of a point set Q drawn randomly from μ_P . For any $S' \subseteq P^*$, let $W_{\mu_P}(S') = \int_{S \in S'} w_{\mu_P}(S) dS$ be the probability that some shape from the subset S' is the summarizing shape of Q drawn from μ_P .

We approximate the sip function at $p \in \mathbb{R}^d$ by returning the fraction $|S_p^\Omega|/m$. The true answer to the sip function at $p \in \mathbb{R}^d$ is $W_{\mu_P}(\mathcal{S}_p)$. Since S^Ω is an ε -sample of (\mathcal{S}, P^*) , then with probability at least $1 - \delta$

$$\left| \frac{|S_p^\Omega|}{m} - \frac{W_{\mu_P}(\mathcal{S}_p)}{1} \right| = \left| \frac{|S_p^\Omega|}{|S^\Omega|} - \frac{W_{\mu_P}(\mathcal{S}_p)}{W_{\mu_P}(P^*)} \right| \leq \varepsilon. \quad \square$$

Representing ε -sip functions by isolines. Shape inclusion probability functions are density functions. A convenient way of visually representing a density function in \mathbb{R}^2 is by drawing the isolines. A γ -*isoline* is a collection of closed curves bounding the regions of the plane where the density function is greater than γ .

In each part of Figure 10 a set of 5 circles correspond to points with a probability distribution. In part (a,c) the probability distribution is uniform over the inside of the circles. In part (b,d) it is drawn from a normal distribution with standard deviation given by the radius. We generate ε -sip functions for the smallest enclosing ball in Figure 10(a,b) and for the smallest axis-aligned rectangle in Figure 10(c,d).

In all figures we draw approximations of $\{.9, .7, .5, .3, .1\}$ -isolines. These drawing are generated by randomly selecting $m = 5000$ (Figure 10(a,b)) or $m = 25000$ (Figure 10(c,d)) shapes, counting the number of inclusions at different points in the plane and interpolating to get the isolines. The innermost and darkest region has probability $> 90\%$, the next one probability $> 70\%$, etc., the outermost region has probability $< 10\%$.

6.2 Deterministic Algorithms

We can also adapt the deterministic algorithms presented in Section 4.1 to deterministically create SIP functions for indecisive (or ε -SIPs for many uncertain) points. We again restrict our attention to a class of LP-type problems, specifically, problems where the output is a (often minimal) summarizing shape $S(Q)$ of a data set Q where the boundaries of the two shapes $S(Q)$ and $S(Q')$ intersect in at most a constant number of locations. An example is the smallest enclosing disk in the plane, where the circles on the boundaries of any two disks intersect at most twice.

As in Section 4.1, this problem has a constant combinatorial dimension β and possible locations of indecisive points can be labeled inside or outside $S(Q)$ using a full violation test. This implies, following Algorithm 4.1, that we can enumerate all $O((nk)^\beta)$ potential bases, and for each determine its weight towards

the SIP function using the full violation test. This procedure generates a set of $O((nk)^\beta)$ weighted shapes in $O((nk)^{\beta+1})$ time. Finally, a query to the SIP function can be evaluated by counting the weighted fraction of shapes that it is contained in.

We can also build a data structure to speed up the query time. Since the boundary of each pair of shapes intersects a constant number of times, then all $O((nk)^{2\beta})$ pairs intersect at most $O((nk)^{2\beta})$ times in total. In \mathbb{R}^2 , the arrangement of these shapes forms a planar subdivision with as many regions as intersection points. We can precompute the weighted fraction of shapes overlapping on each region. Textbook techniques can be used to build a query structure of size $O((nk)^{2\beta})$ that allows for stabbing queries in time $O(\log(nk))$ to determine which region the query lies, and hence what the associated weighted fraction of points is, and what the SIP value is. We summarize these results in the following theorem.

Theorem 6.2. *Consider a set \mathcal{P} of n indecisive point sets of size k each, and an LP-type problem $f_S : \mathcal{P} \rightarrow \mathbb{R}$ with combinatorial dimension β that finds a summarizing shape for which every pair intersects a constant number of times. We can create a data structure of size $O((nk)^\beta)$ in time $O((nk)^{\beta+1})$ that answers SIP queries exactly in $O((nk)^\beta)$ time, and another structure of size $O((nk)^{2\beta})$ in time $O((nk)^{2\beta})$ that answers SIP queries exactly in $O(\log(nk))$ time.*

Furthermore, through specific invocations of Theorem 5.1, we can extend these polynomial deterministic approaches to create ε -SIP data structures for many natural classes of uncertain input point sets.

7 Conclusions

In this paper, we studied the computation and representation of complete probability distributions on the output of single- and multi-valued geometric functions, when the input points are uncertain. We considered randomized and deterministic, exact and approximate approaches for indecisive and probabilistic uncertain points, and presented polynomial-time algorithms as well as hardness results. These results extend to when the output distribution is over the family of low-description-complexity summarizing shapes.

