Coreset Construction and Estimation over Stochastic Data

Dissertation Submitted to

Tsinghua University

in partial fulfillment of the requirement

for the degree of

Doctor of Philosophy

in

Computer Science and Technology

by

Lingxiao Huang

Dissertation Supervisor: Assistant Professor Jian Li

June 2017

Coreset Construction and Estimation over Stochastic Data

by

Lingxiao Huang

Submitted to the Institute for Interdisciplinary Information Sciences
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at

TSINGHUA UNIVERSITY

June 2017

© TSINGHUA UNIVERSITY 2017. All rights reserved.

Author .....
Institute for Interdisciplinary Information Sciences
April 15, 2015

Certified by.....
Jian Li
Assistant Professor
Thesis Supervisor

Coreset Construction and Estimation over Stochastic Data

by

Lingxiao Huang

Submitted to the Tsinghua University
on April 15, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In recent years, designing algorithms for geometric or combinatorial optimization problems over stochastic data have attracted more and more research interest. In this dissertation, we consider two well-known stochastic geometry models. One is the existential model where each point's location is fixed but only occurs with a certain probability. The other one is the locational model where each point has a probability distribution describing its location. Both stochastic geometry models have been widely studied in recent years. In this dissertation, we mainly focus on the following problems, coreset construction for shape fitting problems and estimation for combinatorial optimization problems on these stochastic geometry models.

The first problem concerns with a useful technique for handling large deterministic datasets, called coreset. Roughly speaking, a coreset is a small summary of the original large dataset while guaranteeing that answers for certain queries are provably close to the exact answer of corresponding queries. In this dissertation, we study how to construct coresets on stochastic models. We first extend the concept ε-kernel coresets to stochastic data. 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 and show how to construct such coresets e□ciently. Then we consider two stochastic shape fitting problems, stochastic k-center and stochastic j-flat-center. We propose a new notion called generalized coresets, which is a generalization of coresets. We also provide a framework for constructing generalized coresets of constant size for both the stochastic k-center problem and the stochastic j-flat-center problem. Using these generalized coresets, we give the first PTASs (polynomial time approximation schemes) for both stochastic shape fitting problems.

Secondly, we study the problems of computing the expected lengths of several combinatorial or geometric optimization problems in stochastic geometry models, including closest pair, minimum spanning tree, k-clustering, minimum perfect matching, and minimum cycle cover. Most of the above problems are known to be #P-hard.

In this dissertation, we propose two new techniques, called finding stoch-core and Hierachical Partition Family (HPF). Combining our new techniques and Monte Carlo method, we obtain the first FPRAS (Fully Polynomial Randomized Approximation Scheme) for most of these problems in stochastic geometry models.

Dissertation Supervisor: Assistant Professor Jian Li

Contents

1 Introduction 1

1.1 ε-Kernel Coresets over Stochastic Data 4

1.2 Coreset Construction for Shape Fitting Problems over Stochastic Data 9

1.3 Estimation for Combinatorial Optimization Problems over Stochastic Data 15

2 Preliminaries 20

3 Related Work 21

4 ε-Kernel Coresets over Stochastic Data 25

4.1 ε-Kernel Coresets over Deterministic Data 25

4.2 ε-Kernels for Expectations of Width 27

4.2.1 A Nearly Linear Time Algorithm for Constructing ε-exp-kernels 31

4.2.2 ε-exp-kernel Under the Subset Constraint 33

4.3 ε-Kernels for Probability Distributions of Width 34

4.3.1 A Simple (ε, τ)-quant-kernel Construction 34

4.3.2 Improved (ε, τ)-quant-kernel for Existential Models 37

4.3.3 (ε, τ)-quant-kernel Under the Subset Constraint 51

4.4 (ε, r)-fpow-kernel Under the β-Assumption 52

4.5 Applications 54

4.5.1 Approximating the Extent of Uncertain Functions 54

4.5.2 Stochastic Moving Points 56

4.5.3 Shape Fitting Problems 56

4.5.4 Shape Fitting Problems (Under the β-assumption) 58

4.6 Missing Details in Section 4.2 61

4.6.1 Details for Section 4.2.1 61
4.6.2 Details for Section 4.2.2 62
4.6.3 Locational uncertainty 64
4.7 Missing Details in Section 4.3 66
4.8 Missing Details in Section 4.4 71
4.9 Computing the Expected Direction Width 72
4.9.1 Computing Expected Width for Existential Uncertainty 72
4.9.2 Computing Expected Width for Locational Uncertainty 77

5 Coreset Construction for Stochastic Shape Fitting Problems 82
5.1 Coreset Construction for Deterministic Shape Fitting Problems . . . 82
5.2 Generalized Shape Fitting Problems and Generalized Coresets 85
5.3 Stochastic Minimum k-Center 90
5.3.1 Existential uncertainty model 91
5.3.2 Locational uncertainty model 104
5.4 Stochastic Minimum j-Flat-Center 107
5.4.1 Case 1: $B < \varepsilon$ 107
5.4.2 Case 2: $B \geq \varepsilon$ 108
5.5 Constructing additive ε-coresets 115

6 Estimating the Expected Value of Combinatorial Optimization Problems over Stochastic Data 121
6.1 The Closest Pair Problem 121
6.1.1 Estimating Pr[C $\leq$ 1] 121
6.1.2 Estimating $\mathbb{E}$[C] 125
6.2 k-Clustering 129
6.3 Minimum Spanning Trees 131
6.4 Minimum Perfect Matchings 137
6.5 Minimum Cycle Covers 145
6.6 kth Longest m-Nearest Neighbor 150
6.7 Missing Proofs 152

6.7.1 Closest Pair 152
6.7.2 Minimum Spanning Tree 152
6.7.3 Minimum Perfect Matching 154
6.8 The Closest Pair Problem 155
6.8.1 Estimating kth Closest Pair in the Existential Uncertainty Model 155
6.8.2 Hardness for Closest Pair 156
6.9 Another FPRAS for MST 158

7 Concluding Remarks 160

Acknowledgements 171

List of Figures

4-1 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 $C_1, \ldots, C_5$, by 5 angles $\theta_1, \ldots, \theta_5$. $\vec{u}_{\theta_i}$ is the unit direction corresponding to angle θ_i. Each cone C_i corresponds to a vertex s_i and for any direction $\vec{u} \in C_i$, $f(M, \vec{u}) = \langle \vec{u}, s_i \rangle$ and the vector $\nabla f(M, \vec{u})$ is s_i. (2) Each direction θ_i is perpendicular to an edge of M. $M = \cap_{i=1}^{5} H_i$ where H_i is the supporting halfplane with normal vector $\vec{u}_{\theta_i}$. 31

4-2 The construction of the (ε, τ)-quant-kernel $\mathcal{S}$. The dashed polygon is $\mathcal{H}$. The inner solid polygon is ConvH($\mathcal{E}_{\mathcal{H}}$) and the outer one is $K = (1 + \varepsilon)$ConvH($\mathcal{E}_{\mathcal{H}}$). $\overline{\mathcal{K}}$ is the set of points outside $\mathcal{K}$. 42

4-3 Illustration of the interval graph $\mathcal{I}$. For illustration purpose, co-located points (e.g., points that are split in A) are shown as overlapping points. The arrows indicate the assignment of the segments to the points in B. Theorem 32 ensures that any vertical line can not stab many intervals. 68

4-4 Illustrating the definition of the angle α of: (a) a ray ρ and (b) a line l. 74

4-5 Illustrating the computation of the coordinate $x(s, \vec{u})$ on $l(\vec{u})$: $v(\vec{u})$ is the perpendicular projection of s on $l(\vec{u})$. The length of $\overline{ov}$ is d_s. 76

5-1 An example for Algorithm 1 when k = 2. In this figure, $\mathcal{P} = \{s_1, \ldots, s_{11}\}$ consists of all points, and $S = \{s_3, s_5, s_7\}$ consists of black points. Then by Lemma 70, we have that $\Pr_{P\sim\mathcal{P}}[\mathcal{E}(P) = S] = p_3p_5p_7(1-p_1)(1-p_2)(1-p_4)(1-p_{10})(1-p_{11})$. Now we run Algorithm 1 on S. In Step 1, we first construct a Cartesian grid $G(S)$ as in the figure, and construct a cell collection $\mathcal{C}(S) = \{C_1, C_2, C_3\}$ since $C_4 \cap S = \emptyset$. Note that $\mathcal{E}(S) = S$ (by Lemma 70) and $|S| = 3 > k$. We directly go to Step 3 and want to compute the value $Q(C_i)$ for each cell C_i. For cell C_1, two rectangle points s_1 and s_2 are of smaller index than $s_3 \in S$. So we compute that $Q(C_1) = p_3(1-p_1)(1-p_2)$. Similarly, we compute $Q(C_2) = p_5(1-p_4)$, $Q(C_3) = p_7$, and $Q(C_4) = (1-p_{10})(1-p_{11})$. Finally in Step 4, we output $\Pr_{P\sim\mathcal{P}}[\mathcal{E}(P) = S] = \prod_{C\in G(S)} Q(C) = p_3p_5p_7(1-p_1)(1-p_2)(1-p_4)(1-p_{10})(1-p_{11})$. 95

5-2 In the figure, S_i is the black point set, F^* is the white point set, and F_i^* is the dashed point set. Here, $s_i^* \in S_i$ is the farthest point to F_i^* satisfying $d(s_i^*, F_i^*) = K(S_i, F_i^*)$, and $f_i^* \in F^*$ is the closest point to s_i^* satisfying $d(s_i^*, f_i^*) = d(s_i^*, F^*)$. 97

List of Tables

1.1 Our results for some problems in different stochastic models. 16

Chapter 1 Introduction

In recent years, stochastic data are pervasive in applications. Managing, analyzing and optimizing over such stochastic data have become an increasingly important issue and have attracted significant attentions from several research communities including theoretical computer science, databases, machine learning and sensor networks [29, 35, 100]. Theoretically, any deterministic combinatorial or geometry optimization problem has many uncertain counterparts (corresponding to different uncertainty models). A variety of classic problems in deterministic setting have been well studied, while systematic studies of them under uncertainty have only been initiated. For example, suppose we want to build k facilities to serve a set of uncertain demand points, and our goal is to minimize the expectation of the maximum distance from any realized demand point to its closest facility. This problems is called stochastic k-center which is first considered in this dissertation. Estimation and solving optimization problems over stochastic models and data have recently attracted significant attentions in several research communities (see e.g., [96, 100, 102]).

In the following, we list some application examples to exemplify where the stochastic data comes from, and what types of stochastic data we may encounter, see [79] for more examples.

1. Stochastic shortest path. Consider a traffic problem where we want to arrive at the airport before a specific time. There are several paths leading to the airport. In deterministic setting, each road costs a certain time. Our goal is to find a shortest path which minimizes the travel time. However in reality, we often know the distribution of the travel time of each road rather than the exact time. In this stochastic setting, we want to pick a path that maximizes the probability we can arrive at the airport on time. This problem has been considered a lot since the 1980s [18, 78, 81, 86, 90, 91].

2. Fixed set stochastic knapsack. The knapsack problem is a classic scheduling problem. One natural motivation is as follows. Suppose we are given a series of jobs and a single machine. Each job has a processing time and a profit. Our goal is to pick a subset of jobs which can be finished by a single machine one by one before the deadline and to maximize the total profit of these jobs. However, the processing time of each job is often random and follows from an individual distribution. So a stochastic variant of this problem is called the fixed set stochastic knapsack problem. We still want to choose a set of jobs to maximize the total profit. Since the processing time is random, we have an additional constraint that the probability that finishing all jobs before the deadline is at least some fixed constant $\gamma > 0$. In previous work, many groups of researchers have studied di□erent distributions, including Bernoulli [50, 75], exponential [50] and Gaussian [52].

3. Adaptive stochastic process. Note that the solution is chosen in advance for the above two examples. The adaptive variants of these problems have also been studied extensively. For the adaptive stochastic shortest path problem, we know the exact travel time immediately when we pass through a road. Thus, adaptively choosing the following path may increase the probability that we arrive the airport before the deadline.

   Similarly we can consider adaptive policies for stochastic knapsack. In fact, researchers studied a more complicated version where the processing time is also random, and the precise values of the processing time and the profit are revealed when the job is completed. The goal is to gain as much profit as possible. Dean et al. [33] initially studied this problem and proposed a greedy algorithm. Later on, [24, 81] considered the same problem and improved their results.

   Many other adaptive stochastic problems have been well studied, such as adaptive stochastic matching [17] and adaptive stochastic probing [107].

4. Stochastic geometry. Theoretically, all deterministic computational geom-

etry problems have natural stochastic counterparts. We take the following stochastic variant of the facility location problem as an example. Suppose we want to build k facilities to serve a set of uncertain demand points (i.e., their locations are random), and our goal is to minimize the expectation of the maximum distance from any realized demand point to its closest facility. A variety of classic geometry problems in deterministic setting have been well studied, while systematic studies of them in different stochastic models have been initiated in recent years. Munteanu et al. [88] studied the stochastic minimum enclosing ball problem in fixed dimensional Euclidean space and gave a PTAS. In this dissertation, we also studied some fundamental problems in this area.

In this dissertation, we focus on two well-known stochastic geometry models: the locational uncertainty model and the existential uncertainty model. Both models have been studied extensively for a variety of computational geometry optimization problems or combinatorial optimization problems, such as closest pairs [70], nearest neighbors [5, 70], minimum spanning trees [65, 71], convex hulls [101], maxima [2], perfect matchings [65], clustering [30, 53] , minimum enclosing balls [88] and range queries [1, 4, 80]. We give the formal definitions of these two stochastic geometry models as follows.

1. Locational uncertainty model: We are given a metric space $\mathcal{P}$. The location of each node $v \in \mathcal{V}$ is a random point in the metric space $\mathcal{P}$ and the probability distribution is given as the input. Formally, we use the term nodes to refer to the vertices of the graph, points to describe the locations of the nodes in the metric space. We denote the set of nodes as $\mathcal{V} = \{v_1, \ldots, v_m\}$ and the set of points as $\mathcal{P} = \{s_1, \ldots, s_n\}$, where $m = |\mathcal{V}|$ and $n = |\mathcal{P}|$. A realization r can be represented by an m-dimensional vector $(r_1, \ldots, r_m) \in \mathcal{P}^m$ where point r_i is the location of node v_i for $1 \leq i \leq m$. Let R denote the set of all possible realizations. We assume that the distributions of the locations of nodes in the metric space $\mathcal{P}$ are independent, thus r occurs with probability $\Pr[r] = \prod_{i \in [m]} p_{v_i r_i}$, where p_{vs} represents the probability that the location of

node v is point $s \in \mathcal{P}$. The model is also termed as the locational uncertainty model in [71].

2. Existential uncertainty model: A closely related model is the existential uncertainty model where the location of a node is a fixed point in the given metric space, but the existence of the node is probabilistic. In this model, we use p_{s_i} to denote the probability that node v_i exists (if exists, its location is s_i). For simplicity, we also use p_i to represent p_{s_i}. A realization r can be represented by a subset $P \subset \mathcal{P}$ and Pr[r] = $\prod_{s_i \in P} p_i \prod_{s_i \notin P} (1 - p_i)$.

Now, we introduce and motivate two classes of problems in this dissertation: coreset construction and estimating the expected values of combinatorial optimization problems over stochastic data. We also briefly state our contributions one by one.

1.1 ε-Kernel Coresets over Stochastic Data

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 have become more relevant in the era of big data as they summarize large datasets by datasets with potentially much smaller size, and at the same time guarantee the answer to certain classes of queries to be close to the true answer. The notion of a coreset was studied in 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 notation and briefly review the definition of ε-kernel. For a set P of deterministic points, the support function $f(P, \vec{u})$ is defined to be $f(P, \vec{u}) = \max_{s \in P} \langle \vec{u}, s \rangle$ for $\vec{u} \in \mathbb{R}^d$, where $\langle ., . \rangle$ is the inner product. The directional width of P in direction $u \in \mathbb{R}^d$, denoted by $\omega(P, \vec{u})$, is defined by $\omega(P, \vec{u}) = f(P, \vec{u}) + f(P, -\vec{u})$. It is easy 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 $\vec{u} \in \mathbb{R}^d$, $(1 - \varepsilon)\omega(P, \vec{u}) \leq \omega(Q, \vec{u}) \leq \omega(P, \vec{u})$. For any set of n points, there

is an ε-kernel of size $O(\varepsilon^{-(d-1)/2})$ [7, 8], which can be constructed in $O(n + \varepsilon^{-(d-3/2)})$ time [26, 88].

Our contribution. (Chapter 4) Our main results can be summarized as follows:

1. 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 $\omega(\mathcal{P}, \vec{u}) = \mathbb{E}_{P \sim \mathcal{P}}[\omega(P, \vec{u})]$, where $P \sim \mathcal{P}$ means that P is a (random) realization of $\mathcal{P}$. We first consider how to construct ε-exp-kernel which is defined as follows:

   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 $\vec{u} \in \mathbb{R}^d$,

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

   Our first main result is that an ε-exp-kernel of size $O(\varepsilon^{-(d-1)/2})$ exists for both existential and locational uncertainty model and can be constructed in nearly linear time.

   Theorem 2. $\mathcal{P}$ 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, where n is the number of points (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 16). 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 20).

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 $\mathbb{R}^2$ (Section 4.9).

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 25). 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 26 and Section 4.6).

2. Sometimes it is useful to obtain more than just the expected value (say of the width) on a query; rather one may want to return (an approximation of) a representation of the full probability distribution that the query can take. So we also consider the construction of the following (ε, τ)-quant-kernel.

   Definition 3. 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 $\vec{u}$ and all $x \geq 0$,

$$\Pr_{P \sim \mathcal{P}}\left[\omega(P, \vec{u}) \leq (1-\varepsilon)x\right] - \tau \leq \Pr_{S \sim \mathcal{S}}\left[\omega(S, \vec{u}) \leq x\right] \leq \Pr_{P \sim \mathcal{P}}\left[\omega(P, \vec{u}) \leq (1+\varepsilon)x\right] + \tau.$$

   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 4. An (ε, τ)-quant-kernel of size $\tilde{O}\left(\tau^{-2}\varepsilon^{-3(d-1)/2}\right)$ can be constructed in $\tilde{O}\left(n\tau^{-2}\varepsilon^{-(d-1)}\right)$ time, under both existential and locational uncer-

tainty models.

The algorithm is surprisingly simple. Take a certain number of i.i.d. realizations, compute an ε-kernel for each realization, and then associate each 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 4.3.1.

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

Theorem 5. $\mathcal{P}$ is a set of uncertain points in $\mathbb{R}^d$ with existential uncertainty. There exists an (ε, τ)-quant-kernel for $\mathcal{P}$, which consists of a set of independent uncertain points of cardinality $\tilde{O}(\varepsilon^{-(d-1)}\tau^{-2})\}$. 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 4). 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 coreset. 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 $\overline{K} = \mathcal{P} \setminus K$ can be approximated using the same sampling algorithm as in the first algorithm and we can show that $\lambda(\overline{K})$ is small, thus requiring only a small number of samples. Our algorithm can be easily extended to $\mathbb{R}^d$ for any constant d and the size of the coreset is $\tilde{O}(\tau^{-2}\varepsilon^{-(d-1)})$. In Section 4.3.2, we show such an (ε, τ)-quant-kernel can be computed in $O(n\text{polylog}n)$ time using an

iterative sampling algorithm. Our technique has some interesting connections to other important geometric problems (such as the Tukey depth problem) [87], may be interesting in its own right.

3. The notion (ε, τ)-quant-kernel is also 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 $P^\circ = \{\vec{u} \in \mathbb{R}^d \mid \langle \vec{u}, s\rangle \geq 0, \forall s \in P\}$. Let r be a positive integer. Given a set P of points in $\mathbb{R}^d$ and $\vec{u} \in P^\circ$, we define a function

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

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

Definition 6. 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 $\vec{u} \in \mathcal{P}^\circ$,

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

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 where each point is present with probability above β. The algorithm is almost the same as the construction in Section 4.3.1 except that some parameters are different.

Theorem 7. (Section 4.4) An (ε, r)-fpow-kernel of size $\tilde{O}(\varepsilon^{-(rd-r+2)})$ can be constructed in $\tilde{O}\left(n\varepsilon^{-(rd-r+4)/2}\right)$ time in the existential uncertainty model under the β-assumption.

4. Finally, we show that the above results, combined with the duality and lineariza-

tion 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 e□cient 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 di□erent stochastic settings. We summarize our application results in the following theorems. The details can be found in Section 4.5.

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 $\mathbb{E}[\max_{s\in\mathcal{P}} \|s-c\|^2]$; (2) finding a center point c to minimize $\mathbb{E}[\text{obj}(c)] = \mathbb{E}[\max_{s\in P} \|s-c\|^2 - \min_{s\in P} \|s-c\|^2]$. 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 e□cient 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 $\mathcal{P}$.

1.2 Coreset Construction for Shape Fitting Problems over Stochastic Data

We study two classic geometric optimization problems, the k-center problem and the j-flat center problem in Euclidean spaces. Both problems are important in geometric data analysis. We generalize both problems to the stochastic settings. For the

stochastic k-center problem, we would like to find k points in a fixed dimensional Euclidean space, such that the expected value of the k-center objective is minimized. For the stochastic j-flat-center problem, we seek a j-flat (i.e., a j-dimensional a□ne subspace) such that the expected value of the maximum distance from any point to the j-flat is minimized. One of the motivations for this stochastic version comes from the stochastic variant of the ℓ_∞ regression problem. We still want to construct coresets for these two shape fitting problems.

In the following, we first define the two stochastic shape fitting problems. Then we briefly introduce our contributions and techniques.

Stochastic k-Center. The deterministic Euclidean k-center problem is a central problem in geometric optimization [11, 8]. It asks for a k-point set F in $\mathbb{R}^d$ such that the maximium distance from any of the n given points to its closest point in F is minimized.

Definition 10. For a set of points $P \in \mathbb{R}^d$, and a k-point set $F = \{(f_1, \ldots, f_k) \mid f_i \in \mathbb{R}^d, 1 \leq i \leq k\}$, we define $K(P, F) = \max_{s \in P} \min_{1 \leq i \leq k} d(s, f_i)$ as the k-center value of F w.r.t. P. We use $\mathcal{F}$ to denote the family of all k-point sets in $\mathbb{R}^d$. Given a set $\mathcal{P}$ of n stochastic points (in either the existential or locational uncertainty model) in $\mathbb{R}^d$, and a k-point set $F \in \mathcal{F}$, we define the expected k-center value of F w.r.t $\mathcal{P}$ as

$$K(\mathcal{P}, F) = \mathbb{E}_{P \sim \mathcal{P}}[K(P, F)].$$

In the stochastic minimum k-center problem, our goal is to find a k-point set $F \in \mathcal{F}$ which minimizes $K(\mathcal{P}, F)$. In this dissertation, we assume that both the dimensionality d and k are fixed constants.

Stochastic j-Flat-Center. The deterministic j-flat-center problem is defined as follows: given n points in $\mathbb{R}^d$, we would like to find a j-flat F (i.e., a j-dimensional a□ne subspace) such that the maximum distance from any given point to F is minimized. It is a common generalization of the minimum enclosing ball ($j = 0$), minimum enclosing cylinder ($j = 1$), and minimum width problems ($j = d - 1$), and has been

well studied in computational geometry [8, 45, 106]. Its stochastic version is also naturally motivated by the stochastic variant of the ℓ_∞ regression problem: Suppose we would like to fit a set of points by an a□ne subspace. However, those points may be produced by some machine learning algorithm, which associates some confidence level to each point (i.e., each point has an existential probability). This naturally gives rise to the stochastic j-flat-center problem. Formally, it is defined as follows.

Definition 11. Given a set P of n points in $\mathbb{R}^d$, and a j-flat $F \in \mathcal{F}$ ($0 \leq j \leq d-1$), where $\mathcal{F}$ is the family of all j-flats in $\mathbb{R}^d$, we define the j-flat-center value of F w.r.t. P to be $J(P, F) = \max_{s \in P} d(s, F)$, where $d(s, F) = \min_{f \in F} d(s, f)$ is the distance between point s and j-flat F. Given a set $\mathcal{P}$ of n stochastic points (in either the existential or locational model) in $\mathbb{R}^d$, and a j-flat $F \in \mathcal{F}$ ($0 \leq j \leq d-1$), we define the expected j-flat-center value of F w.r.t. $\mathcal{P}$ to be

$$J(\mathcal{P}, F) = \mathbb{E}_{P \sim \mathcal{P}}[J(P, F)].$$

In the stochastic minimum j-flat-center problem, our goal is to find a j-flat F which minimizes $J(\mathcal{P}, F)$.

Previous Results and Our contributions. Recall that a polynomial time approximation scheme (PTAS) for a minimization problem is an algorithm A that produces a solution whose cost is at most $1 + \varepsilon$ times the optimal cost in polynomial time, for any fixed constant $\varepsilon > 0$.

Stochastic k-Center. Cormode and McGregor [30] first studied the stochastic k-center problem in a finite metric graph under the locational uncertainty model, and obtained a bi-criterion constant approximation. Guha and Munagala [53] improved their result to a single-criterion constant factor approximation. Recently, Wang and Zhang [108] studied the stochastic k-center problem on a line, and proposed an e□-cient exact algorithm. No result better than a constant approximation is known for the Euclidean space $\mathbb{R}^d$ ($d \geq 2$). We obtain the first PTAS for the stochastic k-center problem in $\mathbb{R}^d$.

Theorem 12. Assume that both k and d are fixed constants. There exists a PTAS for the stochastic minimum k-center problem in $\mathbb{R}^d$, under either the existential or the locational uncertainty model.

Our result generalizes the PTAS for stochastic minimum enclosing ball by Munteanu et al. [88]. We remark that the assumption that k is a constant is necessary for getting a PTAS, since even the deterministic Euclidean k-center problem is APX-hard for arbitrary k even in $\mathbb{R}^2$ [42].

Stochastic j-Flat-Center. Our main result for the stochastic j-flat-center is as follows.

Theorem 13. Assume that the dimensionality d is a constant. There exists a PTAS for the stochastic minimum j-flat-center problem, under either the existential or the locational uncertainty model.

This result also generalizes the PTAS for stochastic minimum enclosing ball (i.e., 0-flat-center) by Munteanu et al. [88]. It also generalizes a previous PTAS for the stochastic minimum enclosing cylinder (i.e., 1-flat-center) problem in the existential model where the existential probability of each point is assumed to be lower bounded by a small fixed constant in Chapter 4.

Our techniques. Our techniques for both problems heavily rely on the powerful notion of coresets. In a typical deterministic geometric optimization problem, an instance P is a set of deterministic (weighted) points. A coreset S of P is a set of (weighted) points, such that the solution for the optimization problem over S is a good approximate solution for P. [1] In Chapter 4, we generalize the notion of ε-kernel coreset (for directional width) to stochastic points. However, our techniques can only handle directional width, and extending it to problems such as stochastic minimum enclosing cylinder requires certain technical assumption.

In this dissertation, we introduce a new framework for solving geometric optimization problems over stochastic points. For a stochastic instance $\mathcal{P}$, we consider

[1] It is possible to define coresets for other classes of optimization problems.

$\mathcal{P}$ as a collection of realizations $\mathcal{P} = \{P \mid P \sim \mathcal{P}\}$. Each realization P has a weight Pr[$\models P$], which is its realized probability. Now, we can think the stochastic problem as a certain deterministic problem over (exponential many) all realizations (each being a point set). Our framework constructs an object $\mathcal{S}$ satisfying the following properties.

1. Basically, $\mathcal{S}$ has a constant size description (the constants may depend on d, ε, and k).

2. The objective value for a certain deterministic optimization problem over $\mathcal{S}$ can approximate the objective for the original stochastic problem well. Moreover, the solution to the deterministic optimization over $\mathcal{S}$ is a good approximation for the original problem as well.

At a high level, $\mathcal{S}$ serves very similar roles as the coresets in the deterministic setting. Note that the form of $\mathcal{S}$ may vary for different problems: in stochastic k-center, it is a collection of weighted point sets (we call $\mathcal{S}$ an SKC-Coreset); in stochastic j-flat-center, it is a combination of two collections of weighted point sets for two intermediate problems (we call $\mathcal{S}$ an SJFC-Coreset).

For stochastic k-center under the existential model, we construct an SKC-Coreset $\mathcal{S}$ in two steps. First, we map all realizations to their additive ε-coresets (for deterministic k-centers) [11]. Since there are only a polynomial number of possible additive ε-coresets, the above mapping can partition the space of all realizations into a polynomial number of parts, such that the realizations in each part have very similar objective functions. Moreover, for each additive ε-coresets, it is possible to compute the total probability of the realizations that are mapped to the coreset. In fact, this requires a subtle modification of the construction in [11] so that we can compute the aforementioned probability efficiently. This step has reduced the exponential number of realizations to a polynomial size representation. Next, we define a generalized shape fitting problem, call the generalized k-median problem, over the collection of above additive ε-coresets. Then, we need to properly generalize the previous definition of coreset and the total sensitivity (a notion proposed in the deterministic coreset context by Langberg and Schulman [77]), and prove a constant upper bound for the

generalized total sensitivity by relating it to the total sensitivity of the ordinary k-median problem. The SKC-Coreset $\mathcal{S}$ is a generalized coreset for the generalized k-median problem, which consists of a constant number of weighted point sets.

For stochastic k-center under the locational model, computing the weight for each set in the SKC-Coreset $\mathcal{S}$ is somewhat more complicated. We need to reduce the computational problem to a family of bipartite holant problems, and apply the celebrated result by Jerrum, Sinclair, and Vigoda [67].

