

ε -Kernel Coresets for Stochastic Points

Lingxiao Huang^{*1}, Jian Li², Jeff M. Phillips^{†3}, and Haitao Wang^{‡4}

1 Tsinghua University, Beijing, China

2 Tsinghua University, Beijing, China

lijian83@mail.tsinghua.edu.cn

3 University of Utah, Salt Lake City, UT, USA

jeffp@cs.utah.edu

4 Utah State University, Logan, UT, USA

haitao.wang@usu.edu

Abstract

With the dramatic growth in the number of application domains that generate probabilistic, noisy and uncertain data, there has been an increasing interest in designing algorithms for geometric or combinatorial optimization problems over such data. In this paper, we initiate the study of constructing ε -kernel coresets for uncertain points. We consider uncertainty in the existential model where each point's location is fixed but only occurs with a certain probability, and the locational model where each point has a probability distribution describing its location. An ε -kernel coreset approximates the width of a point set in any direction. We consider approximating the expected width (an ε -EXP-KERNEL), as well as the probability distribution on the width (an (ε, τ) -QUANT-KERNEL) for any direction. We show that there exists a set of $O(\varepsilon^{-(d-1)/2})$ deterministic points which approximate the expected width under the existential and locational models, and we provide efficient algorithms for constructing such coresets. We show, however, it is not always possible to find a subset of the original uncertain points which provides such an approximation. However, if the existential probability of each point is lower bounded by a constant, an ε -EXP-KERNEL is still possible. We also provide efficient algorithms for construct an (ε, τ) -QUANT-KERNEL coreset in nearly linear time. Our techniques utilize or connect to several important notions in probability and geometry, such as Kolmogorov distances, VC uniform convergence and Tukey depth, and may be useful in other geometric optimization problem in stochastic settings. Finally, combining with known techniques, we show a few applications to approximating the extent of uncertain functions, maintaining extent measures for stochastic moving points and some shape fitting problems under uncertainty.

1998 ACM Subject Classification B.2.4 Algorithms, F.2.2 Geometrical problems and computations

Keywords and phrases ε -kernel, coreset, stochastic point, shape fitting

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.50

1 Introduction

Uncertain Data Models: The wide deployment of sensor monitoring infrastructure and increasing prevalence of technologies such as data integration and cleaning [15, 18] have

* LXH and JL are supported in part by the National Basic Research Program of China grants 2015CB358700, 2011CBA00300, 2011CBA00301, and the National NSFC grants 61033001, 61361136003.

† JP thanks support from NSF CCF-1350888, IIS-1251019, ACI-1443046, and CNS-1514520.

‡ HW was supported in part by NSF under Grant CCF-1317143.



© Lingxiao Huang;

licensed under Creative Commons License CC-BY

24th Annual European Symposium on Algorithms (ESA 2016).

Editors: Piotr Sankowski and Christos Zaroliagis; Article No. 50; pp. 50:1–50:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

resulted in an abundance of uncertain, noisy and probabilistic data in many scientific and application domains. Managing, analyzing, and solving optimization problems over such data have become an increasingly important issue and have attracted significant attentions from several research communities including theoretical computer science, databases, machine learning and wireless networks.

In this paper, we focus on two stochastic models, the existential model and locational model. Both models have been studied extensively for a variety of geometric or combinatorial problems, such as closest pairs [34], nearest neighbors [5, 34], minimum spanning trees [32, 35], convex hulls [6, 24, 49, 52], maxima [3], perfect matchings [32], clustering [16, 26], minimum enclosing balls [43] and range queries [1, 4, 38].

1. Existential uncertainty model: In this model, there is a set \mathcal{P} of n points in \mathbb{R}^d . Throughout this paper, we assume that the dimension d is a constant. Each point $v \in \mathcal{P}$ is associated with a real number (called *existential probability*) $p_v \in [0, 1]$ which indicates that v is present independently with probability p_v .
2. Locational uncertainty model: There are a set \mathcal{P} of n points and the existence of each point is certain. However, the location of each point $v \in \mathcal{P}$ is a random location in \mathbb{R}^d . We assume the probability distribution is discrete and independent of other points. For a point $v \in \mathcal{P}$ and a location $s \in \mathbb{R}^d$, we use $p_{v,s}$ to denote the probability that the location of point v is s . Let S be the set of all possible locations, and let $|S| = m$ be the number of all such locations.

In the locational uncertainty model, we distinguish the use of the terms “points” and “locations”; a point refers to the object with uncertain locations and a location refers to a point (in the usual sense) in \mathbb{R}^d . We will use capital letters (e.g., P, S, \dots) to denote sets of deterministic points and calligraphic letters ($\mathcal{P}, \mathcal{S}, \dots$) to denote sets of stochastic points.

Coresets: Given a large dataset P and a class C of queries, a coreset S is a dataset of much smaller size such that for every query $r \in C$, the answer $r(S)$ for the small dataset S is close to the answer $r(P)$ for the original large dataset P . Coresets [46] have become more relevant in the era of big data as they can drastically reduce the size of a dataset while guaranteeing that answers for certain queries are provably close. An early notion of a coreset concerns the directional width problem (in which a coreset is called an ϵ -kernel) and several other geometric shape-fitting problems in the seminal paper [7].

We introduce some notations and review the definition of ϵ -kernel. We assume the dimension d is a constant. For a set P of deterministic points, the *support function* $f(P, u)$ is defined to be $f(P, u) = \max_{p \in P} \langle u, p \rangle$ for $u \in \mathbb{R}^d$, where $\langle \cdot, \cdot \rangle$ is the inner product. The *directional width* of P in direction $u \in \mathbb{R}^d$, denoted by $\omega(P, u)$, is defined by $\omega(P, u) = f(P, u) + f(P, -u)$. It is not hard to see that the support function and the directional width only depend on the convex hull of P . A subset $Q \subseteq P$ is called an ϵ -kernel of P if for each direction $u \in \mathbb{R}^d$, $(1 - \epsilon)\omega(P, u) \leq \omega(Q, u) \leq \omega(P, u)$. For any set of n points, there is an ϵ -kernel of size $O(\epsilon^{-(d-1)/2})$ [7, 8], which can be constructed in $O(n + \epsilon^{-(d-3/2)})$ time [13, 53].

1.1 Problem Formulations

We focus on constructing ϵ -kernel coresets when the input data is uncertain. These results not only provide an understanding of how to compactly represent an approximate convex hull under uncertainty, but can lead to solutions to a variety of other shape-fitting problems.

ε -Kernels for expectations. Suppose \mathcal{P} is a set of stochastic points (in either the existential or locational uncertainty model). Define the *expected* directional width of \mathcal{P} in direction u to be $\omega(\mathcal{P}, u) = \mathbb{E}_{P \sim \mathcal{P}}[\omega(P, u)]$, where $P \sim \mathcal{P}$ means that P is a (random) realization of \mathcal{P} .

► **Definition 1.** For a constant $\varepsilon > 0$, a set S of (deterministic or stochastic) points in \mathbb{R}^d is called an ε -EXP-KERNEL of \mathcal{P} , if for all directions $u \in \mathbb{R}^d$,

$$(1 - \varepsilon)\omega(\mathcal{P}, u) \leq \omega(S, u) \leq \omega(\mathcal{P}, u).$$

Recall in the deterministic setting, we require that the ε -kernel S be a subset of the original point set (we call this *the subset constraint*). It is important to consider the subset constraint since it can reveal how concisely arbitrary uncertain point sets can be represented with just a few uncertain points (size depending only on ε). For ε -kernels on deterministic points, $\Omega(\varepsilon^{-(d-1)/2})$ points may be required and can always be found under the subset constraint [7, 8]. However, in the stochastic setting, we will show this is no longer true. Yet, coresets without the subset constraint, in fact made of deterministic points, can sometimes be obtained when no coreset with the subset constraint is possible.

ε -Kernels for probability distributions. Sometimes it is useful to obtain more than just the expected value (say of the width) on a query; rather one may want (an approximation of) a representation of the full probability distribution that the query can take.

► **Definition 2.** For a constant $\varepsilon, \tau > 0$, a set \mathcal{S} of stochastic points in \mathbb{R}^d is called an (ε, τ) -QUANT-KERNEL of \mathcal{P} , if for all directions u and all $x \geq 0$,

$$\begin{aligned} \Pr_{P \sim \mathcal{P}}[\omega(P, u) \leq (1 - \varepsilon)x] - \tau &\leq \Pr_{S \sim \mathcal{S}}[\omega(S, u) \leq x] \\ &\leq \Pr_{P \sim \mathcal{P}}[\omega(P, u) \leq (1 + \varepsilon)x] + \tau. \end{aligned} \quad (1)$$

In the above definition, we do not require the points in \mathcal{S} are independent. So when they are correlated, we will specify the distribution of \mathcal{S} . If all points in \mathcal{P} are deterministic and $\tau < 0.5$, the above definition essentially boils down to requiring $(1 - \varepsilon)\omega(\mathcal{P}, u) \leq \omega(\mathcal{S}, u) \leq (1 + \varepsilon)\omega(\mathcal{P}, u)$. Assuming the coordinates of the input points are bounded, an (ε, τ) -QUANT-KERNEL ensures that for any choice of u , the cumulative distribution function of $\omega(\mathcal{S}, u)$ is within a distance ε under the Lévy metric, to that of $\omega(\mathcal{P}, u)$.