We draw two main conclusions. Firstly, we observe that the tractability of exact computations on indecisive points really depends on the problem at hand. On the one hand, the output distribution of LP-type problems can be represented concisely and computed efficiently. On the other hand, even computing a single value of the output distribution of the diameter problem is already $\#P$ -hard. Secondly, we showed that computing approximate quantizations deterministically is often possible in polynomial time. However, the polynomials in question are of rather high degree, and while it is conceivable that these degrees can be reduced further, this will require some new ideas. In the mean time, the randomized alternatives remain more practical.

We believe that the problem of representing and approximating distributions of more complicated objects (especially when they are somehow representing uncertain points), is an important direction for further study.

Acknowledgments

The authors would like to thank Joachim Gudmundsson and Pankaj Agarwal for helpful discussions in early phases of this work, Sariel Har-Peled for discussions about wedges, and Suresh Venkatasubramanian for organizational tips.

This research was partially supported by the Netherlands Organisation for Scientific Research (NWO) through the project GOGO and under grant 639.021.123, by the Office of Naval Research under MURI grant N00014-08-1-1015, and by a subaward to the University of Utah under NSF award 0937060 to CRA.

References

- [1] C. C. Agarwal and P. S. Yu, editors. *Privacy Preserving Data Mining: Models and Algorithms*. Springer, 2008.
- [2] P. K. Agarwal, S.-W. Cheng, Y. Tao, and K. Yi. Indexing uncertain data. *Proceedings of ACM Principals of Database Systems*, 2009.
- [3] P. K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximations via coresets. *C. Trends Comb. and Comp. Geom. (E. Welzl)*, 2007.
- [4] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measure of points. *Journal of ACM* 51(4):2004, 2004.
- [5] P. K. Agarwal, C. M. Procopiuc, and K. R. Varadarajan. Approximation algorithms for k -line center. *Proceedings European Symposium on Algorithms*, pp. 54–63, 2002.
- [6] R. Agarwal and R. Srikant. Privacy-preserving data mining. *ACM SIGMOD Record* 29:439–450, 2000.
- [7] P. Agrawal, O. Benjelloun, A. D. Sarma, C. Hayworth, S. Nabar, T. Sugihara, and J. Widom. Trio: A system for data, uncertainty, and lineage. *Proceedings ACM Principals of Database Systems*, 2006.
- [8] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: Nearest neighbor relations for imprecise points. *Proceedings on ACM-SIAM Symposium on Discrete Algorithms*, 2004.
- [9] N. Bansal. Constructive algorithms for discrepancy minimization. *Proceedings 51st Annual IEEE Symposium on Foundations of Computer Science*, pp. 407–414, 2010.
- [10] J. Bi and T. Zhang. Support vector classification with input data uncertainty. *Proceedings Neural Information Processing Systems*, 2004.
- [11] M. Bădoiu and K. Clarkson. Smaller core-sets for balls. *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2003.
- [12] T. Chan. Faster core-set constructions and data-stream algorithms in fixed dimensions. *Computational Geometry: Theory and Applications* 35:20–35, 2006.
- [13] B. Chazelle and J. Matousek. On linear-time deterministic algorithms for optimization problems in fixed dimensions. *J. Algorithms* 21:579–597, 1996.
- [14] R. Cheng, J. Chen, M. Mokbel, and C.-Y. Chow. Probability verifiers: Evaluating constrained nearest-neighbor queries over uncertain data. *Proceedings International Conference on Data Engineering*, 2008.
- [15] R. Cheng, D. V. Kalashnikov, and S. Prabhakar. Evaluating probabilistic queries over imprecise data. *Proceedings 2003 ACM SIGMOD International Conference on Management of Data*, 2003.
- [16] G. Cormode and M. Garafalakis. Histograms and wavelets of probabilistic data. *IEEE 25th International Conference on Data Engineering*, 2009.
- [17] G. Cormode, F. Li, and K. Yi. Semantics of ranking queries for probabilistic data and expected ranks. *IEEE 25th International Conference on Data Engineering*, 2009.
- [18] G. Cormode and A. McGregor. Approximation algorithms for clustering uncertain data. *PODS*, 2008.
- [19] N. Dalvi and D. Suciu. Efficient query evaluation on probabilistic databases. *The VLDB Journal* 16:523–544, 2007.