For the stochastic minimum j-flat-center problem, we proposed an e□cient algorithm for constructing an SJFC-Coreset. We utilize several ideas in Chapter 4, as well as prior results on the shape fitting problem. We first partition the realizations $P \sim \mathcal{P}$ into two parts through a construction similar to the (ε, τ)-quant-kernel construction in Chapter 4. Roughly speaking, after linearization, we need to find a convex set $\mathcal{K}$ in a higher dimensional space such that the total probability of any point falling outside $\mathcal{K}$ is small, but not so small such that in each direction the expected directional width of $\mathcal{P}$ is comparable to that of $\mathcal{K}$. Then, for those points inside $\mathcal{K}$, it is possible to use a slight modification of the construction in Chapter 4 to construct a collection of weighted point sets. For the points outside $\mathcal{K}$, since the total probability is small, we reduce the problem to a weighted j-flat-median problem, and use the coreset in [106] (this step is similar to that in [88]). By combining the two collections, we obtain the SJFC-Coreset $\mathcal{S}$ for the problem, which is of constant size. Then, we can easily obtain a PTAS by solving a constant size polynomial system defined by $\mathcal{S}$.

We remark that our overall approach is very di□erent from that in Munteanu et al. [88] (except one aforementioned step and that they also crucially used some machinery from the coreset literature). Munteanu et al. [88] defined a near-metric distance measure $m(A, B) = \max_{a \in A, b \in B} d(a, b)$ for two non-empty point sets A, B. This near-metric measure satisfies many metric properties, like non-negativity, symmetry and the triangle inequality. By lifting the problem to the space defined by such metric and utilizing a previous coreset result for clustering, they obtained a PTAS for the problem. However, in the more general stochastic minimum k-center problem and

stochastic minimum j-flat-center problem, it is unclear how to translate the distance function between point sets and k-centers or point sets and j-flat sets to a near-metric distance (and still satisfies symmetry and triangle inequality).

1.3 Estimation for Combinatorial Optimization Problems over Stochastic Data

We are interested in the following natural problem over both existential and locational uncertainty models: estimating the expected values of certain statistics of combinatorial objects. In this dissertation, we study several combinatorial or geometry problems in these two models: the closest pair problem, minimum spanning tree, minimum perfect matching (assuming an even number of nodes), k-clustering and minimum cycle cover. We take the minimum spanning tree problem for example. Let MST be the length of the minimum spanning tree (which is a random variable) and MST(r) be the length of the minimum spanning tree spanning all points in the realization r. We would like to estimate the following quantity:

$$\mathbb{E}[\mathrm{MST}] = \sum_{r \in R} \Pr[r] \cdot \mathrm{MST}(r).$$

However, the above formula does not give us an efficient way to estimate the expectation since it involves an exponential number of terms. In fact, computing the exact expected value are either NP-hard or #P-hard. Following many of the theoretical computer science literatures on approximate counting and estimation, our goal is to obtain fully polynomial randomized approximation schemes for computing the expected values.

Our contribution. (Chapter 6) We recall that a fully polynomial randomized approximation scheme (FPRAS) for a problem f is a randomized algorithm A that takes an input instance x, a real number $\varepsilon > 0$, returns $A(x)$ such that $\Pr[(1-\varepsilon)f(x) \leq A(x) \leq (1+\varepsilon)f(x)] \geq \frac{3}{4}$ and its running time is polynomial in both the size of the

input n and $1/\varepsilon$. Our main contributions can be summarized in Table 1.1.

Problems		Existential	Locational
Closest Pair (§6.1)	$\mathbb{E}$[C]	FPRAS	FPRAS
	Pr[C $\leq$ 1]	FPRAS	FPRAS
	Pr[C $\geq$ 1]	Inapprox	Inapprox
Diameter (§6.1)	$\mathbb{E}$[D]	FPRAS	FPRAS
	Pr[D $\leq$ 1]	Inapprox	Inapprox
	Pr[D $\geq$ 1]	FPRAS	FPRAS
Minimum Spanning Tree (§6.3)	$\mathbb{E}$[MST]	FPRAS[71]	FPRAS
k-Clustering (§6.2)	$\mathbb{E}$[kCL]	FPRAS	Open
Perfect Matching (§6.4)	$\mathbb{E}$[PM]	N.A.	FPRAS
kth Closest Pair (§6.8.1)	$\mathbb{E}$[kC]	FPRAS	Open
Cycle Cover (§6.5)	$\mathbb{E}$[CC]	FPRAS	FPRAS
kth Longest m-Nearest Neighbor (§6.6)	$\mathbb{E}$[kmNN]	FPRAS	Open

Table 1.1: Our results for some problems in different stochastic models.

1. Closest Pair: We use C to denote the minimum distance of any pair of two nodes. If a realization has less than two nodes, C is zero. Computing Pr[C $\leq$ 1] exactly in the existential model is known to be #P-hard even in an Euclidean plane [72], but no nontrivial algorithmic result is known before. So is computing Pr[C $\geq$ 1]. In fact, it is not hard to show that computing Pr[C $\geq$ 1] is inapproximable within any factor in a metric space (Section 6.8.2).

   We also consider the problem of computing expected distance $\mathbb{E}$[C] between the closest pair in the same model. We prove that the problem is #P-hard in Section 6.8.2 and give the first known FPRAS in Section 6.1. Note that an FPRAS for computing Pr[C $\leq$ 1] does not imply an FPRAS for computing $\mathbb{E}$[C] [2].

2. Diameter: The problem of computing the expected length of the diameter can be reduced to the closest pair problem as follows. Assume that the longest distance between two points in $\mathcal{P}$ is W. We construct the new instance $\mathcal{P}'$ as follows: for any two points $s, t \in \mathcal{P}$, let their distance be $2W - \mathrm{d}(s,t)$ in $\mathcal{P}'$.

[2] To the contrary, an FPRAS for computing Pr[C $\geq$ 1] or Pr[C = 1] would imply an FPRAS for computing $\mathbb{E}$[C] since $\mathbb{E}[C] = \int_{(s_i,s_j)} \Pr[C = d(s_i,s_j)]d(s_i,s_j) = \int \Pr[C \geq t]dt = \int_{(s_i,s_j)} \Pr[C \geq d(s_i,s_j)](d(s_i,s_j) - d(s'_i,s'_j))$.

The new instance is still a metric. The sum of the distance of closest pair in $\mathcal{P}$ and the diameter in $\mathcal{P}'$ is exactly $2W$ (if there are at least two realized points). Hence, the answer for the diameter can be easily derived from the answer for closest pair in $\mathcal{P}'$.

3. Minimum Spanning Tree: Computing $\mathbb{E}$[MST] exactly in both uncertainty models is known to be #P-hard [71]. Kamousi, Chan, and Suri [71] developed an FPRAS for estimating $\mathbb{E}$[MST] in the existential uncertainty model and a constant factor approximation algorithm in the locational uncertainty model.

   Estimating $\mathbb{E}$[MST] is amendable to several techniques. We obtain an FPRAS for estimating $\mathbb{E}$[MST] in the locational uncertainty model using the stoch-core techinque in Section 6.3. In fact, the idea in [71] can also be extended to give an alternative FPRAS (Section 6.9). It is not clear how to extend their idea to other problems.

4. Clustering (k-clustering): In the deterministic k-clustering problem, we want to partition all points into k disjoint subsets such that the spacing of the partition is maximized, where the spacing is defined to be the minimum of any d(u, v) with u, v in di□erent subsets [74]. In fact, the optimal cost of the problem is the length of the $(k-1)$th most expensive edge in the minimum spanning tree [74]. We show how to estimate $\mathbb{E}$[kCL] using the HPF (hierarchical partition family) technique in Section 6.2.

5. Perfect Matching: We assume that there are even number of nodes to ensure that a perfect matching always exists. Therefore, only the locational uncertainty model is relevant here. We give the first FPRAS for approximating the expected length of minimum perfect matching in Section 6.4 using a more complicated stoch-core technique.

All of our algorithms run in polynomial time. However, we have not attempted to optimize the exact running time.

Our techniques. Perhaps the simplest and the most commonly used technique for estimating the expectation of a random variable is the Monte Carlo method, that is to use the sample average as the estimate. However, the method is only e□cient (i.e., runs in polynomial time) if the variance of the random variable is small (See Lemma 14). To circumvent the di□culty caused by the high variance, a general methodology is to decompose the expectation of the random variable into a convex combination of conditional expectations using the law of total expectation: $\mathbb{E}[X] = \mathbb{E}_Y[\mathbb{E}[X \mid Y]] = \sum_y \Pr[Y = y]\,\mathbb{E}[X \mid Y = y]$. Hopefully, $\Pr[Y = y]$ can be estimated (or calculated exactly) e□ciently, and the random variable X conditioning on each event y has a low variance. However, choosing the events Y to condition on can be tricky.

We develop two new techniques for choosing such events, each being capable of solving a subset of aforementioned problems. In the first technique, we first identify a set $\mathbb{H}$ of points, called the stoch-core of the problem, such that (1): with high probability, all nodes realize in $\mathbb{H}$ and (2): conditioning on event (1), the variance is small. Then, we choose Y to be the number of nodes realized to points not in $\mathbb{H}$. We compute the $(1 \pm \varepsilon)$-estimates for $Y = 0, 1$ using Monte Carlo by (1) and (2). The problematic part is when Y is large, i.e., many nodes realize to points outside $\mathbb{H}$. Even though the probability of such events is very small, the value of X under such events may be considerably large, thus contributing nontrivially. However, we can show that the contribution of such events is dominated by the first few events and thus can be safely ignored. Choosing appropriate stoch-core is easy for some problems, such as closest pair and minimum spanning tree, while it may require additional idea for other problems such as minimum perfect matching.

Our second technique utilizes a notion called Hierarchical Partition Family (HPF). The HPF has n levels, each representing a clustering of all points. For a combinatorial problem, for which the solution is a set of edges, we define Y to be the highest level such that some edge in the solution is an inter-cluster edge. Informally, conditioning on the information of Y, we can essentially bound the variance of X (hence use the Monte Carlo method). To implement Monte Carlo, we need to be able to take samples

e□ciently conditioning on Y. We show that such sampling problems can be reduced to, or have connections to, classical approximate counting and sampling problems, such as approximating permanent, counting knapsack.

Chapter 2 Preliminaries

In this chapter, we recall a useful tool, the Chernoff bound. Suppose we want to estimate $\mathbb{E}[X]$. In each Monte Carlo iteration, we take a sample (a realization of all nodes), and compute the value of X for the sample. At the end, we output the average over all samples. The number of samples required by this algorithm is suggested by the following standard Chernoff bound.

Lemma 14. (Chernoff Bound) Let random variables $X_1, X_2, \ldots, X_N$ be independent random variables taking on values between 0 and U. Let $X = \frac{1}{N}\sum_{i=1}^{N} X_i$ and μ be the expectation of X, for any $\varepsilon > 0$,

$$\Pr[X \in [(1-\varepsilon)\mu, (1+\varepsilon)\mu]] \geq 1 - 2e^{-N\frac{\mu}{U}\varepsilon^2/4}.$$

Therefore, for any $\varepsilon > 0$, in order to get an $(1\pm\varepsilon)$-approximation with probability $1 - \frac{1}{\mathrm{poly}(n)}$, the number of samples needs to be $O(\frac{U}{\mu\varepsilon^2}\log n)$. If $\frac{U}{\mu}$, the ratio between the maximum possible value of X and the expected value $\mathbb{E}[X]$, is bounded by a polynomial size of the input, we can use the above Monte Carlo method to estimate $\mathbb{E}[X]$ with a polynomial number of samples.

Chapter 3 Related Work

In this chapter, we will show some prior work about our research problems. We start with introducing some related work about shape fitting problems in both deterministic and stochastic settings. Next, we discuss related work for coreset construction. Then we review some prior work about computing the expected value of combinatorial optimization problems in different stochastic models. We also briefly mention some other stochastic models that is conceptually or technically related to the dissertation.

Shape fitting problem. A number of theoretical results for the k-center problem have been obtained in the past. In deterministic settings, Agarwal and Procopiuc [11] considered the k-center problem and showed that there exists an additive coreset of a constant size which can represent the whole input point set if k and d are both constants. Har-Peled and Varadarajan [61] improved their result in high dimensions and gave an PTAS if k is a constant. Cormode and McGregor [30] considered the k-center problem for the locational model in a finite metric graph, and obtained a bi-criterion constant approximation. Guha and Munagala [53] improved the result to a true constant factor approximation. Munteanu et al. [88] studied the minimum enclosing ball problem (a.k.a. 1-center problem) for stochastic points in fixed dimensional Euclidean space and gave a PTAS. Coresets were also constructed for imprecise points [85] to help derive results for approximating convex hulls and a variety of other shape-fitting problems. Note that their model is different from the existential or locational models.

Other projective clustering problems have also been studied extensively. If k is a constant, the existence of an additive coreset of a constant size for k-line-center, i.e., for the problem of covering P by k congruent cylinders of the minimum radius, was first proved by Agarwal et al. [11]. Har-Peled and Varadarajan [61] obtained an PTAS for k j-flat-center problem while j and k are both constants. Langberg

and Schulman [77] showed that for the weighted k-median/k-means problem, [1] there exists an ε-coreset of size depending polynomially on d and k by bounding the total sensitivity. Varadarajan and Xiao [106] studied the k-line clustering problem and the (j, k) integer projective clustering problem, and showed that there exists an ε-coreset of a poly-logarithmic size.

Coreset construction. There is a large body of literature [93] on constructing coresets for various problems, such as shape fitting [7, 8], shape fitting with outliers [62], clustering [27, 45, 47, 60, 77], integrals [77], matrix approximation and regression [34, 45] and in di□erent settings, such as geometric data streaming [8, 26] and privacy setting [43]. We have introduced some results for shape fitting problems. In this part, we review other applications of coreset construction. For example, Har-Peled and Wang [62] provided a coreset construction approach for handling outliers. From the dual (function extent) perspective, they want to approximate the distance between two level sets in an arrangement of hyperplanes. In the locational model, coresets are created for range counting queries [1] under the subset constraint, but these techniques do not translate to this setting 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.

Estimation for stochastic combinatorial optimization. Several geometric properties of a set of stochastic points have been studied extensively in the literature under the term stochastic geometry. For instance, Bearwood et al. [21] showed that if there are n points uniformly and independently distributed in $[0, 1]^2$, the minimal traveling salesman tour visiting them has an expected length $\Omega(\sqrt{n})$. Asymptotic results for minimum spanning trees and minimum matchings on n points uniformly distributed in unit balls are established by Bertsimas and van Ryzin [23]. Similar results can be found in e.g., [22, 73, 98]. Compared with results in stochastic geometry, we focus on the e□cient computation of the statistics, instead of giving explicit mathematical formulas.

[1] The k-median/k-means problem in the existential uncertainty model can be considered as a weighted k-median/k-means problem.

Recently, a number of researchers have begun to explore geometric computing under uncertainty and many classical computational geometry problems have been studied in different stochastic/uncertainty models. Agarwal, Cheng, Tao and Yi [3] studied the problem of indexing probabilistic points with continuous distributions for range queries on a line. Agarwal, Efrat, Sankararaman, and Zhang [5] also studied the same problem in the locational uncertainty model under Euclidean metric. The most probable k-nearest neighbor problem and its variants have attracted a lot of attentions in the database community (See e.g., [28]). Several other problems have also been considered recently, such as computing the expected volume of a set of probabilistic rectangles in a Euclidean space [109], convex hulls [6], skylines (Pareto curves) over probabilistic points [2, 15], and shape fitting [83].

Kamousi, Chan and Suri [71] initiated the study of estimating the expected length of combinatorial objects in this model. They showed that computing the expected length of the nearest neighbor (NN) graph, the Gabriel graph (GG), the relative neighborhood graph (RNG), and the Delaunay triangulation (DT) can be solved exactly in polynomial time, while computing $\mathbb{E}$[MST] is #P-hard and there exists a simple FPRAS for approximating $\mathbb{E}$[MST] in the existential model. They also gave a deterministic PTAS for approximating $\mathbb{E}$[MST] in an Euclidean plane. In another paper [72], they studied the closest pair and (approximate) nearest neighbor problems (i.e., finding the point with the smallest expected distance from the query point) in the same model.

The computational/algorithmic aspects of stochastic geometry have also gained a lot of attention in recent years from the area of wireless networking. In many application scenarios, it is common to assume that the nodes (e.g., sensors) are deployed randomly across a certain area, thereby forming a stochastic network. It is of central importance to study various properties in this network, such as connectivity [55], transmission capacity [56]. We refer the interested reader to a recent survey [57] for more references.

Other stochastic models. Besides the stochastic geometry models, geometric uncertain data has also been studied in the imprecise model [16, 64, 76, 84, 89, 92, 104].

In this model, each point is provided with a region where it might be. This originated with the study of imprecision in data representation [54, 94], and can be used to provide upper and lower bounds on several geometric constructs such as the diameter, convex hull, and flow on terrains [36, 104].

Convex hulls have been studied for uncertain points: upper and lower bounds are provided under the imprecise model [41, 85, 89, 104], distributions of circumference and volume are calculated in the locational model [69, 83], 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 [101], and the probability a query point is inside the convex hull [6]. As far as we know, the expected complexity of the convex hull under uncertain points has not been studied, although it has been studied [59] under other random data models.

The randomly weighted graph model where the edge weights are independent non-negative variables has also been studied extensively. Frieze [48] and Steele [99] showed that the expected value of the minimum spanning tree on such a graph with identically and independently distributed edges is $\zeta(3)/D$ where $\zeta(3) = \sum_{j=1}^{\infty} 1/j^3$ and D is the derivative of the distribution at 0. Alexopoulos and Jacobson [13] developed algorithms that compute the distribution of MST and the probability that a particular edge belongs to MST when edge lengths follow discrete distributions. However, the running time of their algorithms may be exponential in the worst cases. Recently, Emek, Korman and Shavitt [40] showed that computing the kth moment of a class of properties, including the diameter, radius and minimum spanning tree, admits an FPRAS for each fixed k.

Chapter 4 ε-Kernel Coresets over Stochastic Data

In this chapter, we initiate the study of constructing ε-kernel coresets for uncertain points. 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. Then 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 stochastic shape fitting problems. We first briefly introduce how to construct an ε-kernel coreset over deterministic point sets.

4.1 ε-Kernel Coresets over Deterministic Data

For a set P of n deterministic points, recall that we define $\mathsf{f}(P, \vec{u})$ to be $\mathsf{f}(P, \vec{u}) = \max_{s \in P} \langle \vec{u}, s \rangle$ for any direction $\vec{u} \in \mathbb{R}^d$ and $\omega(P, \vec{u})$ to be $\omega(P, \vec{u}) = \mathsf{f}(P, \vec{u}) + \mathsf{f}(P, -\vec{u})$. Also recall that a subset $Q \subseteq P$ is called an ε-kernel of P if for any direction $\vec{u} \in \mathbb{R}^d$, $(1-\varepsilon)\omega(P, \vec{u}) \leq \omega(Q, \vec{u}) \leq \omega(P, \vec{u})$. In fact, there exists an ε-kernel of size $O(\varepsilon^{-(d-1)/2})$ with construction time $O(n + \varepsilon^{-(d-3/2)})$, see [26, 110]. In the following, we briefly review the construction.

By Barequet and Har-Peled [19], we first compute a bounding box B in linear time, which has volume at most constant times of the minimum one. By their construction, we also have the property that $\alpha B \subseteq \mathsf{ConvH}(B) \subseteq B$. Here, $\mathsf{ConvH}(B)$ is the convex hull of B and $\alpha > 0$ is some constant depending on d. By applying an affine transformation, we assume that $B = [-1, 1]^d$ without loss of generality. This is because if $M(S)$ is an ε-kernel of $M(P)$ for a non-singular matrix M, then S must be an ε-kernel of P. Since $\alpha B \subseteq \mathsf{ConvH}(B)$, we have $\omega(P, \vec{u}) \geq 2\alpha$ for any direction $\vec{u}$.

Let $\varepsilon' = \sqrt{\varepsilon\alpha}$. We first construct a set $\mathcal{J}$ of size $O(\varepsilon'^{-d+1}) = O(\varepsilon^{-(d-1)/2})$ belonging to the sphere of radius $\sqrt{d}$ + 1 centered at the origin, satisfying that for any point x on this sphere, there exists a point $y \in \mathcal{J}$ such that $\|x - y\| \leq \varepsilon'$. Here, $\|x - y\|$ is the Euclidean distance between x and y. Next, for each point $y \in \mathcal{J}$, we compute a point $\phi(y) \in P$ which minimizes the distance $\|y - \phi(y)\|$. [1] The output is the collection S of all such $\phi(y)$, which is an ε-kernel of P. Note that the size of S is at most $O(\varepsilon^{-(d-1)/2})$.

We then briefly prove the output is an ε-kernel. Fix a direction $\vec{u}$ and let $s \in P$ be the point maximizing $\langle \vec{u}, s \rangle$. Suppose the ray starting from s in direction $\vec{u}$ intersects the sphere at point x. By the construction of $\mathcal{J}$, there exists a point $y \in \mathcal{J}$ with $\|x - y\| \leq \varepsilon'$. We then discuss the following two cases.

1. If $\phi(y) = s$, then $s \in S$ and $\omega(P, \vec{u}) = \omega(S, \vec{u})$.

2. If $\phi(y) \neq s$, we construct a ball of radius $\|y - s\|$ centered at y. Since $\|y - \phi(y)\| \leq \|y - s\|$ by the above algorithm, we have that $\phi(y)$ must locate inside this ball. Assume z is the point minimizing $\langle \vec{u}, z' \rangle$ over all points z' in this ball. By previous work [26, 88], it can be shown that $\langle \vec{u}, s \rangle - \langle \vec{u}, z \rangle \leq \alpha\varepsilon$. Thus, we have

$$\omega(P, \vec{u}) - \omega(S, \vec{u}) \leq \langle \vec{u}, s \rangle - \langle \vec{u}, \phi(y) \rangle \leq \langle \vec{u}, s \rangle - \langle \vec{u}, z \rangle \leq \alpha\varepsilon.$$

The above two cases imply that $\omega(P, \vec{u}) - \omega(S, \vec{u}) \leq 2\alpha\varepsilon \leq \varepsilon\omega(P, \vec{u})$. So the output S is indeed an ε-kernel of P.

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 $(1 - \varepsilon)\omega(\mathcal{P}, \vec{u}) \leq \omega(\mathcal{S}, \vec{u}) \leq (1 + \varepsilon)\omega(\mathcal{P}, \vec{u})$. Assuming the coordinates of the input points are bounded, an (ε, τ)-quant-kernel ensures that for any choice of $\vec{u}$, the cumulative distribution function of $\omega(\mathcal{S}, \vec{u})$ is within a distance ε under the Lévy

[1] In fact, we only need to compute an ε-approximate nearest-neighbor $\phi(y) \in P$ of y, which improves the running time.

metric, to that of $\omega(\mathcal{P}, \vec{u})$. [2].

4.2 ε-Kernels for Expectations of Width

First recall the definition of ε-exp-kernel. 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 $\vec{u}$ to be $\omega(\mathcal{P}, \vec{u}) = \mathbb{E}_{P\sim\mathcal{P}}[\omega(P, \vec{u})]$, where $P \sim \mathcal{P}$ means that P is a (random) realization of $\mathcal{P}$.

Definition 15. 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 $\vec{u} \in \mathbb{R}^d$,

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

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 appendix.

For simplicity of exposition, we assume in this section that all points in $\mathcal{P}$ are in general positions and all p_is are strictly between 0 and 1. For any $s, s' \in \mathbb{R}^d$, we use $\langle s, s'\rangle$ to denote the usual inner product $\sum_{i=1}^{d} s_i s'_i$. For ease of notation, we write $s \succ_s s'$ as a shorthand notation for $\langle s, \hat{s}\rangle > \langle s', \hat{s}\rangle$. For any $s' \in \mathbb{R}^d$, the binary relation $\succ_{s'}$ 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 s. For any two points s and s', we use $d(s, s')$ or $\|s - s'\|$ 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 $\vec{u} \in \mathbb{R}^d$, the support function is $f(P, \vec{u}) = \max_{s\in P} \langle \vec{u}, s\rangle$ and the directional width is $\omega(P, \vec{u}) = f(P, \vec{u}) - f(P, -\vec{u})$. The

[2] Assuming the coordinates of the input points are bounded, the requirement for an (ε, τ)-quant-kernel is in fact stronger than that of Lévy distance being no larger than ε as the former requires a multiplicative error on length, which gives better guarantee when the length is small.

support function and the directional width only depend on the convex hull of P.

Lemma 16. 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 $\omega(\mathcal{P}, \vec{u}) = \omega(S, \vec{u})$ for all $\vec{u} \in \mathbb{R}^d$.

Proof. By the definition of the expected directional width of $\mathcal{P}$, we have that

$$\omega(\mathcal{P}, \vec{u}) = \mathbb{E}_{P\sim\mathcal{P}}[\omega(P, \vec{u})] = \sum_{P\sim\mathcal{P}} \Pr[P] \left[f(P, \vec{u}) + f(P, -\vec{u}) \right].$$

Consider the Minkowski sum $M = M(\mathcal{P}) := \sum_{P\sim\mathcal{P}} \Pr[P]\mathrm{ConvH}(P)$, where $\mathrm{ConvH}(P)$ 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 $\vec{u} \in \mathbb{R}^d$ (see e.g., [95]) $f(M, \vec{u}) = \sum_{P\sim\mathcal{P}} \Pr[P] f(P, \vec{u})$. Hence, $\omega(\mathcal{P}, \vec{u}) = \omega(M, \vec{u})$ for all $\vec{u} \in \mathbb{R}^d$. □

By the result in [7], we know that for any convex body in $\mathbb{R}^d$, there exists an ε-kernel of size $O(\varepsilon^{-(d-1)/2})$. Combining with Lemma 16, we can immediately obtain the following corollary, which is the first half of Theorem 2.

Corollary 17. For any $\varepsilon > 0$, there exists an ε-exp-kernel of size $O(\varepsilon^{-(d-1)/2})$.

Recall that in Lemma 16, the Minkowski sum $M = \sum_{P\sim\mathcal{P}} \Pr[P]\mathrm{ConvH}(P)$. 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 (s, s') of points in $\mathcal{P}$ consider the hyperplane $H_{s,s'}$ that passes through the origin and is orthogonal to the line connecting s and s'. 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 directions $\vec{u} \in \mathbb{R}^d$ in each halfspace, the order of $\langle s, \vec{u}\rangle$ and $\langle s', \vec{u}\rangle$ is the same (i.e., we

have $s \succ_u s'$ in one halfspace and $s' \succ_u s$ in the other). Those hyperplanes in Γ pass through the origin and thus partition $\mathbb{R}^d$ into d-dimensional polyhedral cones. [3] We denote this arrangement as $\mathsf{A}(\Gamma)$.

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

Lemma 18. Assuming the existential model and $p_i \in (0,1)$ for all $s \in \mathcal{P}$, the complexity of M is the same as the cardinality of $\mathsf{A}(\Gamma)$, i.e., $|M| = |\mathsf{A}(\Gamma)|$. Moreover, each cone $C \in \mathsf{A}(\Gamma)$ corresponds to exactly one vertex s of $\mathsf{ConvH}(M)$ in the following sense: the gradient $\nabla \mathsf{f}(M, \vec{u}) = s$ for all $\vec{u} \in \operatorname{int} C$ (note that here s should be understood as a vector).

Proof. We have shown that M is a convex polytope. We first note that the support function uniquely defines a convex body (see e.g., [95]). We need the following well known fact in convex geometry (see e.g., [49]): 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_s corresponds to a vertex s of M, and for each vector $\vec{u} \in C_s$, it holds $\mathsf{f}(M, \vec{u}) = \langle \vec{u}, s \rangle$ (i.e., the maximum of $\mathsf{f}(M, \vec{u}) = \max_{s' \in M} \langle \vec{u}, s' \rangle$ is achieved by s for all $\vec{u} \in C_s$).[4] See Figure 4-1 for an example in $\mathbb{R}^2$. Hence, for each $\vec{u} \in \operatorname{int} C_s$ the gradient of of the support function (as a function of $\vec{u}$) is exactly s:

$$\nabla \mathsf{f}(M, \vec{u}) = \left[\frac{\partial \mathsf{f}(M, \vec{u})}{\partial \vec{u}_j} \right]_{j \in [d]} = \left[\frac{\partial \langle \vec{u}, s \rangle}{\partial \vec{u}_j} \right]_{j \in [d]} = \left[\frac{\partial \sum_{j \in [d]} s_j \vec{u}_j}{\partial \vec{u}_j} \right]_{j \in [d]} = s, \quad (4.1)$$

where $\vec{u}_j$ is the jth coordinate of $\vec{u}$. With a bit abuse of notation, we denote the set of cones defined above by $\mathsf{A}(M)$.

Now, consider a cone $C \in \mathsf{A}(\Gamma)$. We show that for all $\vec{u} \in \operatorname{int} C$, $\nabla \mathsf{f}(M, \vec{u})$ is a distinct constant vector independent of $\vec{u}$. In fact, we know that $\mathsf{f}(M, \vec{u}) = \mathsf{f}(\mathcal{P}, \vec{u}) =$

[3] We ignore the lower dimensional cells in the arrangement.

[4] One intuitive way to see this is as follows: 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. In fact, we call such a cone C_s the outer normal cones.