ε -Kernels for expected fractional powers. Sometimes, the notion ε -EXP-KERNEL is not powerful enough for certain shape fitting problems (e.g., the minimum enclosing cylinder problem and the minimum spherical shell problem) in the stochastic setting. The main reason is the appearance of the l_2 -norm in the objective function. So we need to be able to handle the fractional powers in the objective function. For a set P of points in \mathbb{R}^d , the polar set of P is defined to be $P^\star = \{u \in \mathbb{R}^d \mid \langle u, v \rangle \geq 0, \forall v \in P\}$. Let r be a positive integer. Given a set P of points in \mathbb{R}^d and $u \in P^\star$, we define a function

$$T_r(P, u) = \max_{v \in P} \langle u, v \rangle^{1/r} - \min_{v \in P} \langle u, v \rangle^{1/r}.$$

We only care about the directions in \mathcal{P}^\star (i.e., the polar of the points in \mathcal{P}) for which $T_r(P, u), \forall P \sim \mathcal{P}$ is well defined.

► **Definition 3.** For a constant $\varepsilon > 0$, a positive integer r , a set \mathcal{S} of stochastic points in \mathbb{R}^d is called an (ε, r) -FPOW-KERNEL of \mathcal{P} , if for all directions $u \in \mathcal{P}^\star$,

$$(1 - \varepsilon)\mathbb{E}_{P \sim \mathcal{P}}[T_r(P, u)] \leq \mathbb{E}_{S \sim \mathcal{S}}[T_r(S, u)] \leq (1 + \varepsilon)\mathbb{E}_{P \sim \mathcal{P}}[T_r(P, u)].$$

1.2 Our Results

Now, we discuss the main technical results of the paper.

ε -Kernels for expectations. First, we consider ε -EXP-KERNELS under various constraints. Our first main result is that an ε -EXP-KERNEL of size $O(\varepsilon^{-(d-1)/2})$ exists for both existential and locational uncertainty models and can be constructed in nearly linear time.

► **Theorem 4.** *\mathcal{P} is a set of n uncertain points in \mathbb{R}^d (in either locational uncertainty model or existential uncertainty model). There exists an ε -EXP-KERNEL of size $O(\varepsilon^{-(d-1)/2})$ for \mathcal{P} . For existential uncertainty model (locational uncertainty model resp.), such an ε -EXP-KERNEL can be constructed in $O(\varepsilon^{-(d-1)}n \log n)$ time, ($O(\varepsilon^{-(d-1)}m \log m)$ time resp.), where n is the number of points and m is the total number of possible locations.*

The existential result is a simple Minkowski sum argument. We first show that there exists a convex polytope M such that for any direction, the directional width of M is exactly the same as the expected directional width of \mathcal{P} (Lemma 10). This immediately implies the existence of a ε -EXP-KERNEL consisting $O(\varepsilon^{-(d-1)/2})$ deterministic points (using the result in [7]), but without the subset constraint. The Minkowski sum argument seems to suggest that the complexity of M is exponential. However, we show that the complexity of M is in fact polynomial $O(n^{2d-2})$ and we can construct it explicitly in $O(n^{2d-1} \log n)$ time (Theorem 14).

Although the complexity of M is polynomial, we cannot afford to construct it explicitly if we are to construct an ε -EXP-KERNEL in nearly linear time. Thus we construct the ε -EXP-KERNEL without explicitly constructing M . In particular, we show that it is possible to find the extreme vertex of M in a given direction in nearly linear time, by computing the gradient of the support function of M . We also provide quadratic-size data structures that can calculate the exact width $\omega(\mathcal{P}, \cdot)$ in logarithmic time under both models in \mathbb{R}^2 .

We also show that under subset constraint (i.e., the ε -EXP-KERNEL is required to be a subset of the original point set, with the same probability distribution for each chosen point), there is no ε -EXP-KERNEL of sublinear size (Lemma 15). However, if there is a constant lower bound $\beta > 0$ on the existential probabilities (called β -assumption), we can construct an ε -EXP-KERNEL of constant size (Theorem 16).

ε -Kernels for probability distributions. Now, we describe our main results for (ε, τ) -QUANT-KERNELS. We first propose a quite simple but general algorithm for constructing (ε, τ) -QUANT-KERNELS, which achieves the following guarantee.

► **Theorem 5.** *An (ε, τ) -QUANT-KERNEL of size $\tilde{O}(\tau^{-2}\varepsilon^{-3(d-1)/2})$ can be constructed in $\tilde{O}(n\tau^{-2}\varepsilon^{-(d-1)})$ time, under both existential and locational uncertainty models.*

The algorithm is surprisingly simple. Take a certain number N of i.i.d. realizations, compute an ε -kernel for each realization, and then associate each ε -kernel with probability $1/N$ (so the points are not independent). The analysis requires the VC uniform convergence bound for unions of halfspaces. The details can be found in Section 3.1.

For existential uncertainty model, we can improve the size bound as follows.

► **Theorem 6.** *\mathcal{P} is a set of uncertain points in \mathbb{R}^d with existential uncertainty. Let $\lambda = \sum_{v \in \mathcal{P}} (-\ln(1 - p_v))$. There exists an (ε, τ) -QUANT-KERNEL for \mathcal{P} , which consists of a set of independent uncertain points of cardinality $\min\{\tilde{O}(\tau^{-2} \max\{\lambda^2, \lambda^4\}), \tilde{O}(\varepsilon^{-(d-1)}\tau^{-2})\}$. The algorithm for constructing such a coreset runs in $\tilde{O}(n \log^{O(d)} n)$ time.*

We note that another advantage of the improved construction is that the (ε, τ) -QUANT-KERNEL is a set of independent stochastic points (rather than correlated points as in Theorem 5). We achieve the improvement by two algorithms. The first algorithm transforms the Bernoulli distributed variables into Poisson distributed random variables and creates a probability distribution using the parameters of the Poissons, from which we take a number of i.i.d. samples as the coresets. Our analysis leverages the additivity of Poisson distributions and the VC uniform convergence bound (for halfspaces). However, the number of samples required depends on $\lambda(\mathcal{P})$, so the first algorithm only works when $\lambda(\mathcal{P})$ is small. The second algorithm complements the first one by identifying a convex set K that lies in the convex hull of \mathcal{P} with high probability (K exists when $\lambda(\mathcal{P})$ is large) and uses a small size deterministic ε -kernel to approximate K . The points in $\bar{K} = \mathcal{P} \setminus K$ can be approximated using the same sampling algorithm as in the first algorithm and we can show that $\lambda(\bar{K})$ is small, thus requiring only a small number of samples. In the appendix (Section 3.2.3), we show such an (ε, τ) -QUANT-KERNEL can be computed in $O(n \cdot \text{polylog} n)$ time using an iterative sampling algorithm. Our technique has some interesting connections to other important geometric problems (such as the Tukey depth problem) [42], may be interesting in its own right.

ε -Kernels for expected fractional powers. For (ε, r) -FPOW-KERNELS, we provide a linear time algorithm for constructing an (ε, r) -FPOW-KERNEL of size $\tilde{O}(\varepsilon^{-(rd-r+2)})$ in the existential uncertainty model under the β -assumption. The algorithm is almost the same as the construction in Section 3.1 except that some parameters are different.

► **Theorem 7** (Section 4). *An (ε, r) -FPOW-KERNEL of size $\tilde{O}(\varepsilon^{-(rd-r+2)})$ can be constructed in $\tilde{O}(n\varepsilon^{-(rd-r+4)/2})$ time in the existential uncertainty model under the β -assumption.*

Applications to Uncertain Function Approximation and Shape Fitting. Finally, we show that the above results, combined with the duality and linearization arguments [7], can be used to obtain constant size coresets for the function extent problem in the stochastic setting, and to maintain extent measures for stochastic moving points.

Using the above results, we also obtain efficient approximation schemes for various shape-fitting problems in the stochastic setting, such as minimum enclosing ball, minimum spherical shell, minimum enclosing cylinder and minimum cylindrical shell in different stochastic settings. We summarize our application results in the following theorems. The details can be found in Section 5.

► **Theorem 8.** *Suppose \mathcal{P} is a set of n independent stochastic points in \mathbb{R}^d under either existential or locational uncertainty model. There are linear time approximation schemes for the following problems: (1) finding a center point c to minimize $\mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\|^2]$; (2) finding a center point c to minimize $\mathbb{E}[\text{obj}(c)] = \mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\|^2 - \min_{v \in \mathcal{P}} \|v - c\|^2]$. Note that when $d = 2$ the above two problems correspond to minimizing the expected areas of the enclosing ball and the enclosing annulus, respectively.*

Under β -assumption, we can obtain efficient approximation schemes for the following shape fitting problems.

► **Theorem 9.** *Suppose \mathcal{P} is a set of n independent stochastic points in \mathbb{R}^d , each appearing with probability at least β , for some fixed constant $\beta > 0$. There are linear time approximation schemes for minimizing the expected radius (or width) for the minimum spherical shell, minimum enclosing cylinder, minimum cylindrical shell problems over \mathcal{P} .*

1.3 Other Related Work

Besides the stochastic models mentioned above, geometric uncertain data has also been studied in the *imprecise* model [11, 31, 36, 40, 44, 45, 50]. In this model, each point is provided with a region where it might be. This originated with the study of imprecision in data representation [27, 47], and can be used to provide upper and lower bounds on several geometric constructs such as the diameter, convex hull, and flow on terrains [19, 50].