- [20] A. Deshpande, C. Guestrin, S. R. Madden, J. M. Hellerstein, and W. Hong. Model-driven data acquisition in sensor networks. *Proceedings 30th International Conference on Very Large Data Bases*, 2004.
- [21] A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Annals of Mathematical Statistics* 27:642–669, 1956.
- [22] A. Eliazar and R. Parr. Dp-slam 2.0. *IEEE International Conference on Robotics and Automation*, 2004.
- [23] G. N. Frederickson and D. B. Johnson. Generalized selection and ranking: Sorted matrices. *SIAM Journal on Computing* 13:14–30, 1984.
- [24] M. Fürer. Faster integer multiplication. *SIAM J. Computing* 39:979–1005, 2009.
- [25] B. Gärtner. Fast and robust smallest enclosing balls. *Proceedings European Symposium on Algorithms*, 1999.
- [26] L. J. Guibas, D. Salesin, and J. Stolfi. Epsilon geometry: building robust algorithms from imprecise computations. *Proceedings Symposium on Computational Geometry*, pp. 208–217, 1989.
- [27] L. J. Guibas, D. Salesin, and J. Stolfi. Constructing strongly convex approximate hulls with inaccurate primitives. *Algorithmica* 9:534–560, 1993.
- [28] R. H. Güting and M. Schneider. *Moving Object Databases*. Morgan Kaufmann, San Francisco, 2005.
- [29] S. Har-Peled. Chapter 5: On complexity, sampling, and ε -nets and ε -samples. http://valis.cs.uiuc.edu/~sariel/teach/notes/aprx/lec/05_vc_dim.pdf, May 2010.
- [30] D. Haussler and E. Welzl. epsilon-nets and simplex range queries. *Discrete & Computational Geometry* 2:127–151, 1987.
- [31] M. Held and J. S. B. Mitchell. Triangulating input-constrained planar point sets. *Information Processing Letters* 109:54–56, 2008.
- [32] D. V. Kalashnikov, Y. Ma, S. Mehrotra, and R. Hariharan. Index for fast retrieval of uncertain spatial point data. *Proceedings 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2008.
- [33] R. E. Kalman. A new approach to linear filtering and prediction problem. *Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.
- [34] P. Kamousi, T. M. Chan, and S. Suri. The stochastic closest pair problem and nearest neighbor search. *Proceedings 12th Algorithms and Data Structure Symposium*, vol. LNCS 6844, pp. 548–559, 2011.
- [35] P. Kamousi, T. M. Chan, and S. Suri. Stochastic minimum spanning trees in euclidean spaces. *Proceedings of the 27th Symposium on Computational Geometry*, pp. 65–74, 2011.
- [36] J. Kiefer and J. Wolfowitz. On the deviations of the empirical distribution function of vector chance variables. *Transactions of the American Mathematical Society* 87:173–186, 1958.
- [37] M. van Kreveld and M. Löffler. Largest bounding box, smallest diameter, and related problems on imprecise points. *Computational Geometry Theory and Applications*, 2009.
- [38] H. Kruger. Basic measures for imprecise point sets in \mathbb{R}^d . Master’s thesis, Utrecht University, 2008.
- [39] Y. Li, P. M. Long, and A. Srinivasan. Improved bounds on the samples complexity of learning. *J. Comp. and Sys. Sci.* 62:516–527, 2001.
- [40] T. M. Lillesand, R. W. Kiefer, and J. W. Chipman. *Remote Sensing and Image Interpretation*. John Wiley & Sons, 2004.

- [41] M. Löffler and J. Snoeyink. Delaunay triangulations of imprecise points in linear time after preprocessing. *SoCG*, pp. 298–304, 2008, <http://dx.doi.org/10.1145/1377676.1377727>.
- [42] S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. Tech. rep., arXiv:1203.5747, 2012.
- [43] J. Matousek. Approximations and optimal geometric divide-and-conquer. *Proceedings ACM Symposium on Theory of Computation*, pp. 505–511, 1991.
- [44] J. Matousek. *Geometric Discrepancy; An Illustrated Guide*. Springer, 1999.
- [45] J. Matousek, E. Welzl, and L. Wernisch. Discrepancy and approximations for bounded vc-dimension. *Combinatorica* 13(4):455–466, 1993.
- [46] R. van der Merwe, A. Doucet, N. de Freitas, and E. Wan. The unscented particle filter. *Advances in Neural Information Processing Systems*, vol. 8, pp. 351–357, 2000.
- [47] T. Nagai and N. Tokura. Tight error bounds of geometric problems on convex objects with imprecise coordinates. *Proceedings Japanese Conference on Discrete and Computational Geometry*, 2000.
- [48] Y. Ostrovsky-Berman and L. Jostkowicz. Uncertainty envelopes. *Abstracts 21st European Workshop on Comput. Geom.*, pp. 175–178, 2005.
- [49] J. M. Phillips. Algorithms for ε -approximations of terrains. *Proceedings International Conference on Automata, Languages and Programming*, 2008.
- [50] J. M. Phillips. *Small and Stable Descriptors of Distributions for Geometric Statistical Problems*. Ph.D. thesis, Duke University, 2009.
- [51] S. Potluri, A. K. Yan, J. J. Chou, B. R. Donald, and C. Baily-Kellogg. Structure determination of symmetric homo-oligomers by complete search of symmetry configuration space, using nmr restraints and van der Waals packing. *Proteins* 65:203–219, 2006.
- [52] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. *Proceedings Symposium on Theoretical Aspects of Computer Science*, 1992.
- [53] S. Shekhar and S. Chawla. *Spatial Databases: A Tour*. Pearsons, 2001.
- [54] Y. Tao, R. Cheng, X. Xiao, W. K. Ngai, B. Kao, and S. Prabhakar. Indexing multi-dimensional uncertain data with arbitrary probability density functions. *Proceedings Conference on Very Large Data Bases*, 2005.
- [55] S. Thrun. Robotic mapping: A survey. *Exploring Artificial Intelligence in the New Millenium*, 2002.
- [56] L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8:410–421, 1979.
- [57] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications* 16:264–280, 1971.
- [58] H. Yu, P. K. Agarwal, R. Poreddy, and K. R. Varadarajan. Practical methods for shape fitting and kinetic data structures using coresets. *Proceedings ACM Symposium on Computational Geometry*, 2004.
- [59] Y. Zou and K. Chakrabarty. Uncertainty-aware and coverage-oriented deployment of sensor networks. *Journal of Parallel and Distributed Computing*, 2004.