$\sum_{s\in\mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(s,\vec{u})\langle s,\vec{u}\rangle$, where $\mathsf{Pr}^{\mathsf{R}}(s,\vec{u}) = \prod_{s'\succ_{\vec{u}} s}(1-p_{s'})p_s$. For all $\vec{u}\in \mathsf{int}\, C$, the $\mathsf{Pr}^{\mathsf{R}}(s,\vec{u})$ value is the same since the value only depends on the canonical order with respect to $\vec{u}$, which is the same for all $\vec{u}\in C$. Hence, we can get that for all $\vec{u}\in \mathsf{int}\, C$,

$$\nabla \mathsf{f}(M,\vec{u}) = \sum_{s\in\mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(s,\vec{u})s, \tag{4.2}$$

which is a constant independent of $\vec{u}$. We prove the lemma by showing that the gradient $\nabla \mathsf{f}(M,\vec{u})$ must be different for two adjacent cones C_1, C_2 (separated by some hyperplane in Γ) in $\mathsf{A}(\Gamma)$. Suppose $\vec{u}_1\in \mathsf{int}\, C_1$ and $\vec{u}_2\in \mathsf{int}\, C_2$. Consider the canonical orders O_1 and O_2 of $\mathcal{P}$ with respect to $\vec{u}_1$ and $\vec{u}_2$ respectively. Since C_1 and C_2 are adjacently, O_1 and O_2 only differ by one swap of adjacent vertices. W.l.o.g., assume that $O_1 = \{s_1,\ldots,s_i,s_{i+1},\ldots,s_n\}$ and $O_1 = \{s_1,\ldots,s_{i+1},s_i,\ldots,s_n\}$. Using (4.2), we can get that

$$\begin{aligned}\nabla \mathsf{f}(M,\vec{u}_1) - \nabla \mathsf{f}(M,\vec{u}_2) &= \mathsf{Pr}^{\mathsf{R}}(s_i,\vec{u}_1)s_i + \mathsf{Pr}^{\mathsf{R}}(s_{i+1},\vec{u}_1)s_{i+1} - \mathsf{Pr}^{\mathsf{R}}(s_i,\vec{u}_2)s_i - \mathsf{Pr}^{\mathsf{R}}(s_{i+1},\vec{u}_2)s_{i+1}\\ &= D\cdot(p_i s_i + (1-p_i)p_{i+1}s_{i+1} - p_{i+1}s_{i+1} - (1-p_{i+1})p_i s_i)\\ &= D\cdot p_i p_{i+1}(s_i - s_{i+1}) \neq 0\end{aligned}$$

where $D = \prod_{j=1}^{i-1}(1-p_j) \neq 0$.

In summary, we have shown in the first paragraph that $\nabla \mathsf{f}(M,\vec{u})$ is piecewise constant, with a distinct constant in each cone in $\mathsf{A}(M)$. The same also holds for $\mathsf{A}(\Gamma)$. This is only possible if $\mathsf{A}(\Gamma)$ (thinking as a partition of $\mathbb{R}^d$) partitions $\mathbb{R}^d$ exactly the same way as $\mathsf{A}(M)$ does. Hence, we have $\mathsf{A}(\Gamma) = \mathsf{A}(M)$ 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., [12]), we immediately obtain the following corollary.

Corollary 19. It holds that $|M| \leq O(\binom{n^2}{d-1}) = O(n^{2d-2})$.

The proof of Lemma 18 can be easily made constructive. We only need to compute

Note that $\mathcal{K}_{i-1} \subseteq \mathcal{H}_i$. Compute a deterministic ε_2-kernel $C_{\mathcal{H}_i}$ of $\mathcal{H}_i$ and let $\mathcal{K}_i = (1+\varepsilon_1)\mathrm{ConvH}(C_{\mathcal{H}_i})$ (hence $\mathrm{ConvH}(C_{\mathcal{H}_i}) \subseteq \mathcal{H}_i \subseteq \mathcal{K}_i$). Then, we delete all points in $\mathcal{P} \cap \mathcal{K}_i$ and add all vertices of $\mathcal{K}_i$ (denoted as $\mathrm{Ver}(\mathcal{K}_i)$). Let $\mathcal{P}_{i+1}$ be the remaining points in $\mathcal{P} \cap \overline{\mathcal{K}}_i$.

Our algorithm terminates when $\lambda(\mathcal{P}_i) \leq \tilde{\Theta}(1/\sqrt{\varepsilon})$. Suppose the last round is z. Finally, we let $\mathcal{H} = \frac{1}{(1+\varepsilon_1)^z}\mathcal{H}_z$ and $\mathcal{K} = \mathcal{K}_z$.

First we show the algorithm terminates after at most a logarithmic number of rounds.

Lemma 41. $z = O(\log n)$.

Proof. If $\mathcal{H}_i = \mathrm{TK}(\mathcal{E}_i \cup \mathrm{Ver}(\mathcal{K}_{i-1}), 2\ln(2/\tau))$, then we stop since $\lambda(\mathcal{P}_{i+1}) \leq \tilde{O}(1/\sqrt{\varepsilon})$ by Lemma 37. Thus, we only need to bound the number of iterations where $\mathcal{H}_i = \mathrm{TK}(\mathcal{E}_i \cup \mathrm{Ver}(\mathcal{K}_{i-1}), 4\varepsilon_2\lambda(\mathcal{P}_i))$.

Initially, it is not hard to see that $\lambda(\mathcal{P}_1) \leq n\ln(2/\tau)$. Using Lemma 37, we can see that $\lambda(\mathcal{P}_{i+1}) = \lambda(\mathcal{P} \cap \overline{\mathcal{K}}_i) \leq O(5\varepsilon_2\lambda(\mathcal{P}_i)/\sqrt{\varepsilon_1}) \leq \lambda(\mathcal{P}_i)/2$ for the constant defining ε_1 sufficiently large. Hence, $\lambda(\mathcal{P}_i) \leq \lambda(\mathcal{P})/2^i$. □

We need to show $\mathcal{H}$ and $\mathcal{K}$ satisfy P1, P2 and P3. P3 is quite obvious by our algorithm. It is also not hard to see P1 since $(1+\varepsilon_1)^{z+1} \leq 1+\varepsilon$ and

$$\frac{1}{(1+\varepsilon)}\mathcal{K}_z \subseteq \frac{1}{(1+\varepsilon_1)^{z+1}}\mathcal{K}_z \subseteq \frac{1}{(1+\varepsilon_1)^z}\mathcal{H}_z = \mathcal{H} \subseteq \mathcal{H}_z \subseteq \mathcal{K}_z = \mathcal{K}.$$

The most difficult part is to show P2 holds: For every direction $\vec{u} \in \mathbb{S}^1$, $\lambda(H(\mathcal{H}, \vec{u})) = \sum_{s \in H(\mathcal{H}, \vec{u})} \lambda_s \geq \ln(2/\tau)$. In fact, we show that $\mathcal{H} \subseteq \mathcal{H}^{\square}$, from which P2 follows trivially, which suffices to prove the following lemma.

Lemma 42. $\frac{1}{(1+\varepsilon_1)^i}\mathcal{H}_i \subseteq \mathcal{H}^{\square}$ for all $0 \leq i \leq z$. In particular, $\mathcal{H} \subseteq \mathcal{H}^{\square}$.

Proof. We prove the lemma by induction. $\mathcal{H}_0 \subseteq \mathcal{H}^{\square}$ clearly satisfies the lemma. For ease of notation, we let $\eta = 1/(1+\varepsilon_1)$. Suppose the lemma is true for $\mathcal{H}_{i-1}$, from which we can see that

$$\eta^i\mathcal{K}_{i-1} \subseteq \eta^{i-1}\mathcal{H}_{i-1} \subseteq \mathcal{H}^{\square}.$$

Now we show the lemma holds for $\mathcal{H}_i$. Consider the ith round. Let $\mathcal{E}_i$ be an $\varepsilon_2\lambda(\mathcal{P}_i)$ -approximation for $\mathcal{P}_i = \mathcal{P} \cap \overline{\mathcal{K}}_{i-1}$ and $\mathcal{H}_i = \mathsf{TK}(\mathcal{E}_i \cup \mathsf{Ver}(\mathcal{K}_{i-1}), 4\varepsilon_2\lambda(\mathcal{P}_i))$. Fix an arbitrary direction $\vec{u} \in \mathbb{S}^1$ (w.l.o.g., assume that $u = (0, -1)$, i.e., the downward direction), let $H(\eta^i\mathcal{H}_i, \vec{u})$ be a halfplane whose boundary is tangent to $\eta^i\mathcal{H}_i$. It suffices to show that $\lambda(H(\eta^i\mathcal{H}_i, \vec{u})) = \sum_{s \in H(\eta^i\mathcal{H}_i, \vec{u})} \lambda_s \geq \ln(2/\tau)$. We move a sweep line $\ell_{\vec{u}}$ orthogonal to $\vec{u}$, along the direction $\vec{u}$ (i.e., from top to bottom), to sweep through the points in $\mathcal{P}_i \cap \mathsf{Ver}(\mathcal{K}_{i-1})$ until the total weight we have swept is at least $\ln(2/\tau)$. We distinguish two cases:

1. $\ell_{\vec{u}}$ hits a point s in $\mathsf{Ver}(\mathcal{K}_{i-1})$ (recall the weight for such point is $+\infty$). We can see that s is the topmost point of $\mathcal{K}_{i-1}$ and $\mathcal{H}_i$ (or equivalently, $\ell_{\vec{u}}$ is also a supporting line for $\mathcal{H}_i$). Since $\eta^i\mathcal{K}_{i-1} \subseteq \mathcal{H}^*$ by the induction hypothesis, the topmost point of $\eta^i\mathcal{K}_{i-1}$ is lower than that for $\mathcal{H}^*$. The topmost point of $\eta^i\mathcal{K}_{i-1}$ is also the highest point of $\eta^i\mathcal{H}_i$, from which we can see $H(\eta^i\mathcal{H}_i, \vec{u})$ is lower than $H(\mathcal{H}^*, \vec{u})$, which implies that $\lambda(H(\eta^i\mathcal{H}_i, \vec{u})) \geq \ln(2/\tau)$.

2. $\ell_{\vec{u}}$ stops moving when it hits an original point in $\mathcal{P}_i$. Since $\max\{3\varepsilon_2\lambda(\mathcal{P}_i), 2\ln(2/\tau) - \varepsilon_2\lambda(\mathcal{P}_i)\} > \ln(2/\tau)$, by definition of $\mathcal{H}_i$, $H(\mathcal{H}_i, \vec{u})$ can not be higher than $\ell_{\vec{u}}$. The boundary of $H(\eta^i\mathcal{H}_i, \vec{u})$ is even lower, from which we can see $\lambda(H(\eta^i\mathcal{H}_i, \vec{u})) \geq \ln(2/\tau)$.

Hence, every point in $\eta^i\mathcal{H}_i$ has Tukey depth at least $\ln(2/\tau)$, which implies the lemma. □

Running time. In each round, we compute in linear time an ε_2-approximation $\mathcal{E}_i$ of size $O(\varepsilon_2^{-2}\log(1/\delta)) = \mathrm{polylog}(n)$ (with $\delta = \mathrm{poly}(n)$ to ensure each probabilistic event succeeds with high probability). $\mathcal{K}_i$ is a dilation of an ε_1-kernel. So the size of $\mathsf{Ver}(\mathcal{K}_i)$ is at most $1/\sqrt{\varepsilon_1} = O(\log^{1/2} n)$. Deciding whether a point is inside $\mathcal{K}_i$ can be solved in $\mathrm{polylog}(n)$ time, by a linear program with $|\mathsf{Ver}(\mathcal{K}_i)|$ variables. To compute $\mathcal{H}_i$, we can use the brute-force algorithm described in Section 4.3.2, which takes $\mathrm{poly}(|\mathcal{E}_i \cap \mathsf{Ver}(\mathcal{K}_{i-1})|) = \mathrm{polylog}(n)$ time. There are logarithmic number of rounds. So the overall running time is $O(n\mathrm{polylog}n)$.

Higher Dimension. Our algorithm can be easily extended to $\mathbb{R}^d$ for any constant $d > 2$. In $\mathbb{R}^d$, we let $\varepsilon_1 = O(\varepsilon/\log n)$ and $\varepsilon_2 = O(\varepsilon/\log n)^{(d-1)/2}$. With the new parameters, we can easily check that Lemma 41 still holds. We can construct an (ε, τ)-quant-kernel of size $\min\{O(\tau^{-2}\max\{\lambda^2, \lambda^4\}\log(1/\delta)), O(\tau^{-2}\varepsilon^{-(d-1)}\log(1/\delta)\text{polylog}(1/\tau)\}$. The first term is from Theorem 36 and the second from the higher-dimensional extension to Theorem 38. Now, let us examine the running time. In $\mathbb{R}^d$, $|\text{Ver}(\mathcal{K}_i)|$ is at most $\varepsilon_1^{-(d-1)/2} = O(\log^{(d-1)/2} n)$. So deciding whether a point is inside $\mathcal{K}_i$ can be solved in $\log^{O(d)}(n)$ time. Computing $\mathcal{H}_i$ takes $\log^{O(d)}(n)$ time using the brute-force algorithm. So the overall running time is $O(n\log^{O(d)} n)$.

In summary, we obtain the following theorem for (ε, τ)-quant-kernel.

Theorem 5. (restated) $\mathcal{P}$ is a set of uncertain points in $\mathbb{R}^d$ with existential uncertainty. Let $\lambda = \sum_{s_i \in \mathcal{P}}(-\ln(1-p_i))$. There exists an (ε, τ)-quant-kernel for $\mathcal{P}$, which consists of a set of independent uncertain points of cardinality $\min\{\bar{O}(\tau^{-2}\max\{\lambda^2, \lambda^4\}), \bar{O}(\varepsilon^{-(d-1)}\tau^{-2})\}$. The algorithm for constructing such a coreset runs in $\bar{O}(n\log^{O(d)} n)$ time.

4.3.3 (ε, τ)-quant-kernel Under the Subset Constraint

We show it is possible to construct an (ε, τ)-quant-kernel in the existential model under the β-assumption: each possible location realizes a point with a probability at least β, where $\beta > 0$ is some fixed constant.

Theorem 43. Under the β-assumption, there is an (ε, τ)-quant-kernel in $\mathbb{R}^d$, which is of size $O(\mu^{-(d-1)/2}\log(1/\mu))$ and satisfies the subset constraint, in the existential uncertainty model, where $\mu = \min\{\varepsilon, \tau\}$.

In fact, the algorithm is exactly the same as constructing an ε-exp-kernel and the proof of the above theorem is implicit in the proof of Theorem 26.

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

In this section, we show an (ε, r)-fpow-kernel exists in the existential uncertainty model under the β-assumption. Recall that the function $T_r(P, \vec{u}) = \max_{s\in P} \langle \vec{u}, s\rangle^{1/r} - \min_{s\in P} \langle \vec{u}, s\rangle^{1/r}$. For ease of notation, we write $\mathbb{E}[T_r(\mathcal{P}, \vec{u})]$ to denote $\mathbb{E}_{P\sim\mathcal{P}}[T_r(P, \vec{u})]$. Our goal is to find a set $\mathcal{S}$ of stochastic points such that for all directions $\vec{u} \in \mathcal{P}^\square$, we have that $\mathbb{E}[T_r(\mathcal{S}, \vec{u})] \in (1 \pm \varepsilon)\mathbb{E}[T_r(\mathcal{P}, \vec{u})]$.

Our construction of $\mathcal{S}$ is almost the same as that in Section 4.3.1. Suppose we sample N (fixed later) independent realizations and take the ε_0-kernel for each of them. Suppose they are $\{\mathcal{E}_1, \ldots, \mathcal{E}_N\}$ and we associate each a probability $1/N$. We denote the resulting (ε, r)-fpow-kernel by $\mathcal{S}$. Hence, for any direction $\vec{u} \in \mathcal{P}^\square$, $\mathbb{E}[T_r(\mathcal{S}, \vec{u})] = \frac{1}{N} \sum_{i=1}^{N} T_r(\mathcal{E}_i, \vec{u})$ and we use this value as the estimation of $\mathbb{E}[T_r(\mathcal{P}, \vec{u})]$. Now, we show $\mathcal{S}$ is indeed an (ε, r)-fpow-kernel.

Recall that we use $\mathcal{E}(P)$ to denote the deterministic ε-kernel for any realization $P \sim \mathcal{P}$. We first compare $\mathcal{P}$ with the random set $\mathcal{E}(P)$.

Lemma 44. For any $t \geq 0$ and any direction $\vec{u} \in \mathcal{P}^\square$, we have that

$$(1 - \varepsilon/2)\mathbb{E}[T_r(\mathcal{P}, \vec{u})] \leq \mathbb{E}_{P\sim\mathcal{P}}[T_r(\mathcal{E}(P), \vec{u}))] \leq \mathbb{E}[T_r(\mathcal{P}, \vec{u})].$$

Proof. By Lemma 4.6 in [7], we have that $(1-\varepsilon/2)T_r(P, \vec{u}) \leq T_r(\mathcal{E}(P), \vec{u}) \leq T_r(P, \vec{u})$. The lemma follows by combining all realizations. □

Now we show that $\mathcal{S}$ is an (ε, r)-fpow-kernel of $\mathcal{E}(P)$. We first prove the following lemma. The proof is almost the same as that of Lemma 30, and can be found in Section 4.8.

Lemma 45. Let $N = O\left(\varepsilon_1^{-2}\varepsilon_0^{-(d-1)/2}\log(1/\varepsilon_0)\right)^\square$, where $\varepsilon_0 = (\varepsilon/4(r-1))^r$, $\varepsilon_1 = \varepsilon\beta^2$. For any $t \geq 0$ and any direction $\vec{u} \in \mathcal{P}^\square$, we have that

$$\Pr_{P\sim\mathcal{S}}[\max_{s\in P}\langle \vec{u}, s\rangle^{1/r} \geq t] \in \Pr_{P\sim\mathcal{P}}[\max_{s\in\mathcal{E}(P)}\langle \vec{u}, s\rangle^{1/r} \geq t)] \pm \varepsilon_1/4, \text{ and}$$

$$\Pr_{P\sim\mathcal{S}}[\min_{s\in P}\langle \vec{u}, s\rangle^{1/r} \geq t] \in \Pr_{P\sim\mathcal{P}}[\min_{s\in\mathcal{E}(P)}\langle \vec{u}, s\rangle^{1/r} \geq t)] \pm \varepsilon_1/4.$$

Lemma 46. Let $N = O\left(\beta^{-4}\varepsilon^{-(rd-r+4)/2}\log(1/\varepsilon)\right)$ and $\varepsilon_0 = (\varepsilon/4(r-1))^r$. $\mathcal{S}$ constructed above is an (ε, r)-fpow-kernel in $\mathbb{R}^d$.

Proof. Fix a direction $\vec{u} \in \mathcal{P}$. Let $A = \max_{s\in\mathcal{P}}\langle\vec{u}, s\rangle^{1/r}$, $B = \min_{s\in\mathcal{P}}\langle\vec{u}, s\rangle^{1/r}$. We observe that $B \le \max_{s\in P}\langle\vec{u}, s\rangle^{1/r} \le A$ for any realization $P \sim \mathcal{P}$. We also need the following basic fact about the expectation: For a random variable X, if $\Pr[X \ge a] = 1$, then $\mathbb{E}[X] = \int_b^\infty \Pr[X \ge x]dx + b$ for any $b \le a$. Thus, we have that

$$\begin{aligned}\mathbb{E}_{P\sim\mathcal{P}}[\max_{s\in\mathcal{E}(P)}\langle\vec{u}, s\rangle^{1/r}] &= \int_A^\infty \Pr_{P\sim\mathcal{P}}[\max_{s\in\mathcal{E}(P)}\langle\vec{u}, s\rangle^{1/r} \ge x]dx + B\\ &\le \int_B^A \Pr_{P\sim\mathcal{S}}[\max_{s\in P}\langle\vec{u}, s\rangle^{1/r} \ge x]dx + B + \varepsilon_1(A-B)/4\\ &= \mathbb{E}_{P\sim\mathcal{S}}[\max_{s\in P}\langle\vec{u}, s\rangle^{1/r}] + \varepsilon_1(A-B)/4,\end{aligned}$$

where the first inequality is due to Lemma 45. Similarly, we can show the following two inequalities:

$$\mathbb{E}_{P\sim\mathcal{S}}[\max_{s\in P}\langle\vec{u}, s\rangle^{1/r}] \in \mathbb{E}_{P\sim\mathcal{P}}[\max_{s\in\mathcal{E}(P)}\langle\vec{u}, s\rangle^{1/r}] \pm \varepsilon_1(A-B)/4,$$

$$\mathbb{E}_{P\sim\mathcal{S}}[\min_{s\in P}\langle\vec{u}, s\rangle^{1/r}] \in \mathbb{E}_{P\sim\mathcal{P}}[\min_{s\in\mathcal{E}(P)}\langle\vec{u}, s\rangle^{1/r}] \pm \varepsilon_1(A-B)/4.$$

Recall that $T_r(P, \vec{u}) = \max_{s\in P}\langle\vec{u}, s\rangle^{1/r} - \min_{s\in P}\langle\vec{u}, s\rangle^{1/r}$. By the linearity of expectation, we conclude that

$$\mathbb{E}[T_r(\mathcal{P}, \vec{u})] \in \mathbb{E}_{P\sim\mathcal{P}}[T_r(\mathcal{E}(P), \vec{u})] \pm \varepsilon_1(A-B)/2.$$

Combining Lemma 44, we have that $\mathbb{E}[T_r(\mathcal{S}, \vec{u})] \in (1\pm\varepsilon/2)\mathbb{E}[T_r(\mathcal{P}, \vec{u})] \pm \varepsilon_1(A-B)/2$. By the β-assumption, we know that $\mathbb{E}[T_r(\mathcal{P}, \vec{u})] \ge \beta^2(A-B)$. Thus, $\varepsilon_1(A-B)/2 \le \frac{\varepsilon}{2}\mathbb{E}[T_r(\mathcal{P}, \vec{u})]$, and $\mathbb{E}[T_r(\mathcal{S}, \vec{u})] \in (1\pm\varepsilon)\mathbb{E}[T_r(\mathcal{P}, \vec{u})]$. □

Running time. In each sample, the size of a deterministic ε_0-kernel $\mathcal{E}_i$ is at most $O\left(\varepsilon_0^{-(d-1)/2}\right)$. Note that constructing an ε_0-kernel can be solved in linear time. We take $O\left(\varepsilon_1^{-2}\varepsilon_0^{-(d-1)/2}\log(1/\varepsilon_0)\right)$ samples in total. So the overall running time is $O\left(n\beta^{-4}\varepsilon^{-(rd-r+4)/2}\log(1/\varepsilon) + \mathrm{poly}(1/\varepsilon)\right) = \tilde{O}\left(n\varepsilon^{-(rd-r+4)/2}\right)$.

Note that each ε_0-kernel contains $O(\varepsilon^{-r(d-1)/2})$ points. We take $N = O\left(\varepsilon_1^{-2}\varepsilon_0^{-(d-1)/2}\log(1/\varepsilon_0)\right)$ independent samples. So the total size of (ε, r)-fpow-kernel is $O(\beta^{-4}\varepsilon^{-(rd-r+2)}\log(1/\varepsilon))$.
In summary, we obtain the following theorem.

Theorem 7. (restated) An (ε, r)-fpow-kernel of size $\tilde{O}(\varepsilon^{-(rd-r+2)})$ can be constructed in $\tilde{O}\left(n\varepsilon^{-(rd-r+4)/2}\right)$ time in the existential uncertainty model under the β-assumption. In particular, the (ε, r)-fpow-kernel consists of $N = \tilde{O}(\varepsilon^{-(rd-r+4)/2})$ point sets, each occuring with probability $1/N$ and containing $O(\varepsilon^{-r(d-1)/2})$ deterministic points.

4.5 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.

4.5.1 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 $\{h_1, h_2, \ldots\}$ 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 $\mathsf{E}_{\mathcal{H}}(x) = \max_{h\in\mathcal{H}} h(x) - \min_{h\in\mathcal{H}} h(x)$ 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 $(1-\varepsilon)\mathsf{E}_{\mathcal{H}}(x) \leq \mathsf{E}_{\mathcal{S}}(x) \leq \mathsf{E}_{\mathcal{H}}(x)$ for any $x \in \mathbb{R}^d$. We say $\mathcal{S}$ is the (ε,τ)-quant-kernel for $\mathcal{H}$ if $\Pr_{S\sim\mathcal{S}}[\mathsf{E}_S(x) \leq t] \in \Pr_{H\sim\mathcal{H}}[\mathsf{E}_H(x) \leq (1\pm\varepsilon)t] \pm\phi$. 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 $y = a_1x_1 + \ldots + a_dx_d + a_{d+1}$ to the point $(a_1,\ldots,a_{d+1}) \in \mathbb{R}^{d+1}$, 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 20 and Corollary 17, we can construct a set S of $O(n^{2d})$ deterministic linear functions in $\mathbb{R}^d$, such that $\mathsf{E}_S(x) = \mathbb{E}[\mathsf{E}_{\mathcal{H}}(x)]$ 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 47. 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 $\mathsf{E}_M(x) = \mathbb{E}[\mathsf{E}_{\mathcal{H}}(x)]$ 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 $\lambda = \max_{h\in\mathcal{H}}(-\ln(1-p_h))$.

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 $\vec{u} \in \mathcal{H}$, we have $u(x) \geq 0$. For each point $x \in \mathcal{H}^*$, we let the random variable $\mathsf{E}_{r,\mathcal{H}}(x) = \max_{h\in\mathcal{H}} h(x)^{1/r} - \min_{h\in\mathcal{H}} h(x)^{1/r}$. We say another random set $\mathcal{S}$ of functions is the (ε,r)-fpow-kernel for $\mathcal{H}$ if $(1-\varepsilon)\mathsf{E}_{r,\mathcal{H}}(x) \leq \mathsf{E}_{r,\mathcal{S}}(x) \leq \mathsf{E}_{r,\mathcal{H}}(x)$ for any $x \in \mathcal{H}^*$. By the duality transformation and Theorem 7, we can obtain the following result.

Theorem 48. 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 $\tilde{\Theta}(\varepsilon^{-(rk-r+2)})$. Furthermore, the (ε, r)-fpow-kernel consists of $N = O\left(\varepsilon^{-(rk-r+4)/2}\right)$ sets, each occurring with probability $1/N$ and containing $O\left(\varepsilon^{-r(k-1)/2}\right)$ deterministic fractional polynomials.

4.5.2 Stochastic Moving Points

We can extend our stochastic models to moving points. In the existential model, each point s is present with probability p_s and follows a trajectory $s(t)$ in $\mathbb{R}^d$ when present ($s(t)$ is the position of s at time t). In the locational model, each point s is associated with a distribution of trajectories (the support size is finite) and the actual trajectory of s is a random sample for the distribution. Such uncertain trajectory models have been used in several applications in spatial databases [111]. 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 s, any direction $\vec{u}$ and time t, define the polynomial $f_s(\vec{u}, t) = \langle s(t), \vec{u} \rangle$ and let $\mathcal{H}$ include f_s with probability p_s. For a set $\mathcal{P}$ of points, the directional width at time t is $\mathsf{E}_{\mathcal{H}}(\vec{u}, t) = \max_{s \in \mathcal{P}} f_s(\vec{u}, t) - \min_{s \in \mathcal{P}} f_s(\vec{u}, t)$. Each polynomial f_s admits a linearization of dimension $k = (r+1)d - 1$. Using Theorem 47, 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 $\vec{u}$ is the same as the expected directional width of $\mathcal{P}$ in direction $\vec{u}$. 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 $\tilde{\Theta}(\varepsilon^{-k}\tau^{-2})$.

4.5.3 Shape Fitting Problems

Theorem 47 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 $\mathbb{E}[\max_{s\in\mathcal{P}} \|s-c\|^2]$ 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 s is present with probability p_s. For each point $s \in P$, define the polynomial $h_s(x) = \|x\|^2 - 2\langle x, s\rangle + \|s\|^2$, which admits a linearization of dimension $d+1$ [7]. Let $\mathcal{H}$ be the family of uncertain polynomials $\{h_s\}_{s\in\mathcal{P}}$ (h_s exists with probability p_s). We can see that for any $x \in \mathbb{R}^d$, $\max_{s\in\mathcal{P}} \|x-s\|^2 = \max_{h_s\in\mathcal{H}} h_s(x)$. Using Theorem 47, [11] we can see that there is a set M of $O(n^{2d+2})$ deterministic polynomials such that $\max_{h\in M} h(x) = \mathbb{E}[\max_{s\in\mathcal{P}} \|x-s\|^2]$ for any $x \in \mathbb{R}^d$ and a set S of $O(\varepsilon^{-(d+1)/2})$ deterministic polynomials such that $(1-\varepsilon)\mathbb{E}[\max_{s\in\mathcal{P}} \|x-s\|^2] \le \max_{h\in S} h(x) \le \mathbb{E}[\max_{s\in\mathcal{P}} \|x-s\|^2]$ 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 s, return the expected length of the furthest point from s. 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., [12]). 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 $\mathbb{E}[\text{obj}(c)] = \mathbb{E}[\max_{s\in P} \|s-c\|^2 - \min_{s\in P} \|s-c\|^2]$ 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 $(1-\varepsilon)\mathbb{E}[\text{obj}(c)] \le E_S(x) \le \mathbb{E}[\text{obj}(c)]$ for any $x \in \mathbb{R}^d$. We would like to make a few remarks here.

[11] We can see from the proof that all results that hold for width/extent also hold for support function/maximum.

1. Let us 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 $h(x) = \|x\|^2 - 2\langle x, s\rangle + \|s\|^2$, 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 $\mathbb{E}[\max_{s\in P} d(s, c)]$ and the expected shell width $\mathbb{E}[\max_{s\in P} d(s, c) - \min_{s\in P} d(s, c)]$. 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. [12] 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.

Theorem 8. (restated) 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 $\mathbb{E}[\max_{s\in\mathcal{P}} \|s - c\|^2]$; (2) finding a center point c to minimize $\mathbb{E}[\text{obj}(c)] = \mathbb{E}[\max_{s\in P} \|s - c\|^2 - \min_{s\in P} \|s - c\|^2]$. Note that when $d = 2$ the above two problems correspond to minimizing the expected areas of the enclosing ball and the enclosing annulus, respectively.

4.5.4 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 48 to obtain linear time approximation schemes for those problems.