Convex hulls have been studied for uncertain points: upper and lower bounds are provided under the imprecise model [20, 41, 44, 50], distributions of circumference and volume are calculated in the locational model [33, 39], the most likely convex hull is found in the existential model in \mathbb{R}^2 and shown NP-hard for \mathbb{R}^d for $d > 2$ and in the locational model [49], and the probability a query point is inside the convex hull [6, 24, 52]. As far as we know, the expected complexity of the convex hull under uncertain points has not been studied, although it has been studied [28] under other random data models.

There is a large body of literature [46] on constructing coresets for various problems, such as shape fitting [7, 8], shape fitting with outliers [30], clustering [14, 22, 23, 29, 37], integrals [37], matrix approximation and regression [17, 22] and in different settings, such as geometric data streaming [8, 13] and privacy setting [21]. Coresets were constructed for imprecise points [41] to help derive results for approximating convex hulls and a variety of other shape-fitting problems, but because of the difference in models, these approaches do not translate to existential or locational models. In the locational model, coresets are created for range counting queries [1] under the subset constraint, but again these techniques do not translate because ϵ -kernel coresets in general cannot be constructed from a density-preserving subset of the data, as is preserved for the range counting coresets. Also in the locational model (and directly translating to the existential model) Löffler and Phillips [39] show how a large set of uncertain points can be approximated with a set of deterministic point sets, where each certain point set can be an ϵ -kernel. This can provide approximations similar to the (ϵ, τ) -QUANT-KERNEL with space $O(\epsilon^{-(d+3)/2} \log(1/\delta))$ with probability at least $1 - \delta$. However it is not a coreset of the data, and answering width queries requires querying $O(\epsilon^{-2} \log(1/\delta))$ deterministic point sets.

Recently, Munteanu et al. [43] studied the minimum enclosing ball problem over stochastic points, and obtained an efficient approximation scheme. Their algorithm and analysis utilize the results from the deterministic coreset literature [2]. However, they do not directly address the problem of constructing coresets for stochastic points and it is also unclear how to extend their technique to other shape fitting problems, such as minimum spherical shells.

Technically, our (ϵ, τ) -QUANT-KERNEL construction bears some similarity to the coreset by Har-Peled and Wang [30] for handling outliers. From the dual (function extent) perspective, they want to approximate the distance between two level sets in an arrangement of hyperplanes, and (the dual of) \mathcal{H} in Section 3.2.2 also needs to be (approximately) sandwiched by two fractional level sets (our hyperplanes have weights). However, we have an important requirement that the total weight outside $(1 + \epsilon)\mathcal{H}$ must be small, which cannot be addressed by their technique.

2 ϵ -Kernels for Expectations of Width

We first state our results in this section for the existential uncertainty model. All results can be extended to the locational uncertainty model, with slightly different bounds (essentially replacing the number of points n with the number of locations m) or assumptions. We describe the difference for locational model in the full version.

For simplicity of exposition, we assume in this section that all points in \mathcal{P} are in general positions and all p_v s are strictly between 0 and 1. For any $u, v \in \mathbb{R}^d$, we use $\langle u, v \rangle$ to denote the usual inner product $\sum_{i=1}^d u_i v_i$. For ease of notation, we write $v \succ_u w$ as a shorthand notation for $\langle u, v \rangle > \langle u, w \rangle$. For any $u \in \mathbb{R}^d$, the binary relation \succ_u defines a total order of all vertices in \mathcal{P} . (Ties should be broken in an arbitrary but consistent manner.) We call this order the *canonical order of \mathcal{P} with respect to u* . For any two points u and v , we use $d(u, v)$ or $\|v - u\|$ to denote their Euclidean distance. For any two sets of points, A and B , the Minkowski sum of A and B is defined as $A \oplus B := \{a + b \mid a \in A, b \in B\}$. Recall the definitions for a set P of deterministic points and a direction $u \in \mathbb{R}^d$, the support function is $f(P, u) = \max_{p \in P} \langle u, p \rangle$ and the *directional width* is $\omega(P, u) = f(P, u) - f(P, -u)$. The support function and the directional width only depend on the convex hull of P .

► **Lemma 10.** *Consider a set \mathcal{P} of uncertain points in \mathbb{R}^d (in either locational uncertainty model or existential uncertainty model). There exists a set S of deterministic points in \mathbb{R}^d (which may not be a subset of \mathcal{P}) such that $\omega(u, \mathcal{P}) = \omega(u, S)$ for all $u \in \mathbb{R}^d$.*

Proof. By the definition of the expected directional width of \mathcal{P} , we have that $\omega(\mathcal{P}, u) = \mathbb{E}_{P \sim \mathcal{P}}[\omega(P, u)] = \sum_{P \sim \mathcal{P}} \Pr[P] (f(P, u) + f(P, -u))$. Consider the Minkowski sum $M = M(\mathcal{P}) := \sum_{P \sim \mathcal{P}} \Pr[P] \text{ConvH}(P)$, where $\text{ConvH}(P)$ is the convex hull of P (including the interior). It is well known that the Minkowski sum of a set of convex sets is also convex. Moreover, it also holds that for all $u \in \mathbb{R}^d$ (see e.g., [48]) $f(M, u) = \sum_{P \sim \mathcal{P}} \Pr[P] f(P, u)$. Hence, $\omega(\mathcal{P}, u) = \omega(M, u)$ for all $u \in \mathbb{R}^d$. ◀

By the result in [7], for any convex body in \mathbb{R}^d , there exists an ε -kernel of size $O(\varepsilon^{-(d-1)/2})$. Combining with Lemma 10, we can immediately obtain the following corollary.

► **Corollary 11.** *For any $\varepsilon > 0$, there exists an ε -EXP-KERNEL of size $O(\varepsilon^{-(d-1)/2})$.*


Recall that in Lemma 10, the Minkowski sum $M = \sum_{P \sim \mathcal{P}} \Pr[P] \text{ConvH}(P)$. Since M is the Minkowski sum of exponential many convex polytopes, so M is also a convex polytope. At first sight, the complexity of M (i.e., number of vertices) could be exponential. However, as we will show shortly, the complexity of M is in fact polynomial.

We need some notations first. For each pair (r, w) of points in \mathcal{P} consider the hyperplane $H_{r,w}$ that passes through the origin and is orthogonal to the line connecting r and w . We call these $\binom{n}{2}$ hyperplanes *the separating hyperplanes induced by \mathcal{P}* and use Γ to denote the set. Each such hyperplane divides \mathbb{R}^d into 2 halfspaces. For all vectors $u \in \mathbb{R}^d$ in each halfspace, the order of $\langle r, u \rangle$ and $\langle w, u \rangle$ is the same (i.e., we have $r \succ_u w$ in one halfspace and $w \succ_u r$ in the other). Those hyperplanes in Γ pass through the origin and thus partition \mathbb{R}^d into d -dimensional polyhedral cones. We denote this *arrangement* as $\mathbb{A}(\Gamma)$.

Consider an arbitrary cone $C \in \mathbb{A}(\Gamma)$. Let $\text{int } C$ denote the interior of C . We can see that for all vectors $u \in \text{int } C$, the canonical order of \mathcal{P} with respect to u is the same (since all vector $u \in \text{int } C$ lie in the same set of halfspaces). We use $|M|$ to denote the complexity of M , i.e., the number of vertices in $\text{ConvH}(M)$.

► **Lemma 12.** *Assuming the existential model and $p_v \in (0, 1)$ for all $v \in \mathcal{P}$, the complexity of M is the same as the cardinality of $\mathbb{A}(\Gamma)$, i.e., $|M| = |\mathbb{A}(\Gamma)|$. Moreover, each cone $C \in \mathbb{A}(\Gamma)$ corresponds to exactly one vertex v of $\text{ConvH}(M)$ in the following sense: the gradient $\nabla f(M, u) = v$ for all $u \in \text{int } C$ (note that here v should be understood as a vector).*

Proof. (sketch) We have shown that M is a convex polytope. We first note that the support function uniquely defines a convex body (see e.g., [48]). We need the following well known



■ **Figure 1** The figure depicts a pentagon M in \mathbb{R}^2 to illustrate some intuitive facts in convex geometry. (1) The plane can be divided into 5 cones C_1, \dots, C_5 , by 5 angles $\theta_1, \dots, \theta_5$. u_{θ_i} is the unit vector corresponding to angle θ_i . Each cone C_i corresponds to a vertex v_i and for any direction $u \in C_i$, $f(M, u) = \langle u, v_i \rangle$ and the vector $\nabla f(M, u)$ is v_i . (2) Each direction θ_i is perpendicular to an edge of M . $M = \bigcap_{i=1}^5 H_i$ where H_i is the supporting halfplane with normal vector u_{θ_i} .

fact in convex geometry (see e.g., [25]): For any convex polytope M , \mathbb{R}^d can be divided into exactly $|M|$ polyhedral cones (of dimension d , ignoring the boundaries), such that each such cone C_v corresponds to a vertex v of M , and for each vector $u \in C_v$, it holds $f(M, u) = \langle u, v \rangle$ (i.e., the maximum of $f(M, u) = \max_{v' \in M} \langle u, v' \rangle$ is achieved by v for all $u \in C_v$)¹. See Figure 1 for an example in \mathbb{R}^2 . Hence, for each $u \in \text{int } C_v$ the gradient of the support function (as a function of u) is exactly v :

$$\nabla f(M, u) = \left\{ \frac{\partial f(M, u)}{\partial u_j} \right\}_{j \in [d]} = \left\{ \frac{\partial \langle u, v \rangle}{\partial u_j} \right\}_{j \in [d]} = \left\{ \frac{\partial \sum_{j \in [d]} v_j u_j}{\partial u_j} \right\}_{j \in [d]} = v, \quad (2)$$

where u_j is the j th coordinate of u . With a bit abuse of notation, we denote the set of cones defined above by $\mathbb{A}(M)$.