A ε -Samples of Distributions

In this section we explore conditions for continuous distributions such that they can be approximated with bounded error by discrete distributions (point sets). We state specific results for multi-variate normal distributions.

We say a subset $W \subset \mathbb{R}^d$ is *polygonal approximable* if there exists a polygonal shape $S \subset \mathbb{R}^d$ with m facets such that $\phi(W \setminus S) + \phi(S \setminus W) \leq \varepsilon\phi(W)$ for any $\varepsilon > 0$. Usually, m is dependent on ε , for instance for a d -variate normal distribution $m = O((1/\varepsilon^{d+1}) \log(1/\varepsilon))$ [49, 50]. In turn, such a polygonal shape S describes a continuous point set where (S, \mathcal{A}) can be given an ε -sample Q using $O((1/\varepsilon^2) \log(1/\varepsilon))$ points if (S, \mathcal{A}) has bounded VC-dimension [44] or using $O((1/\varepsilon) \log^{2k}(1/\varepsilon))$ points if \mathcal{A} is defined by a constant k number of directions [49]. For instance, where $\mathcal{A} = \mathcal{B}$ is the set of all balls then the first case applies, and when $\mathcal{A} = \mathcal{R}_2$ is the set of all axis-aligned rectangles then either case applies.

A shape $W \subset \mathbb{R}^{d+1}$ may describe a distribution $\mu : \mathbb{R}^d \rightarrow [0, 1]$. For instance for a range space (μ, \mathcal{B}) , then the range space of the associated shape W_μ is $(W_\mu, \mathcal{B} \times \mathbb{R})$ where $\mathcal{B} \times \mathbb{R}$ describes balls in \mathbb{R}^d for the first d coordinates and any points in the $(d+1)$ th coordinate.

The general scheme to create an ε -sample for (S, \mathcal{A}) , where $S \in \mathbb{R}^d$ is a polygonal shape, is to use a lattice Λ of points. A *lattice* Λ in \mathbb{R}^d is an infinite set of points defined such that for d vectors $\{v_1, \dots, v_d\}$ that form a basis, for any point $p \in \Lambda$, $p + v_i$ and $p - v_i$ are also in Λ for any $i \in [1, d]$. We first create a discrete $(\varepsilon/2)$ -sample $M = \Lambda \cap S$ of (S, \mathcal{A}) and then create an $(\varepsilon/2)$ -sample Q of (M, \mathcal{A}) using standard techniques [13, 49]. Then Q is an ε -sample of (S, \mathcal{A}) . For a shape S with m $(d-1)$ -faces on its boundary, any subset $A' \subset \mathbb{R}^d$ that is described by a subset from (S, \mathcal{A}) is an intersection $A' = A \cap S$ for some $A \in \mathcal{A}$. Since S has m $(d-1)$ -dimensional faces, we can bound the VC-dimension of (S, \mathcal{A}) as $\nu = O((m + \nu_{\mathcal{A}}) \log(m + \nu_{\mathcal{A}}))$ where $\nu_{\mathcal{A}}$ is the VC-dimension of $(\mathbb{R}^d, \mathcal{A})$. Finally the set $M = S \cap \Lambda$ is determined by choosing an arbitrary initial origin point in Λ and then uniformly scaling all vectors $\{v_1, \dots, v_d\}$ until $|M| = \Theta((\nu/\varepsilon^2) \log(\nu/\varepsilon))$ [44]. This construction follows a less general but smaller construction in Phillips [49].

It follows that we can create such an ε -sample K of (S, \mathcal{A}) of size $|M|$ in time $O(|M|m \log |M|)$ by starting with a scaling of the lattice so a constant number of points are in S and then doubling the scale until we get to within a factor of d of $|M|$. If there are n points inside S , it takes $O(nm)$ time to count them. We can then take another ε -sample of (K, \mathcal{A}) of size $O((\nu_{\mathcal{A}}/\varepsilon^2) \log(\nu_{\mathcal{A}}/\varepsilon))$ in time $O(\nu_{\mathcal{A}}^{3\nu_{\mathcal{A}}} |M| ((1/\varepsilon^2) \log(\nu_{\mathcal{A}}/\varepsilon))^{\nu_{\mathcal{A}}})$.

