Coresets for (k, ℓ)-Median Clustering under the Fréchet Distance

Maike Buchin [0000-0002-3446-4343] and Dennis Rohde [0000-0001-8984-1962] ✉

Faculty of Informatics, Ruhr-University Bochum
`maike.buchin@rub.de`
`dennis.rohde-t1b@rub.de`

Abstract. We present an algorithm for computing ε-coresets for (k, ℓ)-median clustering of polygonal curves in $\mathbb{R}^d$ under the Fréchet distance. This type of clustering is an adaption of Euclidean k-median clustering: we are given a set of n polygonal curves in $\mathbb{R}^d$, each of complexity (number of vertices) at most m, and want to compute k median curves such that the sum of distances from the given curves to their closest median curve is minimal. Additionally, we restrict the complexity of the median curves to be at most ℓ each, to suppress overfitting, a problem specific for sequential data. Our algorithm has running time linear in n, sub-quartic in m and quadratic in ε^{-1}. With high probability it returns ε-coresets of size quadratic in ε^{-1} and logarithmic in n and m. We achieve this result by applying the improved ε-coreset framework by Langberg and Feldman to a generalized k-median problem over an arbitrary metric space. Later we combine this result with the recent result by Driemel et al. on the VC dimension of metric balls under the Fréchet distance. Furthermore, our framework yields ε-coresets for any generalized k-median problem where the range space induced by the open metric balls of the underlying space has bounded VC dimension, which is of independent interest. Finally, we show that our ε-coresets can be used to improve the running time of an existing approximation algorithm for $(1, \ell)$-median clustering.

Keywords: clustering, coresets, median, polygonal curves

1 Introduction

At the present time even efficient approximation algorithms are often incapable of handling massive data sets, which have become common. Here, we need efficient methods to reduce data while (approximately) maintaining the core properties of the data. A popular approach to this topic are ε-coresets; see for example [27,14] for comprehensive surveys. An ε-coreset is a small (weighted) set that aggregates certain properties of a given (massive) data set up to some small error. ε-coresets are very popular in the field of clustering, cf. [20,19,11,22] and they are becoming a topic in other fields, too, cf. [28,16]. The technique for computing an ε-coreset for a given data set highly depends on the application at hand, but mostly ε-coresets are computed by filtering the given data set.

While ε-coresets can be computed efficiently for k-clustering of points in the Euclidean space, less is known for clustering of curves. In particular, only little effort (cf. [10]) has been made in designing methods for computing ε-coresets for (k, ℓ)-clustering of polygonal curves in $\mathbb{R}^d$ under the Fréchet distance. This type of clustering, which has recently drawn increasing popularity due to a growing number of applications, see [12,5,9,6], is an adaption of the Euclidean k-clustering: we are given a set of polygonal curves and seek to compute k center curves that minimize either the maximum Fréchet distance (center objective), or the sum of Fréchet distances (median objective), among the given curves and their closest center curve. In addition, we restrict the complexity—the number of vertices—of each center curve to be at most ℓ to suppress overfitting, a problem specific for sequential data. This means that input curves and center curves are in general of different complexities, which is the reason why we need a specialized algorithm for computing ε-coresets and can not apply ε-coreset algorithms for discrete metric spaces (cf. [15]) on the input.

The Fréchet distance is a natural dissimilarity measure for curves that is a pseudo-metric and can be computed efficiently [1]. Unlike other measures for curves, like the dynamic time warping distance (or the discrete version of the Fréchet distance), it takes the whole course of the curves into account, not only the pairwise distances among their vertices. This can be particularly useful, e.g. when the input consists of irregularly sampled trajectories, cf. [12]. Unfortunately, since the Fréchet distance is a bottleneck distance measure, i.e., it boils down to a single distance between two points on the curves, it is sensitive to outliers, which may negatively affect its applications. In clustering, we can counteract by choosing an appropriate clustering objective and indeed, the (k, ℓ)-median objective is a good choice, because the median is a robust measure of central tendency. However, the state of the art (k, ℓ)-median clustering algorithms (cf. [12,9]) have exponential running time dependencies and cannot be used in practice, while the practical algorithms for (k, ℓ)-clustering (cf. [5,6]) rely on the (k, ℓ)-center objective, which is not robust and therefore amplifies the sensitivity on outliers.

In this work, we present an algorithm for computing ε-coresets for (k, ℓ)-median clustering under the Fréchet distance and improve an $(1, \ell)$-median clustering algorithm by Buchin et al. [8], using ε-coresets and rendering it much more practical.

1.1   Related Work

(k, ℓ)-clustering of polygonal curves was introduced by Driemel et al. [12]. They developed the first approximation schemes for (k, ℓ)-center and (k, ℓ)-median clustering of polygonal curves in $\mathbb{R}$, which run in near-linear time. They proved that both problems are NP-hard, when k is part of the input. Further, they showed that the doubling dimension of the space of polygonal curves under the Fréchet distance is unbounded, even when the curves are of bounded complexity. Subsequently, Buchin et al. [5] presented a constant factor approximation algorithm for (k, ℓ)-center clustering of polygonal curves in $\mathbb{R}^d$, with running time

linear in the number of given curves and polynomial in their maximum complexity. Also they showed that (k, ℓ)-center clustering is NP-hard and NP-hard to approximate within a factor of $(1.5 - \varepsilon)$ for curves in $\mathbb{R}$, respectively $(2.25 - \varepsilon)$ for curves in $\mathbb{R}^d$ with $d \geq 2$, even for $k = 1$. Buchin et al. [6] provided practical algorithms for (k, ℓ) clustering under the Fréchet distance and thereby introduce a new technique, the so called Fréchet centering, for computing better cluster centers. Also, Meintrup et al. [26] provided a practical $(1 + \varepsilon)$-approximation algorithm for discrete (k, ℓ)-median clustering under the presence of a certain number of outliers. Buchin et al. [7] proved that (k, ℓ)-median clustering is also NP-hard, even for $k = 1$. Furthermore, they presented polynomial-time approximation schemes for (k, ℓ)-center and (k, ℓ)-median clustering of polygonal curves under the discrete Fréchet distance. Nath and Taylor [29] gave a near-linear time approximation scheme for (k, ℓ)-median clustering of polygonal curves in $\mathbb{R}^d$ under the discrete Fréchet distance and a polynomial-time approximation scheme for k-median clustering of sets of points from $\mathbb{R}^d$ under the Hausdorff distance. Furthermore, they showed that k-median clustering of point sets under the Hausdorff distance is NP-hard (for constant k). Recently, Buchin et al. [9] developed an approximation scheme for (k, ℓ)-median clustering under the Fréchet distance with running time linear in the number of curves and polynomial in their complexity, where the computed centers have complexity up to $2\ell - 2$.

Langberg and Schulman [25] developed a framework for computing relative error approximations of integrals over any function from a given family of unbounded and non-negative real functions. In particular, this framework can be used to compute ε-coresets for k-clustering of points in $\mathbb{R}^d$ with objective functions based on sums of distances among the points and their closest center. The idea of their framework is to sub-sample the input with respect to a certain non-uniform probability distribution, which is computed using an approximate solution to the problem. More precisely, the approximate solution is used to compute an upper bound on the sensitivity of each data element. The sensitivity is the maximum fraction of cost that the element may cause for any possible solution. It is a notion of the data elements *importance* for the problem and the probability distribution is set up such that each element has probability proportional to its importance. A sample of a certain size drawn from this distribution and properly weighted, is an ε-coreset for the underlying clustering problem with high probability. Feldman and Langberg [15] developed a unified framework for approximate clustering, which is largely based on ε-coresets. They combine the techniques by Langberg and Schulman [25] with ε-approximations, which stem from the framework of range spaces and VC dimension developed in statistical learning theory. As a result, they address a spectrum of clustering problems, such as k-median clustering, k-line median clustering, projective clustering and also other problems like subspace approximation. Braverman et al. [3] improved the aforementioned framework by switching to (ε, η)-approximations, which leads to substantially smaller sample sizes in many cases. Also, they simplified and further generalized the framework and applied it to k-means clustering of points in $\mathbb{R}^d$. Following, Feldman et al. [16] improved this framework by switching to an-

other range space, thereby obtaining smaller coresets for k-means, k-line means and affine subspace clustering.

1.2 Our Contributions

In this work we develop an algorithm for computing ε-coresets for (k, ℓ)-median clustering of polygonal curves in $\mathbb{R}^d$ under the Fréchet distance (where in the following we assume k, ℓ and d to be constant):

Theorem 1. *There exists an algorithm that, given a set of n polygonal curves in $\mathbb{R}^d$ of complexity at most m each and a parameter $\varepsilon \in (0, 1)$, returns with constant positive probability an ε-coreset for (k, ℓ)-median clustering under the Fréchet distance of size $O(k^2 \log^2(k)\varepsilon^{-2} \log(m) \log(kn))$, in time*

$$O(nm \log(m) + nm^3 \log(m) + \varepsilon^{-2} \log(m) \log(n))$$

for $k > 1$ and

$$O(nm \log(m) + m^3 \log(m) + \varepsilon^{-2} \log(m) \log(n))$$

for $k = 1$.

Also we show that ε-coresets can be used to improve the running time of an existing $(1, \ell)$-median $(5 + \varepsilon)$-approximation algorithm [8], thereby facilitating its application in practice.

We start by defining *generalized k-median clustering*, where input and centers come from a subset (not necessarily the same) of an underlying metric space, each, and then derive our ε-coreset result in this setting. This notion captures (k, ℓ)-median clustering under the Fréchet distance in particular, but the analysis holds for any metric space. In doing so, we first give a universal bound on the so called sensitivity of the elements of the given data set and their total sensitivity, i.e., the sum of their sensitivities. The sensitivities are a measure of the data elements importance, i.e., the maximum fraction of the cost an element might cause for any center set, and later they determine the sample probabilities. Our analysis is based on the analysis of Langberg and Schulman [25].

Next, we apply the improved ε-coreset framework by Feldman and Langberg [15]. Here, our analysis is based on the analysis of Feldman et al. [16], but our sample size depends on the VC-dimension of the range space induced by the open metric balls. The open metric balls form a basis of the metric topology, hence it is more natural to study the VC dimension of their associated range space in a geometric setting. Indeed, for the ℓ_p^d spaces these range spaces have already been studied [17, Theorem 2.2] and recently, results for the (continuous and discrete) Fréchet, weak Fréchet and Hausdorff distance were obtained [13], enabling our main result. Finally, we show how an existing $(1, \ell)$-median $(5+\varepsilon)$-approximation algorithm [8] can be improved by means of our ε-coresets.