Now, consider a cone $C \in \mathbb{A}(\Gamma)$. We show that for all $u \in \text{int } C$, $\nabla f(M, u)$ is a distinct constant vector independent of u . In fact, we know that $f(M, u) = f(\mathcal{P}, u) = \sum_{v \in \mathcal{P}} \text{Pr}^R(v, u) \langle v, u \rangle$, where $\text{Pr}^R(v, u) = \prod_{v' \succ_u v} (1 - p_{v'}) p_v$. For all $u \in \text{int } C$, the $\text{Pr}^R(v, u)$ value is the same since the value only depends on the canonical order with respect to u , which is the same for all $u \in C$. Hence, we can get that for all $u \in \text{int } C$, $\nabla f(M, u) = \sum_{v \in \mathcal{P}} \text{Pr}^R(v, u) v$, which is a constant independent of u . We can also show that the gradient $\nabla f(M, u)$ must be different for two adjacent cones C_1, C_2 (separated by some hyperplane in Γ) in $\mathbb{A}(\Gamma)$. So $\nabla f(M, u)$ is piecewise constant, with a distinct constant in each cone in $\mathbb{A}(M)$. The same also holds for $\mathbb{A}(\Gamma)$. This is only possible if $\mathbb{A}(\Gamma)$ (thinking as a partition of \mathbb{R}^d) partitions \mathbb{R}^d exactly the same way as $\mathbb{A}(M)$ does. Hence, we have $\mathbb{A}(\Gamma) = \mathbb{A}(M)$ and the lemma follows immediately. ◀

Since $O(n^2)$ hyperplanes passing through the origin can divide \mathbb{R}^d into at most $O(\binom{n^2}{d-1})$ d -dimensional polyhedral cones (see e.g., [9]), we immediately obtain the following corollary.

► **Corollary 13.** *It holds that $|M| \leq O(\binom{n^2}{d-1}) = O(n^{2d-2})$.*

¹ The support function for a polytope is just the upper envelope of a finite set of linear functions, thus a piecewise linear function, and the domain of each piece is a polyhedral cone.

► **Theorem 14.** *In \mathbb{R}^d for constant d , the polytope M which defines $f(\mathcal{P}, u)$ for any direction u can be described with $O(n^{2d-2})$ vertices in \mathbb{R}^d , and can be computed in $O(n^{2d-1} \log n)$ time. In \mathbb{R}^2 , the runtime can be improved to $O(n^2 \log n)$.*

In fact, we can reduce the construction time to nearly linear time, which leads to Theorem 4. The details can be found in the full version.

2.1 ε -exp-kernel Under the Subset Constraint

First, we show that under the subset constraint (i.e., the ε -EXP-KERNEL is required to be a subset of the original point set, with the same probability distribution for each chosen point), there exists no ε -EXP-KERNEL with small size in general.

► **Lemma 15.** *For some constant $\varepsilon > 0$, there exist a set \mathcal{P} of stochastic points such that no $o(n)$ size ε -EXP-KERNEL exists for \mathcal{P} under the subset constraint (for both locational model and existential model).*

In light of the above negative result, we make the following β -assumption: we assume each possible location realizes a point with probability at least β , for a constant $\beta > 0$.

► **Theorem 16.** *Under the β -assumption, in the existential uncertainty model, there is an ε -EXP-KERNEL in \mathbb{R}^d of size $O(\beta^{-(d-1)} \varepsilon^{-(d-1)/2} \log(1/\varepsilon))$ that satisfies the subset constraint.*

3 ε -Kernels for Probability Distributions of Width

Recall \mathcal{S} is an (ε, τ) -QUANT-KERNEL if for all $x \geq 0$, $\Pr_{P \sim \mathcal{P}}[\omega(P, u) \leq (1 - \varepsilon)x] - \tau \leq \Pr_{P \sim \mathcal{S}}[\omega(\mathcal{S}, u) \leq x] \leq \Pr_{P \sim \mathcal{P}}[\omega(P, u) \leq (1 + \varepsilon)x] + \tau$. For ease of notation, we sometimes write $\Pr[\omega(\mathcal{P}, u) \leq t]$ to denote $\Pr_{P \sim \mathcal{P}}[\omega(P, u) \leq t]$, and abbreviate the above as $\Pr[\omega(\mathcal{S}, u) \leq x] \in \Pr[\omega(\mathcal{P}, u) \leq (1 \pm \varepsilon)x] \pm \tau$. We first provide a simple linear time algorithm for constructing an (ε, τ) -QUANT-KERNEL for both existential and locational models, in Section 3.1. The points in the constructed kernel are not independent. Then, for existential models, we provide a nearly linear time (ε, τ) -QUANT-KERNEL construction where all stochastic points in the kernel are independent in Section 3.2.

3.1 A Simple (ε, τ) -quant-kernel Construction

In this section, we show a linear time algorithm for constructing an (ε, τ) -QUANT-KERNEL for any stochastic model if we can sample a realization from the model in linear time (which is true for both locational and existential uncertainty models).

Algorithm: Let $N = O(\tau^{-2} \varepsilon^{-(d-1)} \log(1/\varepsilon))$. We sample N independent realizations from the stochastic model. Let \mathcal{H}_i be the convex hull of the present points in the i th realization. For \mathcal{H}_i , we use the algorithm in [7] to find a deterministic ε -kernel \mathcal{E}_i of size $O(\varepsilon^{-(d-1)/2})$. Our (ε, τ) -QUANT-KERNEL \mathcal{S} is the following simple stochastic model: with probability $1/N$, all points in \mathcal{E}_i are present. Hence, \mathcal{S} consists of $O(\tau^{-2} \varepsilon^{-3(d-1)/2} \log(1/\varepsilon))$ points (two such points either co-exist or are mutually exclusive). Hence, for any direction u , $\Pr[\omega(\mathcal{S}, u) \leq t] = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\omega(\mathcal{E}_i, u) \leq t)$, where $\mathbb{I}(\cdot)$ is the indicator function.

For a realization $P \sim \mathcal{P}$, we use $\mathcal{E}(P)$ to denote the deterministic ε -kernel for P . So, $\mathcal{E}(P)$ is a random set of points, and we can think of $\mathcal{E}_1, \dots, \mathcal{E}_N$ as samples from the random set. We first show that \mathcal{S} is an (ε, τ) -QUANT-KERNEL for $\mathcal{E}(P)$.

► **Lemma 17.** *Let $N = O(\tau^{-2}\varepsilon^{-(d-1)}\log(1/\varepsilon))$. For any $t \geq 0$ and any direction u ,*

$$\Pr[\omega(\mathcal{S}, u) \leq t] \in \Pr_{P \sim \mathcal{P}}[\omega(\mathcal{E}(P), u) \leq t] \pm \tau.$$

Proof (Sketch). Let $L = O(\varepsilon^{-(d-1)/2})$. We first build a mapping g that maps each realization $\mathcal{E}(P)$ to a point in \mathbb{R}^{dL} , as follows: Consider a realization P of \mathcal{P} . Suppose $\mathcal{E}(P) = \{(x_1^1, \dots, x_d^1), \dots, (x_1^L, \dots, x_d^L)\}$ (if $|\mathcal{E}(P)| < L$, we pad it with $(0, \dots, 0)$). We let $g(\mathcal{E}(P)) = (x_1^1, \dots, x_d^1, \dots, x_1^L, \dots, x_d^L) \in \mathbb{R}^{dL}$. For any $t \geq 0$ and any direction $u \in \mathbb{R}^d$, note that $\omega(\mathcal{E}(P), u) \geq t$ holds if and only if there exists some $1 \leq i, j \leq |\mathcal{E}(P)|, i \neq j$ satisfies that $\sum_{k=1}^d (x_k^i - x_k^j)u_k \geq t$, which is equivalent to saying that point $g(\mathcal{E}(P))$ is in the union of the those $O(|\mathcal{E}(P)|^2)$ halfspaces. We can show that the VC dimension of the union of $O(|\mathcal{E}(P)|^2)$ such halfspaces is bounded by $O(\varepsilon^{-(d-1)}\log(1/\varepsilon))$. Then, the lemma follows from the VC uniform convergence theorem [51, 10]. ◀

From the above lemma, it is not difficult to obtain Theorem 5.

3.2 Improved (ε, τ) -quant-kernel for Existential Models

In this section, we show an (ε, τ) -QUANT-KERNEL \mathcal{S} can be constructed in nearly linear time for the existential model, and all points in \mathcal{S} are independent of each other. The size bound $\tilde{O}(\tau^{-2}\varepsilon^{-(d-1)})$ (see Theorem 6) is better than that in Theorem 5 for the general case, and the independence property may be useful in certain applications. Moreover, some of the insights developed in this section may be of independent interest (e.g., the connection to Tukey depth). Due to the independence requirement, the construction is somewhat more involved. For ease of the description, we assume the Euclidean plane first. All results can be easily extended to \mathbb{R}^d . We also assume that all probability values are strictly between 0 and 1 and $0 < \varepsilon, \tau \leq 1/2$ is a fixed constant.