Theorem A.1. *For a polygonal shape $S \subset \mathbb{R}^d$ with m (constant size) facets, we can construct an ε -sample for (S, \mathcal{A}) of size $O((\nu/\varepsilon^2) \log(\nu/\varepsilon))$ in time $O(m(\nu/\varepsilon^2) \log^2(\nu/\varepsilon))$, where (S, \mathcal{A}) has VC-dimension $\nu_{\mathcal{A}}$ and $\nu = O((\nu_{\mathcal{A}} + m) \log(\nu_{\mathcal{A}} + m))$.*

This can be reduced to size $O((\nu_{\mathcal{A}}/\varepsilon^2) \log(\nu_{\mathcal{A}}/\varepsilon))$ in time $O((1/\varepsilon^{2\nu_{\mathcal{A}}+2})(m+\nu_{\mathcal{A}}) \log^{\nu_{\mathcal{A}}}((m+\nu_{\mathcal{A}})/\varepsilon) \log(m/\varepsilon))$.

We can consider the specific case of when $W \subset \mathbb{R}^3$ is a d -variate normal distribution $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$. Then $m = O((1/\varepsilon^d) \log(1/\varepsilon))$ and $|M| = O((m/\varepsilon^2) \log m \log(m/\varepsilon)) = O((1/\varepsilon^{d+2}) \log^3(1/\varepsilon))$.

Corollary A.2. *Let $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$ be a d -variate normal distribution with constant standard deviation. We can construct an ε -sample of (μ, \mathcal{A}) with VC-dimension $\nu_{\mathcal{A}} \geq 2$ (and where $d \leq \nu_{\mathcal{A}} \leq 1/\varepsilon$) of size $O((\nu_{\mathcal{A}}/\varepsilon^2) \log(\nu_{\mathcal{A}}/\varepsilon))$ in time $O((1/\varepsilon^{2\nu_{\mathcal{A}}+d+2}) \log^{\nu_{\mathcal{A}}+1}(1/\varepsilon))$.*

For convenience, we restate a tighter, but less general theorem from Phillips, here slightly generalized.

Theorem A.3 ([49, 50]). *Let $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$ be a d -variate normal distribution with constant standard deviation. Let (μ, \mathcal{Q}_k) be a range space where the ranges are defined as the intersection of k slabs with fixed normal directions. We can construct an ε -sample of (μ, \mathcal{Q}_k) of size $O((1/\varepsilon) \log^{2k}(1/\varepsilon))$ in time $O((1/\varepsilon^{d+4}) \log^{6k+3}(1/\varepsilon))$.*

A.1 Avoiding Degeneracy

An important part of the above construction is the arbitrary choice of the origin points of the lattice Λ . This allows us to arbitrarily shift the lattice defining M and thus the set Q . In Section 4.1 we need to construct n

ε -samples $\{Q_1, \dots, Q_n\}$ for n range spaces $\{(S_1, \mathcal{A}), \dots, (S_n, \mathcal{A})\}$. In Algorithm 4.1 we examine sets of $\nu_{\mathcal{A}}$ points, each from separate ε -samples that define a minimal shape $A \in \mathcal{A}$. It is important that we do not have two such (possibly not disjoint) sets of $\nu_{\mathcal{A}}$ points that define the same minimal shape $A \in \mathcal{A}$. (Note, this does not include cases where say two points are antipodal on a disk and any other point in the disk added to a set of $\nu_{\mathcal{A}} = 3$ points forms such a set; it refers to cases where say four points lie (degenerately) on the boundary of a disc.) We can guarantee this by enforcing a property on all pairs of origin points p and q for (S_i, \mathcal{A}) and (S_j, \mathcal{A}) . For the purpose of construction, it is easiest to consider only the l th coordinates p_l and q_l for any pair of origin points or lattice vectors (where the same lattice vectors are used for each lattice). We enforce a specific property on every such pair p_l and q_l , for all l and all distributions and lattice vectors.

First, consider the case where $\mathcal{A} = \mathcal{R}_d$ describes axis-aligned bounding boxes. It is easy to see that if for all pairs p_l and q_l that $(p_l - q_l)$ is irrational, then we cannot have $> 2d$ points on the boundary of an axis-aligned bounding box, hence the desired property is satisfied.

Now consider the more complicated case where $\mathcal{A} = \mathcal{B}$ describes smallest enclosing balls. There is a polynomial of degree 2 that describes the boundary of the ball, so we can enforce that for all pairs p_l and q_l that $(p_l - q_l)$ is of the form $c_1(r_{p_l})^{1/3} + c_2(r_{q_l})^{1/3}$ where c_1 and c_2 are rational coefficients and r_{p_l} and r_{q_l} are distinct integers that are not multiple of cubes. Now if $\nu = d + 1$ such points satisfy (and in fact define) the equation of the boundary of a ball, then no $(d + 2)$ th point which has this property with respect to the first $d + 1$ can also satisfy this equation.