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 $\mathbb{E}[\max_{s\in P} \|s - c\| - \min_{s\in P} \|s - c\|]$ is minimized. For each point $s \in P$, let $h_s(x) = \|x\|^2 - 2\langle x, s\rangle + \|s\|^2$, which admits a

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

linearization of dimension $d+1$. It is not hard to see that $\mathbb{E}[\max_{s\in P}\|s-c\|] = \mathbb{E}[\max_{s\in P}\sqrt{h_s(c)}]$ and $\mathbb{E}[\min_{s\in P}\|s-c\|] = \mathbb{E}[\min_{s\in P}\sqrt{h_s(c)}]$. Using Theorem 48, we can see that there are $N = \tilde{O}(\varepsilon^{-(d+3)})$ sets S_i, each containing $O(\varepsilon^{-(d+1)})$ fractional polynomial $\sqrt{h_s}$s such that for all $x\in\mathbb{R}^d$,

$$\frac{1}{N}\sum_{i\in[N]}(\max_{S_i}\sqrt{h_s(x)}-\min_{S_i}\sqrt{h_s(x)})\in(1\pm\varepsilon)(\mathbb{E}[\max_{s\in P}\|s-x\|]-\mathbb{E}[\min_{s\in P}\|s-x\|]). \tag{4.5}$$

Note that our (ε,r)-fpow-kernel satisfies the subset constraint. Hence, each function $\sqrt{h_s}$ 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 (4.5), 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 $s\in C$, 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 $s_i = \arg\max_{s\in P_i}\|s-x\|$ and $s_i' = \arg\min_{s\in P_i}\|s-x\|$. Hence, our problem can be formulated as the following optimization problem:

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

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.,[20]).

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 $d(\ell, s)$ denote the distance between a point $s \in \mathbb{R}^d$ and a line $\ell \subset \mathbb{R}^d$. The goal for the minimum enclosing cylinder problem is to find a line ℓ such that $\mathbb{E}[\max_{s\in\mathcal{P}} d(\ell, s)]$ is minimized, while that for the minimum cylindrical shell problem is to minimize $\mathbb{E}[\max_{s\in\mathcal{P}} d(\ell, s) - \min_{s\in\mathcal{P}} d(\ell, s)]$. 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 $(x_1, \ldots, x_{2d-1}) \in \mathbb{R}^{2d-1}$: $\ell = \{p + tq \mid t \in \mathbb{R}\}$, where $p = (x_1, \cdots, x_{d-1}, 0)$ is the intersection point of ℓ with the hyperplane $x_d = 0$ and $q = (x_d, \ldots, x_{2d-1})$, $\|q\|^2 = 1$ is the orientation of ℓ. Then for any point $s \in \mathbb{R}^d$, we have that

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

   where the polynomial $d^2(\ell, s)$ admits a linearization of dimension $O(d^2)$. Now, proceeding as for the minimum enclosing ball problem and using Theorem 48, we can obtain a coreset $\mathcal{S}$ consisting $N = O\left(\varepsilon^{-O(d^2)}\right)$ 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 $\|(x_d, \ldots, x_{2d-1})\| = 1$ represents a line in $\mathbb{R}^d$) into $O\left(\varepsilon^{-O(d^2)}\right)$ 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 $s_i = \arg\max_{s\in P_i} d(\ell, s_i)$ for all $i \in [N]$, where ℓ is an arbitrary line in C. We can

formulate the problem as the following polynomial system:

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

Again the polynomial system has a constant number of variables and constraints. Thus, we can compute the optimum in constant time.

Theorem 9. (restated) 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 $\mathcal{P}$.

4.6 Missing Details in Section 4.2

4.6.1 Details for Section 4.2.1

Lemma 23. We find an affine transform T in $O(2^{O(d)} n \log n)$ time, such that the convex polytope $M' = T(M)$ is α-fat for some constant α.

Proof. By the results in [19], we only need to construct an approximate bounding box, which can be done as follows: We first identify two points y_1 and y_2 in M such that their distance is a constant approximation of the diameter of M. Then we project the points in M to a hyperplane $H \in \mathbb{R}^{d-1}$ perpendicular to the line through y_1 and y_2, and recursively identify two points among the projected points as the approximate diameter. Hence, it suffices to show how to identify such two points y_1 and y_2. Let $\delta = \arccos(1/2)$. Suppose we are working on $\mathbb{R}^d$. We compute a set $\mathcal{I}$ of $O(\delta^{-(d-1)})$ points on the unit sphere $\mathbb{S}^{d-1}$ such that for any point $s \in \mathbb{S}^{d-1}$, there is a point $\vec{u} \in \mathcal{I}$ such that $\angle(\vec{u}, s) \leq \delta$ (see e.g., [10, 25]). From Lemma 22, we know that we can compute for each direction $\vec{u} \in \mathbb{S}^{d-1}$, the point $x(\vec{u}) \in M$ that maximizes $\langle \vec{u}, x(\vec{u}) \rangle$ in $O(n \log n)$ time. For each $\vec{u} \in \mathcal{I}$, compute both $x(\vec{u})$ and $x(-\vec{u})$, and pick the pair that maximizes $\|x(\vec{u}) - x(-\vec{u})\|$. Now, we argue this

is a constant approximation of the diameter. Suppose the diameter of M is (y_1, y_2) where $y_1, y_2 \in M$. Consider the direction $\vec{u}' = (y_1 - y_2)/\|y_1 - y_2\|$. Without loss of generality, assume $y_1 = \arg\max_y \langle y, \vec{u}' \rangle$ and $y_2 = \arg\max_y \langle y, -\vec{u}' \rangle$. Moreover, there is a direction $\vec{u} \in \mathcal{I}$ such that $\angle(\vec{u}, s) \leq \delta$. Therefore, we can get that

$$\begin{aligned}\omega(M, \vec{u}) &= \mathsf{f}(M, \vec{u}) + \mathsf{f}(M, -\vec{u}) \geq \langle y_1, \vec{u} \rangle + \langle y_2, -\vec{u} \rangle \\ &= \langle \vec{u}, y_1 - y_2 \rangle = \|y_1 - y_2\| \cos\angle(\vec{u}, \vec{u}') \geq \|y_1 - y_2\|/2.\end{aligned}$$

In the third equation, we use the simple fact that $\cos\angle(\vec{u}, \vec{u}') = \langle \vec{u}, \vec{u}' \rangle / \|\vec{u}\| \|\vec{u}'\|$. □

Lemma 24. $S = \{x(\vec{u})\}_{\vec{u} \in \mathcal{I}}$ is an ε-kernel for M'.

Proof. Consider an arbitrary vector $s \in \mathbb{S}^{d-1}$ with $\|s\| = 1$. Suppose the point $a \in M'$ maximizes $\langle s, a \rangle$ and $b \in M'$ maximizes $\langle -s, b \rangle$. Hence, $\omega(M', s) = \langle s, a \rangle - \langle s, b \rangle = \langle s, a - b \rangle$. By the construction of $\mathcal{I}$, there is a direction $\vec{u} \in \mathcal{I}$ (with $\|\vec{u}\| = 1$) such that $\|\vec{u} - s\| \leq \delta$. Then, we can see that

$$\begin{aligned}\omega(S, s) &\geq \langle s, x(\vec{u}) \rangle - \langle s, x(-\vec{u}) \rangle = \langle s, x(\vec{u}) - x(-\vec{u}) \rangle \\ &= \langle \vec{u}, x(\vec{u}) - x(-\vec{u}) \rangle + \langle s - \vec{u}, x(\vec{u}) - x(-\vec{u}) \rangle \\ &\geq \langle \vec{u}, a - b \rangle - \|s - \vec{u}\| \|x(\vec{u}) - x(-\vec{u})\| \\ &= \langle s, a - b \rangle + \langle \vec{u} - s, a - b \rangle - \|s - \vec{u}\| \|x(\vec{u}) - x(-\vec{u})\| \\ &\geq \langle s, a - b \rangle - \|s - \vec{u}\| \|x(\vec{u}) - x(-\vec{u})\| - \|\vec{u} - s\| \|a - b\| \\ &\geq \omega(M', s) - O(\delta d) \geq (1 - \varepsilon)\omega(M', s)\end{aligned}$$

In the last and 2nd to last inequalities, we use the fact that M' is α-fat (i.e., $\alpha\overline{\mathbb{C}} \subset M' \subset \overline{\mathbb{C}}$). □

4.6.2 Details for Section 4.2.2

Theorem 26. Under the β-assumption, there is an ε-exp-kernel in $\mathbb{R}^d$ (for $d = O(1)$), which is of size $O(\beta^{-(d-1)}\varepsilon^{-(d-1)/2} \log(1/\varepsilon))$ and satisfies the subset constraint, in the existential uncertainty model.

Proof. Our algorithm is inspired by the peeling idea in [9]. Let $\varepsilon_1 = \varepsilon\alpha\beta^2/4\sqrt{d}$, where α is a constant defined later. We repeat the following for $L = O(\log_{1-\beta}\varepsilon_1) = O(\log(1/\varepsilon))$ rounds: In round i, we first compute an $(\varepsilon_1/\sqrt{d})$-kernel $\mathcal{S}_i$ (of size $O((\sqrt{d}/\varepsilon_1)^{(d-1)/2}) = O(\beta^{-(d-1)}\varepsilon^{-(d-1)/2})$) for the remaining points (in the deterministic sense) and then delete all points of $\mathcal{S}_i$. Let $\mathcal{S} = \cup_i\mathcal{S}_i$. Now, we show that $\mathcal{S}$ is an ε-exp-kernel for $\mathcal{P}$.

We first establish a lower bound of $\omega(\mathcal{P}, \vec{u})$ for any unit vector $\vec{u} \in \mathbb{S}^{d-1}$. Assume without loss of generality that $\alpha\overline{\mathbb{C}} \subset \mathsf{ConvH}(\mathcal{P}) \subset \overline{\mathbb{C}}$ where $\overline{\mathbb{C}} = [-1, 1]^d$ and α is a constant only depending on d. Since $\alpha\overline{\mathbb{C}} \subset \mathsf{ConvH}(P)$, we know there is a point $s \in \mathsf{ConvH}(P)$ such that $\langle\vec{u}, s\rangle \geq \alpha$ and a different point $s' \in \mathsf{ConvH}(P)$ such that $\langle\vec{u}, s'\rangle \leq -\alpha$. Hence, we have that

$$\omega(\mathcal{P}, \vec{u}) \geq \beta^2(\langle\vec{u}, s\rangle - \langle\vec{u}, s'\rangle) \geq 2\alpha\beta^2.$$

Fix an arbitrary direction $\vec{u} \in \mathbb{S}^{d-1}$. Now, we bound the difference between $\mathsf{f}(\mathcal{P}, \vec{u})$ and $\mathsf{f}(\mathcal{S}, \vec{u})$. We show that for any real value $x \in [-\sqrt{d}, \sqrt{d}]$,

$$\Pr_{P\sim\mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x] \leq \Pr_{S\sim\mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x - \varepsilon_1] + \varepsilon_1. \tag{4.6}$$

In fact, a proof of the above statement provides a proof for Theorem 43 (i.e., $\mathcal{S}$ is an (ε, τ)-quant-kernel as well).

Let $\mathcal{L}_\mathcal{P} = \{s_1, s_2, \ldots, s_L\}$ be the set of L points $s \in \mathcal{P}$ that maximize $\langle s, \vec{u}\rangle$ (i.e., the first L vertices in the canonical order w.r.t. $\vec{u}$). Similarly, let $\mathcal{L}_\mathcal{S} = \{s'_1, s'_2, \ldots, s'_L\}$ be the set of L points $s' \in \mathcal{S}$ that maximize $\langle s', \vec{u}\rangle$. We distinguish two cases:

1. $\mathcal{L}_\mathcal{P} = \mathcal{L}_\mathcal{S}$: If $x \geq \langle\vec{u}, s_L\rangle$, we can see that $\Pr_{P\sim\mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x] = \Pr_{S\sim\mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x]$. If $x < \langle\vec{u}, s_L\rangle$, both $\Pr_{P\sim\mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x]$ and $\Pr_{S\sim\mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x]$ are at least $1 - \prod_{s\in\mathcal{L}_\mathcal{P}}(1 - p_s) \geq 1 - (1 - \beta)^L \geq 1 - \varepsilon_1$.

2. Suppose j is the smallest index such that $s_j \neq s'_j$. For $x > \langle\vec{u}, s_j\rangle$, we can see that $\Pr_{P\sim\mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x] = \Pr_{S\sim\mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x]$. Now, we focus on the case where $x \leq \langle\vec{u}, s_j\rangle$. From the construction of $\mathcal{S}$, we can see that $\langle s'_j, \vec{u}\rangle \geq \langle s_j, \vec{u}\rangle - \varepsilon_1$

for all $j' \geq j$. [13] Hence, for $x \leq \langle \vec{u}, s_j \rangle$, we can see that

$$\Pr_{S \sim \mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x - \varepsilon_1] \geq 1 - \prod_{s \in \mathcal{L}_{\mathcal{S}}} (1 - p_s) \geq 1 - \varepsilon_1.$$

So, in either case, (4.6) is satisfied. We also need the following basic fact about the expectation: For a random variable X, if $\Pr[X \geq a] = 1$, then $\mathbb{E}[X] = \int_b^\infty \Pr[X \geq x]dx + b$ for any $b \leq a$. Since $-\sqrt{d} \leq \mathsf{f}(P, \vec{u}) \leq \sqrt{d}$ for any realization P, we have

$$\begin{aligned}
\mathsf{f}(\mathcal{P}, \vec{u}) &= \int_{-\sqrt{d}}^{\infty} \Pr_{P \sim \mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x]dx - \sqrt{d} \\
&\leq \int_{-\sqrt{d}}^{\infty} \Pr_{S \sim \mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x - \varepsilon_1]dx + 2\sqrt{d}\varepsilon_1 - \sqrt{d} \\
&\leq \int_{-\sqrt{d}-\varepsilon_1}^{\infty} \Pr_{S \sim \mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x]dx - \sqrt{d} - \varepsilon_1 + 3\sqrt{d}\varepsilon_1 \\
&= \mathsf{f}(\mathcal{S}, \vec{u}) + 3\sqrt{d}\varepsilon_1,
\end{aligned}$$

where the only inequality is due to (4.6) and the fact that $\Pr_{P \sim \mathcal{P}}[\mathsf{f}(P, \vec{u}) \geq x] = \Pr_{S \sim \mathcal{S}}[\mathsf{f}(S, \vec{u}) \geq x] = 0$ for $x > 1$. Similarly, we can get that $\mathsf{f}(\mathcal{S}, -\vec{u}) \geq \mathsf{f}(\mathcal{P}, -\vec{u}) - 3\varepsilon_1\sqrt{d}$. By the choice of ε_1, we have that $6\sqrt{d}\varepsilon_1 \leq \varepsilon \cdot 2\alpha\beta^2 \leq \varepsilon\omega(\mathcal{P}, \vec{u})$. Hence, $\omega(\mathcal{S}, \vec{u}) \geq \omega(\mathcal{P}, \vec{u}) - 6\sqrt{d}\varepsilon_1 \geq (1 - \varepsilon)\omega(\mathcal{P}, \vec{u})$. □

4.6.3 Locational uncertainty

Similar results are possible for uncertain points with locational uncertainty. Let $\mathcal{V} = \{v_1, \ldots, v_m\}$ be the set of nodes and $\mathcal{P} = \{s_1, \ldots, s_n\}$ be the collection of possible locations. Now there are n possible locations, and thus $\binom{n}{2}$ hyperplanes Γ that partition $\mathbb{R}^d$. We can replicate all bounds in this setting, except that m replaces n in each bound. The main difficulty is in replicating Lemma 49 that given a direction $\vec{u}$ calculates the vertex of M; for locational uncertain points this is described in Lemma 51. Moreover, the $O(n^2 \log n)$ bound for $\mathbb{R}^2$ is also carefully described in Lemma 52.

In the locational uncertainty model, Lemma 18 also holds with a stronger general

[13] To see this, consider the round in which $s'_{j'}$ is chosen. Let $\hat{s}$ be the vertex minimizing $\langle \hat{s}, \vec{u} \rangle$. As s_j is not chosen, we must have $\langle s'_{j'}, \vec{u} \rangle - \langle \hat{s}, \vec{u} \rangle \geq (1 - \varepsilon_1/\sqrt{d})(\langle s_j, \vec{u} \rangle - \langle \hat{s}, \vec{u} \rangle)$.

position assumption. With the new general position assumption, it is straightforward to show that the gradient vector is different for two adjacent cones in $\mathsf{A}(\Gamma)$. Other parts of the proof is essentially the same as Lemma 18. The details can be found below. Theorem 2 also holds for the locational model without any change in the proof (the running time becomes $O(n \log^2 n)$).

Now, we prove that Lemma 18 also holds for the locational model. For this purpose, we need a stronger general position assumption: (1) For any $v \in \mathcal{V}$, $\prod_{s \in \mathcal{P}} p_{vs} \in (0, 1)$. This suggests that we need to consider the model with both existential and locational uncertainty. We can make this assumption hold by subtracting an infinitesimal value from each probability value without affecting the directional width in any essential way. (2) For any two nodes $v_1, v_2 \in \mathcal{V}$, two locations $s_1, s_2 \in \mathcal{P}$ and two subsets of locations $S_1, S_2 \subseteq \mathcal{P}$, $p_{v_1 s_1}(\prod_{s \in S_2} p_{v_2 s}) s_1 \neq p_{v_2 s_2}(\prod_{s \in S_1} p_{v_1 s}) s_2$ (this is indeed a general position assumption since we only have a finite number of equations to exclude but uncountable number of choices of the positions).

Lemma 18. (for the locational model). Assuming the locational model and the above general position assumption, the complexity of M is the same as the cardinality of $\mathsf{A}(\Gamma)$, i.e., $|M| = |\mathsf{A}(\mathcal{P})|$. Moreover, each cone $C \in \mathsf{A}(\mathcal{P})$ corresponds to exactly one vertex s of $\mathsf{ConvH}(M)$ in the following sense: $\nabla \mathsf{f}(M, \vec{u}) = s$ for all $\vec{u} \in \operatorname{int} C$.

Proof. The proof is almost the same as that for Lemma 18 except that we need to show $\mathsf{f}(M, \vec{u})$ is different for two adjacent cones in $\mathsf{A}(\Gamma)$. Again, let C_1, C_2 be two adjacent cones separated by some hyperplane in Γ. Suppose $\vec{u}_1 \in \operatorname{int} C_1$ and $\vec{u}_2 \in \operatorname{int} C_2$. Consider the canonical orders O_1 and O_2 of $\mathcal{P}$ with respect to $\vec{u}_1$ and $\vec{u}_2$ respectively. W.l.o.g., assume that $O_1 = \{s_1, \ldots, s_i, s_{i+1}, \ldots, s_n\}$ and $O_1 = \{s_1, \ldots, s_{i+1}, s_i, \ldots, s_n\}$.

Let $\mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u})$ be the probability that the largest point along $\vec{u}$ is uncertain node $v \in \mathcal{V}$ at location $s \in \mathcal{P}$. Using the notations from Lemma 51, $\mathsf{f}(\mathcal{V}, \vec{u})$ can computed by $\prod_{v \in \mathcal{V}, s \in \mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u}) \langle s, \vec{u} \rangle$. Hence, $\nabla \mathsf{f}(\mathcal{V}, \vec{u}) = \prod_{v \in \mathcal{V}, s \in \mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u}) s$.

Suppose s_i is a possible location for v_1 and s_{i+1} is a possible location for v_2. Denote by $\mathcal{P}^{\mathsf{R}}(s, \vec{u})$ the subset of $s' \in \mathcal{P}$ such that $\langle s', \vec{u} \rangle > \langle s, \vec{u} \rangle$ and denote by $\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, \vec{u})$ as the probability that no node $v \in \mathcal{V}$ appears at a larger location than

$s \in \mathcal{P}$ along direction $\vec{u}$. If $v_1 \neq v_2$, we have $\nabla f(\mathcal{V}, \vec{u}_1) - \nabla f(\mathcal{V}, \vec{u}_2) = \Pr_{\emptyset}^{R}(s_i, \vec{u}_1) \cdot \left(s_i p_{v_1 s_i} \prod_{s' \in \mathcal{P}^R(s, \vec{u}_1)} p_{v_2 s'} - s_{i+1} p_{v_2 s_{i+1}} \prod_{s' \in \mathcal{P}^R(s, \vec{u}_2)} p_{v_1 s'} \right) \neq 0$. If $v_1 = v_2 = v$, we have $\nabla f(\mathcal{V}, \vec{u}_1) - \nabla f(\mathcal{V}, \vec{u}_2) = \Pr_{\emptyset}^{R}(s_i, \vec{u}_1) \cdot \left(s_i p_{v s_i} - s_{i+1} p_{v s_{i+1}} \right) \neq 0$. □

4.7 Missing Details in Section 4.3

Theorem 36. 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 $\vec{u}$, we have that $\Pr\left[\omega(\mathcal{S}, \vec{u}) \leq t \right] \in \Pr\left[\omega(\mathcal{P}, \vec{u}) \leq t \right] \pm \tau$.

Proof. Fix an arbitrary direction $\vec{u}$ (w.l.o.g., say it is the x-axis) and rename all points in $\mathcal{P}$ as $s_1, s_2, \ldots, s_n$ as before. Consider the Poissonized instance of $\mathcal{P}$. Let $s'_1, \ldots, s'_N$ be the N points in $\mathcal{S}$ (also sorted in nondecreasing order of their x-coordinates). Now, we create a coupling between all mass in A and that in B, as follows. We process all points in A from left to right, starting with s_1. The process has N rounds. In each round, we assign exactly $1/N$ units of mass in A to a point in B. In the first round, if s_1 contains less than $1/N$ units of mass, we proceed to $s_2, s_3, \ldots s_i$ until we reach $1/N$ units collectively. We split the last point s_i into two points s_{i1} and s_{i2} so that the mass contained in $s_1, \ldots, s_{i-1}, s_{i1}$ is exactly $1/N$, and we assign those points to s'_1. We start the next round with s_{i2}. If s_1 contains more than $1/N$ units of mass, we split s_1 into s_{11} (s_{11} contains $1/N$ units) and s_{12} and we start the second round with s_{12}. We repeat this process until all mass in A is assigned.

The above coupling can be viewed as a mass transportation from A to B. We will need one simple but useful property about this transportation: for any vertical line $x = t$, at most τ_1 units of mass are transported across the vertical line (by Theorem 32).

In the construction of the coupling, many points in A may be split. We rename them to be $s_1, \ldots, s_L$ (according to the order in which they are processed). The sequence $s_1, \ldots, s_L$ can be divided into N segments, each assigned to a point in $\mathcal{S}$. For a point s'_i in $\mathcal{S}$, let $\text{seg}(i)$ be the segment (the set of points) assigned to s'_i. For any node s and real $t > 0$, we use $H(s, t)$ to denote the right open halfplane defined

by the vertical line $x = x(s) + t$, where $x(s)$ is the x-coordinate of s (see Figure 4-3).

Let X_i (Y_i resp.) be the Poisson distributed random variable corresponding to s_i (s'_i resp.) (i.e., $X_i \sim \text{Pois}(\lambda_{s_i})$ and $Y_i \sim \text{Pois}(\lambda/N)$) for all i. For any $H \subset \mathbb{R}^2$, we write $X(H) = \sum_{s_i \in H \cap \mathcal{P}} X_i$ and $Y(H) = \sum_{s'_i \in H \cap \mathcal{S}} Y_i$. We can rewrite $\Pr[\omega(\mathcal{S}, \vec{u}) \le t]$ as follows:

$$\begin{aligned}\Pr[\omega(\mathcal{S}, \vec{u}) \le t] &= \sum_{i=1}^{N} \Pr[s'_i \text{ is the leftmost point and } \omega(\mathcal{S}, \vec{u}) \le t] + \Pr[\text{no point in } \mathcal{S} \text{ appears}] \\ &= \sum_{i=1}^{N} \Pr[Y_i \neq 0] \Pr\Big[\sum_{j=1}^{i-1} Y_j = 0\Big] \Pr[Y(H(s'_i, t)) = 0] + \Pr\Big[\sum_{s'_i \in \mathcal{S}} Y_i = 0\Big]\end{aligned} \tag{4.7}$$

Similarly, we can write that [14]

$$\begin{aligned}\Pr[\omega(\mathcal{P}, \vec{u}) \le t] &= \sum_{i=1}^{m} \Pr[X_i \neq 0] \Pr\Big[\sum_{j=1}^{i-1} X_j = 0\Big] \Pr[X(H(s_i, t)) = 0] + \Pr\Big[\sum_{s_i \in \mathcal{P}} X_i = 0\Big] \\ &= \sum_{i=1}^{N} \sum_{k \in \text{seg}(i)} \Pr[X_k \neq 0] \Pr\Big[\sum_{j=1}^{k-1} X_j = 0\Big] \Pr[X(H(s_k, t)) = 0] + \Pr\Big[\sum_{s_i \in \mathcal{P}} X_i = 0\Big]\end{aligned} \tag{4.8}$$

We proceed by attempting to show each each summand of (4.8) is close to the corresponding one in (4.7). First, we can see that $\Pr[\sum_{s'_i \in \mathcal{S}} Y_i = 0] = \Pr[\sum_{s_i \in \mathcal{P}} X_i = 0]$ since both $\sum_{s'_i \in \mathcal{S}} Y_i$ and $\sum_{s_i \in \mathcal{P}} X_i$ follow the Poisson distribution $\text{Pois}(\lambda)$.

For any segment i, we can see that $\sum_{k \in \text{seg}(i)} \lambda_{s_k} = \lambda/N$. Moreover, we have $\lambda_{s_k} \le \lambda/N \le \tau/32$, thus $\exp(-\lambda_{s_k}) \in (1 - \lambda_{s_k}, (1 + \tau/16)(1 - \lambda_{s_k}))$.

$$\begin{aligned}\sum_{k \in \text{seg}(i)} \Pr[X_k \neq 0] &= \sum_{k \in \text{seg}(i)} (1 - \exp(-\lambda_{s_k})) \in (1 \pm \frac{\tau}{16}) \sum_{k \in \text{seg}(i)} \lambda_{s_k} \\ &\subset (1 \pm \frac{\tau}{8})(1 - \exp(\frac{\lambda}{N})) = (1 \pm \frac{\tau}{8})\Pr[Y_i \neq 0].\end{aligned} \tag{4.9}$$

[14] Note that splitting nodes does not change the distribution of $\omega(\mathcal{P}, \vec{u})$: Suppose a node s (corresponding to r.v. X) was spit to two nodes s_1 and s_2 (corresponding to X_1 and X_2 resp.). We can see that $\Pr[X \neq 0] = \Pr[X_1 \neq 0 \text{ and } X_2 \neq 0] = \Pr[X_1 + X_2 \neq 0]$.

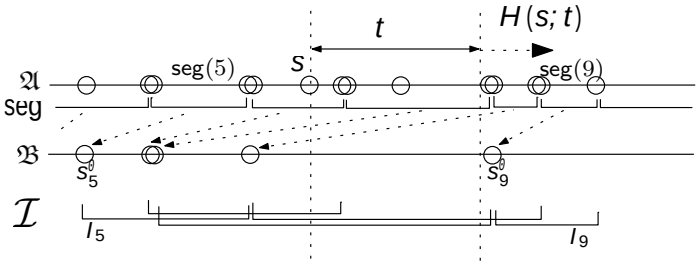

Figure 4-3: Illustration of the interval graph $\mathcal{I}$. For illustration purpose, co-located points (e.g., points that are split in A) are shown as overlapping points. The arrows indicate the assignment of the segments to the points in B. Theorem 32 ensures that any vertical line can not stab many intervals.

Then, we notice that for any $k \in \mathsf{seg}(i)$ (i.e., s_k is in the segment assigned to s'_i), it holds that

$$\Pr\left[\sum_{j=1}^{k} X_j = 0\right] \in [e^{-i\lambda/N}, e^{-\lambda(i-1)/N}] \subset (1 \pm \frac{\tau}{8})e^{-\lambda(i-1)/N} = (1 \pm \frac{\tau}{8})\Pr\left[\sum_{j=1}^{i-1} Y_j = 0\right]. \tag{4.10}$$

The first inequality holds because $\sum_{j=1}^{k} X_j \sim \mathsf{Pois}\left(\sum_{j=1}^{k} \lambda_{s_j}\right)$ and $\lambda(i-1)/N \leq \sum_{j=1}^{k} \lambda_{s_j} \leq \lambda i/N$.

If we can show that $\Pr[X(H(s_k, t)) = 0]$ is close to $\Pr[Y(H(s'_i, t)) = 0]$ for $k \in \mathsf{seg}(i)$, we can finish the proof easily since each summand of (4.8) would be close to the corresponding one in (4.7). However, this is in general not true and we have to be more careful.

Recall that the sequence $s_1, \ldots, s_n$ is divided into N segments. Let $K = \lambda/\tau$. We say that the ith segment (say $\mathsf{seg}(i) = \{s_j, s_{j+1}, \ldots, s_k\}$) is a good segment if

$$E_i = \max\left\{\left|\mathsf{B}(H(s'_i, t)) - \mathsf{A}(H(s_j, t))\right|, \left|\mathsf{B}(H(s'_i, t)) - \mathsf{A}(H(s_k, t))\right|\right\} \leq \frac{1}{K}.$$

Otherwise, the segment is bad. For a good segment $\mathsf{seg}(i)$ and any $k \in \mathsf{seg}(i)$,

$$\Pr[X(H(s_k, t)) = 0] = \exp\left(-\lambda \mathsf{A}(H(s_k, t))\right) \in \exp\left(-\lambda \mathsf{B}(H(s'_i, t)) \pm \lambda/K\right)$$

$$\subset \Pr[Y(H(s'_i,t))=0]e^{\pm\lambda/K} \subset \Pr[Y(H(s'_i,t))=0](1\pm\tau/8). \tag{4.11}$$

We use Gs to denote the set of good segments and Bs the set of bad segments. Now, we consider the summations in both (4.7) and (4.8) with only good segments. We have that

$$\begin{aligned}
&\sum_{i\in Gs}\sum_{k\in seg(i)} \Pr[X_k\neq 0]\Pr\Big[\sum_{j=1}^{k-1} X_j = 0\Big]\Pr[X(H(s_k,t))=0] \\
\in &\sum_{i\in Gs}\Pr\Big[\sum_{j=1}^{i-1} Y_j = 0\Big](1\pm\tau/8)\Pr[Y(H(s'_i,t))=0](1\pm\tau/8)\Pr[Y_i\neq 0](1\pm\tau/8) \\
\subset &\sum_{i\in Gs}\Pr[Y_i\neq 0]\Pr\Big[\sum_{j=1}^{i-1} Y_j = 0\Big]\Pr[Y(H(s'_i,t))=0]\pm\tau/2,
\end{aligned}$$

where the first inequality is due to (4.10) and (4.11) and the second holds because (4.9)

Now, we show the total contributions of bad segments to both (4.7) and (4.8) are small. We partition all of the bad segments into $\log(1/\tau)+1$ different sets $B_0,\ldots,B_{\log\frac{1}{\tau}-1}$. Let $B_i = \{i \mid \frac{2^i}{K} < E_i \le \frac{2^{i+1}}{K}\}$ for $0\le i\le \log(1/\tau)-1$. Let $B_{\log\frac{1}{\tau}} = \{i \mid E_i > \frac{1}{\lambda}\}$.