Theorem 2. *There exists an algorithm that, given a set T of n polygonal curves of complexity at most m each, and a parameter $\varepsilon \in (0, 1/2]$, computes a polygonal curve c of complexity $2\ell - 2$, such that with constant positive probability, it holds that*

$$\mathrm{cost}(T, \{c\}) = \sum_{\tau \in T} d_F(\tau, c) \leq (5 + \varepsilon) \sum_{\tau \in T} d_F(\tau, c^*) = (5 + \varepsilon)\, \mathrm{cost}(T, \{c^*\}),$$

where c^* is an optimal $(1, \ell)$-median for T under the Fréchet distance. The algorithm has running time

$$O\left(nm \log(m) + m^2 \log(m) + m^{2\ell-1}\varepsilon^{-2\ell d + 2d - 2} \log^2(m) \log(n)\right).$$

Theorems 1 and 2 will follow from Theorems 7 and 9, respectively. Note that although we do not present algorithms for computing ε-coresets for the weak and discrete Fréchet and Hausdorff distance, our results also imply the existence of ε-coresets of similar size for these metrics.

1.3   Organization

In Section 2 we give the results for general metric spaces: we derive a universal bound on the sensitivities in Section 2.1 and the ε-coreset result in Section 2.2. In Section 3 we present the algorithm for computing ε-coresets for (k, ℓ)-median clustering. Finally, in Section 4 we demonstrate the use of ε-coresets in an existing $(5 + \varepsilon)$-approximation algorithm for $(1, \ell)$-median clustering.

2   Coresets for Generalized k-Median Clustering in Metric Spaces

In this section, we first derive general results for ε-coreset based on the sensitivity sampling framework [25,15]. In the following $d \in \mathbb{N}$ is an arbitrary constant. By $\|\cdot\|$ we denote the Euclidean norm and for $n \in \mathbb{N}$, we define $[n] = \{1, \ldots, n\}$. For a closed logical formula Ψ we define by $\mathbb{1}(\Psi)$ the function that is 1 if Ψ is true and 0 otherwise.

Let $\mathcal{X} = (X, \rho)$ be an arbitrary metric space, where X is any non-empty set and $\rho \colon X \times X \to \mathbb{R}_{\geq 0}$ is a distance function. We introduce a generalized definition of k-median clustering, where the input is restricted to come from a predefined subset $Y \subseteq X$ and the medians are restricted to come from a predefined subset $Z \subseteq X$.

Definition 1. *The generalized k-median clustering problem is defined as follows, where $k \in \mathbb{N}$ is a fixed (constant) parameter of the problem: given a finite and non-empty set $T = \{\tau_1, \ldots, \tau_n\} \subseteq Y$, compute a set C of k elements from Z, such that $\mathrm{cost}(T, C) = \sum_{\tau \in T} \min_{c \in C} \rho(\tau, c)$ is minimal.*

We analyze the problem in terms of functions. This allows us to apply the improved ε-coreset framework by Feldman and Langberg [15]. Therefore, given a set $T = \{\tau_1, \ldots, \tau_n\} \subseteq Y$ we define $F = \{f_1, \ldots, f_n\}$ to be a set of functions with $f_i \colon 2^Z \setminus \{\emptyset\} \to \mathbb{R}_{\geq 0}$, $C \mapsto \min_{c \in C} \rho(c, \tau_i)$. For each $C \in 2^Z \setminus \{\emptyset\}$ we now have $\mathrm{cost}(T, C) = \sum_{i=1}^n f_i(C)$.

In the following, we bound the sensitivity of each $\tau \in T$. That is the maximum fraction of $\mathrm{cost}(T, C)$ that is caused by τ, for all C. To comply with the k-median problem we only take into account the k-subsets $C \subseteq Z$.

2.1   Sensitivity Bound

First, we formally define the sensitivities of the inputs $\tau \in T$ in terms of the respective functions.

Definition 2 ([15]). *Let F be a finite and non-empty set of functions $f \colon 2^Z \setminus \{\emptyset\} \to \mathbb{R}_{\geq 0}$. For $f \in F$ we define the sensitivity with respect to F:*

$$\mathfrak{s}(f, F) = \sup_{\substack{C = \{c_1, \ldots, c_k\} \subseteq Z \\ \sum_{g \in F} g(C) > 0}} \frac{f(C)}{\sum_{g \in F} g(C)}.$$

We define the total sensitivity of F as $\mathfrak{S}(F) = \sum_{f \in F} \mathfrak{s}(f, F)$.

We now prove a bound on the sensitivity of all $f \in F$, which then yields a bound on the total sensitivity of F. Later, our coreset will be a weighted sample from a distribution whose probabilities are determined by the derived bounds. To compute the bounds, any (bi-criteria) approximate solution to the generalized k-median problem can be used. Our analysis is an adaption of the analysis of the sensitivities for sum-based k-clustering of points in $\mathbb{R}^d$ by Langberg and Schulman [25]. We note that similar bounds have already been derived in the literature, see e.g., [30].

Lemma 1. *Let $k' \in \mathbb{N}$, $C^* = \{c_1^*, \ldots, c_k^*\} \subseteq Z$ with $\Delta^* = \sum_{i=1}^n f_i(C^*)$ minimal and $\hat{C} = \{\hat{c}_1, \ldots, \hat{c}_{k'}\} \subseteq X$ with $\hat{\Delta} = \sum_{i=1}^n f(\hat{C}) \leq \alpha \cdot \Delta^*$ for an $\alpha \in [1, \infty)$. Breaking ties arbitrarily, we assume that every $\tau \in T$ has a unique nearest neighbor in $\hat{C}$ and for $i \in [k']$, we define $\hat{V}_i = \{\tau \in T \mid \forall j \in [k'] : \rho(\tau, \hat{c}_i) \leq \rho(\tau, \hat{c}_j)\}$ to be the Voronoi cell of $\hat{c}_i$ and $\hat{\Delta}_i = \sum_{\tau \in \hat{V}_i} \rho(\tau, \hat{c}_i)$ to be its cost. For each $i \in [k']$ and $\tau_j \in \hat{V}_i$ it holds that*

$$\gamma(f_j) = \left(1 + \sqrt{\frac{2k'}{3\alpha}}\right) \left(\frac{\alpha \rho(\tau_j, \hat{c}_i)}{\hat{\Delta}} + \frac{2\alpha \hat{\Delta}_i}{\hat{\Delta} |\hat{V}_i|}\right) + \left(1 + \sqrt{\frac{3\alpha}{2k'}}\right) \frac{2}{|\hat{V}_i|} \geq \mathfrak{s}(f_j, F)$$

and $\Gamma = \sum_{f \in F} \gamma(f) = 2k' + 2\sqrt{6\alpha k'} + 3\alpha \geq \mathfrak{S}(F)$.

Proof. We assume that $\hat{\Delta} > 0$. By assumption $\hat{V}_1, \ldots, \hat{V}_{k'}$ form a partition of T and by definition $\hat{\Delta} = \sum_{i=1}^{k'} \hat{\Delta}_i$ as well as $\sum_{f \in F} f(C) \geq \hat{\Delta}/\alpha$ for each $C =$

$\{c_1, \ldots, c_k\} \subseteq Z$. For $i \in [k']$, we let $\hat{B}_i = \{\tau \in \hat{V}_i \mid \rho(\tau, \hat{c}_i) \leq 2\hat{\Delta}_i/|\hat{V}_i|\}$. It holds that $|\hat{B}_i| \geq |\hat{V}_i|/2$, since otherwise $\sum_{\tau \in \hat{V}_i \setminus \hat{B}_i} \rho(\tau, \hat{c}_i) > \hat{\Delta}_i$, which is a contradiction. By the triangle-inequality, we have for each $\{c_1, \ldots, c_k\} \subseteq Z$, $i \in [k']$, $j \in [k]$ and $\tau \in T$:

$$\rho(\hat{c}_i, c_j) \leq \rho(\hat{c}_i, \tau) + \rho(\tau, c_j) \iff \rho(\tau, c_j) \geq \rho(\hat{c}_i, c_j) - \rho(\hat{c}_i, \tau).$$

Furthermore, since ρ is non-negative: $\rho(\tau, c_j) \geq \max\{0, \rho(\hat{c}_i, c_j) - \rho(\hat{c}_i, \tau)\}$.

For each $C = \{c_1, \ldots, c_k\} \subseteq Z$, $i \in [k']$ and $\beta \in [0, 1]$ we now have the following bound:

$$\sum_{f \in F} f(C) \geq \beta \sum_{\tau \in \hat{B}_i} \min_{j \in [k]} \rho(\tau, c_j) + (1 - \beta)\frac{\hat{\Delta}}{\alpha}$$

$$\geq \beta \max\left\{0, \rho(\hat{c}_i, c_j) - \frac{2\hat{\Delta}_i}{|\hat{V}_i|}\right\} \frac{|\hat{V}_i|}{2} + (1 - \beta)\frac{\hat{\Delta}}{\alpha},$$

where among $\{c_1, \ldots, c_k\}$, c_j is closest to $\hat{c}_i$. Using the triangle-inequality and the above bound yields for each $\beta \in [0, 1)$, $i \in [k']$ and $\tau_m \in \hat{V}_i$:

$$\mathfrak{s}(f_m, F) \leq \sup_{\{c_1, \ldots, c_k\} \subseteq Z} \frac{\rho(\tau_m, \hat{c}_i) + \rho(\hat{c}_i, c_j)}{\beta \max\left\{0, \rho(\hat{c}_i, c_j) - \frac{2\hat{\Delta}_i}{|\hat{V}_i|}\right\} \frac{|\hat{V}_i|}{2} + (1 - \beta)\frac{\hat{\Delta}}{\alpha}}$$

$$\leq \sup_{\substack{\{c_1, \ldots, c_k\} \subseteq Z \\ \rho(\hat{c}_i, c_j) \geq 2\hat{\Delta}_i/|\hat{V}_i|}} \frac{\rho(\tau_m, \hat{c}_i) + \rho(\hat{c}_i, c_j)}{\beta\left(\frac{|\hat{V}_i|\rho(\hat{c}_i, c_j)}{2} - \hat{\Delta}_i\right) + (1 - \beta)\frac{\hat{\Delta}}{\alpha}}$$