More generally, if \mathcal{A} can be described with a polynomial of degree p with ν variables, then enforce that every pair of coordinates are the sum of $(p + 1)$ -roots. This ensures that no $\nu + 1$ points can satisfy the equation, and the undesired situation cannot occur.

B Computing Other Measures on Uncertain Points

We generalize this machinery to other LP-type problems f defined on a set of points in \mathbb{R}^d and with constant combinatorial dimension. Although, in some cases (like smallest axis-aligned bounding box by perimeter) we are able to show that $(\mathbb{R}^d, \mathcal{A}_{f,n})$ has constant VC-dimension, for other cases (like radius of the smallest enclosing disk) we cannot and need to first decompose each range $A(Z, w) \in \mathcal{A}_{f,n}$ into a set of disjoint “wedges” from a family of ranges \mathcal{W}_f .

To simplify the already large polynomial runtimes below we replace the runtime bound in Theorem C.2 (below, which extends Theorem 4.2 to deal with large input sizes) with $O((nk)^{\beta+1} \log^4(nk))$.

Lemma B.1. *If the disjoint union of m shapes from \mathcal{W}_f can form any shape from $\mathcal{A}_{f,n}$, then an (ε/m) -sample of (M, \mathcal{W}_f) is an ε -sample of $(M, \mathcal{A}_{f,n})$.*

Proof. For any shape $A \in \mathcal{A}_{f,n}$ we can create a set of m shapes $\{W_1, \dots, W_n\} \subset \mathcal{W}_f$ whose disjoint union is A . Since each range of \mathcal{W}_f may have error ε/m , their union has error at most ε . \square

We study several example cases for which we can deterministically compute ε -quantizations. For each case we show an example element of $\mathcal{A}_{f,n}$ on an example of 7 points.

To facilitate the analysis, we define the notion of shatter dimension, which is similar to VC-dimension. The shatter function $\pi_T(m)$ of a range space $T = (Y, \mathcal{A})$ is the maximum number of sets in T where $|Y| = m$. The *shatter dimension* σ_T of a range space $T = (Y, \mathcal{A})$ is the minimum value such that $\pi_T(m) = O(m^{\sigma_T})$. If a range $A \in \mathcal{A}$ is defined by k points, then $k \geq \sigma_T$. It can be shown [29] that $\sigma_T \leq \nu_T$ and $\nu_T = O(\sigma_T \log \sigma_T)$. And, in general, the basis size of the related LP-type problem is bounded $\beta \leq \sigma_T$.

Directional Width. We first consider the problem of finding the width along a particular direction u (dwid). Given a point set P , $f(P)$ is the width of the minimum slab containing P , as in Figure 11(a). This can be thought of as a one-dimensional problem by projecting all points P using the operation $\langle \cdot, u \rangle$. The directional width is then just the difference between the largest point and the smallest point. As such, the VC-dimension of $(\mathbb{R}^d, \mathcal{A}_f)$ is 2. Furthermore, $\mathcal{A}_{f,n} = \mathcal{A}_f$ in this case, so $(\mathbb{R}^d, \mathcal{A}_{f,n})$ also has VC-dimension 2. For $1 \leq i \leq n$, we can then create an (ε/n) -sample Q_i of $(\mu_i, \mathcal{A}_{f,n})$ of size $k = O(n/\varepsilon)$ in $O((n/\varepsilon) \log(n/\varepsilon))$ time given




Fig. 11: (a) Directional (vertical) width. (b) Axis-aligned bounding box, measured by area. The curves are hyperbola parts.

basic knowledge of the distribution μ_i . We can then apply Theorem 4.2 to build an ε -quantization in time $O((kn)^{2d+1} \log^4(n^2/\varepsilon)) = O(n^6/\varepsilon^3 \log^4(n/\varepsilon))$ for the *dwid* case.

We can actually evaluate $\tilde{f}(Q, r)$ for all values of r faster using a series of sweep lines. Each of the $O(n^4/\varepsilon^2)$ potential bases are defined by a left and right end point. Each of the $O(n^2/\varepsilon)$ points in $\bigcup_i Q_i$ could be a left or right end point. We only need to find the $O(1/\varepsilon)$ widths (defined by pairs of end points) that wind up in the final ε -quantization. Using a Frederickson and Johnson approach [23], we can search for each width in $O(\log(n/\varepsilon))$ steps. At each step we are given a width ω and need to decide what fraction of supports have width at most ω . We can scan from each of the possible left end points and count the number of supports that have width at most ω . For each ω , this can be performed in $O(n^2/\varepsilon)$ time with a pair of simultaneous sweep lines. The total runtime is $O(1/\varepsilon) \cdot O(\log(n/\varepsilon)) \cdot O(n^2/\varepsilon) = O((n^2/\varepsilon^2) \log(n/\varepsilon))$.