With the above notations, we prove the following crucial inequality:

$$\sum_{i=0}^{\log\frac{1}{\tau}} |B_i|\cdot 2^i = O(\tau_1 NK). \tag{4.12}$$

Now, we prove (4.12). Consider all points $s_1,\ldots,s_L$ and $s'_1,\ldots,s'_L$ lying on the same x-axis. For each i (with $seg(i)=\{s_j,s_{j+1},\ldots,s_k\}$), we draw the minimal interval I_i that contains s'_i, s_j and s_k. If the ith segment is bad and belongs to B_j, we also say I_i is a bad interval of label j. All intervals $\{I_i\}_i$ define an interval graph $\mathcal{I}$. We can see that any vertical line can stab at most $\tau_1 N+1$ intervals, because at most τ_1 unit of mass can be transported across the vertical line, and each interval is responsible for a transportation of exactly $1/N$ units of mass (except the one that

intersects the vertical line). Hence, the interval graph $\mathcal{I}$ can be colored with at most $\tau_1 N + 1$ colors (this is because the clique number of $\mathcal{I}$ is at most $\tau_1 N + 1$ and the chromatic number of an interval graph is the same as its clique number). Consider a color class C (which consists of a set of non-overlapping intervals). Imagine we move an interval I of length t along the x-axis from left to right. When the left endpoint of I passes through an bad interval of label j in C, by the definition of bad segments, the right endpoint of I passes through $\Omega(2^j N/K)$ segments. Suppose the color class C contains b_j bad segments in B_j. Since the right endpoint of I can pass through at most N segments, we have the following inequalities by summing over all labels:

$$\sum_{j=0}^{\log \frac{1}{\tau}-1} b_j \cdot 2^j N/K \leq N.$$

Summing up all color classes, we obtain (4.12).

For B_j ($0 \leq j \leq \log \frac{1}{\tau}$), we can bound the total contribution as follows. By the definition of B_j, we can see that

$$\begin{aligned}
&\sum_{i \in B_j} \sum_{k \in seg(i)} \Pr[X_k \neq 0] \Pr\left[\sum_{j=1}^{k-1} X_j = 0\right] \Pr[X(H(s_k, t)) = 0] \\
&\subset (1 \pm 5 \cdot 2^j \tau) \sum_{i \in B_j} \Pr[Y_i \neq 0] \Pr\left[\sum_{j=1}^{i-1} Y_j = 0\right] \Pr[Y(H(s'_i, t)) = 0],
\end{aligned}$$

Thus, the total contribution of bad segments in B_j ($0 \leq j \leq \log \frac{1}{\tau} - 1$) to the corresponding summands in ((4.7)-(4.8)) is at most

$$5 \cdot 2^j \tau \sum_{i \in B_j} \Pr[Y_i \neq 0] = 5|B_j| \cdot 2^j \tau \times (1 - \exp(-\frac{\lambda}{N})) = O(|B_i| 2^j \tau \lambda / N),$$

where $\Pr[Y_i \neq 0] = 1 - \exp(-\frac{\lambda}{N})$ (since $Y_i \sim \text{Pois}(\frac{\lambda}{N})$).

For $B_{\log \frac{1}{\tau}}$, the total contribution is bounded by the following.

$$\sum_{i \in B_{\log \frac{1}{\tau}}} \sum_{k \in seg(i)} \Pr[X_k \neq 0] \Pr\left[\sum_{j=1}^{k-1} X_j = 0\right] \Pr[X(H(s_k, t)) = 0]$$

$$- \Pr[Y_i \neq 0] \Pr\Big[\sum_{j=1}^{i-1} Y_j = 0\Big] \Pr[Y(H(s'_i, t)) = 0]\Big)$$

$$\leq \sum_{i \in B_{\log \frac{1}{\tau}}} \sum_{k \in \mathrm{seg}(i)} \Pr[X_k \neq 0] + \sum_{i \in B_{\log \frac{1}{\tau}}} \Pr[Y_i \neq 0] \leq 3 \sum_{i \in B_{\log \frac{1}{\tau}}} \Pr[Y_i \neq 0]$$

$$\leq 3|B_{\log \frac{1}{\tau}}| \times (1 - \exp(-\frac{\lambda}{N}) = O(|B_{\log \frac{1}{\tau}}|\lambda/N)$$

Summing up all j and using (4.12), we obtain the following inequality .

$$\sum_{j=0}^{\log \frac{1}{\tau}} O(|B_j| 2^j \tau \lambda / N) = O(\tau_1 \tau \lambda K) \leq \frac{\tau}{4}.$$

This finishes the proof. □

4.8 Missing Details in Section 4.4

Lemma 45. Let $N = O\big(\varepsilon_1^{-2} \varepsilon_0^{-(d-1)/2} \log(1/\varepsilon_0)\big)$, where $\varepsilon_0 = (\varepsilon/4(r-1))^r$, $\varepsilon_1 = \varepsilon\beta^2$. For any $t \geq 0$ and any direction $\vec{u} \in \mathcal{P}$, we have that

$$\Pr_{P \sim \mathcal{S}}\Big[\max_{s \in P} \langle \vec{u}, s\rangle^{1/r} \geq t\Big] \in \Big(\Pr_{P \sim \mathcal{P}}\Big[\max_{s \in \mathcal{E}(P)} \langle \vec{u}, s\rangle^{1/r} \geq t\Big) \pm \varepsilon_1/4, \text{ and}$$

$$\Pr_{P \sim \mathcal{S}}\Big[\min_{s \in P} \langle \vec{u}, s\rangle^{1/r} \geq t\Big] \in \Big(\Pr_{P \sim \mathcal{P}}\Big[\min_{s \in \mathcal{E}(P)} \langle \vec{u}, s\rangle^{1/r} \geq t\Big) \pm \varepsilon_1/4.$$

Proof. The argument is almost the same as that in Lemma 30. Let $L = O(\varepsilon_0^{-(d-1)/2})$. We still 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 $\mathcal{E}(P) = \{(x_1^1, \ldots, x_d^1), \ldots, (x_1^L, \ldots, x_d^L)\}$ (if $|\mathcal{E}(P)| < L$, we pad it with $(0, \ldots, 0)$). We let $g(\mathcal{E}(P)) = (x_1^1, \ldots, x_d^1, \ldots, x_1^L, \ldots, x_d^L) \in \mathbb{R}^{dL}$. For any $t \geq 0$ and any direction $\vec{u} \in \mathcal{P}$, note that $\max_{s \in \mathcal{E}(P)} \langle \vec{u}, s\rangle^{1/r} \geq t$ holds if and only if there exists some $1 \leq i \leq |\mathcal{E}(P)|$ satisfies that $\sum_{j=1}^{d} x_j^i \vec{u}_j \geq t^r$, which is equivalent to saying that point $g(\mathcal{E}(P))$ is in the union of the those $|\mathcal{E}(P)|$ half-spaces.

Let X be the image set of g. Let $(X, \mathcal{R}^i)$ $(1 \leq i \leq L))$ be a range space, where $\mathcal{R}^i$ is the set of half spaces $\{\sum_{j=1}^{d} x_j^i \vec{u}_j \geq t \mid u = (\vec{u}_1, \ldots, \vec{u}_d) \in \mathbb{R}^d, t \geq 0\}$. Let

is denoted p_{vs}. To simplify analysis and discussion, we assume each location $s \in \mathcal{P}$ only has the potential to be realized by any one uncertain point $v \in \mathcal{V}$.

We now replicate the lemmas in the previous section for this setting. We use the same notation and structure when possible.

Lemma 51. For any direction $\vec{u}$, we can compute $\omega(\mathcal{V}, \vec{u})$, $\mathsf{f}(\mathcal{V}, \vec{u})$, and $\nabla\mathsf{f}(\mathcal{V}, \vec{u})$ in $O(n \log n)$ time; if the locations of $\mathcal{P}$ are already sorted along the direction $\vec{u}$, then we can compute them in $O(n)$ time.

Proof. Again, we first compute $\mathsf{f}(\mathcal{V}, \vec{u})$ since $\omega(\mathcal{V}, \vec{u}) = \mathsf{f}(\mathcal{V}, \vec{u}) - \mathsf{f}(\mathcal{V}, -\vec{u})$ and $\mathsf{f}(\mathcal{V}, -\vec{u})$ can be computed similarly.

We follow the structure and proof of Lemma 49 and just note the changes. The first change is that we need to keep a bit more structure since there is now dependence between the different locations of each uncertain node v. Recall that $\mathcal{P}^{\mathsf{R}}(s, \vec{u})$ is the subset of $s' \in \mathcal{P}$ such that $\langle s', \vec{u}\rangle > \langle s, \vec{u}\rangle$ and $\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, \vec{u})$ is the probability that no node $v \in \mathcal{V}$ appears at a larger location than $s \in \mathcal{P}$ along direction $\vec{u}$. To describe this probability we first define a vector A_{v} indexed by s as $A_{\mathsf{v}}[s] = 1 - \sum_{s' \in \mathcal{P}^R(s,\vec{u})} p_{vs'}$ as the probability that uncertain node v does not appear in any of its possible locations which are after s along direction $\vec{u}$. Now we can define

$$\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, u) = \prod_{v \in \mathcal{P}} A_{\mathsf{v}}[s].$$

Also recall that $\mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u})$ is the probability that the largest point along $\vec{u}$ is uncertain point $v \in \mathcal{V}$ at location $s \in \mathcal{P}$. This updates equation (4.13) to be

$$\mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u}) = p_{vs} \cdot \mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, \vec{u}) / A_{\mathsf{v}}[s].$$

Note the two key differences. First we need to sum the probabilities for each location of v since they are mutually exclusive. Second, value $A_{\mathsf{v}}[s]$ needs to be factored out of $\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, u)$ because it is already accounted for in p_{vs} locating s at s, again since they are mutually exclusive.

It follows that $\mathsf{f}(\mathcal{V}, \vec{u}) = \sum_{v \in \mathcal{V}, s \in \mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u}) \langle s, \vec{u}\rangle$.

To compute $\mathsf{f}(\mathcal{V}, \vec{u})$ we again start by projecting all points P onto $\rho(\vec{u})$ obtaining coordinates $\langle \vec{u}, s \rangle$ and sorting, if needed. This takes $O(n \log n)$ time. Given these coordinates, sorted, it now takes a bit more work in the locational setting to show that $\mathsf{f}(\mathcal{V}, \vec{u})$ can be computed in $O(n)$ additional time. We focus on computing all n values $\mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u})$; from there is straightforward to compute $\mathsf{f}(\mathcal{V}, \vec{u})$ in $O(n)$ time.

We sweep over the locations $s \in \mathcal{P}$ from largest $\langle s, \vec{u} \rangle$ value to smallest, and we maintain each $A_{\mathsf{v}}[s]$ and $\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, \vec{u})$ along the way. Given these, it is not hard to calculate $\mathsf{Pr}^{\mathsf{R}}(v, s', \vec{u})$ in constant time with $p_{\mathsf{v}s'}$ for $s' \in \mathcal{P}$ as the next smallest value $\langle s', \vec{u} \rangle$. The important observation is that we only need to update $A_{\mathsf{v}}[s]^{\text{new}} = A_{\mathsf{v}}[s]^{\text{old}} - p_{\mathsf{v}s}$ if $p_{\mathsf{v}s} > 0$ (which by assumption holds for only one $v \in \mathcal{P}$). Then $\mathsf{Pr}^{\mathsf{R}}_{\emptyset}(s, u)$ is updated by multiplying by $A_{\mathsf{v}}[s]^{\text{new}} / A_{\mathsf{v}}[s]^{\text{old}}$. Both operations can be done in constant time as needed to complete the proof.

It remains to compute $\nabla\mathsf{f}(\mathcal{V}, \vec{u})$. It is not hard to see that in the locational model

$$\nabla\mathsf{f}(\mathcal{V}, \vec{u}) = \sum_{\mathsf{v} \in \mathcal{V}, \mathsf{s} \in \mathcal{P}} \mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u}) s.$$

Note that the above has already computed the n values $\mathsf{Pr}^{\mathsf{R}}(v, s, \vec{u})$ for all $v \in \mathcal{V}$ and $s \in \mathcal{P}$. Therefore, $\nabla\mathsf{f}(\mathcal{V}, \vec{u})$ can be computed in additional $O(n)$ time. □

Lemma 52. We can build a data structure of $O(n^2)$ size in $O(n^2 \log n)$ time that can compute $\omega(\mathcal{V}, \vec{u})$, $\mathsf{f}(\mathcal{V}, \vec{u})$, and $\nabla\mathsf{f}(\mathcal{V}, \vec{u})$ in $O(\log n)$ time for any query direction $\vec{u}$. Further, we can construct M explicitly in $O(n^2 \log n)$ time.

Proof. Again we first discuss the case for computing $\mathsf{f}(\mathcal{V}, \vec{u})$. For ease of discussion, we assume the angle $\theta_{\vec{u}}$ is in $[0, \pi)$. We again follow the structure of Lemma 50. The geometry is largely the same, except that there are $h = O(n^2)$ angles $\beta_1, \beta_2, \ldots, \beta_{\mathsf{h}}$ since each pair $s, s' \in \mathcal{P}$ now defines an angle $\beta(s, s')$. But it remains to compute $\mathsf{f}(\mathcal{V}, \vec{u}) = a \cdot \cos(\theta_{\vec{u}}) + b \cdot \sin(\theta_{\vec{u}})$ for some constants a and b for any $\theta_{\vec{u}}$ in each $(\beta_{\mathsf{i}}, \beta_{\mathsf{i}+1})$. The argument is virtually the same, replacing $\mathsf{Pr}^{\mathsf{R}}(s, \vec{u})$ with $\mathsf{Pr}^{\mathsf{R}}(v, s, u)$.

It remains to show that we can calculate the constants $a(I_{\mathsf{i}+1})$ and $b(I_{\mathsf{i}+1})$ for an interval $I_{\mathsf{i}+1} = (\beta_{\mathsf{i}+1}, \beta_{\mathsf{i}+2})$ e□ciently, given the values for interval $I_{\mathsf{i}} = (\beta_{\mathsf{i}}, \beta_{\mathsf{i}+1})$.

Assume β_{i+1} is defined for two points $s_1, s_2 \in \mathcal{P}$, where $\langle s_1, \vec{u}\rangle < \langle s_2, \vec{u}\rangle$ for $\theta_{\bar{u}} \in I_i$ and $\langle s_1, \vec{u}'\rangle > \langle s_2, \vec{u}'\rangle$ for $\theta_{\bar{u}'} \in I_{i+1}$. By definition, the ordering among all other pairs of points is unchanged within (β_i, β_{i+2}). Let only v_1 take location s_1 with positive probability and only v_2 take location s_2 with positive probability. We focus on the more general case where $v_1 \neq v_2$; when $v_1 = v_2$, it is easier to update.

We again focus on updating a and the algorithm for b is symmetric. By the $v_1 \neq v_2$ assumption $A_{v_1}[s_1]$ and $A_{v_2}[s_2]$ are unchanged in interval (β_i, β_{i+2}). However from directions $\vec{u}$ to $\vec{u}'$ $A_{v_2}[s_1]$ increases by $p_{v_2 s_2}$ and $A_{v_1}[s_2]$ decreases by $p_{v_1 s_1}$; all other such values are unchanged. Let $A_v[s]^{\bar{u}}$ denote the value in direction $\vec{u}$. Hence

$$\Pr^R_\emptyset(s_1, \vec{u}') = \Pr^R_\emptyset(s_1, \vec{u}) \frac{A^{\bar{u}'}_{v_2}[s_1]}{A^{\bar{u}}_{v_2}[s_1]},$$

and so if we can update $A_{v_2}[s_1]$, we can update $\Pr^R_\emptyset(s_1, \vec{u}')$ in constant time. Only $A_{v_1}[s_2]$ and $\Pr^R_\emptyset(s_2, \vec{u}')$ also need to be updated. And then

$$\Pr^R(v_1, s_1, \vec{u}') = p_{v_1 s_1} \cdot \Pr^R_\emptyset(s_1, \vec{u}')/A_{v_1}[s_1] = \Pr^R(v_1, s_1, \vec{u}) \frac{\Pr^R_\emptyset(s_1, \vec{u}')}{\Pr^R_\emptyset(s_1, \vec{u})}$$

can also be updated in constant time, and similar for $\Pr^R(v_2, s_2, \vec{u}')$. Thus the only remaining difficulty is accessing and updating $A_{v_1}[s_2]$ and $A_{v_2}[s_1]$. We can easily do this for if we store the full size n array $A_v[\cdot]$ for each $v \in \mathcal{V}$. Note that this takes $O(n \cdot m)$ space, but since the output is a structure of size $O(n^2)$ and $m \leq n$, this is not prohibitive. (We note that these full arrays are not explicitly required in Lemma 51, which only requires $O(n)$ space.)

Finally, we can update $a(I_{i+1})$ from $a(I_i)$ similarly to equation (4.14) using $\Pr^R(v, s, \vec{u})$ in place of $\Pr^R(s, \vec{u})$. Thus in $O(m^2)$ time, after sorting all interval breakpoints in $O(n^2 \log n)$ time, we can build a data structure that allows calculation of $f(\mathcal{V}, \vec{u})$ for any $\vec{u}$ in $O(\log n)$ time.

Next, to compute $\nabla f(\mathcal{V}, \vec{u})$, recall that in the locational model $\nabla f(\mathcal{V}, \vec{u}) = \sum_{v \in \mathcal{V}, s \in \mathcal{P}} \Pr^R(v, s, \vec{u}) s$ by the proof of Lemma 18. As preprocessing, we compute the value $\sum_{v \in \mathcal{V}, s \in \mathcal{P}} \Pr^R(s, \vec{u}) s$ for each interval $\theta_{\bar{u}} \in (\beta_i, \beta_{i+1})$ for $i = 0, 1, \ldots, h$. This can be done in $O(n^2)$ time

(after we sort all angles), by using the similar idea as above. The argument is similar to that in Lemma 50 and we ignore the details. Due to the above preprocessing, given any direction $\vec{u}$, we can compute $\nabla \mathsf{f}(\mathcal{V}, \vec{u})$ in $O(\log n)$ time by binary search. Further, the above preprocessing essentially computes M, in totally $O(n^2 \log n)$ time. The lemma thus follows. □

Chapter 5 Coreset Construction for Stochastic Shape Fitting Problems

Solving geometric optimization problems over uncertain data has become increasingly important in many applications and has attracted a lot of attentions in recent years. In this chapter, we study two important geometric optimization problems, the k-center problem and the j-flat-center problem, over stochastic/uncertain data points in Euclidean spaces. We consider both problems under two popular stochastic geometric models, the existential uncertainty model and the locational uncertainty model. We provide the first PTAS (Polynomial Time Approximation Scheme) for both problems under the two models. Our results generalize the previous results for stochastic minimum enclosing ball and stochastic enclosing cylinder.

5.1 Coreset Construction for Deterministic Shape Fitting Problems

In this section, we introduce an important class of problems in computational geometry, called the shape fitting problems (see e.g., [106]). We also review two useful concepts to bound the size of coresets, called total sensitivity and VC-dimension. Finally, we briefly introduce a framework of constructing coresets for deterministic shape fitting problems.

Definition 53. (Shape fitting problems) A shape fitting problem is specified by a triple $(\mathbb{R}^d, \mathcal{F}, d)$. Here the set $\mathcal{F}$ of shapes is a family of subsets of $\mathbb{R}^d$ (e.g., all k-point sets, or all j-flats), and $d : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^{\geq 0}$ is a symmetric distance function. Define the distance of a point $s \in \mathbb{R}^d$ to a shape $F \in \mathcal{F}$ to be $d(s, F) = \min_{s' \in F} d(s, s')$. An instance P of the shape fitting problem is a (weighted) point set $\{s_1, \ldots, s_n\}$ ($s_i \in \mathbb{R}^d$),

and each s_i has a positive weight $w_i \in \mathbb{R}^+$. The goal is find a shape which best fits P, that is, a shape minimizing $\sum_{s_i \in P} w_i \cdot d(s_i, F)$ over all shapes $F \in \mathcal{F}$.

In this dissertation, if we consider the Euclidean space $\mathbb{R}^d$, we let the function $d(\cdot, \cdot)$ be the Euclidean distance function. Next, we introduce a powerful technique for deterministic shape fitting problems.

Definition 54. (Coreset for shape fitting problems) Given a (weighted) instance P of a shape fitting problem $(\mathbb{R}^d, \mathcal{F}, d)$ with a weight function $w : P \to \mathbb{R}^+$, an ε-coreset of S is a (weighted) point set together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$, such that for any shape $F \in \mathcal{F}$, we have that

$$\sum_{s_i \in S} w'_i \cdot d(s_i, F) \in (1 \pm \varepsilon) \sum_{s_i \in P} w_i \cdot d(s_i, F).^1$$

Total sensitivity and dimension. Feldman and Langberg [45] proposed a framework to construct coresets for a variety of deterministic shape fitting problems. We briefly introduce their framework in the following. To bound the size of coresets for deterministic shape fitting problems, a useful notion is called total sensitivity, originally introduced in [77].

Definition 55. (Total sensitivity of a shape fitting instance). Given an instance $P = \{s_1, \ldots, s_n\}$ of a shape fitting problem $(\mathbb{R}^d, \mathcal{F}, d)$, with a weight function $w : P \to \mathbb{R}^+$, the sensitivity of $s_i \in P$ is $\sigma_P(s_i) := \inf\{\beta \geq 0 \mid w_i \cdot d(s_i, F) \leq \beta \sum_{j \in [n]} w_j \cdot d(s_j, F), \forall F \in \mathcal{F}\}$. The total sensitivity of P is defined by $G_P = \sum_{s_i \in P} \sigma_P(s_i)$.

We also recall the definition of VC-dimension and shattering dimension and show their relations.

Definition 56. (VC-dimension and shattering dimension) Let $P = \{s_1, \ldots, s_n\}$ be an instance of a shape fitting problem $(\mathbb{R}^d, \mathcal{F}, d)$. Suppose w_i is the weight of s_i. We consider the range space $(P, \mathcal{R})$, where $\mathcal{R}$ is a family of subsets $R_{F,r}$ of P defined as follows: given an $F \in \mathcal{F}$ and $r \geq 0$, let $R_{F,r} = \{s_i \in P \mid w_i \cdot d(s_i, F) \geq r\} \in \mathcal{R}$ consist

[1] The notation $(1 \pm \varepsilon)B$ means the interval $[(1 - \varepsilon)B, (1 + \varepsilon)B]$.

of those points s_i whose weighted distance to the shape F is at least r. We denote the VC-dimension of the instance of the instance P by $\dim_{VC}(P)$, to be the largest integer m, such that there exists a weight function w and a subset $A \subseteq P$ satisfying the property $|\{A \cap R_{F,r} \mid F \in \mathcal{F}, r \geq 0\}| = 2^{|A|}$. We also denote the shattering dimension of the instance P by $\dim(P)$, to be the smallest integer m, such that for any weight function w and $A \subseteq P$ of size $|A| = a \geq 2$, we have $|\{A \cap R_{F,r} \mid F \in \mathcal{F}, r \geq 0\}| \leq a^m$.

The next lemma offers a natural way to bound the VC-dimension of a range space by bounding the shattering dimension.

Lemma 57. ([58, Lemma 5.14]) If $(X, \mathcal{R})$ is a range space with shattering dimension d, then its VC-dimension is bounded by $O(d \log d)$.

A framework for constructing coresets of shape fitting instances. The main idea of the framework in [45] is based on the following lemma.

Lemma 58. Given any instance $P = \{s_1, \ldots, s_n\}$ of a shape fitting problem $(\mathbb{R}^d, \mathcal{F}, d)$, any weight function $w : P \to \mathbb{R}^+$, and any $\varepsilon \in (0, 1]$, there exists an ε-coreset for P of size $O((\frac{G_P}{\varepsilon})^2 \dim_{VC}(P))$.

We take 1-median problem as an example. In this example, $\mathcal{F} = \mathbb{R}^d$ and d is the Euclidean distance function. The framework mainly consists of the following steps.

1. Compute $F \subseteq \mathcal{F}$ which is a constant-factor approximate shape in $\mathcal{F}$. For 1-median, we only need to choose $F = s^* = \arg\min_{s \in P} \sum_{s_i \in P} w_i \cdot d(s_i, s)$. It is not hard to verify that $F = s^*$ is a 2-approximate shape.

2. For each $s_i \in P$, compute its projection $s'_i = \arg\min_{s \in F} d(s_i, s)$. Let $P' = \{s'_i : s_i \in P\}$ be the collection of all projection points. For 1-median, we observe that $s'_i = s^*$ for all $i \in [n]$.

3. For each $s_i \in P$, compute an upper bound $\sigma_P(s_i)$ of the sensitivity (see [106, Theorem 7] for details). Compute $G_P = \sum_{i \in [n]} \sigma_P(s_i)$ as an upper bound of the total sensitivity. For 1-median, we let $\sigma_P(s_i) = \frac{w_i \cdot d(s_i, s)}{\sum_{j \in [n]} w_j \cdot d(s_j, s)} + \frac{3}{n}$ as an upper bound [106]. By these values, we have $G_P = \sum_{i \in [n]} \sigma_P(s_i) = 4$.

4. Sample points from P with probabilities proportional to $\sigma_P(s_i)$. Formally speaking, we modify the weight w_i to be $\frac{w_i \cdot G_P}{\sigma_P(s_i)}$ and sample the point s_i with probability $\frac{\sigma_P(s_i)}{G_P}$. This technique is called importance sampling, see [45, Section 4.1] for more details. By Lemma 58, we need to take $O((\frac{G_P}{\varepsilon})^2 \dim_{VC}(P))$ samples. For 1-median, the VC-dimension is $d+1$. Thus, there exists an ε-coreset for P of size $O(d\varepsilon^{-2})$.

5.2 Generalized Shape Fitting Problems and Generalized Coresets

We recall the definitions of the two stochastic shape fitting problems in this chapter.

Definition 59. For a set of points $P \in \mathbb{R}^d$, and a k-point set $F = \{f_1, \ldots, f_k) \mid f_i \in \mathbb{R}^d, 1 \leq i \leq k\}$, we define $K(P,F) = \max_{s \in P} \min_{1 \leq i \leq k} d(s, f_i)$ as the k-center value of F w.r.t. P. We use $\mathcal{F}$ to denote the family of all k-point sets in $\mathbb{R}^d$. Given a set $\mathcal{P}$ of n stochastic points (in either the existential or locational uncertainty model) in $\mathbb{R}^d$, and a k-point set $F \in \mathcal{F}$, we define the expected k-center value of F w.r.t $\mathcal{P}$ as

$$K(\mathcal{P}, F) = \mathbb{E}_{P \sim \mathcal{P}}[K(P,F)].$$

In the stochastic minimum k-center problem, our goal is to find a k-point set $F \in \mathcal{F}$ which minimizes $K(\mathcal{P}, F)$. In this dissertation, we assume that both the dimensionality d and k are fixed constants.

Definition 60. Given a set P of n points in $\mathbb{R}^d$, and a j-flat $F \in \mathcal{F}$ ($0 \leq j \leq d-1$), where $\mathcal{F}$ is the family of all j-flats in $\mathbb{R}^d$, we define the j-flat-center value of F w.r.t. P to be $J(P,F) = \max_{s \in P} d(s,F)$, where $d(s,F) = \min_{f \in F} d(s,f)$ is the distance between point s and j-flat F. Given a set $\mathcal{P}$ of n stochastic points (in either the existential or locational model) in $\mathbb{R}^d$, and a j-flat $F \in \mathcal{F}$ ($0 \leq j \leq d-1$), we define the expected j-flat-center value of F w.r.t. $\mathcal{P}$ to be

$$J(\mathcal{P}, F) = \mathbb{E}_{P \sim \mathcal{P}}[J(P,F)].$$

In the stochastic minimum j-flat-center problem, our goal is to find a j-flat F which minimizes J($\mathcal{P}, F$).

In the following, we define generalized shape fitting problems. In fact, we can consider stochastic k-center and stochastic j-flat-center as generalized shape fitting problems. We also give the definition of generalized coresets and introduce a framework for constructing generalized coresets.

Generalized Shape Fitting Problems and Generalized Coresets. In this section, we define the generalized shape fitting problems, which are defined over a collection of (weighted) point sets, (recall the traditional shape fitting problems are defined over a set of (weighted) points). We use $\mathbb{R}^d$ to denote the d-dimensional Euclidean space. Let d(s, s') denote the Euclidean distance between point s and s' and d(s, F) = $\min_{s' \in F}$ d(s, s') for any $F \subset \mathbb{R}^d$. Let $\mathbb{U}^d = \{P \mid P \subset \mathbb{R}^d, |P| \text{ is finite}\}$ be the collection of all finite discrete point sets in $\mathbb{R}^d$.

Definition 61. (Generalized shape fitting problems) A generalized shape fitting problem is specified by a triple ($\mathbb{R}^d, \mathcal{F}$, dist). Here the set $\mathcal{F}$ of shapes is a family of subsets of $\mathbb{R}^d$ (e.g., all k-point sets, or all j-flats), and dist : $\mathbb{U}^d \times \mathcal{F} \to \mathbb{R}^{\geq 0}$ is a generalized distance function, defined as dist(P, F) = $\max_{s \in P}$ d(s, F) for a point set $P \in \mathbb{U}^d$ and a shape $F \in \mathcal{F}$. [2] An instance S of the generalized shape fitting problem is a (weighted) collection $\{S_1, \ldots, S_m\}$ ($S_i \in \mathbb{U}^d$) of point sets, and each S_i has a positive weight $w_i \in \mathbb{R}^+$. For any shape $F \in \mathcal{F}$, define the total generalized distance from S to F to be dist(S, F) = $\sum_{S_i \in S} w_i \cdot$ dist(S_i, F). Given an instance S, our goal is to find a shape $F \in \mathcal{F}$, which minimizes the total generalized distance dist(S, F).