Here, again c_j is closest to $\hat{c}_i$ among $\{c_1, \ldots, c_k\}$ and the last inequality follows because it can be observed that the term takes smaller values for $\rho(\hat{c}_i, c_j) < 2\hat{\Delta}_i/|\hat{V}_i|$ than for $\rho(\hat{c}_i, c_j) \geq 2\hat{\Delta}_i/|\hat{V}_i|$, independent of β. Now, to obtain a bound that is independent of c_j, we substitute $\rho(\hat{c}_i, c_j)$ by a free variable x and let

$$h\colon [2\hat{\Delta}_i/|\hat{V}_i|, \infty) \to \mathbb{R}_{\geq 0}, \ x \mapsto \frac{\rho(\tau_m, \hat{c}_i) + x}{\beta\left(\frac{|\hat{V}_i|x}{2} - \hat{\Delta}_i\right) + (1 - \beta)\frac{\hat{\Delta}}{\alpha}}.$$

The derivative of h is

$$\frac{(1 - \beta)\frac{\hat{\Delta}}{\alpha} - \beta\left(\frac{|\hat{V}_i|\rho(\tau_m, \hat{c}_i)}{2} + \hat{\Delta}_i\right)}{\left(\beta\left(\frac{|\hat{V}_i|x}{2} - \hat{\Delta}_i\right) + (1 - \beta)\frac{\hat{\Delta}}{\alpha}\right)^2}$$

and it can be observed that the sign of this function is independent of x. Therefore, h is a monotonic function and is thus either maximized at $x = 2\hat{\Delta}_i/|\hat{V}_i|$ or when $x \to \infty$. Using l'Hôspital's rule we obtain

$$\mathfrak{s}(f_m, F) \leq \max\left\{\frac{\rho(\tau_m, \hat{c}_i) + \frac{2\hat{\Delta}_i}{|\hat{V}_i|}}{(1 - \beta)\frac{\hat{\Delta}}{\alpha}}, \frac{1}{\beta\frac{|\hat{V}_i|}{2}}\right\} \leq \frac{\alpha\rho(\tau_m, \hat{c}_i)}{(1 - \beta)\hat{\Delta}} + \frac{2\alpha\hat{\Delta}_i}{(1 - \beta)\hat{\Delta}|\hat{V}_i|} + \frac{2}{\beta|\hat{V}_i|}.$$

Therefore,

$$\mathfrak{S}(F) \leq \sum_{i=1}^{k'} \sum_{\tau_m \in \hat{V}_i} \left(\frac{\alpha \rho(\tau_m, \hat{c}_i)}{(1-\beta)\hat{\Delta}} + \frac{2\alpha \hat{\Delta}_i}{(1-\beta)\hat{\Delta}|\hat{V}_i|} + \frac{2}{\beta|\hat{V}_i|} \right) = \frac{3\alpha}{1-\beta} + \frac{2k'}{\beta}.$$

By simple calculus, this bound is minimized at $\beta = \frac{1}{1+\sqrt{\frac{3\alpha}{2k'}}} < 1$. $\qquad\qquad\square$

2.2   Coresets by Sensitivity Sampling

We apply the framework of Feldman and Langberg [15]. First, we formally define ε-coresets for generalized k-median clustering.

Definition 3. *Given $\varepsilon \in (0,1)$ and a finite non-empty set $T \subseteq Y$, a (multi-)set $S \subseteq X$ together with a weight function $w \colon S \to \mathbb{R}_{>0}$ is a weighted ε-coreset for k-median clustering of T, if for all $C \subseteq Z$ with $|C| = k$ it holds that*

$$(1-\varepsilon)\operatorname{cost}(T, C) \leq \operatorname{cost}_w(S, C) \leq (1+\varepsilon)\operatorname{cost}(T, C),$$

where $\operatorname{cost}_w(S, C) = \sum_{s \in S} w(s) \cdot \min_{c \in C} \rho(s, c)$.

We define range spaces and the associated concepts.

Definition 4. *A range space is a pair $(X, \mathcal{R})$, where X is a set, called ground set and $\mathcal{R}$ is a set of subsets $R \subseteq X$, called ranges.*

The projection of a range space $(X, \mathcal{R})$ onto a subset $Y \subseteq X$ is the range space $(Y, \{Y \cap R \mid R \in \mathcal{R}\})$. Furthermore, for each range space there exists a complementary range space.

Definition 5. *Let $F = (X, \mathcal{R})$ be a range space. We call $\overline{F} = (X, \overline{\mathcal{R}})$, the range space over $\overline{\mathcal{R}} = \{X \setminus R \mid R \in \mathcal{R}\}$, the complementary range space of F.*

A measure of the combinatorial complexity of a range space is the VC dimension.

Definition 6. *The VC dimension of a range space $(X, \mathcal{R})$ is the cardinality of a maximum cardinality subset $Y \subseteq X$, such that $|\{Y \cap R \mid R \in \mathcal{R}\}| = 2^{|Y|}$.*

Note that F and $\overline{F}$ have equal VC dimension and for any $Y \subseteq X$, the projection of F onto Y has VC dimension at most the VC dimension of F, see for example [18]. We define (ε, η)-approximations of range spaces.

Definition 7 ([21, Definition 2.3]). *Let $\varepsilon, \eta \in (0,1)$ and $(X, \mathcal{R})$ be a range space with finite non-empty ground set. An (η, ε)-approximation of $(X, \mathcal{R})$ is a set $S \subseteq X$, such that for all $R \in \mathcal{R}$*

$$\left| \frac{|R \cap X|}{|X|} - \frac{|R \cap S|}{|S|} \right| \leq \begin{cases} \varepsilon \cdot \frac{|R \cap X|}{|X|}, & \text{if } |R \cap X| \geq \eta \cdot |X| \\ \varepsilon \cdot \eta, & \text{else.} \end{cases}$$

The following theorem is useful for obtaining (ε, η)-approximations.

Theorem 3 ([21, Theorem 2.11]). *Let $(X, \mathcal{R})$ be a range space with finite non-empty ground set and VC dimension $\mathcal{D}$. Also, let $\varepsilon, \delta, \eta \in (0, 1)$. There is an absolute constant $c \in \mathbb{R}_{>0}$ such that a sample of*

$$\frac{c}{\eta \cdot \varepsilon^2} \cdot \left(\mathcal{D} \log \left(\frac{1}{\eta} \right) + \log \left(\frac{1}{\delta} \right) \right)$$

elements drawn independently and uniformly at random with replacement from X is a (η, ε)-approximation for $(X, \mathcal{R})$ with probability at least $1 - \delta$.

We define open metric balls, which are the ranges used to derive our result.

Definition 8. *For $r \in \mathbb{R}_{\geq 0}$, $z \in Z$ and $Y \subseteq X$ we denote by $\mathrm{B}(z, r, Y) = \{ y \in Y \mid \rho(y, z) < r \}$ the open metric ball with center z and radius r. We denote the set of all open metric balls by $\mathbb{B}(Y, Z) = \{ \mathrm{B}(z, r, Y) \mid z \in Z, r \in \mathbb{R}_{\geq 0} \}$.*

Now, we are ready to analyze the computation of the actual ε-coresets. We use the reduction to uniform sampling, introduced by Feldman and Langberg [15] and improved by Braverman et al. [3] (using Theorem 3). Preferably we would apply Theorem 31 by Feldman et al. [16], which however is not possible since it depends on a range space where each function $f \in F$ may be assigned a distinct scaling factor. This is incompatible with the range space induced by the open metric balls we use to obtain our result. However, by adapting and modifying the proof of their theorem we can derive the desired and more versatile result. To handle necessary scaling factors still involved in the analysis, we incorporate results by Munteanu et al. [28] for bounding the VC dimension.

Theorem 4. *For $f \in F$ we let $\lambda(f) = \left\lceil |F| \cdot 2^{\lceil \log_2(\gamma(f)) \rceil} \right\rceil / |F|$, $\Lambda = \sum_{f \in F} \lambda(f)$, $\psi(f) = \frac{\lambda(f)}{\Lambda}$ and $\mathcal{D}$ be the VC dimension of the range space $(Y, \mathbb{B}(Y, Z))$. Let $\delta, \varepsilon \in (0, 1)$. A sample S of $\Theta \left(\varepsilon^{-2} \alpha k' (\mathcal{D} k \log(k) \log(\alpha k' n) \log(\alpha k') + \log(1/\delta)) \right)$ elements τ_i from T, drawn independently with replacement with probability $\psi(f_i)$ and weighted by $w(f_i) = \frac{\Lambda}{|S| \lambda(f_i)}$ is an ε-coreset with probability at least $1 - \delta$.*

Proof. We define for $C \subseteq Z$ with $|C| = k$ the estimator

$$\widehat{\mathrm{cost}}(S, C) = \sum_{\tau_i \in S} w(f_i) \cdot \min_{c \in C} \rho(\tau_i, c) = \sum_{\tau_i \in S} w(f_i) \cdot f_i(C) = \sum_{\tau_i \in S} \frac{\Lambda}{|S| \lambda(f_i)} f_i(C)$$

for $\mathrm{cost}(T, C)$. We see that

$$\mathrm{E}[\widehat{\mathrm{cost}}(S, C)] = \sum_{i=1}^{|S|} \sum_{\tau_j \in T} \frac{\Lambda}{|S| \lambda(f_j)} f_j(C) \frac{\lambda(f_j)}{\Lambda} = \sum_{\tau_j \in T} f_j(C) = \mathrm{cost}(T, C),$$

thus $\widehat{\mathrm{cost}}(S, C)$ is unbiased. We want to bound the error of $\widehat{\mathrm{cost}}(S, C)$ by applying Theorem 3. To do so, we reduce the sensitivity sampling to uniform sampling

as follows: We let G be a multiset that is a copy of F, where each $f \in F$ is contained $|F|\lambda(f)$ times and is scaled by $\frac{1}{|F|\lambda(f)}$. Note that $|G| = |F|\Lambda$. Also, $\psi(f) = \frac{|F|\lambda(f)}{|G|}$, $|F|\lambda(f)$ is integral for each $f \in F$ and

$$\sum_{g \in G} g(C) = \sum_{f \in F} \frac{|F|\lambda(f)}{|F|\lambda(f)} f(C) = \sum_{f \in F} f(C) = \text{cost}(T, C).$$