Let $\lambda(\mathcal{P}) = \sum_{v \in \mathcal{P}} (-\ln(1 - p_v))$. In the following, we present two algorithms. The first algorithm works for any $\lambda(\mathcal{P})$ and produces an (ε, τ) -QUANT-KERNEL \mathcal{S} whose size depends on $\lambda(\mathcal{P})$. In Section 3.2.2, we present the second algorithm that only works for $\lambda(\mathcal{P}) \geq 3 \ln(2/\tau)$ but produces an (ε, τ) -QUANT-KERNEL \mathcal{S} with a constant size (the constant only depends on ε, τ and δ). Thus, we can get a constant size (ε, τ) -QUANT-KERNEL by running the first algorithm when $\lambda(\mathcal{P}) \leq 3 \ln(2/\tau)$ and running the second algorithm otherwise.

3.2.1 Algorithm 1: For Any $\lambda(\mathcal{P})$

In this section, we present the first algorithm which works for any $\lambda(\mathcal{P})$. We can think of each point v associated with a Bernoulli random variable X_v that takes value 1 with probability p_v and 0 otherwise. Now, we replace the Bernoulli random variable X_v by a Poisson distributed random variable \tilde{X}_v with parameter $\lambda_v = -\ln(1 - p_v)$ (denoted by $\text{Pois}(\lambda_v)$), i.e., $\Pr[\tilde{X}_v = k] = \frac{1}{k!} \lambda_v^k e^{-\lambda_v}$, for $k = 0, 1, 2, \dots$. Here, $\tilde{X}_v = k$ means that there are k realized points located at the position of v . We call the new instance *the Poissonized instance corresponding to \mathcal{P}* . We can check that $\Pr[\tilde{X}_v = 0] = e^{-\lambda_v} = 1 - p_v = \Pr[X_v = 0]$. Also note that co-locating points does not affect any directional width, so the Poissonized instance is essentially equivalent to the original instance for our problem.

The construction of the (ε, τ) -QUANT-KERNEL \mathcal{S} is as follows: Let \mathfrak{A} be the probability measure over all points in \mathcal{P} defined by $\mathfrak{A}(\{v\}) = \lambda_v/\lambda$ for every $v \in \mathcal{P}$, where $\lambda := \lambda(\mathcal{P}) = \sum_{v \in \mathcal{P}} \lambda_v$. Let τ_1 be a small positive constant to be fixed later. We take $N = \tilde{O}(\tau_1^{-2})$ independent samples from \mathfrak{A} (we allow more than one point to be co-located at the same position), and let \mathfrak{B} be the empirical measure, i.e., each sample point having probability

$1/N$. The coresets \mathcal{S} consists of the N sample points in \mathfrak{B} , each with the same existential probability $1 - \exp(-\lambda/N)$.

► **Theorem 18.** *Let $\tau_1 = O(\frac{\tau}{\max\{\lambda, \lambda^2\}})$ and $N = O(\frac{1}{\tau_1^2} \log \frac{1}{\delta}) = O(\frac{\max\{\lambda^2, \lambda^4\}}{\tau^2} \log \frac{1}{\delta})$. With probability at least $1 - \delta$, for any $t \geq 0$ and any direction u , we have that $\Pr[\omega(\mathcal{S}, u) \leq t] \in \Pr[\omega(\mathcal{P}, u) \leq t] \pm \tau$.*

3.2.2 Algorithm 2: For $\lambda(\mathcal{P}) > 3 \ln(2/\tau)$

In the second algorithm, we assume that $\lambda(\mathcal{P}) = \sum_{v \in \mathcal{P}} \lambda_v > 3 \ln(2/\tau)$. When $\lambda(\mathcal{P})$ is large, we cannot directly use the sampling technique in the previous section since it requires a large number of samples. However, the condition $\lambda(\mathcal{P}) \geq 3 \ln(2/\tau)$ implies there is a nonempty convex region \mathcal{K} inside the convex hull of \mathcal{P} with high probability. Moreover, we can show the sum of λ_v values in $\bar{\mathcal{K}} = \mathbb{R}^2 \setminus \mathcal{K}$ is small. Hence, we can use the sampling technique just for $\bar{\mathcal{K}}$ and use the deterministic ε -kernel construction for \mathcal{K} .

Again consider the Poissonized instance of \mathcal{P} . Imagine the following process. Fix a direction u . We move a sweep line ℓ_u orthogonal to u , along the direction u , to sweep through the points in \mathcal{P} . We use H_u to denote the halfplane defined by ℓ_u (with normal vector u) and \bar{H}_u denote its complement. So $\mathcal{P}(\bar{H}_u) = \mathcal{P} \cap \bar{H}_u$ is the set of points that have been swept so far. We stop the movement of ℓ_u at the first point such that $\sum_{v \in \bar{H}_u} \lambda_v \geq \ln(2/\tau)$ (ties should be broken in an arbitrary but consistent manner). One important property about \bar{H}_u is that $\Pr[\bar{H}_u \models 0] \leq \tau/2$. We repeat the above process for all directions u and let $\mathcal{H} = \cap_u \bar{H}_u$. Since $\lambda(\mathcal{P}) > 3 \ln(2/\tau)$, by Helly's theorem, \mathcal{H} is nonempty. A careful examination of the above process reveals that \mathcal{H} is in fact a convex polytope and each edge of the polytope is defined by two points in \mathcal{P} .

The construction of the (ε, τ) -QUANT-KERNEL \mathcal{S} is as follows. First, we use the algorithm in [7] to find a deterministic ε -kernel $\mathcal{E}_{\mathcal{H}}$ of size $O(\varepsilon^{-1/2})$ for \mathcal{H} . One useful property of the algorithm in [7] is that $\mathcal{E}_{\mathcal{H}}$ is a subset of the vertices of \mathcal{H} . Hence the convex polytope $\text{ConvH}(\mathcal{E}_{\mathcal{H}})$ is contained in \mathcal{H} . Since $\mathcal{E}_{\mathcal{H}}$ is an ε -kernel, $(1 + \varepsilon)\text{ConvH}(\mathcal{E}_{\mathcal{H}})$ (properly shifted) contains \mathcal{H} . Let $\mathcal{K} = (1 + \varepsilon)\text{ConvH}(\mathcal{E}_{\mathcal{H}})$ and $\bar{\mathcal{K}} = \mathcal{P} \setminus \mathcal{K}$. See Figure 2.

Now, we apply the random sampling construction over $\bar{\mathcal{K}}$. More specifically, let $\lambda := \lambda(\bar{\mathcal{K}}) = \sum_{v \in \bar{\mathcal{K}} \cap \mathcal{P}} \lambda_v$. Let \mathfrak{A} be the probability measure over $\mathcal{P} \cap \bar{\mathcal{K}}$ defined by $\mathfrak{A}(\{v\}) = \lambda_v/\lambda$ for every $v \in \mathcal{P} \cap \bar{\mathcal{K}}$. Let $\tau_1 = O(\tau/\lambda)$. We take $N = O(\tau_1^{-2} \log(1/\delta))$ independent samples from \mathfrak{A} and let \mathfrak{B} be the empirical distribution with each sample point having probability $1/N$. The (ε, τ) -QUANT-KERNEL \mathcal{S} consists of the N points in \mathfrak{B} , each with the same existential probability $1 - \exp(-\lambda/N)$, as well as all vertices of \mathcal{K} , each with probability 1.


► **Theorem 19.** *Let $\lambda = \lambda(\bar{\mathcal{K}})$ and $\tau_1 = O(\tau/\lambda)$, and $N = O(\frac{1}{\tau_1^2} \log \frac{1}{\delta}) = O(\frac{\ln^2 1/\tau}{\varepsilon \tau^2} \log \frac{1}{\delta})$. With probability at least $1 - \delta$, for any $t \geq 0$ and any direction u , we have that*

$$\Pr[\omega(\mathcal{S}, u) \leq t] \in \Pr[\omega(\mathcal{P}, u) \leq (1 \pm \varepsilon)t] \pm \tau. \quad (3)$$

In summary, we obtain Theorem 6 combining Theorem 18 and Theorem 19.

4 (ε, r) -fpow-kernel Under the β -Assumption

We now show an (ε, r) -FPOW-KERNEL exists in the existential uncertainty model under the β -assumption. Recall that the function $T_r(P, u) = \max_{v \in P} \langle u, v \rangle^{1/r} - \min_{v \in P} \langle u, v \rangle^{1/r}$. For



■ **Figure 2** Construction of (ε, τ) -QUANT-KERNEL \mathcal{S} . The dashed polygon is \mathcal{H} . The inner solid polygon is $\text{ConvH}(\mathcal{E}_{\mathcal{H}})$ and the outer one is $K = (1 + \varepsilon)\text{ConvH}(\mathcal{E}_{\mathcal{H}})$. \bar{K} is the set of points outside K .

ease of notation, we write $\mathbb{E}[T_r(\mathcal{P}, u)]$ to denote $\mathbb{E}_{P \sim \mathcal{P}}[T_r(P, u)]$. Our goal is to find a set \mathcal{S} of stochastic points so for all directions $u \in \mathcal{P}^*$, that $\mathbb{E}[T_r(\mathcal{S}, u)] \in (1 \pm \varepsilon)\mathbb{E}[T_r(\mathcal{P}, u)]$.