If we replace $\mathbb{U}^d$ with $\mathbb{R}^d$, the above definition reduces to the traditional shape fitting problem, see Definition 53. Here, we give an example for Definition 61.

Example. Consider a generalized shape fitting problem where $\mathcal{F}$ is the collection of all 2-point sets in $\mathbb{R}^2$. In this case, for a point $s \in \mathbb{R}^2$ and a 2-point set $F \in \mathcal{F}$, the function d(s, F) = $\min_{f \in F}$ d(s, f) is the Euclidean distance between s and its

[2] Note that dist may not be a metric in general.

generalized k-median problem (GKM) is to find a k-point set F which minimizes the total generalized distance $\mathsf{K}(\mathsf{S}, F)$.

Recall that a generalized ε-coreset is a sub-collection $\mathcal{S} \subseteq \mathsf{S}$ of point sets, together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$, such that for any k-point set $F \in \mathcal{F}$, we have $\prod_{S \in \mathcal{S}} w'(S) \cdot \mathsf{K}(S, F) \in (1 \pm \varepsilon) \prod_{S \in \mathsf{S}} w(S) \cdot \mathsf{K}(S, F)$ (or $\mathsf{K}(\mathcal{S}, F) \in (1 \pm \varepsilon)\mathsf{K}(\mathsf{S}, F)$). This generalized coreset will serve as the SKC-Coreset for the original stochastic k-center problem.

Our main lemma asserts that a constant sized generalized coreset exists, as follows.

Lemma 73. (main lemma) Given an instance $\mathcal{P}$ of n stochastic points in $\mathbb{R}^{\mathsf{d}}$, let $\mathcal{E}(\mathcal{P})$ be the collection of all additive ε-coresets. There exists a generalized ε-coreset $\mathcal{S} \subseteq \mathcal{E}(\mathcal{P})$ of cardinality $|\mathcal{S}| = O(\varepsilon^{-(\mathsf{d}+2)} dk^4)$, together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$, which satisfies that for any k-point set $F \in \mathcal{F}$,

$$\prod_{S \in \mathcal{S}} w'(S) \cdot \mathsf{K}(S, F) \in (1 \pm \varepsilon) \prod_{S \in \mathcal{E}(\mathcal{P})} \Pr_{\mathsf{P} \sim \mathcal{P}}[\mathcal{E}(P) = S] \cdot \mathsf{K}(S, F).$$

Now, we prove Lemma 73 by showing a constant upper bound on the cardinality of a generalized ε-coreset. This is done by applying Lemma 65 and providing constant upper bounds for both the total sensitivity and the generalized dimension of the generalized k-median instance.

Given an instance $\mathsf{S} = \{S_{\mathsf{i}} \mid S_{\mathsf{i}} \in \mathbb{U}^{\mathsf{d}}, 1 \leq i \leq n\}$ of a generalized k-median problem with a weight function $w : \mathsf{S} \to \mathbb{R}^+$, we denote F^* to be the k-point set which minimizes the total generalized distance $\mathsf{K}(\mathsf{S}, F) = \prod_{S \in \mathsf{S}} w(S) \cdot \mathsf{K}(S, F)$ over all $F \in \mathcal{F}$. W.l.o.g., we assume that $\mathsf{K}(\mathsf{S}, F^*) > 0$. Since if $\mathsf{K}(\mathsf{S}, F^*) = 0$, there are at most k different points in the instance.

We first construct a projection instance P^* of a weighted k-median problem for S, and relate the total sensitivity G_{S} to $\mathsf{G}_{\mathsf{P}^*}$. Recall that $\mathsf{G}_{\mathsf{S}} = \prod_{S \in \mathsf{S}} \sigma_{\mathsf{S}}(S)$ is the total sensitivity of S. Our construction of P^* is as follows. For each point set $S_{\mathsf{i}} \in \mathsf{S}$, assume that $F_{\mathsf{i}}^* \in \mathcal{F}$ is the k-point set satisfying that $F_{\mathsf{i}}^* = \mathsf{argmax}_{\mathsf{F}} \frac{\mathsf{w}(\mathsf{S}_i) \cdot \mathrm{K}(\mathsf{S}_i, \mathsf{F})}{\mathrm{K}(\mathsf{S}, \mathsf{F})}$, i.e., the sensitivity $\sigma_{\mathsf{S}}(S_{\mathsf{i}})$ of S_{i} is equal to $\frac{\mathsf{w}(\mathsf{S}_i)\mathrm{K}(\mathsf{S}_i, \mathsf{F}_i^*)}{\mathrm{K}(\mathsf{S}, \mathsf{F}_i^*)}$. Let $s_{\mathsf{i}}^* \in S_{\mathsf{i}}$ denote the point

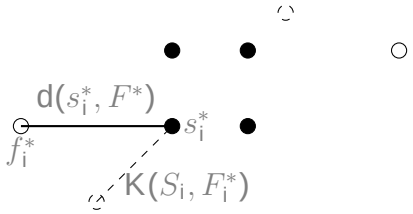

Figure 5-2: In the figure, S_i is the black point set, F^* is the white point set, and F_i^* is the dashed point set. Here, $s_i^* \in S_i$ is the farthest point to F_i^* satisfying $d(s_i^*, F_i^*) = K(S_i, F_i^*)$, and $f_i^* \in F^*$ is the closest point to s_i^* satisfying $d(s_i^*, f_i^*) = d(s_i^*, F^*)$.

farthest to F_i^* (breaking ties arbitrarily). Let $f_i^* \in F^*$ denote the point closest to s_i^* (breaking ties arbitrarily). Denote P^* to be the multi-set $\{f_i^* \mid S_i \in S\}$, and denote the weight function $w' : P^* \to \mathbb{R}^+$ to be $w'(f_i^*) = w(S_i)$ for any $i \in [n]$. Thus, P^* is a weighted k-median instance in $\mathbb{R}^d$ with a weight function w'. See Figure 5-2 for an example of the construction of P^*.

Lemma 74. Given an instance $S = \{S_i \mid S_i \in \mathbb{U}^d, 1 \le i \le n\}$ of a generalized k-median problem in $\mathbb{R}^d$ with a weight function $w : S \to \mathbb{R}^+$, let P^* be its projection instance. Then, we have $G_S \le 2G_{P^*} + 1$.

Proof. First note that we have the following fact. Given $i, j \in [n]$, recall that $s_j^* \in S_j$ is the farthest point to F_j^*, and $f_j^* \in F^*$ is the closest point to s_j^*. Let $f \in F_i^*$ be the point closest to s_j^*.

$$\begin{aligned} K(S_j, F_i^*) + K(S_j, F^*) &\ge d(s_j^*, F_i^*) + d(s_j^*, F^*) = d(s_j^*, F_i^*) + d(s_j^*, f_j^*) \\ &= d(s_j^*, f) + d(s_j^*, f_j^*) \ge d(f_j^*, f) \ge d(f_j^*, F_i^*), \end{aligned} \tag{5.3}$$

The first inequality follows from the definitions of $K(S_j, F_i^*)$ and $K(S_j, F^*)$. The first equality follows from the definition of f_j^*. The second inequality follows from the triangle inequality, and the last inequality is by the definition of $d(f_j^*, F_i^*)$.

Then we have the following fact:

$$\begin{aligned} \sum_{f \in P^*} w'(f) \cdot d(f, F_i^*) &= \sum_{f_j^* \in P^*} w'(f_j^*) \cdot d(f_j^*, F_i^*) \le \sum_{S_j \in S} w(S_j) \cdot \left[K(S_j, F^*) + K(S_j, F_i^*) \right] \\ &= K(S, F^*) + K(S, F_i^*) \le 2K(S, F_i^*), \end{aligned} \tag{5.4}$$

since $\mathsf{K}(S, F^*) \leq \mathsf{K}(S, F_i^*)$ and Inequality (5.3).

Let $f' \in F_i^*$ be the point closest to f_i^*. We also notice the following fact:

$$\begin{aligned}\mathsf{K}(S_i, F^*) + \mathsf{d}(f_i^*, F_i^*) &\geq \mathsf{d}(s_i^*, f_i^*) + \mathsf{d}(f_i^*, F_i^*) = \mathsf{d}(s_i^*, f_i^*) + \mathsf{d}(f_i^*, f') \\ &\geq \mathsf{d}(s_i^*, f') \geq \mathsf{d}(s_i^*, F_i^*) = \mathsf{K}(S_i, F_i^*). \end{aligned} \tag{5.5}$$

The first inequality follows from the definition of f_i^*, the second inequality follows from the triangle inequality, and the last inequality follows from the definition of $\mathsf{d}(s_i^*, F_i^*)$.

Now we are ready to analyze $\sigma_S(S_i)$ for some $S_i \in S$. We can see that

$$\begin{aligned} w(S_i) \cdot \mathsf{K}(S_i, F_i^*) &\leq w(S_i) \cdot \mathsf{K}(S_i, F^*) + w(S_i) \cdot \mathsf{d}(f_i^*, F_i^*) && \text{[by (5.5)]} \\ &\leq w(S_i) \cdot \mathsf{K}(S_i, F^*) + \sigma_{P^*}(f_i^*) \cdot \sum_{f \in P^*} w'(f) \cdot \mathsf{d}(f, F_i^*) && \text{[by the definition of } \sigma_{P^*}\text{]} \\ &\leq w(S_i) \cdot \mathsf{K}(S_i, F^*) + 2\sigma_{P^*}(f_i^*) \cdot \mathsf{K}(S, F_i^*) && \text{[by (5.4)]} \\ &= \frac{w(S_i) \cdot \mathsf{K}(S_i, F^*)}{\mathsf{K}(S, F_i^*)} \cdot \mathsf{K}(S, F_i^*) + 2\sigma_{P^*}(f_i^*) \cdot \mathsf{K}(S, F_i^*) \\ &\leq \left(\frac{w(S_i) \cdot \mathsf{K}(S_i, F^*)}{\mathsf{K}(S, F^*)} + 2\sigma_{P^*}(f_i^*) \right) \mathsf{K}(S, F_i^*). && \text{[by } \mathsf{K}(S, F_i^*) \geq \mathsf{K}(S, F^*)\text{]} \end{aligned}$$

Finally, we bound the total sensitivity as follows:

$$\mathsf{G}_S = \sum_{S_i \in S} \sigma_S(S_i) \leq \sum_{S_i \in S} \left(\frac{w(S_i) \cdot \mathsf{K}(S_i, F^*)}{\mathsf{K}(S, F^*)} + 2\sigma_{P^*}(f_i^*) \right) = 1 + 2\mathsf{G}_{P^*}.$$

This finishes the proof of the lemma. □

Since P^* is an instance of a weighted k-median problem, we know that the total sensitivity G_{P^*} is at most $2k + 1$, by [77, Theorem 9]. [6] Then combining Lemma 74, we have the following lemma which bounds the total sensitivity of G_S.

Lemma 75. Consider an instance S of a generalized k-median problem $(\mathbb{R}^d, \mathcal{F}, \mathsf{K})$. The total sensitivity G_S is at most $4k + 3$.

[6] Theorem 9 in [77] bounds the total sensitivity for the unweighted version. However, the proof can be extended to the weighted version in a straightforward way.

time, we enumerate at most $O(n^{\tilde{O}(\varepsilon^{-(2d+2)}d^2k^5)})$ weighted sub-collections. Moreover, computing the optimal k-point set for each sub-collection costs constant time. Then the total running time is at most $O(n^{\tilde{O}(\varepsilon^{-(2d+2)}d^2k^5)})$. Thus, we have the following corollary.

Corollary 79. If both k and d are constants, given an instance $\mathcal{P}$ of n stochastic points in $\mathbb{R}^d$ in the existential uncertainty model, there exists a PTAS for the stochastic minimum k-center problem in $O(n^{\tilde{O}(\varepsilon^{-(2d+2)}d^2k^5)})$ time.

Enumerating possible generalized ε-coresets. Given an instance $S = \{S_i \mid S_i \in \mathbb{U}^d, 1 \le i \le N\}$ of a generalized k-median problem in $\mathbb{R}^d$ with a weight function $w : S \to \mathbb{R}^+$, now we show how to enumerate polynomially many sub-collections $\mathcal{S}_i \subseteq S$ together with their weight functions, such that there exists a generalized ε-coreset of S. Recall that $\sigma_S(S_i)$ is the sensitivity of S_i, and $G_S = \sum_{i\in[N]} \sigma_S(S_i)$ is the total sensitivity. Also recall that dim(S) is the generalized dimension of S. Define $q(S_i) = \sigma_S(S_i) + 1/N$ for $1 \le i \le M$, and define $q_S = \sum_{1\le i\le N} q(S_i)$. Note that $q_S = G_S + 1 \le 4k + 4$ by Lemma 75. Our algorithm is as follows.

1. Let $M = O((\frac{q_S}{\varepsilon})^2 \dim(S))$. Let $L = \frac{10}{\varepsilon}(\log M + \log N + \log k)$.
2. Enumerate all collections $\mathcal{S}_i \subseteq S$ of cardinality at most M. Note that we only need to enumerate at most N^M collections.
3. For a collection $\mathcal{S} \subseteq S$, w.l.o.g., assume that $\mathcal{S} = \{S_1, S_2, \cdots, S_m\}$ $(m \le M)$. Enumerate all sequences $(1+\varepsilon)^{a_1}, \ldots, (1+\varepsilon)^{a_m}$ where each $0 \le a_i \le L$ is an integer.
4. Given a collection $\mathcal{S} = \{S_1, S_2, \cdots, S_m\}$ and a sequence $(1+\varepsilon)^{a_1}, \ldots, (1+\varepsilon)^{a_m}$, we construct a weight function $w' : \mathcal{S} \to \mathbb{R}^+$ as follows: for a point set $S_i \in \mathcal{S}$, denote $w'(S_i)$ to be $(1+\varepsilon)^{a_i} \cdot w(S_i)/M$. Recall that $w(S_i)$ is the weight of $S_i \in S$.

Analysis. Recall that given an instance $\mathcal{P}$ of a stochastic minimum k-center problem, we first reduce to an instance $S = \mathcal{E}(\mathcal{P})$ of a generalized k-median problem. Note

that the cardinality of S is at most $n^{O(k/\varepsilon^d)}$, and the cardinality of a generalized ε-coreset is at most $M = \tilde{O}(\varepsilon^{-(d+2)}d^2k^4)$ by Theorem 77. Thus, we enumerate at most $N^M = n^{\tilde{O}(\varepsilon^{-(2d+2)}d^2k^5)}$ polynomially many sub-collections $\mathcal{S}_i \subseteq S$. For each collection $\mathcal{S}_i$, we construct at most $M^{L+1} = O(n^{O(k/\varepsilon^d)})$ polynomially many weight functions. In total, we enumerate $N^M \cdot M^{L+1} = O(n^{\tilde{O}(\varepsilon^{-(2d+2)}d^2k^5)})$ polynomially many weighted sub-collections.

It remains to show that there exists a generalized ε-coreset of S. We first have the following lemma.

Lemma 80. Given an instance $S = \{S_i \mid S_i \in \mathbb{U}^d, 1 \le i \le N\}$ of a generalized k-median problem in $\mathbb{R}^d$ with a weight function $w : S \to \mathbb{R}^+$, there exists a generalized ε-coreset $\mathcal{S} \subseteq S$ with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$, such that

$$\sum_{S \in \mathcal{S}} w'(S) \cdot K(S, F) \in (1 \pm \varepsilon) \sum_{S \in S} w(S) \cdot K(S, F).$$

The cardinality of $\mathcal{S}$ is at most $M = O((\frac{q_S}{\varepsilon})^2 \dim(S))$. Moreover, each weight $w'(S)$ ($S \in \mathcal{S}$) has the form that $w'(S) = \frac{c \cdot q_S \cdot w(S)}{q(S) \cdot M}$, where $1 \le c \le M$ is an integer.

Proof. For each $S \in S$, let $g_S : \mathcal{F} \to \mathbb{R}^+$ be defined as $g_S(F) = w(S) \cdot K(S, F)/q(S)$. Let $D = \{g_S \mid S \in S\}$ be a collection, together with a weight function $w'' : D \to \mathbb{R}^+$ defined as $w''(g_S) = q(S)$. Note that for any k-point set $F \in \mathcal{F}$, we have that

$$\sum_{g_S \in G} w''(g_S) \cdot g_S(F) = \sum_{S \in S} w(S) \cdot K(S, F) = K(S, F).$$

By Theorem 4.1 in [45], we can randomly sample (with replacement) a collection $\mathcal{S} \subseteq D$ of cardinality at most $M = O((\frac{q_S}{\varepsilon})^2 \dim(S))$, together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$ defined as $w'(g_S) = q_S/M$. Then the multi-set $\mathcal{S}$ satisfies that for every $F \in \mathcal{F}$,

$$\sum_{g_S \in \mathcal{S}} w'(g_S) \cdot g_S(F) \in (1 \pm \varepsilon) \sum_{g_S \in G} w''(g_S) \cdot g_S(F) = (1 \pm \varepsilon) K(S, F).$$

By the definition of g_S and w', we prove the lemma. □

We are ready to prove the following lemma.

Lemma 81. Among all sub-collections $\mathcal{S} \subseteq \mathsf{S}$ of cardinality at most $M = O((\frac{q_S}{\varepsilon})^2 \dim(\mathsf{S}))$, together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$ of the form $w'(S_i) = (1+\varepsilon)^{a_i} \cdot w(S_i)/M$ ($0 \leq a_i \leq \frac{10(\log M + \log N + \log k)}{\varepsilon}$ is an integer), there exists a generalized ε-coreset of S.

Proof. By Lemma 80, there exists a generalized ε-coreset $\mathcal{S} \subseteq \mathcal{S}$ of cardinality at most M together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$ defined as follows: each weight $w'(S)$ ($S \in \mathcal{S}$) has the form that $w'(S) = \frac{c_S \cdot q_S \cdot w(S)}{q(S) \cdot M}$ for some integer $1 \leq c_S \leq M$. W.l.o.g., we assume that $\mathcal{S} = \{S_1, S_2, \ldots, S_m \mid S_i \in \mathsf{S}\}$ ($m \leq M$).

By the definition of $q(S)$, we have that $1/N \leq q(S) \leq q_S = G_S + 1 \leq 4k + 4$. Then we conclude that for each $S \in \mathcal{S}$,

$$1 \leq \frac{c_S \cdot q_S}{q(S)} \leq (4k+4)MN.$$

For $1 \leq i \leq m$, let $a_i = \lfloor \log_{1+\varepsilon}(\frac{c_{S_i} \cdot q_S}{q(S_i)}) \rfloor$. Note that each a_i satisfies that $0 \leq a_i \leq \frac{10(\log M + \log N + \log k)}{\varepsilon}$. Thus, we have enumerated the following sub-collection $\mathcal{S} = \{S_1, S_2, \cdot, S_m \mid S_i \in \mathsf{S}\}$ with a weight function $w'' : \mathcal{S} \to \mathbb{R}^+$, such that $w''(S_i) = (1+\varepsilon)^{a_i} \cdot w(S_i)/M$. Moreover, for any k-point set F, we have the following inequality.

$$\begin{aligned}\square_{1 \leq i \leq m} w''(S_i) \cdot K(S_i, F) &= \square_{1 \leq i \leq m} \frac{(1+\varepsilon)^{a_i} \cdot w(S_i)}{M} \cdot K(S_i, F) \in (1 \pm \varepsilon) \square_{1 \leq i \leq m} \frac{c_S \cdot q_S \cdot w(S)}{q(S_i) \cdot M} \cdot K(S_i, F) \\ &= (1 \pm \varepsilon) \square_{1 \leq i \leq m} w'(S_i) \cdot K(S_i, F) \in (1 \pm 3\varepsilon) \square_{S \in \mathsf{S}} w(S) \cdot K(S, F).\end{aligned}$$

The last inequality is due to the assumption that the sub-collection $\mathcal{S}$ with a weight function w' is a generalized ε-coreset of S. Let $\varepsilon' = \varepsilon/3$, we prove the lemma. □

5.3.2 Locational uncertainty model

Next, we consider the stochastic minimum k-center problem in the locational uncertainty model. Given an instance of m nodes $v_1, \ldots, v_m$ which may locate in the point set $\mathcal{P} = \{s_1, \ldots, s_n \mid s_i \in \mathbb{R}^d, 1 \leq i \leq n\}$, our construction of additive ε-coresets and the method for bounding the total sensitivity is exactly the same as in the existential

uncertainty model. The only difference is that for an additive ε-coreset S, how to compute the probability $\Pr_{P\sim\mathcal{P}}[\mathcal{E}(P) = S] = \sum_{P:P\sim\mathcal{P},\mathcal{E}(P)=S} \Pr[\models P]$. Here, $P \sim \mathcal{P}$ is a realized point set according to the probability distribution of $\mathcal{P}$. Run $\mathbb{A}(S)$, and construct a Cartesian grid $G(S)$. Denote $T(S) = \left(\cup_{P:P\sim\mathcal{P},\mathcal{E}(P)=S} P\right) \setminus S$ to be the collection of all points s which might be contained in some realization $P \sim \mathcal{P}$ with $\mathcal{E}(P) = S$. Recall that $\mathcal{C}(S) = \{C \in G \mid |C \cap S| = 1\}$ is the collection of d-dimensional Cartesian cells C which contains a point $s^C \in S$. By Lemma 70, for any realization P with $\mathcal{E}(P) = S$, we have the following observations.

1. For any cell $C \notin \mathcal{C}(S)$, $C \cap P = \emptyset$. It means that for any point $s \in C \cap \mathcal{P}$, we have $s \notin T(S)$.

2. For any cell $C \in \mathcal{C}(S)$ and any point $s' \in C \cap \mathcal{P}$ with a smaller index than that of s^C, we have $s' \notin P$. It means that $s' \notin T(S)$.

By the above observations, we conclude that $T(S)$ is the collection of those points s' belonging to some cell $C \in \mathcal{C}(S)$ and with a larger index than that of s^C.

Then we reduce the counting problem $\Pr_{P\sim\mathcal{P}}[\mathcal{E}(P) = S]$ to a family of bipartite holant problems. We first give the definition of holant problems.

Definition 82. An instance of a holant problem is a tuple $\Lambda = \left(G(V, E), (g_u)_{u\in V}, (w_e)_{e\in E}\right)$, where for every $u \in V$, $g_u : \{0,1\}^{E_u} \to \mathbb{R}^+$ is a function, where E_u is the set of edges incident to $\vec{u}$. For every assignment $\sigma \in \{0,1\}^E$, we define the weight of σ as

$$w_\Lambda(\sigma) \triangleq \prod_{u\in V} g_u(\sigma\mid_{E_u}) \prod_{e\in\sigma} w_e.$$

Here $\sigma\mid_{E_u}$ is the assignment of E_u with respect to the assignment σ. We denote the value of the holant problem $Z(\Lambda) \triangleq \sum_{\sigma\in\{0,1\}^E} w_\Lambda(\sigma)$.

For a counting problem $\Pr_{P\sim\mathcal{P}}[\mathcal{E}(P) = S]$, w.l.o.g., we assume that $S = \{s_1, \ldots, s_{|S|}\}$. Then we construct a family of holant instance $\Lambda_{\mathcal{L}}$ as follows.

1. Enumerate all integer sequences $\mathcal{L} = (l_1, \ldots, l_{|S|}, l_t)$ such that $\sum_{1\le i\le|S|} l_i + l_t = n$, $l_i \ge 1$ $(1 \le i \le |S|)$, and $l_t \ge 0$. Let L be the collection of all these integer sequences $\mathcal{L}$.

2. For a sequence $\mathcal{L}$, assume that $\Lambda_{\mathcal{L}} = \left(G(U, V, E), (g_u)_{u \in U \cup V}\right)$ is a holant instance on a bipartite graph, where $V = \{v_1, \ldots, v_n\}$, and $U = S \cup \{t\}$ (we use vertex t to represent the collection $T(S)$).

3. The weight function $w : E \to \mathbb{R}^+$ is defined as follows:

   (a) For a vertex $v_i \in V$ and a vertex $s_j \in S$, $w_{ij} = p_{ij}$.

   (b) For a vertex $v_i \in V$ and $t \in V$, $w_{it} = \prod_{s_j \in T(S)} p_{ij}$.

4. For each vertex $v \in V$, the function $g_v = (= 1)$. [10] For each vertex $s_i \in S$, the function $g_{s_i} = (= l_i)$, and the function $g_t = (= l_t)$.

Since each $S \in \mathcal{E}(\mathcal{P})$ is of constant size, we only need to enumerate at most $O(m^{|S|+1}) = \text{poly}(n)$ integer sequences $\mathcal{L}$. Given an integer sequence $\mathcal{L} = (l_1, \ldots, l_{|S|}, l_t)$, we can see that $Z(\Lambda_{\mathcal{L}})$ is exactly the probability that l_i nodes are realized at point $s_i \in S$ ($\forall 1 \leq i \leq |S|$), and l_t nodes are realized inside the point set $T(S)$. Then by Lemma 70, we have the following equality:

$$\Pr_{P \sim \mathcal{P}}[\mathcal{E}(P) = S] = \prod_{\mathcal{L} \in L} Z(\Lambda_{\mathcal{L}}).$$

It remains to show that we can compute each $Z(\Lambda_{\mathcal{L}})$ efficiently. Fortunately, we have the following lemma.

Lemma 83. ([67],[103]) For any bipartite graph $\Lambda_{\mathcal{L}}$ with a specified integer sequence $\mathcal{L}$, there exists an FPRAS to compute the holant value $Z(\Lambda_{\mathcal{L}})$.

Thus, we have the following theorem.

Theorem 84. If both k and d are constants, given an instance of m stochastic nodes in $\mathbb{R}^d$ in the locational uncertainty model, there exists a PTAS for the stochastic minimum k-center problem.

Combining Theorem 77 and 84, we obtain the main result Theorem 12.

[10] Here the function $g_v = (= i)$ means that the function value g_v is 1 if exactly i edges incident to v are of value 1 in the assignment. Otherwise, $g_v = 0$

5.4 Stochastic Minimum j-Flat-Center

In this section, we consider a generalized shape fitting problem, the minimum j-flat-center problem in the stochastic models. Let $\mathcal{F}$ be the family of all j-flats in $\mathbb{R}^d$. Our main technique is to construct an SJFC-Coreset of constant size, which satisfies that for any j-flat $F \in \mathcal{F}$, we can use the SJFC-Coreset to obtain a $(1\pm\varepsilon)$-estimation for the expected j-flat-center value J$(\mathcal{P}, F)$. Then since the SJFC-Coreset is of constant size, we have a polynomial system of constant size to compute the optimum in constant time.

Let $B = \sum_{1\le i\le n} p_i$ be the total probability. We discuss two different cases. If $B < \varepsilon$, we reduce the problem to a weighted j-flat-median problem, which has been studied in [106]. If $B \ge \varepsilon$, the construction of an SJFC-Coreset can be divided into two parts. We first construct a convex hull, such that with high probability (say $1-\varepsilon$) that all points are realized inside the convex hull. Then we construct a collection of point sets to estimate the contribution of points insider the convex hull. On the other hand, for the case that some point appears outside the convex hull, we again reduce the problem to a weighted j-flat-median problem. The definition of the weighted j-flat-median problem is as follows.

Definition 85. For some $0 \le j \le d-1$, let $\mathcal{F}$ be the family of all j-flats in $\mathbb{R}^d$. Given a set P of n points in $\mathbb{R}^d$ together with a weight function $w : P \to \mathbb{R}^+$, denote cost$(P, F) = \sum_{s_i\in P} w_i \cdot d(s_i, F)$. A weighted j-flat-median problem is to find a shape $F \in \mathcal{F}$ which minimizes the value cost(P, F).

5.4.1 Case 1: $B < \varepsilon$

In the first case, we show that the minimum j-flat-center problem can be reduced to a weighted j-flat-median problem. We need the following lemmas.

Lemma 86. If $B < \varepsilon$, for any j-flat $F \in \mathcal{F}$, we have $\sum_{s_i\in\mathcal{P}} p_i \cdot d(s_i, F) \in (1\pm\varepsilon)\cdot$ J$(\mathcal{P}, F)$.

Proof. For a j-flat $F \in \mathbb{R}^d$, w.l.o.g., we assume that $d(s_i, F)$ is non-decreasing in i.

Thus, we have

$$J(\mathcal{P}, F) = \prod_{i \in [n]} p_i \cdot d(s_i, F) \cdot \prod_{j > i} (1 - p_j)$$

Since $B < \varepsilon$, for any $i \in [n]$, we have that $1 - \varepsilon \leq 1 - \prod_{j \in [n]} p_j \leq \prod_{j > i}(1 - p_j) \leq 1$. So we prove the lemma. □

By Lemma 86, we reduce the original problem to a weighted j-flat-median problem, where each point $s_i \in \mathcal{P}$ has weight p_i. We then need the following lemma to bound the total sensitivity.

Lemma 87. (Theorem 18 in [106]) [11] Consider the weighted j-flat-median problem where $\mathcal{F}$ is the set of all j-flats in $\mathbb{R}^d$. The total sensitivity of any weighted n-point set is $O(j^{1.5})$.

On the other hand, we know that the dimension of the weighted j-flat-median problem is $O(jd)$ by [45]. Then by Lemma 65, there exists an ε-coreset $\mathcal{S} \subseteq \mathcal{P}$ of cardinality $O(j^4 d\varepsilon^{-2} \log(jd)) = \tilde{O}(j^4 d\varepsilon^{-2})$ to estimate the j-flat-median value $\prod_{s_i \in \mathcal{P}} p_i \cdot d(s_i, F)$ for any j-flat $F \in \mathcal{F}$. [12] Moreover, we can compute a constant approximation j-flat in $O(ndj^{O(j^2)})$ time by [44]. Then by [106], we can construct an ε-coreseet $\mathcal{S}$ in $O(ndj^{O(j^2)})$ time. Combining Lemma 86, we conclude the main lemma in this subsection.