Given a sample S', with $|S'| = |S|$, drawn independently and uniformly at random with replacement from G, for $C \subseteq Z$ with $|C| = k$ we define the estimator

$$\widetilde{\text{cost}}(S', C) = \frac{|G|}{|S'|} \sum_{g \in S'} g(C)$$

for $\text{cost}(T, C)$. We see that

$$\text{E}\left[\widetilde{\text{cost}}(S', C)\right] = \frac{|G|}{|S'|} \sum_{i=1}^{|S'|} \sum_{f \in F} \frac{f(C)}{|F|\lambda(f)} \frac{|F|\lambda(f)}{|G|} = \frac{1}{|S'|} \sum_{i=1}^{|S'|} \sum_{f \in F} f(C) = \sum_{g \in G} g(C).$$

Thus, $\widetilde{\text{cost}}(S', C)$ is unbiased, too. We now assume that $S' = \left\{ \frac{1}{|F|\lambda(f_i)} \cdot f_i \mid \tau_i \in S \right\}$. Then,

$$\widetilde{\text{cost}}(S', C) = \frac{|G|}{|S'|} \sum_{g \in S'} g(C) = \frac{|F|\Lambda}{|S'|} \sum_{\tau_i \in S} \frac{1}{|F|\lambda(f_i)} f_i(C) = \sum_{\tau_i \in S} \frac{\Lambda}{|S|\lambda(f_i)} f_i(C)$$

$$= \widehat{\text{cost}}(S, C),$$

so the error bound for $\widetilde{\text{cost}}(S', C)$, that we derive in the following, also applies to $\widehat{\text{cost}}(S, C)$, hence S together with w is a weighted ε-coreset (see Definition 3). We now apply Theorem 3 with the given δ, $\varepsilon/2$ and $\eta = 1/\Lambda$, so the overall error is at most $\varepsilon \cdot \text{cost}(T, C)$ for each $C \subseteq Z$ with $|C| = k$.

For $H \subseteq G$, $C \subseteq Z$ and $r \in \mathbb{R}_{\geq 0}$, we let $\text{range}(H, C, r) = \{g \in H \mid g(C) \geq r\}$. Now, we let $(G, \mathcal{R})$ be a range space over G, where $\mathcal{R} = \{\text{range}(G, C, r) \mid r \in \mathbb{R}_{\geq 0}, C \subseteq Z, |C| = k\}$. For all $C \subseteq Z$ with $|C| = k$ and all $H \subseteq G$ we have that

$$\sum_{g \in H} g(C) = \sum_{g \in H} \int_0^\infty \mathbb{1}(g(C) \geq r) \, dr = \int_0^\infty \sum_{g \in H} \mathbb{1}(g(C) \geq r) \, dr$$

$$= \int_0^\infty |\text{range}(H, C, r)| \, dr. \tag{I}$$

Note that the indicator function is integrable under this circumstances and $|\text{range}(H, C, r)|$ is a step function and is integrable, too. Using this identity, for all $C \subseteq Z$ with $|C| = k$ we now bound the error introduced by $\widehat{\text{cost}}(S, C)$:

$$\left| \text{cost}(T, C) - \widehat{\text{cost}}(S, C) \right| = \left| \text{cost}(T, C) - \widetilde{\text{cost}}(S', C) \right| = \left| \sum_{g \in G} g(C) - \frac{|G|}{|S'|} \sum_{g \in S'} g(C) \right|$$

$$= \left| \int_0^\infty |\text{range}(G, C, r)| \, dr - \frac{|G|}{|S'|} \int_0^\infty |\text{range}(S', C, r)| \, dr \right|$$

$$= \left| \int_0^\infty |\text{range}(G, C, r)| - \frac{|G|}{|S'|} |\text{range}(S', C, r)| \, dr \right|$$

$$\leq \int_0^\infty \left| |\text{range}(G, C, r)| - \frac{|G|}{|S'|} |\text{range}(S', C, r)| \right| \, dr.$$

Here the second equation follows from Eq. (I).

In the following, let $\text{rerror}(C, r) = \left| |\text{range}(G, C, r)| - \frac{|G|}{|S'|} |\text{range}(S', C, r)| \right|$, $r_u(C) = \max_{g \in G} g(C)$, $R_1(C) = \{ r \in \mathbb{R}_{\geq 0} \mid |\text{range}(G, C, r)| \geq \eta \cdot |G| \}$ and $R_2(C) = \mathbb{R}_{\geq 0} \setminus R_1(C)$. Note that $R_1(C)$ and $R_2(C)$ are intervals due to the monotonicity of $|\text{range}(G, C, r)|$. Furthermore, for $r \in (r_u(C), \infty)$ it holds that $|\text{range}(G, C, r)| = 0$. Using these facts, we further derive:

$$\int_0^\infty \text{rerror}(C, r) \, dr = \int_{R_1} \text{rerror}(C, r) \, dr + \int_{R_2} \text{rerror}(C, r) \, dr$$

$$\leq \int_{R_1} \frac{\varepsilon}{2} |\text{range}(G, C, r)| \, dr + \int_{R_2} \frac{\varepsilon}{2} \eta |G| \, dr$$

$$\leq \frac{\varepsilon}{2} \int_0^\infty |\text{range}(G, C, r)| \, dr + \frac{\varepsilon \eta |G|}{2} \int_0^{r_u(C)} dr$$

$$= \frac{\varepsilon}{2} \sum_{g \in G} g(C) + \frac{\varepsilon \eta |G| r_u(C)}{2}. \tag{III}$$

Here, the first inequality follows from Definition 7 and in the last equation we use Eq. (I). Finally, we bound the last summand in Eq. (III). First note that we have for each $g \in G$:

$$\frac{g(C)}{\sum_{h \in G} h(C)} = \frac{\frac{1}{|F| \lambda(f)} f(C)}{\sum_{h \in F} h(C)} \leq \frac{\lambda(f)}{|F| \lambda(f)} \iff \frac{g(C)}{\sum_{h \in G} h(C)} \leq \frac{1}{|F|},$$

where $f \in F$ is the function that g is a copy of and the inequality follows from Definition 2. This implies $r_u(C) \leq \frac{1}{|F|} \sum_{h \in G} h(C)$. We now further derive:

$$\frac{\varepsilon \eta |G| r_u(C)}{2} \leq \frac{\varepsilon}{2} \frac{1}{\Lambda} |F| \Lambda \frac{1}{|F|} \sum_{g \in G} g(C) = \frac{\varepsilon}{2} \sum_{g \in G} g(C).$$

So, all in all $|\text{cost}(T, C) - \widehat{\text{cost}}(S, C)| \leq \varepsilon \cdot \text{cost}(T, C)$ for all $C \subseteq Z$ with $|C| = k$.

The claim now follows from the facts that

- $\gamma(f) \leq \lambda(f) \leq 2 \cdot \gamma(f) + \frac{1}{|F|}$ for each $f \in F$, thus $\Lambda \leq 2 \cdot \Gamma(F) + 1$ and $\Gamma(F) \in O(\alpha k')$ by Lemma 1 and
- $(G, \mathcal{R})$ has VC dimension in $O(\mathcal{D}k \log(k) \log(\alpha k' n))$ by the following Lemma 2.

□

Lemma 2. $(G, \mathcal{R})$ *has VC dimension* $O(\mathcal{D}k \log(k) \log(\alpha k' n))$, *where* $\mathcal{D}$ *is the VC dimension of the range space* $(Y, \mathbb{B}(Y, Z))$.

Proof. First, we assume that there exists a $\varphi \in \mathbb{R}$, such that $\lambda(f) = \varphi$ for all $f \in F$. Then all functions in G are scaled uniformly and we can completely neglect the scaling. In this case $(G, \mathcal{R})$ has equal VC dimension as

$$Q_1 = (T, \{T \setminus (\mathrm{B}(c_1, r, T) \cup \cdots \cup \mathrm{B}(c_k, r, T)) \mid \{c_1, \ldots, c_k\} \subseteq Z, r \in \mathbb{R}_{\geq 0}\}),$$

the VC dimension of Q_1 is at most the VC dimension of

$$Q_2 = (T, \{T \setminus (B_1 \cup \cdots \cup B_k) \mid B_1, \ldots, B_k \in \mathbb{B}(T, Z)\}),$$

Q_2 is the projection of

$$Q_3 = (Y, \{Y \setminus (B_1 \cup \cdots \cup B_k) \mid B_1, \ldots, B_k \in \mathbb{B}(Y, Z)\})$$

onto T and has thus at most the VC dimension of Q_3 and finally, Q_3 is the complementary range space of

$$Q_4 = (Y, \{B_1 \cup \cdots \cup B_k \mid B_1, \ldots, B_k \in \mathbb{B}(Y, Z)\})$$

and has thus equal VC dimension as Q_4. By the k-fold union theorem [2, Lemma 2.3.2] Q_4 has VC dimension $O(\mathcal{D} \cdot k \log(k))$, where $\mathcal{D}$ is the VC dimension of $(Y, \mathbb{B}(Y, Z))$. For the following, let c be the constant hidden in this O-notation.

Contrary to the former case, if there are $t > 1$ distinct values $\Phi = \{\varphi_1, \ldots, \varphi_t\} \subset \mathbb{R}$, such that $\lambda(f) \in \Phi$ for each $f \in F$ and $\forall i \in [t] \exists f \in F : \lambda(f) = \varphi_i$, we apply the techniques of Munteanu et al. [28] (see Lemma 11 and Theorem 15 therein).

First, assume that the VC dimension of $(G, \mathcal{R})$ is greater than $t \cdot c \cdot \mathcal{D} \cdot k \log(k)$. Hence there exists a set $G' \subseteq G$ with $|G'| > t \cdot c \cdot \mathcal{D} \cdot k \log(k)$, such that $|\{G' \cap R \mid R \in \mathcal{R}\}| = 2^{|G'|}$. Let $\{G_1, \ldots, G_t\}$ be a partition of G, such that for each $g \in G_i$ we have $g = \frac{1}{|F|\lambda(f)}f = \frac{1}{|F|\varphi_i}f$ for an $f \in F$. Furthermore, for $i \in [t]$, let $G_i' = G' \cap G_i$.