Our construction of \mathcal{S} is almost the same as that in Section 3.1. Suppose we sample N (fixed later) independent realizations and take the ε_0 -kernel for each of them. Suppose they are $\{\mathcal{E}_1, \dots, \mathcal{E}_N\}$ and we associate each a probability $1/N$. We denote the resulting (ε, r) -FPOW-KERNEL by \mathcal{S} . Hence, for any direction $u \in \mathcal{P}^*$, $\mathbb{E}[T_r(\mathcal{S}, u)] = \frac{1}{N} \sum_{i=1}^N T_r(\mathcal{E}_i, u)$ and we use this value as the estimation of $\mathbb{E}[T_r(\mathcal{P}, u)]$. Our result is summarized by Theorem 7.

5 Applications

In this section, we show that our coreset results for the directional width problem readily imply several coreset results for other stochastic problems, just as in the deterministic setting. We introduce these stochastic problems and briefly summarize our results below.

5.1 Approximating the Extent of Uncertain Functions

We first consider the problem of approximating the extent of a set \mathcal{H} of uncertain functions. As before, we consider both the existential model and the locational model of uncertain functions.

1. In the existential model, each uncertain function h is a function in \mathbb{R}^d associated with a existential probability p_f , which indicates the probability that h presents in a random realization.
2. In the locational model, each uncertain function h is associated with a finite set $\{h_1, h_2, \dots\}$ of deterministic functions in \mathbb{R}^d . Each h_i is associated with a probability value $p(h_i)$, such that $\sum_i p(h_i) = 1$. In a random realization, h is independently realized to some h_i , with probability $p(h_i)$.

We use \mathcal{H} to denote the random instance, that is a random set of functions. We use $h \in \mathcal{H}$ to denote the event that the deterministic function h is present in the instance. For each point $x \in \mathbb{R}^d$, we let the random variable $\mathfrak{E}_{\mathcal{H}}(x) = \max_{h \in \mathcal{H}} h(x) - \min_{h \in \mathcal{H}} h(x)$ be the extent of \mathcal{H} at point x . Suppose \mathcal{S} is another set of uncertain functions. We say \mathcal{S} is the ε -EXP-KERNEL for \mathcal{H} if $(1 - \varepsilon)\mathfrak{E}_{\mathcal{H}}(x) \leq \mathfrak{E}_{\mathcal{S}}(x) \leq \mathfrak{E}_{\mathcal{H}}(x)$ for any $x \in \mathbb{R}^d$. We say \mathcal{S} is the (ε, τ) -QUANT-KERNEL for \mathcal{H} if $\Pr_{\mathcal{S} \sim \mathcal{S}}[\mathfrak{E}_{\mathcal{S}}(x) \leq t] \in \Pr_{\mathcal{H} \sim \mathcal{H}}[\mathfrak{E}_{\mathcal{H}}(x) \leq (1 \pm \varepsilon)t] \pm \phi$, for any $t \geq 0$ and any $x \in \mathbb{R}^d$.

Let us first focus on linear functions in \mathbb{R}^d . Using the *duality transformation* that maps linear function $y = a_1x_1 + \dots + a_dx_d + a_{d+1}$ to the point $(a_1, \dots, a_{d+1}) \in \mathbb{R}^{d+1}$, we can reduce the extent problem to the directional width problem in \mathbb{R}^{d+1} . Let \mathcal{H} be a set of uncertain linear functions (under either existential or locational model) in \mathbb{R}^d for constant d . From Theorem 14 and Corollary 11, we can construct a set S of $O(n^{2d})$ deterministic linear functions in \mathbb{R}^d , such that $\mathfrak{E}_S(x) = \mathbb{E}[\mathfrak{E}_{\mathcal{H}}(x)]$ for any $x \in \mathbb{R}^d$. Moreover, for any $\varepsilon > 0$, there exists an ε -EXP-KERNEL of size $O(\varepsilon^{-d/2})$ and an (ε, τ) -QUANT-KERNEL of size $\tilde{O}(\tau^{-2}\varepsilon^{-d})$. Using the standard linearization technique [7], we can obtain the following generalization for uncertain polynomials.

► **Theorem 20.** *Let \mathcal{H} be a family of uncertain polynomials in \mathbb{R}^d (under either existential or locational model) that admits linearization of dimension k . We can construct a set M of $O(n^{2k})$ deterministic polynomials, such that $\mathfrak{E}_M(x) = \mathbb{E}[\mathfrak{E}_{\mathcal{H}}(x)]$ for any $x \in \mathbb{R}^d$. Moreover, for any $\varepsilon > 0$, there exists an ε -EXP-KERNEL of size $O(\varepsilon^{-k/2})$ and an (ε, τ) -QUANT-KERNEL of size $\min\{\tilde{O}(\tau^{-2} \max\{\lambda^2, \lambda^4\}), \tilde{O}(\varepsilon^{-k}\tau^{-2})\}$. Here $\lambda = \sum_{h \in \mathcal{H}} (-\ln(1 - p_h))$.*

Now, we consider functions of the form $u(x) = p(x)^{1/r}$ where $p(x)$ is a polynomial and r is a positive integer. We call such a function a *fractional polynomial*. We still use \mathcal{H} to denote the random set of fractional polynomials. Let $\mathcal{H}^* \subseteq \mathbb{R}^d$ be the set of points such that for any points $x \in \mathcal{H}^*$ and any function $u \in \mathcal{H}$, we have $u(x) \geq 0$. For each point $x \in \mathcal{H}^*$, we let the random variable $\mathfrak{E}_{r, \mathcal{H}}(x) = \max_{h \in \mathcal{H}} h(x)^{1/r} - \min_{h \in \mathcal{H}} h(x)^{1/r}$. We say another random set \mathcal{S} of functions is the (ε, r) -FPOW-KERNEL for \mathcal{H} if $(1 - \varepsilon)\mathfrak{E}_{r, \mathcal{H}}(x) \leq \mathfrak{E}_{r, \mathcal{S}}(x) \leq \mathfrak{E}_{r, \mathcal{H}}(x)$ for any $x \in \mathcal{H}^*$. By the duality transformation and Theorem 7, we can obtain the following result.

► **Theorem 21.** *Let \mathcal{H} be a family of uncertain fractional polynomials in \mathbb{R}^d in the existential uncertainty model under the β -assumption. Further assume that each polynomial admits a linearization of dimension k . For any $\varepsilon > 0$, there exists an (ε, r) -FPOW-KERNEL of size $\tilde{O}(\varepsilon^{-(rk-r+2)})$. Furthermore, the (ε, r) -FPOW-KERNEL consists of $N = O(\varepsilon^{-(rk-r+4)/2})$ sets, each occurring with probability $1/N$ and containing $O(\varepsilon^{-r(k-1)/2})$ deterministic fractional polynomials.*

5.2 Stochastic Moving Points

We can extend our stochastic models to moving points. In the existential model, each point v is present with probability p_v and follows a trajectory $v(t)$ in \mathbb{R}^d when present ($v(t)$ is the position of v at time t). In the locational model, each point v is associated with a distribution of trajectories (the support size is finite) and the actual trajectory of v is a random sample for the distribution. Such uncertain trajectory models have been used in several applications in spatial databases [54]. For ease of exposition, we assume the existential model in the following. Suppose each trajectory is a polynomial of t with degree at most r . For each point v , any direction u and time t , define the polynomial $f_v(u, t) = \langle v(t), u \rangle$ and let \mathcal{H} include f_v with probability p_v . For a set \mathcal{P} of points, the directional width at time t is $\mathfrak{E}_{\mathcal{H}}(u, t) = \max_{v \in \mathcal{P}} f_v(u, t) - \min_{v \in \mathcal{P}} f_v(u, t)$. Each polynomial f_v admits a linearization of dimension $k = (r + 1)d - 1$. Using Theorem 20, we can see that there is a set M of $O(n^{2k})$ deterministic moving points, such that the directional width of M in any direction u is the same as the expected directional width of \mathcal{P} in direction u . Moreover, for any $\varepsilon > 0$, there exists an ε -EXP-KERNEL (which consists of only deterministic moving points) of size $O(\varepsilon^{-(k-1)/2})$ and an (ε, τ) -QUANT-KERNEL (which consists of both deterministic and stochastic moving points) of size $\tilde{O}(\varepsilon^{-k}\tau^{-2})$.

5.3 Shape Fitting Problems

Theorem 20 can be also applied to some stochastic variants of certain shape fitting problems. We first consider the following variant of the minimum enclosing ball problem over stochastic points. We are given a set \mathcal{P} of stochastic points (under either existential or locational model), find the center point c such that $\mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\|^2]$ is minimized. It is not hard to see that the problem is equivalent to minimizing the expected area of the enclosing ball in \mathbb{R}^2 . For ease of exposition, we assume the existential model where v is present with probability p_v . For each point $v \in \mathcal{P}$, define the polynomial $h_v(x) = \|x\|^2 - 2\langle x, v \rangle + \|v\|^2$, which admits a linearization of dimension $d + 1$ [7]. Let \mathcal{H} be the family of uncertain polynomials $\{h_v\}_{v \in \mathcal{P}}$ (h_v exists with probability p_v). We can see that for any $x \in \mathbb{R}^d$, $\max_{v \in \mathcal{P}} \|x - v\|^2 = \max_{h_v \in \mathcal{H}} h_v(x)$. Using Theorem 20, we can see that there is a set M of $O(n^{2d+2})$ deterministic polynomials such that $\max_{h \in M} h(x) = \mathbb{E}[\max_{v \in \mathcal{P}} \|x - v\|^2]$ for any $x \in \mathbb{R}^d$ and a set S of $O(\varepsilon^{-(d+1)/2})$ deterministic polynomials such that $(1 - \varepsilon)\mathbb{E}[\max_{v \in \mathcal{P}} \|x - v\|^2] \leq \max_{h \in S} h(x) \leq \mathbb{E}[\max_{v \in \mathcal{P}} \|x - v\|^2]$ for any $x \in \mathbb{R}^d$. We can store the set S instead of the original point set in order to answer the following queries: given a point v , return the expected length of the furthest point from v . The problem of finding the optimal center c can be also carried out over S , which can be done in $O(\varepsilon^{-O(d^2)})$ time: We can decompose the arrangement of n semialgebraic surfaces in \mathbb{R}^d into $O(n^{O(d+k)})$ cells of constant description complexity, where k is the linearization dimension (see e.g., [9]). By enumerating all those cells in the arrangement of S , we know which polynomials lie in the upper envelopes, and we can compute the minimum value in each such cell in constant time when d is constant.