Theorem B.2. *We can create an ε -quantization of size $O(1/\varepsilon)$ for the *dwid* problem in $O((n^2/\varepsilon^2) \log(n/\varepsilon))$ time.*

Axis-aligned bounding box. We now consider the set of problems related to axis-aligned bounding boxes in \mathbb{R}^d . For a point set P , we minimize $f(P)$, which either represents the d -dimensional volume of $S(P)$ (the *aabbv* case — minimizes the area in \mathbb{R}^2) or the $(d-1)$ -dimensional volume of the boundary of $S(P)$ (the *aabbp* case — minimizes the perimeter in \mathbb{R}^2). Figures 9(a) and 11(b) show two examples of elements of $\mathcal{A}_{f,n}$ for the *aabbp* case and the *aabbv* case in \mathbb{R}^2 . For both $(\mathbb{R}^2, \mathcal{A}_{f,n})$ has a shatter dimension of 4 because the shape is determined by the x -coordinates of 2 points and the y -coordinates of 2 points. This generalizes to a shatter dimension of $2d$ for $(\mathbb{R}^d, \mathcal{A}_{f,n})$, and hence a VC-dimension of $O(d \log d)$. The smaller VC-dimension in the *aabbp* case discussed in detail above can be extended to higher dimensions.

Hence, for $1 \leq i \leq n$, for both cases we can create an (ε/n) -sample Q_i of $(\mu_i, \mathcal{A}_{f,n})$, each of size $k = O((n^2/\varepsilon^2) \log(n/\varepsilon))$ in total time $O(((n/\varepsilon) \log(n/\varepsilon))^{O(d \log d)})$ via Corollary A.2. In \mathbb{R}^d , we can construct the ε -quantization in $O((kn)^{2d+1} \log^4(nk)) = O((n^{6d+3}/\varepsilon^{4d+2}) \log^{2d+1}(n/\varepsilon))$ time via Theorem 4.2.

Theorem B.3. *We can create an ε -quantization of size $O(1/\varepsilon)$ for the *aabbp* or *aabbv* problem on n d -variate normal distributions in $O(((n/\varepsilon) \log(n/\varepsilon))^{O(d \log d)})$ time.*

Smallest enclosing ball. Figure 12 shows example elements of $\mathcal{A}_{f,n}$ for smallest enclosing ball, for metrics L_∞ (the *seb_∞* case) and L_1 (the *seb₁* case) in \mathbb{R}^2 . An example element of $\mathcal{A}_{f,n}$ for smallest enclosing ball for the L_2 metric (the *seb₂* case) was shown in Figure 9(b). For *seb_∞* and *seb₁*, $(\mathbb{R}^d, \mathcal{A}_{f,n})$ has VC-dimension $2d$ because the shapes are defined by the intersection of halfspaces from d predefined normal directions. For *seb₁* and *seb_∞*, we can create n (ε/n) -samples Q_i of each $(\mu_i, \mathcal{A}_{f,n})$ of size $k = O((n/\varepsilon) \log^{2d}(n/\varepsilon))$ in total time $O(n(n/\varepsilon)^{d+4} \log^{6k+3}(n/\varepsilon))$ via Theorem A.3. We can then create an ε -quantization in $O((nk)^{2d+1} \log^4(nk)) = O((n^{4d+2}/\varepsilon^{2d+1}) \log^{4d^2+2d+4}(n/\varepsilon))$ time via Theorem 4.2.




Fig. 12: (a) Smallest enclosing ball, L_∞ metric. (b) Smallest enclosing ball, L_1 metric.

Theorem B.4. *We can create an ε -quantization of size $O(1/\varepsilon)$ for the seb_1 or seb_∞ problem on n d -variate normal distributions in $O((n^{4d+2}/\varepsilon^{2d+1}) \log^{4d^2+2d+4}(n/\varepsilon))$ time.*

For the seb_2 case in \mathbb{R}^2 , $(\mathbb{R}^2, \mathcal{A}_{f,n})$ has infinite VC-dimension, but $(\mathbb{R}^2, \mathcal{W}_f)$ has VC-dimension at most 9 because it is the intersection of 2 halfspaces and one disc. Any shape $A(T, w) \in \mathcal{A}_{f,n}$ can be formed from the disjoint union of $2n$ wedges. Choosing a point in the convex hull of T as the vertex of the wedges will ensure that each wedge is completely inside the ball that defines part of its boundary. Thus, in \mathbb{R}^2 the n (ε/n) -samples of each $(\mu_{p_i}, \mathcal{A}_{f,n})$ are of size $\lambda_f(n, \varepsilon) = O((n^4/\varepsilon^2) \log(n/\varepsilon))$ and can all be calculated in total time $O((n^5/\varepsilon^2) \log^2(n/\varepsilon))$. And then the ε -quantization can be calculated in $O(n^{56/3}/\varepsilon^{22/3} \log^{11/3}(n/\varepsilon))$ time by assuming general position of all Q_i and then using range counting data structures. We conjecture this technique can be extended to \mathbb{R}^d , but we cannot figure how to decompose a shape $A \in \mathcal{A}_{f,n}$ into a polynomial number of wedges with constant VC-dimension.