By disjointness, we have $|\{G_i' \cap R_i \mid R_i = (R \cap G_i), R \in \mathcal{R}\}| = 2^{|G_i'|}$ for each $i \in [t]$ and also there must exist at least one $j \in [t]$, such that $|G_j'| \geq \frac{|G'|}{t} > \frac{t \cdot c \cdot \mathcal{D} \cdot k \log(k)}{t} = c \cdot \mathcal{D} \cdot k \log(k)$, hence the projection of $(G, \mathcal{R})$ on G_j has VC dimension greater than $c \cdot \mathcal{D} \cdot k \log(k)$. This is a contradiction to the former case of uniformly scaled functions in G, thus $(G, \mathcal{R})$ has VC dimension $O(t \cdot \mathcal{D} \cdot k \log(k))$ in this case.

Now, we bound t. Recall that for each $f \in F$, $\lambda(f) = \lceil |F| 2^{\lceil \log_2(\gamma(f)) \rceil} \rceil / |F|$. Furthermore for each $i \in [k']$ and $\tau_j \in \hat{V}_i$ (see Lemma 1),

$$\left(1 + \sqrt{\frac{3\alpha}{2k'}}\right) \frac{2}{|\hat{V}_i|} \leq \gamma(f_j) \leq \left(1 + \sqrt{\frac{2k'}{3\alpha}}\right) \left(\alpha + \frac{2\alpha}{|\hat{V}_i|}\right) + \left(1 + \sqrt{\frac{3\alpha}{2k'}}\right) \frac{2}{|\hat{V}_i|}.$$

Therefore, there can be at most

$$\log_2 \left(\frac{\left(1 + \sqrt{\frac{2k'}{3\alpha}}\right)\left(\alpha + \frac{2\alpha}{|\hat{V}_i|}\right) + \left(1 + \sqrt{\frac{3\alpha}{2k'}}\right)\frac{2}{|\hat{V}_i|}}{\left(1 + \sqrt{\frac{3\alpha}{2k'}}\right)\frac{2}{|\hat{V}_i|}} \right)$$

$$\leq \log_2 \left(\frac{\left(1 + \sqrt{\frac{2k'}{3\alpha}}\right)}{\left(1 + \sqrt{\frac{3\alpha}{2k'}}\right)}(\alpha n/2 + \alpha) + 1 \right)$$

$$\leq \log_2(9\alpha k'(\alpha n/2 + \alpha) + 1)$$

distinct values of $2^{\lceil \log_2(\gamma(f)) \rceil}$, which upper bounds the number of distinct values of $\lambda(f)$.

We conclude that the VC dimension of $(G, \mathcal{R})$ is in $O(\mathcal{D}k \log(k) \log(\alpha k' n))$.

$\square$

3 Coresets for (k, ℓ)-Median Clustering under the Fréchet Distance

Now we present an algorithm for computing ε-coresets for (k, ℓ)-median clustering of polygonal curves under the Fréchet distance. We start by defining polygonal curves.

Definition 9. *A (parameterized) curve is a continuous mapping $\tau \colon [0,1] \to \mathbb{R}^d$. A curve τ is polygonal, iff there exist $v_1, \ldots, v_m \in \mathbb{R}^d$, no three consecutive on a line, called τ's vertices, and $t_1, \ldots, t_m \in [0,1]$ with $t_1 < \cdots < t_m$, $t_1 = 0$ and $t_m = 1$, called τ's instants, such that τ connects every two consecutive vertices $v_i = \tau(t_i), v_{i+1} = \tau(t_{i+1})$ by a line segment.*

We call the segments $\overline{v_1 v_2}, \ldots, \overline{v_{m-1} v_m}$ edges of τ and m the complexity of τ, denoted by $|\tau|$.

Definition 10. *Let $\mathcal{H}$ denote the set of all continuous bijections $h \colon [0,1] \to [0,1]$ with $h(0) = 0$ and $h(1) = 1$, which we call reparameterizations. The Fréchet distance between curves σ and τ is $d_F(\sigma, \tau) = \inf_{h \in \mathcal{H}} \max_{t \in [0,1]} \|\sigma(t) - \tau(h(t))\|$.*

Now we introduce the classes of curves we are interested in.

Definition 11. *For $d \in \mathbb{N}$, we define by $\mathbb{X}^d$ the set of equivalence classes of polygonal curves (where two curves are equivalent, iff they can be made identical by a reparameterization) in ambient space $\mathbb{R}^d$. For $m \in \mathbb{N}$ we define by $\mathbb{X}_m^d$ the subclass of polygonal curves of complexity at most m.*

Finally, we define the (k, ℓ)-median clustering problem for polygonal curves.

Definition 12. *The (k, ℓ)-median clustering problem is defined as follows, where $k, \ell \in \mathbb{N}$ are fixed (constant) parameters of the problem: given a set $T \subset \mathbb{X}_m^d$ of n polygonal curves, compute a set of k curves $C^* \subset \mathbb{X}_\ell^d$, such that $\mathrm{cost}(T, C^*) = \sum_{\tau \in T} \min_{c^* \in C^*} d_F(\tau, c^*)$ is minimal.*

We bound the VC dimension of metric balls under the Fréchet distance by showing that a result of Driemel et al. [13] holds also in our setting.

Theorem 5. *The VC dimension of $(\mathbb{X}_m^d, \mathbb{B}(X_m^d, X_\ell^d))$ is $O\left(\ell^2 \log(\ell m)\right)$.*

Proof. We argue that the claim follows from Theorem 18 by Driemel et al. [13]. First, in their paper polygonal curves do not need to adhere the restriction that no three consecutive vertices may be collinear and they define $\mathbb{X}_m^d$ to be the polygonal curves of exactly m vertices. However, our definitions match by simulating the addition of collinear vertices to those curves in $\mathbb{X}_m^d$ with less than m vertices.

Now, looking into their proof, we can slightly modify the geometric primitives by letting $B_r(p) = \{x \in \mathbb{R}^d \mid \|x - p\| < r\}$, $D_r(\overline{st}) = \{x \in \mathbb{R}^d \mid \exists p \in \overline{st} : \|p - x\| < r\}$, $C_r(\overline{st}) = \{x \in \mathbb{R}^d \mid \exists p \in \ell(\overline{st}) : \|p - x\| < r\}$ and $R_r(\overline{st}) = \{p + u \mid p \in \overline{st}, u \in \mathbb{R}^d, \langle t - s, u \rangle = 0, \|u\| < r\}$, which does not affect the remainder of the proof and thus yields the same bound on the VC dimension. $\square$

To compute ε-coresets for (k, ℓ)-median clustering under the Fréchet distance, we first need to compute the sensitivities and to do so, we utilize constant factor approximation algorithms. We use [9, Algorithm 1], which only works for $k = 1$ but is very efficient in this case. For $k > 1$ we use Algorithm 1, a modification of [12, Algorithm 3], which we now present. This algorithm uses (approximate) minimum-error ℓ-simplifications, which we now define.

Definition 13. *An α-approximate minimum-error ℓ-simplification of a polygonal curve $\tau \in \mathbb{X}^d$ is a curve $\sigma \in \mathbb{X}_\ell^d$ with $d_F(\tau, \sigma) \leq \alpha \cdot d_F(\tau, \sigma')$ for all $\sigma' \in \mathbb{X}_\ell^d$.*

The following lemma is useful to obtain simplifications.

Lemma 3 ([5, Lemma 7.1]). *Given a curve $\sigma \in \mathbb{X}_m^d$, a 4-approximate minimum-error ℓ-simplification can be computed in $O(m^3 \log m)$ time, by combining the algorithms by Alt and Godau [1] and Imai and Iri [23].*

We now present the constant factor approximation algorithm.

Algorithm 1 Constant Factor Approximation for (k, ℓ)-Median Clustering

---

1: **procedure** (k, ℓ)-MEDIAN-96-APPROXIMATION($T = \{\tau_1, \ldots, \tau_n\}$)
2:     **for** $i = 1, \ldots, n$ **do**
3:         $\hat{\tau}_i \leftarrow$ approximate minimum-error ℓ-simplification of τ_i
4:     $C \leftarrow$ Chen's algorithm with $\varepsilon = 0.5, \lambda = \delta$ on $\{\hat{\tau}_1, \ldots, \hat{\tau}_n\}$ [11, Theorem 6.2]
5:     **return** C

We prove the correctness and analyze the running time of Algorithm 1.

Theorem 6. *Given $\delta \in (0,1)$ and $T = \{\tau_1, \ldots, \tau_n\} \subset \mathbb{X}_m^d$, Algorithm 1 returns with probability at least $1 - \delta$ a 109-approximate (k, ℓ)-median solution for T in time $O(nm \log(1/\delta) \log(m) + nm^3 \log(m))$.*

Proof. We assume that the approximate minimum-error
ℓ-simplifications are computed combining the algorithms by Alt and Godau [1] and Imai and Iri [23], so the approximation factor is 4 by Lemma 3.

For $i \in [n]$, let $\hat{\tau}_i$ be the simplification of τ_i and let $\hat{T} = \{\hat{\tau}_1, \ldots, \hat{\tau}_n\}$. Note that $d_F(\tau_i, \hat{\tau}_i) \leq 4 \cdot d_F(\tau, \sigma)$ for all $i \in [n]$ and $\sigma \in \mathbb{X}_\ell^d$.