Lemma 88. Given an instance $\mathcal{P}$ of n stochastic points in $\mathbb{R}^d$, if the total probability $\prod_i p_i < \varepsilon$, there exists an SJFC-Coreset of cardinality $\tilde{O}(j^4 d\varepsilon^{-2})$ for the minimum j-flat-center problem. Moreover, we have an $O(ndj^{O(j^2)})$ time algorithm to compute the SJFC-Coreset.

5.4.2 Case 2: $B \geq \varepsilon$

Note that if F is a j-flat, the function $d(x, F)^2$ has a linearization. Here, a linearization is to map the function $d(x, F)^2$ to a k-variate linear function through variate

[11] Theorem 18 in [106] bounds the total sensitivity for the unweighted version. However, the proof can be extended to the weighted version in a straightforward manner.

[12] We remark that for the j-flat-median problem, Feldman and Langberg [45] showed that there exists a coreset of size $O(jd\varepsilon^{-2})$. However, it is unclear how to generalize their technique to weighted version.

embedding. The number k is called the dimension of the linearization, see [8]. We have the following lemma to bound the dimension of the linearization.

Lemma 89. ([46]) Suppose F is a j-flat in $\mathbb{R}^d$, the function $d(x,F)^2$ ($x \in \mathbb{R}^d$) has a linerization. Let D be the dimension of the linearization. If $j = 0$, we have $D = d+1$. If $j = 1$, we have $D = O(d^2)$. Otherwise, for $2 \leq j \leq d-1$, we have $D = O(j^2d^3)$.

Suppose $\mathcal{P}$ is an instance of n stochastic points in $\mathbb{R}^d$. For each j-flat $F \in \mathbb{R}^d$, let $h_F(x) = d(x,F)^2$ ($x \in \mathbb{R}^d$), which admits a linearization of dimension $O(j^2d^3)$ by Lemma 89. Now, we map each point $s \in \mathcal{P}$ into an $O(j^2d^3)$ dimensional point s' and map each j-flat $F \in \mathbb{R}^d$ into an $O(j^2d^3)$ dimensional direction $\vec{u}$, such that $d(s,F) = \langle s', \vec{u}\rangle^{1/2}$. For convenience, we still use $\mathcal{P}$ to represent the collection of points after linearization. Recall that $\Pr[\vDash P]$ is the realized probability of the realization $P \sim \mathcal{P}$. By this mapping, we translate our goal into finding a direction $\vec{u} \in \mathbb{R}^{O(j^2d^3)}$, which minimizes the expected value $\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle \vec{u}, x\rangle^{1/2}] = \sum_{P\sim\mathcal{P}} \Pr[\vDash P] \cdot \max_{x\in P}\langle \vec{u}, x\rangle^{1/2}$. We also denote $\mathcal{P}^\circ = \{\vec{u} \in \mathbb{R}^d \mid \langle \vec{u}, s\rangle \geq 0, \forall s \in \mathcal{P}\}$ to be the polar set of $\mathcal{P}$. We only care about the directions in the polar set $\mathcal{P}^\circ$ for which $\langle \vec{u}, s\rangle^{1/2}$, $\forall s \in \mathcal{P}$ is well defined.

We first construct a convex hull $\mathcal{H}$ to partition the realizations into two parts. Our construction uses the method of (ε,τ)-quant-kernel construction in Chapter 4. For any normal vector (direction) $\vec{u}$, we move a sweep line $l_{\vec{u}}$ orthogonal to $\vec{u}$, along the direction $\vec{u}$, to sweep through the points in $\mathcal{P}$. Stop the movement of $\ell_{\vec{u}}$ at the first point such that $\Pr[\mathcal{P} \cap \overline{H}_{\vec{u}})] \geq \varepsilon'$, where $\varepsilon' = \varepsilon^{O(j^2d^3)}$ is a fixed constant. Denote $H_{\vec{u}}$ to be the halfplane defined by the sweep line $\ell_{\vec{u}}$ (orthogonal to the normal vector $\vec{u}$) and $\overline{H}_{\vec{u}}$ to be its complement. Denote $\mathcal{P}(\overline{H}_{\vec{u}}) = \mathcal{P} \cap \overline{H}_{\vec{u}}$ to be the set of points swept by the sweep line $l_{\vec{u}}$. We repeat the above process for all normal vectors (directions) $\vec{u}$, and let $\mathcal{H} = \cap_{\vec{u}} H_{\vec{u}}$. Since the total probability $B \geq \varepsilon$, $\mathcal{H}$ is nonempty by Helly's theorem. We also know that $\mathcal{H}$ is a convex hull by Chapter 4. Moreover, we have the following lemma.

Lemma 90. (Lemma 39 and Theorem 5) Suppose the dimensionality is d. There is a convex set $\mathcal{K}$, which is an intersection of $O(\varepsilon^{-(d-1)/2})$ halfspaces and satisfies

$(1-\varepsilon)\mathcal{K} \subseteq \mathcal{H} \subseteq \mathcal{K}$. Moreover, $\mathcal{K}$ can be constructed in $O(n \log^{O(d)} n)$ time.

By the above lemma, we construct a convex set $\mathcal{K} = \cap_{\vec{u}} \mathcal{K}_{\vec{u}}$, which is the intersection of $O(\varepsilon^{-O(j^2d^3)})$ halfspaces $\mathcal{K}_{\vec{u}}$ ($\vec{u}$ is the direction orthogonal to the halfspace $\mathcal{K}_{\vec{u}}$). Let $\overline{\mathcal{K}}_{\vec{u}}$ be the complement of $\mathcal{K}_{\vec{u}}$, and let $\mathcal{P}(\overline{\mathcal{K}}_{\vec{u}}) = \mathcal{P} \cap \overline{\mathcal{K}}_{\vec{u}}$ be the set of points in $\overline{\mathcal{K}}_{\vec{u}}$. Denote $\mathcal{P}(\overline{\mathcal{K}})$ to be the set of points outside the convex set $\mathcal{K}$. Then we have the following lemma, which shows that the total probability outside $\mathcal{K}$ is very small.

Lemma 91. Let $\mathcal{K}$ be a convex set constructed as in Lemma 90. The total probability $\Pr[\mathcal{P}(\overline{\mathcal{K}})] \leq \varepsilon$.

Proof. Assume that $\mathcal{K} = \cap_{\vec{u}} \mathcal{K}_{\vec{u}}$. Consider a halfspace $\mathcal{K}_{\vec{u}}$. By Lemma 90, the convex set $\mathcal{K}$ satisfies that $\mathcal{H} \subseteq \mathcal{K}$. Thus, we have that $\Pr[\mathcal{P}(\overline{\mathcal{K}}_{\vec{u}})] \leq \Pr[\mathcal{P}(\overline{H}_{\vec{u}})] \leq \varepsilon'$ by the definition of $\overline{H}_{\vec{u}}$.

Note that $\Pr[\mathcal{P}(\overline{\mathcal{K}})]$ is upper bounded by the multiplication of ε' and the number of halfspaces of $\mathcal{K}$. By Lemma 90, there are at most $O(\varepsilon^{-O(j^2d^3)})$ halfspaces $\mathcal{K}_{\vec{u}}$. Thus, we have that $\Pr[\mathcal{P}(\overline{\mathcal{K}})] \leq \varepsilon$. □

Our construction of SJFC-Coreset is consist of two parts. For points inside $\mathcal{K}$, we construct a collection $\mathcal{S}_1$. Our construction is almost the same as (ε, r)-fpow-kernel construction in Chapter 4, except that the cardinality of the collection $\mathcal{S}_1$ is different. For completeness, we provide the details of the construction here. Let $\mathcal{P}(\mathcal{K})$ be the collection of points in $\mathcal{K} \cap \mathcal{P}$, then $\mathcal{P}(\mathcal{K})$ is also an instance of a stochastic minimum j-flat-center problem. We show that we can estimate $\mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}]$ by $\mathcal{S}_1$. For the rest points outside $\mathcal{K}$, we show that the contribution for the objective function $\mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}]$ is almost linear and can be reduced to a weighted j-flat-median problem as in Case 1.

We first show how to construct $\mathcal{S}_1$ for points inside $\mathcal{K}$ as follows.

1. Sample $N = O((\varepsilon'\varepsilon)^{-2}\varepsilon^{-O(j^2d^3)} \log(1/\varepsilon)) = O(\varepsilon^{-O(j^2d^3)})$ independent realizations restricted to $\mathcal{P}(\mathcal{K})$.

2. For each realization S_i, use the algorithm in [7] to find a deterministic ε-kernel $\mathcal{E}_i$ of size $O(\varepsilon^{-O(j^2d^3)})$. Here, a deterministic ε-kernel $\mathcal{E}_i$ satisfies that $(1-$

$\varepsilon)CH(S_i) \subseteq CH(\mathcal{E}_i) \subseteq CH(S_i)$, where $CH(\cdot)$ is the convex hull of the point set.

3. Let $\mathcal{S}_1 = \{\mathcal{E}_i \mid 1 \leq i \leq N\}$ be the collection of all ε-kernels, and each ε-kernel $\mathcal{E}_i$ has a weight $1/N$.

Hence, the total size of $\mathcal{S}_1$ is $O(\varepsilon^{-O(j^2d^3)})$. For any direction $\vec{u} \in \mathcal{P}^{\square}$, we use $\frac{1}{N} \sum_{\mathcal{E}_i \in \mathcal{S}_1} \max_{x \in \mathcal{E}_i} \langle \vec{u}, x \rangle^{1/2}$ as an estimation of $\mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}]$. By Chapter 4, we have the following lemma.

Lemma 92. (Lemma 44-46) For any direction $\vec{u} \in \mathcal{P}^{\square}$, let $M_{\vec{u}} = \max_{x \in \mathcal{P}(\mathcal{K})} \langle \vec{u}, x \rangle^{1/2}$. We have that

$$\frac{1}{N} \sum_{\mathcal{E}_i \in \mathcal{S}_1} \max_{x \in \mathcal{E}_i} \langle \vec{u}, x \rangle^{1/2} \in (1 \pm \varepsilon/2) \mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}] \pm \varepsilon' \varepsilon (1 - \varepsilon) M_{\vec{u}}/4$$

Now we are ready to prove the following lemma.

Lemma 93. For any direction $\vec{u} \in \mathcal{P}^{\square}$, we have the following property.

$$\frac{1}{N} \sum_{\mathcal{E}_i \in \mathcal{S}_1} \max_{x \in \mathcal{E}_i} \langle \vec{u}, x \rangle^{1/2} + \sum_{s_i \in \mathcal{P}(\overline{\mathcal{K}})} p_i \cdot \langle \vec{u}, s_i \rangle^{1/2} \in (1 \pm 4\varepsilon) \mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}].$$

Proof. Let E be the event that no point is present in $\overline{\mathcal{K}}$. By the fact $\Pr[\overline{\mathcal{K}}] \leq \varepsilon$, we have that $\Pr[E] = \Pi_{s_i \in \mathcal{P}(\overline{\mathcal{K}})}(1 - p_i) \geq 1 - \sum_{s_i \in \mathcal{P}(\overline{\mathcal{K}})} p_i \geq 1 - \varepsilon$. Thus, we conclude that $1 - \varepsilon \leq \Pr[E] \leq 1$ We first rewrite $\mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}]$ as follows:

$$\mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}] = \Pr[E] \cdot \mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2} \mid E] + \Pr[\overline{E}] \cdot \mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2} \mid \overline{E}]$$

$$= \Pr[E] \cdot \mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}] + \Pr[\overline{E}] \cdot \mathbb{E}_{P \sim \mathcal{P}}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2} \mid \overline{E}]$$

For event E, we bound the term $\Pr[E] \cdot \mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}]$ via the collection $\mathcal{S}_1$. Let $M_{\vec{u}} = \max_{x \in \mathcal{P}(\mathcal{K})} \langle \vec{u}, x \rangle^{1/2}$. By Lemma 92, for any direction $\vec{u} \in \mathcal{P}^{\square}$, we have that

$$\frac{1}{N} \sum_{\mathcal{E}_i \in \mathcal{S}_1} \max_{x \in \mathcal{E}_i} \langle \vec{u}, x \rangle^{1/2} \in (1 \pm \varepsilon/2) \mathbb{E}_{P \sim \mathcal{P}(\mathcal{K})}[\max_{x \in P} \langle \vec{u}, x \rangle^{1/2}] \pm \varepsilon' \varepsilon (1 - \varepsilon) M_{\vec{u}}/4$$

By Lemma 90, we have that $(1-\varepsilon)\mathcal{K} \subseteq \mathcal{H}$. Then by the construction of $\mathcal{H}_{\bar{u}}$, we have that $\Pr[\mathcal{P} \cap (1-\varepsilon)\overline{\mathcal{K}}_{\bar{u}}] \geq \varepsilon'$. Thus, we obtain that

$$\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}] \geq \varepsilon'(1-\varepsilon)\max_{x\in\mathcal{P}(\mathcal{K})}\langle\vec{u},x\rangle^{1/2} = \varepsilon'(1-\varepsilon)M_{\bar{u}}.$$

So we conclude that

$$(1-2\varepsilon)\Pr[E]\cdot\mathbb{E}_{P\sim\mathcal{P}(\mathcal{K})}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}] - \varepsilon\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}] \leq \frac{1}{N}\sum_{\mathcal{E}_i\in\mathcal{S}_1}\max_{x\in\mathcal{E}_i}\langle\vec{u},x\rangle^{1/2}$$

$$\leq (1+2\varepsilon)\Pr[E]\cdot\mathbb{E}_{P\sim\mathcal{P}(\mathcal{K})}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}] + \varepsilon\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}], \tag{5.6}$$

since $1-\varepsilon \leq \Pr[E] \leq 1$.

For event $\overline{E}$, without loss of generality, we assume that the n points $s_1,\ldots,s_n$ in $\mathcal{P}$ are sorted in nondecreasing order according to the inner product $\langle\vec{u},s_i\rangle$. Assume that $s_{i_1},\ldots,s_{i_l}$ $(i_1 < i_2 < \ldots < i_l)$ are points in $\mathcal{P}(\overline{\mathcal{K}})$. Let E_j be the event that point s_{i_j} is present and all points s_{i_k} are not present for $k > j$. We have that

$$\Pr[\overline{E}]\cdot\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid \overline{E}] = \sum_{j\in[l]}\Pr[E_j]\cdot\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid E_j]$$

$$= \sum_{j\in[l]} p_{i_j}\cdot\prod_{j+1\leq k\leq l}(1-p_{i_k})\cdot\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid E_j].$$

By the above equality, on one hand, we have that

$$\Pr[\overline{E}]\cdot\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid \overline{E}] \geq (1-\varepsilon)\sum_{j\in[l]} p_{i_j}\cdot\langle\vec{u},s_{i_j}\rangle^{1/2}, \tag{5.7}$$

since $\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \geq \langle\vec{u},s_{i_j}\rangle^{1/2}$ if event E_j happens. On the other hand, the

following inequality also holds.

$$\begin{aligned}
&\Pr[\overline{E}] \cdot \mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid \overline{E}] = \sum_{j\in[l]} \Pr[E_j] \cdot \mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2} \mid E_j] \\
\leq &\sum_{j\in[l]} \Pr[E_j] \cdot \mathbb{E}_{P\sim\mathcal{P}}[\langle\vec{u},s_{i_j}\rangle^{1/2} + \max_{x\in P\cap\mathcal{P}(\mathcal{K})}\langle\vec{u},x\rangle^{1/2} \mid E_j] \\
\leq &\sum_{j\in[l]} p_{i_j} \cdot \left(\mathbb{E}_{P\sim\mathcal{P}}[\langle\vec{u},s_{i_j}\rangle^{1/2} \mid E_j] + \mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P\cap\mathcal{P}(\mathcal{K})}\langle\vec{u},x\rangle^{1/2} \mid E_j]\right) \\
\leq &\sum_{j\in[l]} p_{i_j} \cdot \langle\vec{u},s_{i_j}\rangle^{1/2} + \sum_{j\in[l]} p_{i_j} \cdot \mathbb{E}_{P\sim\mathcal{P}(\mathcal{K})}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}] \leq \sum_{j\in[l]} p_{i_j} \cdot \langle\vec{u},s_{i_j}\rangle^{1/2} + \varepsilon \cdot \mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}].
\end{aligned} \tag{5.8}$$

The last inequality holds since that $\sum_{j\in[l]} p_{i_j} = \Pr[\mathcal{P}(\overline{\mathcal{K}})] \leq \varepsilon$ by Lemma 91. Combining Inequalities (5.6), (5.7) and (5.8), we prove the lemma. □

By Lemma 88, we construct a point set $\mathcal{S}_2$ to estimate $\sum_{s_i\in\mathcal{P}(\overline{\mathcal{K}})} p_i \cdot d(s_i, F)$ with a weight function $w' : \mathcal{S}_2 \to \mathbb{R}$. We have that the size of $\mathcal{S}_2$ can be bounded by $O(j^4 d\varepsilon^{-2})$. Then $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$ is a collection of constant size, which satisfies the following property:

$$\frac{1}{N}\sum_{\mathcal{E}_i\in\mathcal{S}_1} \max_{x\in\mathcal{E}_i}\langle\vec{u},x\rangle^{1/2} + \sum_{s_i\in\mathcal{S}_2} w'_i \cdot \langle\vec{u},s_i\rangle^{1/2} \in (1+O(\varepsilon))\mathbb{E}_{P\sim\mathcal{P}}[\max_{x\in P}\langle\vec{u},x\rangle^{1/2}]. \tag{5.9}$$

Here w'_i is the weight of s_i in $\mathcal{S}_2$. We can think $\mathcal{S}_2 = \{\{s_i\} \mid 1 \leq s_i \leq |\mathcal{S}_2|\}$ as a collection of singleton point sets $\{s_i\}$. Then by Inequality 5.9, we have that $\mathcal{S}$ is a generalized ε-coreset satisfying Definition 62. We conclude the following lemma.

Lemma 94. Given an instance $\mathcal{P}$ of n stochastic points of the stochastic minimum j-flat-center problem in the existential model, if the total probability $\sum_i p_i \geq \varepsilon$, there exists an SJFC-Coreset $\mathcal{S}$ containing $O(\varepsilon^{-O(j^2d^3)} + j^4 d\varepsilon^{-2})$ point sets of size at most $O(\varepsilon^{-O(j^2d^3)})$, together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$, which satisfies that for any j-flat $F \in \mathcal{F}$,

$$\sum_{S\in\mathcal{S}} w'(S) \cdot J(S,F) \in (1\pm\varepsilon)J(\mathcal{P},F).$$

Combining Lemma 88 and Lemma 94, we can obtain the following theorem.

Theorem 95. Given an instance $\mathcal{P}$ of n stochastic points in the existential model, there is an SJFC-Coreset of size $O(\varepsilon^{-O(j^2d^3)} + j^4d\varepsilon^{-2})$ for the minimum j-flat-center problem. Moreover, we have an $O(n\log^{O(d)} n + \varepsilon^{-O(j^2d^3)}n)$ time algorithm to compute the SJFC-Coreset.

Proof. We only need to prove the running time. Recall that the SJFC-Coreset $\mathcal{S}$ can be divided into two parts $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$. For the first part $\mathcal{S}_1$, we construct the convex hull $\mathcal{K}$ in $O(n\log^{O(d)} n)$ by Lemma 90. Then we construct $\mathcal{S}_1$ by taking $O(\varepsilon^{-O(j^2d^3)})$ independent realizations restricted to $\mathcal{P}(\mathcal{K})$. For each sample, we construct a deterministic ε-kernel in $O(n + \varepsilon^{-(d-3/2)})$ by [110]. So the total time for constructing $\mathcal{S}_1$ is $O(n\log^{O(d)} n + \varepsilon^{-O(j^2d^3)}n)$. On the other hand, we can construct $\mathcal{S}_2$ in $O(ndj^{O(j^2)})$ time by Lemma 88. Thus, we prove the theorem. □

PTAS for stochastic minimum j-flat-center. Given an SJFC-Coreset $\mathcal{S}$ together with a weight function $w' : \mathcal{S} \to \mathbb{R}^+$ by Theorem 95, it remains to show how to compute the optimal j-flat for $\mathcal{S}$. Our goal is to find the optimal j-flat F^* such that the total generalized distance $\sum_{S\in\mathcal{S}} w'(S)\cdot J(S, F^*)$ is minimized. The argument is similar to the stochastic minimum k-center problem.

We first divide the family $\mathcal{F}$ of j-flats into a constant number of sub-families. In each sub-family $\mathcal{F}' \subseteq \mathcal{F}$, we have the following property: for each $S_i \in \mathcal{S}$, and each j-flat $F \in \mathcal{F}'$, the point $s^i = \arg\max_{s\in S_i} d(s, F)$ is fixed. By Lemma 41, we have that $h_F(x) = d(x, F)^2$ ($x \in \mathbb{R}^d$) admits a linearization of dimension $O(j^2d^3)$. For each sub-family $\mathcal{F}'$, we can formulate the optimization problem as a polynomial system of constant degree, a constant number of variables, and a constant number of constraints. Then we can compute the optimal j-flat in constant time for all sub-families $\mathcal{F}' \subseteq \mathcal{F}$. Thus, we can compute the optimal j-flat-center for the SJFC-Coreset $\mathcal{S}$ in constant time. We then have the following corollary.

Corollary 96. If the dimensionality d is a constant, given an instance of n stochastic points in $\mathbb{R}^d$ in the existential uncertainty model, there exists a PTAS for the stochastic minimum j-flat-center problem in $O(n\log^{O(d)} n + \varepsilon^{-O(j^2d^3)}n)$ time.

Locational Uncertainty Model. Note that in the locational uncertainty model, we only need to consider Case 2. We use the same construction as in the existential model. Let $p_i = \prod_j p_{v_j s_i}$. Similarly, we make a linearization for the function $d(x, F)^2$, where $x \in \mathbb{R}^d$ and $F \in \mathcal{F}$ is a j-flat. Using this linearization, we also map $\mathcal{P}$ into $O(j^2d^3)$-dimensional points. For the jth node and a set P of points, we denote $p_j(P) = \prod_{s_i \in P} p_{v_j s_i}$ to be the total probability that the jth node locates inside P.

By the condition $\Pr[\mathcal{P}(\overline{\mathcal{K}})] \leq \varepsilon$, we have that $\Pr[E] = 1 - \prod_{j \in [m]}(1 - p_j(\overline{\mathcal{K}})) \leq 1 - (1 - \varepsilon) = \varepsilon$, where event E represents that there exists a point present in $\overline{\mathcal{K}}$. So we can regard those points outside $\mathcal{K}$ independent. On the other hand, for any direction $\vec{u}$, since $\Pr[\mathcal{P} \cap (1 - \varepsilon)\overline{\mathcal{H}}_{\vec{u}}] \geq \varepsilon'$, we have that $\Pr[E_{\vec{u}}] = 1 - \prod_{j \in [m]}(1 - p_j(\mathcal{P} \cap (1 - \varepsilon)\overline{H}_{\vec{u}})) \geq 1 - (1 - \frac{\varepsilon'}{m})^m \geq \varepsilon'/2$, where event $E_{\vec{u}}$ represents that there exists a point present in $\mathcal{P} \cap (1 - \varepsilon)\overline{H}_{\vec{u}}$. Moreover, we can use the same method to construct a collection $\mathcal{S}_1$ as an estimation for the point set $\mathcal{P}(\mathcal{K})$ in the locational uncertainty model. So Lemma 93 still holds. Then by Lemma 94, we can construct an SJFC-Coreset of constant size.

Theorem 97. Given an instance $\mathcal{P}$ of n stochastic points in the locational uncertainty model, there is an SJFC-Coreset of cardinality $O(\varepsilon^{-O(j^2d^3)} + j^4 d\varepsilon^{-2})$ for the minimum j-flat-center problem. Moreover, we have a polynomial time algorithm to compute the gerneralized ε-coreset.

By a similar argument as in the existential model, we can give a PTAS for the locational uncertainty model. Then combining with Corollary 96, we prove the main result Theorem 13.

5.5 Constructing additive ε-coresets

In this section, we first give the algorithm for constructing an additive ε-coreset. We construct Cartesian grids and maintain one point from each nonempty grid cell, which is similar to [11]. However, our algorithm is more complicated. See Algorithm 2 for details.

Now we analyze the algorithm.

We estimate the sum for all $s \in B$. Notice that the sum is in fact $\Pr[\mathrm{Cl}(v)] \cdot \mathbb{E}[\mathrm{MST} \mid \mathrm{Cl}(v)]$. We can see that $\Pr[\mathrm{Cl}(v)]$ can be computed exactly in linear time. We argue that the quality of the estimation taken on $N_1 = O\left(\frac{mn^2}{\varepsilon^5} \ln m\right)$ samples is sufficient by considering the following two cases:

(a) Assume that $\mathbb{E}[\mathrm{MST} \mid \mathrm{Cl}(v)] \geq \frac{1}{2}\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle] \geq \Omega\left(\frac{\varepsilon^2}{n^2}\right) \mathrm{diam}(\mathbb{H})$. In this case, we have a poly-bounded random variable. This is because under the condition $\mathrm{Cl}(v)$, the maximum possible length of any minimum spanning tree is $O(\frac{m}{\varepsilon}\mathrm{diam}(\mathbb{H}))$. Hence we can use Monte Carlo to get a $(1 \pm \varepsilon)$-approximation of $\mathbb{E}[\mathrm{MST} \mid \mathrm{Cl}(v)]$ with $O\left(\frac{mn^2}{\varepsilon^5} \ln m\right)$ samples.

(b) Otherwise, we assume that $\mathbb{E}[\mathrm{MST} \mid \mathrm{Cl}(v)] \leq \frac{1}{2}\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle]]$. Let V_0 be the collection of these nodes. The probability that the sample average is larger than $\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle]]$ is at most $\mathrm{poly}(\frac{1}{m})$ by Chernoff Bound. The probability that for all nodes $v \in V_0$, the sample average are at most $\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle]]$ is at least $1 - \mathrm{poly}(\frac{1}{m})$ by union bound. If this is the case, we can see their total contribution to the final estimation of $\mathbb{E}[\mathrm{MST}]$ is less than $\varepsilon\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle]]\Pr[\mathbb{H}\langle m\rangle]$. In fact, this is because

$$\sum_{v \in V_0} \Pr[\mathrm{Cl}(v)] \cdot T_v \leq \sum_{v \in V_0} \Pr[\mathrm{Cl}(v)] \cdot \mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle] < \varepsilon\mathbb{E}[\mathrm{MST} \mid \mathbb{H}\langle m\rangle]\Pr[\mathbb{H}\langle m\rangle].$$

The second inequality is due to the fact that $\sum_{v \in V_0} \Pr[\mathrm{Cl}(v)] \leq m - p(\mathbb{H}) < \varepsilon/16 < \varepsilon\Pr[\mathbb{H}\langle m\rangle]$.

2. Point $s \in \mathbb{F} \setminus B$, each term has $\mathrm{d}(s, \mathbb{H}) > \frac{m}{\varepsilon} \cdot \mathrm{diam}(\mathbb{H})$.

We just use $\mathrm{d}(s, \mathbb{H})$ as the estimation of $\mathbb{E}[\mathrm{MST} \mid \mathbb{F}\langle v\rangle, v \vDash s]$. This is because the length of MST is always at least $\mathrm{d}(s, \mathbb{H})$ and at most $\mathrm{d}(s, \mathbb{H}) + m \cdot \mathrm{diam}(\mathbb{H}) \leq (1 + \varepsilon)\mathrm{d}(s, \mathbb{H})$.

□

6.7.3 Minimum Perfect Matching

Lemma 119. Algorithm 6.4 produces a $(1 \pm \varepsilon)$-estimate for the second term with high probability.

Proof. To compute the second term, we first rewrite it as follows:

$$\mathbb{E}[\mathrm{PM} \mid \mathbb{F}\langle 1\rangle] \cdot \Pr[\mathbb{H}\langle 1\rangle] = \sum_{v \in \mathcal{V}} \sum_{s \notin \mathbb{H}(v)} \Pr[\mathbb{F}\langle v\rangle \wedge v \vDash s]\, \mathbb{E}[\mathrm{PM} \mid \mathbb{F}\langle v\rangle, v \vDash s].$$

Fix a particular node v. To estimate $\sum_{s \in \mathbb{F}} \Pr[\mathbb{F}\langle v\rangle \wedge v \vDash s]\, \mathbb{E}[\mathrm{PM} \mid \mathbb{F}\langle v\rangle, v \vDash s]$, we consider the following two situations:

1. Point $s \in B_v$, i,e, $d(s, \mathbb{H}(v)) < \frac{4mD}{\varepsilon}$.

   We estimate the sum for all $s \in B^v$. Notice that the sum is in fact $\Pr[\mathrm{Cl}(v)] \cdot \mathbb{E}[\mathrm{PM} \mid \mathrm{Cl}(v)]$. We can see that $\Pr[\mathrm{Cl}(v)]$ can be computed exactly in linear time. We argue that the quality of the estimation taken on $N_2 = O\left(\frac{m^2 n^5}{\varepsilon^4} \ln m\right)$ samples is poly-bounded by considering the following two cases:

   (a) Assume that $\mathbb{E}[\mathrm{PM} \mid \mathrm{Cl}(v)] \geq \frac{1}{2}\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle] = \Omega\left(\frac{\varepsilon D}{mn^5}\right)$. In this case, our estimation is poly-bounded. This is because under the condition $\mathrm{Cl}(v)$, the maximum possible length of any minimum perfect matching is $O(\frac{mD}{\varepsilon})$. Hence we can use Monte Carlo to get a $(1 \pm \varepsilon)$-approximation of $\mathbb{E}[\mathrm{PM} \mid \mathrm{Cl}(v)]$ with $O\left(\frac{m^2 n^5}{\varepsilon^4} \ln m\right)$ samples.