Tab. 1: ε -Samples for Summarizing Shape Family $\mathcal{A}_{f,n}$.

case	k	ν_T	$(nk)^{\nu_T}$	$\text{RC}(k, \mathcal{A}_f)$	runtime
dwid	$O(n/\varepsilon)$	2	$O(n^4/\varepsilon^2)$	$\tilde{O}(1)$	$\tilde{O}(n^5/\varepsilon^2)$
aabbbp	$\tilde{O}(n^2/\varepsilon^2)$	$2d$	$\tilde{O}(n^{6d}/\varepsilon^{4d})$	$\tilde{O}(1)$	$\tilde{O}(n^{6d+1}/\varepsilon^{4d})$
aabbbv	$\tilde{O}(n^2/\varepsilon^2)$	$2d$	$\tilde{O}(n^{6d}/\varepsilon^{4d})$	$\tilde{O}(1)$	$\tilde{O}(n^{6d+1}/\varepsilon^{4d})$
seb_∞	$\tilde{O}(n/\varepsilon)$	$d+1$	$\tilde{O}(n^{2d+2}/\varepsilon^{d+1})$	$\tilde{O}(1)$	$\tilde{O}(n^{2d+3}/\varepsilon^{d+1})$
seb_1	$\tilde{O}(n/\varepsilon)$	$d+1$	$\tilde{O}(n^{2d+2}/\varepsilon^{d+1})$	$\tilde{O}(1)$	$\tilde{O}(n^{2d+3}/\varepsilon^{d+1})$
$\text{seb}_2 \in \mathbb{R}^2$	$\tilde{O}(n^4/\varepsilon^2)$	3	$\tilde{O}(n^{15}/\varepsilon^6)$	$\tilde{O}(n^{8/3}/\varepsilon^{4/3})$	$\tilde{O}(n^{56/3}/\varepsilon^{22/3})$
diam $\in \mathbb{R}^2$	$\tilde{O}(n^4/\varepsilon^2)$	n	$\tilde{O}(n^5/\varepsilon^2)^n$	$\tilde{O}(n^4/\varepsilon^2)$	$\tilde{O}((n^5/\varepsilon^2)^{n+1})$

$\tilde{O}(f(n, \varepsilon))$ ignores poly-logarithmic factors $(\log(n/\varepsilon))^{O(\text{poly}(d))}$.

C Algorithm 4.1 in RAM model

In this section we will analyze Algorithm 4.1 without the assumption that the integer k^n can be stored in $O(1)$ words. To simplify the results, we assume a RAM model where a word size contains $O(\log(nk))$ bits, each weight $w(q_{i,j})$ can be stored in one word, and the weight of any basis $f(B)$, the product of n $O(1)$ word weights, can thus be stored in $O(n)$ words. The following lemma describes the main result we will need relating to large numbers.

Lemma C.1. *In the RAM model where a word size has b bits we can calculate the product of n numbers where each is described by $O(b)$ bits in $(nb \log^2 n) 2^{O(\log^* n)}$ time.*

Proof. Using Fürer's recent result [24] we can multiply two m -bit numbers in $m \log m 2^{O(\log^* m)}$ bit operations.

The product of n m -bit numbers has $O(\log((2^m)^n)) = O(nm)$ bits, and can be accomplished with $n - 1$ pairwise multiplications.

We can calculate this product efficiently from the bottom up, starting with $n/2$ multiplications of two $m = O(b)$ bit numbers. Then we perform $n/4$ multiplications of two $2m$ bit numbers, and so on. Since each operation on $O(b)$ bits takes $O(1)$ time in our model, Fürer's result clearly upper bounds the RAM result. The total cost of this can be written as

$$\sum_{i=0}^{\log n} \frac{n}{2^{i-1}} (2^i m) \log(2^i m) 2^{O(\log^*(2^i m))} = 2nb \sum_{i=0}^{\log n} (i + \log b) 2^{O(\log^*(2^i \log b))} = nb \log n \log(nb) 2^{O(\log^*(nb))}.$$

Since we assume $b = O(\log n)$ we can simplify this bound to $(nb \log^2 n) 2^{O(\log^* n)}$. □

This implies that Algorithm 4.1 takes $O((nk)^{\beta+1}) + ((nk)^\beta n \log(nk) \log^2 n) 2^{O(\log^* n)}$ time where each w_i can be described in $O(b) = O(\log(nk))$ bits. This dominates the single division and all other operations. Now we can rewrite Theorem 4.2 without the restriction that k^n can be stored in $O(1)$ words.

Theorem C.2. *Given a set \mathcal{Q} of n indecisive point sets of size k each, and given an LP-type problem $f : \mathcal{Q} \rightarrow \mathbb{R}$ with combinatorial dimension β , we can create the distribution of f over \mathcal{Q} in $O((nk)^{\beta+1}) + ((nk)^\beta n \log(nk) \log^2 n) 2^{O(\log^* n)}$ time. The size of the distribution is $O(n(nk)^\beta)$.*