The above argument can also be applied to the following variant of the spherical shell for stochastic points. We are given a set \mathcal{P} of stochastic points (under either existential or locational model). Our objective is to find the center point c such that $\mathbb{E}[\text{obj}(c)] = \mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\|^2 - \min_{v \in \mathcal{P}} \|v - c\|^2]$ is minimized. The problem is equivalent to minimizing the expected area of the enclosing annulus in \mathbb{R}^2 . The objective can be represented as a polynomial of linearization dimension $k = d + 1$. Proceeding as for the enclosing balls, we can show there is a set S of $O(\varepsilon^{-(k-1)/2})$ deterministic polynomials such that $(1 - \varepsilon)\mathbb{E}[\text{obj}(c)] \leq \mathfrak{C}_S(x) \leq \mathbb{E}[\text{obj}(c)]$ for any $x \in \mathbb{R}^d$. We summarize our results by Theorem 8. We would like to make a few remarks here.

1. We take the minimum enclosing ball for example. If we examine the construction of set S , each polynomial $h \in S$ may *not* be of the form $h(x) = \|x\|^2 - 2\langle x, v \rangle + \|v\|^2$, therefore does not translate back to a minimum enclosing ball problem over deterministic points.
2. Another natural objective function for the minimum enclosing ball and the spherical shell problem would be the expected radius $\mathbb{E}[\max_{v \in \mathcal{P}} d(v, c)]$ and the expected shell width $\mathbb{E}[\max_{v \in \mathcal{P}} d(v, c) - \min_{v \in \mathcal{P}} d(v, c)]$. However, due to the fractional powers (square roots) in the objectives, simply using an ε -EXP-KERNEL does not work. This is unlike the deterministic setting.² We leave the problem of finding small coresets for the spherical shell problem as an interesting open problem. However, under the β -assumption, we can use (ε, r) -FPOW-KERNELS to handle such fractional powers, as in the next subsection.

5.4 Shape Fitting Problems (Under the β -assumption)

In this subsection, we consider several shape fitting problems in the existential model *under the β -assumption*. We show how to use Theorem 21 to obtain linear time approximation schemes for those problems.

² In particular, there is no stochastic analogue of Lemma 4.6 in [7].

1. (Minimum spherical shell) We first consider the minimum spherical shell problem. Given a set \mathcal{P} of stochastic points (under the β -assumption), our goal is to find the center point c such that $\mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\| - \min_{v \in \mathcal{P}} \|v - c\|]$ is minimized. For each point $v \in \mathcal{P}$, let $h_v(x) = \|x\|^2 - 2\langle x, v \rangle + \|v\|^2$, which admits a linearization of dimension $d + 1$. It is not hard to see that $\mathbb{E}[\max_{v \in \mathcal{P}} \|v - c\|] = \mathbb{E}[\max_{v \in \mathcal{P}} \sqrt{h_v(c)}]$ and $\mathbb{E}[\min_{v \in \mathcal{P}} \|v - c\|] = \mathbb{E}[\min_{v \in \mathcal{P}} \sqrt{h_v(c)}]$. Using Theorem 21, we can see that there are $N = \tilde{O}(\varepsilon^{-(d+3)})$ sets S_i , each containing $O(\varepsilon^{-(d+1)})$ fractional polynomial $\sqrt{h_v}$ s such that for all $x \in \mathbb{R}^d$,

$$\frac{1}{N} \sum_{i \in [N]} (\max_{S_i} \sqrt{h_v(x)} - \min_{S_i} \sqrt{h_v(x)}) \in (1 \pm \varepsilon)(\mathbb{E}[\max_{v \in \mathcal{P}} \|v - x\|] - \mathbb{E}[\min_{v \in \mathcal{P}} \|v - x\|]). \quad (4)$$

Note that our (ε, r) -FPOW-KERNEL satisfies the subset constraint. Hence, each function $\sqrt{h_v}$ corresponds to an original point in \mathcal{P} . So, we can store N point sets $P_i \subseteq \mathcal{P}$, with $|P_i| = O(\varepsilon^{-d})$ as the coreset for the original point set. By (4), an optimal solution for the coreset is an $(1 + \varepsilon)$ -approximation for the original problem.

Now, we briefly sketch how to compute the optimal solution for the coreset. Consider all points in $\cup_i P_i$. Consider the arrangement of $O(\varepsilon^{-O(d)})$ hyperplanes, each bisecting a pair of points in $\cup_i P_i$. For each cell C of the arrangement, for any point $v \in C$, the ordering of all points in $\cup_i P_i$ is fixed. We then enumerate all those cells in the arrangement and try to find the optimal center in each cell. Fix a cell C . For any point set P_i , we know which point is the furthest one and which point is the closest one from points in C_0 . Say they are $v_i = \arg \max_{v \in P_i} \|v - x\|$ and $v'_i = \arg \min_{v \in P_i} \|v - x\|$. Hence, our problem can be formulated as the following optimization problem:

$$\min_x \frac{1}{N} \sum_i (d_i - d'_i), \quad \text{s.t.} \quad d_i^2 = \|v_i - x\|^2, d'_i{}^2 = \|v'_i - x\|^2, d_i, d'_i \geq 0, \forall i \in [N]; x \in C_0.$$

The polynomial system has a constant number of variables and constraints, hence can be solved in constant time. More specifically, we can introduce a new variable t and let $t = \frac{1}{N} \sum_i (d_i - d'_i)$. All polynomial constraints define a semi-algebraic set. By using constructive version of Tarski-Seidenberg theorem, we can project out all variables except t and the resulting set is still a semi-algebraic set (which would be a finite collection of points and intervals in \mathbb{R}^1) (See e.g., [12]).

2. (Minimum enclosing cylinder, Minimum cylindrical shell) Let \mathcal{P} be a set of stochastic points in the existential uncertainty model under the β -assumption. Let $d(\ell, v)$ denote the distance between a point $v \in \mathbb{R}^d$ and a line $\ell \subset \mathbb{R}^d$. The goal for the minimum enclosing cylinder problem is to find a line ℓ such that $\mathbb{E}[\max_{v \in \mathcal{P}} d(\ell, v)]$ is minimized, while that for the minimum cylindrical shell problem is to minimize $\mathbb{E}[\max_{v \in \mathcal{P}} d(\ell, v) - \min_{v \in \mathcal{P}} d(\ell, v)]$. The algorithms for both problems are almost the same and we only sketch the one for the minimum enclosing cylinder problem.

We follow the approach in [7]. We represent a line $\ell \in \mathbb{R}^d$ by a $(2d - 1)$ -tuple $(x_1, \dots, x_{2d-1}) \in \mathbb{R}^{2d-1}$: $\ell = \{p + tq \mid t \in \mathbb{R}\}$, where $p = (x_1, \dots, x_{d-1}, 0)$ is the intersection point of ℓ with the hyperplane $x_d = 0$ and $q = (x_d, \dots, x_{2d-1})$, $\|q\|^2 = 1$ is the orientation of ℓ . Then for any point $v \in \mathbb{R}^d$, we have that

$$d(\ell, v) = \|(p - v) - \langle p - v, q \rangle q\|,$$

where the polynomial $d^2(\ell, v)$ admits a linearization of dimension $O(d^2)$. Now, proceeding as for the minimum enclosing ball problem and using Theorem 21, we can obtain a coreset \mathcal{S} consisting $N = O(\varepsilon^{-O(d^2)})$ deterministic point sets $P_i \subseteq \mathcal{P}$.

We briefly sketch how to obtain the optimal solution for the coreset. We can also decompose \mathbb{R}^{2d-1} (a point x in the space with $\|(x_d, \dots, x_{2d-1})\| = 1$ represents a line in \mathbb{R}^d) into $O(\epsilon^{-O(d^2)})$ semi-algebraic cells such that for each cell, the ordering of the points in \mathcal{S} (by their distances to a line in the cell) is fixed. Note that such a cell is a semi-algebraic cell. For a cell C , assume that $v_i = \arg \max_{v \in P_i} d(\ell, v_i)$ for all $i \in [N]$, where ℓ is an arbitrary line in C . We can formulate the problem as the following polynomial system:

$$\min_l \frac{1}{N} \sum_i d_i, \quad \text{s.t.} \quad d_i^2 = d^2(\ell, v_i), d_i \geq 0, \forall i \in [N]; \ell = (p, q) \in C_0, \|q\|^2 = 1.$$

Again the polynomial system has a constant number of variables and constraints. Thus, we can compute the optimum in constant time. We summarize our results by Theorem 9.