Let $\hat{C}^* = \{\hat{c}_1^*, \ldots, \hat{c}_k^*\} \subset \mathbb{X}_\ell^d$ be an optimal (k, ℓ)-median solution for $\hat{T}$ and let $C' = \{c_1', \ldots, c_k'\} \subseteq \hat{T}$ be an optimal solution to the discrete (k, ℓ)-median problem for $\hat{T}$, i.e. the centers are chosen among the input. Breaking ties arbitrarily, we assume that every $\hat{\tau} \in \hat{T}$ has a unique nearest neighbor in $\hat{C}^*$ and for $i \in [k]$ we define $\hat{T}_i = \{\hat{\tau} \in \hat{T} \mid \forall j \in [k] : d_F(\hat{\tau}, \hat{c}_i^*) \leq d_F(\hat{\tau}, \hat{c}_j^*)\}$, such that $\hat{T}_1, \ldots, \hat{T}_k$ form a partition of $\hat{T}$. By the triangle inequality:

$$\text{cost}\left(\hat{T}, C'\right) = \min_{\substack{C \subseteq \hat{T} \\ |C|=k}} \sum_{i=1}^{k} \sum_{\hat{\tau} \in \hat{T}_i} \min_{c \in C} d_F(\hat{\tau}, c) \leq \min_{\substack{C \subseteq \hat{T} \\ |C|=k}} \sum_{i=1}^{k} \sum_{\hat{\tau} \in \hat{T}_i} (d_F(\hat{\tau}, \hat{c}_i^*) + \min_{c \in C} d_F(\hat{c}_i^*, c))$$

$$= \text{cost}\left(\hat{T}, \hat{C}^*\right) + \min_{\substack{C \subseteq \hat{T} \\ |C|=k}} \sum_{i=1}^{k} \sum_{\hat{\tau} \in \hat{T}_i} \min_{c \in C} d_F(\hat{c}_i^*, c)$$

$$= \text{cost}\left(\hat{T}, \hat{C}^*\right) + \sum_{i=1}^{k} \sum_{\hat{\tau} \in \hat{T}_i} \min_{\hat{\sigma} \in \hat{T}_i} d_F(\hat{c}_i^*, \hat{\sigma}).$$

For each $i \in [k]$ there must exist a $\hat{\sigma} \in \hat{T}_i$ with $d_F(\hat{\sigma}, \hat{c}_i^*) \leq \sum_{\hat{\tau} \in \hat{T}_i} d_F(\hat{\tau}, \hat{c}_i^*)/|\hat{T}_i|$, since otherwise $\sum_{i=1}^{k} \sum_{\hat{\tau} \in \hat{T}_i} \min_{\hat{\sigma} \in \hat{T}_i} d_F(\hat{c}_i^*, \hat{\sigma}) > \text{cost}\left(\hat{T}, \hat{C}^*\right)$, which is a contradiction. We conclude that $\text{cost}\left(\hat{T}, C'\right) \leq 2 \text{cost}\left(\hat{T}, \hat{C}^*\right)$. Also, by [11, Theorem 6.2] $\text{cost}\left(\hat{T}, C\right) \leq 10.5 \text{cost}\left(\hat{T}, C'\right)$.

Now, let $C^* = \{c_1^*, \ldots, c_k^*\} \subset \mathbb{X}_\ell^d$ be an optimal (k, ℓ)-median solution for T and $C = \{c_1, \ldots, c_k\}$ be a solution returned by Algorithm 1 for $\hat{T}$. We derive:

$$\text{cost}(T, C) \leq \sum_{i=1}^{n} (d_F(\tau_i, \hat{\tau}_i) + \min_{j \in [k]} d_F(\hat{\tau}_i, c_j)) = \sum_{i=1}^{n} d_F(\tau_i, \hat{\tau}_i) + \text{cost}\left(\hat{T}, C\right)$$

$$\leq 4 \text{cost}(T, C^*) + 21 \text{cost}\left(\hat{T}, \hat{C}^*\right) \leq 4 \text{cost}(T, C^*) + 21 \text{cost}\left(\hat{T}, C^*\right)$$

$$\leq 4 \text{cost}(T, C^*) + 21 \left(\sum_{i=1}^{n} (d_F(\hat{\tau}_i, \tau_i) + \min_{j \in [k]} d_F(\tau_i, c_j^*)) \right)$$

$$\leq 4 \text{cost}(T, C^*) + 84 \text{cost}(T, C^*) + 21 \text{cost}(T, C^*) \leq 109 \text{cost}(T, C^*).$$

Computing the simplifications takes time $O(nm^3 \log(m))$, see Lemma 3. Further, we incorporate the given probability of failure (see [11, Theorem 3.6]) into

the running time stated in [11, Theorem 6.2]. Hence, Chen's algorithm can be run in time $O(nm \log(1/\delta) \log(m))$ when the distances are computed using Alt and Godau's algorithm [1]. $\square$

We now present the algorithm for computing weighted ε-coresets for (k,ℓ)-median clustering.

Algorithm 2 Coresets for (k,ℓ)-Median Clustering

1: **procedure** (k,ℓ)-MEDIAN-CORESET$(T = \{\tau_1, \ldots, \tau_n\}, \delta, \varepsilon)$
2:    **if** $k = 1$ **then**
3:        $\hat{c} \leftarrow \ell$-Median-34-Approximation$(T, \delta/2)$ [9, Algorithm 1]
4:        $\hat{C} = \{\hat{c}\}$
5:    **else**
6:        $\hat{C} = \{\hat{c}_1, \ldots, \hat{c}_k\} \leftarrow$ Algorithm 1$(T, \delta/2)$
7:    compute $\hat{V}_1, \ldots, \hat{V}_k, \hat{\Delta}_1, \ldots, \hat{\Delta}_k$ and γ w.r.t. $\hat{C}$ (cf. Lemma 1)
8:    compute λ, Λ w.r.t. γ and ψ w.r.t. λ (cf. Theorem 4)
9:    $S \leftarrow$ sample $\Theta(k\varepsilon^{-2}(d^2\ell^2 k \log(d\ell m) \log(kn) \log^2(k) + \log(1/(2\delta))))$
          elements from T independently with replacement with respect to ψ
10:    compute w w.r.t. λ, Λ and S (cf. Theorem 4)
11:    **return** S and w

We prove the correctness and analyze the running time of Algorithm 2. Also, we analyze the size of the resulting ε-coreset.

Theorem 7. *Given a set $T = \{\tau_1, \ldots, \tau_n\} \subset \mathbb{X}_m^d$ and $\delta, \varepsilon \in (0,1)$, Algorithm 2 computes a weighted ε-coreset of size $O(\varepsilon^{-2}(\log(m)\log(n) + \log(1/\delta)))$ for (k,ℓ)- median clustering with probability at least $1 - \delta$, in time*

$$O(nm \log(m) \log(1/\delta) + nm^3 \log(m) + \varepsilon^{-2}(\log(m)\log(n) + \log(1/\delta)))$$

for $k > 1$ and

$$O(nm \log(m) + m^2 \log(m) \log^2(1/\delta) + m^3 \log(m) + \varepsilon^{-2}(\log(m)\log(n) + \log(1/\delta)))$$

for $k = 1$.

Proof. First note that by a union bound, with probability at least $1 - \delta$ Algorithm 1, respectively [9, Algorithm 1], and the sampling are simultaneously successful (see Theorems 4 and 6 and [9, Corollary 3.1]). The correctness of the algorithm follows from the observations that the (k,ℓ)-median clustering objective fits the generalized k-median clustering objective with $X = \mathbb{X}^d$, $Y = \mathbb{X}_m^d \subset X$ and $Z = \mathbb{X}_\ell^d \subset X$, therefore Lemma 1 and Theorem 4 can be applied, and the VC dimension of $(X, \mathbb{B}(Y, Z))$ is $O(\ell^2 \log(\ell m))$ by Theorem 5.

We now analyze the running time. $\hat{V}_1, \ldots, \hat{V}_k, \hat{\Delta}_1, \ldots, \hat{\Delta}_k$ and γ can be computed in time $O(nm \log(m))$ using Alt and Godau's algorithm [1]. λ, Λ and ψ can be computed in time $O(n)$ and the sampling can be carried out in time

$O(\varepsilon^{-2}(\log(m)\log(n) + \log(1/\delta)))$. Finally, w can be computed in time $O(n)$. If $k > 1$ we run Algorithm 1 in time $O(nm\log(1/\delta)\log(m) + nm^3\log(m))$, see Theorem 6. Else, we run [9, Algorithm 1] in time $(m^2\log(m)\log^2(1/\delta) + m^3\log(m))$, see [9, Corollary 3.1]. All in all, the running time is then

$$O(nm\log(1/\delta)\log(m) + nm^3\log(m) + \varepsilon^{-2}(\log(m)\log(n) + \log(1/\delta)))$$

for $k > 1$ and

$$O(nm\log(m) + m^2\log(m)\log^2(1/\delta) + m^3\log(m) + \varepsilon^{-2}(\log(m) + \log(1/\delta)))$$

for $k = 1$. □

4 Towards Practical $(1, \ell)$-Median Approximation Algorithms

In this section, we present a modification of Algorithm 3 from [8]. Our modification uses ε-coresets to improve the running time of the algorithm, rendering it more tractable in a big data setting. We start by giving some definitions. For $p \in \mathbb{R}^d$ and $r \in \mathbb{R}_{\geq 0}$ we denote by $B(p, r) = \{q \in \mathbb{R}^d \mid \|p - q\| \leq r\}$ the closed Euclidean ball of radius r with center p. We give a standard definition of grids.

Definition 14 (grid). *Given a number $r \in \mathbb{R}_{>0}$, for $(p_1, \ldots, p_d) \in \mathbb{R}^d$ we define by $G(p, r) = (\lfloor p_1/r \rfloor \cdot r, \ldots, \lfloor p_d/r \rfloor \cdot r)$ the r-grid-point of p. Let $P \subseteq \mathbb{R}^d$ be a subset of $\mathbb{R}^d$. The grid of cell width r that covers P is the set $\mathbb{G}(P, r) = \{G(p, r) \mid p \in P\}$.*

Such a grid partitions the set P into cubic regions and for each $r \in \mathbb{R}_{>0}$ and $p \in P$ we have that $\|p - G(p, r)\| \leq \sqrt{d}r$. The following theorem by Indyk [24] is useful for evaluating the cost of a curve at hand.

Theorem 8 ([24, Theorem 31]). *Let $\varepsilon \in (0, 1]$ and $T \subset \mathbb{X}^d$ be a set of polygonal curves. Further let W be a non-empty sample, drawn uniformly and independently at random from T, with replacement. For $\tau, \sigma \in T$ with $\mathrm{cost}(T, \tau) > (1 + \varepsilon)\mathrm{cost}(T, \sigma)$ it holds that $\Pr[\mathrm{cost}(W, \tau) \leq \mathrm{cost}(W, \sigma)] < \exp\left(-\varepsilon^2|W|/64\right)$.*

The following theorem, which we combine with fine-tuned grids, allows us to obtain low-complexity center curves.