   (b) Otherwise, we assume that $\mathbb{E}[\mathrm{PM} \mid \mathrm{Cl}(v)] \leq \frac{1}{2}\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]]$. Let V_0 be the collection of these nodes. The probability that the sample average is larger than $\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]]$ is at most $\mathrm{poly}(\frac{1}{m})$ by Cherno Bound. The probability that for each node $v \in V_0$, the sample average is at most $\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]]$ is at least $1 - \mathrm{poly}(\frac{1}{m})$ by union bound. If this is the case, we can see their total contribution to the final estimation of $\mathbb{E}[\mathrm{PM}]$ is less than $\varepsilon\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]]\Pr[\mathbb{H}\langle m\rangle]$. In fact, this is because

   $$\sum_{v \in V_0} \Pr[\mathrm{Cl}(v)] \cdot T_v \leq \sum_{v \in V_0} \Pr[\mathrm{Cl}(v)] \cdot \mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]] < \varepsilon\mathbb{E}[\mathrm{PM} \mid \mathbb{H}\langle m\rangle]]\Pr[\mathbb{H}\langle m\rangle].$$

The second inequality is due to the fact that $\sum_{v \in V} \Pr[\mathrm{Cl}(v)] \leq m - \sum_{v \in V_0} p_v(\mathbb{H}(v)) \leq \frac{\varepsilon}{n^3} < \varepsilon \Pr[\mathbb{H}\langle m \rangle]$.

2. Point $s \in \mathcal{P} \setminus (B_v \cup \mathbb{H}(v))$, each term has $d(s, \mathbb{H}(v)) > \frac{4mD}{\varepsilon}$. The algorithm uses $d(s, \mathbb{H}(v))$ as the estimation of $\mathbb{E}[\mathrm{PM} \mid \mathbb{F}\langle v \rangle, v \vDash s]$. Note that the length of PM is always at least $d(s, \mathbb{H}(v)) - mD \geq (1 - \frac{\varepsilon}{4}) d(s, \mathbb{H}(v))$. This is because such an instance PM contains a path from s to some point $t \in \mathbb{H}(v)$ deleting no more than m segments of length at most D (each segment is in some $\mathbb{H}_i$). On the other hand, the length of PM is at most $d(s, \mathbb{H}(v)) + mD \leq (1 + \frac{\varepsilon}{4}) d(s, \mathbb{H}(v))$. So it is a $(1 \pm \varepsilon)$-estimation.

□

6.8 The Closest Pair Problem

6.8.1 Estimating kth Closest Pair in the Existential Uncertainty Model

Again, we construct the HPF $\Psi(\mathcal{P})$. Let the random variable Y be the largest integer i such that there are at least k point collisions in Γ_i. Here we use a point collision to denote that a pair of points are present in the same component. Note that if there are exactly i points in a component, the amount of point collisions in this component is $\binom{i}{2}$. We denote as $\Gamma\langle \# j \rangle$ the event that there are exactly j point collisions among the partition Γ on $\mathcal{P}$. Similarly, we can rewrite $\mathbb{E}[\mathrm{kC}]$ by $\mathbb{E}[\mathrm{kC}] = \sum_{i=1}^{m-1} \Pr[Y = i] \mathbb{E}[\mathrm{kC} \mid Y = i]$.

We use dynamic programming technique to achieve an FPRAS for computing $\mathbb{E}[\mathrm{kC}]$. Note that conditioning on $Y = i$, the value of kC is between $d(e_i)$ and $m \cdot d(e_i)$. So we only need to show the following lemma.

Lemma 129. We can compute $\Pr[Y = i]$ in polynomial time. Moreover, there exists a polynomial time sampler conditioning on $Y = i$.

Proof. We denote $E[a, b]$ $(1 \leq a \leq i - 1,\ b \leq k)$ the event that among the first a components in Γ'_i, there are exactly $b \leq k$ point collisions. We denote the probability

of $E[a, b]$ by $\Pr[a, b]$. We give the dynamic programming as follows.

1. If $\sum_{1 \le j \le a} \binom{|C_j|}{2} < b$, $\Pr[a, b] = 0$. If $b = 0$, $\Pr[a, b] = \prod_{1 \le j \le a} \Pr[C_j \langle \le 1 \rangle]$. If $b < 0$, $\Pr[a, b] = 0$.
2. If $\sum_{1 \le j \le a} \binom{|C_j|}{2} \ge b$, $1 \le b \le k$, $\Pr[a, b] = \sum_{0 \le l \le n_a} \Pr[C_a \langle l \rangle] \cdot \Pr[a - 1, b - \binom{l}{2}]$.

By the above dynamic programming, we can compute $\Pr[i - 1, l]$ for $0 \le l \le k - 1$ in polynomial time.

By the definition of $Y = i$, it is no hard to see that we can rewrite $\Pr[Y = i]$ as follows:

$$\Pr[Y = i] = \sum_{1 \le n_1 \le |\mu'_{i+1}|, 1 \le n_2 \le |\mu''_{i+1}|} \Pr[\mu'_{i+1}\langle n_1 \rangle] \cdot \Pr[\mu''_{i+1}\langle n_2 \rangle] \cdot \left[\sum_{k - \binom{n_1 + n_2}{2} \le l \le k - 1 - \binom{n_1}{2} - \binom{n_2}{2}} \Pr[\Gamma'_i \langle \# l \rangle] \right]$$

Note that we can compute $\Pr[Y = i]$ in polynomial time. We need to describe our sampler conditioning on $Y = i$. We first sample the event $\mu'_{i+1}\langle n_1 \rangle \wedge \mu''_{i+1}\langle n_2 \rangle$ with probability $\Pr[\mu'_{i+1}\langle n_1 \rangle \wedge \mu''_{i+1}\langle n_2 \rangle \mid Y = i]$. Then conditioning on $k - \binom{n_1 + n_2}{2} \le l \le k - 1 - \binom{n_1}{2} - \binom{n_2}{2}$, we sample the total number of point collisions in Γ'_i. Then we sample the number of present points in each component in Γ'_i using the dynamic programming. Finally, based on the number of present points in each component, we sample the present points by Lemma 108.

Using the Monte Carlo method, we only need to take $O(\frac{n}{\varepsilon^2} \ln n)$ independent samples for estimating $\mathbb{E}[k_C \mid Y = i]$. Thus, we totally take $O(\frac{n^2}{\varepsilon^2} \ln n)$ independent samples.

□

Theorem 130. There is an FPRAS for estimating the expected distance between the kth closest pair in the existential uncertainty model.

6.8.2 Hardness for Closest Pair

Theorem 131. Computing $\Pr[C \ge 1]$ is #P-hard to approximate within any factor in a metric space in both the existential and locational uncertainty models.

Proof. First consider the existential uncertainty model. Consider a metric graph G with edge weights being either 0.9 or 1.8. Each vertex in this graph exists with probability 1/2. Let G' be the unweighted graph with the same number of vertices. G' contains only those edges corresponding to edges with weight 0.9 in G. It is not hard to see that

$$\Pr[C \geq 1] = \#\text{independent sets of size at least two in } G' \cdot \frac{1}{2^n}.$$

The right hand side is well known to be inapproximable for arbitrary graphs [97].

For the locational model, let the instance be G (with n vertices $s_1, \ldots, s_n$) with n additional vertices $t_1, \ldots, t_n$ which are far away from each other and any vertex in G. Let the probability distribution of node v_i be $p_{v_i s_i} = 1/2$, and $p_{v_i t_i} = 1/2$. We can see that in this locational uncertainty model, the value $\Pr[C \geq 1]$ is the same as that in the corresponding existential model G.

□

Theorem 132. Computing $\mathbb{E}[C]$ exactly in both the existential and locational uncertainty models is #P-hard in a metric space.

Proof. Consider a metric graph G with edge weights being either 1 or 2. Each vertex in this graph exists with probability 1/2. Note that

$$\mathbb{E}[C] = \Pr[C = 1] + 2\Pr[C = 2] = (\Pr[C \leq 1] - \Pr[C = 0]) + 2(1 - \Pr[C \leq 1])$$

Computing $\Pr[C = 0]$ can be easily done in polynomial time. Computing $\Pr[C \leq 1]$ in such a graph is as hard as counting independent sets in general graphs, hence is also #P-hard (as in Theorem 131). So, computing $\mathbb{E}[C]$ is #P-hard as well.

For the locational model, let the instance be G (with n vertices $s_1, \ldots, s_n$) with n additional vertices $t_1, \ldots, t_n$ which satisfies $d(s_i, t_j) = d(t_i, t_j) = 5$ ($1 \leq i, j \leq m, i \neq j$). Let the probability distribution of node v_i be $p_{v_i s_i} = 1/2$, and $p_{v_i t_i} = 1/2$. It is not hard to see that in this locational uncertainty model, the value $\mathbb{E}[C]$ is linearly related to the value $\mathbb{E}[C]$ in the existential model G. Therefore, computing $\mathbb{E}[C]$ is

also #P-hard in the locational uncertainty model.

□

6.9 Another FPRAS for MST

W.l.o.g., we assume that for each point, there is only one node that may be realized to this point. Our algorithm is a slight generalization of the one proposed in [71]. Let $\mathbb{E}[i]$ be the expected MST length conditioned on the event that all nodes $\{v_1, \ldots, v_m\}$ are realized to points in $\{s_i, \ldots, s_n\}$ (denote the event by $\ln(i, n)$). Let $\mathbb{E}'[i]$ be the expected MST length conditioned on the event that all nodes $\{v_1, \ldots, v_n\}$ are realized to $\{s_i, \ldots, s_n\}$ and at least one node is realized to s_i. We use $s \vDash s$ to denote the event that node v is realized to point s. Note that

$$\mathbb{E}[i] = \mathbb{E}'[i]\Pr[\exists v, v \vDash s_i \mid \ln(i, m)] + \mathbb{E}[i+1]\Pr[\nexists v, v \vDash s_i \mid \ln(i, m)]$$

For a particular point s_i, we reorder the points $\{s_i, \ldots, s_n\}$ as $\{s_i = r_i, \ldots, r_n\}$ in increasing order of distance from s_i. Let $\mathbb{E}'[i, j]$ be the expected MST length for all nodes conditioned on the event that all nodes are realized to $\{r_i, \ldots, r_j\}$ (denoted as $\ln'(i, j)$) and $\exists v, v \vDash s_i$. Let $\mathbb{E}''[i, j]$ be the expected MST length for all nodes conditioned on the event $\ln'(i, j) \wedge (\exists v, v \vDash r_i) \wedge (\exists s', s' \vDash r_j)$. We can see that

$$\begin{aligned}\mathbb{E}'[i, j] =& \mathbb{E}''[i, j]\Pr[\exists v', v' \vDash r_j \mid \ln'(i, j), \exists v, v \vDash r_i] \\ &+ \mathbb{E}'[i, j-1]\Pr[\nexists v, v \vDash r_i \mid \ln'(i, j), \exists v, v \vDash r_i]\end{aligned}$$

It is not difficult to see the probability $\Pr[\exists v', v' \vDash r_j \mid \ln'(i, j), \exists v, v \vDash r_i]$ can be computed in polynomial time. Here we use the assumption that for each point, only one node that may realize to it. Moreover, we can also take samples conditioning on event $\ln'(i, j) \wedge (\exists v, v \vDash r_i) \wedge (\exists v', v' \vDash r_j)$. Therefore $\mathbb{E}''[i, j]$ can be approximated within a factor of $(1 \pm \varepsilon)$ using the Monte Carlo method in polynomial time since it is poly-bounded. The number of samples needed can be bounded by $O\left(\frac{mn^2}{\varepsilon^2} \ln m\right)$.

We can easily generalize the above algorithm to the case where $\sum_{j=1}^{n} p_{ij} \leq 1$,

i.e., node i may not be present with some certainty. Indeed, this can be done by generalizing the definition of $\mathsf{In}(i,j)$ (and similarly $\mathsf{In}'(i,j)$) to be the event that each node is either absent or realized to some point in $\{r_{\mathrm{i}},\ldots,r_{\mathrm{j}}\}$.

Chapter 7 Concluding Remarks

In this dissertation, we study two famous stochastic geometry models, the locational uncertainty model and the existential uncertainty model.

In the first part of the dissertation, we study how to construct coresets for different stochastic problems. We initiate the study of constructing ε-kernel coresets in stochastic geometry models. 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 provide e□cient algorithms for constructing such ε-kernel coresets in nearly linear time. Our ε-kernel coresets have a few applications, including approximating the extent of uncertain functions, maintaining extent measures for stochastic moving points and giving PTAS for some stochastic shape fitting problems.

We also study another two important stochastic geometric optimization problems, the k-center problem and the j-flat-center problem in stochastic geometry models in Euclidean spaces. We first think each of the stochastic problems as a certain deterministic problem over (exponential many) all possible realizations (each being a point set). By this view, we introduce a new notion called generalized coreset, which is a collection of realizations (instead of points for coresets). We also propose a new framework for generalized coreset construction. By the framework, we provide the first PTAS (Polynomial Time Approximation Scheme) for both stochastic geometry optimization problems, which generalize the previous results for stochastic minimum enclosing ball [88] and stochastic enclosing cylinder [66].

The second part of the dissertation is to estimate the expected value of a variety combinatorial objects over stochastic data. Several geometric properties of a set of stochastic points have been studied extensively in the literature under the term stochastic geometry. For instance, it is well known that if there are n points uniformly and

independently distributed in $[0,1]^2$, the minimal traveling salesman tour/minimum spanning tree/minimum matching visiting them has an expected length $\Theta(\sqrt{n})$ [21, 28]. Compared with results in stochastic geometry, we focus on the efficient computation of the statistics, instead of giving explicit mathematical formulas.

We study the problems of computing the expected lengths of several combinatorial or geometric optimization problems in both models, including closest pair, minimum spanning tree, k-clustering, minimum perfect matching, and minimum cycle cover. We also consider the problem of estimating the probability that the length of closest pair, or the diameter, is at most, or at least, a given threshold. Most of the above problems are known to be #P-hard. Caused by the high variance, we can not directly use Monte Carlo method to estimate the expected value. Thus, we develop two new techniques: stoch-core and Hierarchical Partition Family (HPF). Both techniques are used to decompose the expectation of certain random variable into a convex combination of conditional expectations, such that each conditional expectation has a low variance. Combining our new techniques and Monte Carlo method, we obtain FPRAS (Fully Polynomial Randomized Approximation Scheme) for most of these problems in stochastic geometry models.

There are still many open optimization problems over different stochastic models. In this dissertation, we study the discrete stochastic models. In practice, the distribution of point locations often follows some continuous distribution in practice, such as GPS system, robot control, and so on. Many fundamental issues in this domain, such as many classic geometry computation and optimization problems are still not well understood by researchers. We believe it is a fruitful direction for further research.

Bibliography

[1] A. Abdullah, S. Daruki, and J.M. Phillips. Range counting coresets for uncertain data. In Proceedings 29th ACM Syposium on Computational Geometry, pages 223–232, 2013.

[2] 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.

[3] P.K. Agarwal, S.W. Cheng, Y. Tao, and K. Yi. Indexing uncertain data. In Proceedings of the twenty-eighth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pages 137–146. ACM, 2009.

[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. Varadarajan. Approximating extent measures of points. Journal of the ACM, 51(4):606–635, 2004.

[8] P.K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximation via coresets. Combinatorial and Computational Geometry, 52:1–30, 2005.

[9] P.K. Agarwal, S. Har-Peled, and H. Yu. Robust shape fitting via peeling and grating coresets. Discrete & Computational Geometry, 39(1-3):38–58, 2008.

[10] P.K. Agarwal, J. Matoušek, and S. Suri. Farthest neighbors, maximum spanning trees and related problems in higher dimensions. Computational Geometry - Theory and Applications, 1(4):189–201, 1992.

[11] P.K. Agarwal and C.M. Procopiuc. Exact and approximation algorithms for clustering. Algorithmica, 33(2):201–226, 2002.

[12] 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.

[13] C. Alexopoulos and J.A. Jacobson. State space partition algorithms for stochastic systems with applications to minimum spanning trees. Networks, 35(2):118–138, 2000.

[14] M. Anthony and P.L. Bartlett. Neural network learning: Theoretical foundations. cambridge university press, 2009.

[15] M.J. Atallah, Y. Qi, and H. Yuan. Asymptotically e□cient algorithms for skyline probabilities of uncertain data. ACM Trans. Datab. Syst, 32(2):12, 2011.

[16] 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.

[17] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When lp is the cure for your matching woes: Improved bounds for stochastic matchings. In European Symposium on Algorithms, pages 218–229. Springer, 2010.

[18] J.F. Bard and J.E. Bennett. Arc reduction and path preference in stochastic acyclic networks. Management Science, 37(2):198–215, 1991.

[19] G. Barequet and S. Har-Peled. E□ciently approximating the minimum-volume bounding box of a point set in three dimensions. Journal of Algorithms, 38(1):91–109, 2001.

[20] S. Basu, R. Pollack, and M. Roy. Algorithms in real algebraic geometry. AMC, 10:12, 2011.

[21] J. Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. In Proc. Cambridge Philos. Soc, pages 55:299–327, 1959.

[22] M. W. Bern and D. Eppstein. Worst-case bounds for suadditive geometric graphs. In Symposium on Computational Geometry, pages 183–188, 1993.

[23] D.J. Bertsimas and G. van Ryzin. An asymptotic determination of the minimum spanning tree and minimum matching constants in geometrical probability. Operations Research Letters, 9(4):223–231, 1990.

[24] A. Bhalgat. A (2+ε)-approximation algorithm for the stochastic knapsack problem. Unpublished manuscript, 2011.

[25] T.M. Chan. Approximating the diameter, width, smallest enclosing cylinder, and minimum-width annulus. In Proceedings of the 16th Annual Son Computational Geometry, pages 300–309, 2000.

[26] T.M. Chan. Faster core-set constructions and data-stream algorithms in fixed dimensions. Computational Geometry: Theory and Applications, 35:20–35, 2006.

[27] 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.

[28] R. Cheng, J. Chen, M. Mokbel, and C. Chow. Probabilistic verifiers: Evaluating constrained nearest-neighbor queries over uncertain data. In ICDE, 2008.

[29] R. Cheng, J. Chen, and X. Xie. Cleaning uncertain data with quality guarantees. Proceedings of the VLDB Endowment, 1(1):722–735, 2008.

[30] 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.

[31] G. Cormode and S. Muthukrishnan. Radial histograms for spatial streams. Technical Report 2003-11, Center for Discrete Mathematics and Computer Science (DIMACS), 2003.

[32] G. Cornuejols and W. Pulleyblank. A matching problem with side constraints. Discrete Math., 29, 1980.

[33] B.C. Dean, M.X. Goemans, and J. Vondrdk. Approximating the stochastic knapsack problem: The benefit of adaptivity. In Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on, pages 208–217. IEEE, 2004.

[34] 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.

[35] 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.

[36] A. Driemel, H. HAverkort, M. Lö□er, and R.I. Silveira. Flow computations on imprecise terrains. Journal of Computational Geometry, 4:38–78, 2013.

[37] R.M. Dudley. Metric entropy of some classes of sets with di□erentiable boundaries. Journal of Approximation Theory, 10(3):227–236, 1974.

[38] M. Dyer. Approximate counting by dynamic programming. In ACM Symposium on Theory of Computing, pages 693–699, 2003.

[39] H. Edelsbrunner, J. O'Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. SIAM Journal on Computing, 15(2):341–363, 1986.

[40] Y. Emek, A. Korman, and Y. Shavitt. Approximating the statistics of various properties in randomly weighted graphs. In Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1455–1467. SIAM, 2011.

[41] W. Evans and J. Sember. The possible hull of imprecise points. In Proceedings of the 23rd Canadian Conference on Computational Geometry, 2011.

[42] T. Feder and D. Greene. Optimal algorithms for approximate clustering. In Proceedings of the twentieth annual ACM symposium on Theory of computing, pages 434–444. ACM, 1988.

[43] 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.

[44] D. Feldman, A. Fiat, and M. Sharir. Coresets forweighted facilities and their applications. In 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), pages 315–324. IEEE, 2006.

[45] 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.

[46] D. Feldman, M. Schmidt, and C. Sohler. Turning big data into tiny data: Constant-size coresets for k-means, pca and projective clustering. In Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1434–1453. SIAM, 2013.

[47] D. Feldman and L.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.

[48] A.M. Frieze. On the value of a random minimum spanning tree problem. Discrete Applied Mathematics, 10(1):47–56, 1985.

[49] 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.

[50] A. Goel and P. Indyk. Stochastic load balancing and related problems. In Foundations of Computer Science, 1999. 40th Annual Symposium on, pages 579–586. IEEE, 1999.

[51] T.F. Gonzalez. Clustering to minimize the maximum intercluster distance. Theoretical Computer Science, 38:293–306, 1985.

[52] V. Goyal and R. Ravi. A ptas for the chance-constrained knapsack problem with random item sizes. Operations Research Letters, 38(3):161–164, 2010.

[53] 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.

[54] L.J. Guibas, D. Salesin, and J. Stolfi. Constructing strongly convex approximate hulls with inaccurate primitives. Algorithmica, 9:534–560, 1993.

[55] P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity. In Proceedings of the 37th IEEE Conference on Decision and Control, volume 1, pages 1106–1110. IEEE, 1998.

[56] P. Gupta and P.R. Kumar. The capacity of wireless networks. IEEE Transactions on Information Theory, 46(2):388–404, 2000.

[57] M. Haenggi, J.G. Andrews, F. Baccelli, O. Dousse, and M. Franceschetti. Stochastic geometry and random graphs for the analysis and design of wireless networks. IEEE Journal on Selected Areas in Communications, 27(7):1029–1046, 2009.

[58] S. Har-Peled. Geometric approximation algorithms, volume 173. American mathematical society Providence, 2011.

[59] S. Har-Peled. On the expected complexity of random convex hulls. arXiv:1111.5340, 2011.

[60] 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.

[61] S. Har-Peled and K. Varadarajan. Projective clustering in high dimensions using core-sets. In Proceedings of the eighteenth annual symposium on Computational geometry, pages 312–318. ACM, 2002.

[62] S. Har-Peled and Y. Wang. Shape fitting with outliers. SIAM Journal on Computing, 33(2):269–285, 2004.

[63] D. Hartvigsen. An extension of matching theory. phd thesis, carnegie-mellon university. 1984.

[64] M. Held and J.S.B. Mitchell. Triangulating input-constrained planar point sets. Information Processing Letters, 109(1):54–56, 2008.

[65] L. Huang and J. Li. Approximating the expected values for combinatorial optimization problems over stochastic points. In Automata, Languages, and Programming, pages 910–921. Springer, 2015.

[66] Lingxiao Huang, Jian Li, Je M Phillips, and Haitao Wang. ε-kernel coresets for stochastic points. In European Symposium on Algorithms. Springer, 2016.

[67] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. Journal of the ACM (JACM), 51(4):671–697, 2004.

[68] M. Jerrum, L.G. Valiant, and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. Theoretical Computer Science, 43:169–188, 1986.

[69] A.G. Jørgensen, M. Lö□er, and J.M. Phillips. Geometric computation on indecisive points. In Proceedings of the 12th Algorithms and Data Structure Symposium, pages 536–547, 2011.

[70] 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.

[71] P. Kamousi, T.M. Chan, and S. Suri. Stochastic minimum spanning trees in Euclidean spaces. In Proceedings of the 27th annual ACM symposium on Computational Geometry, pages 65–74. ACM, 2011.

[72] P. Kamousi, T.M. Chan, and S. Suri. Closest pair and the post o□ce problem for stochastic points. Computational Geometry, 47(2):214–223, 2014.

[73] H.J. Karlo□. How long can a Euclidean traveling salesman tour be? In J. Discrete Math, page 2(1). SIAM, 1989.

[74] J. Kleinberg and T. Eva. Algorithm design. Pearson Education India, 2006.

[75] J. Kleinberg, Y. Rabani, and É. Tardos. Allocating bandwidth for bursty connections. SIAM Journal on Computing, 30(1):191–217, 2000.

[76] H. Kruger. Basic measures for imprecise point sets in $\mathbb{R}^d$. Master's thesis, Utrecht University, 2008.

[77] M. Langberg and L.J. Schulman. Universal ε-approximators for integrals. In Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms, 2010.

[78] J. Li and A. Deshpande. Maximizing expected utility for stochastic combinatorial optimization problems. In Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on, pages 797–806. IEEE, 2011.

[79] J. Li and Y. Liu. Approximation algorithms for stochastic combinatorial optimization problems. Journal of the Operations Research Society of China, 4(1):1–47, 2016.

[80] J. Li and H. Wang. Range queries on uncertain data. Theoretical Computer Science, 609:32–48, 2016.

[81] J. Li and W. Yuan. Stochastic combinatorial optimization via poisson approximation. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 971–980. ACM, 2013.

[82] Y. Li, P.M. Long, and A. Srinivasan. Improved bounds on the samples complexity of learning. Journal of Computer and System Sciences, 62:516–527, 2001.

[83] M. Lö□er and J. Phillips. Shape fitting on point sets with probability distributions. In Proceedings of the 17th European Symposium on Algorithms, pages 313–324, 2009.

[84] M. Lö□er and J. Snoeyink. Delaunay triangulations of imprecise points in linear time after preprocessing. In Proceedings of the 24th Sympoium on Computational Geometry, pages 298–304, 2008.

[85] M. Lö□er and M. van Kreveld. Approximating largest convex hulls for imprecise points. Journal of Discrete Algorithms, 6:583–594, 2008.

[86] R.P. Loui. Optimal paths in graphs with stochastic or multidimensional weights. Communications of the ACM, 26(9):670–676, 1983.

[87] J. Matoušek. Computing the center of planar point sets. Discrete and Computational Geometry, 6:221, 1991.

[88] A. Munteanu, C. Sohler, and D. Feldman. Smallest enclosing ball for probabilistic data. In Proceedings of the thirtieth annual symposium on Computational geometry, page 214. ACM, 2014.

[89] 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.

[90] E. Nikolova. Approximation algorithms for reliable stochastic combinatorial optimization. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pages 338–351. Springer, 2010.

[91] E. Nikolova, J.A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In European Symposium on Algorithms, pages 552–563. Springer, 2006.

[92] Y. Ostrovsky-Berman and L. Joskowicz. Uncertainty envelopes. In Abstracts of the 21st European Workshop on Comput. Geom., pages 175–178, 2005.

[93] J.M. Phillips. Coresets and sketches. In Handbook of Discrete and Computational Geometry, number Chapter 49. CRC Press, 3rd edition, 2016.

[94] 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.

[95] R. Schneider. Convex bodies: the Brunn-Minkowski theory, volume 44. Cambridge University Press, 1993.

[96] A. Shapiro, D. Dentcheva, and A. Ruszczyński. Lectures on stochastic programming: modeling and theory, volume 16. SIAM, 2014.

[97] A. Sly. Computational transition at the uniqueness threshold. In Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on, pages 287–296. IEEE, 2010.

[98] T.L. Snyder and J.M. Steele. A priori bounds on the Euclidean traveling salesman. In J. Comput, page 24(3). SIAM, 1995.

[99] J.M. Steele. On frieze's $\zeta(3)$ limit for lengths of minimal spanning trees. Discrete Applied Mathematics, 18(1):99–103, 1987.

[100] D. Suciu, D. Olteanu, C. Ré, and C. Koch. Probabilistic databases. Synthesis Lectures on Data Management, 3(2):1–180, 2011.

[101] S. Suri, K. Verbeek, and H. Yıldız. On the most likely convex hull of uncertain points. In Proceedings of the 21st European Symposium on Algorithms, pages 791–802, 2013.

[102] C. Swamy and D. B. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. pages 37(1):33–46, 2006.

[103] W.T. Tutte. A short proof of the factor theorem for finite graphs. Canad. J. Math., 6, 1954.

[104] M. van Kreveld and M. Lö□er. Largest bounding box, smallest diameter, and related problems on imprecise points. Computational Geometry: Theory and Applications, 43:419–433, 2010.

[105] 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.

[106] K. Varadarajan and X. Xiao. On the sensitivity of shape fitting problems. In 32nd International Conference on Foundations of Software Technology and Theoretical Computer Science, page 486, 2012.

[107] J. Vondrak, C. Chekuri, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 783–792. ACM, 2011.

[108] H. Wang and J. Zhang. One-dimensional k-center on uncertain data. Theoretical Computer Science, 602:114–124, 2015.

[109] H. Yıldız, L. Foschini, J. Hershberger, and S. Suri. The union of probabilistic boxes: Maintaining the volume. European Symposia on Algorithms, pages 591–602, 2011.

[110] H. Yu, P.K. Agarwal, R. Poreddy, and K. Varadarajan. Practical methods for shape fitting and kinetic data structures using coresets. Algorithmica, 52(3):378–402, 2008.

[111] 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.

Acknowledgements

First, I would like to thank my advisor, Jian Li, for his invaluable help and patient instructions. He always o□ers me advices of great value and inspiration of new ideas. I enjoyed discussing with him and have learnt a lot from him. It is my honor to be his academic student.

I would like to take this opportunity to thank all of my collaborators: Pingzhong Tang, Yicheng Liu, Lingqing Ai, Xian Wu, Longbo Huang, Qicai Shi, Je□ M. Phillips, Haitao Wang, Pinyan Lu, Chihao Zhang, Hu Ding, Yu Liu, Yifei Jin and Lunjia Hu. I have benefited a lot and improved my research skills while working with them. I am indebted to Professor Pinyan Lu, who was my mentor during my visit to MSRA in Shanghai in the spring of 2015, and Chihao Zhang, who helped me a lot during my internship at MSRA. My PhD life in IIIS was so enjoyable with my dear friends. I would like to thank my friends, Yu, Jianan, Ye, Ruichuang, Linyun, Xian, Lingqing, Yifei, Wei, Mengwen, Dong, Hao, Zhize and Qicai (I have been moving a lot). We have so many colorful shared memory of learning and playing together. I would also like to thank my Dota-mates who won the champion of 9cg in Tsinghua with me: Lingqing, Xian, Linyun, Jianan and Ruichuang.

Finally, I must give my most special thanks to my mother Xifeng and my girl friend Zhongjun, who always accompanied me. Whenever I have met di□culties, they always support me and encourage me. Their selfless love has made this all possible.