References

- 1 A. Abdullah, S. Daruki, and J.M. Phillips. Range counting coresets for uncertain data. In *Proceedings 29th ACM Symposium on Computational Geometry*, pages 223–232, 2013.
- 2 Marcel R Ackermann, Johannes Blömer, and Christian Sohler. Clustering for metric and nonmetric distance measures. *ACM Transactions on Algorithms (TALG)*, 6(4):59, 2010.
- 3 P. Afshani, P.K. Agarwal, L. Arge, K.G. Larsen, and J.M. Phillips. (Approximate) uncertain skylines. In *Proceedings of the 14th International Conference on Database Theory*, pages 186–196, 2011.
- 4 P.K. Agarwal, S.-W. Cheng, and K. Yi. Range searching on uncertain data. *ACM Transactions on Algorithms (TALG)*, 8(4):43, 2012.
- 5 P.K. Agarwal, A. Efrat, S. Sankararaman, and W. Zhang. Nearest-neighbor searching under uncertainty. In *Proceedings of the 31st Symposium on Principles of Database Systems*, pages 225–236, 2012.
- 6 P.K. Agarwal, S. Har-Peled, S. Suri, H. Yıldız, and W. Zhang. Convex hulls under uncertainty. In *Proceedings of the 22nd Annual European Symposium on Algorithms*, pages 37–48, 2014.
- 7 P.K. Agarwal, S. Har-Peled, and K.R. Varadarajan. Approximating extent measures of points. *Journal of the ACM*, 51(4):606–635, 2004.
- 8 P.K. Agarwal, S. Har-Peled, and K.R. Varadarajan. Geometric approximation via coresets. *Combinatorial and Computational Geometry*, 52:1–30, 2005.
- 9 P.K. Agarwal and M. Sharir. Arrangements and their applications. *Handbook of Computational Geometry*, J. Sack and J. Urrutia (eds.), pages 49–119. Elsevier, Amsterdam, The Netherlands, 2000.
- 10 Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.
- 11 D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: Nearest neighbor relations for imprecise points. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, pages 410–419, 2004.
- 12 Saugata Basu, Richard Pollack, and M Roy. Algorithms in real algebraic geometry. *AMC*, 10:12, 2011.
- 13 T.M. Chan. Faster core-set constructions and data-stream algorithms in fixed dimensions. *Computational Geometry: Theory and Applications*, 35:20–35, 2006.
- 14 K. Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.
- 15 R. Cheng, J. Chen, and X. Xie. Cleaning uncertain data with quality guarantees. *Proceedings of the VLDB Endowment*, 1(1):722–735, 2008.

- 16 G. Cormode and A. McGregor. Approximation algorithms for clustering uncertain data. In *Proceedings of the 27th Symposium on Principles of Database Systems*, pages 191–200, 2008.
- 17 A. Deshpande, L. Rademacher, S. Vempala, and G. Wang. Matrix approximation and projective clustering via volume sampling. In *Proceedings of the 17th ACM-SIAM symposium on Discrete algorithm*, pages 1117–1126, 2006.
- 18 X. Dong, A.Y. Halevy, and C. Yu. Data integration with uncertainty. In *Proceedings of the 33rd International Conference on Very Large Data Bases*, pages 687–698, 2007.
- 19 A. Driemel, H. Haverkort, M. Löffler, and R.I. Silveira. Flow computations on imprecise terrains. *Journal of Computational Geometry*, 4:38–78, 2013.
- 20 W. Evans and J. Sember. The possible hull of imprecise points. In *Proceedings of the 23rd Canadian Conference on Computational Geometry*, 2011.
- 21 D. Feldman, A. Fiat, H. Kaplan, and K. Nissim. Private coresets. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 361–370, 2009.
- 22 D. Feldman and M. Langberg. A unified framework for approximating and clustering data. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 569–578, 2011.
- 23 Dan Feldman and Leonard J Schulman. Data reduction for weighted and outlier-resistant clustering. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1343–1354. SIAM, 2012.
- 24 Martin Fink, John Hershberger, Nirman Kumar, and Subhash Suri. Hyperplane separability and convexity of probabilistic point sets. In *Proceedings Symposium on Computational Geometry*, 2016.
- 25 P.K. Ghosh and K.V. Kumar. Support function representation of convex bodies, its application in geometric computing, and some related representations. *Computer Vision and Image Understanding*, 72(3):379–403, 1998.
- 26 S. Guha and K. Munagala. Exceeding expectations and clustering uncertain data. In *Proceedings of the 28th Symposium on Principles of Database Systems*, pages 269–278, 2009.
- 27 L.J. Guibas, D. Salesin, and J. Stolfi. Constructing strongly convex approximate hulls with inaccurate primitives. *Algorithmica*, 9:534–560, 1993.
- 28 S. Har-Peled. On the expected complexity of random convex hulls. *arXiv:1111.5340*, 2011.
- 29 S. Har-Peled and S. Mazumdar. On coresets for k-means and k-median clustering. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 291–300, 2004.
- 30 Sariel Har-Peled and Yusu Wang. Shape fitting with outliers. *SIAM Journal on Computing*, 33(2):269–285, 2004.
- 31 M. Held and J.S.B. Mitchell. Triangulating input-constrained planar point sets. *Information Processing Letters*, 109(1):54–56, 2008.
- 32 Lingxiao Huang and Jian Li. Approximating the expected values for combinatorial optimization problems over stochastic points. In *The 42nd International Colloquium on Automata, Languages, and Programming*, pages 910–921. Springer, 2015.
- 33 A.G. Jørgensen, M. Löffler, and J.M. Phillips. Geometric computation on indecisive points. In *Proceedings of the 12th Algorithms and Data Structure Symposium*, pages 536–547, 2011.
- 34 P. Kamousi, T.M. Chan, and S. Suri. The stochastic closest pair problem and nearest neighbor search. In *Proceedings of the 12th Algorithms and Data Structure Symposium*, pages 548–559, 2011.
- 35 P. Kamousi, T.M. Chan, and S. Suri. Stochastic minimum spanning trees in euclidean spaces. In *Proceedings of the 27th Symposium on Computational Geometry*, pages 65–74, 2011.

- 36 H. Kruger. Basic measures for imprecise point sets in \mathbb{R}^d . Master's thesis, Utrecht University, 2008.
- 37 M. Langberg and L.J. Schulman. Universal ϵ -approximators for integrals. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, 2010.
- 38 J. Li and H. Wang. Range queries on uncertain data. *Theoretical Computer Science*, 609(1):32–48, 2016.
- 39 M. Löffler and J. Phillips. Shape fitting on point sets with probability distributions. In *Proceedings of the 17th European Symposium on Algorithms*, pages 313–324, 2009.
- 40 M. Löffler and J. Snoeyink. Delaunay triangulations of imprecise points in linear time after preprocessing. In *Proceedings of the 24th Symposium on Computational Geometry*, pages 298–304, 2008. doi:10.1145/1377676.1377727.
- 41 M. Löffler and M. van Kreveld. Approximating largest convex hulls for imprecise points. *Journal of Discrete Algorithms*, 6:583–594, 2008.
- 42 J. Matoušek. Computing the center of planar point sets. *Discrete and Computational Geometry*, 6:221, 1991.
- 43 A. Munteanu, C. Sohler, and D. Feldman. Smallest enclosing ball for probabilistic data. In *Proceedings of the 30th Annual Symposium on Computational Geometry*, 2014.
- 44 T. Nagai and N. Tokura. Tight error bounds of geometric problems on convex objects with imprecise coordinates. In *Jap. Conf. on Discrete and Comput. Geom.*, LNCS 2098, pages 252–263, 2000.
- 45 Y. Ostrovsky-Berman and L. Jostkowicz. Uncertainty envelopes. In *Abstracts of the 21st European Workshop on Comput. Geom.*, pages 175–178, 2005.
- 46 Jeff M. Phillips. Coresets and sketches. In *Handbook of Discrete and Computational Geometry*. CRC Press, 3rd edition, 2016. Chapter 49.
- 47 D. Salesin, J. Stolfi, and L.J. Guibas. Epsilon geometry: building robust algorithms from imprecise computations. In *Proceedings of the 5th Symposium on Computational Geometry*, pages 208–217, 1989.
- 48 R. Schneider. *Convex bodies: the Brunn-Minkowski theory*, volume 44. Cambridge University Press, 1993.
- 49 S. Suri, K. Verbeek, and H. Yildiz. On the most likely convex hull of uncertain points. In *Proceedings of the 21st European Symposium on Algorithms*, pages 791–802, 2013.
- 50 M. van Kreveld and M. Löffler. Largest bounding box, smallest diameter, and related problems on imprecise points. *Computational Geometry: Theory and Applications*, 43:419–433, 2010. doi:10.1016/j.comgeo.2009.03.007.
- 51 V.N. Vapnik and A.Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- 52 Jie Xue, Yuan Li, and Ravi Janardan. On the separability of stochastic geometric objects, with applications. In *Proceedings Symposium on Computational Geometry*, 2016.
- 53 H. Yu, P.K. Agarwal, R. Poredy, and K. Varadarajan. Practical methods for shape fitting and kinetic data structures using coresets. *Algorithmica*, 52(378-402), 2008.
- 54 K. Zheng, G. Trajcevski, X. Zhou, and P. Scheuermann. Probabilistic range queries for uncertain trajectories on road networks. In *Proceedings of the 14th International Conference on Extending Database Technology*, pages 283–294, 2011.