Lemma 4 ([8, Lemma 4.1]). *Let $\sigma, \tau \in \mathbb{X}^d$ be polygonal curves. Let $v_1^\tau, \ldots, v_{|\tau|}^\tau$ be the vertices of τ and let $r = d_F(\sigma, \tau)$. There exists a polygonal curve $\sigma' \in \mathbb{X}^d$ with every vertex contained in at least one of $B(v_1^\tau, r), \ldots, B(v_{|\tau|}^\tau, r)$, $d_F(\sigma', \tau) \leq d_F(\sigma, \tau)$ and $|\sigma'| \leq 2|\sigma| - 2$.*

Finally, we present our improved modification of Algorithm 3 from [8]. This algorithm uses ε-coresets every time it has to evaluate the cost of a center set. The dramatic effect of this small modification is that we nearly lose the original

linear running time dependency on n in the most time consuming part of the algorithm, rendering it practical in the setting where we have a lot of curves of much smaller complexity than number ($\ell < m \ll n$).

Algorithm 3 $(1, \ell)$-Median by Simple Shortcutting and ε-Coreset

---

1: **procedure** $(1, \ell)$-MEDIAN-$(5 + \varepsilon)$-APPROXIMATION$(T = \{\tau_1, \ldots, \tau_n\}, \delta, \varepsilon)$
2:      $\hat{c} \leftarrow (1, \ell)$-Median-34-Approximation$(T, \delta/4)$ [9, Algorithm 1]
3:      $\varepsilon' \leftarrow \varepsilon/67$, $P \leftarrow \emptyset$
4:      $(T', w) \leftarrow (1, 2\ell - 2)$-Median-Coreset$(T, \delta/4, \varepsilon')$
5:      $\Delta \leftarrow \text{cost}_w(T', \{\hat{c}\})$, $\Delta_u \leftarrow \Delta/(1 - \varepsilon')$, $\Delta_l \leftarrow \Delta/((1 + \varepsilon')34)$
6:      $S \leftarrow$ sample $\lceil -2(\varepsilon')^{-1}(\ln(\delta) - \ln(4)) \rceil$ curves from T uniformly
             and independently with replacement
7:      $W \leftarrow$ sample $\lceil -64(\varepsilon')^{-2}(\ln(\delta) - \ln(\lceil -8(\varepsilon')^{-1}(\ln(\delta) - \ln(4)) \rceil)) \rceil$ curves
             from T uniformly and independently with replacement
8:      $c \leftarrow$ arbitrary element from $\arg\min_{s \in S} \text{cost}(W, s)$
9:      **for** $i = 1, \ldots, |c|$ **do**
10:         $P \leftarrow P \cup \mathbb{G}\left(B\left(v_i^c, (3 + 4\varepsilon')\Delta_u/n\right), \varepsilon'\Delta_l/(n\sqrt{d})\right)$ (v_i^c: i^{th} vertex of c)
11:      $C \leftarrow$ set of all polygonal curves with $2\ell - 2$ vertices from P
12:      **return** $\arg\min_{c' \in C} \text{cost}_w(T', \{c'\})$

We show the correctness and analyze the running time of Algorithm 3.

Theorem 9. *Given two parameters $\delta \in (0, 1)$, $\varepsilon \in (0, 1/2]$ and a set $T = \{\tau_1, \ldots, \tau_n\} \subset \mathbb{X}_m^d$ of polygonal curves, with probability at least $1 - \delta$ Algorithm 3 returns a $(5 + \varepsilon)$-approximate $(1, \ell)$-median for T with $2\ell - 2$ vertices, in time*

$$O\left(nm\log(m) + m^2\log(m)\log^2(1/\delta) + m^{2\ell-1}\frac{\log(m)\log(n) + \log(1/\delta))\log(m)}{\varepsilon^{2\ell d - 2d + 2}}\right).$$

Proof. Let $c^* \in \arg\min_{c \in \mathbb{X}_\ell^d} \text{cost}(T, \{c\})$ be an optimal $(1, \ell)$-median for T. The expected distance between $s \in S$ and c^* is

$$\mathrm{E}\left[d_F(s, c^*)\right] = \sum_{i=1}^{n} d_F(\tau_i, c^*) \cdot \frac{1}{n} = \frac{\text{cost}(T, \{c^*\})}{n}.$$

Now, using Markov's inequality, for every $s \in S$ we have

$$\Pr[d_F(s, c^*) > (1 + \varepsilon)\text{cost}(T, \{c^*\})/n] \leq \frac{\text{cost}(T, \{c^*\})n^{-1}}{(1 + \varepsilon)\text{cost}(T, \{c^*\})n^{-1}} = \frac{1}{1 + \varepsilon},$$

therefore by independence

$$\Pr\left[\bigwedge_{s \in S} (d_F(s, c^*) > (1 + \varepsilon)\text{cost}(T, \{c^*\})/n)\right] \leq \frac{1}{(1 + \varepsilon)^{|S|}} \leq \exp\left(-\frac{\varepsilon|S|}{2}\right).$$

Hence, with probability at most $\exp\left(-\varepsilon\frac{\left\lceil-\frac{2(\ln(\delta)-\ln(4))}{\varepsilon}\right\rceil}{2}\right) \le \delta/4$ there is no $s \in S$ with $d_F(s, c^*) \le (1+\varepsilon)\frac{\mathrm{cost}(T,\{c^*\})}{n}$. Now, assume there is a $s \in S$ with $d_F(s, c^*) \le (1+\varepsilon)\mathrm{cost}(T, \{c^*\})/n$. We do not want any $t \in S \setminus \{s\}$ with $\mathrm{cost}(T, \{t\}) > (1+\varepsilon)\mathrm{cost}(T, \{s\})$ to have $\mathrm{cost}(W, \{t\}) \le \mathrm{cost}(W, \{s\})$. Using Theorem 8, we conclude that this happens with probability at most

$$\exp\left(-\frac{\varepsilon^2\lceil-64\varepsilon^{-2}(\ln(\delta) - \ln(\lceil-8(\varepsilon')^{-1}(\ln(\delta) - \ln(4))\rceil))\rceil}{64}\right)$$

$$\le \frac{\delta}{\lceil-8(\varepsilon')^{-1}(\ln(\delta) - \ln(4))\rceil}$$

$$\le \frac{\delta}{4|S|},$$

for each $t \in S \setminus \{s\}$. Also, with probability at most $\delta/4$ [9, Algorithm 1] fails to compute a 34-approximate $(1, \ell)$-median $\hat{c} \in \mathbb{X}_\ell^d$ for T, see [9, Corollary 3.1], and with probability at most $\delta/4$, Algorithm 2 fails to compute a weighted ε-coreset for T, see Theorem 4.

Using a union bound over these bad events, we conclude that with probability at least $1 - \delta$,

- Algorithm 3 samples a curve $s \in S$ with $d_F(s, c^*) \le (1 + \varepsilon)\mathrm{cost}(T, \{c^*\})/n$,
- Algorithm 3 samples a curve $t \in S$ with $\mathrm{cost}(T, \{t\}) \le (1 + \varepsilon)\mathrm{cost}(T, \{s\})$,
- [9, Algorithm 1] computes a 34-approximate $(1, \ell)$-median $\hat{c} \in \mathbb{X}_\ell^d$ for T, i.e., $\mathrm{cost}(T, \{c^*\}) \le \mathrm{cost}(T, \{\hat{c}\}) \le 34\,\mathrm{cost}(T, \{c^*\})$
- and Algorithm 2 computes a weighted ε-coreset for T.

Using the triangle-inequality yields

$$\sum_{\tau \in T}(d_F(t, c^*) - d_F(\tau, c^*)) \le \sum_{\tau \in T} d_F(t, \tau) \le (1 + \varepsilon)\sum_{\tau \in T} d_F(s, \tau)$$

$$\le (1 + \varepsilon)\sum_{\tau \in T}(d_F(\tau, c^*) + d_F(c^*, s)),$$

which is equivalent to

$$n \cdot d_F(t, c^*) \le (2 + \varepsilon)\mathrm{cost}(T, c^*) + (1 + \varepsilon)n(1 + \varepsilon)\mathrm{cost}(T, c^*)/n$$
$$\Leftrightarrow d_F(t, c^*) \le (3 + 4\varepsilon)\mathrm{cost}(T, c^*)/n.$$

Let $v_1^t, \ldots, v_{|t|}^t$ be the vertices of t. By Lemma 4 there exists a polygonal curve $c' \in \mathbb{X}_{2\ell-2}^d$ with every vertex contained in one of $B(v_1^t, d_F(c^*, t)), \ldots, B(v_{|t|}^t, d_F(c^*, t))$ and $d_F(t, c') \le d_F(t, c^*)$. We have $d_F(t, c') \le d_F(t, c^*) \le (3+4\varepsilon)\mathrm{cost}(T, \{c^*\})/n$. Furthermore, by the ε-coreset guarantee, see Definition 3, we have $|\Delta - \mathrm{cost}(T, \{\hat{c}\})| \le \varepsilon\,\mathrm{cost}(T, \{\hat{c}\})$. Therefore, $\Delta_l = \Delta/(34(1 + \varepsilon)) \le \mathrm{cost}(T, \{c^*\}) \le \Delta_u = \Delta/(1 - \varepsilon)$ and $d_F(t, c') \le (3+4\varepsilon)\Delta_u/n$. We conclude that the set C of all curves with up to $2\ell - 2$ vertices

from P, the union of the grid covers, contains a curve $c'' \in \mathbb{X}_{2\ell-2}^d$ with distance at most $\frac{\varepsilon \Delta_l}{n} \leq \varepsilon \frac{\text{cost}(T, \{c^*\})}{n}$ between every corresponding pair of vertices of c' and c'', thus $d_F(t, c'') \leq (3 + 5\varepsilon) \text{cost}(T, \{c^*\})/n$.

In the last step, Algorithm 3 returns a curve $\bar{c} \in C$, that evaluates best against the ε-coreset. By the ε-coreset guarantee and the range of ε, we know that $\text{cost}(T, \{\bar{c}\}) \leq (1 + \varepsilon)/(1 - \varepsilon) \text{cost}(T, \{c''\}) \leq (1 + 4\varepsilon) \text{cost}(T, \{c''\})$. We can now bound the cost of $\bar{c}$ as follows:

$$
\begin{aligned}
\text{cost}(T, \{\bar{c}\}) &\leq (1 + 4\varepsilon) \sum_{\tau \in T} d_F(\tau, c'') \leq (1 + 4\varepsilon) \sum_{\tau \in T} (d_F(\tau, t) + d_F(t, c'')) \\
&\leq (1 + 4\varepsilon) \text{cost}(T, \{t\}) + (1 + 4\varepsilon)(3 + 5\varepsilon) \text{cost}(T, \{c^*\}) \\
&\leq (1 + \varepsilon)(1 + 4\varepsilon) \text{cost}(T, \{s\}) + (3 + 37\varepsilon) \text{cost}(T, \{c^*\}) \\
&\leq (1 + 9\varepsilon) \sum_{\tau \in T} (d_F(\tau, c^*) + d_F(c^*, s)) + (3 + 37\varepsilon) \text{cost}(T, \{c^*\}) \\
&\leq (4 + 48\varepsilon) \text{cost}(T, \{c^*\}) + (1 + \varepsilon)(1 + 9\varepsilon) \text{cost}(T, \{c^*\}) \\
&\leq (5 + 67\varepsilon) \text{cost}(T, \{c^*\})
\end{aligned}
$$

Finally, we rescale ε by $\frac{1}{67}$ to obtain the desired approximation guarantee.

We now discuss the running time. [9, Algorithm 1] has running time $O(m^2 \log(m) \log^2(1/\delta) + m^3 \log(m))$, see [9, Corollary 3.1] and Algorithm 2 has running time

$$
O(nm \log(m) + m^2 \log(m) \log^2(1/\delta) + m^3 \log(m) + \varepsilon^{-2}(\log(m) \log(n) + \log(1/\delta))),
$$

see Theorem 4. The ε-coreset has size $O(\varepsilon^{-2}(\log(m) \log(n) + \log(1/\delta)))$, therefore $\text{cost}_w(T', \hat{c})$ can be evaluated in time $O(m\varepsilon^{-2} \log(m)(\log(m) \log(n) + \log(1/\delta)))$, using Alt and Godau's algorithm [1] to compute the distances.

The sample S has size $O\left(\frac{-\ln(\delta)}{\varepsilon}\right)$ and the sample W has size $O\left(\frac{-\ln(\delta)}{\varepsilon^2}\right)$. Evaluating each curve of S against W takes time $O\left(\frac{m^2 \log(m) \log^2(1/\delta)}{\varepsilon^3}\right)$, using Alt and Godau's algorithm [1] to compute the distances.

Now, c has up to m vertices and every grid consists of
$$\left(\frac{\frac{2(3+4\varepsilon')\Delta_u}{n}}{\frac{2\varepsilon'\Delta_l}{n\sqrt{d}}}\right)^d = \left(\frac{(3+4\varepsilon')\sqrt{d}}{\varepsilon'}34(1+\varepsilon)\right)^d \in O\left(\frac{1}{\varepsilon^d}\right)$$ points (note that $\Delta_u/\Delta_l = (1 + \varepsilon')/(1 - \varepsilon')34 \leq 34(1 + \varepsilon)$). Therefore, we have $O\left(\frac{m}{\varepsilon^d}\right)$ points in P and Algorithm 3 enumerates all combinations of $2\ell - 2$ points from P taking time $O\left(\frac{m^{2\ell-2}}{\varepsilon^{(2\ell-2)d}}\right)$. Afterwards, these candidates are evaluated against the ε-coreset, which takes time

$$
O\left(\frac{m^{2\ell-1}(\log(m) \log(n) + \log(1/\delta)) \log(m)}{\varepsilon^{2\ell d - 2d + 2}}\right),
$$

using Alt and Godau's algorithm [1] to compute the distances. All in all, we then have running time

$$O\left(nm\log(m) + m^2\log(m)\log^2(1/\delta) + \frac{m^{2\ell-1}(\log(m)\log(n) + \log(1/\delta))\log(m)}{\varepsilon^{2\ell d - 2d + 2}}\right).$$

$\square$

5   Conclusion

We presented an algorithm for computing ε-coresets for (k, ℓ)-median clustering of polygonal curves under the Fréchet distance and used these to improve the running time of an existing approximation algorithm for $(1, \ell)$-median clustering. Unfortunately, it was not possible to improve the existing (k, ℓ)-median approximation algorithms in [12,9] by means of ε-coresets. This is due to the recursive approximation scheme used in these works, where the candidate center sets are not necessarily evaluated against the input, but against subsets of the input. Thus, we would need an ε-coreset for any subset of the input, which is not practical. We note that, to the best of our knowledge, no (k, ℓ)-median clustering algorithm exists that do not employ this approximation scheme.

It is still an interesting open problem whether there exist sublinear size ε-coresets for weighted sets of polygonal curves. To derive such a result one may need a sublinear bound on the VC dimension of the range space of metric balls under scaled Fréchet distances, which is not evident at the moment. We note that such a result would enable the use of the iterative size reduction technique recently introduced by Braverman et al. [4].

References

1. Alt, H., Godau, M.: Computing the Fréchet distance between two polygonal curves. International Journal of Computational Geometry & Applications 5, 75–91 (1995)
2. Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K.: Learnability and the Vapnik-Chervonenkis dimension. Journal of the ACM 36(4), 929–965 (1989)
3. Braverman, V., Feldman, D., Lang, H.: New Frameworks for Offline and Streaming Coreset Constructions. CoRR abs/1612.00889 (2016)
4. Braverman, V., Jiang, S.H., Krauthgamer, R., Wu, X.: Coresets for Clustering in Excluded-minor Graphs and Beyond. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms. pp. 2679–2696. SIAM (2021)
5. Buchin, K., Driemel, A., Gudmundsson, J., Horton, M., Kostitsyna, I., Löffler, M., Struijs, M.: Approximating (k, l)-center clustering for curves. In: Proceedings of the 2019 ACM-SIAM Symposium on Discrete Algorithms. pp. 2922–2938 (2019)
6. Buchin, K., Driemel, A., van de L'Isle, N., Nusser, A.: klcluster: Center-based Clustering of Trajectories. In: Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. pp. 496–499 (2019)

7. Buchin, K., Driemel, A., Struijs, M.: On the Hardness of Computing an Average Curve. In: 17th Scandinavian Symposium and Workshops on Algorithm Theory. LIPIcs, vol. 162, pp. 19:1–19:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)

8. Buchin, M., Driemel, A., Rohde, D.: Approximating (k, ℓ)-Median Clustering for Polygonal Curves. arXiv e-prints arXiv:2009.01488 (2020)

9. Buchin, M., Driemel, A., Rohde, D.: Approximating (k, ℓ)-Median Clustering for Polygonal Curves. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms. pp. 2697–2717. SIAM (2021)

10. Buchin, M., Rohde, D.: Coresets for (k, l)-Clustering under the Fréchet Distance. In: Proceedings of the 35th European Workshop on Computational Geometry (2019)

11. Chen, K.: 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)

12. Driemel, A., Krivosija, A., Sohler, C.: Clustering time series under the Fréchet distance. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 766–785 (2016)

13. Driemel, A., Phillips, J.M., Psarros, I.: The VC dimension of metric balls under Fréchet and Hausdorff distances. In: 35th International Symposium on Computational Geometry. pp. 28:1–28:16 (2019)

14. Feldman, D.: Introduction to Core-sets: an Updated Survey. CoRR abs/2011.09384 (2020)

15. Feldman, D., Langberg, M.: A unified framework for approximating and clustering data. In: Proceedings of the 43rd ACM Symposium on Theory of Computing. pp. 569–578. ACM (2011)

16. Feldman, D., Schmidt, M., Sohler, C.: Turning Big Data Into Tiny Data: Constant-Size Coresets for k-Means, PCA, and Projective Clustering. SIAM Journal on Computing 49(3), 601–657 (2020)

17. Goldberg, P.W., Jerrum, M.: Bounding the Vapnik-Chervonenkis Dimension of Concept Classes Parameterized by Real Numbers. Mach. Learn. 18(2-3), 131–148 (1995)

18. Har-Peled, S.: Geometric Approximation Algorithms. American Mathematical Society (2011)

19. Har-Peled, S., Kushal, A.: Smaller Coresets for k-Median and k-Means Clustering. Discret. Comput. Geom. 37(1), 3–19 (2007)

20. Har-Peled, S., Mazumdar, S.: On coresets for k-means and k-median clustering. In: Proceedings of the 2004 ACM Symposium on Theory of Computing. pp. 291–300 (2004)

21. Har-Peled, S., Sharir, M.: Relative (p,ε)-Approximations in Geometry. Discrete & Computational Geometry 45(3), 462–496 (2011)

22. Huang, L., Jiang, S.H., Li, J., Wu, X.: Epsilon-Coresets for Clustering (with Outliers) in Doubling Metrics. In: 59th IEEE Annual Symposium on Foundations of Computer Science. pp. 814–825. IEEE Computer Society (2018)

23. Imai, H., Iri, M.: Polygonal Approximations of a Curve — Formulations and Algorithms. Machine Intelligence and Pattern Recognition 6, 71–86 (1988)

24. Indyk, P.: High-dimensional Computational Geometry. Ph.D. thesis, Stanford University, CA, USA (2000)

25. Langberg, M., Schulman, L.J.: Universal epsilon-approximators for Integrals. In: Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 598–607 (2010)

26. Meintrup, S., Munteanu, A., Rohde, D.: Random Projections and Sampling Algorithms for Clustering of High-Dimensional Polygonal Curves. In: Advances in Neural Information Processing Systems 32. pp. 12807–12817 (2019)
27. Munteanu, A., Schwiegelshohn, C.: Coresets-Methods and History: A Theoreticians Design Pattern for Approximation and Streaming Algorithms. Künstliche Intell. 32(1), 37–53 (2018)
28. Munteanu, A., Schwiegelshohn, C., Sohler, C., Woodruff, D.P.: On Coresets for Logistic Regression. In: Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018. pp. 6562–6571 (2018)
29. Nath, A., Taylor, E.: k-Median Clustering Under Discrete Fréchet and Hausdorff Distances. In: 36th International Symposium on Computational Geometry. LIPIcs, vol. 164, pp. 58:1–58:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
30. Varadarajan, K.R., Xiao, X.: On the Sensitivity of Shape Fitting Problems. In: IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. LIPIcs, vol. 18, pp. 486–497 (2012)